



Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
<http://www.econ.upenn.edu/pier>

PIER Working Paper 07-034

“Reputation Effects”

by

George J. Mailath

<http://ssrn.com/abstract=1023658>

Reputation Effects*[†]

George J. Mailath[‡]

University of Pennsylvania

October 19, 2007

Abstract

This article gives a brief introduction to reputation effects. A canonical model is described, the reputation bound result of Fudenberg and Levine (1989, 1992) and the temporary reputation result of Cripps, Mailath, and Samuelson (2004, 2007) are stated and discussed. *Journal of Economic Literature Classification Numbers* C70, C78. **Keywords:** commitment, incomplete information, reputation bound, reputation effects.

Glossary:

action type a type of player who is committed to playing a particular action, also called a *commitment type* or *behavioral type*.

complete information characteristics of all players are common knowledge.

flow payoff stage game payoff

imperfect monitoring past actions of all players are not public information.

*Prepared for the *Encyclopedia of Complexity and System Science*, Robert A. Meyers (Editor-in-Chief), to be published by Springer.

[†]I thank Eduardo Faingold, KyungMin Kim, Antonio Penta, and Larry Samuelson for helpful comments.

[‡]Walter H. Annenberg Professor in the Social Sciences, Department of Economics, gmailath@econ.upenn.edu

incomplete information characteristics of some player are not common knowledge.

long-lived player player subject to intertemporal incentives, typically has the same horizon as length of the game.

myopic optimum an action maximizing stage game payoffs.

Nash equilibrium a strategy profile from which no player has a profitable unilateral deviation (i.e., it is self-enforcing).

Nash reversion in a repeated game, permanent play of a stage game Nash equilibrium.

normalized discounted value The discounted sum of an infinite sequence $\{a_t\}_{t \geq 0}$, calculated as $(1 - \delta) \sum_{t \geq 0} \delta^t a_t$, where $\delta \in (0, 1)$ is the discount value.

perfect monitoring past actions of all players are public information.

repeated game the finite or infinite repetition of a stage game.

reputation bound the lower bound on equilibrium payoffs of a player that the other player(s) believe may be a simple action type (typically the Stackelberg type).

short-lived player player not subject to intertemporal incentives, having a one-period horizon and so is myopically optimizing.

simple action type an action who plays the same (pure or mixed) stage-game action in every period, regardless of history.

stage game a game played in one period.

Stackelberg action in a stage game, the action a player would commit to, if that player had the chance to do so, i.e., the optimal commitment action.

Stackelberg type a simple action type that plays the Stackelberg action.

subgame in a repeated game with perfect monitoring, the game following any history.

subgame perfect equilibrium a strategy profile that induces a Nash equilibrium on every subgame of the original game.

type the characteristic of a player that is not common knowledge.

1 Definition

Repeated games have many equilibria, including the repetition of stage game Nash equilibria. At the same time, particularly when monitoring is imperfect, certain plausible outcomes are not consistent with equilibrium. *Reputation effects* is the term used for the impact upon the set of equilibria (typically of a repeated game) of perturbing the game by introducing incomplete information of a particular kind. Specifically, the characteristics of a player are not public information, and the other players believe it is possible that the distinguished player is a type that necessarily plays some action (typically the Stackelberg action). Reputation effects fall into two classes: “Plausible” phenomena that are not equilibria of the original repeated game are equilibrium phenomena in the presence of incomplete information, and “implausible” equilibria of the original game are not equilibria of the incomplete information game. As such, reputation effects provide an important qualification to the general indeterminacy of equilibria.

2 Introduction

Repeating play of a stage game often allows for equilibrium behavior inconsistent with equilibrium of that stage game. If the stage game has multiple Nash equilibrium payoffs, a large finite number of repetitions provide sufficient intertemporal incentives for behavior inconsistent with stage-game Nash equilibria to arise in some subgame perfect equilibria. However, many classic games do not have multiple Nash equilibria. For example, mutual defection DD is the unique Nash equilibrium of the prisoners’ dilemma, illustrated in Figure 1.

	C	D
C	2, 2	-1, 3
D	3, -1	0, 0

Figure 1: The prisoners’ dilemma. The cooperative action is labeled C , while defect is labeled D .

A standard argument shows that the finitely repeated prisoner’s dilemma has a unique subgame perfect equilibrium, and in this equilibrium, DD is played in every period: In any subgame perfect equilibrium, in the last period, DD must be played independently of history, since the stage game

has a unique Nash equilibrium. Then, since play in the last period is independent of history, there are no intertemporal incentives in the penultimate period, and so DD must again be played independently of history. Proceeding recursively, DD must be played in every period independently of history. (In fact, the finitely repeated prisoners' dilemma has a unique Nash equilibrium outcome, given by DD in every period.)

This contrasts with intuition, which suggests that if the prisoners' dilemma were repeated a sufficiently large (though finite) number of times, the two players would find a way to play cooperatively (C) at least in the initial stages. In response, Kreps, Milgrom, Roberts, and Wilson (1982) argued that intuition can be rescued in the finitely repeated prisoners' dilemma by introducing incomplete information. In particular, suppose each player assigns some probability to their opponent being a behavioral type who mechanically plays tit-for-tat (i.e., plays C in the first period or if the opponent had played C in the previous period, and plays D if the opponent had played D in the previous period) rather than being a rational player. No matter how small the probability, if the number of repetitions is large enough, the rational players will play C in early periods, and the fraction of periods in which CC is played is close to one.

This is the first example of a *reputation effect*: a small degree of incomplete information (of the right kind) both rescues the intuitive CC for many periods as an equilibrium outcome, and eliminates the unintuitive always DD as one. In the same issue of the *Journal of Economic Theory* containing Kreps, Milgrom, Roberts, and Wilson (1982), Kreps and Wilson (1982) and Milgrom and Roberts (1982) explored reputation effects in the finite chain store of Selten (1978), showing that intuition is again rescued, this time by introducing the possibility that the chain store is a "tough" type who always fights entry.

Reputation effects describe the impact upon the set of equilibria of the introduction of small amounts of incomplete information of a particular form into repeated games (and other dynamic games). Reputation effects fall into two classes: "Plausible" phenomena that are not equilibria of the complete information game are equilibrium phenomena in the presence of incomplete information, and "implausible" equilibria of the complete information game are not equilibria of the incomplete information game.

Reputation effects are distinct from the equilibrium phenomenon in complete information repeated games that are sometimes described as capturing *reputations*. In this latter use, an equilibrium of the complete information repeated game is selected, involving actions along the equilibrium path that are not Nash equilibria of the stage game. As usual, incentives to choose

	h	ℓ
H	2, 3	0, 2
L	3, 0	1, 1

Figure 2: The product-choice game.

these actions are created by attaching less favorable continuation paths to deviations. Players who choose the equilibrium actions are then interpreted as maintaining a reputation for doing so, with a punishment-triggering deviation interpreted as causing the loss of one’s reputation. For example, players who cooperate in the *infinitely* repeated prisoners’ dilemma are interpreted as having (or maintaining) a cooperative reputation, with any defection destroying that reputation. In this usage, the link between past behavior and expectations of future behavior is an equilibrium phenomenon, holding in some equilibria, but not in others. The notion of reputation is used to interpret an equilibrium strategy profile, but otherwise adds nothing to the formal analysis.

In contrast, the approach underlying reputation effects begins with the assumption that a player is uncertain about key aspects of her opponent. For example, player 2 may not know player 1’s payoffs, or may be uncertain about what constraints player 1 faces on his ability to choose various actions. This incomplete information is a device that introduces an intrinsic connection between past behavior and expectations of future behavior. Since incomplete information about players’ characteristics can have dramatic effects on the *set* of equilibrium payoffs, reputations in this approach do not describe certain equilibria, but rather place constraints on the set of possible equilibria.

2.1 An Example

While reputation effects were first studied in a symmetric example with two long-lived players, they arise in their purest form in infinitely repeated games with one long-lived player playing against a sequence of short-lived players. The chain store game of Selten (1978) is a finitely repeated game in which a chain store (the long-lived player) faces a finite sequence of potential entrants in its different markets. Since each entrant only cares about its own decision, it is short-lived.

Consider the “product-choice” game of Figure 2. The row player (player

1), who is long-lived, is a firm choosing between high (H) and low (L) effort, while the column player (player 2), who is short-lived, is a customer choosing between a high (h) or low (ℓ) priced product. (Mailath and Samuelson (2006) illustrate various aspects of repeated games and reputation effects using this example.) Player 2 prefers the high-priced product if the firm has exerted high effort, but prefers the low-priced product if the firm has not. The firm prefers that customers purchase the high-priced product and is willing to commit to high effort to induce that choice by the customer. In a simultaneous move game, however, the firm cannot observably choose effort before the customer chooses the product. Since high effort is costly, the firm prefers low effort, no matter the choice of the customer.

The stage game has a unique Nash equilibrium, in which the firm exerts low effort and the customer purchases the low-priced product. Suppose the game is played infinitely often, with *perfect monitoring* (i.e., the history of play is public information). The firm is *long-lived* and discounts flow profits by the discount factor $\delta \in (0, 1)$, and is *patient* if δ is close to 1. The role of the customer is taken by a succession of *short-lived* players, each of whom plays the game only once (and so myopically optimizes). It is standard to abuse language by treating the collection of short-lived players as a single myopically optimizing player.

When the firm is sufficiently patient, there is an equilibrium outcome in the repeated game in which the firm always exerts high effort and customers always purchase the high-priced product. The firm is deterred from taking the immediate myopically optimal action of low effort by the prospect of future customers then purchasing the low-priced product. Purchasing the high-priced product is a best response for the customer to high effort, so that no incentive issues arise concerning the customer's behavior. In this equilibrium, the long-lived player's payoff is 2 (the firm's payoffs are calculated as the normalized discounted sum, i.e., as the discounted sum of flow payoffs normalized by $(1 - \delta)$, so that payoffs in the infinite horizon game are comparable to flow payoffs). However, there are many other equilibria, including one in which low effort is exerted and low price purchased in every period, leading to a payoff of 1 for the long-lived player. Indeed, for $\delta \geq 1/2$, the set of pure-strategy subgame-perfect-equilibrium player 1 payoffs is given by the entire interval $[1, 2]$.

Reputation effects effectively rule out any payoff less than 2 as an equilibrium payoff for player 1. Suppose customers are not entirely certain of the characteristics of the firm. More specifically, suppose they attach high probability to the firm's being "normal," that is, having the payoffs given above, but they also entertain some (possibly very small) probability that

they face a firm who fortuitously has a technology or some other characteristic that ensures high effort. Refer to the latter as the “ H -action” type of firm. Since such a type necessarily plays H in every period, it is a type described by behavior (not payoffs), and such a type is often called a *behavioral* or *commitment* type.

This is now a game of *incomplete information*, with the customers uncertain of the firm’s type. Since the customers assign high probability to the firm being “normal,” the game is in some sense close to the game of complete information. None the less, reputation effects are present: For a sufficiently patient firm, in any Nash equilibrium of the repeated game, the firm’s payoff cannot be significantly less than 2. This result holds no matter how unlikely customers think the H -action type to be, though increasing patience is required from the normal firm as the action type becomes less likely.

The intuition behind this result is most easily seen by considering pure strategy Nash equilibria of the incomplete information game where the customers believe the firm is either the normal or the H -action type. In that case, there is no pure strategy Nash equilibrium with a payoff less than 2δ (which is clearly close to 2 for δ close to 1). In the pure strategy Nash equilibrium, either the firm always plays H , (in which case, the customers always play h and the firm’s payoff is 2), or there is a first period (say t) in which the firm plays L , revealing to future customers that he is the normal type (since the action type plays H in every period). In such an equilibrium, customers play h before t (since both types of firm are choosing H). After observing H in period t , customers conclude the firm is the H -action type. Consequently, as long as H is always chosen thereafter, customers subsequently play h (since they continue to believe the firm is the H -action type, and so necessarily plays H). An easy lower bound on the normal firm’s equilibrium payoff is then obtained by observing that the normal firm’s payoff must be at least the payoff from mimicking the action type in every period. The payoff from such behavior is at least as large as

$$\underbrace{(1 - \delta) \sum_{\tau=0}^{t-1} \delta^\tau 2}_{\text{payoff in } \tau < t \text{ from pooling with } H\text{-action type}} \quad + \quad \underbrace{(1 - \delta)\delta^t \times 0}_{\text{payoff in } t \text{ from playing } H \text{ when } L \text{ may be myopically optimal}}$$

$$+ \underbrace{(1 - \delta) \sum_{\tau=t+1}^{\infty} \delta^{\tau} 2}_{\text{payoff in } \tau > t \text{ from playing like}} \\ \text{and being treated as the } H\text{-action type}$$

$$= (1 - \delta^t)2 + \delta^{t+1}2 \\ = 2 - 2\delta^t(1 - \delta) \\ \geq 2 - 2(1 - \delta) = 2\delta.$$

The outcome in which the stage game Nash equilibrium $L\ell$ is played in every period is thus eliminated.

Since reputation effects are motivated by the hypothesis that the short-lived players are uncertain about some aspect of the long-lived player's characteristics, it is important that the results are not sensitive to the precise nature of that uncertainty. In particular, the lower bound on payoffs should not require that the short-lived players *only* assign positive probability to the normal and the H -action type (as in the game just analyzed). And it does not: The customers in the example may assign positive probability to the firm being an action type that plays H on even periods, and L on odd periods, as well as to an action type that plays H in every period before some period t' (that can depend on history), and then always plays L . Yet, as long as the customers assign positive probability to the H -action type, for a sufficiently patient firm, in any Nash equilibrium of the repeated game, the firm's payoff cannot be significantly less than 2.

Reputation effects are more powerful in the presence of imperfect monitoring. Suppose that the firm's choice of H or L is not observed by the customers. Instead, the customers observe a public signal $y \in \{\underline{y}, \bar{y}\}$ at the end of each period, where the signal \bar{y} is realized with probability $p \in (0, 1)$ if the firm chose H , and with the smaller probability $q \in (0, p)$ if the firm chose L . Interpret \bar{y} as a good meal: while customers do not observe effort, they do observe a noisy signal (the quality of the meal) of that effort, with high effort leading to a good meal with higher probability. In the game with complete information, the largest equilibrium payoff to the firm is now given by

$$\bar{v}_1 \equiv 2 - \frac{1 - p}{p - q}, \quad (1)$$

reflecting the imperfect monitoring of the firm's actions (the firm is said to be subject to binding moral hazard, see Mailath and Samuelson (2006,

section 7.6)). Since deviations from H cannot be detected for sure, there are no equilibria with the deterministic outcome path of Hh in every period. In some periods after some histories, $L\ell$ must be played in order to provide the appropriate intertemporal incentives to the firm.

As under perfect monitoring, as long as customers assign positive probability to the H -action type in the incomplete information game with imperfect monitoring, for a sufficiently patient firm, in any Nash equilibrium of the repeated game, the firm's payoff cannot be significantly less than 2 (in particular, this lower bound exceeds \bar{v}_1). Thus, in this case, reputation effects provide an intuitive lower bound on equilibrium payoffs that both rules out "bad" equilibrium payoffs, as well as rescues outcomes in which Hh occurs in most periods.

Proving that a reputation bound holds in the imperfect monitoring case is considerably more involved than in the perfect monitoring case. In perfect-monitoring games, it is only necessary to analyze the evolution of the customers' beliefs when always observing H , the action of the H -action type. In contrast, imperfect monitoring requires consideration of belief evolution on all histories that arise with positive probability.

None the less, the intuition is the same: Consider a putative equilibrium in which the normal firm receives a payoff less than $2 - \varepsilon$. Then the normal and action types must be making different choices over the course of the repeated game, since an equilibrium in which they behave identically would induce customers to choose h and would yield a payoff of 2. As in the perfect monitoring case, the normal firm has the option of mimicking the behavior of the H -action type. Suppose the normal firm does so. Since the customers expect the normal type of firm to behave differently from the H -action type, they will more often see signals indicative of the H -action type (rather than the normal type), and so must eventually become convinced that the firm is the H -action type. Hence, in response to this deviation, the customers will eventually play their best response to H of h . While "eventually" may take a while, that time is independent of the equilibrium (indeed of the discount factor), depending only on the imperfection in the monitoring and the prior probability assigned to the H -action type. Then, if the firm is sufficiently patient, the payoff from mimicking the H -action type is arbitrarily close to 2, contradicting the existence of an equilibrium in which the firm's payoff fell short of $2 - \varepsilon$.

At the same time, because monitoring is imperfect, as discussed in Section 3.5, the reputation effects are necessarily transient. Under general conditions in imperfect-monitoring games, the incomplete information that is at the core of reputation effects is a short-run phenomenon. Player 2 must

eventually come to learn player 1's type and continuation play must converge to an equilibrium of the complete information game.

Reputation effects arise for very general specifications of the incomplete information as long as the customers assign strictly positive probability to the H -action type. It is critical, however, that the customers do assign strictly positive probability to the H -action type. For example, in the product-choice game, the set of Nash equilibria of the repeated game is not significantly impacted by the possibility that the firm is either normal or the L -action type only. While reputation effects per se do not arise from the L -action type, it is still of interest to investigate the impact of such uncertainty on behavior using stronger equilibrium notions, such as Markov perfection (see Mailath and Samuelson (2001)).

3 A Canonical Model

3.1 The Stage Game

The stage game is a two-player simultaneous-move finite game of public monitoring. Player i has action set A_i , $i = 1, 2$. Pure actions for player i are denoted by $a_i \in A_i$, and mixed actions are denoted by $\alpha_i \in \Delta(A_i)$, where $\Delta(A_i)$ is the set of probability distributions over A_i . Player 2's actions are public, while player 1's are potentially private. The public signal of player 1's action, denoted by y is drawn from a finite set Y , with the probability that y is realized under the pure action profile $a \in A \equiv A_1 \times A_2$ denoted by $\rho(y | a)$. Player 1's ex post payoff from the action profile a and signal realization y is $r_1(y, a)$, and so the ex ante (or expected) flow payoff is $u_1(a) \equiv \sum_y r_1(y, a)\rho(y | a)$. Player 2's ex post payoff from the action profile a and signal realization y is $r_2(y, a_2)$, and so the ex ante (or expected) flow payoff is $u_2(a) \equiv \sum_y r_2(y, a_2)\rho(y | a)$. Since player 2's ex post payoff is independent of player 1's actions, player 1's actions only affect player 2's payoffs through the impact on the distribution of the signals and so on ex ante payoffs. While the ex post payoffs r_i play no explicit role in the analysis, they justify the informational assumptions to be made. In particular, the model requires that histories of signals and past actions are the only information players receive, and so it is important that stage game payoffs u_i are not informative about the action choice (and this is the critical feature delivered by the assumptions that ex ante payoffs are not observable and that payer 2's ex post payoffs do *not* depend on a_1).

Perfect monitoring is the special case where $Y = A_1$ and $\rho(y | a) = 1$ if $y = a_1$, and 0 otherwise.

The results in this section hold under significantly weaker monitoring assumptions. In particular, it is not necessary that the actions of player 2 be public. If these are also imperfectly monitored, then the ex post payoff for player 1 is independent of player 2 actions. Since player 2 is short-lived, when player 2 are not public, it is then natural to also assume that the period t player 2 does not know earlier player 2's actions.

3.2 The Complete Information Repeated Game

The stage game is infinitely repeated. Player 1 is long-lived, with payoffs given by the normalized discounted value $(1 - \delta) \sum_{t=0}^{\infty} \delta^t u_1^t$, where $\delta \in (0, 1)$ is the discount factor and u_1^t is player 1's period t flow payoff. Player 1 is *patient* if δ is close to 1. As in our example, the role of player 2 is taken by a succession of *short-lived* players, each of whom plays the game only once (and so myopically optimizes).

Player 1's set of private histories is $\mathcal{H}_1 \equiv \cup_{t=0}^{\infty} (Y \times A)^t$ and the set of public histories (which coincides with the set of player 2's histories) is $\mathcal{H} \equiv \cup_{t=0}^{\infty} (Y \times A_2)^t$. If the game has perfect monitoring, histories $h = (y^0, a^0; y^1, a^1; \dots; y^{t-1}, a^{t-1})$ in which $y^\tau \neq a_1^\tau$ for some $\tau \leq t - 1$ arise with zero probability, independently of behavior, and so can be ignored. A strategy σ_1 for player 1 specifies a probability distribution over 1's pure action set for each possible private history, i.e., $\sigma_1 : \mathcal{H}_1 \rightarrow \Delta(A_1)$. A strategy σ_2 for player 2 specifies a probability distribution over 2's pure action set for each possible public history, i.e., $\sigma_2 : \mathcal{H} \rightarrow \Delta(A_2)$.

Definition 1 *The strategy profile (σ_1^*, σ_2^*) is a Nash equilibrium if*

1. *there does not exist a strategy σ_1 yielding a strictly higher payoff for player 1 when player 2 plays σ_2^* , and*
2. *in all periods t , after any history $h^t \in \mathcal{H}$ arising with positive probability under (σ_1^*, σ_2^*) , $\sigma_2^*(h^t)$ maximizes $E[u_2(\sigma_1^*(h_1^t), a_1) \mid h^t]$, where the expectation is taken over the period t -private histories that player 1 may have observed.*

3.3 The Incomplete Information Repeated Game

In the incomplete information game, the type of player 1 is unknown to player 2. A possible type of player 1 is denoted by $\xi \in \Xi$, where Ξ is a finite or countable set (see Fudenberg and Levine (1992) for the uncountable case). Player 2's prior belief about 1's type is given by the distribution μ ,

with support Ξ . The set of types is partitioned into a set of *payoff types* Ξ_1 , and a set of *action types* $\Xi_2 \equiv \Xi \setminus \Xi_1$. Payoff types maximize the average discounted value of payoffs, which depend on their type and which may be nonstationary,

$$u_1 : A_1 \times A_2 \times \Xi_1 \times \mathbb{N}_0 \rightarrow \mathbb{R}.$$

Type $\xi_0 \in \Xi_1$ is the *normal type* of player 1, who happens to have a stationary payoff function, given by the stage game in the benchmark game of complete information,

$$u_1(a, \xi_0, t) = u_1(a) \quad \forall a \in A, \forall t \in \mathbb{N}_0.$$

It is standard to think of the prior probability $\mu(\xi_0)$ as being relatively large, so the games of incomplete information are a seemingly small departure from the underlying game of complete information, though there is no requirement that this be the case.

Action types (also called *commitment* or *behavioral types*) do not have payoffs, and simply play a specified repeated game strategy. For any repeated-game strategy from the complete information game, $\hat{\sigma}_1 : \mathcal{H}_1 \rightarrow \Delta(A_1)$, denote by $\xi(\hat{\sigma}_1)$ the action type committed to the strategy $\hat{\sigma}_1$. In general, a commitment type of player 1 can be committed to any strategy in the repeated game. If the strategy in question plays the same (pure or mixed) stage-game action in every period, regardless of history, that type is called a *simple action type*. For example, the *H-action type* in the product-choice game is a simple action type. The (simple action) type that plays the pure action a_1 in every period is denoted by $\xi(a_1)$ and similarly the simple action type committed to $\alpha_1 \in \Delta(A_1)$ is denoted by $\xi(\alpha_1)$. As will be seen soon, allowing for mixed action types is an important generalization from simple pure types.

A strategy for player 1, also denoted by $\sigma_1 : \mathcal{H}_1 \times \Xi \rightarrow \Delta(A_1)$, specifies for each type $\xi \in \Xi$ a repeated game strategy such that for all $\xi(\hat{\sigma}_1) \in \Xi_2$, the strategy $\hat{\sigma}_1$ is specified. A strategy σ_2 for player 2 is as in the complete information game, i.e., $\sigma_2 : \mathcal{H} \rightarrow \Delta(A_2)$.

Definition 2 *The strategy profile (σ_1^*, σ_2^*) is a Nash equilibrium of the incomplete information game if*

1. *for all $\xi \in \Xi_1$, there does not exist a repeated game strategy σ_1 yielding a strictly higher payoff for payoff type ξ of player 1 when player 2 plays σ_2^* , and*

2. in all periods t , after any history $h^t \in \mathcal{H}$ arising with positive probability under (σ_1^*, σ_2^*) and μ , $\sigma_2^*(h^t)$ maximizes $E[u_2(\sigma_1^*(h^t, \xi), a_1) \mid h^t]$, where the expectation is taken over both the period t -private histories that player 1 may have observed and player 1's type.

Example 1 Consider the product-choice game (Figure 2) under perfect monitoring. The firm is willing to commit to H to induce h from customers. This incentive to commit is best illustrated by considering a sequential version of the product-choice game: The firm first publicly commits to an effort, and then the customer chooses between h and ℓ , knowing the firm's choice. In this sequential game, the firm chooses H in the unique subgame perfect equilibrium. Since Stackelberg (1934) was the first investigation of such leader-follower interactions, it is traditional to call H the Stackelberg action, and the H -action type of player 1 the Stackelberg type, with associated Stackelberg payoff 2. Suppose $\Xi = \{\xi_0, \xi(H), \xi(L)\}$. For $\delta \geq 1/2$, the grim trigger strategy profile of always playing Hh , with deviations punished by Nash reversion, is a subgame perfect equilibrium of the complete information game. Consider the following adaptation of this profile in the incomplete information game:

$$\sigma_1(h^t, \xi) = \begin{cases} H, & \text{if } \xi = \xi(H), \\ & \text{or } \xi = \xi_0 \text{ and } a^\tau = Hh \text{ for all } \tau < t, \\ L, & \text{otherwise,} \end{cases}$$

and

$$\sigma_2(h^t) = \begin{cases} h, & \text{if } a^\tau = Hh \text{ for all } \tau < t, \\ \ell, & \text{otherwise.} \end{cases}$$

In other words, player 2 and the normal type of player 1 follow the strategies from the Nash-reversion equilibrium in the complete information game, and the action types $\xi(H)$ and $\xi(L)$ play their actions.

This is a Nash equilibrium for $\delta \geq 1/2$ and $\mu(\xi(L)) < 1/2$. The restriction on $\mu(\xi(L))$ ensures that player 2 finds h optimal in period 0. Should player 2 ever observe L , then Bayes' rule causes her to place probability 1 on type $\xi(L)$ (if L is observed in the first period) or the normal type (if L is first played in a subsequent period), making her participation in Nash reversion optimal. The restriction on δ ensures that Nash reversion provides sufficient incentive to make H optimal for the normal player 1. After observing $a_1^0 = H$ in period 0, player 2 assigns zero probability to $\xi = \xi(L)$.

However, the posterior probability that 2 assigns to the Stackelberg type does not converge to 1. In period 0, the prior probability is $\mu(\xi(H))$. After one observation of H , the posterior increases to $\mu(\xi^*)/[\mu(\xi^*) + \mu(\xi_0)]$, after which it is constant. By stipulating that an observation of H in a history in which L has previously been observed causes player 2 to place probability one on the normal type of player 1, a specification of player 2's beliefs that is consistent with sequentiality is obtained.

As seen in the introduction, for δ close to 1, $\sigma_1(h^t, \xi_0) = L$ for all h^t is not part of any Nash equilibrium. ★

3.4 The Reputation Bound

Which type would the normal type most like to be treated as? Player 1's *pure-action Stackelberg payoff* is defined as

$$v_1^* = \sup_{a_1 \in A_1} \min_{\alpha_2 \in B(a_1)} u_1(a_1, \alpha_2), \quad (2)$$

where $B(a_1) = \arg \max_{a_2} u_2(a_1, a_2)$ is the set of player 2 myopic best replies to a_1 . If the supremum is achieved by some action a_1^* , that action is an associated Stackelberg action,

$$a_1^* \in \arg \max_{a_1 \in A_1} \min_{\alpha_2 \in B(a_1)} u_1(a_1, \alpha_2).$$

This is a pure action to which player 1 would commit, if player 1 had the chance to do so (and hence the name “Stackelberg” action, see the discussion in Example 1), given that such a commitment induces a best response from player 2. If there is more than one such action for player 1, the action can be chosen arbitrarily.

However, player 1 would typically prefer to commit to a mixed action. In the product-choice game, for example, a commitment by player 1 to mixing between H and L , with slightly larger probability on H , still induces player 2 to choose h and gives player 1 a larger payoff than a commitment to H . Define the mixed-action Stackelberg payoff as

$$v_1^{**} \equiv \sup_{\alpha_1 \in \Delta(A_1)} \min_{\alpha_2 \in B(\alpha_1)} u_1(\alpha_1, \alpha_2), \quad (3)$$

where $B(\alpha_1) = \arg \max_{a_2} u_2(\alpha_1, a_2)$ is the set of player 2's best responses to α_1 . In the product-choice game, $v_1^* = 2$, while $v_1^{**} = 5/2$. Typically, the

supremum is not achieved by any mixed action, and so there is no mixed-action Stackelberg type. However, there are mixed action types that, if player 2 is convinced she is facing such a type, will yield payoffs arbitrarily close to the mixed-action Stackelberg payoff.

As with imperfect monitoring, simple mixed action types under perfect monitoring raise issues of monitoring, since a deviation by the normal type from the distribution α_1 of a mixed action type $\xi(\alpha_1)$, to some action in the support cannot be detected. However, when monitoring of the pure actions is perfect, it is possible to statistically detect deviations, and this will be enough to imply the appropriate reputation lower bound.

When monitoring is imperfect, the public signals are statistically informative about the actions of the long-lived player under the next assumption (Lemma 1).

Assumption 1 *For all $a_2 \in A_2$, the collection of probability distributions $\{\rho(y | (a_1, a_2)) : a_1 \in A_1\}$ is linearly independent.*

This assumption is trivially satisfied in the perfect monitoring case. Reputation effects still exist when this assumption fails, but the bounds are more complicated to calculate (see Fudenberg and Levine (1992) or Mailath and Samuelson (2006, Section 15.4.1)).

Fixing an action for player 2, a_2 , the mixed action α_1 implies the signal distribution $\sum_{a_1} \rho(y | (a_1, a_2))\alpha_1(a_1)$.

Lemma 1 *Suppose ρ satisfies Assumption 1. Then, if for some a_2 ,*

$$\sum_{a_1} \rho(y | (a_1, a_2))\alpha_1(a_1) = \sum_{a_1} \rho(y | (a_1, a_2))\alpha'_1(a_1), \quad \forall y, \quad (4)$$

then $\alpha_1 = \alpha'_1$.

Proof. Suppose (4) holds for some a_2 . Let R denote the $|Y| \times |A_1|$ matrix whose y - a_1 element is given by $\rho(y | (a_1, a_2))$ (so that the a_1 -column is the probability distribution on Y implied by the action profile $a_1 a_2$). Then, (4) can be written as $R\alpha_1 = R\alpha'_1$, or more simply as $R(\alpha_1 - \alpha'_1) = 0$. By Assumption 1, R has full column rank, and so $x = 0$ is the only vector $x \in \mathbb{R}^{|A_1|}$ solving $Rx = 0$. ■

Consequently, if player 2 believes that the long-lived player's behavior implies a distribution over the signals close to the distribution implied by some particular action α'_1 , then player 2 must believe that the long-lived player's action is also close to α'_1 . Since A_2 is finite, this then implies

that when player 2 is best responding to some belief about the long-lived player's behavior implying a distribution over signals sufficiently close to the distribution implied by α'_1 , then player 2 is in fact best responding to α'_1 .

We are now in a position to state the main reputation bound result. Let $v_1(\xi_0, \mu, \delta)$ be the infimum over the set of the normal player 1's payoffs in any (pure or mixed) Nash equilibrium in the incomplete information repeated game, given the distribution μ over types and the discount factor δ .

Proposition 1 (Fudenberg and Levine (1989, 1992)) *Suppose ρ satisfies Assumption 1 and let $\hat{\xi}$ denote the simple action type that always plays $\hat{\alpha}_1 \in \Delta(A_1)$. Suppose $\mu(\xi_0), \mu(\hat{\xi}) > 0$. For every $\eta > 0$, there is a value K such that for all δ ,*

$$v_1(\xi_0, \mu, \delta) \geq (1 - \eta)\delta^K \min_{\alpha_2 \in B(\hat{\alpha}_1)} u_1(\hat{\alpha}_1, \alpha_2) + (1 - (1 - \eta)\delta^K) \min_{a \in A} u_1(a). \quad (5)$$

This immediately yields the pure action Stackelberg reputation bound. Fix $\varepsilon > 0$. Taking $\hat{\alpha}_1$ in the proposition as the degenerate mixture that plays the Stackelberg action a_1^* with probability 1, equation (5) becomes

$$\begin{aligned} v_1(\xi_0, \mu, \delta) &\geq (1 - \eta)\delta^K v_1^* + (1 - (1 - \eta)\delta^K) \min_{a \in A} u_1(a) \\ &\geq v_1^* - (1 - (1 - \eta)\delta^K)2M, \end{aligned}$$

where $M \equiv \max_{a_1} |u_1(a)|$. This last expression is at least as large as $v_1^* - \varepsilon$ when $\eta < \varepsilon/(2M)$ and δ is sufficiently close to 1.

The mixed action Stackelberg reputation bound is also covered:

Corollary 1 *Suppose ρ satisfies Assumption 1 and μ assigns positive probability to some sequence of simple types $\{\xi(\alpha_1^k)\}_{k=1}^\infty$ with each α_1^k in $\Delta(A_1)$ satisfying*

$$v_1^{**} = \lim_{k \rightarrow \infty} \min_{\alpha_2 \in B(\alpha_1^k)} u_1(\alpha_1^k, \alpha_2).$$

For all $\varepsilon' > 0$, there exists $\underline{\delta} < 1$ such that for all $\delta \in (\underline{\delta}, 1)$,

$$v_1(\xi_0, \mu, \delta) \geq v_1^{**} - \varepsilon'.$$

The remainder of this subsection outlines a proof of Proposition 1. Fix a strategy profile (σ_1, σ_2) (which may be Nash, but at this point of the discussion, need not be). The beliefs μ then induce a probability distribution \mathbf{P} on the set of outcomes, which is the set of possible infinite histories (denoted by h^∞) and realized types, $(Y \times A)^\infty \times \Xi \equiv \Omega$. The probability

measure \mathbf{P} describes how the short-lived players believe the game will evolve, given their prior beliefs μ about the types of the long-lived player. Let $\hat{\mathbf{P}}$ denote the probability distribution on the set of outcomes induced by (σ_1, σ_2) and the action type $\hat{\xi}$. The probability measure $\tilde{\mathbf{P}}$ describes how the short-lived players believe the game will evolve if the long-lived player's type is $\hat{\xi}$. Finally, let $\tilde{\mathbf{P}}$ denote the probability distribution on the set of outcomes induced by (σ_1, σ_2) conditioning on the long-lived player's type not being the action type $\hat{\xi}$. Then, $\mathbf{P} \equiv \hat{\mu}\tilde{\mathbf{P}} + (1 - \hat{\mu})\tilde{\mathbf{P}}$, where $\hat{\mu} \equiv \mu(\hat{\xi})$.

The discussion after Lemma 1 implies that the optimal behavior of the short-lived player in period t is determined by that player's beliefs over the signal realizations in that period. These beliefs can be viewed as a *one-step ahead prediction* of the signal y that will be realized conditional on the history h^t , $\mathbf{P}(y | h^t)$. Let $\hat{\mu}^t(h^t) = \mathbf{P}(\hat{\xi} | h^t)$ denote the posterior probability after observing h^t that the short-lived player assigns to the long-lived player having type $\hat{\xi}$. Note also that if the long-lived player is the action type $\hat{\xi}$, then the true probability of the signal y is $\hat{\mathbf{P}}(y | h^t) = \rho(y | (H, \sigma_2(h^t)))$. Then,

$$\mathbf{P}(y | h^t) = \hat{\mu}^t(h^t)\hat{\mathbf{P}}(y | h^t) + (1 - \hat{\mu}^t(h^t))\tilde{\mathbf{P}}(y | h^t).$$

The key step in the proof of Proposition 1 is a statistical result on merging. The following lemma essentially says that the short-lived players cannot be surprised too many times. Note first that an infinite public history h^∞ can be thought of a sequence of ever longer finite public histories h^t . Consider the collection of infinite public histories with the property that player 2 often sees histories h^t that lead to very different one-step ahead predictions about the signals under $\tilde{\mathbf{P}}$ and under $\hat{\mathbf{P}}$ and have a “low” posterior that the long-lived player is $\hat{\xi}$. The lemma asserts that if the long-lived player is in fact the action type $\hat{\xi}$, this collection of infinite public histories has low probability. Seeing the signals more likely under $\hat{\xi}$ leads the short-lived players to increase the posterior probability on $\hat{\xi}$. The posterior probability fails to converge to 1 under $\hat{\mathbf{P}}$ only if the play of the types different from $\hat{\xi}$ leads, on average, to a signal distribution similar to that implied by $\hat{\xi}$. For the purely statistical statement and its proof, see Mailath and Samuelson (2006, Section 15.4.2).

Lemma 2 *For all $\eta, \psi > 0$ and $\mu^\dagger \in (0, 1]$, there exists a positive integer K such that for all $\mu(\hat{\xi}) \in [\mu^\dagger, 1]$, for every strategy $\sigma_1 : \mathcal{H}_1 \times \Xi \rightarrow \Delta(A_1)$ and $\sigma_2 : \mathcal{H} \rightarrow \Delta(A_2)$,*

$$\hat{\mathbf{P}} \left(h^\infty : \left| \left\{ t \geq 1 : (1 - \hat{\mu}^t(h^t)) \max_y \left| \tilde{\mathbf{P}}(y | h^t) - \hat{\mathbf{P}}(y | h^t) \right| \geq \psi \right\} \right| \geq K \right)$$

$$\leq \eta. \quad (6)$$

Note that the bound K holds for all strategy profiles (σ_1, σ_2) and all prior probabilities $\mu(\hat{\xi}) \in [\mu^\dagger, 1)$. This allows us to bound equilibrium payoffs.

Proof of Proposition 1.

Fix $\eta > 0$. From Lemma 1, by choosing ψ sufficiently small in Lemma 2, with $\hat{\mathbf{P}}$ -probability at least $1 - \eta$, there are at most K periods in which the short-lived players are not best responding to $\hat{\alpha}_1$.

Since a deviation by the long-lived player to the simple strategy of always playing $\hat{\alpha}_1$ induces the same distribution on public histories as $\hat{\mathbf{P}}$, the long-lived player's expected payoff from such a deviation is bounded below by the right side of (5). ■

3.5 Temporary Reputation Effects

Under perfect monitoring, there are often pooling equilibria in which the normal and some action type of player 1 behave identically on the equilibrium path (as in Example 1). Deviations on the part of the normal player 1 are deterred by the prospect of the resulting punishment. Under imperfect monitoring, such pooling equilibria do not exist. The normal and action types may play identically for a long period of time, but the normal type always eventually has an incentive to cheat at least a little on the commitment strategy, contradicting player 2's belief that player 1 will exhibit commitment behavior. Player 2 must then eventually learn player 1's type.

In addition to Assumption 1, disappearing reputation effects require full support monitoring.

Assumption 2 *For all $a \in A$, $y \in Y$, $\rho(y | a) > 0$.*

This assumption implies that Bayes' rule determines the beliefs of player 2 about the type of player 1 after *all* histories.

Suppose there are only two types of player 1, the normal type ξ_0 and a simple action type $\hat{\xi}$, where $\hat{\xi} = \xi(\hat{\alpha}_1)$ for some $\hat{\alpha}_1 \in \Delta(A_1)$. The analysis is extended to many commitment types in Cripps, Mailath, and Samuelson (2004, Section 6.1). It is convenient to denote a strategy for player 1 as a pair of functions $\tilde{\sigma}_1$ and $\hat{\sigma}_1$ (so $\hat{\sigma}_1(h_1^t) = \hat{\alpha}_1$ for all $h_1^t \in \mathcal{H}_1$), the former for the normal type and the latter for the action type.

Recall that $\mathbf{P} \in \Delta(\Omega)$ is the unconditional probability measure induced by the prior μ , and the strategy profile $(\hat{\sigma}_1, \tilde{\sigma}_1, \sigma_2)$, while $\hat{\mathbf{P}}$ is the measure

induced by conditioning on $\hat{\xi}$. Since $\{\xi_0\} = \Xi \setminus \{\hat{\xi}\}$, $\tilde{\mathbf{P}}$ is the measure induced by conditioning on ξ_0 . That is, $\tilde{\mathbf{P}}$ is induced by the strategy profile $\hat{\sigma} = (\hat{\sigma}_1, \sigma_2)$ and $\tilde{\mathbf{P}}$ by $\tilde{\sigma} = (\tilde{\sigma}_1, \sigma_2)$, describing how play evolves when player 1 is the commitment and normal type, respectively.

The action of the commitment type satisfies the following assumption.

Assumption 3 *Player 2 has a unique stage-game best response to $\hat{\alpha}_1$ (denoted by \hat{a}_2) and $\hat{\alpha} \equiv (\hat{\alpha}_1, \hat{a}_2)$ is not a stage-game Nash equilibrium.*

Let $\hat{\sigma}_2$ denote the strategy of playing the unique best response \hat{a}_2 to $\hat{\alpha}_1$ in each period independently of history. Since $\hat{\alpha}$ is not a stage-game Nash equilibrium, $(\hat{\sigma}_1, \hat{\sigma}_2)$ is not a Nash equilibrium of the complete information infinite horizon game.

Proposition 2 (Cripps, Mailath, and Samuelson (2004)) *Suppose the monitoring distribution ρ satisfies Assumptions 1 and 2, and the commitment action $\hat{\alpha}_1$ satisfies Assumption 3. In any Nash equilibrium of the game with incomplete information, the posterior probability assigned by player 2 to the commitment type, $\hat{\mu}^t$, converges to zero under $\tilde{\mathbf{P}}$, i.e.,*

$$\hat{\mu}^t(h^t) \rightarrow 0, \quad \tilde{\mathbf{P}}\text{-a.s.}$$

The intuition is straightforward: Suppose there is a Nash equilibrium of the incomplete information game in which both the normal and the action type receive positive probability in the limit (on a positive probability set of histories). On this set of histories, player 2 cannot distinguish between signals generated by the two types (otherwise player 2 could ascertain which type she is facing), and hence must believe that the normal and action types are playing the same strategies on average. But then player 2 must play a best response to this strategy, and hence to the action type. Since the action type's behavior is not a best response for the normal type (to this player 2 behavior), player 1 must eventually find it optimal to *not* play the action-type strategy, contradicting player 2's beliefs.

Assumption 3 requires a unique best response to $\hat{\alpha}_1$. For example, in the product-choice game, every action for player 2 is a best response to player 1's mixture α'_1 that assigns equal probability to H and L . This indifference can be exploited to construct an equilibrium in which (the normal) player 1 plays α'_1 after every history (Mailath and Samuelson, 2006, Section 7.6.2). This will still be an equilibrium in the game of incomplete information in which the commitment type plays α'_1 , with the identical play of the normal and commitment types ensuring that player 2 never learns player 1's type.

In contrast, player 2 has a unique best response to any other mixture on the part of player 1. Therefore, if the commitment type is committed to any mixed action other than α'_1 , player 2 will eventually learn player 1's type.

As in Proposition 1, a key step in the proof of Proposition 2 is a purely statistical result on updating. Either player 2's expectation (given her history) of the strategy played by the normal type ($\tilde{E}[\tilde{\sigma}_1^t | h^t]$, where \tilde{E} denotes expectation with respect to $\tilde{\mathbf{P}}$) is in the limit identical to the strategy played by the action type ($\hat{\alpha}_1$), or player 2's posterior probability that player 1 is the action type ($\hat{\mu}^t(h^t)$) converges to zero (given that player 1 is indeed normal). This is a merging argument and closely related to Lemma 2. If the distributions generating player 2's signals are different for the normal and action type, then these signals provide information that player 2 will use in updating her posterior beliefs about the type she faces. This (converging, since beliefs are a martingale) belief can converge to an interior probability only if the distributions generating the signals are asymptotically uninformative, which requires that they be asymptotically identical.

Lemma 3 *Suppose the monitoring distribution ρ satisfies Assumptions 1 and 2. Then in any Nash equilibrium,*

$$\lim_{t \rightarrow \infty} \hat{\mu}^t \max_{a_1} \left| \hat{\alpha}_1(a_1) - \tilde{E}[\tilde{\sigma}_1^t(a_1) | h^t] \right| = 0, \quad \tilde{\mathbf{P}}\text{-a.s.} \quad (7)$$

Given Proposition 2, it should be expected that continuation play converges to an equilibrium of the complete information game, and this is indeed the case. See Cripps, Mailath, and Samuelson (2004, Theorem 2) for the formal statement.

Proposition 2 leaves open the possibility that for any period T , there may be equilibria in which uncertainty about player 1's type survives beyond T , even though such uncertainty asymptotically disappears in any equilibrium. This possibility cannot arise. The existence of a sequence of Nash equilibria with uncertainty about player 1's type persisting beyond period $T \rightarrow \infty$ would imply the (contradictory) existence of a limiting Nash equilibrium in which uncertainty about player 1's type persists.

Proposition 3 (Cripps, Mailath, and Samuelson (2007)) *Suppose the monitoring distribution ρ satisfies Assumptions 1 and 2, and the commitment action $\hat{\alpha}_1$ satisfies Assumption 3. For all $\varepsilon > 0$, there exists T such that for any Nash equilibrium of the game with incomplete information,*

$$\tilde{\mathbf{P}}(\hat{\mu}^t < \varepsilon, \forall t > T) > 1 - \varepsilon.$$

Example 2 Recall that in the product-choice game, the unique player 2 best response to the H is to play h , and Hh is not a stage-game Nash equilibrium. Proposition 1 ensures that the normal player 1's expected value in the repeated game of incomplete information with the H -action type is arbitrarily close to 2, when player 1 is very patient. In particular, if the normal player 1 plays H in every period, then player 2 will at least eventually play her best response of h . If the normal player 1 persisted in mimicking the action type by playing H in each period, this behavior would persist indefinitely. It is the feasibility of such a strategy that lies at the heart of the reputation bounds on expected payoffs. However, this strategy is not optimal. Instead, player 1 does even better by attaching some probability to L , occasionally reaping the rewards of his reputation by earning a stage-game payoff even larger than 2. The result of such equilibrium behavior, however, is that player 2 must eventually learn player 1's type. The *continuation* payoff is then bounded below 2 (recall (1)). ★

Reputation effects arise when player 2 is uncertain about player 1's type, and there may well be a long period of time during which player 2 is sufficiently uncertain of player 1's type (relative to the discount factor), and in which play does not resemble an equilibrium of the complete information game. Eventually, however, such behavior must give way to a regime in which player 2 is (correctly) convinced of player 1's type.

For any prior probability $\hat{\mu}$ that the long-lived player is the commitment type and for any $\varepsilon > 0$, there is a discount factor δ sufficiently large that player 1's expected payoff is close to the commitment-type payoff. This holds no matter how small $\hat{\mu}$. However, for any fixed δ and in any equilibrium, there is a time at which the posterior probability attached to the commitment type has dropped below the corresponding critical value of $\hat{\mu}$, becoming too small (relative to δ) for reputation effects to operate.

A reasonable response to the results on disappearing reputation effects is that a model of *long-run* reputations should incorporate some mechanism by which the uncertainty about types is continually replenished. For example, Holmström (1982), Cole, Dow, and English (1995), Mailath and Samuelson (2001), and Phelan (2006) assume that the type of the long-lived player is governed by a stochastic process rather than being determined once and for all at the beginning of the game. In such a situation, reputation effects can indeed have long-run implications.

3.6 Reputation as a State

The posterior probability that short-lived players assign to player 1 being $\hat{\xi}$ is sometimes interpreted as player 1's reputation, particularly if $\hat{\xi}$ is the Stackelberg type. When Ξ contains only the normal type and $\hat{\xi}$, the posterior belief $\hat{\mu}^t$ is a state variable of the game, and attention is sometimes restricted to Markov strategies (i.e., strategies that only depend on histories through their impact on the posterior beliefs of the short-lived players). An informative example is Benabou and Laroque (1992), who study the Markov perfect equilibria of a game in which the uninformed players respond continuously to their beliefs. They show that the informed player eventually reveals his type in any Markov perfect equilibrium. On the other hand, Markov equilibria need not exist in finitely repeated reputation games (Mailath and Samuelson, 2006, Section 17.3).

The literature on reputation effects has typically not restricted attention to Markov strategies, since the results do not require the restriction.

4 Two Long-Lived Players

The introduction of nontrivial intertemporal incentives for the uninformed player significantly reduces reputation effects. For example, when only simple Stackelberg types are considered, the Stackelberg payoff may not bound equilibrium payoffs. The situation is further complicated by the possibility of non-simple commitment types (i.e., types that follow nonstationary strategies).

Consider applying the logic from Section 3.4 to obtain the Stackelberg reputation bound when both players are long-lived and player 1's characteristics are unknown, under perfect monitoring. The first step is to demonstrate that, if the normal player 1 persistently plays the Stackelberg action and there exists a type committed to that action, then player 2 must eventually attach high probability to the event that the Stackelberg action is played in the future. This argument, a simple version of Lemma 2, depends only upon the properties of Bayesian belief revision, independently of whether the person holding the beliefs is a long-lived or short-lived player.

When player 2 is short-lived, the next step is to note that if she expects the Stackelberg action, then she will play a best response to this action. If player 2 is instead a long-lived player, she may have an incentive to play something other than a best response to the Stackelberg type.

The key step when working with two long-lived players is thus to establish conditions under which, as player 2 becomes increasingly convinced

that the Stackelberg action will appear, player 2 must eventually play a best response to that action. One might begin such an argument by observing that, as long as player 2 discounts, any losses from not playing a current best response must be recouped within a finite length of time. But if player 2 is “very” convinced that the Stackelberg action will be played not only now but for sufficiently many periods to come, there will be no opportunity to accumulate subsequent gains, and hence player 2 might just as well play a stage-game best response.

Once it is shown that player 2 is best responding to the Stackelberg action, the remainder of the argument proceeds as in the case of a short-lived player 2. The normal player 1 must eventually receive very nearly the Stackelberg payoff in each period of the repeated game. By making player 1 sufficiently patient (*relative to player 2*, so that discount factors differ), this consideration dominates player 1’s payoffs, putting a lower bound on the latter. Hence, the obvious handling of discount factors is to fix player 2’s discount factor δ_2 , and to consider the limit as player 1 becomes patient, i.e., δ_1 approaching one.

This intuition misses the following possibility. Player 2 may be choosing something other than a best response to the Stackelberg action out of fear that a current best response may trigger a disastrous future punishment. This punishment would not appear if player 2 faced the Stackelberg type, but player 2 can be made confident only that she faces the Stackelberg *action*, not the Stackelberg type. The fact that the punishment lies off the equilibrium path makes it difficult to assuage player 2’s fear of such punishments. Short-lived players in the same situation are similarly uncertain about the future ramifications of best responding, but being short-lived, this uncertainty does not affect their behavior.

Consequently, reputation effects are typically weak with two long-lived players under perfect monitoring: Celentani, Fudenberg, Levine, and Pendorfer (1996) and Cripps and Thomas (1997), describe examples with only the normal and the Stackelberg types of player 1, in which the future play of the normal player 1 is used to punish player 2 for choosing a best response to the Stackelberg action when she is not supposed to, and player 1’s payoff is significantly below the Stackelberg payoff. Moreover, the robustness of reputation effects to additional types beyond the Stackelberg type, a crucial feature of settings with one long-lived player, does not hold with two long-lived players. Schmidt (1993b) showed that the possibility of a “punishment” type can prevent player 2 best responding to the Stackelberg action, while Evans and Thomas (1997) showed that the Stackelberg bound is valid if in addition to the Stackelberg type, there is an action type who punishes

player 2 for *not* behaving appropriately (see Mailath and Samuelson (2006, Sections 16.1 and 16.5)).

Imperfect monitoring (of *both* players' actions), on the other hand, rescues reputation effects. With a sufficiently rich set of commitment types, player 1 can be assured of at least his Stackelberg payoff. Indeed, player 1 can often be assured of an even higher payoff, in the presence of commitment types who play nonstationary strategies (Celentani, Fudenberg, Levine, and Pesendorfer, 1996). At the same time, these reputation effects are temporary (Cripps, Mailath, and Samuelson, 2007, Theorem 2).

Finally, there is a literature on reputation effects in bargaining games (see Chatterjee and Samuelson (1987, 1988), Schmidt (1993a), and Abreu and Gul (2000)), where the issues described above are further complicated by the need to deal with the bargaining model itself.

5 Future Directions

The detailed structure of equilibria of the incomplete information game is not well understood, even for the canonical game of Section 3. A more complete description of the structure of equilibria is needed.

While much of the discussion was phrased in terms of the Stackelberg type, Proposition 1 provides a reputation bound for *any* action type. While in some settings, it is natural that the uninformed players assign strictly positive probability to the Stackelberg type, it is not natural in other settings. A model endogenizing the nature of action types would be an important addition to the reputation literature.

Finally, while the results on reputation effects with two long-lived players are discouraging, there is still the possibility that some modification of the model will rescue reputation effects in this important setting.

References

- ABREU, D., AND F. GUL (2000): "Bargaining and Reputation," *Econometrica*, 68(1), 85–117.
- BENABOU, R., AND G. LAROQUE (1992): "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility," *Quarterly Journal of Economics*, 107(3), 921–958.
- CELENTANI, M., D. FUDENBERG, D. K. LEVINE, AND W. PESENDORFER

- (1996): “Maintaining a Reputation Against a Long-Lived Opponent,” *Econometrica*, 64(3), 691–704.
- CHATTERJEE, K., AND L. SAMUELSON (1987): “Bargaining with Two-Sided Incomplete Information: An Infinite Horizon Model with Alternating Offers,” *Review of Economic Studies*, 54(2), 175–192.
- (1988): “Bargaining with Two-Sided Incomplete Information: The Unrestricted Offers Case,” *Operations Research*, 36(4), 605–638.
- COLE, H. L., J. DOW, AND W. B. ENGLISH (1995): “Default, Settlement, and Signalling: Lending Resumption in a Reputational Model of Sovereign Debt,” *International Economic Review*, 36(2), 365–385.
- CRIPPS, M. W., G. J. MAILATH, AND L. SAMUELSON (2004): “Imperfect Monitoring and Impermanent Reputations,” *Econometrica*, 72(2), 407–432.
- (2007): “Disappearing Private Reputations in Long-Run Relationships,” *Journal of Economic Theory*, 134(1), 287–316.
- CRIPPS, M. W., AND J. P. THOMAS (1997): “Reputation and Perfection in Repeated Common Interest Games,” *Games and Economic Behavior*, 18(2), 141–158.
- EVANS, R., AND J. P. THOMAS (1997): “Reputation and Experimentation in Repeated Games with Two Long-Run Players,” *Econometrica*, 65(5), 1153–1173.
- FUDENBERG, D., AND D. K. LEVINE (1989): “Reputation and Equilibrium Selection in Games with a Patient Player,” *Econometrica*, 57(4), 759–778.
- (1992): “Maintaining a Reputation When Strategies Are Imperfectly Observed,” *Review of Economic Studies*, 59(3), 561–579.
- HOLMSTRÖM, B. (1982): “Managerial Incentive Problems: A Dynamic Perspective,” in *Essays in Economics and Management in Honour of Lars Wahlbeck*, pp. 209–230. Swedish School of Economics and Business Administration, Helsinki, Published in *Review of Economic Studies* 66 (1), January 1999, 169–182.
- KREPS, D., P. R. MILGROM, D. J. ROBERTS, AND R. WILSON (1982): “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma,” *Journal of Economic Theory*, 27, 245–252.

- KREPS, D., AND R. WILSON (1982): “Reputation and Imperfect Information,” *Journal of Economic Theory*, 27, 253–279.
- MAILATH, G. J., AND L. SAMUELSON (2001): “Who Wants a Good Reputation?,” *Review of Economic Studies*, 68(2), 415–441.
- (2006): *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, New York, NY.
- MILGROM, P. R., AND D. J. ROBERTS (1982): “Limit pricing and entry under incomplete information: An equilibrium analysis,” *Econometrica*, 50, 443–459.
- PHELAN, C. (2006): “Public Trust and Government Betrayal,” *Journal of Economic Theory*, 130(1), 27–43.
- SCHMIDT, K. M. (1993a): “Commitment Through Incomplete Information in a Simple Repeated Bargaining Game,” *Journal of Economic Theory*, 60(1), 114–139.
- (1993b): “Reputation and Equilibrium Characterization in Repeated Games of Conflicting Interests,” *Econometrica*, 61(2), 325–351.
- SELTEN, R. (1978): “Chain-store paradox,” *Theory and Decision*, 9, 127–59.
- STACKELBERG, H. V. (1934): *Marktform und Gleichgewicht*. Julius Springer, Vienna.