The Reputation Trap[1]

David K. Levine[2]

Abstract

Few want to do business with a partner who has a bad reputation. Consequently once a bad reputation is established it can be difficult to get rid of. This leads on the one hand to the intuitive idea that a good reputation is easy to lose and hard to gain. On the other hand it can lead to a strong form of history dependence in which a single beneficial or adverse event can cast a shadow over a very long period of time. It gives rise to a reputational trap where an agent rationally chooses not to invest in a good reputation because the chances others will find out is too low. Never-the-less the same agent with a good reputation will make every effort to maintain it. Here a simple reputational model is constructed and the conditions for there to be a unique equilibrium that constitutes a reputation trap are characterized.

[1] I am indebted to Drew Fudenberg, Juan Block, Rohan Dutta, Mustafa Kaba, John Mair, Andrea Mattozzi, Salvatore Modica, Guillermo Ordonez, Nicole Stoelinga, to the 24th Spring Meeting of Young Economists, to the Lancaster Game Theory Conference and to the editor Asher Wolinsky and four anonymous referees. Financial support from the EUI Research Council and MIUR PRIN 2017 n. 2017H5KPLL_01 is gratefully acknowledged.

[2] Department of Economics and RSCAS at the EUI, emeritus professor at WUSTL; email: david@dklevine.com

"Glass, china, and reputation are easily cracked, and never well mended."

sometimes attributed to Benjamin Franklin.

1. Introduction

It is conventional to think that a good reputation is easy to lose and hard to gain. One reason we suspect this might be the case is that if you have a good reputation people will be eager to do business with you – hence if they are cheated it will quickly become known. On the other hand if you have a bad reputation few will do business with you so even if you are honest few will find out. In such a setting it is intuitive that history matters. If an adverse event causes a loss of reputation the difficulty of restoring it provides little incentive for honesty, so the bad reputation will deservedly remain so long after the circumstances that caused it are gone. On the other hand, there are reasons for honesty besides reputation – if circumstances dictate honesty it will take a long time before others find out, but once they do reputation will be restored – and even after the circumstances dictating honesty are gone it will be desirable to continue to be honest to avoid losing reputation. In other words, once reputation is restored it will also persist. Consequently, two otherwise identical individuals may find themselves with entirely different incentives for honesty because of an adverse or beneficial event that happened in the distant past.

This paper examines that intuition in an entry game between a long-run and short-run player prototypical of those used in the reputational literature. It follows in the tradition of the gang-of-four, Kreps and Wilson [1982] and Milgrom and Roberts [1982], who studied good equilibria in which the long-run player is always honest and showed that with behavioral types if the long-run player is sufficiently patient not only does such an equilibrium exist but it is necessary, that is, a good equilibrium is the only equilibrium. This paper studies trap equilibrium of the type described in the first paragraph in which a long-run player with a good reputation is honest and retains a good reputation while a long-run player with a bad reputation is dishonest and retains a bad reputation. As we indicate in our subsequent literature review it is known that without behavioral types such an equilibrium can exist if the long-run payer is sufficiently patient. This paper moves beyond that by using behavioral types to characterize which particular equilibrium we should expect to see. In line with the existing literature we show that for sufficient patience there can be no trap. The crucial new finding is that for an intermediate range of patience not only does does a trap exist but it is necessary, that is, a trap equilibrium is the only equilibrium.

In the entry game we study the short-run player prefers to enter if the long-run player provides costly effort and not otherwise, and the long-run player prefers effort and entry to the short-run player staying out. As indicated the model is driven by behavioral types: we have both good types corresponding to beneficial events as in the gang-of-four and bad types corresponding to adverse events as in Mailath and Samuelson [2001]. These types are persistent but not completely so as in Mailath and Samuelson [2001] and Horner [2002].[3] Finally, we insist that the information generated about long-run player behavior is greater if the short-run player chooses to enter than if not.[4] This observational asymmetry leads to an important change from the Mailath and Samuelson [2001] model where good and bad events are symmetric and reputation is equally easily lost or restored.

This model leads to a unique trap if we add an additional assumption concerning the short-run player. If short-run players stay out and no information is generated it eventually becomes likely that the long-run player has migrated back to a "normal" type. It is now possible for the short-run players and long-run player to coordinate. On a particular date it is common knowledge that if the long-run player is normal honest behavior will take place and that the short-run player will enter. This is then a self-fulfilling prophecy.[5] It is not, however, a very compelling one: it requires that both players agree about the exact timing of events in the long-distant past and that they agree that "today is the day." To rule this out we assume that agents know only about events that took place during their lifetime and that short-run player strategies and beliefs are independent of calendar time.

2. The Model

A dynamic game is played between overlapping generations of finitely lived players. There are two player roles: player 1 is a long-run player who lives many periods and player 2 represents a mass of short-run players who live a single period. Each period $t = 1, 2, \ldots$ a stage game is played. In the stage game long-run player must first choose whether or not to provide effort. Let $a_1 \in \{0, 1\}$ denote the decision of the long-run player with 1 meaning to provide effort and the cost being ca_1 where $0 < c < 1$. The short-run player moves second and without observing the effort choice of the long-run player[6] decides whether to enter $a_2 = 1$ or stay out $a_2 = 0$. The short-run player receives utility 0 for staying out, utility -1 for entering when no effort has been made and utility $V > 0$ for entering when effort is provided. There are three privately known types $\tau \in \{b, n, g\}$ of long-run player where g means "good" (a beneficial event), b means "bad" (an adverse event), and n means "normal." Player type is fixed during the lifetime of the player. The good and bad types are behavioral types: the good type always provides effort and the bad type never does. The stage game payoff of the normal type is given by $a_2 - ca_1$. Players care only about expected average utility during their lifetime.

The life of a long-run player is stochastic: with probability δ the player continues for another period, and with probability $1 - \delta$ is replaced. This re-

[3]As Cripps, Mailath and Samuelson [2004] show this is essential if we are to have reputations restored as well as decline.

[4]Fudenberg and Levine [1989] show how this limits the possibilities for reputation building.

[5]In Acemoglu and Wolitzky [2012] this induces a cycle.

[6]Meaning the game is simultaneous move.

placement is not observed by the short-run player. When a long-run player is replaced the type may change. The probability type τ is replaced by a type $\sigma \neq \tau$ is $Q_{\tau\sigma}\epsilon/(1-\delta)$ where $Q_{\tau\sigma} > 0$. Note that the scaling by $1-\delta$ implies that $1/\epsilon$ is a measure of the number of long-run player lifetimes before a type transition. We are interested in the case in which types are persistent - that is, in which ϵ is small.

At the beginning of each period a public signal z of what occurred in the previous period is observed and takes on one of three values: $1, 0, N$. If entry took place last period the signal is equal to last-period long-run player effort decision. If the short-run player stayed out last period then with probability $1 \geq \pi > 0$ the signal is equal to last period long-run player effort decision and with probability $1-\pi$ the signal is N. Here we are to think of "1" as a good signal (effort was observed), "0" as a bad signal (it was observed that there was no effort) and "N" as no signal.

There are two features of this information technology. First, even when the short-run player stays out some information is generated. Second, when the short-run players enter information is perfect. Subsequently we will model more closely investment and information and demonstrate the robustness of our results when information upon entry is less than perfect.

The game begins with an initial draw of the public signal $z(1)$ and private type $\tau(1)$ from the common knowledge distribution $\mu_{z\tau}(1)$.

Players are only aware of events that occur during their lifetime. The long-run player also knows their own generation T.[7] Let h denote a finite history for a long-run player. A strategy for the normal type of long-run player is a choice of effort probability $\alpha_1(h, t, T)$ as a function of privately known history, calendar time, and generation T. A strategy for the short-run player is a probability of entering $\alpha_2(z, t)$ as a function of the beginning of period signal and calendar time.

We study Nash equilibria of this game.

Throughout the paper we will assume *generic cost* in the sense that

$$c \notin \left\{ \delta, \frac{\delta}{2-\pi}, \frac{\delta\pi}{1-\delta+\delta\pi}, \frac{\delta\pi(\pi-\delta\pi)}{(1-\delta\pi)(1-\delta)) + \delta\pi(\pi-\delta\pi)} \right\}.$$

Short-run Player Beliefs and Time Invariant Equilibrium

If players know calendar time, as indicated in the introduction, they can use this information to coordinate their play in an implausible way. Hence we wish to assume that short-run player strategies and beliefs are independent of calendar time.[8] Notice that this same assumption is implicit in the definition of a Markov equilibrium, but is weaker since long-run player strategies may depend on the entire lifetime history of events as well as generation and calendar time.

[7] That is, how many replacement events have taken place since the beginning of the game.

[8] See Clark, Fudenberg and Wolitzky [2019] for the consequences of a similar information restriction in an overlapping generations setting.

For brevity all references to a decision problem of the long-run player should be understood to refer to the normal type. A strategy for a short-run player is a now a time invariant probability of entering $\alpha_2(z) \in [0,1]$ as a function of the beginning of period signal. Given such a strategy the normal type faces a well-posed Markov decision problem. It depends only on the probability α_2 with which the short-run player enters. Let $V(\alpha_2)$ denote the corresponding expected average value of utility. First period utility is $\alpha_2 - ca_1$. With probability δ the game continues and the probability of the next signal is $P(z'|z,a_1)$ where $P(1|z,1) = P(0|z,0) = \alpha_2(z)+(1-\alpha_2(z))\pi$ and $P(N|z,a_1) = (1-\alpha_2(z))(1-\pi)$. Hence the Bellman equation is

$$V(\alpha_2) = \max_{a_1}(1-\delta)\left[\alpha_2 - ca_1\right] + \delta \sum_{z'} P(z'|z,a_1)V(\alpha_2(z')).$$

As usual, this has a unique solution. The set of best responses, for the normal type, then, is determined entirely by the current state through $\alpha_2(z)$. Hence at time t with signal z_t any best response of the normal type $\alpha_1(y_t,t,T_t)$ must lie in this set. Time invariant beliefs of the short-run player about the effort probability of the normal type, which we denote by $\alpha_1(z)$, are then a weighted average of the best responses $\alpha_1(y_t,t,T_t)$ - and so must also be a best response and lie in this set.

Prior to observing the signal z_t the short-run player at time t has unconditional beliefs about the joint distribution $\mu_{z\tau}(t)$ from which the signal and type of the long-run player are drawn. After observing z_t short-run player beliefs about long-run player type are given by the conditional probability $\mu_{\tau|z_t}(t)$. This together with beliefs about the normal type effort $\alpha_1(z_t)$ determines $\mu^1(z_t,t)$ the overall beliefs about the probability of long-run player effort. The short-run player strategy $\alpha_2(z_t)$ must then be a best response to those beliefs.

The evolution of $\mu_{z\tau}(t)$ depends upon the initial condition $\mu_{z\tau}(1)$ and the beliefs of the short-run player about the probabilities with which earlier normal-type long-run and short-run players chose actions $\alpha_1(z)$, $\alpha_2(z)$. It does not depend on the actual choice of those actions or the earlier signals, none of which are observed. This has two consequences. First, no action or deviation by the long-run player has any effect on the evolution of $\mu_{z\tau}(t)$. Second, the evolution of $\mu_{z\tau}(t)$ is deterministic as it does not depend on the stochastic realization of actions, signals or types. The stochastic nature of short-run player beliefs are due to the single stochastic variable they observe, the signal, that is, $\mu_{\tau|z_t}(t)$ is stochastic because z_τ is.

Since $\mu_{z\tau}(t)$ follows a deterministic law of motion if we let $\overrightarrow{\mu}(t)$ denote the vector with components $\mu_{z\tau}(t)$ that law is $\overrightarrow{\mu}(t+1) = \mathbf{A}\overrightarrow{\mu}(t)$ where $\mathbf{A}$ is a Markov transition matrix the coefficients of which are determined by $\alpha_1(z), \alpha_2(z)$ and π, Q, ϵ.[9] To have an equilibrium with time invariant beliefs it must be that $\overrightarrow{\mu}(t+1) = \overrightarrow{\mu}(t)$ and this is true if and only if the initial condition $\mu_{z\tau}(1)$ is

[9]This is computed in the Appendix.

a stationary distribution of $\mathbf{A}$. For time invariance we cannot have arbitrary initial short-run player beliefs $\mu_{zt}(1)$, but only initial beliefs that are consistent with the strategies of the players and the passage of time.

We take our object of study, then, to be *time invariant equilibrium.* This is a Nash equilibrium in which the initial beliefs of the short-run players are determined endogenously to be the stationary distribution that arises from the equilibrium strategies. It is conveniently described as a triple $(\alpha_1(z), \alpha_2(z), \mu_{z\tau})$ where $\alpha_1(z)$ and $\mu_{z\tau}$ are time invariant beliefs of the short-run player and $\alpha_2(z)$ is the strategy of the short-run players. The conditions for equilibrium are that $\alpha_1(z)$ is a solution to the Markov decision problem induced by the short-run player strategy $\alpha_2(z)$, that $\mu_{z\tau}$ is a stationary distribution of the Markov transition matrix $\mathbf{A}$ determined by $\alpha_1(z), \alpha_2(z)$, and Q, ϵ, and that $\alpha_2(z)$ is a best response to beliefs about long-run player action $\mu^1(z)$ determined from $\alpha_1(z), \mu_{z\tau}$.

Let $z(y)$ be the most recently observed signal by the long-run player in the history y. We may conveniently summarize the discussion:

Theorem 1. *If $(\alpha_1(z), \alpha_2(z), \mu_{z\tau})$ is a time invariant equilibrium then the strategies $\alpha_1(y,t,T) = \alpha_1(z(y)), \alpha_2(z,t) = \alpha_2(z)$ are a Nash equilibrium with respect to the initial condition $\mu_{z\tau}(1) = \mu_{z\tau}$. Conversely if $\alpha_1(y,t,T), \alpha_2(z,t)$ is a Nash equilibrium that satisfies the time invariant short-run player condition that the short-run player equilibrium beliefs $\alpha_1(z,t) = \alpha_1(z), \mu_{z\tau}(t) = \mu_{z\tau}$ and equilibrium strategy $\alpha_2(z,t) = \alpha_2(z)$ then $(\alpha_1(z), \alpha_2(z), \mu_{z\tau})$ is a time invariant equilibrium.*

Hereafter by equilibrium we mean time invariant equilibrium.

3. Characterization of Equilibrium

Our main result characterizes when a trap does and does not occur. It shows that there is a single pure strategy equilibrium that is one of three types and give conditions under which that equilibrium is unique. In reading the theorem, note that $1 - \delta + \delta\pi$ is a weighted average of 1 and π so is strictly greater than π.

Theorem 2. *For given V, Q there exists an $\underline{\epsilon} > 0$ such that for all $\epsilon \in (0, \underline{\epsilon}\pi^2(1-\pi))$ and*

i. [bad] *If*

$$c > \delta$$

then there is a unique equilibrium, it is strict and in pure strategies, there is no effort by the normal type, and the short-run player enters only on the good signal.

ii. [trap] *If*

$$\delta > c > \delta\pi/(1 - \delta + \delta\pi)$$

then there is exactly one pure strategy equilibrium, it is strict, the normal type provides effort only on the good signal, and the short-run player enters only on the good signal. *If in addition $c > \delta/(1+\delta(1-\pi))$ this is the unique equilibrium.*

iii. [good] *If*

$$c < \delta\pi/(1-\delta+\delta\pi)$$

then there is exactly one pure strategy equilibrium, the normal type always provides effort, and the short-run player enters only on the good signal.

Note that the boundary cases are ruled out by the generic cost assumption.[10]

This result is described in terms of the comparative statics of entry cost c: it shows how the set of equilibria changes as c is reduced. As all of the cutoffs $\delta\pi/(1-\delta+\delta\pi) = \delta\pi/(1-\delta(1-\pi))$ and $\delta/(1+\delta(1-\pi))$ are strictly increasing in δ the results may equally be described in terms of increasing the discount factor δ, with the (more complicated) cutoffs described in terms of c .

The proof is outlined below with the detailed computations in the Appendix. The result has two main parts: the characterization of pure strategy equilibria and the uniqueness of pure strategy equilibria. We will discuss each of these in turn.

The pure strategy equilibrium is relatively intuitive. The assumption that ϵ is small means that types are highly persistent so the short-run player does not put much weight on the possibility of the type changing. Given the possible strategies of the long-run player the signal 0 indicates either a bad type or a normal type who will not provide effort if entry is not anticipated. Hence it makes sense for the short-run player not to enter in the face of bad signal. Similarly the signal 1 indicates either a good type or a normal type who will provide effort if entry is anticipated, so it makes sense for the short-run player to enter in the face of a good signal.

More subtle is the inference of the short-run player when the signal N is observed. The short-run player can infer that the previous short-run player chose not to enter - hence must have received the bad signal or was in the same boat with the signal N. As a result while less decisive than the signal 0 the signal N also indicates past bad behavior by the long-run player, so staying out is a good idea.

For the long-run player the choice is whether to provide effort when entry is anticipated and when it is not. The difference between the two cases lies in the probability that effort results in a good reputation which we may denote by $p = 1$ when entry is anticipated and $p = \pi$ when it is not. It is useful to consider the problem for general values of p: when the cost c is incurred there is a probability p of successfully establishing a good reputation and gaining $1-c$ in the future and probability $1-p$ of failing to establish a good reputation and starting over again. Here the expected average present value of the gain from effort is $\Gamma = -(1-\delta)c + \delta p(1-c) + \delta(1-p)\Gamma$ or

$$\Gamma = \frac{\delta p(1-c) - (1-\delta)c}{1-\delta(1-p)}.$$

[10] There is also a fourth case: if $c < \delta/(1+\delta(1-\pi))$ and there are "enough" normal types then there are at least two mixed strategy equilibria. As this result is not central it is discussed only in the online Appendix.

If this is negative, that is $\delta p(1-c) < (1-\delta)c$, then it is best not to provide effort and conversely. Take first the case where information is revealed immediately, that is $p = 1$. This is the situation most conducive to effort. The condition for not wishing to provide effort is $c > \delta$ so when this is the case there will be no effort. This is a standard case, corresponding to part (i) of the Theorem in which the long-run player is impatient and does not find it worthwhile to give up c for a future gain of $1 - c$. In this case effort will be provided only occasionally during beneficial events when the good type provides effort for non-reputational reasons.

When $c < \delta$ it is worth it to maintain a reputation when the short-run player enters as indeed in this case $p = 1$. The remaining question is whether it is also worth it to provide effort when the short-run player does not enter. In this case $p = \pi$, and the condition for effort is that given in (iii). If c is very small then it is worth providing effort even when the short-run player does not enter. This good equilibrium corresponds to the "usual" reputational case, for example in Kreps and Wilson [1982], Milgrom and Roberts [1982], Fudenberg and Levine [1989], Fudenberg and Levine [1992] or Mailath and Samuelson [2001]. There the long-run player is always is willing to provide effort over the relevant horizon.[11] Here, as in Mailath and Samuelson [2001], occasionally an adverse event occurs and the bad type does not provide effort regardless of reputational consequences so there is no effort until another normal or good type arrives.

The new and the interesting case is the trap equilibrium in case (ii) where $\delta > c$ so the cost of effort is low enough to maintain a reputation, but $c > \delta\pi/(1 - \delta + \delta\pi)$ so it is not worth it to try to acquire a reputation. Here we have strong history dependence. Depending on the history a normal type will be in one of two very different situations. A normal type that follows a history of good signals, will provide effort, have a good reputation and have a wealthy and satisfactory life with an income of $1 - c$. A normal type that has the ill-luck to follow a history in which the last signal was bad or there was no signal will not provide effort , will have a (deservedly) bad reputation, and have an impoverished life with an income of 0. This is a reputational trap. The only difference between these normal types is an event that took place in the far distant past: did the last behavioral type correspond to an adverse or beneficial event? Looked at another way, adverse and beneficial events, rare as they are, cast a very long shadow. After a beneficial event there will be many lives of prosperous normal types - indeed until an adverse event occurs. Contrariwise, following an adverse event normal types will be mired in the reputation trap until they are fortunate enough to have a beneficial event.

Observe that $\delta\pi/(1 - \delta + \delta\pi)$ is increasing in π so as π increases and news spreads quickly the range of costs for the reputation trap diminishes and we are more likely to see the "usual" good reputation case. More important, although

[11] In models without type replacement eventually effort stops and the equilibrium collapses permanently into a no effort trap. Mailath and Samuelson [2001] show that with type replacement there is always effort.

we will defer discussion of mixed strategies, is the condition

$$\delta > c > \delta \max \left\{ \frac{\pi}{1-\delta+\delta\pi}, \frac{1}{1+\delta(1-\pi)} \right\}$$

in which the trap equilibrium is the only equilibrium: that is, in this case not only does the pure strategy equilibrium constitute a trap but there is no other equilibrium. Here the crucial fact is that both $\pi/(1-\delta+\delta\pi)$ and $1/(1+\delta(1-\pi)$ are both strictly less than one, so there is always a range of costs c in which the trap is the unique equilibrium.

4. Discussion

We place this result in the literature then give the idea of the proof.

Literature Review and the Role of Behavioral Types

There are two distinct strands of the reputation literature. The first follows the gang-of-four Kreps and Wilson [1982] and Milgrom and Roberts [1982] and uses behavioral types. It focuses not only on the existence of equilibria, but on the uniqueness of equilibrium. The second follows the repeated long-run short run player game (without types) literature starting with Fudenberg, Kreps and Maskin [1990] who show that many types of equilibria are possible.

In the literature with behavioral types the possibility of unique history dependent equilibrium has been studied, but the type of equilibrium that has been studied is cyclic. In a cyclic equilibrium long-run players with a good reputation exploit it by providing low effort and reduce their reputation while those with a bad reputation provide high effort in an attempt to rebuild their reputation. This is the opposite of a trap equilibrium where those with a good reputation work to preserve it and those with a bad reputation choose not to rebuild it. Like a trap equilibrium a cyclic equilibrium alternates between good and bad reputation, but a player with a bad reputation is by no means trapped: that player has a bad reputation through the earlier choice of running it down and is actively working to rebuild it. Cyclic equilibria are studied by Liu [2011] and Liu and Skrzypacz [2014], and earlier work by Benabou and Laroque [1992] points in the same direction. A related analysis can be found in Board and Meyer-ter-Vehn [2013]'s good news case. Here it can clearly be seen the the informational assumption is the opposite of the one that leads to a trap: when low effort is provided information leaks out slowly even when the short-run player enters. A related result is Phelan [2006], who also examines reputation that is gradually rebuilt, albeit this is driven by the normal type playing a mixed strategy.[12]

In the literature without behavioral types trap equilibria have been studied but there are no uniqueness results and trap equilibria are but one of many. The idea can be understood by examining the model here without the behavioral

[12]Mathevet, Pearce and Stachetti [2019] examine mixing by the behavioral type.

types. As usual the bad equilibrium, the static Nash equilibrium of always stay out and never provide effort, is a subgame perfect equilibrium. In the high cost/low discount case (i) of Theorem 2 this is the only equilibrium regardless of whether their are behavioral types. For higher discount factors both the trap and good strategies are also Nash equilibria[13] What enables us to pin down a particular equilibrium are the behavioral types. In the usual way in the gang-of-four literature the presence of good types eliminates the static Nash equilibrium once the discount factor is high enough. The bad types, however, are key in selecting between the trap and good equilibria, and this is the new result of this paper. The presence of behavioral types insures that the ergodic distribution is unique and that all signals (except possibly N) are present - so acts somewhat like trembles. The good equilibrium is eliminated in the intermediate case (ii) and the trap equilibrium in low cost/high discount factor case (iii) because play must be optimal following a signal of no effort.[14]

The result that without types their can be multiple equilibria including trap equilibria is well established in the literature. Rob and Fishman [2005] use an information structure similar to that here and establish the existence of a trap equilibrium in which those with a good reputation provide effort and those with a bad reputation do not. However, equilibrium is their model is certainly not unique and indeed they "note the existence of a trivial equilibrium, which replicates the static equilibrium under a one-shot interaction." Along the same lines is the bad-news case of Board and Meyer-ter-Vehn [2013]. Their model differs from the standard reputation model. In the standard reputation model reputation is analyzed as a substitute for commitment. By contrast Board and Meyer-ter-Vehn [2013] allow partial commitment in the sense that actions by the long-run player once taken persist for some length of time. In this setting with a bad news information structure similar to the one here they show that trap equilibria exist. They do not establish uniqueness, but in the opposite direction they do give a sufficient condition for a continuum of equilibria to exist.

One paper that does combine the information structure of this paper with (good) behavioral types is Ordonez [2007]. That paper, however, is not focused on uniqueness, but rather introduces a second dimension of long-run player action, how many groups to serve, and focuses on the issue how the number of groups served depends on reputation and whether or not it is efficient.

Outline of the Proof: Pure Strategies

The proof of the main theorem involves the interplay between the strategy of the long-run player and the beliefs of the short-run player. The detailed calculations are given in the Appendix through a series of Lemmas. Lemma 1 analyzes the optimum of the long-run player. It shows that regardless of the

[13] Only the trap equilibrium is subgame perfect, however.

[14] The good equilibrium is chosen in case (iii) despite the fact that it is not subgame perfect. Although normal types always provide effort it is optimal for the short-run player to stay out on a signal of no effort: this is because such a signal indicates a bad type.

strategy of the short-run player the long-run player must provide effort when entry is is anticipated if she is willing to do so when entry is not anticipated. It shows in addition that unless the short-run player enters on the good signal and stays out on the bad signal the long-run player should never provide effort. This information is subsequently used to rule out many combinations of long-run and short-run player strategies.

The next series of steps are to characterize the ergodic beliefs of the short-run player about the long-run player. Lemma 2 examines the marginal ergodic beliefs of the short-run player about the type of long-run player. As these transition probabilities are exogenous it is straightforward to show that these beliefs do not depend on ϵ and are bounded away from zero.

The key to showing that the unique equilibrium strategy of the short-run player is to enter only on a good signal is to characterize the ergodic beliefs of the short-run player about the type of long-run player conditional on the signal. Let $\overline{B}$ be the probability of effort that makes the short-run player indifferent to entering, that is, $\overline{B}V = (1 - \overline{B})$. Recall that $\mu^1(z)$ is the ergodic belief of the short-run player about the probability that the long-run player will provide effort. If $\mu^1(z) > \overline{B}$ it is strictly optimal to enter, and if it is less than this, strictly optimal to stay out. If we can show that

$$\mu^1(1) \geq 1 - K\frac{\epsilon}{\min\{\pi, 1-\pi\}}$$

and

$$\mu^1(0), \mu^1(N) \leq K\frac{\epsilon}{\min\{\pi, 1-\pi\}}$$

for some positive constant K depending only on Q then it follows that for

$$K\frac{\epsilon}{\min\{\pi, 1-\pi\}} < \min\{\overline{B}, 1-\overline{B}\}$$

it is strictly optimal for the short-run player to stay out on a bad or no signal and to enter on a good signal. This then gives the main theorem with $\underline{\epsilon} = \min\{\overline{B}, 1-\overline{B}\}/K$.

The derivation of the bounds requires several steps. Lemma 3 shows that to a good approximation the beliefs of the short-run player about the type of long-run player are the same at the beginning of a period where the type may have changed as they were at the end of the previous period. This enables us to compute approximate conditional beliefs about types and signals from the simpler problem in which types are persistent. We then want to apply Bayes law to compute the probability of types conditional on signals. To implement this we need to know a lower bound on the marginal probability of the signals: in the case of the good and bad signal this follows from the fact that the good and bad types are playing the good and bad action; the crucial case of no signal is addressed in Lemma 4 using ergodic calculations simplified by Lemma 3. Lemma 5 then uses Bayes law for the special case in which the long-run player takes an action independent of signal (as is the case for the behavioral types).

At this point there are three possible strategies for the long-run player and eight for the short-run. It is now possible to check each of the twenty four combinations to find the ergodic beliefs and show that the only best response for the short-run player to a best response of the long-run player is to enter on a good signal and stay out for all others. Fortunately many combinations can be checked at once. This is done in Proposition 1 using the previously established bounds and partial characterization of optimal strategies.

Finally, now that we know the unique strategy of the short-run player, we must calculate the best response of the long-run player: this is the computation with Γ above.

Intuition of the Main Result: Mixed Strategies

The important result is that there is a range of c for which there is a reputation trap and also no other equilibria. Why must this be the case? The reason is that the equilibrium short-run player pure strategy of staying out on a bad or no signal $z \in \{0, N\}$and entering on a good signal $z = 1$ provides the greatest incentive for the normal type to provide effort. If $c > \delta$ this is not enough, so weakening the incentive to provide effort by mixing does not help and the only equilibrium is the one in which the normal type never provides effort.

In the Appendix it is shown that if the short-run player uses a pure strategy the long-run player must do so as well. To understand why the short-run player strategy must remain pure even for $c < \delta$ (but not too small) consider that at $c = \delta$ the normal type strictly prefers to not to provide effort on a bad or no signal and is indifferent to effort on a good signal. When c is lowered slightly the normal type now strictly prefers to provide effort on a good signal, while of course the strict preference on bad and no signals remain. Can there be an equilibrium in which the short-run player mixes only "a little?" That cannot happen on a bad or no signal since to get the short-run player to mix the normal type would have to mix "a lot" and this in turn would require the short-run player to mix "a lot."

What about the good signal? Here with c a little less than δ "a little" mixing by the short-run player gets the normal type back to indifference. Without types this can be an equilibrium - but not with types. The reason is tied to the ergodic distribution of types and signals. With the normal type providing no effort on a bad or no signal once those states are reached the normal type will no longer get the good signal. With the short-run player mixing on the good signal there is a positive probability that the normal type will get no signal: this "drains" the normal types from the good signal so that in the ergodic distribution of types and signals conditional on a good signal it is extremely likely the short-run player is facing a good type. Consequently, the short-run player will not mix on a good signal - rather the short-run player will enter for certain.

The conclusion is that mixed strategy equilibria require the short-run player to mix "a lot." Formally it is shown in Lemma 14 that in any mixed equilibrium the short-run player must be at least as likely to enter on no signal as on a good signal. This provides substantially less incentive for the normal type to provide effort than the short-run player equilibrium pure strategy in which the short-run player is a lot less likely to enter on no signal than on a good signal. Hence the value of c that is low enough to provide adequate incentive for effort is higher for a pure strategy equilibrium than for any mixed strategy equilibrium.

5. Robustness

As indicated, we made two key assumptions about the information technology: first that some information is generated even when the short-run player stays out, and second that there is perfect information when the short-run player enters. To focus thinking it is useful to think of the short-run player as choosing between a single investment or purchase or making a large number K of identical investments or purchases. Each is subject to an idiosyncratic shock. In particular, for each investment/purchase, we may imagine that there is an independent probability π that the behavior of the long-run player is observed. Hence with a single investment/purchase - "staying out" - the probability of observation is π as in the base model. If there are K investment/purchases then the probability that the behavior of the long-run player is observed is $1-(1-\pi)^K$. Hence the base model corresponds to the limit in which there are many investments or purchases in which case the probability of observation is one.

In this context, it is important to know that our results are robust to K large but finite. This is straightforward because Theorem 2 shows that the pure strategy equilibria are strict for both the long-run and short-run player. The equilibrium conditions in the Appendix consist of finitely many continuous equalities and inequalities. Hence by standard arguments the equilibrium correspondence is upper-hemi-continuous as $K \to \infty$. In the crucial case in which equilibrium is unique, since it is strict, for K sufficiently large, the equilibrium strategies are unique and exactly those described in Theorem 2: the key result about a unique trap holds for K sufficiently large.

There is a second issue of importance, and that is the timing of information. We have assumed that the effort of the long-run player is observed by the short-run player only after entry - although of course it is not the timing that matters, but the fact that neither player knows the action of the other when the decision is taking. If the short-run player observes the effort of the long-run player before the entry decision is taken then the long-run player is a Stackelberg leader in the stage-game and the normal type will always invest: this is standard - there is no need for reputation as a substitute for commitment when commitment is possible in the stage game.

There is also the opposite timing: the long-run player observes the entry decision of the short-run player before deciding whether or not to provide effort. This is the case in Veugelers [1993], who studies a long-run government facing a series of short-run foreign investors using a conventional reputational model with a single good type and unlimited memory with the conventional result that if the long-run player is sufficiently patient near first best results are obtained. Veugelers [1993] is interested in the case of a government with low state capacity unable to provide a rule of law that must decide *ex post* whether or not to expropriate. While complete analysis of this case is beyond the scope of this paper, it is easy to see that if the long-run player instead of observing the signal of the short-run player observes whether or not the short-run player entered there can be no trap. The strategy spaces of the players and the belief dynamics of the short-run players remain unchanged. Hence for the trap parameters the long-run player provides effort when the short-run player enters, and does not do so when the short-run player stays out. If a bad or no signal means that the short-run player very likely faces a normal type it follows that the short-run player should enter knowing that the long-run player will respond by providing effort. This contradicts the supposition that the short-run player stays out on these signals.

References

[1] Acemoglu, D., and Wolitzky, A. (2012). Cycles of Distrust: An Economic Model (No. w18257). NBER.

[2] Benabou, R. and Laroque, G. (1992). Using privileged information to manipulate markets: Insiders, gurus, and credibility. Quarterly Journal of Economics 107: 921-958..

[3] Board, S., and Meyer-ter-Vehn, M. (2013). Reputation for quality. *Econometrica*, 81(6), 2381-2462.

[4] Block, J. I., Fudenberg, D. and Levine, D. K. (2018). Learning dynamics with social comparisons and limited memory. *Theoretical Economics*.

[5] Clark, D., D. Fudenberg, and A. Wolitzky (2019). Robust Cooperation with First-Order Information. mimeo MIT.

[6]

[7] Cripps, M. W., Mailath, G. J., and Samuelson, L. (2004). Imperfect monitoring and impermanent reputations. *Econometrica*, 407-432.

[8] Dutta, R., D. K. Levine, and S. Modica (2019). Interventions with Sticky Social Norms: A Critique. EUI.

[9] Ekmekci, M., Gossner, O., and Wilson, A. (2012). Impermanent types and permanent reputations. *Journal of Economic Theory*, 147(1), 162-178.

[10] Fudenberg, D., D. M. Kreps and E. S. Maskin (1990). Repeated games with long-run and short-run players. *The Review of Economic Studies*, 57(4), 555-573.

[11] Fudenberg, D., and Levine, D. K. (1989). Reputation and Equilibrium Selection in Games with a Single Long-Run Player. *Econometrica*, 57.

[12] Fudenberg, D., and Levine, D. K.. (1992). Maintaining a Reputation when Strategies are Imperfectly. *Review of Economic Studies*, 59, 561-579.

[13] Hörner, J. (2002). Reputation and competition. *American economic review*, 92(3), 644-663.

[14] Kreps, D. M., and Wilson, R. (1982). Reputation and imperfect information. *Journal of economic theory*, 27(2), 253-279.

[15] Liu, Q. (2011). Information Acquisition and Reputation Dynamics. *Review of Economic Studies* 78: 1400–1425.

[16] Liu, Q. and Skrzypacz, A. (2014). Limited records and reputation bubbles. *Journal of Economic Theory* 151: 2-29.

[17] Mailath, G. J., and Samuelson, L. (2001). Who wants a good reputation? *The Review of Economic Studies*, 68(2), 415-441.

[18] Mathevet, L., Pearce, D. and Stachetti, E. (2019). Reputation and Information Design, NYU

[19] Milgrom, P., and Roberts, J. (1982). Predation, reputation, and entry deterrence. *Journal of economic theory*, 27(2), 280-312.

[20] Ordonez, G. (2007).“Hard to Get, Easy to Lose: Endogenous Reputation and Underinvestment, mimeo Penn.

[21] Phelan, C. (2006). Public trust and government betrayal. *Journal of Economic Theory*, 130(1), 27-43.

[22] Rob, R., and A. Fishman (2005). Is bigger better? Customer base expansion through word-of-mouth reputation, *Journal of Political Economy*, 113, 1146-1162.

[23] Veugelers, R. (1993). Reputation as a mechanism alleviating opportunistic host government behavior against MNEs. *The Journal of Industrial Economics*, 1-17.

Appendix

For brevity and clarity only the results of lengthy computations are reported here. The interested reader can find the computations themselves in the online version of this appendix.

Problem of the long-run Player

We examine the problem of the normal type of long-run player. Recall the Bellman equation

$$V(\alpha_2) = \max_{a_1}(1-\delta)\left[\alpha_2 - ca_1\right] + \delta \sum_{z'} P(z'|z, a_1)V(\alpha_2(z')).$$

We may write this out as

$$V(\alpha_2) =$$

$$\max_{a_1}(1-\delta)\left[\alpha_2 - ca_1\right] + \delta\left[\left(\alpha_2 + (1-\alpha_2)\pi\right)V(\alpha_2(a_1)) + (1-\alpha_2)(1-\pi)V(\alpha_2(N))\right].$$

Lemma 1. *The optimum for the normal type of long-run-player depends on the state only through α_2 and one of three cases applies:*

(i) $V(\alpha_2(1)) - V(\alpha_2(0)) < c(1-\delta)/\delta$: it is strictly optimal to provide no effort in every state. In particular if $\alpha_2(1) = \alpha_2(0)$ this is the case.

(ii) $V(\alpha_2(1)) - V(\alpha_2(0)) > c(1-\delta)/(\delta\pi)$: it is strictly optimal to provide effort in every state

Defining

$$\tilde{\alpha}_2 = \frac{1-\delta}{\delta(1-\pi)\left(V(\alpha_2(1)) - V(\alpha_2(0))\right)}c - \frac{\pi}{1-\pi}$$

(iii) it is strictly optimal to provide effort if $\alpha_2(z) > \tilde{\alpha}_2$ and conversely. In particular the strategy $\alpha_1(0) > \alpha_1(1)$ is never optimal.

In addition

(iv) if $\alpha_2(0) = 1$ then it is strictly optimal to provide no effort in every state.

Finally, if the short-run player uses a pure strategy then the optimum of the long-run player is strict and pure.

Proof. The argmax is derived from:

$$\max_{a_1} -(1-\delta)ca_1 + \delta\left(\alpha_2 + (1-\alpha_2)\pi\right)V(\alpha_2(a_1)).$$

The gain to providing no effort is

$$G(\alpha_2) = (1-\delta)c - \delta\left(\alpha_2 + (1-\alpha_2\pi\right)\left[V(\alpha_2(1)) - V(\alpha_2(0))\right].$$

We then solve this equation form α_2 to see when effort is and is not optimal.

Finally, we analyze best response of the long-run player when the short-run player uses a pure strategy. From (i) and (iv) if $\alpha_2(0) \geq \alpha_2(1)$ it is strictly best to provide no effort . That leaves only the case $\alpha_2(a_1) = a_1$, or rather two cases, depending on $\alpha_2(N)$. This is a matter of solving the Bellman equations for each case to determine the value of c (if any) there can be a tie. This are the "non-generic" values listed in the text. $\square$

Ergodic Beliefs of the Short-Run Player

Next we examine the beliefs of the short-run player. For given pure strategies of both players the signal type pairs (z, τ) are a Markov chain with transition probabilities independent of δ and depending only on ϵ, π and the strategies of the two players. Excluding the state N in case the short-run player always enters the chain is irreducible and aperiodic so it has a unique ergodic distribution $\mu_{z\tau}$. We first analyze the marginals μ_τ and μ_z.

Lemma 2. *The marginals μ_τ are independent of ϵ. Let $\underline{\mu} = \min_{\tau \neq n} \mu_\tau$. Then $\underline{\mu} > 0$, $\mu_0, \mu_1 \geq \pi\underline{\mu}$, if $\alpha_2(0) = \alpha_2(1) = 1$ then $\mu_N = 0$, otherwise if the short-run player plays a pure strategy then $\mu_N \geq (1-\pi)\underline{\mu}$.*

Proof. The type transitions are independent of the signals, so we analyze those first. For $\epsilon > 0$ we have $\mu_\tau > 0$ since every type transition has positive probability. This ergodic distribution is the unique fixed point of the 3×3 transition matrix A, which is to say given by the intersection of the null space of $I - A$ with the unit simplex. Since $A = I + Q\epsilon$ it follows that it is given by the intersection of the null space of $Q\epsilon$ with the unit simplex. As the null space of $Q\epsilon$ is independent of ϵ the marginals μ_τ are independent of ϵ as well.

For the signals we have $\mu_1 \geq \pi\mu_g$ and $\mu_0 \geq \pi\mu_b$. If if $a_2(0) = a_2(1) = 1$ then the state N is transient. If $\alpha_2(1) = 0$ then $\mu_N \geq (1-\pi)\mu_g$ while if $\alpha_2(0) = 0$ then $\mu_N \geq (1-\pi)\mu_b$. $\square$

It will be convenient to normalize so that $\max(\mu_\sigma/\mu_\tau)Q_{\tau\sigma} = 1$. Next we show how the conditional probabilities $\mu_{z|\tau}$ can be computed approximately by using the ergodic conditions for $\epsilon = 0$.

Lemma 3. *When $z = N$*

$$\mu_{N|\tau} = (1-\pi)\left(\sum_y (1-\alpha_2(y))\mu_{y|\tau} + \epsilon H_{N\tau}\right)$$

when $z \neq N$

$$\mu_{z|\tau} =$$

$$\sum_y 1\left((z=1)\alpha_1(\tau, y) + 1(z=0)(1-\alpha_1(\tau, y))\right)\left[\alpha_2(y) + \pi(1-\alpha_2(y))\right]\mu_{y|\tau} + \epsilon H_{z\tau}.$$

where $|H_{z\tau}| \leq 2$ for all z.

Proof. The idea is that the process for types is exogenous, so the stationary probabilities can be computed directly. This enables us to find a linear recursive relationship for the conditionals where the coefficients depend upon the strategies and the (already known) marginals over types. We then show that when ϵ is small to a good approximation we can do the computation for $\epsilon = 0$, that is, ignoring the type transitions, with the result above showing how good the approximation is for given ϵ. $\square$

To apply Bayes Law we will need to bound marginal probabilities of signals from below. The hard case is that of no signal where we must solve the equations for the conditionals simultaneously. Here we analyze the short-run pure strategy case. If the short-run player enters for both $z = 0, 1$ then no signals are unlikely as they are generated only from type transitions, so we rule that out.

Lemma 4. *Suppose $\alpha_2(a_1) = 0$ for some $a_1 \in \{0, 1\}$. Then*

$$\mu_N \geq \frac{1-\pi}{2}\left(1 - \frac{4\epsilon}{\pi}\right)\underline{\mu}.$$

Proof. Let τ be the type that plays a_1. We have

$$\mu_{a_1|\tau} = \sum_y \left[\alpha_2(y) + \pi(1-\alpha_2(y))\right]\mu_{y|\tau} + \epsilon H_{a_1\tau}$$

$$\mu_{N|\tau} = (1-\pi)\left(\sum_y (1-\alpha_2(y))\mu_{y|\tau} + \epsilon H_{N\tau}\right)$$

These imply the inequalities

$$\mu_{a_1|\tau} \geq \pi(1-\mu_{N|\tau}) + \left[\alpha_2(N) + \pi(1-\alpha_2(N))\right]\mu_{N|\tau} + \epsilon H_{a_1\tau}$$

$$\mu_{N|\tau} \geq (1-\pi)\left((1-\alpha_2(N))\mu_{N|\tau} + \mu_{a_1|\tau} + \epsilon H_{N\tau}\right).$$

Hence

$$\mu_{N|\tau} \geq (1-\pi)\left(\pi + (1-\pi)\mu_{N|\tau} + \epsilon H_{N\tau} + \epsilon H_{a_1\tau}\right).$$

It follows that

$$\mu_{N|\tau} \geq \frac{1-\pi}{2}\left(1 - \frac{4\epsilon}{\pi}\right).$$

The result now follows from $\mu_N \geq \mu_{N|\tau}\mu_\tau \geq \mu_{N|\tau}\underline{\mu}$. $\square$

Finally we compute bounds on beliefs about types that play the same action independent of the signal. Here we combine bounds from the equations for the conditionals with Bayes Law.

Lemma 5. *A long-run type τ that plays the pure action a_1 regardless of the signal has*

$$\mu_{\tau|-a_1} \leq \frac{2}{\underline{\mu}}\left(\frac{\epsilon}{\pi}\right)$$

and if $\alpha_2(1) = 1$ and $\alpha_2(0) = 0$ then a type τ that plays the action 1 regardless of signal has

$$\mu_{\tau|N} \leq \frac{8}{\left(1 - 4\left(\frac{\epsilon}{\pi}\right)\right)\underline{\mu}}\left(\frac{\epsilon}{\pi}\right).$$

Proof. If long-run type τ plays the pure action a_1 from Lemma 3 $\mu_{-a_1|\tau} = \epsilon H_{-a_1\tau} \leq 2\epsilon$. From Lemma 2 $\mu_{-a_1} \geq \pi\underline{\mu}$ and Bayes law then implies

$$\mu_{\tau|-a_1} \leq \frac{\epsilon 2}{\pi\underline{\mu}}.$$

For the second part we have from Lemma 3

$$\mu_{N|\tau} = (1-\pi)\left(\mu_{0|\tau} + [1-\alpha_2(N)]\,\mu_{N|\tau}\right) + (1-\pi)\epsilon H_{N\tau}.$$

$$\mu_{0|\tau} = \epsilon H_{0\tau}.$$

Plugging in $\mu_{N|\tau} \leq (1-\pi)\mu_{N|\tau} + (1-\pi)\epsilon H_{0\tau} + (1-\pi)\epsilon H_{N\tau}$ so

$$\mu_{N|\tau} \leq \frac{(1-\pi)4\epsilon}{\pi}.$$

From Lemma 4

$$\mu_N \geq \frac{1-\pi}{2}\left(1-\frac{4\epsilon}{\pi}\right)\underline{\mu}.$$

Hence Bayes law implies

$$\mu_{\tau|N} \leq \frac{8\epsilon}{\pi\left(1-\frac{4\epsilon}{\pi}\right)\underline{\mu}}.$$

□

Short-Run Player Optimality

Recall that $\mu^1(z)$ is the probability of $a_1 = 1$ in state z and that $\overline{B} = 1/(V+1)$ is the critical value of $\mu^1(z)$ such that

Lemma 6. *If $\mu^1(z) > \overline{B}$ the short-run player strictly prefers to enter; if $\mu^1(z) < \overline{B}$ the short-run player strictly prefers to stay out, and if $\mu^1(z) = \overline{B}$ the short-run player is indifferent.*

We next show that it cannot be optimal for the short-run player always to enter. Set $B \equiv \underline{\mu}\min\{\pi, 1-\pi\}\min\{\overline{B}, 1-\overline{B}\}$.

Lemma 7. *For $\epsilon < (1/2)B$ always enter $a_2(z) = 1$ for all z is not an equilibrium.*

Proof. By Lemma 1 always enter implies no effort by the normal long-run player. As there are few good types at $z = 0$ we show that this forces the short-run player to stay out there so the short-run player should not in fact enter. □

Lemma 8. *For $\epsilon < (1/16)B$ the strict equilibrium response to never providing effort is to enter only on $z = 1$ and do so with probability 1.*

Proof. As the normal and bad types never provide effort the signal $z = 1$ implies a good type with high probability so the short-run player should enter there. This means that the long-run player can have the signal $z = 1, N$ only through a type transition. In particular the bad signal is dominated by normal and bad types so the short-run player should stay out. This in turn means that most of the N signals are generated by normal and bad types, so the short-run player should stay out there too. □

Lemma 9. *For $\epsilon < (1/16)B$ there is no equilibrium in which $\alpha_2(0) = 1$.*

Proof. By Lemma 1 $\alpha_2(0) = 1$ implies never provide effort so by Lemma 8 $\alpha_2(0) = 0$ a contradiction. □

Lemma 10. *For $\epsilon < (1/32)B$ the unique equilibrium response to always provide effort is to enter only on $z = 1$ and do so with probability* 1.

Proof. This is basically the opposite of Lemma 8. Now at $z = 1$ there are mainly good and normal types so it is optimal for the short-run player to enter. While at $z = 0$ there are mainly bad types so it is optimal for the short-run player to stay out. Hence no-signal is generated by bad types from $z = 0$ so it is optimal for the short-run player to stay out there too. □

Lemma 11. *If $\epsilon < (1/2)B$ and for some a_1 we have $\alpha_1(a_1) = a_1$ then $\alpha_2(a_1) = a_1$.*

Proof. If $\alpha_1(0) = 0$ then from Lemmas 3 and 2 $\mu^1(0) = \mu_{0|g}\mu_g/\mu_0 = \epsilon H_{0g}\mu_g/\mu_0 \leq 2\epsilon/(\pi\underline{\mu})$. If $\alpha_1(1) = 1$ then $1-\mu^1(1) = \mu_{1|b}\mu_b/\mu_1 = \epsilon H_{1b}\mu_b/\mu_1 \leq 2\epsilon/(\pi\underline{\mu})$. Hence for $\epsilon/\pi < \overline{B}\underline{\mu}/2$ it follows that $\alpha_2(a_1) = a_1$. □

Uniqueness of Short-Run Pure Equilibria

We define an *equilibrium response* of the short-run player to a strategy of the long-run player to be a best response to $\mu_{z\tau}$ induced by the long-run player strategy and itself.

Proposition 1. *There exists an $\underline{\epsilon} > 0$ depending only on V such that for any ϵ satisfying*

$$\underline{\epsilon} > \frac{\epsilon}{\underline{\mu}\min\{\pi, 1-\pi\}} > 0$$

in any short-run pure equilibrium the short-run player must enter on the good signal and only on the good signal. Moreover this is a strict equilibrium response.

Proof. We rule out all other possibilities:

(a) A*lways enter $a_2(z) = 1$ for all z is not an equilibrium.* By Lemma 7

(b) *The unique equilibrium response to never provide effort is to enter only on $z = 1$.* From Lemma 7.

(c) *A equilibrium response requires $a_2(1) = 1, a_2(0) = 0$.* Any other strategy satisfies $a_2(0) \geq a_2(1)$. From Lemma 1 this implies no effort by the long-run player. Part (b) then forces $0 = a_2(0) < a_2(1) = 1$ a contradiction.

(d) *The unique equilibrium response to always provide effort is to enter only on* $z = 1$. From Lemma 10.

This leaves only the strategy $\tilde{a}$ in which the long-run player plays $a_1 = 1$ on entry and $a_1 = 0$ if the short-run player stays out. As we know that $\alpha_2(1) = 1, \alpha_2(0) = 0$ there are two possibilities $\alpha_2(N) = 1$ and $\alpha_2(N) = 0$. The former is ruled out because it leads to primarily bad types at $z = N$, and the latter is a strict best response by the short-run player because there are few good types at $z = N$. □

Mixing

Recall that all of the Lemmas concerning short-run optimality hold for $\epsilon \leq B/32$ (and the remaining Lemmas do not place restrictions on ϵ) where $B = \underline{\mu} \min\{\pi, 1-\pi\} \min\{\overline{B}, 1-\overline{B}\}$. Recall also the notion of a fundamental bound: it may depend on the fundamentals of the game π, V, δ, c but not on the type dynamics Q, ϵ. Define the fundamental bound $\overline{A} \equiv \pi^2(1-\pi)\min\{\overline{B}, 1-\overline{B}\}$ and observe that if $\epsilon \leq \underline{\mu}\overline{A}/32$ then also $\epsilon \leq B/32$. We shall assume $\epsilon \leq \underline{\mu}\overline{A}/32$ hereafter.

Lemma 12. *There is no non-pure equilibrium with* $\alpha_1(1) = 1$.

Proof. By Lemma 2 $\mu_{1|b} = \epsilon H_{1b} \leq 2\epsilon$. Hence for $\epsilon < \overline{B}/2$ by Lemma 6 $\alpha_2(1) = 1$. Then by Lemma 2 $\mu_{1|n} = \mu_{1|n} + \sum_{y \in \{0,N\}} \alpha_1(y) \left[\alpha_2(y) + \pi(1-\alpha_2(y))\right] \mu_{y|n} + \epsilon H_{z\tau}$.I t follows that

$$\sum_{y \in \{0,N\}} \alpha_1(y)\mu_{y|n} \leq 2(\epsilon/\pi) \text{ so } \max_{y \in \{0,N\}} \alpha_1(y)\mu_{y|n} \leq 2(\epsilon/\pi).$$

Moreover for $z \in \{0, N\}$ we have $\mu_{z|g} = \epsilon H_{zg} \leq 2\epsilon$. Hence

$$\mu^1(0) = \frac{\mu_{0|g}\mu_g + \alpha_1(0)\mu_{0|n}\mu_n}{\mu_0} \leq 2(\epsilon/\pi)(\mu_g + \mu_n)/(\pi\underline{\mu}) \leq 2(\epsilon/\pi)/(\pi\underline{\mu}).$$

So for $\epsilon/\pi^2 < \overline{B}\underline{\mu}/2$ (this is why π^2 appears in $\overline{A}$) by Lemma 6 we have $\alpha_2(0) = 0$. This implies by Lemma 4 that

$$\mu^1(N) = \frac{\mu_{N|g}\mu_g + \alpha_1(N)\mu_{N|n}\mu_n}{\mu_N} \leq 2(\epsilon/\pi)(\mu_g + \mu_n)/\mu_N$$
$$\leq \frac{8(\epsilon/\pi)}{(1-\pi)\left(1 - \frac{4\epsilon}{\pi}\right)\underline{\mu}}.$$

So when this is less than or equal $\overline{B}$ by Lemma 6 we have $\alpha_2(N) = 0$. For $\epsilon \leq \overline{A}/8$ this is

$$\frac{16\epsilon}{\pi(1-\pi)\underline{\mu}} \leq \overline{B}$$

so holds for $\epsilon < \underline{\mu}\overline{A}/16$ which was assumed. □

Lemma 13. *In any equilibrium* $\alpha_1(0) = \alpha_2(0) = 0$.

Proof. We already know this to be true in any pure equilibrium, so we may assume the equilibrium is not pure. From Lemma 11 if $\alpha_1(0) = 0$ then $\alpha_2(0) = 0$ so we may assume this is not the case, that is $\alpha_1(0) > 0$. From Lemma 12 we know that $\alpha_1(1) < 1$. It cannot be that the normal type is indifferent at both $z = 0, 1$ for then by Lemma 1 it must be that $\alpha_2(1) = \alpha_2(0) = \tilde{\alpha}_2$ so that $V_1 = V(\tilde{\alpha}) = V_0$ and that the normal type never provides effort in which case by Lemma 8 we would have a pure strategy equilibrium. Hence either the normal type strictly prefers to provide no effort at $z = 1$ and is willing to provide effort at $z = 0$ or the normal type is indifferent at $z = 1$ and strictly prefers to provide effort at $z = 0$. In either case from Lemma 1 we must have $\alpha_2(1) < \alpha_2(0)$.

The key point is that having the short-run player enter when there is no effort is kind of like winning the lottery - you get something for nothing. If that happens in the state 0 it is particularly good because you are guaranteed that you get to play again. Since $\alpha_2(1) < \alpha_2(0)$ we can write $\alpha_2(0) = \beta + (1-\beta)\alpha_2(1)$ where $\beta > 0$ meaning that in the state $z = 0$ there is a better chance of winning the lottery. We will use this to show that $V(\alpha_2(0)) \geq V(\alpha_2(1))$ so that never provide effort is optimal and the equilibrium must be pure by Lemma 8. □

Lemma 14. *In any non-pure equilibrium* $0 < \alpha_2(1) < 1$, $\alpha_1(N) > 0$, *and* $\alpha_2(N) \geq \alpha_2(1)$.

Proof. First suppose that $\alpha_2(1) = 1$. Since the short-run player must be mixing and by Lemma 13 is not doing so at $z = 0$ the short-run player must be mixing at $z = N$, that is, that $0 < \alpha_2(N) < 1$. Lemma 12 implies that at $z = 1$ the normal type does not strictly prefer to provide effort. Since $\alpha_2(N) < \alpha_2(1)$ Lemma 1 implies that at $z = N$ normal type strictly prefers not to provide effort, so $\alpha_1(N) = 0$. Hence $\mu^1(N) = \mu_{N|g}\mu_g/\mu_N = \epsilon H_{0g}\mu_g/\mu_N$. As $\alpha_2(0) = 0$ by Lemma 13 it follows from Lemma 4 that

$$\mu^1(N) \leq \frac{4\epsilon}{(1-\pi)\left(1 - \frac{4\epsilon}{\pi}\right)\underline{\mu}}$$

as the RHS this is less than $\overline{B}$ by assumption we have $\alpha_2(N) = 0$ a contradiction.

Next suppose that $\alpha_2(1) = 0$. By Lemma 13 we also have $\alpha_2(0) = 0$ so by Lemma 1 the long-run player never provides effort. Hence $\alpha_2(1) > 0$ follows from Lemma 8, a contradiction. We have now shown strict mixing the the short-run player at $z = 1$.

Now we show that since the short-run player is strictly mixing at $z = 1$ then $\alpha_1(N) > 0$. Strict mixing by the short-run player at $z = 1$ implies from Lemma 6 $1 - \overline{B} = 1 - \mu^1(1) = \left([1 - \alpha_1(1)]\mu_{1|n}\mu_n + \mu_{1|b}\mu_b\right)/\mu_1$. From Lemma 3 and Lemma 13 if $\alpha_1(N) = 0$ we have $\mu_{1|n} \leq \alpha_1(1)\mu_{1|n} + 2\epsilon$ and $\mu_{1|b} \leq 2\epsilon$. Hence by Lemma 2 $1 - \mu^1(1) \leq 2\epsilon/(\pi\underline{\mu})$, so for $2\epsilon/(\pi\underline{\mu}) < 1 - \overline{B}$ this is a contradiction.

Since $\alpha_2(N) > 0$ the normal type weakly prefers to provide effort at $z = N$. If $\alpha_2(1) > \alpha_2(N)$ by Lemma 1 this implies the normal type would strictly prefer to provide effort at $z = 1$ contradicting Lemma 12. □

Signal Jamming

Define the *auxiliary system* with respect to $0 \leq \lambda, \gamma \leq 1$ as

$$V_1 = (1-\delta)\tilde{\alpha}_2 + \delta\left[(\tilde{\alpha}_2 + (1-\tilde{\alpha}_2)\pi) V_0 + (1-\tilde{\alpha}_2)(1-\pi)V_N\right]$$

$$V_N = (1-\gamma)(\lambda - c) + \gamma V_1$$

$$V_0 = \frac{\delta(1-\pi)}{1-\delta\pi} V_N.$$

Since in a mixed equilibrium we know from Lemma 12 that $\alpha_1(1) < 1$ so that at $z = 1$ the long-run player must be willing to provide no effort. This system corresponds to providing no effort at $z = 0, 1$. From the contraction mapping fixed point theorem this has a unique solution V_1, V_N, V_0. Define the function $\underline{\Delta}(\tilde{\alpha}_2) \equiv V_1 - V_0$.

Lemma 15. *We have*

$$V_1 = \frac{\delta(1-\pi)(1-\gamma)(\lambda-c) + (1-\delta)\left[1-\delta\pi-\delta(1-\pi)(1-\gamma)(\lambda-c)\right]\tilde{\alpha}_2}{(1-\delta\pi-\gamma\delta(1-\pi)) + \gamma\delta(1-\pi)(1-\delta)\tilde{\alpha}_2}$$

strictly increasing in $\tilde{\alpha}_2$.

Proof. Here we simply solve the linear system and determine the sign of the derivative of V_1. □

Lemma 16. $\underline{\Delta}(\tilde{\alpha}_2)$ *is strictly increasing. There is a solution* $0 < \hat{\alpha}_2 < 1$ *to*

$$\underline{\Delta}(\tilde{\alpha}_2) = \overline{\Delta}(\tilde{\alpha}_2) \equiv \frac{1-\delta}{\delta\left(\tilde{\alpha}_2 + (1-\tilde{\alpha}_2)\pi\right)} c,$$

it and only if

$$c < \delta\frac{(1-\delta\pi-\delta(1-\pi)\left[\gamma+\lambda(1-\gamma)\right])}{1-\delta\pi-\delta^2(1-\pi)},$$

in which case it is unique.

Proof. Here solve V_0 as a function of V_1 from the system. We subtract this from V_1 and find that $\underline{\Delta}(\tilde{\alpha}_2)$ is strictly increasing in V_1. Hence we may apply Lemma 15. Since $\overline{\Delta}(\tilde{\alpha}_2)$ is decreasing there will be a unique intersection if and only if $\overline{\Delta}(0) > \underline{\Delta}(0)$ and $\overline{\Delta}(1) < \underline{\Delta}(1)$. By computation we show that the first condition is always satisfied and the second is the condition on c given as the result. □

Proposition 2. *If* $\epsilon < \underline{\mu}\pi^2(1-\pi)\min\{\overline{B}, 1-\overline{B}\}/32$ *and*

$$c \geq \delta\frac{1}{1+\delta(1-\pi)}.$$

all equilibria are in pure strategies.

Proof. Suppose that $\alpha_1(z), \alpha_2(z)$ is a non-pure equilibrium. If the normal type is willing to provide effort at $z = 1$ we take $\hat{\alpha}_2 = \alpha_2(1)$. If the long-run player strictly prefers to provide no effort at $z = 1$ we show how to construct a $1 > \hat{\alpha}_2 > \alpha_2(1)$ for which the long-run player is indifferent at $z = 1$ and strictly prefers to provide effort at $z = N$. We show that $1-c \geq V(\alpha_2(N)) \geq V(\hat{\alpha}_2)$ and use this to show that at $\hat{\alpha}_2$ we must have $\underline{\Delta}(\hat{\alpha}_2) = \overline{\Delta}(\hat{\alpha}_2)$ for $\lambda = 1$. Applying Lemma 16 then yields the desired condition. $\square$