

Reputation and the Flow of Information in Repeated Games*

Eduardo Faingold

Department of Economics

Yale University

eduardo.faingold@yale.edu

October 2013

Abstract

This paper derives equilibrium payoff bounds from reputation effects for repeated moral hazard games in which a long-run player interacts *frequently* with a population of short-run players and the monitoring technology varies with the length of the period of interaction. The bounds depend on the monitoring technology through the *flow of information*, a measure of signal informativeness (per unit of time) based on relative entropy. Examples are shown where, under complete information, the set of equilibrium payoffs of the long-run player converges, as the period length tends to zero, to the set of *static* equilibrium payoffs, whereas when the game is perturbed by a small ex ante probability on commitment types, reputation effects remain powerful in the frequent interaction limit.

Keywords: Reputation, commitment, imperfect monitoring, frequent interactions.

JEL Classification: C70, C72.

*I thank Mehmet Ekmekci, Drew Fudenberg, Olivier Gossner and George Mailath for stimulating discussions and helpful comments, and the National Science Foundation (grant SES-1227593) for financial support. This paper subsumes and improves upon the earlier analysis of reputations under frequent interactions in [Faingold \(2006\)](#).

1 Introduction

This paper studies reputation effects (in the sense of [Kreps and Wilson \(1982\)](#), [Milgrom and Roberts \(1982\)](#)) in repeated games in which a long-run player interacts *frequently* with a population of short-run players. The main result provides upper and lower bounds (akin to those of [Fudenberg and Levine \(1989, 1992\)](#)) on the set of Nash equilibrium payoffs of the long-run player in the limit as the discount rate and the period length tend to zero (in any order).

In typical studies of repeated games, it is common to interpret comparative statics with the discount factor as an exercise of varying the length of the period of interaction. The limit when the discount factor tends to one is often interpreted as the limit when the period length tends to zero and the monitoring technology is independent of the period length. But, when players observe noisy signals about each other's actions (as in many applications of interest) it may be more natural to assume that the informativeness of those signals is constant over each unit of *real time*. Since this requires the signal distributions to vary with the period length ([Abreu, Milgrom, and Pearce, 1991](#)), it leads to a comparative statics that is different from the traditional one: as the periods shrink, the discount factor tends to one *and* the informativeness of the signals (over the period of interaction) deteriorates. It is now well known that allowing the monitoring technology to vary with the period length in this way can have a dramatic effect on the equilibrium analysis of repeated games: in some cases, the set of (perfect public) equilibrium payoffs of the repeated game converges, as the period length tends to zero, to the (convex hull of the) set of *static* Nash equilibrium payoffs ([Abreu, Milgrom, and Pearce \(1991\)](#), [Faingold and Sannikov \(2011\)](#), [Fudenberg and Levine \(2007, 2009\)](#), [Sannikov and Skrzypacz \(2007\)](#)).

How about when the repeated game is perturbed by a small *ex ante* probability that the long-run player may be a commitment type? Do reputation effects survive in the limit as the period length shrinks to zero and the signal informativeness deteriorates? The main result of this paper, [Theorem 2](#), answers this question in the affirmative, by providing upper and lower bounds on the set of Nash equilibrium payoffs of the long-run player in the limit as the period length and the discount rate tend to zero (in any order). Examples are shown in which the bounds are tight and equal the Stackelberg payoff, whereas without uncertainty over types the equilibrium set would collapse to static Nash.

The intuition for the dramatic difference between the complete and incomplete information results is simple. In the perturbed game, due to the learning process of the uninformed short-run players, the relevant time frame for reputation effects is the “long run”—a reputation is not won in a day—and in the long run the many (poorly informative) signals can be aggregated to become arbitrarily informative. By contrast, in the unperturbed complete

information game, there is no learning about types and non-myopic incentives can only be created by using each period's signals to adjust continuation play to deter deviations. Thus, unlike the perturbed game, the relevant period for incentives in the complete information game is the period of interaction itself—over which the signals are becoming noisier and noisier—, hence the possibility of collapse to static Nash equilibrium.¹

In spite of this clean intuition, the result is more subtle than it may appear at first glance, for it cannot be proved by appealing directly to the classical Fudenberg-Levine bounds (henceforth FL bounds). To see why, let us first briefly review the argument of [Fudenberg and Levine \(1992\)](#). Assume (for simplicity) that there is a unique best-reply to the Stackelberg action in the stage game and that the signal distributions statistically identify the action of the long-run player. Fix an $\varepsilon > 0$ and an arbitrary Nash equilibrium of the perturbed game and consider what happens under a deviation in which the normal type mimics the behavior of the Stackelberg type (i.e, the behavioral type who is committed to playing the Stackelberg action in every period). [Fudenberg and Levine \(1992\)](#) show that, under this deviation, on a set of infinite histories that has probability at least $1 - \varepsilon$, there is an upper bound K on the number of periods in which the short-run players are not playing a best-reply to the Stackelberg action, where K depends on the monitoring technology and on ε , but is independent of the particular equilibrium fixed and hence of of the discount factor. Due to this uniformity, by letting the discount factor tend to 1 and ε tend to zero in this order (assuming a fixed monitoring structure), the expected average discounted payoff of the normal type (under the deviation) converges to the Stackelberg payoff. Hence, the Stackelberg payoff must be a lower bound on (the patient limit of) the set of Nash equilibrium payoffs of the normal type, because it is a lower bound on the expected payoff from a deviation. Similar ideas apply to the upper bound.

The argument above assumes a fixed monitoring technology, and thus a fixed period of interaction Δ . The reason why it is not possible to derive the reputation bounds for games with frequent interactions by taking the limit of the FL bounds is related to how the uniform bound K depends on Δ (through the monitoring technology). In [Section 4](#), a simple example with Poisson signals is provided which shows that K need not scale with Δ . That is, if the period length is cut in half, one might expect the number of periods of non-best replies to (at least approximately) double, but this is not the case for K , which is only an upper bound on the number of non-best replies (albeit a uniform one).

To circumvent this problem the proof of [Theorem 2](#) result relies on a recent result of [Gossner \(2011\)](#), who uses entropy techniques to derive reputation bounds that improve on the FL bounds for fixed discount factor and fixed monitoring. The bounds provided in

¹See a recent paper by [Bohren \(2011\)](#) for an alternative channel by which intertemporal incentives can be created in the frequent interaction limit which does not rely on uncertainty over types.

Theorem 2, which are the limits of the Gossner bounds as the period length tends to zero, rely on a notion of self-confirmed best-reply that is based on an infinitesimal form of relative entropy, which can be interpreted as a measure of the *flow of information* in the game. Several examples are provided to illustrate how this measure of information flow (and hence the bounds) can be calculated, including an example in which the upper and lower bounds converge to the Stackelberg payoff as the period length and the discount rate tend to zero in this order, while applying the same order of limits to the classical Fudenberg-Levine discrete-time result yields uninformative bounds.^{2,3} While the Gossner bounds and the FL bounds always coincide in the patient limit when the monitoring technology is fixed, the limits can be very different when the monitoring and the discount factor vary simultaneously to adjust for the period length.⁴

Reputation effects under frequent interactions have also been studied in [Faingold and Sannikov \(2011\)](#), where it is assumed that the period length is exactly zero, i.e., the interaction takes place directly in continuous time. Unlike this note, which focuses on frequent-interaction bounds on equilibrium payoffs for patient long-run players under general assumptions on the monitoring structure and on the type space, [Faingold and Sannikov \(2011\)](#) provide a characterization of equilibrium behavior (directly in continuous time) for fixed discount rates, by restricting attention to games with Brownian signals and assuming that there is a single commitment type. Another closely related paper is [Faingold \(2013\)](#), which derives analogues of the FL bounds for continuous-time games with signals driven by a Lévy process (which includes games with Brownian and Poisson signals as special cases.). A pair of examples in Section 3 demonstrates that the reputation result for continuous-time games with Lévy signals does not imply the limit of discrete-time games result of the present paper.⁵

²That is, the FL upper and lower bounds converge to the greatest and least feasible payoff of the large player in the stage game, respectively.

³Such an example is not at all pathological: the gap between the limit of the FL bounds as the period length shrinks and the frequent interaction bounds based on the information flow always arises for (nontrivial) games with signals that are sampled from a Poisson process (as in [Abreu, Milgrom, and Pearce \(1991\)](#)).

⁴See also [Ekmekci, Gossner, and Wilson \(2012\)](#) for another reputation result that relies crucially on the improvement of the Gossner bound over the FL bound.

⁵This seemingly pathological failure of upper hemi-continuity (from discrete to continuous time) is due to the nonexistence of a topology on the space of continuous-time behavior strategies under which expected payoffs are continuous and the space of behavior strategies is compact.

2 Baseline Model

2.1. Discrete-time repeated game. We recall the canonical reputation model of [Fudenberg and Levine \(1992\)](#), in which a long-run player faces a sequence of short-run players in an infinitely repeated game with imperfect monitoring in discrete time. At the beginning of period $n = 0, 1, 2, \dots$, the long-run player and the current short-run player simultaneously, and privately, choose actions $a_1^n \in A_1$ and $a_2^n \in A_2$ respectively, where A_1 and A_2 are finite action spaces. A publicly observable signal y^n is then drawn from a measurable space Y according to a probability distribution $q(\cdot | a_1^n, a_2^n)$, where $q : A \rightarrow \Delta(Y)$ and $A := A^1 \times A^2$. At the end of the period, the current short-run player departs from the game, being replaced by a new short-run player at the beginning of the following period. The new short-run player knows the entire past history of public signals.

The long-run player is privately informed of his *type* $\omega \in \Omega$, where Ω is a countable type space. The short-run players are uncertain as to which type of long-run player they face and share a common prior $\mu \in \Delta(\Omega)$, which satisfies $\mu(\omega) > 0$ for all $\omega \in \Omega$. The type space Ω contains *strategic* (or *normal*) types and *behavioral* types. The set of strategic types is denoted $\Omega^s \subset \Omega$ and the overall payoff of a strategic type ω is

$$\sum_{n=0}^{\infty} (1 - \delta) \delta^n g_1(a_1^n, a_2^n, \omega),$$

where $g_1 : A_1 \times A_2 \times \Omega^s \rightarrow \mathbb{R}$ is the stage-game payoff function of the long-run player and $\delta \in (0, 1)$ is his discount factor. A behavioral type is a mechanistic type of long-run player who is committed to playing some mixed action $\alpha_1 \in \Delta(A_1)$ in every period, irrespective of history.⁶ The set of behavioral types, denoted $\Omega^b := \Omega \setminus \Omega^s$, is then naturally identified with a countable subset of $\Delta(A_1)$. Finally, the payoff function of the short-run players is common knowledge and denoted $g_2 : A_1 \times A_2 \rightarrow \mathbb{R}$.

2.2. Strategies and equilibrium. The space of n -period histories of the long-run player is denoted $H_1^n := \Omega \times (A_1 \times Y)^n$ and is equipped with the product σ -field, where $H_1^0 := \Omega$. Likewise, the space of n -period histories of the short-run players is denoted $H_2^n := Y^n$ and is endowed with the product σ -field, where H_2^0 is the one-point set. A (*behavior*) *strategy* of player i ($i = 1, 2$) is a sequence of measurable maps $\sigma_i^n : H_i^n \rightarrow \Delta(A_i)$, $n \geq 0$, such that $\sigma_1(\omega, \cdot) = \omega$ for every behavioral type $\omega \in \Omega^b$. The set of strategies of player i is denoted Σ_i and the set of strategy profiles is denoted $\Sigma := \Sigma_1 \times \Sigma_2$. Given the prior μ , each strategy profile $\sigma \in \Sigma$ induces a unique probability measure $P_{\mu, \sigma}$ over the space of plays $H^\infty := \Omega \times (A_1 \times A_2 \times Y)^\infty$ endowed with the product σ -field. For each $\omega \in \Omega$, the symbol $P_{\omega, \sigma}$ designates the induced probability measure over plays conditional on ω . Expectations

⁶The payoff functions of behavioral types are left unmodeled

with respect to $P_{\mu,\sigma}$ and $P_{\omega,\sigma}$ are denoted $E_{\mu,\sigma}$ and $E_{\omega,\sigma}$, respectively. A profile $\sigma \in \Sigma$ is a (Bayes-)Nash equilibrium of the repeated game $\Gamma_\delta = ((\Omega, \mu), (A_i, g_i)_{i=1,2}, (Y, q), \delta)$ if for every $\sigma'_1 \in \Sigma_1$,

$$E_{\mu,\sigma} \left[\sum_{n=1}^{\infty} (1-\delta)\delta^n g_1(a_1^n, a_2^n) \right] \geq E_{\mu,(\sigma'_1, \sigma_2)} \left[\sum_{n=1}^{\infty} (1-\delta)\delta^n g_1(a_1^n, a_2^n) \right],$$

and for every $n \geq 0$ and $\sigma'_2 \in \Sigma_2$ such that $\sigma'^m_2 = \sigma^m_2$ for every $m \neq n$,

$$E_{\mu,\sigma} [g_2(a_1^n, a_2^n)] \geq E_{\mu,(\sigma_1, \sigma'_2)} [g_2(a_1^n, a_2^n)].$$

For each $\omega \in \Omega^s$, let

$$N_1(\omega, \delta) := \left\{ E_{\omega,\sigma} \left[\sum_{n=1}^{\infty} (1-\delta)\delta^n g_1(a_1^n, a_2^n) \right] : \sigma \text{ is a Nash equilibrium of } \Gamma_\delta \right\},$$

the set of Nash equilibrium interim payoffs of type ω . Finally, define

$$\underline{N}_1(\omega, \delta) := \inf N_1(\omega, \delta) \quad \text{and} \quad \bar{N}_1(\omega, \delta) := \sup N_1(\omega, \delta).$$

2.3. Reputation effects. We review the classical result of [Fudenberg and Levine \(1992\)](#), which provides upper and lower bounds on the set of Nash equilibrium payoffs of patient strategic types. Recall that a mixed action $\alpha_2 \in \Delta(A_2)$ is a *self-confirmed best-reply* to $\alpha_1 \in \Delta(A_1)$, written $\alpha_2 \in B_2(\alpha_1)$, if α_2 is not weakly dominated and $\alpha_2 \in \arg \max_{\alpha'_2} g_2(\alpha'_1, \alpha'_2)$ for some $\alpha'_1 \in \Delta(A_1)$ with $q(\cdot | \alpha'_1, \alpha_2) = q(\cdot | \alpha_1, \alpha_2)$. For each $\omega \in \Omega^s$ define

$$\underline{g}_1(\omega) := \sup_{\alpha_1 \in \Delta(A_1)} \inf_{\alpha_2 \in B_2(\alpha_1)} g_1(\alpha_1, \alpha_2, \omega)$$

and

$$\bar{g}_1(\omega) := \sup_{\alpha_1 \in \Delta(A_1)} \sup_{\alpha_2 \in B_2(\alpha_1)} g_1(\alpha_1, \alpha_2, \omega).$$

The latter is called the *generalized Stackelberg payoff*, which can be greater than the traditional Stackelberg payoff due to imperfect observability. Indeed the set of self-confirmed best-replies can be greater than the set of best-replies, but the two concepts coincide when the long-run player's actions are *identified*, i.e. when $q(\cdot | \alpha_1, \alpha_2) = q(\cdot | \alpha'_1, \alpha_2)$ implies $\alpha_1 = \alpha'_1$ for all $\alpha_1, \alpha'_1 \in \Delta(A_1)$ and all $\alpha_2 \in \Delta(A_2)$ that are not weakly dominated. Finally, the stage game is called *non-degenerate* if there is no undominated pure action $a_2 \in A_2$ with $g_1(\cdot, a_2) = g_1(\cdot, \alpha_2)$ for some $\alpha_2 \in \Delta(A_2) \setminus \{a_2\}$.

Theorem 1 ([Fudenberg and Levine \(1992\)](#)). *Suppose that behavioral types have full support, i.e. Ω^b is a dense subset of $\Delta(A_1)$. Then, for every strategic type ω with $\mu(\{\omega\}) > 0$,*

$$\underline{g}_1(\omega) \leq \liminf_{\delta \rightarrow 1} \underline{N}_1(\omega, \delta) \leq \limsup_{\delta \rightarrow 1} \bar{N}_1(\omega, \delta) \leq \bar{g}_1(\omega).$$

Moreover, if the stage game is non-degenerate and the long-run player's action are identified, then $\bar{g}_1(\omega) = \underline{g}_1(\omega)$ and hence,

$$\lim_{\delta \rightarrow 1} N_1(\omega, \delta) = \{\bar{g}_1(\omega)\}.$$

3 Reputation under Frequent Interactions

Turning to the analysis of reputation effects in games with frequent interactions, embed the discrete-time model of the previous section in continuous time, so each period of interaction has length $\Delta > 0$. The long-run player discounts his flow payoffs exponentially at rate $r > 0$, so his discount factor across periods is $\delta = e^{-r\Delta}$. As in [Abreu, Milgrom, and Pearce \(1991\)](#), the monitoring structure (Y^Δ, q^Δ) is allowed to vary with the period length Δ . In particular, the informativeness of each period's signals may deteriorate as the periods shrink, as when the signals are sampled from an underlying continuous-time process, such as when the noise is driven by a Poisson process (as in [Abreu, Milgrom, and Pearce \(1991\)](#)), a Brownian motion (as in [Sannikov and Skrzypacz \(2007\)](#) and [Fudenberg and Levine \(2007, 2009\)](#)) or, more generally, a Lévy process (as in [Sannikov and Skrzypacz \(2010\)](#)).

The ensuing analysis relies on [Gossner \(2011\)](#), who uses methods from information theory ([Cover and Thomas, 2006](#)) to derive payoff bounds that improve on those of [Fudenberg and Levine \(1992\)](#) for fixed discount factor and monitoring structure. Recall the definition of the *relative entropy* (also called *Kullback-Leibler divergence*) between two probability measures P and Q on a measurable space X :

$$H(Q\|P) := \begin{cases} \int_X \ln\left(\frac{dQ}{dP}\right) dQ, & \text{if } Q \ll P, \\ \infty & \text{otherwise,} \end{cases}$$

where dQ/dP denotes the Radon-Nikodym derivative. The relative entropy is known to be non-negative, equal to zero if $Q = P$ and lower semi-continuous when viewed as a function defined on $\Delta(X) \times \Delta(X)$ mapping into $\mathbb{R} \cup \{+\infty\}$.⁷

For each $\alpha_1, \alpha'_1 \in \Delta(A_1)$ and $\alpha_2 \in \Delta(A_2)$ define

$$h(\alpha_1, \alpha'_1; \alpha_2) := \liminf_{\Delta \rightarrow 0} \frac{1}{\Delta} H(q^\Delta(\cdot|\alpha_1, \alpha_2) \| q^\Delta(\cdot|\alpha'_1, \alpha_2)), \quad (1)$$

which inherits the properties of the relative entropy and is thus non-negative, equal to zero if $\alpha_1 = \alpha'_1$ for any α_2 , and lower semi-continuous when viewed as a function defined on $\Delta(A_1) \times \Delta(A_1) \times \Delta(A_2)$ taking values in $\mathbb{R} \cup \{+\infty\}$. The function h quantifies the *flow*

⁷The lower semi-continuity is relative to the norm topology on $\Delta(X)$, not the weak*-topology.

of information that the short-run players receive about the actions of the long-run player in the limit when the interactions become arbitrarily frequent.

A mixed action $\alpha_2 \in \Delta(A_2)$ is an *information-flow confirmed best-reply* to $\alpha_1 \in \Delta(A_1)$, written $\alpha_2 \in B_2^*(\alpha_1)$, if α_2 is not weakly dominated and $\alpha_2 \in \arg \max_{\alpha_2'} g_2(\alpha_1, \alpha_2')$ for some $\alpha_1' \in \Delta(A_1)$ with $h(\alpha_1, \alpha_1'; \alpha_2) = 0$. The long-run player's actions are *information-flow identified* if $h(\alpha_1, \alpha_1'; \alpha_2) = 0$ implies $\alpha_1 = \alpha_1'$ for all $\alpha_1, \alpha_1' \in \Delta(A_1)$ and undominated $\alpha_2 \in \Delta(A_2)$. For each $\omega \in \Omega^s$ define

$$\underline{g}_1^*(\omega) := \sup_{\alpha_1 \in \Delta(A_1)} \inf_{\alpha_2 \in B_2^*(\alpha_1)} g_1(\alpha_1, \alpha_2, \omega)$$

and

$$\bar{g}_1^*(\omega) := \sup_{\alpha_1 \in \Delta(A_1)} \sup_{\alpha_2 \in B_2^*(\alpha_1)} g_1(\alpha_1, \alpha_2, \omega).$$

Finally, denote by $\underline{N}_1(\omega, r, \Delta)$ and $\bar{N}_1(\omega, r, \Delta)$ the infimum and supremum, respectively, of the set of Nash equilibrium payoffs of type ω in the repeated game with discount rate r and period length Δ (i.e. the repeated game of the previous section with discount factor $e^{-r\Delta}$ and monitoring structure (Y^Δ, q^Δ)).

The following result is the analogue of Theorem 1 for games with frequent interactions. The proof, relegated to the appendix, appeals to the main result in Gossner (2011).

Theorem 2. *Suppose that behavioral types have full support, i.e. Ω^b is a dense subset of $\Delta(A_1)$. Then, for every strategic type ω with $\mu(\{\omega\}) > 0$,*

$$\underline{g}_1^*(\omega) \leq \liminf_{r \rightarrow 0} \liminf_{\Delta \rightarrow 0} \underline{N}_1(\omega, r, \Delta) \leq \limsup_{r \rightarrow 0} \limsup_{\Delta \rightarrow 0} \bar{N}_1(\omega, r, \Delta) \leq \bar{g}_1^*(\omega).$$

Moreover, if the stage game is non-degenerate and the long-run player's action are information-flow identified, then $\bar{g}_1^*(\omega) = \underline{g}_1^*(\omega)$ and hence,

$$\lim_{r \rightarrow 0} \lim_{\Delta \rightarrow 0} N_1(\omega, r, \Delta) = \{\bar{g}_1^*(\omega)\}.$$

What is learned from this result, especially in conjunction with the examples below, is that the appropriate notion of signal informativeness to study reputation effects in the frequent interaction limit is the one given by the information flow h , the infinitesimal version of relative entropy. This is key, for if other notions of informativeness were used, such as the one based on the total variation distance used by Fudenberg and Levine (1992) in discrete time, the resulting payoff bounds may become completely uninformative in the frequent interaction limit, as discussed in the next section.

To illustrate the result, explicit calculations of the information flow h are provided for a few examples of monitoring structures with frequent interactions. The examples will

also be useful to clarify the connection between the frequent-interaction limit and repeated games that are played directly in continuous time.

Example 1 (Sampling from a Poisson process). As in [Abreu, Milgrom, and Pearce \(1991\)](#), assume there is an underlying counting process, $N(t)$, whose instantaneous arrival rate is a function of the action profile. The players can adjust their actions at discrete times $t = 0, \Delta, 2\Delta, \dots$, and at any time t in this grid the value of $N(t)$ is publicly observed before the players get to choose their time t actions. Thus, in the formalism of Section 2, the public signals are the increments $y^n = N(n\Delta) - N((n-1)\Delta)$, so that

$$Y^\Delta = \{0, 1, 2, \dots\} \text{ and } q^\Delta(y|a) = e^{-\lambda(a)\Delta} (\lambda(a)\Delta)^y / y! \text{ for all } y \in \{0, 1, \dots\},$$

where the function $\lambda : A \rightarrow [0, \infty)$ gives the arrival rate. Then, the information flow can be explicitly calculated:

$$h(\alpha_1, \alpha'_1; \alpha_2) = \lambda(\alpha'_1, \alpha_2) - \lambda(\alpha_1, \alpha_2) - \lambda(\alpha_1, \alpha_2) \ln \left(\frac{\lambda(\alpha'_1, \alpha_2)}{\lambda(\alpha_1, \alpha_2)} \right),$$

for every $\alpha_1, \alpha'_1 \in \Delta(A_1)$ and $\alpha_2 \in \Delta(A_2)$. By the concavity of the logarithm, $h(\alpha_1, \alpha'_1; \alpha_2)$ is always nonnegative, and it can be readily verified that $h(\alpha_1, \alpha'_1; \alpha_2) = 0$ if and only if $\lambda(\alpha'_1, \alpha_2) = \lambda(\alpha_1, \alpha_2)$. In particular, if the long-run player has two pure actions, then the game is information-flow identified if and only if the arrival rates induced by the two pure actions are different regardless of the short-run player's mixed action; if the long-run player has three or more pure actions, then the game cannot be information-flow identified.

Thus, if the long-run player has two pure actions, which induce different arrival rates, then the set of Nash equilibrium payoffs of a strategic type will converge to his Stackelberg payoff as the period length and the discount rate tend to zero (assuming a non-degenerate stage game and a prior with full support). By contrast, consider the repeated product-choice game where the long-run player is a firm who chooses whether to provide high or low quality, and the short-run players are consumers who decide whether to buy the high-end or the low-end product of the firm; the stage-game payoffs are:

| | High-end | Low-end |
|------|----------|---------|
| High | 2, 3 | 0, 2 |
| Low | 3, 0 | 1, 1 |

Then, if the Poisson arrivals are “good news,” i.e. $\lambda(\text{High}, a_2) > \lambda(\text{Low}, a_2)$ for every a_2 , then the only sequential equilibrium of the complete information game (when Δ is sufficiently small) is the repetition of the static equilibrium (Low, Low-end) after every history, regardless of the discount rate, by an argument similar to that in [Abreu, Milgrom,](#)

and Pearce (1991).⁸ ◆

Example 2 (Sampling from a Brownian motion). As in Sannikov and Skrzypacz (2007), assume there is an underlying diffusion process, $X(t)$, which evolves according to

$$dX(t) = \mu(a(t)) dt + dZ(t), \quad t \in [0, \infty),$$

where Z is a standard Brownian motion and the drift is given by $\mu : A \rightarrow \mathbb{R}$. Players can adjust their actions at times $t = 0, \Delta, 2\Delta, \dots$, and at any time t in this grid the value of $X(t)$ is publicly observed before the players get to choose their actions $a(t) = (a_1(t), a_2(t))$, so that

$$Y^\Delta = \mathbb{R} \text{ and } q^\Delta(\cdot|a) \sim \mathcal{N}(\mu(a)\Delta, \Delta).$$

Thus, for any $\Delta > 0$ and any mixed-action profile, the induced distribution over public signals is a *mixture* of Gaussian distributions. Unfortunately, the relative entropy of a pair of mixtures of Gaussians does not admit a representation in terms of the relative entropies of the pairs of Gaussians in the support, and no closed-form expression is known. But we are only interested in the flow of information, that is, the limit as $\Delta \rightarrow 0$ of the quotient between the divergence and the period length Δ , and that turns out to have a simple expression:

$$h(\alpha_1, \alpha'_1; \alpha_2) = \frac{1}{2} (\mu(\alpha'_1, \alpha_2) - \mu(\alpha_1, \alpha_2))^2 \text{ for all } \alpha_1, \alpha'_1 \in \Delta(A_1), \alpha_2 \in \Delta(A_2),$$

as shown in the appendix. Thus, $h(\alpha_1, \alpha'_1; \alpha_2) = 0$ if and only if $\mu(\alpha'_1, \alpha_2) = \mu(\alpha_1, \alpha_2)$, hence, if the long-run player has only two pure actions, then the game is information-flow identified if and only if the drifts induced by the two pure actions are different regardless of the mixed action of the short-run player; if the long-run player has three or more actions, then the game cannot be information-flow identified.

Thus, also in this example, the incomplete information game features strong reputation effects, while the equilibrium of the complete information game becomes degenerate as the period length tends to zero, as shown in Fudenberg and Levine (2007, 2009). ◆

Example 3 (Binomial approximation of Brownian motion I). Consider the monitoring technology (Y^Δ, q^Δ) , where

$$Y^\Delta = \{\sqrt{\Delta}, -\sqrt{\Delta}\} \text{ and } q^\Delta(\pm \sqrt{\Delta}|a) = \frac{1}{2} \pm \mu(a)\sqrt{\Delta},$$

⁸Although Abreu, Milgrom, and Pearce (1991) only examine repeated games between two long-run players, similar arguments apply to games between a long-run player and a sequence of short-run players.

for given functions $\mu : A \rightarrow \mathbb{R}$ and $\sigma : A_2 \rightarrow (0, \infty)$. This relates to Brownian motion via Donsker's Invariance Principle: For any fixed $a \in A$, the family of stochastic processes $(S_\Delta)_{\Delta>0}$,

$$S_\Delta(t) = \sum_{n=0}^{\lfloor t/\Delta \rfloor} y_\Delta^n, \quad t \in [0, \infty),$$

where $y_\Delta^1, y_\Delta^2, \dots$ are i.i.d. random variables with distribution $q^\Delta(\cdot|a)$, converges weakly to a diffusion process with drift $\mu(a)$ and volatility 1, as $\Delta \rightarrow 0$. Here, as expected, the information flow h takes exactly the same form as in the previous example, as can be readily verified. However, this is not true for all discrete-time approximations of Brownian motion, as shown by the next two examples. \blacklozenge

Example 4 (Binomial approximation of Brownian motion II). Consider the monitoring technology (Y^Δ, q^Δ) , where

$$Y^\Delta = \{\mu(a)\Delta \pm \sqrt{\Delta} : a \in A\} \quad \text{and} \quad q^\Delta(\mu(a)\Delta \pm \sqrt{\Delta}|a) = \frac{1}{2}$$

for given $\mu : A \rightarrow \mathbb{R}$. Here, too, the cumulative process S_Δ (defined as in the previous example) satisfies the conditions of Donsker's Theorem, and so it converges weakly to a diffusion with drift $\mu(a)$ and volatility 1. Suppose that the long-run player has only two actions and that those two actions induce different values of μ irrespective of the short-run player's action. Then, for every positive Δ , the short-run players can *perfectly* monitor the long-run player (so that $h = \infty$). In particular, the long-run player's actions are identified for every positive Δ . Thus, given any non-degenerate stage game and any prior with full support on behavioral types, the set of Nash equilibrium payoffs of any strategic type converges to his Stackelberg payoff as Δ tends to 0, *irrespective of patience* (i.e., for any fixed positive discount rate r).

Nonetheless, the continuous-time limit game has full support imperfect monitoring: changes of drift induce absolutely continuous changes of the probability measure over the trajectories of the diffusion (over any finite horizon). Indeed, in the continuous-time game, the set of Nash equilibrium payoffs of a strategic type must converge to his Stackelberg payoff *as the discount rate r tends to 0*, as shown in [Faingold \(2013\)](#). For fixed discount rates, however, the equilibrium payoffs in the continuous-time game are typically lower than the Stackelberg payoff ([Faingold and Sannikov, 2011](#)), unlike the frequent interaction limit discussed above, which does not require $r \rightarrow 0$ to yield a reputation effect. \blacklozenge

Example 5 (Binomial approximation of Brownian motion III). Suppose $A_1 = \{0, 1\}$ and consider the monitoring technology (Y^Δ, q^Δ) , where

$$Y^\Delta = \{\pm \sqrt{\Delta}, \pm \sqrt{\Delta}(1 - \Delta)\} \quad \text{and} \quad q^\Delta(\pm \sqrt{\Delta}(1 - a_1\Delta)|a) = \frac{1}{2} \quad \text{for } a_1 \in \{0, 1\}.$$

Then, the cumulative process S_Δ , defined as in Example 3, converges weakly to a standard Brownian motion *irrespective of a* .⁹ Thus, in the continuous time limit the signal is completely uninformative, and therefore only static equilibria are possible. Nonetheless, for each positive Δ , the signal distributions under the two pure actions have disjoint support, hence there is perfect monitoring for each $\Delta > 0$. Just like in the previous example, here a reputation result obtains in the limit as the period length Δ tends to zero *for any fixed discount rate*. But, unlike the previous example, only static equilibria are possible in the continuous-time game, irrespective of patience. So, here, the discontinuity between discrete and continuous time is even more dramatic.¹⁰ ♦

4 Discussion

This section discusses the connection between the frequent interaction bounds based on the information flow h and the limit of the FL bounds as the period length tends to zero. Specifically, an example is provided in which the frequent interaction limit of the FL bounds is completely uninformative, while the frequent interaction bounds based on the information flow h are tight and equal the Stackelberg payoff in the patient limit.¹¹

For the discussion, it is useful to recall how the FL bounds are calculated. Let us first review the main building block in Fudenberg and Levine's (1992) proof of their reputation bounds, which is a Bayesian learning result related to Blackwell and Dubins (1962)'s classical merging theorem.¹² Let (Ω, \mathcal{F}) be a measurable space and $(\mathcal{F}_n)_{n \geq 0}$ a filtration. Given P and Q probability measures on (Ω, \mathcal{F}) with $Q \ll P$, define

$$d_n(Q, P) = \text{ess sup}_{B \in \mathcal{F}_{n+1}} |P(B|\mathcal{F}_n) - Q(B|\mathcal{F}_n)|,$$

where ess sup denotes the essential supremum of a family of measurable functions.¹³

⁹It can be readily verified that the conditions of Donsker's Theorem are satisfied.

¹⁰The reader might be, at first, unsatisfied with this example (and the previous one), because the approximation does not have full support. But the example can be slightly modified to deliver a full support approximation with essentially the same features. Indeed, suppose that when player 1 chooses $a_1 \in \{0, 1\}$ then with probability $1 - \Delta$ the public signal is drawn from $\{\pm\sqrt{\Delta}(1 - a_1\Delta)\}$ with equal probabilities, and with probability Δ it is drawn from $\{\pm\sqrt{\Delta}(1 - (1 - a_1)\Delta)\}$ also with equal probabilities. Such a quadrinomial approximation also satisfies the conditions of Donsker's Theorem, but it has support independent of a_1 for any Δ . It can be readily verified that $h = \infty$ also in this case, and hence a reputation result obtains in the frequent-interaction limit, although in the continuous-time game the monitoring is pure noise.

¹¹This is unlike the discrete-time case with fixed period length, where the improvement of the Gossner bounds over the FL bounds can only be useful when the long-run player is impatient.

¹²See Sorin (1999) for further connections between merging and reputation.

¹³The ess sup operator gives the smallest measurable function that is an almost sure upper bound on a given

Uniform Merging Theorem (Fudenberg and Levine (1992)).¹⁴ For every $\phi \in (0, 1]$, $\varepsilon > 0$ and $\eta > 0$, there exists a non-negative integer K such that, for every P , Q and Q' probability measures on (Ω, \mathcal{F}) with $P = \phi Q + (1 - \phi)Q'$, one has

$$Q\left(\#\{n \geq 0 : d_n(Q, P) \geq \varepsilon\} \geq K\right) \leq \eta. \quad (2)$$

For concreteness, let us review how this result relates to the FL bounds in the context of the product-choice game with Poisson signals of Example 1. In fact, to make things even simpler, let us assume that only the long-run player controls the arrival rate λ . Fix the period length $\Delta > 0$ and consider a behavioral type $\hat{\omega}$ committed to playing High with probability strictly greater than $1/2$. Denote the mixed action of this behavioral type by $\hat{\alpha}_1$. Let σ be an arbitrary Nash equilibrium of the repeated game with period length Δ . Thus, following the notation of Section 2,

$$P_{\mu, \sigma} = \mu(\hat{\omega})P_{\hat{\omega}, \sigma} + (1 - \mu(\hat{\omega}))P_{-\hat{\omega}, \sigma},$$

where

$$P_{-\hat{\omega}, \sigma} = \frac{\sum_{\omega \neq \hat{\omega}} \mu(\omega) P_{\omega, \sigma}}{1 - \mu(\hat{\omega})}.$$

Next, choose $\varepsilon(\Delta) > 0$ small enough that $\sup_y |q^\Delta(y|\alpha_1) - q^\Delta(y|\hat{\alpha}_1)| \leq \varepsilon(\Delta)$ implies $\alpha_1(\text{High-end}) > 1/2$. A fact that will turn out to be useful later is that for any choice of $\varepsilon(\Delta)$ for which the latter implication holds, one has $\limsup_{\Delta \rightarrow 0} \varepsilon(\Delta)/\Delta < \infty$.¹⁵

If the normal firm deviates and mimics the behavioral type $\hat{\omega}$, the payoff from this deviation must be lower than the payoff the firm receives in equilibrium. The uniform merging theorem will help finding a lower bound on the firm's payoff under this deviation. Fix $\eta > 0$ and let $K(\Delta, \eta, \hat{\omega})$ designate the smallest integer K such that inequality (2) holds for $\varepsilon = \varepsilon(\Delta)$, $\phi = \mu(\hat{\omega})$, $P = P_{\mu, \sigma}$, $Q = P_{\hat{\omega}, \sigma}$, $Q' = P_{-\hat{\omega}, \sigma}$ and \mathcal{F}_n equal to the product σ -field on H_2^n . Thus, under the deviation, with probability at least $1 - \eta$ there are at most $K(\Delta, \eta, \hat{\omega})$ periods in which the consumers are expecting the firm to choose high quality

family of measurable functions. The essential supremum above is well-defined because $Q \ll P$ (Neveu, 1975).

¹⁴The Q -almost sure convergence $d_n(Q, P) \rightarrow 0$ is an immediate implication of Blackwell and Dubins's (1962) merging theorem. The uniform merging theorem of Fudenberg and Levine (1992) strengthens the conclusion that $d_n(Q, P) \rightarrow 0$ Q -a.s., under a stronger form of absolute continuity which came to be known as the *grain of truth* condition: the theorem establishes the uniformity of the rate of convergence K over all P , Q and Q' that satisfy $P = \phi Q + (1 - \phi)Q'$ for a fixed size ϕ of the "grain of truth." Merging also plays a central role in the literature on rational learning in games (Kalai and Lehrer (1993)).

¹⁵This follows from the fact that $\sup_y |q^\Delta(y|\alpha_1) - q^\Delta(y|\hat{\alpha}_1)| = |\hat{\alpha}_1(\text{High}) - \alpha_1(\text{High})| \cdot |\lambda(\text{Low}) - \lambda(\text{High})|\Delta + o(\Delta)$ and $\hat{\alpha}_1(\text{High}) > 1/2$.

with probability less than or equal to $1/2$. Hence, with probability at least $1 - \eta$ under the deviation, there are also at most $K(\Delta, \eta, \hat{\omega})$ periods in which the consumers are not consuming the high-end product.¹⁶ Assume, conservatively, that such exceptional periods are precisely the $K(\Delta, \eta, \hat{\omega})$ initial periods. This yields the following lower bound for the payoff the firm receives in any Nash equilibrium:

$$\text{FL}(r, \Delta, \eta, \hat{\omega}) := (1 - \eta)(3 - \hat{\alpha}_1(\text{High}))e^{-r\Delta K(\Delta, \eta, \hat{\omega})}. \quad (3)$$

Recall that $K(\Delta, \eta, \hat{\omega})$ depends only on the period length Δ (through $\varepsilon(\Delta)$), on η and on the behavioral type $\hat{\omega}$. In particular, it is independent of the discount rate r . Thus, letting $r \rightarrow 0$ first, then $\eta \rightarrow 0$ and then $\hat{\alpha}_1(\text{High}) \rightarrow 1/2$, the lower bound $\text{FL}(r, \Delta, \eta, \hat{\omega})$ converges to the Stackelberg payoff of 2.5. This is the argument of [Fudenberg and Levine \(1992\)](#).

Now, consider what happens to the FL bound, $\text{FL}(r, \Delta, \eta, \hat{\omega})$, when we send Δ to 0 before sending r to 0. To figure out this limit, we need to understand how $K(\Delta, \eta, \hat{\omega})$ varies with Δ . Or, more generally, how K in the uniform merging theorem varies with ε . The following result (proved in the appendix) provides sharp estimates:

Lemma 1. *For each $\phi \in (0, 1]$ and $\varepsilon, \eta > 0$, let $K^*(\varepsilon, \eta, \phi)$ be the smallest non-negative integer such that inequality (2) holds for all probability measures P, Q and $Q' = \phi Q + (1 - \phi)Q'$. Then, for every $\phi \in (0, 1]$ and $\eta > 0$,*

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon^2 K^*(\varepsilon, \eta, \phi) < \infty.$$

Furthermore, if $\mathcal{F}_n \subsetneq \mathcal{F}_{n+1}$ for all n , then for every $\gamma > 0$ and $\phi \in (0, 1]$, there exists $\bar{\eta} > 0$ such that for every $0 < \eta < \bar{\eta}$,

$$\lim_{\varepsilon \rightarrow 0} \varepsilon^{2-\gamma} K^*(\varepsilon, \eta, \phi) = \infty.$$

Now, consider the FL bound (3) with $K(\Delta, \eta, \hat{\omega}) := K^*(\varepsilon(\Delta), \eta, \mu(\{\hat{\omega}\}))$. Then, the second part of the lemma above (with $\gamma = 1$) yields $\varepsilon(\Delta)K(\Delta, \eta, \hat{\omega}) \rightarrow \infty$ as $\Delta \rightarrow 0$. But since $\limsup_{\Delta \rightarrow 0} \varepsilon(\Delta)/\Delta < \infty$ (cf. footnote 15), it follows that the *real time* $\Delta \times K(\Delta, \eta, \hat{\omega})$ tends to ∞ as Δ shrinks to 0, and hence

$$\lim_{\Delta \rightarrow 0} \text{FL}(r, \Delta, \eta, \hat{\omega}) = 0,$$

for every $\eta > 0$ small enough and every $\hat{\omega}$ with $\hat{\alpha}_1(\text{High}) > 1/2$. Thus, for the product-choice game with Poisson signals of Example 1, the FL bounds become completely uninformative in the frequent-interaction limit.

¹⁶Recall that whenever the consumers assign probability greater than $\frac{1}{2}$ to high quality, their best-reply is to consume the high-end product.

A Appendix

A.1 Proof of Theorem 2

The main result of [Gossner \(2011\)](#) yields the following lower bound on $\underline{N}_1(\omega, r, \Delta)$:

$$\sup_{\hat{\omega} \in \Omega^b} \underline{w}_{\hat{\omega}}^{\Delta}((1 - e^{-r\Delta}) \ln \mu(\hat{\omega})) = \sup_{\hat{\omega} \in \Omega^b} \underline{w}_{\hat{\omega}}^{\Delta}(r\Delta \ln \mu(\hat{\omega}) + o(\Delta)),$$

where, for each $\hat{\omega} \in \Omega^b$, the function $\varepsilon \mapsto \underline{w}_{\hat{\omega}}^{\Delta}(\varepsilon)$ is the pointwise supremum over all convex functions that are below

$$\varepsilon \mapsto \inf \{g_1(\hat{\omega}, \alpha_2; \omega) : \alpha_2 \in B_{2,\varepsilon}^{\Delta}(\hat{\omega})\},$$

and

$$B_{2,\varepsilon}^{\Delta}(\hat{\omega}) = \{\alpha_2 : \alpha_2 \text{ is undominated and } \exists \alpha_1 \text{ s.t. } H(q^{\Delta}(\cdot|\hat{\omega}, \alpha_2)\|q^{\Delta}(\cdot|\alpha_1, \alpha_2)) \leq \varepsilon, \alpha_2 \in \arg \max_{\alpha_2'} g_2(\alpha_1, \alpha_2')\}.$$

Since for all $\hat{\omega}$, α_1 and α_2 ,

$$H(q^{\Delta}(\cdot|\hat{\omega}, \alpha_2)\|q^{\Delta}(\cdot|\alpha_1, \alpha_2)) \geq h(\hat{\omega}, \alpha_1; \alpha_2)\Delta + o(\Delta),$$

we must have

$$B_{2,(1-e^{-r\Delta})\ln \mu(\hat{\omega})}^{\Delta}(\hat{\omega}) \subseteq B_{2,r \ln \mu(\hat{\omega})+o(1)}^*(\hat{\omega}),$$

and hence,

$$\sup_{\hat{\omega} \in \Omega^b} \underline{w}_{\hat{\omega}}^{\Delta}((1 - e^{-r\Delta}) \ln \mu(\hat{\omega})) \geq \sup_{\hat{\omega} \in \Omega^b} \underline{w}_{\hat{\omega}}^*(r \ln \mu(\hat{\omega}) + o(1)),$$

where $o(1)$ is a function that converges to zero as $\Delta \rightarrow 0$, and $\varepsilon \mapsto \underline{w}_{\hat{\omega}}^*(\varepsilon)$ is the pointwise supremum over all convex functions that are below

$$\varepsilon \mapsto \inf \{u_1(\hat{\omega}, \alpha_2) : \alpha_2 \in B_{2,\varepsilon}^*(\hat{\omega})\}.$$

The result then follows from the lower semi-continuity of the function h , by taking the limit as $\Delta \rightarrow 0$ and $r \rightarrow 0$.¹⁷ The proof for the upper bound is analogous, and hence omitted. Also omitted is the proof that if the game is non-degenerate and the action of the long-run player is information-flow identified then $\underline{g}_1^*(\omega) = \bar{g}_1^*(\omega)$ for every strategic type ω , as this is similar to the proof of Theorem 3.3 of [Fudenberg and Levine \(1992\)](#).

¹⁷See [Gossner \(2011, Lemma 8\)](#) for details of a similar continuity argument.

A.2 Proof for Example 1

Fix $\alpha_1, \alpha'_1 \in \Delta(A_1)$ and $\alpha_2 \in \Delta(A_2)$ and write $\alpha = (\alpha_1, \alpha_2)$, $\alpha' = (\alpha'_1, \alpha_2)$. We have

$$\begin{aligned} H(q^\Delta(\cdot|\alpha)\|q^\Delta(\cdot|\alpha')) &= \sum_{\bar{a}} \alpha(\bar{a}) \exp(-\lambda(\bar{a})\Delta) \sum_{k=0}^{\infty} \frac{\lambda(\bar{a})^k \Delta^k}{k!} \log \frac{\sum_a \alpha(a) \lambda(a)^k \exp(-\lambda(a)\Delta)}{\sum_{a'} \alpha'(a') \lambda(a')^k \exp(-\lambda(a')\Delta)} \\ &= \sum_{\bar{a}} \alpha(\bar{a}) \exp(-\lambda(\bar{a})\Delta) \left(\log \frac{\sum_a \alpha(a) \exp(-\lambda(a)\Delta)}{\sum_{a'} \alpha'(a') \exp(-\lambda(a')\Delta)} \right. \\ &\quad \left. + \Delta \lambda(\bar{a}) \log \frac{\sum_a \alpha(a) \lambda(a) \exp(-\lambda(a)\Delta)}{\sum_{a'} \alpha'(a') \lambda(a') \exp(-\lambda(a')\Delta)} + o(\Delta) \right), \end{aligned}$$

and hence,

$$\begin{aligned} h(\alpha_1, \alpha'_1; \alpha_2) &= \sum_{\bar{a}} \alpha(\bar{a}) \left(\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \log \frac{\sum_a \alpha(a) \exp(-\lambda(a)\Delta)}{\sum_{a'} \alpha'(a') \exp(-\lambda(a')\Delta)} + \lambda(\bar{a}) \log \frac{\lambda(\alpha)}{\lambda(\alpha')} \right) \\ &= \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \log \frac{\sum_a \alpha(a) \exp(-\lambda(a)\Delta)}{\sum_{a'} \alpha'(a') \exp(-\lambda(a')\Delta)} + \lambda(\alpha) \log \frac{\lambda(\alpha)}{\lambda(\alpha')}, \end{aligned}$$

where, by L'Hospital rule,

$$\begin{aligned} \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \log \frac{\sum_a \alpha(a) \exp(-\lambda(a)\Delta)}{\sum_{a'} \alpha'(a') \exp(-\lambda(a')\Delta)} &= \frac{1}{\sum_a \alpha(a) \exp(-\lambda(a)\Delta) \sum_{a'} \alpha'(a') \exp(-\lambda(a')\Delta)} \\ &\quad \times \left[- \sum_a \alpha(a) \lambda(a) \exp(-\lambda(a)\Delta) \sum_{a'} \alpha'(a') \exp(-\lambda(a')\Delta) \right. \\ &\quad \left. + \sum_{a'} \alpha'(a') \lambda(a') \exp(-\lambda(a')\Delta) \sum_a \alpha(a) \exp(-\lambda(a)\Delta) \right] = \lambda(\alpha') - \lambda(\alpha), \end{aligned}$$

which is the desired result.

A.3 Proof for Example 2

Fix $\alpha_1, \alpha'_1 \in \Delta(A_1)$ and $\alpha_2 \in \Delta(A_2)$ and write $\alpha = (\alpha_1, \alpha_2)$, $\alpha' = (\alpha'_1, \alpha_2)$. We have

$$H(q^\Delta(\cdot|\alpha)\|q^\Delta(\cdot|\alpha')) = \int_{-\infty}^{\infty} \log \frac{f^\Delta(y|\alpha)}{f^\Delta(y|\alpha')} f^\Delta(y|\alpha) dy$$

where $f^\Delta(\cdot|a)$ denotes the density function of a Gaussian random variable with mean $\mu(a)\Delta$ and variance Δ . For any $\bar{a} \in A$,

$$\begin{aligned} \frac{f^\Delta(y|\alpha)}{f^\Delta(y|\alpha')} &= \frac{\sum_a \alpha(a) \exp\{-(y - \mu(a)\Delta)^2/(2\Delta)\}}{\sum_{a'} \alpha'(a') \exp\{-(y - \mu(a')\Delta)^2/(2\Delta)\}} \\ &= \frac{\sum_a \alpha(a) \exp\{-(y - \mu(\bar{a})\Delta + (\mu(\bar{a}) - \mu(a))\Delta)^2/(2\Delta)\}}{\sum_{a'} \alpha'(a') \exp\{-(y - \mu(\bar{a})\Delta + (\mu(\bar{a}) - \mu(a'))\Delta)^2/(2\Delta)\}} \\ &= \frac{\sum_a \alpha(a) \exp\{(\mu(a) - \mu(\bar{a}))(y - \mu(\bar{a})\Delta) - (\mu(a) - \mu(\bar{a}))^2\Delta/2\}}{\sum_{a'} \alpha'(a') \exp\{(\mu(a') - \mu(\bar{a}))(y - \mu(\bar{a})\Delta) - (\mu(a') - \mu(\bar{a}))^2\Delta/2\}} \\ &= \frac{\sum_a \alpha(a) \exp\{\sqrt{\Delta}s(a, \bar{a})(y - \mu(\bar{a})\Delta)/\sqrt{\Delta} - \Delta s(a, \bar{a})^2/2\}}{\sum_{a'} \alpha'(a') \exp\{\sqrt{\Delta}s(a', \bar{a})(y - \mu(\bar{a})\Delta)/\sqrt{\Delta} - \Delta s(a', \bar{a})^2/2\}}, \end{aligned}$$

where

$$s(a, a') := \mu(a) - \mu(a') \quad \forall a, a' \in A.$$

Thus,

$$\begin{aligned} H(q^\Delta(\cdot|\alpha)\|q^\Delta(\cdot|\alpha')) &= \\ &= \sum_{\bar{a}} \alpha(\bar{a}) \int \log \left\{ \sum_a \alpha(a) \exp \left[\sqrt{\Delta}s(a, \bar{a})(y - \mu(\bar{a})\Delta)/\sqrt{\Delta} - \Delta s(a, \bar{a})^2/2 \right] \right\} f^\Delta(y|\bar{a}) dy \\ &\quad - \sum_{\bar{a}} \alpha(\bar{a}) \int \log \left\{ \sum_{a'} \alpha'(a') \exp \left[\sqrt{\Delta}s(a', \bar{a})(y - \mu(\bar{a})\Delta)/\sqrt{\Delta} - \Delta s(a', \bar{a})^2/2 \right] \right\} f^\Delta(y|\bar{a}) dy, \end{aligned}$$

hence the change of variables $z := (y - \mu(\bar{a})\Delta)/\sqrt{\Delta}$ yields

$$\begin{aligned} H(q^\Delta(\cdot|\alpha)\|q^\Delta(\cdot|\alpha')) &= \\ &= \sum_{\bar{a}} \alpha(\bar{a}) \int \log \left\{ \sum_a \alpha(a) \exp \left[\sqrt{\Delta}s(a, \bar{a})z - \Delta s(a, \bar{a})^2/2 \right] \right\} f(z) dz \\ &\quad - \sum_{\bar{a}} \alpha(\bar{a}) \int \log \left\{ \sum_{a'} \alpha'(a') \exp \left[\sqrt{\Delta}s(a', \bar{a})z - \Delta s(a', \bar{a})^2/2 \right] \right\} f(z) dz, \end{aligned}$$

where f denotes the density of a Gaussian random variable with mean 0 and variance 1. To compute the integral

$$I_{\alpha, \bar{a}}^\Delta := \int g_{\alpha, \bar{a}}^\Delta(z) f(z) dz,$$

where

$$g_{\alpha, \bar{a}}^\Delta(z) := \log \left\{ \sum_a \alpha(a) \exp \left[\sqrt{\Delta}s(a, \bar{a})z - \Delta s(a, \bar{a})^2/2 \right] \right\},$$

consider a second-order expansion of $g_{\alpha, \bar{a}}^\Delta(z)$ around $z = 0$. By Taylor's formula with Lagrange form of remainder,

$$I_{\alpha, \bar{a}}^\Delta = g_{\alpha, \bar{a}}^\Delta(0) + \frac{1}{2} (g_{\alpha, \bar{a}}^\Delta)''(0) + \frac{1}{6} \int (g_{\alpha, \bar{a}}^\Delta)'''(\theta^\Delta(z)) z^3 f(z) dz,$$

where $\theta^\Delta(\cdot)$ is a function that satisfies $0 < \theta^\Delta(z)/z < 1$ for every $z \neq 0$. We have

$$g_{\alpha, \bar{a}}^\Delta(0) = \log \left\{ \sum_a \alpha(a) \exp \left[-\Delta s(a, \bar{a})^2/2 \right] \right\},$$

and therefore, by L'Hospital rule,

$$\lim_{\Delta \rightarrow 0} \frac{g_{\alpha, \bar{a}}^\Delta(0)}{\Delta} = -\frac{1}{2} \sum_a \alpha(a) s(a, \bar{a})^2,$$

Moreover, it can be readily verified that

$$(g_{\alpha, \bar{a}}^\Delta)''(0) = \frac{\sum_a \alpha(a) s(a, \bar{a})^2 \exp \left[-\Delta s(a, \bar{a})^2/2 \right]}{\sum_a \alpha(a) \exp \left[-\Delta s(a, \bar{a})^2/2 \right]} \Delta - \frac{\left(\sum_a \alpha(a) s(a, \bar{a}) \exp \left[\Delta s(a, \bar{a})^2/2 \right] \right)^2}{\left(\sum_a \alpha(a) \exp \left[-\Delta s(a, \bar{a})^2/2 \right] \right)^2} \Delta,$$

and

$$(g_{\alpha, \bar{a}}^\Delta)'''(\theta^\Delta(z)) = O(\Delta^{3/2}), \quad \text{uniformly in } z.$$

Thus,

$$\lim_{\Delta \rightarrow 0} \frac{(g_{\alpha, \bar{a}}^\Delta)''(0)}{\Delta} = \sum_a \alpha(a) s(a, \bar{a})^2 - \left(\sum_a \alpha(a) s(a, \bar{a}) \right)^2,$$

and

$$\frac{(g_{\alpha, \bar{a}}^\Delta)'''(\theta^\Delta(z))}{\Delta} = O(\Delta^{1/2}), \quad \text{uniformly in } z,$$

and therefore, by Lebesgue's Dominated Convergence Theorem,

$$\begin{aligned} h(\alpha_1, \alpha'_1; \alpha_2) &= \lim_{\Delta \rightarrow 0} \frac{\sum_a \alpha(\bar{a})(I_{\alpha, \bar{a}}^\Delta - I_{\alpha', \bar{a}}^\Delta)}{\Delta} \\ &= \frac{1}{2} \left(\sum_a \alpha'(a) s(a, \bar{a}) \right)^2 - \left(\sum_a \alpha(a) s(a, \bar{a}) \right)^2 \\ &= \frac{1}{2} \sum_{\bar{a}} \alpha(\bar{a}) [(\mu(\alpha') - \mu(\bar{a}))^2 - (\mu(\alpha) - \mu(\bar{a}))^2] \\ &= \frac{1}{2} \sum_{\bar{a}} \alpha(\bar{a}) (\mu(\alpha') - \mu(\alpha)) (\mu(\alpha') + \mu(\alpha) - 2\mu(\bar{a})) \\ &= \frac{1}{2} (\mu(\alpha') - \mu(\alpha))^2, \end{aligned}$$

as was to be shown.

A.4 Proof of Lemma 1

We begin with a few lemmas that are needed in the proof of Lemma 1. The first lemma below is related to the Borel-Cantelli lemma. Under a stronger assumption than the latter, it asserts that, in addition to the events happening infinitely often with probability one, there is positive probability that the upper density of the indices for which the event obtains is bounded above 0.

Lemma 2. *Let A_1, A_2, \dots be a sequence of independent events on a probability space (Ω, \mathcal{F}, P) . Suppose that $c = \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n PA_k > 0$. Then $P(A_n \text{ i.o.}) = 1$ and, in addition,*

$$P \left(\limsup_{N \rightarrow \infty} \frac{\#\{n \leq N : A_n \text{ obtains}\}}{N} \geq c \right) \geq \liminf_{N \rightarrow \infty} P \left(\frac{\#\{n \leq N : A_n \text{ obtains}\}}{N} \geq c \right) \geq \frac{1}{2}.$$

Proof. Since the assumption implies that $\sum_n PA_n = \infty$, it follows from Borel-Cantelli that $P(A_n \text{ i.o.}) = 1$. Let us turn to the statement concerning the upper density. By the Central Limit Theorem,

$$\frac{\sum_{k=1}^n I_{A_k} - PA_k}{\sqrt{\sum_{k=1}^n PA_k}} \rightarrow \mathcal{N}(0, 1),$$

where the convergence is in distribution. Let $0 < c < \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n PA_k$. We have:

$$\begin{aligned} P \left(\frac{1}{n} \sum_{k=1}^n I_{A_k} \geq c \right) &= P \left(\frac{\sum_{k=1}^n I_{A_k} - PA_k}{\sqrt{\sum_{k=1}^n PA_k}} \geq \frac{cn - \sum_{k=1}^n PA_k}{\sqrt{\sum_{k=1}^n PA_k}} \right) \\ &\geq P \left(\frac{\sum_{k=1}^n I_{A_k} - PA_k}{\sqrt{\sum_{k=1}^n PA_k}} \geq 0 \right) \quad \text{for } n \text{ sufficiently large.} \end{aligned}$$

Therefore, denoting by Φ the standard normal distribution function, we have

$$\liminf_n P \left(\frac{1}{n} \sum_{k=1}^n I_{A_k} \geq c \right) \geq 1 - \Phi(0) = 1/2,$$

as required. ■

Lemma 3. *If $p_k = \frac{1}{2}(1 + 1/k^\alpha)$ with $\alpha > 1/2$, then*

$$\sum_{k=1}^{\infty} \ln \left(\frac{1}{4p_k(1-p_k)} \right) < \infty \quad \text{and} \quad \sum_{k=1}^{\infty} \left(\ln \left(\frac{p_k}{1-p_k} \right) \right)^2 < \infty.$$

Proof. Follows directly from the definition of p_k , the inequality $\ln(1+x) < x$ for $x > 0$, and the fact that $\sum_{k=1}^{\infty} k^{-\nu} < \infty$ for all $\nu > 1$. ■

Lemma 4. *Let P, Q and Q' be such that $P = \phi Q + (1 - \phi)Q'$. Let ϕ_n be the posterior probability on the event that the process follows Q , conditional on \mathcal{F}_n . Then $\forall c > 0$,*

$$Q(\phi_n \leq c\phi_0) \leq c.$$

Proof. Since the odds ratio $(1 - \phi_n)/\phi_n$ is a martingale under Q , we have:

$$\begin{aligned} Q(\phi_n \leq c\phi_0) &= Q\left(\frac{1 - \phi_n}{\phi_n} \geq \frac{1 - c\phi_0}{c\phi_0}\right) \\ &\leq \frac{E_Q[(1 - \phi_n)/\phi_n]}{(1 - c\phi_0)/(c\phi_0)} \\ &= \frac{(1 - \phi_0)/\phi_0}{(1 - c\phi_0)/(c\phi_0)} \leq c, \end{aligned}$$

as was to be shown. ■

Lemma 5. *Suppose that $(\mathcal{F}_n)_{n \geq 1}$ is a strictly increasing sequence of atomic σ -algebras. Then, $\forall \nu > 1/2$ and $\forall \phi \in (0, 1]$, $\exists P, Q$ and Q' with $P = \phi Q + (1 - \phi)Q'$, such that $\forall c > 0$,*

$$\liminf_{N \rightarrow \infty} Q(\#\{n \geq 1 : d_n(P, Q) \geq c/N^\nu\} \geq N) \geq 1/2,$$

Proof. First, we prove the result for the case in which the measurable space is the set of all infinite sequences of Heads and Tails and the σ -algebra \mathcal{F}_n is that generated by all histories of length n . Let Q be the probability measure corresponding to i.i.d. draws of an unbiased coin. Let Q' be the probability corresponding to independent draws of a coin such that the probability of Heads in period n is $p_n = 1/2(1 + 1/n^\alpha)$, where $1/2 < \alpha < \nu$. Let $P = \phi Q + (1 - \phi)Q'$ and ϕ_n be the posterior probability that the stochastic process follows Q conditional on period n information.

Notice that $d_n(P, Q) = (p_n - 1/2)(1 - \phi_n)$. Denote by χ_n the $\{0, 1\}$ -valued random variable that takes value 1 if and only if the coin turns out Heads at time n . Explicit

We have thus proved that,

$$\liminf_{n \rightarrow \infty} Q(d_n(P, Q) \geq K/n^\nu) \geq \frac{1}{4}.$$

Therefore, the events $\{d_n(P, Q) \geq K/n^\nu\}$ satisfy the assumption of Lemma 2 and we have

$$\liminf_{N \rightarrow \infty} Q\left(\frac{\#\{n \leq N : d_n(P, Q) \geq K/n^\nu\}}{N} \geq \frac{1}{4}\right) \geq \frac{1}{2}.$$

Since

$$Q\left(\frac{\#\{n \leq N : d_n(P, Q) \geq K/N^\nu\}}{N} \geq \frac{1}{4}\right) \geq Q\left(\frac{\#\{n \leq N : d_n(P, Q) \geq K/n^\nu\}}{N} \geq \frac{1}{4}\right),$$

it follows that

$$\liminf_N Q\left(\#\{n : d_n(P, Q) \geq K/N^\nu\} \geq \frac{1}{4}N\right) \geq \frac{1}{2},$$

which proves the desired result for the case in which $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n \geq 1})$ is the canonical coin tossing space.

If (Ω, \mathcal{F}) is an arbitrary space with a strictly increasing atomic filtration (\mathcal{F}_n) , then, for every n , there are at least two disjoint atomic events A_n, B_n such that $A_n, B_n \in \mathcal{F}_n \setminus \mathcal{F}_{n-1}$. Proceed as in the previous construction, identifying the event A_n with ‘‘Heads’’ and B_n with ‘‘Tails’’. ■

We are now ready to prove Lemma 1.

First, a close look at S. Sorin’s proof of the Fudenberg-Levine uniform merging theorem (Sorin, 1999, Lemma 2.5) reveals that the integer K of inequality (2) can be taken as C/ε^2 , where $C > 0$ is a constant that depends only on η and ϕ . Hence, we have $\limsup_{\varepsilon \rightarrow 0} \varepsilon^2 K^*(\varepsilon, \eta, \phi) \leq \limsup_{\varepsilon \rightarrow 0} \varepsilon^2 K = C < \infty$, which is the first claim.

Consider the second claim. Assume the filtration is strictly increasing. Suppose, towards a contradiction, that for some $q > 0$ and $\phi \in (0, 1]$ there exist a function $(\varepsilon, \eta) \mapsto N(\varepsilon, \eta)$ and a sequence $\eta_k \rightarrow 0$ with

$$\limsup_{\varepsilon \rightarrow 0} \varepsilon^{2-q} N(\varepsilon, \eta_k) < \infty,$$

and such that for all P, Q and Q' with $P = \phi Q + (1 - \phi)Q'$ one has:

$$Q(\#\{n \geq 1 : d_n(P, Q) \geq \varepsilon\} \geq N(\varepsilon, \eta_k)) \leq \eta_k$$

for all $k \geq 1$ and for all $\varepsilon > 0$ sufficiently small.

For each $k \geq 1$, let $M_k = 1 + \limsup_{\varepsilon \rightarrow 0} \varepsilon^{2-q} N(\varepsilon, \eta_k)$. Define a double sequence $\varepsilon_{m,k} = (M_k/m)^{1/(2-q)}$ for all $k, m \geq 1$. Hence we have for all $k \geq 1$,

$$\limsup_{m \rightarrow \infty} N(\varepsilon_{m,k}, \eta_k)/m < 1$$

By Lemma 5, there exist probability measures P, Q and Q' , with $P = \phi Q + (1 - \phi)Q'$, such that for all $k \geq 1$,

$$\liminf_{m \rightarrow \infty} Q(\#\{n \geq 1 : d_n(P, Q) \geq \varepsilon_{m,k}\} \geq N(\varepsilon_{m,k}, \eta_k)) \geq 1/2$$

Since $\eta_k < 1/2$ for k large enough, it follows that for every $\bar{\varepsilon} > 0$ and every k sufficiently large $\exists \varepsilon \in (0, \bar{\varepsilon}]$ such that:

$$Q(\#\{n \geq 1 : d_n(P, Q) \geq \varepsilon\} \geq N(\varepsilon, \eta_k)) > \eta_k,$$

which is a contradiction. The contradiction shows that

$$\lim_{\varepsilon \rightarrow 0} \varepsilon^{2-q} K^*(\varepsilon, \eta, \phi) = \infty$$

for every $q > 0, \phi \in (0, 1]$ and η sufficiently small, as was to be shown.

References

- Abreu, Dilip, Paul Milgrom, and David Pearce (1991), "Information and timing in repeated partnerships." *Econometrica*, 59, 1713–1733.
- Blackwell, D. and L. Dubins (1962), "Merging of opinions with increasing information." *Ann. Math. Statist.*, 38, 882–886.
- Bohren, Aislinn (2011), "Stochastic games in continuous time: Persistent actions in long-run relationships." Univeristy of California, San Diego.
- Cover, Thomas M. and Joy A. Thomas (2006), *Elements of Information Theory*. Wiley, second edition.
- Ekmekci, Mehmet, Olivier Gossner, and Andrea Wilson (2012), "Impermanent types and permanent reputations." *Journal of Economic Theory*, 147, 142–178.
- Faingold, Eduardo (2006), "Building a reputation under frequent decisions." Mimeo., University of Pennsylvania.
- Faingold, Eduardo (2013), "Maintaining a reputation in continuous time." Yale University.

- Faingold, Eduardo and Yuliy Sannikov (2011), “Reputation in continuous-time games.” *Econometrica*, 79, 773–876.
- Fudenberg, Drew and David K. Levine (1989), “Reputation and equilibrium selection in games with a single long-run player.” *Econometrica*, 57, 759–778.
- Fudenberg, Drew and David K. Levine (1992), “Maintaining a reputation when strategies are imperfectly observed.” *Review of Economic Studies*, 59, 561–579.
- Fudenberg, Drew and David K. Levine (2007), “Continuous time models of repeated games with imperfect public monitoring.” *Review of Economic Dynamics*, 10, 173–192.
- Fudenberg, Drew and David K. Levine (2009), “Repeated games with frequent signals.” *Quarterly Journal of Economics*, 124, 233–265.
- Gossner, Olivier (2011), “Simple bounds on the value of a reputation.” *Econometrica*, 79, 1627–1641.
- Kalai, Ehud and Ehud Lehrer (1993), “Rational learning leads to nash equilibrium.” *Econometrica*, 61, 1019–1045.
- Kreps, David and Robert Wilson (1982), “Reputation and imperfect information.” *Journal of Economic Theory*, 27, 253–279.
- Milgrom, Paul and John Roberts (1982), “Predation, reputation and entry deterrence.” *Journal of Economic Theory*.
- Neveu, J. (1975), *Discrete-Parameter Martingales*. Elsevier Science.
- Sannikov, Yuliy and Andrzej Skrzypacz (2007), “Impossibility of collusion under imperfect monitoring with flexible production.” *American Economic Review*, 97, 1794–1823.
- Sannikov, Yuliy and Andrzej Skrzypacz (2010), “The role of information in repeated games with frequent actions.” *Econometrica*, 78, 847–882.
- Sorin, Sylvain (1999), “Merging, reputation and repeated games with incomplete information.” *Games and Economic Behavior*, 29, 274–308.