

# Privacy, Patience, and Protection\*

Ronen Gradwohl<sup>†</sup>      Rann Smorodinsky<sup>‡</sup>

## Abstract

We analyze repeated games in which players have private information about their levels of patience and in which they would like to maintain the privacy of this information vis-à-vis third parties. We show that privacy protection in the form of shielding players' actions from outside observers is harmful, as it limits and sometimes eliminates the possibility of attaining Pareto-optimal payoffs.

**JEL Classification:** C72, D82

**Keywords:** Privacy, privacy protection, perception games, signaling games

---

\*Gradwohl gratefully acknowledges the support of National Science Foundation grant 1718670. Smorodinsky gratefully acknowledges United States-Israel Binational Science Foundation and National Science Foundation grant 2016734, the German-Israel Foundation grant I-1419-118.4/2017, the Ministry of Science and Technology grant 19400214, Technion VPR grants, and the Bernard M. Gordon Center for Systems Engineering at the Technion. Declarations of interest: None. Data sharing not applicable to this article as no datasets were generated or analyzed during the current study.

<sup>†</sup>Department of Economics and Business Administration, Ariel University. Email: [roneng@ariel.ac.il](mailto:roneng@ariel.ac.il).

<sup>‡</sup>Faculty of Industrial Engineering and Management, The Technion – Israel Institute of Technology. Email: [rann@ie.technion.ac.il](mailto:rann@ie.technion.ac.il).

# 1 Introduction

Questions surrounding consumer privacy have taken center stage in recent years, as have debates surrounding privacy protection. These debates feature at least two distinct ways of thinking about the value of privacy. First, within economics, privacy is generally viewed as having some instrumental value, affecting present and future interactions by changing the informational landscape. Privacy protection alters market interactions by changing behavior and equilibrium, and this may have positive or negative effects. For example, in the online shopping market, when there is no protection and information about online purchases is public, retailers can better discern consumers' preferences and the market is more efficient (Fudenberg and Villas-Boas, 2006). In contrast, with repeated bargaining or auctions, outcomes can be more efficient when privacy protection is in place and individuals' behavior is not observed; examples include Hörner and Vieille (2009), Bergemann and Hörner (2018), and Chaves (2019).

A second approach is to view privacy as having some intrinsic value, where the leakage of an individual's private information is directly associated with a decrease in well-being.<sup>1</sup> This is the typical approach within the computer science literature, manifested in the large body of work on "differential privacy" (DP).<sup>2</sup> However, while this approach generally promotes maximal privacy protection, it often ignores the effects of such protection on behavior.

In this paper we examine the impacts of privacy protection while accounting for *both* the instrumental and intrinsic values of privacy. Privacy protection consists of some technology measure (e.g., encryption) or proper regulation that bans the observation or collection of information, and thus shields indi-

---

<sup>1</sup>Such an association is corroborated by surveys of individuals' privacy concerns, such as Rainie et al. (2013) and Madden et al. (2014).

<sup>2</sup>Loosely, a mechanism satisfies  $\epsilon$ -DP for some individual  $i$  if the distribution over the outcomes of the mechanism with  $i$  present is  $\epsilon$ -close to the distribution with  $i$  absent (Dinur and Nissim, 2003). See Dwork (2008) for a survey and Abowd and Schmutte (2019) for a discussion of DP from an economic point of view.

viduals' actions from outside observers. For example, in the online shopping market, privacy can be protected by add-ons that erase cookies from the user's browser to hide her actions, as well as by legal restrictions on collecting or selling individual browsing or shopping history. Our main conclusion is that, even (and especially) when privacy has large intrinsic value, such privacy protection may be harmful because of its effects on behavior.

In a previous paper (Gradwohl and Smorodinsky, 2017) we show that for single-agent decision-making problems such privacy protection is always beneficial. This is quite intuitive, as such protection provides a dual benefit to the decision maker: it eliminates the cost associated with information leakage, while also allowing the decision maker to choose an efficient action without worrying about its privacy-related implications. This straightforward observation, and in particular the second benefit, does not carry over to multi-player games, and so the question about the desirability of privacy protection remains.

In this paper we study this question in a multi-player setting in which, absent privacy concerns, efficiency can also be attained in equilibrium. In particular, we consider repeated games, where patient players can cooperate and reach an efficient outcome in equilibrium.<sup>3</sup> For concreteness, and to illustrate our model and results, consider an online shopping scenario in which two buyers repeatedly compete for an item in a first-price auction. At each stage they both value the product at 6 and are allowed to bid either 0 or 2. Furthermore, when the buyers bid equally one is randomly awarded the item. Note that the resulting stage game is none other than the Prisoner's Dilemma (PD), depicted in Figure 1a. The folk theorem implies that mutual cooperation can be sustained in equilibrium when players are sufficiently long-sighted.

Next, we modify the game by incorporating both private information and privacy concerns about the revelation of this information. More specifically, let the private information (types) of the individuals be their levels of patience.

---

<sup>3</sup>This is captured by the celebrated folk theorem—see, for example, Theorem 13.17 of Maschler et al. (2013).

<table border="1" style="border-collapse: collapse; width: 100%; text-align: center;"> <tr> <td></td> <td>C</td> <td>D</td> </tr> <tr> <td>C</td> <td>3, 3</td> <td>0, 4</td> </tr> <tr> <td>D</td> <td>4, 0</td> <td>2, 2</td> </tr> </table> <p>(a) Prisoner's Dilemma</p>		C	D	C	3, 3	0, 4	D	4, 0	2, 2	<table border="1" style="border-collapse: collapse; width: 100%; text-align: center;"> <tr> <td></td> <td>S</td> <td>H</td> </tr> <tr> <td>S</td> <td>2, 2</td> <td>0, 1</td> </tr> <tr> <td>H</td> <td>1, 0</td> <td>1, 1</td> </tr> </table> <p>(b) Stag Hunt</p>		S	H	S	2, 2	0, 1	H	1, 0	1, 1	<table border="1" style="border-collapse: collapse; width: 100%; text-align: center;"> <tr> <td></td> <td>L</td> <td>C</td> <td>R</td> </tr> <tr> <td>U</td> <td>0, 0</td> <td>0, 1</td> <td>1, 0</td> </tr> <tr> <td>M</td> <td>1, 0</td> <td>0, 0</td> <td>0, 1</td> </tr> <tr> <td>D</td> <td>0, 1</td> <td>1, 0</td> <td>0, 0</td> </tr> </table> <p>(c) Shapley Game</p>		L	C	R	U	0, 0	0, 1	1, 0	M	1, 0	0, 0	0, 1	D	0, 1	1, 0	0, 0
	C	D																																		
C	3, 3	0, 4																																		
D	4, 0	2, 2																																		
	S	H																																		
S	2, 2	0, 1																																		
H	1, 0	1, 1																																		
	L	C	R																																	
U	0, 0	0, 1	1, 0																																	
M	1, 0	0, 0	0, 1																																	
D	0, 1	1, 0	0, 0																																	

Figure 1: Examples of games

Each player can be one of two types—a (very) patient type, who evaluates the stream of payments from the stage game via the limit-of-means criterion, or a (very) impatient player, who only cares about the outcome in the current stage.<sup>4</sup> In addition, the players place some intrinsic value on maintaining the privacy of their types. We capture this value by extending the players' utility functions to account for how much information about their type is revealed, above and beyond what is known by the prior, from the way they play the repeated PD. One interpretation of this intrinsic value is that the players anticipate playing some future game against some unknown player, nicknamed Big Brother (BB).<sup>5</sup>

In this context, privacy protection is a technology or regulation that prevents BB from observing the players' actions and so deducing anything about their types. Under privacy protection, then, privacy concerns play no role, and players play the repeated PD game. Without privacy protection, however, players' actions are observed. As a result, BB can draw inferences about

---

<sup>4</sup>See Section 5 for a discussion of this assumption on the extreme nature of patience and impatience.

<sup>5</sup>This future game may take the form of a bargaining game, à la Rubinstein (Rubinstein, 1982), in which case they would have an advantage if they were to be perceived as patient. Alternatively, the future game could take the form of the buyer–seller bargaining model of Fudenberg and Tirole (1983), where, in equilibrium, impatient buyers obtain the goods at lower prices. Given the ambiguity of the future interaction the buyers would rather not have their type revealed in the current game.

their types, which impacts the utilities they derive from the game. Our primary question is whether society (in this case the two players) would be better off with or without privacy protection.

Perhaps counter-intuitively, if privacy has high intrinsic value—specifically, if the cost associated with the revelation of one’s type is more than one—then privacy protection turns out to be detrimental to welfare in our example. To see this note that whenever both types of a player pool on the same strategy then no information about his type is disclosed to anyone observing players’ actions. In particular, this is true if the types pool on the classic grim trigger strategy. Thus, when actions are observed the grim trigger strategy becomes an equilibrium strategy for *both* types: The long-sighted player does not deviate because he fears being punished, whereas the short-sighted player is motivated to cooperate because a deviation would reveal his type and lead to a privacy cost. In contrast, with privacy protection, when players’ actions are concealed from any third party, the short-sighted player necessarily defects in the first stage, and so his opponent must defect from the second stage onwards. Thus, privacy protection denies the players the option of enjoying the fruits of the Pareto frontier.

The PD is an example in which all Pareto-optimal payoffs can be obtained in equilibrium when there is no privacy protection, regardless of the realized types of the players. With privacy protection, however, this Pareto frontier can be attained only if both realized types are long-sighted. If one or both of the players are short-sighted, then the unique long-run payoffs are the ones associated with mutual defection. This strong result, however, is specific to the structure of the PD, and in particular to the fact that players have strictly dominated actions, which, when played by both, lead to the Pareto-optimal outcomes.

The discrepancy between the two scenarios, with and without privacy protection, in terms of the feasibility of payoffs on the Pareto frontier does not extend generally. For example, in some games such as the Stag Hunt game (see Figure 1b), the Pareto frontier may be attainable in equilibrium even

with privacy protection, regardless of the realized types of players. This holds because the game has a unique Pareto-optimal payoff that is associated with a pure Nash equilibrium: the profile in which both players choose to hunt the  $S(\text{tag})$ .

The Shapley Game illustrated in Figure 1c is an intermediate case. In this game, there is a unique mixed equilibrium whose associated payoffs are not on the Pareto frontier. Thus, when there is privacy protection and both types are short-sighted, the Pareto frontier is unattainable in equilibrium. In contrast, if one of the players is long-sighted and the other short-sighted, then some of the Pareto frontier is attainable in equilibrium (see the end of Section 4.1 for discussion and analysis of this game).

By and large, in this paper we show that privacy concerns may lead to pooling behavior of different types. Introducing privacy protection, therefore, may induce separating behavior. In the context of our repeated game, the pooling behavior will allow the players to attain payoffs on the Pareto frontier of the stage game, via the folk theorem, even if they are impatient. In contrast, the introduction of privacy protection will cause separation between patient and impatient players, which will prevent them from enjoying the Pareto frontier. Thus, there is a cost to privacy protection, even when privacy has intrinsic value.

Slightly more formally, our Theorem 1 states that when there is *no* privacy protection and privacy concerns are large enough to be meaningful, payoffs on the entire Pareto frontier can be attained in equilibria of the repeated game, regardless of players' types. In contrast, our Theorem 2 states that when there *is* privacy protection, some or all of the payoffs on the Pareto frontier cannot be attained in any equilibria of the repeated game whenever at least one of the players is short-sighted.

This statement is simple and requires no complex arguments for the PD game, and in fact it is straightforward to extend Theorem 1 to general games. The general version of Theorem 2, however, is more involved. First, it (necessarily) makes weaker guarantees than those that hold for the PD game;

and second, it requires some intricate arguments, invoking results from the Bayesian learning literature and the reputation literature.

In addition to shedding light on the interplay between privacy protection and privacy concerns, our results also contribute to the literature on repeated games. The folk theorem is central to the analysis of such games as it provides a mechanism by which players can sustain cooperation; however, it is limited to settings in which players are sufficiently long-sighted. Theorem 1 can be interpreted as showing that the presence of privacy concerns can lead to sustained cooperation even when one or both of the players may be short-sighted. Theorem 2 then shows that this channel for cooperation is hindered under privacy protection.

**Comment** We do not necessarily think of patience as the most obvious attribute that should be kept secret, nor that the literature on privacy should focus much of its attention on this issue. Rather, on the one hand, it is a reasonable setting where privacy actually matters and which is reasonably well-motivated, while, on the other hand, it serves our purpose of demonstrating the intricate connection between privacy protection and social welfare.

**Organization** The rest of the paper is organized as follows. The remainder of this section contains a review of the related literature, followed by the model in Section 2. Section 3 contains our results on equilibrium payoffs without privacy protection, and in particular shows that the entire Pareto frontier is attainable. Section 4 then follows with an analysis of payoffs when there is privacy protection and provides conditions under which the Pareto frontier is not attainable, hence demonstrating a cost to privacy protection.<sup>6</sup> Finally, Section 5 concludes the main body of the paper, and the Appendix contains most proofs and some additional results referenced throughout.

---

<sup>6</sup>Section 4 is significantly longer and more complex than Section 3, as the latter is a possibility result that involves an equilibrium construction, whereas the former is an impossibility result that shows that *no* equilibrium attains particular Pareto-optimal outcomes.

## 1.1 Related Literature

Privacy has become a central topic for study by computer scientists. Originally motivated by privacy issues related to data curation, the lion’s share of the literature focuses on a very specific notion of privacy preservation, one that is measured by *differential privacy* (see Dwork and Smith, 2010, for a survey of this literature). Underlying the notion of differential privacy are two implicit assumptions we would like to touch on. First, individuals incur an explicit cost from privacy loss, and so in many of these models we would like to strengthen the differential privacy guarantees. Second, there is a clear separation between the agents’ actions and the measure of privacy. Let us elaborate on this.

The initial literature on differential privacy was motivated by settings where the data are collected involuntarily from individuals, such as the collection of patients’ medical records by hospitals or census data collected by the government (Dinur and Nissim, 2003; Dwork and Nissim, 2004). In such models there are no strategic considerations on the part of the individuals whose data is scrutinized, and so many models asked how one can produce statistical information from the data while maintaining differential privacy guarantees.

More recently there has been a growing literature of differential privacy in models where agents voluntarily choose to share their data (possibly not in a truthful manner). These models typically take one of two forms, both of which maintain a separation between the strategic considerations and the privacy-related considerations. One strand of this literature endows agents with a utility function that does not account for privacy loss and produces mechanism that are incentive compatible while satisfying some differential privacy guarantees (McSherry and Talwar, 2007; Nissim et al., 2012b). The other strand already accounts for the disutility associated with privacy loss in the utility function. However, in these papers the differential privacy guarantee of the mechanism studied (coupled with players’ strategies) serves as a sufficient statistic for the privacy loss (e.g., Ghosh and Roth, 2015; Nissim et al., 2012a; Chen et al., 2013). In particular, the disutility associated with privacy is not a

function of the actual actions played. Thus, whereas a player’s deviation may alter his material utility, it does not change the component related to privacy, and therefore separates strategic consideration from privacy considerations.

In contrast, the literature in economics focuses on privacy in the context of information leakage and its implications on the actions and future utilities of players. Privacy *per se* has no value. This tradition goes back to Posner (1981) but is manifested in many up-to-date studies, many of which are surveyed by Acquisti et al. (2016). In particular, this can be seen through the study of privacy protection in models of repeated games. In such models, hidden actions (e.g., due to privacy protection) in early stages may provide an advantage in later stages. The roots of this approach and the implications of monitoring can be traced back to the monumental work of Aumann et al. (1995), who study an abstract zero-sum setting. In recent years this has been studied in more specific economic models such as monopolist pricing (Taylor, 2004), sequential contracting (Calzolari and Pavan, 2006), repeated signaling games (Chen et al., 2014), repeated bargaining (Hörner and Vieille, 2009; Kaya and Liu, 2015; Chaves, 2019), repeated first-price auctions (Bergemann and Hörner, 2018), and more (we refer the interested reader to the survey of Mailath and Samuelson, 2006). In all these models privacy serves as a means to an end, and players have no intrinsic value for privacy. Some show that welfare is higher when there is no privacy protection, whereas others show that privacy protection increases efficiency.

We bridge the two strands of the literature. On the one hand, we adopt the tradition laid out by the differential-privacy community and introduce an intrinsic predilection for privacy directly into the players’ utility functions. On the other hand, we fully capture the strategic implications and consider equilibrium strategies, where privacy considerations play an important role and are affected by the choice of action. One action could reveal more information about the agent than another. This amalgam is already captured in our previous work (Gradwohl and Smorodinsky, 2017), where we focus on the implications of privacy in one-shot models of decision-making. The previous

paper showed that privacy protection is often beneficial in one-shot settings with a single decision maker, whereas the current paper promotes the idea that such protection may reduce welfare in the repeated setting with multiple players.

Embedding privacy considerations introduces some conceptual modeling difficulties. In particular, a given action may be inferior to another action, but by deviating to that action the associated privacy loss may also change, and this may now render the original action superior. This interplay between actions and their privacy implications is reminiscent of the interplay between messages and actions in signaling games (see, e.g., Sobel, 2009), psychological games (Gilboa and Schmeidler, 1988; Geanakoplos et al., 1989; Battigalli and Dufwenberg, 2009), and models of social image and self image (e.g., Bernheim (1994), Glazer and Konrad (1996), and Ireland (1994) on conformity, charity, and status, respectively; Bénabou and Tirole (2006) on pro-sociality; and Becker (1974) and Rayo (2013) and Friedrichsen (2013) on self-image). Finally, a stronger connection between social image and privacy policies has been studied by Ali and Bénabou (2016). They observe that privacy measures may garble information about the moral values of society and consequently jeopardize overall welfare. We refer the reader to Gradwohl and Smorodinsky (2017) for further discussion of the similarities and distinctions across these models.

In this paper we focus on privacy protection of players' actions while assuming the players' types are necessarily private. This begs the question on the implications of privacy over types. In our model one can easily observe that if player types are not private information then it does not matter whether there is privacy of action, and regardless, we cannot guarantee outcomes on the Pareto frontier whenever short-sighted players are involved. For a related discussion on the nature of information being kept private and the resulting welfare benefits, see Prat (2005). That paper considers an agency setting and makes the distinction between "two types of information that the principal can have about his agent: information about the consequences of the agent's

action and information directly about the action.” In that paper, much in contrast with our findings, it is shown that the lack of privacy on consequences is beneficial, while the lack of privacy on action can have detrimental effects.

Finally, a distinctive feature in our model is that incomplete information is about time preference, rather than payoffs or behavioral type (as in reputation games). In different contexts, Aramendia and Wen (2020) utilize such incomplete information for equilibrium selection in repeated games, and Maor and Solan (2015) study the PD game with uncertainty about discount factors.

## 2 Model

Our model is an extension of the perception games of Gradwohl and Smorodinsky (2017) to a repeated setting. There are two players who play a game and a third party called Big Brother (BB) who may observe the interaction and draw inferences about the players’ types. The two players play an infinitely repeated game of  $G = (A_1, A_2, u_1, u_2)$ , where  $A_i$  is the finite set of player  $i$ ’s actions and  $u_i : A_1 \times A_2 \mapsto \mathbb{R}$  is player  $i$ ’s utility function, in the stage game  $G$ . For a profile  $\alpha$  of mixed actions, we also denote by  $u_i(\alpha)$  the expected utility of player  $i$  under  $\alpha$ .

Each player  $i$  is one of two types,  $t_i = S$  or  $t_i = L$ , with probabilities  $\beta_i$  and  $1 - \beta_i$ , respectively. Prior to the game, each player learns (only) his own type. A (behavioral) strategy,  $\sigma_i$ , for player  $i$  is a function  $\sigma_i : \{S, L\} \times \cup_{k=0}^{\infty} A^k \rightarrow \Delta(A_i)$ , where  $A = A_1 \times A_2$  and  $A^0 = \emptyset$ . The strategy assigns a mixed action in the stage game for each type and finite history of action tuples. A pair of strategies,  $\sigma = (\sigma_1, \sigma_2)$  induces a (random) infinite stream of payoffs. Let  $\text{supp}(\sigma) \subset A^{\infty}$  denote its support.

The types differ in their preferences over infinite streams of payoffs. A player of type  $L$  is long-sighted and evaluates material payoffs by the limit-of-means criterion. Formally, type  $L$  prefers the infinite stream of material payoffs  $\{u_i^k\}_{k=1}^{\infty}$  over  $\{v_i^k\}_{k=1}^{\infty}$  if  $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T u_i^k > \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T v_i^k$ , whenever

both limits exist. Existence is not guaranteed,<sup>7</sup> a technicality that is handled by the notion of equilibrium we will use (Definition 1 below).

In contrast, a player of type  $S$  is short-sighted and prefers the infinite stream of material payoffs  $\{u_i^k\}_{k=1}^\infty$  over  $\{v_i^k\}_{k=1}^\infty$  whenever the former is lexicographically greater than the latter, or formally, whenever there exists some  $K$  such that  $u_i^k = v_i^k$  for all  $k = 1, \dots, K$  and  $u_i^{K+1} > v_i^{K+1}$ .

In what follows we will be interested in evaluating the payoffs of both types of players, from the point of view of a social planner. For the long-sighted player, this evaluation naturally takes the form of the long-run payoff. Regarding the short-sighted player, there is no reason for the social planner to prefer one period over another, and so he can be modeled as evaluating payoffs at a random period. With an infinite horizon this is equivalent to the long-run payoff.<sup>8</sup> Thus, we will use the long-run payoff to evaluate both types' utilities, and so the following notation will be useful:

$$U_i(\sigma : t) \stackrel{\text{def}}{=} \lim_{T \rightarrow \infty} E_\sigma \left[ \frac{\sum_{k=1}^T u_i(a_k)}{T} \middle| t_i = t \right],$$

where  $a_k$  is the action profile played in stage  $k$ . We will be interested in situations in which the limit exists, and will subsequently refer to these payoffs as players' material payoffs.

In addition to these material payoffs, players incur a disutility,  $c_i = c_i(t_i, \beta'_i, \beta''_i)$ , as a bounded function of their type, the ex-ante belief about their type, and the ex-post belief about their type, respectively.<sup>9</sup> That is, players' total utility will be their material payoff minus the disutility incurred due to changes in

---

<sup>7</sup>As an example of nonexistence, suppose payoffs consist of a sequence of  $10^k$  0's, then  $10^{k+1}$  1's, then  $10^{k+2}$  0's, and so on. In this case the  $\liminf$  is 0 whereas the  $\limsup$  is 1, and so the limit does not exist.

<sup>8</sup>If we consider payoffs at a random period, then when there is a finite horizon this would equal the mean payoff. A natural extension of this criterion to the game with an infinite horizon is the limit of the finite horizon means, as we propose.

<sup>9</sup>We will assume that the long-sighted player incurs a disutility that depends on the prior and the belief after the repeated interaction. The short-sighted player can be modeled in one of two ways: either he also incurs a disutility that depends on the prior and on the belief

beliefs about their types. In order to capture the notion of privacy concerns, we impose a bit of structure on the function  $c_i$  as follows: (1)  $c_i(t_i, \beta'_i, \beta''_i) = 0$  whenever  $\beta'_i = \beta''_i$  and (2)  $c_i(L, \beta'_i, \beta''_i) \geq 0$  for all  $\beta'_i \in [0, 1]$ . In words, there is no cost if there is no change in belief and a nonnegative cost for the long-sighted type.<sup>10</sup> Finally, for convenience we will assume that  $c_i$  is bounded: that there is some large  $C$  such that  $|c_i(t_i, \beta'_i, \beta''_i)| \leq C$  for all  $t_i$ ,  $\beta'_i$ , and  $\beta''_i$ .

Our model of payoffs and equilibrium is an extension of the notion of a perception game (Gradwohl and Smorodinsky, 2017) to the current infinitely-repeated setting. To this end, we assign each player a belief function,  $\tau_i$ , which associates an ex-post belief (attributed to BB) over his type for any sequence of action tuples. Formally,  $\tau_i : \bigcup_{k=0}^{\infty} A^k \mapsto [0, 1]$  will denote an ex-post belief that player  $i$  is actually of type  $S$  at the history  $h \in \bigcup_{k=0}^{\infty} A^k$ .  $\tau_i$  is *rational* with respect to  $\sigma$  whenever  $\tau_i(h)$  is computed by Bayes' rule for any finite history  $h$  that is reached with positive probability under  $\sigma$ , and such that  $\tau_i(h) = \tau_i(h')$  if  $\tau_i(h') \in \{0, 1\}$  for some prefix  $h'$  of  $h$ . Note that if  $\tau_i$  is rational with respect to  $\sigma$ , then for an infinite history  $a^\infty \in A^\infty$  with prefixes  $a^0, \dots, a^k \dots$  that are all reached with positive probability under  $\sigma$ , the belief  $\tau_i(a^\infty) = \lim_{k \rightarrow \infty} \tau_i(a^k)$  is well-defined almost surely, by the martingale property of Bayesian updating.

We now define the notion of equilibrium, based on the formulation of Maschler et al. (2013) (Definition 13.16) but extended to incorporate perceptions and different time-discounting types.

**Definition 1** *A perfect perception equilibrium is a quadruple,  $(\sigma_1, \sigma_2, \tau_1, \tau_2)$ , such that:*

- $\tau_i$  is rational w.r.t  $\sigma$  for both  $i = 1, 2$ .

---

after the repeated interaction, or he incurs a disutility that depends on the beliefs before and after *each stage*. We will adopt the latter for simplicity; using the former, however, would not alter our results.

<sup>10</sup>The short-sighted type may incur costs at each stage, and so we do not require his costs to be nonnegative to avoid dynamic inconsistency (see Gradwohl and Smorodinsky, 2017, for discussion).

- For each  $i$ , with probability 1 according to  $\sigma|(t_i = L)$ , the limit  $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T u_i^k$  exists.
- Type  $L$  of either player  $i$  has no incentive to deviate: for any strategy  $\gamma_i$ ,

$$\begin{aligned} & E_{\sigma|(t_i=L)} \left[ \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T u_i^k \right] - E_{\sigma|(t_i=L)} c_i(L, \beta_i, \tau_i(a^\infty)) \\ & \geq E_{(\gamma_i, \sigma_{-i})} \left[ \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T u_i^k \right] - E_{(\gamma_i, \sigma_{-i})} c_i(L, \beta_i, \tau_i(a^\infty)).^{11} \end{aligned}$$

- Type  $S$  of either player  $i$  has no incentive to deviate on any finite history  $h$ : For any  $h \in \cup_{k=0}^\infty A^k$  and strategy  $\gamma_i$  that is identical to  $\sigma_i$  everywhere except at histories that have  $h$  as a prefix,

$$\begin{aligned} & E [u_i(\sigma_i(h), \sigma_{-i}(h)) | t_i = S] - E_{\sigma(S, h)} c_i(S, \tau_i(h), \tau_i((h, a))) \\ & \geq E [u_i(\gamma_i(h), \sigma_{-i}(h)) | t_i = S] - E_{(\gamma_i(h), \sigma_{-i}(h))} c_i(S, \tau_i(h), \tau_i((h, a))). \end{aligned}$$

### 3 No privacy protection

In this section we assume that there is no privacy protection, and so BB observes the players' actions and draws inferences about their types. We will show that, in this case, all Pareto-optimal payoff profiles are attainable in equilibrium. Let

$$\underline{v}_i = \min_{\alpha_{-i}} \max_{a_i} u_i(a_i, \alpha_{-i})$$

be the minimax payoff of player  $i$ , where  $\alpha_{-i}$  is any mixed action of player  $-i$ . A payoff  $v_i$  of player  $i$  is *individually rational (IR)* if  $v_i \geq \underline{v}_i$ , and *strictly IR* if the inequality is strict. Denote by

$$V_p = \{v \in \mathbb{R}^2 : v = (u_1(a), u_2(a)) \text{ for some pure profile } a\},$$

---

<sup>11</sup>The existence of the expectations for the costs is guaranteed by the boundedness of  $c_i$ —see, for example, Theorem 4 of Royden and Fitzpatrick (1988).

let its convex hull  $V = \text{CO}(V_p)$  be the *feasible* set of payoffs, and

$$V^* = \{v \in V : \forall i \ v_i \geq \underline{v}_i\}$$

be the set of feasible, IR payoffs. The standard folk theorem for two long-sighted players states that all payoff pairs in  $V^*$  are attainable as equilibrium payoffs of the repeated game (see, e.g., Theorem 13.17 of Maschler et al., 2013).

The following proposition states that these same payoffs are achievable when there is incomplete information about player types, as long as they have privacy concerns and there is no privacy protection.

**Theorem 1** *Suppose there is no privacy protection. If  $c_i(S, \beta_i, 1)$  is larger than  $\max_{a \in A, a'_i \in A_i} \{u_i(a) - u_i(a'_i, a_{-i})\}$  for both players  $i$ , then for every  $v \in V^*$  there exists a perfect perception equilibrium  $(\sigma, \tau)$  in which the long-run payoff profile is  $v$ . In particular, all feasible, IR, Pareto-optimal payoff profiles are attainable in equilibrium.*

Our result requires that the privacy cost  $c_i(S, \beta_i, 1)$  be sufficiently large. Some comments about this requirement are in order. First, note that the cost must be large relative to the *stage game* payoffs.<sup>12</sup>

Second, some lower bound on the privacy cost is clearly necessary, as a cost of 0 (or rather any amount smaller than the difference between two possible material payoffs of a player) would render privacy costs irrelevant and the claim of Theorem 1 false. Furthermore, in some games the threshold is sharp: in the PD game of Figure 1a, for example, a privacy cost less than 1 would be irrelevant, whereas a privacy cost greater than 1 would suffice to yield the claim of Theorem 1.

Finally, the analysis in Section 4 will show that with privacy protection not all Pareto-optimal payoffs are attainable in equilibrium. This, together with

---

<sup>12</sup>Long-run payoffs in our model are the limit-of-means of stage games, and so on the same order of magnitude, but our results would be unchanged if the long-run payoffs were a large multiplicative factor of the limit-of-means, in which case the privacy cost would be small when compared with the long-run payoffs.

Theorem 1, demonstrates the counter-intuitive insight that privacy protection is harmful specifically when privacy costs are large.

**Proof:** Long-sighted players play the standard grim trigger strategies, iterating over pure strategy profiles that lead to an average payoff profile of  $v$ , with a permanent deviation to the minimax strategy when a deviation of the opponent is detected. Short-sighted players pool with long-sighted players—i.e., they play the exact same strategies. Note that players use pure actions in every stage. Furthermore, for each  $i$  and  $h$ , let  $\tau_i(h) = \beta_i$  if  $h$  is reached with positive probability under  $\sigma$  or, if  $h$  is infinite, if all its prefixes are reached with positive probability under  $\sigma$ . Let  $\tau_i(h) = 1$  otherwise. Note that  $\tau$  is rational with respect to  $\sigma$ .

Observe first that long-sighted players do not gain from a deviation from  $\sigma$ : they will be punished by the other player, who plays grim trigger, and their privacy costs can only increase (since privacy costs are nonnegative). Next, suppose no player has deviated from  $\sigma$  up to stage  $k$ . Does a short-sighted player  $i$  have an incentive to deviate in stage  $k+1$ ? At the beginning of stage  $k+1$ , the belief of BB about his type is  $\beta_i$ , since types have been pooling so far. If a player deviates, beliefs are  $\tau_i = 1$ . Thus, a short-sighted player incurs cost  $c_i(S, \beta_i, 1) \geq \max_{a \in A, a' \in A_i} \{u_i(a) - u_i(a', a_{-i})\}$ . Since this cost is larger than any potential gain from the deviation, the deviation is not profitable. ■

Note that, by Theorem 1, all feasible, IR payoffs profiles are attainable in equilibrium. However, the set of attainable payoffs may be larger, and in particular, may include non-IR payoffs (see Appendix C.1).

A natural question is whether the equilibrium constructed in Theorem 1 survives common equilibrium refinements. In Appendix B we argue that our construction does, indeed, satisfy the Intuitive Criterion of Cho and Kreps (1987). However, it does so in an uninteresting way, as the definition is vacuously true. We then consider a modification of the Intuitive Criterion that has some bite in our model, and provide a strengthening of our theorem. Roughly, for a given profile of strategies  $\sigma$ , we define the notion of  $\sigma$ -intuitive beliefs. We

then show that for every  $v \in V^*$  there is a  $\sigma$  such that for *all*  $\sigma$ -intuitive beliefs  $\tau$ , the profile  $(\sigma, \tau)$  is a perfect perception equilibrium with payoff profile  $v$ .

## 4 Privacy protection

In this section we assume that there is privacy protection, and so BB does not observe the players' actions. Consequently, he cannot draw any inferences about their types, and so his beliefs  $\tau_i$  at every stage are the same, namely the prior  $\beta_i$ . Furthermore, players do not incur any privacy costs, since  $\tau_i = \beta_i$  throughout. We drop the dependence of overall utilities on perception, as these utilities are now equal to the material payoffs.

Our main result is that under privacy protection, the Pareto frontier is not generally attainable in equilibrium. For one part of this result we utilize two minor genericity assumptions on the game  $G$ :<sup>13</sup>

**Assumption 1** *There do not exist two distinct pure-action profiles  $a = (a_1, a_2)$  and  $b = (b_1, b_2)$  and a player  $i$  for which  $u_i(a) = u_i(b)$ .*

**Assumption 2** *There do not exist three distinct pure-action profiles  $a = (a_1, a_2)$ ,  $a' = (a'_1, a'_2)$ ,  $a'' = (a''_1, a''_2)$  for which the payoff pair  $u(a) = (u_1(a), u_2(a))$  is a convex combination of the payoff pairs  $u(a')$  and  $u(a'')$ .*

Before stating our main result, we need a bit more notation. First, denote the Pareto frontier by  $\text{PF}$ , where

$$\text{PF} = \{v \in V^* : \nexists v' \in V^* \text{ s.t. } v' \gg v\}.$$

Strictly speaking,  $\text{PF}$  is the Pareto frontier subject to IR constraints being satisfied.<sup>14</sup> Second, for a stage game  $G$ , denote by  $\text{NE}(G)$  the set of Nash

---

<sup>13</sup>These are needed for Lemma 4, and consequently one claim of Theorem 2. For any game, they hold with probability 1 following a perturbation of the utilities.

<sup>14</sup>An alternative definition is  $\text{PF} = \{v \in V^* : \nexists v' \in V^* \text{ s.t. } v' > v\}$  (the difference is in the inequality). If  $V^*$  is a rectangle, for example, then the question is whether the top and right segments are part of  $\text{PF}$  or just the top-right vertex. The choice of definition does not matter for our results.

equilibria of  $G$ .

We can now state our main result on the unattainability in equilibrium of some payoff profiles on the Pareto frontier under privacy protection. Since the entire Pareto frontier *is* attainable in equilibrium when there is *no* privacy protection, this theorem shows that privacy protection can be harmful.

**Theorem 2** *If  $\text{PF} \not\subseteq \text{CO}\{\text{NE}(G)\}$  and either*

- $t_1 = t_2 = S$ , or
- $G$  is symmetric and satisfies genericity assumptions 1 and 2, and  $t_i = S$  for some  $i$ ,

*then there is some  $v \in \text{PF}$  that is obtained with probability 0 in every perfect perception equilibrium of the repeated incomplete-information game. If  $\text{PF} \cap \text{NE}(G) = \emptyset$  and  $t_1 = t_2 = S$  then every  $v \in \text{PF}$  is obtained with probability 0 in every perfect perception equilibrium.*

The qualification on  $G$  is necessary: if all Pareto-optimal payoffs *are* convex combinations of payoffs attained as Nash equilibria of the stage game  $G$ , then Pareto-optimal payoffs are also attainable in equilibria of the repeated game regardless of player types, simply by iterating over various Nash equilibria of the stage game.<sup>15</sup>

In this section we prove Theorem 2. First, in Section 4.1, we describe the sets of payoffs when the types of players are commonly known and provide conditions under which these payoffs do not contain the entire Pareto frontier. This is followed by Section 4.2, in which we show that the set of payoffs when types are unknown is a subset of the payoffs when the realized types are known. Together, these results imply a cost to privacy protection described by Theorem 2: Without protection, all payoff profiles on the Pareto frontier

---

<sup>15</sup>The symmetry and genericity assumptions in the second bullet are necessary for our proof (in particular, for the proof of Lemma 4), but we do not know if they are necessary for the result to hold.

are attainable in equilibrium, by Theorem 1. With protection, in contrast, a substantial part of the Pareto frontier is unattainable.

Our analysis is tight: in Appendix D we show by construction that when there is privacy protection, the payoffs attained when types are known are also attainable in equilibrium when types are unknown.

## 4.1 Known types

Suppose the types of the players are commonly known. We will describe the set of payoff profiles attainable as equilibria of the repeated game, which, of course, depend on the types of the players. If both players are long-sighted, then the standard folk theorem applies, and the set of attainable payoffs is exactly equal to  $V^*$  (see, e.g., Theorem 13.17 of Maschler et al., 2013). We next consider the other two cases: when one player is long-sighted and the other short-sighted, and when both players are short-sighted.

**LS:** Suppose player  $i$  is long-sighted and player  $-i$  is short-sighted. What are the possible payoffs? Let

$$B'_i = \{(\alpha_1, \alpha_2) : \alpha_{-i} \in \text{BR}(\alpha_i)\},$$

where  $\text{BR}(\alpha_i)$  denotes the set of player  $-i$ 's best responses to the mixed action  $\alpha_i$  of player  $i$ , and

$$B_i = \text{CO}\{(v_i, v_{-i}) \in \mathbb{R}^2 : (v_i, v_{-i}) = (u_i(\alpha), u_{-i}(\alpha)) \text{ for some } \alpha \in B'_i\}.$$

The set of feasible payoffs is a subset of  $B_i$ :

**Lemma 1** *Let  $t_i = L$  and  $t_{-i} = S$ . Then in any equilibrium  $\sigma$  of the repeated game, the corresponding long-term payoffs  $(U_i(\sigma : L), U_{-i}(\sigma : S))$  are contained in  $B_i$ .*

**2S:** Suppose both players are short-sighted. Let

$$V^{2S} = \text{CO}\{v : v = (u_1(\alpha), u_2(\alpha)) \text{ for some NE } \alpha \text{ of } G\}.$$

Then when types are commonly-known, the feasible set of payoffs is a subset of  $V^{2S}$ :

**Lemma 2** *Let  $t_1 = t_2 = S$ . Then in any equilibrium  $\sigma$  of the repeated game, the corresponding long-term payoffs  $(U_i(\sigma : L), U_{-i}(\sigma : S))$  are contained in  $V^{2S}$ .*

**Pareto frontier** We now provide conditions under which the Pareto frontier is not attainable in equilibrium.

A first observation is that if the Pareto frontier is defined by convex combinations of Nash equilibria of the stage game—formally, if  $v \in \text{PF} \Rightarrow v \in \text{CO}\{\bar{v} : \bar{v} = (u_1(\alpha), u_2(\alpha)) \text{ for some NE } \alpha \text{ of } G\}$ —then any payoff profile on the Pareto frontier can be attained as the long-run payoff of a repeated game in which players play a NE at every stage. In this case, the Pareto frontier is attainable in equilibrium regardless of the types of the players, and in particular when both players are short-sighted.

The more interesting case is when the Pareto frontier is not defined only by the NE of the game. For this case we have two lemmas, one for general games about payoffs in  $V^{2S}$  and one for symmetric games about payoffs in  $B_i$ , that show that the entire Pareto frontier is not attainable in equilibrium.

**Lemma 3** *Fix a stage game  $G$ , and suppose  $\text{PF} \not\subseteq \text{CO}\{\text{NE}(G)\}$ . Then there exists a payoff profile  $v \in \text{PF}$  such that  $v \notin V^{2S}$ .*

That is, there is a payoff profile on the Pareto frontier that is not attainable as the long-run payoff in any equilibrium of the repeated interaction of two short-sighted players. The lemma follows almost immediately from the definition of  $V^{2S}$ , but see Appendix A for a formal proof of this and all other lemmas from this section.

Next, consider the case in which player  $i$  is long-sighted and player  $-i$  is short-sighted, and so by Lemma 1 the set of payoffs lies in  $B_i$ . The following lemma shows that in symmetric games, in which  $A_1 = A_2$  and  $u_1(a_1, a_2) = u_2(a_2, a_1)$  for all  $(a_1, a_2) \in A$ , the set  $B_i$  does not contain the entire Pareto frontier.

**Lemma 4** *Fix a symmetric stage game  $G$  that satisfies genericity assumptions 1 and 2, and suppose that  $\text{PF} \not\subseteq \text{CO}\{\text{NE}(G)\}$ . Then there exists a payoff profile  $v \in \text{PF}$  such that  $v \notin B_i$ .*

The main intuition underlying the lemma is as follows. Fix some  $v \in \text{PF}$  that cannot be attained as a NE of the stage game  $G$ , and suppose towards a contradiction that  $v \in B_i$ . Also, suppose for simplicity that the payoff  $v$  results from a pure action profile  $(a_1, a_2)$ . By assumption, in this pure action profile, player  $-i$  best-responds to player  $i$ . Now, by symmetry, if  $v = (v_1, v_2)$  is on the Pareto frontier, then so is  $v' = (v_2, v_1)$ . By a symmetric analysis,  $v'$  can only arise from the pure action profile  $(a_2, a_1)$ , in which player  $i$  best-responds to player  $-i$ . This implies that  $v$  results from a NE of  $G$ , a contradiction. Finally, we note that the more challenging part of the proof is to extend the analysis to the case in which  $v$  does not result from a pure action profile (and this latter part is the one in which we use the genericity assumptions).

Lemmas 3 and 4 showed that  $\text{PF} \not\subseteq V^{2S}$  and  $\text{PF} \not\subseteq B_i$  whenever  $\text{PF} \not\subseteq \text{CO}\{\text{NE}(G)\}$ . The following lemma shows that a stronger conclusion holds if  $\text{PF} \cap \text{NE}(G) = \emptyset$ :

**Lemma 5** *If  $\text{PF} \cap \text{NE}(G) = \emptyset$  then  $V^{2S} \cap \text{PF} = \emptyset$ .*

One might conjecture that there is a stronger lemma to be had also for the LS case: namely, that if no NE lies on the Pareto frontier, then none of the Pareto frontier can be obtained in the LS case either, as in Lemma 5 for the 2S case. However, this conjecture is false, as demonstrated by the game in Figure 1c. This game has a unique NE in which all players mix uniformly, but in which the corresponding payoff profile of  $(1/3, 1/3)$  does not lie on the

Pareto frontier. However, some part of the Pareto frontier can be obtained when the type profile is LS: in particular, if the row player is the  $L$ -type and the column player is the  $S$ -type, then the profile in which the row player mixes uniformly between his first two actions,  $U$  and  $M$ , and the column player plays his third action  $R$  with probability 1, is in  $B_i$ . Furthermore, this profile is Pareto-optimal and has payoffs  $(1/2, 1/2)$ .

## 4.2 Unknown types

In this section we return to our game of incomplete information and show that in any equilibrium of that game and for any pair of realized types, the payoff pair lies in the set of possible payoffs in the counterpart game with complete information in which the realized types are commonly known.

**Lemma 6** *In any equilibrium  $\sigma$  of the incomplete-information repeated game and almost every play path, the realized limit-of-means payoffs of the two players,  $(U_i(\sigma : t), U_{-i}(\sigma : t))$ , lie in the set of payoffs attainable in equilibria of the repeated game when the types are known.*

The proof of Lemma 6 uses a theorem about learning due to Kalai and Lehrer (1993), which states that in a repeated game of incomplete information and finitely many types players' equilibrium strategies and others' beliefs about those strategies get arbitrarily close. In our context, this implies that players eventually play almost the same as they would play if types were commonly known. Thus, the main part in the proof of the lemma is to show that payoffs are also nearly the same.

We are now ready to prove Theorem 2:

**Proof:** Lemma 2 shows that when  $t_1 = t_2 = S$ , the feasible set of payoffs lies in  $V^{2S}$  in the game with complete information, and Lemma 3 shows that when  $\text{PF} \not\subseteq \text{CO}\{\text{NE}(G)\}$  there is some  $v \in \text{PF}$  such that  $v \notin V^{2S}$ . Lemma 5 shows that when  $\text{PF} \cap \text{NE}(G) = \emptyset$  then  $V^{2S} \cap \text{PF} = \emptyset$ .

Lemma 1 shows that when  $t_i = L$  and  $t_{-i} = S$  the feasible set of payoffs lies in  $B_i$ , and Lemma 4 shows that when  $\text{PF} \not\subseteq \text{CO}\{\text{NE}(G)\}$ , the game is symmetric, and the genericity assumptions hold, there is some  $v \in \text{PF}$  such that  $v \notin B_i$ .

Finally, Lemma 6 shows that payoffs in any perfect perception equilibrium of the incomplete-information game lie in the set of payoffs in equilibrium when types are commonly-known. Together, these lemmas imply the claim of the theorem.  $\blacksquare$

## 5 Conclusion

In this paper we argued that privacy protection—in the form of disallowing third parties from observing actions—is harmful to privacy-concerned players engaged in a repeated interaction, as it limits and sometimes eliminates the possibility of attaining Pareto-optimal payoffs.

Our criticism of the potential harm of privacy protection hinges on the prevalence of the “best” equilibria in a game, those that induce payoffs on the Pareto frontier. If, on the other hand, one is concerned with the prevalence of the “worst” equilibria then a similar analysis will argue in favor of privacy protection, as it may eliminate bad equilibria (we refer the reader to Appendix C for the technical analysis).

Our results hold for more general utility functions than the ones we use, with separable additivity between material payoffs and privacy costs. In particular, they hold for any utility function  $u$  that takes as its arguments the action profile and the two beliefs, ex ante and ex post, and that satisfies the following two properties for any profile  $a$  and two beliefs  $0 < \beta' \neq \beta'' < 1$ : (i)  $u(a, \beta', \beta') = u(a, \beta'', \beta'')$ , and (ii)  $u(a, \beta', \beta'') < u(a, \beta', \beta')$ .

While our results apply to general two-player repeated games, they are limited by our assumption about player types. In particular, we assumed that players are either very long-sighted or very short-sighted. While this

assumption is largely for simplicity of exposition, we used it in the proof of Lemma 6, which showed that realized long-run payoffs of players under incomplete information lie in the set of payoffs attainable in equilibrium under their realized types. This proof uses a theorem about learning due to Kalai and Lehrer (1993). If player types were less extreme—say, discounting the future with one of two discount factors—then the *rate* of such learning would play a crucial role, significantly complicating the analysis.

However, while the proofs of our general results do not go through with less extreme player types, they do go through for specific games. For example, our insights on the benefit or harm of privacy protection in the context of the Prisoner’s Dilemma and its variants (see Figure 1a and Appendix C.1) still hold. Suppose long-sighted and short-sighted players are differentiated by different discount factors, say  $\lambda_L > \lambda_S$ , as opposed to limit-of-means versus myopic discounting. Then we can show that mutual cooperation is attainable in equilibrium regardless of the realized types when there is no privacy protection, whereas it is unattainable with privacy protection.

A more general limitation of our analysis is our assumption that a player’s type corresponds to her time discounting. One might consider more general private information, and in particular private information that affects material payoffs of the stage game, as in Gradwohl and Smorodinsky (2017), along with corresponding privacy concerns. We leave the analysis of privacy protection in such a model to future research.

## References

ABOWD, J. M. and SCHMUTTE, I. M. (2019). An economic analysis of privacy protection and statistical accuracy as social choices. *American Economic Review*, **109** 171–202.

ACQUISTI, A., TAYLOR, C. R. and WAGMAN, L. (2016). The economics of privacy. *Journal of Economic Literature*, **54** 442–492.

ALI, S. N. and BÉNABOU, R. (2016). Image versus information: Changing societal norms and optimal privacy. Tech. rep., National Bureau of Economic Research.

ARAMENDIA, M. and WEN, Q. (2020). Myopic perception in repeated games. *Games and Economic Behavior*, **119** 1–14.

AUMANN, R. J., MASCHLER, M. and STEARNS, R. E. (1995). *Repeated games with incomplete information*. MIT press.

BATTIGALLI, P. and DUFWENBERG, M. (2009). Dynamic psychological games. *Journal of Economic Theory*, **144** 1–35.

BECKER, G. S. (1974). A theory of social interactions. *Journal of Political Economy*, **82**.

BÉNABOU, R. and TIROLE, J. (2006). Incentives and prosocial behavior. *The American Economic Review* 1652–1678.

BERGEMANN, D. and HÖRNER, J. (2018). Should first-price auctions be transparent? *American Economic Journal: Microeconomics*, **10** 177–218.

BERNHEIM, B. (1994). A theory of conformity. *Journal of Political Economy* 841–877.

CALZOLARI, G. and PAVAN, A. (2006). On the optimality of privacy in sequential contracting. *Journal of Economic Theory*, **130** 168–204.

CHAVES, I. N. (2019). Privacy in bargaining: The case of endogenous entry. *Available at SSRN 3420766*.

CHEN, Y., CHONG, S., KASH, I. A., MORAN, T. and VADHAN, S. (2013). Truthful mechanisms for agents that value privacy. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. ACM, 215–232.

CHEN, Y., SHEFFET, O. and VADHAN, S. (2014). Privacy games. In *Web and Internet Economics*. Springer, 371–385.

CHO, I.-K. and KREPS, D. M. (1987). Signaling games and stable equilibria. *The Quarterly Journal of Economics*, **102** 179–221.

DINUR, I. and NISSIM, K. (2003). Revealing information while preserving privacy. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. 202–210.

DWORK, C. (2008). Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*. Springer, 1–19.

DWORK, C. and NISSIM, K. (2004). Privacy-preserving datamining on vertically partitioned databases. In *Annual International Cryptology Conference*. Springer, 528–544.

DWORK, C. and SMITH, A. (2010). Differential privacy for statistics: What we know and what we want to learn. *Journal of Privacy and Confidentiality*, **1** 2.

FRIEDRICHSEN, J. (2013). Signals sell: Designing a product line when consumers have social image concerns. Tech. rep., Discussion Paper, Social Science Research Center Berlin (WZB), Research Area Markets and Choice, Research Unit Market Behavior.

FUDENBERG, D., KREPS, D. M. and MASKIN, E. S. (1990). Repeated games with long-run and short-run players. *The Review of Economic Studies*, **57** 555–573.

FUDENBERG, D. and TIROLE, J. (1983). Sequential bargaining with incomplete information. *The Review of Economic Studies*, **50** 221–247.

FUDENBERG, D. and VILLAS-BOAS, J. M. (2006). Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems*, **1** 377–436.

GEANAKOPLOS, J., PEARCE, D. and STACCHETTI, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior*, **1** 60–79.

GHOSH, A. and ROTH, A. (2015). Selling privacy at auction. *Games and Economic Behavior*, **91** 334–346.

GILBOA, I. and SCHMEIDLER, D. (1988). Information dependent games: Can common sense be common knowledge? *Economics Letters*, **27** 215–221.

GLAZER, A. and KONRAD, K. (1996). A signaling explanation for charity. *The American Economic Review*, **86** 1019–1028.

GLICKSBERG, I. L. (1952). A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, **3** 170–174.

GRADWOHL, R. and SMORODINSKY, R. (2014). Subjective perception games and privacy. *CoRR*, *abs/1409.1487*.

GRADWOHL, R. and SMORODINSKY, R. (2017). Perception games and privacy. *Games and Economic Behavior*, **104** 293–308.

HÖRNER, J. and VIEILLE, N. (2009). Public vs. private offers in the market for lemons. *Econometrica*, **77** 29–69.

IRELAND, N. (1994). On limiting the market for status signals. *Journal of Public Economics*, **53** 91–110.

KALAI, E. and LEHRER, E. (1993). Rational learning leads to Nash equilibrium. *Econometrica*, **61** 1019–1045.

KAYA, A. and LIU, Q. (2015). Transparency and price formation. *Theoretical Economics*, **10** 341–383.

MADDEN, M., RAINIE, L., ZICKUHR, K., DUGGAN, M. and SMITH, A. (2014). Public perceptions of privacy and security in the post-snowden era. *Pew Research Center*, **12**.

MAILATH, G. J. and SAMUELSON, L. (2006). Repeated games and reputations: long-run relationships. *OUP Catalogue*.

MAOR, C. and SOLAN, E. (2015). Cooperation under incomplete information on the discount factors. *International Journal of Game Theory*, **44** 321–346.

MASCHLER, M., SOLAN, E. and ZAMIR, S. (2013). *Game Theory*. Cambridge University Press.

MC SHERRY, F. and TALWAR, K. (2007). Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*. IEEE, 94–103.

NISSIM, K., ORLANDI, C. and SMORODINSKY, R. (2012a). Privacy-aware mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 774–789.

NISSIM, K., SMORODINSKY, R. and TENNENHOLTZ, M. (2012b). Approximately optimal mechanism design via differential privacy. In *Proceedings of the 3rd innovations in theoretical computer science conference*. ACM, 203–213.

OSBORNE, M. J. and RUBINSTEIN, A. (1994). *A course in game theory*. MIT press.

POSNER, R. A. (1981). The economics of privacy. *The American Economic Review*, **71** 405–409.

PRAT, A. (2005). The wrong kind of transparency. *American economic review*, **95** 862–877.

RAINIE, L., KIESLER, S., KANG, R., MADDEN, M., DUGGAN, M., BROWN, S. and DABBISH, L. (2013). Anonymity, privacy, and security online. *Pew Research Center*, **5**.

RAYO, L. (2013). Monopolistic signal provision. *The BE Journal of Theoretical Economics*, **13** 27–58.

ROYDEN, H. L. and FITZPATRICK, P. (1988). *Real analysis*, vol. 32. Macmillan New York.

RUBINSTEIN, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society* 97–109.

SOBEL, J. (2009). Signaling games. In *Encyclopedia of Complexity and Systems Science*. Springer, 8125–8139.

TAYLOR, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics*, **35** 631–650.

# Appendix

## A Proofs from Section 4

**Proof of Lemma 1:** In any equilibrium  $\sigma$  and any history  $h$  reached with positive probability by  $\sigma$ , the corresponding mixed actions at  $h$  must belong to  $B'_i$ . Thus, the payoff profile at  $h$  belongs to  $B_i$ . Any infinite realized stream of payoffs to either player  $j$ , namely  $\{u_j^k\}_{k=1}^\infty$ , is such that  $(u_i^k, u_{-i}^k) \in B_i$ , and so  $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T (u_i^k, u_{-i}^k) \in B_i$ .<sup>16</sup> Since  $U_i(\sigma : L)$  and  $U_{-i}(\sigma : S)$  are convex combinations of such limits, they also lie in  $B_i$ .  $\blacksquare$

**Proof of Lemma 2:** In any equilibrium  $\sigma$  and any history  $h$  reached with positive probability by  $\sigma$ , the corresponding mixed actions at  $h$  must be a NE of  $G$ . Thus, the payoff profile at  $h$  belongs to  $V^{2S}$ . Any infinite realized stream of payoffs to either player  $j$ , namely  $\{u_j^k\}_{k=1}^\infty$ , is such that  $(u_i^k, u_{-i}^k) \in V^{2S}$ , and so  $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T (u_i^k, u_{-i}^k) \in V^{2S}$ . Since  $U_i(\sigma : L)$  and  $U_{-i}(\sigma : S)$  are convex combinations of such limits, they also lie in  $V^{2S}$ .  $\blacksquare$

A geometric interpretation of PF will be useful for later. Consider a plot of  $V^*$ , where player  $i$ 's payoffs are on the horizontal axis and player  $-i$ 's on the vertical axis.  $V^*$  is a convex set, whose left and bottom boundaries correspond to the minimax payoffs of players  $i$  and  $-i$ , respectively. Then PF consists of the “top-right” boundary of  $V^*$ , a connected set of line segments of nonpositive slope. Denote by  $e^{-i}$  the top-left endpoint of PF (where player  $-i$  obtains the highest feasible utility), and by  $e^i$  the bottom-right endpoint (where player  $i$  obtains the highest feasible utility). If PF is a singleton, then  $PF = \{e^i\} = \{e^{-i}\}$ . If PF is a line segment, it connects  $e^{-i}$  to  $e^i$ . Otherwise, PF consists of connected line segments: starting at  $e^{-i}$ , proceeding towards some vertex  $v_1$ , then proceeding to another vertex  $v_2$ , and so on, until  $e^i$ .

Note that the vertices  $v_k$  of PF correspond to payoffs of pure-action profiles of the stage game  $G$ . The endpoints  $e^i$  and  $e^{-i}$ , on the other hand, could

---

<sup>16</sup>Recall that the limit exists with probability 1 by Definition 1.

correspond to payoffs of pure-action profiles or could be convex combinations of payoffs of pure-action profiles. In the latter case, it must be that one of the pure-action profiles in the convex combination is not IR for one of the players.

**Proof of Lemma 3:** Any element of  $V^{2S}$  is a convex combination of the payoffs of NE of  $G$ . Since by assumption the Pareto frontier is not defined by the NE of  $G$ , there exists some  $v$  on the Pareto frontier that is not a convex combination of payoffs of NE of  $G$ . Thus,  $v \notin V^{2S}$ .  $\blacksquare$

**Proof of Lemma 4:** Let  $F = \{e^{-i}, v^1, \dots, v^k, e^i\}$  denote the set of endpoints and vertices along PF, and suppose towards a contradiction that all of PF lies in  $B_i$ . Consider the payoff profiles  $F$  that define the Pareto frontier, and note that, by assumption, not all are NE payoffs of the stage game. So there is some payoff profile  $x \in F$  that cannot be attained as the payoff pair of a NE of  $G$ . Furthermore, because  $x \in \text{PF}$  and  $\text{PF} \subseteq B_i$  it follows that  $x \in B_i$ .

We now consider two cases: that  $x \in V_p$ , the set of feasible points attainable by a pure strategy profile, and that it is not. In the former case,  $x$  is a payoff profile that can only be attained as the payoff of some pure-action profile  $a^x = (a_i^x, a_{-i}^x)$ , namely  $x = u(a^x)$ . Furthermore, since  $x \in B_i$  it must be the case that  $a^x \in B'_i$ , and so

$$a_{-i}^x \in \text{BR}(a_i^x). \quad (1)$$

Consider now the payoff profile  $\bar{x} = (x_{-i}, x_i)$ . By symmetry,  $\bar{x} \in V_p$  is a payoff profile that lies in PF that can only be attained as the payoff of the pure-action profile  $a^{\bar{x}} = (a_i^{\bar{x}}, a_{-i}^{\bar{x}}) = (a_{-i}^x, a_i^x)$ . Furthermore, since  $\bar{x} \in \text{PF}$  it must be the case that  $\bar{x} \in B_i$ . This implies that  $a^{\bar{x}} \in B'_i$ , and so  $a_{-i}^{\bar{x}} \in \text{BR}(a_i^{\bar{x}})$ , and thus

$$a_i^x \in \text{BR}(a_{-i}^x). \quad (2)$$

Combining (1) and (2) implies that  $a^x$  is a NE of the game, contradicting the assumption that  $x$  cannot be attained as the payoff profile of a NE.

We now consider the second case, that  $x$  is not an element of  $V_p$ . In particular, this means that  $x \in \{e^i, e^{-i}\}$ , since the other elements of  $F$  are

all vertices of  $V$  and so lie in  $V_p$ . Now,  $x$  is not the payoff of a pure-action profile but can be attained as the convex combination of the payoffs of two pure-action profiles. Denote these payoff profiles by  $y$  and  $z$ , where  $y = u(a^y)$  and  $z = u(a^z)$ . By genericity assumption 2,  $x$  can *only* be attained as the convex combination of these two pure profiles (and no other pure profiles). Furthermore, one element of  $\{y, z\}$  is in PF, and the other is not in PF, and thus is not IR for one of the players. Let  $z$  be the former and  $y$  the latter.

Because the game is symmetric we can assume without loss of generality that  $x = e^{-i}$ . By symmetry, the point  $\bar{x} = e^i$  is also in PF, and furthermore,  $\bar{x}$  is the convex combination of  $\bar{y}$  and  $\bar{z}$ , where  $\bar{y} = (y_{-i}, y_i)$  and  $\bar{z} = (z_{-i}, z_i)$ . Note that  $z$  and  $\bar{z}$  are pure profiles in PF, that  $y$  is not IR for player  $i$ , and that  $\bar{y}$  is not IR for player  $-i$ . As there is no point in  $B'_i$  in which  $-i$  does not best-respond to some action of player  $i$ , the payoff profile  $\bar{y} \notin B_i$ , nor is any payoff of player  $-i$  that is below  $\bar{x}_{-i}$ : these are not IR for player  $-i$ . Thus, the only way  $\bar{x} \in B_i$  is if the *mixed* profile  $a^{\bar{x}}$  satisfies  $a_{-i}^{\bar{x}} \in \text{BR}(a_i^{\bar{x}})$ . However, since  $\bar{x}$  is on the border (and not the interior) of  $V^*$ , genericity assumption 1 implies that  $a^{\bar{x}}$  is such that only one player mixes. Furthermore, since player  $-i$  is short-sighted, and one realization of the mixture yields a payoff below  $-i$ 's IR payoff, it must be player  $i$  who mixes. To conclude, the payoff profile  $\bar{x}$  can only be obtained as the payoff of some mixed action profile  $a^{\bar{x}}$  in which only player  $i$  mixes between two actions, say actions  $b$  and  $c$ . Thus,  $a^{\bar{y}} = (b, a_{-i}^{\bar{x}})$  and  $a^{\bar{z}} = (c, a_{-i}^{\bar{x}})$ , such that  $a_{-i}^{\bar{x}}$  is a best response to the mixture  $a_i^{\bar{x}}$  of  $b$  and  $c$ .

Now consider the payoff profile  $x = e^{-i}$ . Recall that  $x$  is the convex combination of  $y$  and  $z$ , each of which is the utility profile of a unique pure action profile, say  $a^y$  and  $a^z$ . By symmetry,  $a^y = (a_i^x, b)$  and  $a^z = (a_i^x, c)$ , both of which are pure-action profiles. Furthermore,  $a^y, a^z \in B'_i$ , and so both  $b, c \in \text{BR}(a_i^x)$ . But note that  $a_i^x = a_{-i}^{\bar{x}}$ . This, together with the above, implies that  $a^x$  and  $a^{\bar{x}}$  are NE of  $G$ , which is a contradiction. ■

**Proof of Lemma 5:** The set of NE of  $G$  that are in PF is empty, and so

the intersection of  $\text{CO}(\text{NE}(G))$  with  $\text{PF}$  is empty. Since  $V^{2S} = \text{CO}(\text{NE}(G))$ , it follows that  $V^{2S} \cap \text{PF} = \emptyset$ . ■

**Proof of Lemma 6:** For the realized LL case, this is almost immediate: The only limitation on  $V^*$  is that payoffs be IR, and clearly no equilibrium of the repeated game with incomplete information leads to payoffs that are not IR.

Next, fix any equilibrium profile  $\sigma$ . Let  $\sigma_i^L$  denote the strategy of player  $i$  of type  $L$ , and  $\sigma_i^S$  the strategy of player  $i$  of type  $S$ . Furthermore, denote the belief of player  $i$  about player  $-i$ 's strategy as  $\sigma^{-i}$ . Note that initially  $\sigma^{-i}$  places weight  $\beta_i$  on  $\sigma_{-i}^S$  and weight  $1 - \beta_i$  on  $\sigma_{-i}^L$ . Thus,  $\sigma^{-i}$  is absolutely continuous with respect to both  $\sigma_{-i}^L$  and  $\sigma_{-i}^S$ . By Theorem 1 of Kalai and Lehrer (1993), for every  $\varepsilon > 0$  and almost every play path  $h$ , there is a time  $T = T(h, \varepsilon)$  after which the strategy of the realized type of player  $-i$  is  $\varepsilon$ -close to the belief  $\sigma^{-i}$  that player  $i$  has about his opponent's strategy.

First, suppose the realized types are  $L$  and  $S$ . From some point on the  $S$  player  $i$  best responds to his beliefs  $\sigma^{-i}(h)$ , and his beliefs are  $\varepsilon$ -close to the actions of the opponent. That is, from some point on the strategy profile at stage  $t$  is  $\sigma_i^S(h^t) \in \text{BR}(\sigma^{-i}(h^t))$ , where  $\sigma^{-i}(h^t)$  is  $\varepsilon$ -close to  $\sigma_{-i}^L(h^t)$ . This implies that  $i$ 's payoff under  $\sigma_i^S(h^t)$  is  $\varepsilon$ -close to his payoff in some profile in  $V_i^{LS}$ . The long-sighted player's payoff is also in  $V_i^{LS}$ , since at every stage  $t$  his opponent is best-responding to  $\sigma^{-i}(h^t)$ , which is a possible mixed action of player  $-i$ .

Next, suppose the realized types are  $S$  and  $S$ . Note that in  $\sigma$ , players may not be playing a NE of the stage game in early stages, since they are best responding to their beliefs, which place positive weight on the other player being of type  $L$ . At every history  $h$ , player  $i$  is best-responding to his belief  $\sigma^{-i}(H)$  about the other's mixed action at that history. Furthermore, from stage  $T$  onward, the behavior of the other player is  $\varepsilon$ -close to  $\sigma^{-i}$ . From that point on, each player  $i$  is best responding to beliefs that are  $\varepsilon$ -close to the true

strategy of the opponent. We claim that from this stage on their payoffs in each stage must be close to the payoffs of some NE of  $G$ , which will imply that the limit-of-means payoff of the interaction is in  $V^{2S}$ .

To see this, consider some sequence of histories  $(h^t)_{t \geq T}$ . We claim that for every  $\delta > 0$  and sufficiently large  $t$ , the payoff profile  $u(\sigma_i^S(h^t), \sigma_{-i}^S(h^t))$  is within  $\delta$  of  $u(\alpha)$ , where  $\alpha$  is the convex combination of some NE of  $G$ . In fact, we claim something stronger: that all the partial limits of  $(\sigma_i^S(h^t), \sigma_{-i}^S(h^t))_{t \geq T}$  are NE of  $G$ .

For suppose towards a contradiction that this is false. Then there exists some partial limit  $(x_i, x_{-i})$ , some  $x'_i \in \Delta(A_i)$ , and some  $\gamma > 0$  such that  $u_i(x'_i, x_{-i}) > u_i(x_i, x_{-i}) + \gamma$ . Let  $(x_i^t, x_{-i}^t)_t$  be the subsequence of  $(\sigma_i^S(h^t), \sigma_{-i}^S(h^t))_{t \geq T}$  that converges to  $(x_i, x_{-i})$ . For all sufficiently large  $t$ , Theorem 1 of Kalai and Lehrer (1993) implies that  $x_i^t \in \text{BR}(y_i^t)$ , where  $y_i^t$  is  $\frac{\delta}{2(\max_a u_i(a) - \min_a u_i(a))}$ -close to  $x_{-i}^t$ . But this implies that  $x_i^t$  is a  $\delta/2$ -best response to  $x_{-i}^t$ . This contradicts the existence of the  $x'_i$  above. Thus, all the limit points are NE, and so the limit-of-means payoff is in the convex combination of all NE payoffs, namely  $V^{2S}$ . ■

## B Equilibrium refinements

In this section we consider strengthening Theorem 1 with a refinement of perfect perception equilibrium. One possibility would be to consider perfect perception equilibria that satisfy the Intuitive Criterion of Cho and Kreps (1987). Gradwohl and Smorodinsky (2014) define a variant of the Intuitive Criterion that applies to perception games, and this definition is easily amenable to our repeated setting.

In our setting, the Intuitive Criterion can be described as follows. Fix a perfect perception equilibrium  $\sigma$  and a history  $h$  reached with positive probability by  $\sigma$ . Consider a player  $i$ , type  $t_i$ , and a deviation by this type at  $h$  to an action that has probability 0 under  $\sigma$ , and let the perception following

that action place weight 1 on type  $t_i$ . Then if type  $t_i$  strictly gains from this deviation, whereas type  $t_{-i}$  weakly prefers the equilibrium strategy, then  $\sigma$  does not satisfy the Intuitive Criterion.

It is straightforward to show that the perfect perception equilibrium construction of Theorem 1 satisfies the Intuitive Criterion. However, it does so in an uninteresting way, as the definition is vacuously true: no player will ever gain from a deviation that leads to a perception that places weight 1 on his type, because of the associated privacy cost. Thus, in no perfect perception equilibrium will the Intuitive Criterion have any bite.

In the remainder of this section we consider a modification of the Intuitive Criterion that does have some bite in our model, and then provide a strengthening of our theorem that utilizes this refinement. Roughly, for a given profile of strategies  $\sigma$ , we define the notion of  $\sigma$ -intuitive beliefs. We then show that for every  $v \in V^*$  there is a  $\sigma$  such that for *all*  $\sigma$ -intuitive beliefs  $\tau$ , the profile  $(\sigma, \tau)$  is a perfect perception equilibrium with payoff profile  $v$ .

We begin with a definition of our refinement, followed by the theorem and intuition.

**Definition of refinement** Denote by

$$U_i(L, \gamma_i, \sigma_{-i}, \tau_i) \stackrel{\text{def}}{=} U_i(\gamma_i, \sigma_{-i} : L) - E_{(\gamma_i, \sigma_{-i})} c_i(L, \beta_i, \tau_i(a^\infty))$$

the utility of a long-sighted player, when strategies are  $(\gamma_i, \sigma_{-i})$  and beliefs are  $\tau_i$ . For any history  $h$ , denote by

$$U_i(S, \gamma_i, \sigma_{-i}, \tau_i, h) \stackrel{\text{def}}{=} E(u_i(\gamma_i(h), \sigma_{-i}(h)) | t_i = S) - E_{(\gamma_i(h), \sigma_{-i}(h))} c_i(S, \tau_i(h), \tau_i((h, a)))$$

the utility of a short-sighted player at  $h$ , when strategies are  $(\gamma_i, \sigma_{-i})$  and beliefs are  $\tau_i$ .

The following refinement restricts BB's beliefs on profiles for which players assign probability zero. In particular, it states that if one type of player will not gain from deviating *regardless* of the beliefs at a deviation, whereas another

type will gain by deviating given *some* belief, then the support of BB's belief must consist only of the latter type.

**Definition 2** For a given strategy profile  $\sigma$ , beliefs  $\tau$  are  $\sigma$ -intuitive if they are rational w.r.t.  $\sigma$ , and if they satisfy the following for every  $i \in \{1, 2\}$ ,  $h \in \bigcup_{k=0}^{\infty} A^k$ , and strategy  $\gamma_i$  that is identical to  $\sigma_i$  everywhere except at histories that have  $h$  as a prefix,:

- If

1.  $U_i(L, \sigma, \tau_i) \geq U_i(L, \gamma_i, \sigma_{-i}, \tau'_i)$  for every  $\tau'_i$  for which  $\tau_i(h') = \tau'_i(h')$  at every  $h'$  that is reached with positive probability under  $\sigma$ ;
2.  $U_i(S, \sigma, \tau_i, h) \leq U_i(S, \gamma_i, \sigma_{-i}, \tau'_i, h)$  for some  $\tau'_i$  for which  $\tau_i(h') = \tau'_i(h')$  at every  $h'$  that is reached with positive probability under  $\sigma$ ; and
3. at least one of the inequalities above is strict,

then  $\tau_i(h, a) = 1$  for every  $a \in \text{supp}(\gamma_i(h)) \setminus \text{supp}(\sigma_i(h))$ .

- If

1.  $U_i(S, \sigma, \tau_i, h) \geq U_i(S, \gamma_i, \sigma_{-i}, \tau'_i, h)$  for every  $\tau'_i$  for which  $\tau_i(h') = \tau'_i(h')$  at every  $h'$  that is reached with positive probability under  $\sigma$ ;
2.  $U_i(L, \sigma, \tau_i) \leq U_i(L, \gamma_i, \sigma_{-i}, \tau'_i)$  for some  $\tau'_i$  for which  $\tau_i(h') = \tau'_i(h')$  at every  $h'$  that is reached with positive probability under  $\sigma$ ; and
3. at least one of the inequalities above is strict,

then  $\tau_i(h, a) = 0$  for every  $a \in \text{supp}(\gamma_i(h)) \setminus \text{supp}(\sigma_i(h))$ .

**Strengthening of Theorem 1** We are now ready to state our theorem.

**Theorem 3** If  $c_i(S, p, 1)$  is sufficiently large, then for every  $v \in V^*$  there exists a strategy profile  $\sigma$  such that for **all**  $\sigma$ -intuitive beliefs  $\tau$ , the profile  $(\sigma, \tau)$  is a perfect perception equilibrium in which the long-run payoff profile is  $v$ . Furthermore, there exists a  $\sigma$ -intuitive  $\tau$ .

**Proof:** Let  $\sigma$  be the profile in which both types of both players play the standard grim trigger strategy leading to payoff  $v$ , as in Theorem 1. The  $L$  type will never gain by deviating, since this will lead to (long-term) punishment and a nonnegative privacy cost. For an  $S$  type, if at some history  $h$  he has a profitable deviation (for some belief), then the perception must place weight 1 on type  $S$  after that deviation at  $h$ —formally,  $\tau_i(h') = 1$  whenever  $h$  is a prefix of  $h'$ —and so the deviation will not be profitable for him with the given belief. If at some  $h$  the  $S$  type has no profitable deviation for any belief, then the refinement has no bite at that deviation, and so the perception at the deviation can be anything: there will be no profitable deviation since neither type gains from deviating at  $h$ , by assumption. ■

## C The benefits of privacy protection

In Section 4 we argued that privacy protection is harmful, as it may hinder the ability of players to obtain Pareto-optimal payoff profiles. That is, when comparing the best payoffs attainable in equilibrium, privacy protection is harmful. In this section we show that there may be benefits to privacy protection. In particular, we show that in some games privacy protection can prevent the players from obtaining suboptimal payoffs. More specifically, we show that when comparing the *worst* payoffs in equilibrium, privacy protection can be beneficial. We illustrate two distinct such benefits, both for the particular case in which the stage game  $G$  is the Prisoner’s Dilemma (PD) from Figure 1a, or a small modification thereof (although it will be clear that the ideas extend to other games as well).

### C.1 Avoiding non-IR payoffs

When there is no privacy protection, Theorem 1 states that all of  $V^*$  can be obtained in equilibrium. However, it may also be possible to obtain lower payoffs for some player. For example, suppose  $G$  is the PD, and that both

$c_1(S, \beta_i, 1)$  and  $c_1(L, \beta_i, 1)$  are high. That is, both types of player 1 incur a high cost to the belief that they are short-sighted type (whether or not this is true).

Here there is an equilibrium in which both types of player 1 always play  $C$ , while both types of player 2 always play  $D$ . This is not IR for player 1, and he obtains a low payoff of 0. However, the equilibrium can be sustained by perceptions  $\tau_1(h) = 1$  for all histories  $h$  that are not on the equilibrium path. Note that under privacy protection, such a low payoff to player 1 is impossible, as he will always obtain at least his minimax payoff of 1.

Two additional notes are in order. First, the “bad” equilibrium above does not satisfy the refinement of Definition 2. However, other (more complicated) equilibria can be constructed that do satisfy the refinement and which also lead to non-IR payoffs.

Second, equilibria with such low payoffs are not always possible, and their existence depends on the specifics of the privacy cost function. For example, in the PD, if only the short-sighted type incurs privacy costs—formally, if  $c_i(L, \beta_i, 0) \equiv 0$ —then no player will ever get non-IR payoffs in equilibrium.

## C.2 Higher minimax values

When there is no privacy protection, players get at least the minimax payoffs  $\underline{v}_i = \min_{\alpha_{-i}} \max_{a_i} u_i(a_i, \alpha_{-i})$ , regardless of whether their opponent is long-sighted or short-sighted. With privacy protection, however, a player  $i$  facing a short-sighted opponent has a different minimax payoff, namely  $\underline{v}'_i = \min_{\alpha \in B_i} \max_{a_i} u_i(a_i, \alpha_{-i})$ : this is the minimax value under the additional condition that the short-sighted player  $-i$  is best responding to player  $i$ . In some games,  $\underline{v}'_i > \underline{v}_i$ . In such games, the worst-case (over all equilibria) payoff of a player without privacy protection will be lower than her worst-case payoff with privacy protection.

A simple example of a game in which  $\underline{v}'_i > \underline{v}_i$  is the PD, but where each player has a third option  $B$ , to set off a bomb. Payoffs are such that if one

player chooses  $B$ , both players get  $-1000$ , and if both players choose  $B$ , they both get  $-1001$ . In this game the minimax payoff is  $\underline{v}_i = -1000$ , whereas the minimax payoff of a player facing a short-sighted player is  $\underline{v}'_i = 1$ . This is because choosing  $B$  is a strictly dominated action, and so an  $S$  type will never choose it. In particular, he cannot use it to threaten punishment on the opponent.

## D Construction

### D.1 Preliminaries

**LS:** Suppose player  $i$  is long-sighted and player  $-i$  is short-sighted. What are the possible payoffs? Recall that

$$B_i = \{(\alpha_1, \alpha_2) : \alpha_{-i} \in \text{BR}(\alpha_i)\}$$

is the set of feasible mixed actions, and

$$\underline{v}'_i = \min_{\alpha \in B_{-i}} \max_{a_i} u_i(a_i, \alpha_{-i})$$

the minimax payoff of player  $i$ . Next, let

$$V'_i = \text{CO}\{(v_i, v_{-i}) \in \mathbb{R}^2 : (v_i, v_{-i}) = (u_i(\alpha), u_{-i}(\alpha)) \text{ for some } \alpha \in B_{-i}\},$$

and

$$V_i^{LS} = \{(v_i, v_{-i}) \in V'_i : v_i \geq \underline{v}'_i\}.$$

When types are known, Fudenberg et al. show that all player  $i$  payoffs in  $V_i^{LS}$ , and only those, are attainable as equilibrium payoffs of the repeated game (Fudenberg et al., 1990, Proposition 5). Their proof can be extended to show that, in fact, all payoff pairs in  $V_i^{LS}$  can be attained in equilibrium:

**Lemma 7** *Suppose player  $i$  is long-sighted and player  $-i$  short-sighted. For every  $v \in V_i^{LS}$ , there exists an equilibrium strategy profile  $\sigma$  such that the long-run average payoffs of the players are  $v$ .*

**Proof:** Since  $v \in V_i^{LS}$ , there exist three mixed-action profiles  $\alpha^1, \alpha^2, \alpha^3 \in B_{-i}$  such that  $v \in \text{CO}\{u(\alpha^1), u(\alpha^2), u(\alpha^3)\}$ . Furthermore, by standard folk theorem arguments, there is an infinite sequence of alternations among the three mixed-action profiles, for which the long-run average payoffs are exactly  $v$ . Denote this sequence by  $\{s^k\}_{k=1}^{\infty}$ , where each  $s^k \in \{\alpha^1, \alpha^2, \alpha^3\}$ .

Modify the construction of  $\sigma$  from the proof of Fudenberg et al.'s Proposition 5 as follows: Whenever  $I_t \leq 0$ , let  $\sigma^t$  be the first unplayed  $s^k$  in the sequence. If  $I_t > 0$  let  $\sigma^t = m^i$ , the profile that minimaxes player  $i$ .<sup>17</sup> The remainder of the proof is the same as in Fudenberg et al. (1990).  $\blacksquare$

**2S:** Suppose both players are short-sighted. Recall that

$$V^{2S} = \text{CO}\{v : v = (u_1(\alpha), u_2(\alpha)) \text{ for some NE } \alpha \text{ of } G\}.$$

When types are known, the long-run payoffs of the repeated game for the players are in  $V^{2S}$ . That is, for any equilibrium  $\sigma$ , both  $U_i(\sigma : S) \in V^{2S}$ . Then:

**Lemma 8** *Suppose both players are short-sighted. For every  $v \in V^{2S}$ , and only such  $v$ , there exists an equilibrium strategy profile  $\sigma$  such that the long-run average payoffs of the players are  $v$ .*

**Proof:** If  $v \in V^{2S}$ , then there exist three Nash equilibrium profiles  $\alpha^1, \alpha^2, \alpha^3$  such that  $v \in \text{CO}\{u(\alpha^1), u(\alpha^2), u(\alpha^3)\}$ . By standard folk theorem arguments, there is an infinite sequence of alternations among the three mixed-action profiles, for which the long-run average payoffs are exactly  $v$ . Let  $\sigma$  be the strategy profile that alternates between these profiles in this manner.

Now suppose  $v \notin V^{2S}$ , but that there is an equilibrium  $\sigma$  with long-run payoffs equal to  $v$ . If under  $\sigma$  in every stage of the game both players play a NE of  $G$ , then  $v \in V^{2S}$ . Thus, in some stage of the game players must play a non-NE action profile. This, however, cannot be an equilibrium for two short-sighted players.  $\blacksquare$

---

<sup>17</sup>See the proof of Proposition 5 in Fudenberg et al. (1990) for details and definitions of  $I_t$  and  $m^i$ .

## D.2 The construction

We have the following construction:

**Theorem 4** *For any  $(v, v_1^{LS}, v_2^{LS}, v^{2S}) \in V^* \times V_1^{LS} \times V_2^{LS} \times V^{2S}$ , there exists an equilibrium  $\sigma = \sigma(t_1, t_2)$  such that:*

- if  $t_1 = t_2 = L$ , then  $(U_1(L, \sigma), U_2(L, \sigma)) = v$ ;
- if  $t_i = L$  and  $t_{-i} = S$ , then  $(U_i(L, \sigma), \bar{U}_{-i}(S, \sigma)) = v_i^{LS}$ ; and
- if  $t_1 = t_2 = S$ , then  $(\bar{U}_1(S, \sigma), \bar{U}_2(S, \sigma)) = v^{2S}$ .

The proof is by construction of a strategy profile that has the following structure. First, the players play a series of stage games in which player 1 “reveals” his type to player 2. Then, they play a series of stage games in which player 2 “reveals” his type to player 1. Finally, the players play folk theorem strategies corresponding to their now-commonly known types. The challenge lies in constructing the two revelation phases in such a way that they will be part of the equilibrium of the repeated game.

For each player  $i$ , we will consider three cases for the revelation phase. The first case is easiest and applies to stage games  $G$  in which player  $i$  has a dominant action. The second case applies when there is a NE of the stage game  $G$  in which player  $i$  plays a mixed action. Finally, the third and most involved case applies when neither of the first two cases does, namely, when there is no dominant action and, in all NE of  $G$ , player  $i$  plays a pure action.<sup>18</sup>

For each kind of revelation phase, we will argue that it can be part of an equilibrium of the repeated game. This requires that the  $S$  type of each player play a best response to the other player in every stage game. Additionally, the  $L$  type of each player must either best-respond in a stage game or play a suboptimal action, but can do the latter only finitely many times. Finally, in

---

<sup>18</sup>Note that the cases are not mutually exclusive, and that the latter two would suffice. We include the first for illustrative purposes and because it involves a shorter revelation phase.

order to obtain the claimed long-run payoff, the revelation phases must end with probability 1 in finitely many stages.

### D.2.1 Dominant action

Let  $a$  be a dominant action of player  $i$ . The revelation phase will last one round, in which the short-sighted type of player  $i$  plays action  $a$ , and the long-sighted type of player  $i$  plays some other action  $b \neq a$ . Both types of player  $-i$  play the same action  $c$  that is a best response to the mixed action that plays  $a$  with probability  $\pi$  and  $b$  with probability  $1 - \pi$ , where  $\pi$  is the probability that player  $i$  is the short-sighted type.

This clearly reveals player  $i$ 's type, as  $S$  and  $L$  play different actions. It is also a best response for the  $S$  type, since he plays a dominant action, and for both types of player  $-i$ , since they play a best response. Thus, this revelation phase can be part of an equilibrium of the repeated game.

### D.2.2 Mixed NE action

Suppose now that player  $i$  has no dominant action in  $G$ . Let  $\alpha$  be a NE of  $G$  in which player  $i$  plays a mixed strategy, and let  $a$  be the action played by player  $i$  with minimal but positive probability in this equilibrium. Suppose  $\alpha_i(a) = q$ . Also, suppose that the probability that player  $i$  is the short-sighted type is  $\pi$ . The type-revelation phase for player  $i$  will consist of a sequence of stage games  $G$ , where both types of player  $-i$  play  $\alpha_{-i}$ . The two types of player  $i$  play differently, as follows.

Repeat the following until the posterior  $\pi$  on type  $S$  is 0 or 1:

1. If  $\pi \leq q$ :
  - The short-sighted type of player  $i$  plays action  $a$ .

- The long-sighted type of player  $i$  plays action  $a$  such that the total (unconditional) probability of  $a$  is  $q$ , and other actions so that the distribution over actions is  $\alpha_i$ .
- If after this game player  $i$  played an action other than  $a$ , then the posterior becomes  $\pi = 0$ . Otherwise, the posterior  $\pi$  increases to  $\pi/q$ .

2. If  $\pi > q$ :

- The short-sighted type of player  $i$  plays  $a$  so that the unconditional probability of  $a$  is  $q$ , and plays the other actions with probabilities proportional to their probabilities under  $\alpha_i$ .
- The long-sighted type of player  $i$  plays all actions except  $a$  with probabilities proportional to their probabilities under  $\alpha_i$ .
- If after this game player  $i$  played action  $a$ , then the posterior on  $S$  is  $\pi = 1$ . Otherwise, the posterior  $\pi$  decreases to  $(\pi - q)/(1 - q)$ .

Note that this type-revelation phase can be part of an equilibrium of the repeated game, since both types of both players play actions that are part of a NE.

Furthermore, this phase leads to the revelation of player  $i$ 's type with probability 1. Suppose first that he is short-sighted. Whenever the players play (2) above, his type is revealed with probability  $q/\pi$ . When they play (1) above, his type will not be revealed, but the posterior on  $\pi$  increases by a factor of  $1/q$ . Within a finite number of stages, then,  $\pi$  will once again be greater than  $q$ , and they will play (2) again, and so on. Thus, player  $i$ 's type will be revealed in a finite number of rounds, and this revelation phase will end. More formally, for any  $\varepsilon > 0$ , there is a  $K$  such that after  $K$  repetitions, the probability that player  $i$ 's type will be revealed as  $S$  is at least  $1 - \varepsilon$ .

Finally, a similar argument holds when player  $i$  is long-sighted: in that case he will fully reveal his type when players play (1), and they can play (2) at most a finite number of times for each time they play (1).

### D.2.3 Pure NE action only

The last case to consider is when player  $i$  has no dominant action in  $G$  and when all NE of  $G$  are such that  $i$  plays a pure action. This case is more involved, but the basic strategy will be to construct a type-revelation phase that lasts at most  $k$  rounds. In each stage game of the phase, both types of player  $-i$  will best-respond to player  $i$ , and the  $S$ -type of player  $i$  will best-respond to player  $-i$ . The  $L$ -type of player  $i$ , however, will play a different action with some small probability, in order to allow for separation. The construction of such a strategy is closely related to the notion of a trembling-hand perfect equilibrium and requires some additional definitions.

An  $\varepsilon$ -mixed action is a mixed action that places weight at least  $\varepsilon$  on each pure action. Furthermore, recall that a trembling-hand perfect equilibrium (THPE)  $\alpha$  of a game  $G$  is a mixed-action profile such that there exists a sequence  $(\varepsilon^k)_{k \geq 0}$  that converges to 0 and a sequence  $(\alpha^k)_{k \geq 0}$  that converges to  $\alpha$ , and for which each  $\alpha^k$  is  $\varepsilon^k$ -mixed, such that for each player  $i$ , the mixed action  $\alpha_i$  is a best response to  $\alpha_{-i}^k$  for all  $k$ .

**Definition 3** *A one-sided THPE for player  $i$  is a THPE where only player  $i$  trembles. Formally, it is a mixed-action profile,  $\alpha$ , such that there exists a sequence  $(\varepsilon^k)_{k \geq 0}$  that converges to 0 and a sequence  $(\alpha_i^k)_{k \geq 0}$  that converges to  $\alpha_i$  for which the following hold:*

- each  $\alpha_i^k$  is  $\varepsilon^k$ -mixed;
- $\alpha_i$  is a best response to  $\alpha_{-i}$ ; and
- $\alpha_{-i}$  is a best response to  $\alpha_i^k$  for all  $k$ .

Note that any (one-sided) THPE is also a NE.

**Lemma 9** *In any game, there exists a one-sided THPE for player  $i$ .*

**Lemma 10** *For any sequence  $\{\varepsilon^m\}_{m \geq 0}$  that converges to 0 there exists a convergent subsequence  $\{\varepsilon^k\}_{k \geq 0}$  and a one-sided THPE  $\alpha$  for player  $i$  with a corresponding sequence  $\{\alpha_i^k\}_{k \geq 0}$ , such that each strategy  $\alpha_i^k$  is  $\varepsilon^k$ -mixed.*

**Proof:** The proof is analogous to the standard proof for the existence of a THPE (see, e.g., Proposition 249.1 in Osborne and Rubinstein, 1994), with the proper modifications for a one-sided THPE for player  $i$ . We include it here for completeness.

For each  $m$ , define the normal-game  $G_m$  to be the one in which player  $i$ 's actions are the set of all  $\varepsilon_i^m$ -mixed actions of player  $i$  in  $G$ , and player  $-i$ 's actions are all his mixed actions in  $G$ . By Glicksberg (1952), each such game has a Nash equilibrium  $\alpha^m$ . By Bolzano-Weierstrass,  $\{\alpha^m\}_{m \geq 0}$  has a convergent subsequence  $\{\alpha^k\}_{k \geq 0}$ , which converges to some  $\alpha$ . It is straightforward to verify that  $\alpha$  is a one-sided THPE for player  $i$ , with corresponding sequence  $\{\alpha_i^k\}_{k \geq 0}$  in which each  $\alpha_i^k$  is  $\varepsilon^k$ -mixed.  $\blacksquare$

Consider a NE of the stage game,  $\alpha$ , where player  $i$  plays a pure action,  $a$ . Assume that at the given stage, the prior probability that  $i$  is of type  $L$  is  $q = q(0)$ . We now construct an auxiliary strategy profile  $\alpha^q(\varepsilon)$  for any  $0 < \varepsilon \leq q$ . For player  $-i$ , let  $\alpha^q(\varepsilon)_{-i}(S) = \alpha^q(\varepsilon)_{-i}(L) = \alpha_{-i}$ . That is, both types of player  $-i$  play their equilibrium action  $\alpha_{-i}$ . On the other hand, type  $S$  of player  $i$  plays the pure action  $a$  ( $\alpha^q(\varepsilon)_i(S, a) = 1$ ), and type  $L$  mixes as follows: he plays all actions other than  $a$  with equal probabilities, and action  $a$  with some probability, such that the prior probability (not knowing the type) that  $a$  is not played equals  $\varepsilon$ . This is possible because  $0 < \varepsilon \leq q$ .

Assume we play this stage game action profile once. If  $i$  plays an action other than  $a$ , then this reveals that the type of  $i$  is  $L$ . Otherwise, if  $a$  is played, then the probability that the player is of type  $L$  decreases. Denote this posterior probability by  $q(1)$ . If it is still greater than or equal to  $\varepsilon$ , then the strategy  $\alpha^{q(1)}(\varepsilon)$  is well-defined.

We can repeat this iteratively until one of the following occurs: Either, at some stage, the action  $a$  is not played, in which case player  $-i$  learns that  $i$  is of type  $L$ . Otherwise, action  $a$  is repeatedly played until, at some stage  $k$ , the posterior probability that player  $i$  is of type  $L$ , denoted  $q(k)$ , is eventually less than  $\varepsilon$ .

**Lemma 11** *There is a sequence  $\{\varepsilon^k\}_{k \geq 0}$  converging to 0 such that the following holds for each  $k$ : if we start with the prior on  $L$  equal to  $q(0) = 1 - \beta_i$ , players play the mixed actions  $\alpha^{q(\cdot)}(\varepsilon^k)$ , and  $a$  is not played in one of the first  $k - 1$  stages, then at the beginning of the  $k$ 'th stage, the probability that player  $i$  is of type  $L$  is exactly  $\varepsilon^k$ .*

**Proof:** By construction, for each  $\varepsilon$  there is some  $k(\varepsilon)$  such that if we start with prior  $q(0) = 1 - \beta_i$ , players play the mixed actions  $\alpha^{q(\cdot)}(\varepsilon)$ , and  $a$  is not played in one of the first  $k(\varepsilon) - 1$  stages, then at the  $k(\varepsilon)^{th}$  stage the posterior of player  $i$  being of type  $L$  is  $q(k(\varepsilon) - 1) \leq \varepsilon$ . Denote by  $g(k, \varepsilon)$  the probability that player  $i$  is an  $L$ -type at the beginning of round  $k$ , given that players play the mixed actions  $\alpha^{q(\cdot)}(\varepsilon)$ . Observe that for any  $\varepsilon' < \varepsilon$ , if  $k(\varepsilon') = k(\varepsilon)$ , then  $g(k(\varepsilon'), \varepsilon') > g(k(\varepsilon), \varepsilon)$ . Furthermore,  $g(k(\varepsilon), \varepsilon)$  changes continuously with  $\varepsilon$ , conditional on  $k(\varepsilon)$  remaining fixed.

We now iteratively construct the sequence  $\{\varepsilon^k\}_{k \geq 0}$ . First, fix  $\varepsilon^0 = 1 - \beta_i$  and the corresponding  $k = 0$ . Next, suppose we have constructed the sequences for all  $k \leq \bar{k}$ , and consider the case  $k = \bar{k} + 1$ . If we decrease  $\varepsilon^{\bar{k}}$  to some  $\varepsilon < \varepsilon^{\bar{k}}$ , then in round  $\bar{k}$  we will have  $g(\bar{k}, \varepsilon) > g(\bar{k}, \varepsilon^{\bar{k}})$ . In particular  $g(\bar{k}, \varepsilon) > \varepsilon^{\bar{k}}$ , since  $g(\bar{k}, \varepsilon^{\bar{k}}) \geq \varepsilon^{\bar{k}}$ , and also  $g(\bar{k}, \varepsilon) > \varepsilon$  (since  $\varepsilon < \varepsilon^{\bar{k}}$ ). Thus, we can add another round of the stage game. As we consider smaller and smaller  $\varepsilon$ , the posterior on  $L$  in round  $k = \bar{k} + 1$  will be higher and higher. Since both  $\varepsilon$  and  $g(\bar{k} + 1, \varepsilon)$  change continuously, at some  $\varepsilon$  they will be equal. Set  $\varepsilon^{\bar{k}+1} = \varepsilon$ . ■

When a mixed action is  $\varepsilon$ -mixed, it may be the case that the mixing player plays all actions with probability strictly greater than  $\varepsilon$ . The following lemma shows that when  $\varepsilon$  is small enough this is no longer the case.

**Lemma 12** Suppose  $G$  is such that in every NE, player  $i$  plays a pure strategy. For any one-sided THPE  $\alpha$  for player  $i$  with corresponding sequence  $\{\alpha_i^k\}_{k \geq 0}$  that is  $\varepsilon^k$ -mixed for each  $k$ , the mixed action  $\alpha_i^k$  places weight exactly  $\varepsilon^k$  on all but at most one action.

**Proof:** Suppose towards a contradiction that for some  $k$ , the mixed action  $\alpha_i^k$  places weight greater than  $\varepsilon^k$  on more than one action, say actions  $a$  (the pure equilibrium action) and  $b$ . Recall from the proof of Lemma 10 that  $\alpha_i^k$  is a best-response to  $\alpha_{-i}$  out of the set of all  $\varepsilon^k$ -mixed actions. But since both  $a$  and  $b$  are played with probability greater than  $\varepsilon^k$ , it holds that both  $a$  and  $b$  are best responses of player  $i$  to  $\alpha_{-i}$ .

Now, since  $\alpha$  is a NE,  $\alpha_{-i}$  is a best-response of player  $-i$  to the pure action  $a$  of player  $i$ . By genericity assumption 1 on  $G$ ,  $\alpha_{-i}$  must also be a pure action, as there are no distinct pure action  $c$  and  $c'$  for which  $u_{-i}(a, c) = u_{-i}(a, c')$ . Thus, as  $\alpha_i^k$  is a best response to  $\alpha_{-i}$ , it is a best response to a pure action, say action  $c$ . Again by the genericity assumption, it cannot be the case that  $u_i(a, c) = u_i(b, c)$ . Thus, it is impossible for both  $a$  and  $b$  to be best responses to  $\alpha_{-i}$ , a contradiction. ■

Combining Lemmas 10, 11, and 12 yields the following. There exists a  $k$  with the following properties:

1.  $\alpha_i^k$  places weight exactly  $\varepsilon^k$  on all actions other than  $a$ , and the action  $a$  is a best response of player  $i$  to  $\alpha_{-i}$ .
2.  $\alpha_{-i}$  is a best response of player  $-i$  to  $\alpha_i^k$ .
3. If players play the mixed actions  $\alpha^{q(\cdot)}(\varepsilon^k)$ , and  $a$  is not played in one of the first  $k - 1$  stages, then at the beginning of the  $k$ 'th stage, the probability that player  $i$  is of type  $L$  is exactly  $\varepsilon^k$ . In the  $k$ 'th stage, the  $L$ -type does not play action  $a$ , and so his type will be revealed with certainty.

This is thus a type-revelation phase that lasts at most  $k$  rounds. Note that in each stage game of the phase both types of player  $-i$  best-respond to player  $i$ , and the  $S$ -type of player  $i$  best-responds to player  $-i$ . Finally, after at most  $k$  rounds, the type of player  $i$  is revealed with certainty.