Reputation and Information Lag

Ichiro Obara
University of California, Los Angeles

V. Bhaskar
University of Texas Austin

October 25, 2019

Abstract

This paper studies the effect of information lag in dynamic moral hazard problems/repeated games. As our basic setting, we take up a canonical repeated product choice game with imperfect public monitoring where one long-run player faces a sequence of short run players. A distinguishing feature of our model is that the public signals about unobservable actions of the long-run player realize with a random lag.

We characterize the best eequilibrium payoff as a function of the lag distribution, which can be achieved exactly when the long-run player is patient enough. The best eequilibrium payoff is an increasing function of the expected waiting time for the first arrival of any public signal, and converges to the first best payoff in the limit as the expected waiting time goes to infinity (with the long-run player's discount rate going to 0). Our closed-form expression of the best eequilibrium payoff implies that the best eequilibrium payoff is higher for one lag distribution than another when the former distribution second order stochastically dominates the latter. So not only the expected length of lag matters, but also the variance of the lag distribution matters to support efficient outcomes.

1 Introduction

The models of repeated interactions help us understand how cooperative behavior, which is difficult to sustain in one-shot interaction, may emerge as an eequilibrium behavior among self-interested individuals. The underlying logic behind this is extremely simple. An agent is willing to cooperate today because she does not like to be punished in the future. She just does not want to waste her future profit by ruining the long-term relationship with her current partner now. This simple logic has been employed to explain so many things such as: why firms can collude by keeping the price high, why a country may not initiate a protective trade policy against other countries, and why some workers are paid more than the market wage, and how a monopoly firm can keep its reputation and enjoy the large price premium etc.

For this mechanism of cooperation and punishment to work, it is essential that each agent's behavior is monitored accurately and promptly. Thus it is important to understand how those long-term relationships are affected when the feedback of information is not perfect or slow.

The effect of information accuracy has been studied extensively. For example, Green and Porter [6] considers a model of dynamic Cournot competition where a firm's production decision is not directly observable by other firms and the market price is used as an imperfect signal of the productions of the firms. They show that the eequilibrium behavior with imperfect monitoring would be drastically different with the eequilibrium behavior with perfect monitoring. In particular, collusion must be associated with an episode of price war in eequilibrium when the firm's action is imperfectly observed. Kandori [7] considers repeated games with imperfect public monitoring and shows that the set of eequilibrium payoffs expands as the information structure improves in the sense of Blackwell.

The timing of information has not attracted as much attention as the accuracy of information, but is equally important. We can think of many situations where information critical for long-term relationships becomes available with lag. For example, consider a manager/CEO who makes an investment decision on multiple long-term projects sequentially. It would take some time before the success or failure of each project becomes apparent. Is it better for the manager to wait before making another investment until we know the outcome of the current project? Or is it OK to let the manager to make another investment decision before any feedback about the current project becomes publicly available? As an another example, consider a central bank who makes a policy decision based on its private information and suppose that a noisy signal about the private information becomes publicly available a few months later. How would the reputation of the central bank be affected by the timing of information? These are all important questions.

In this paper, we take up a canonical repeated product choice game (Mailath and Samuelson, 1.5 [9]) as our benchmark model, where one long-run player ("firm") faces a sequence of short run players ("consumers"). The firm produces a good, which may be of high quality or low quality. The firm can choose either high effort or low effort, which is not observable directly by the customers and determines the distribution of qualities of the goods they consume.

The distinguishing feature of our model is the the quality of each good may not be observed by anybody immediately and becomes publicly available with lag. Moreover the timing of each signal realization could be random. For example, although each consumer knows the quality of the product immediately, this information can be circulated via word-of-mouth first, then becomes a public information only after a certain time has passed. So each consumer decides whether to buy a good or not given his belief about the quality, based on the observed qualities of some random past goods. For example, a customer may know the quality of the goods up to 3 weeks ago, but does not know the quality of any good in most recent two weeks.

Given a distribution of lag, which we assume to be invariant across time, we obtain the closed-form expression of the best eequilibrium payoff, which is exactly achieved when the firm is patient enough. We show that it is an increasing function of the expected waiting time for the first arrival of any public signal, and converges to the first best payoff in the limit as the expected waiting time goes to infinity (and discounting factor δ goes to 1). Our main findings follow from this formula. We show that the best eequilibrium payoff is higher for one lag distribution than another when the former distribution second order stochastically dominates the latter. In particular, this implies that the best eequilibrium payoff is increasing in the order of first order stochastic dominance. This finding echoes and confirms an observation in the literature that information lag is sometimes useful to mitigate the incentive problem. Our paper provides a precise characterization about when information lag is beneficial in a more complex environment where information arrives with lag randomly and identifies the variance of lag as an important factor in achieving more efficient outcome.

Technically the difficulty with our model of stochastic information lag is that it does not have the standard recursive structure ([1]), as a deviating long-run player carries private information about the distribution of delayed signals. We approach this problem by first deriving an upper bound of the eequilibrium payoffs, which does have some recursive structure, then explicitly construct an eequilibrium to achieve the bound.

Related Literature

Fudenberg, Ishii, and Kominers [4] studies general repeated games with private monitoring and stochastic information lag. Fudenberg and Olszewski [5] considers a game between a long-run player and short-run players where the players observe an evolving state variable with random private lag. Both papers focus on the limit case: the former paper proves a folk theorem as $\delta \to 1$ and the latter paper characterizes the eequilibrium payoff in the high frequency limit. On the other hand, we consider a random rag of imperfect public monitoring and obtain a closed-form expression of the best eequilibrium payoff for a given level of discount factor. The idea that information lag is helpful to reduce the incentive cost has appeared in Abreu, Milgrom and Pearce [1]. They study a repeated prisoner's dilemma with imperfect public monitoring and with a specific form of deterministic information lag where the most recent n signals are observed publicly only every n periods. They construct a trigger strategy eequilibrium for each n, show that the eequilibrium payoff increases in n and approximates the full efficient outcome as $n \to \infty$.[1] Our paper delivers similar results with a variety of general stochastic information lag structures, but we allow for a long-run player to carry a private information even in public eequilibrium and we identify the importance of the dispersion of stochastic lag in addition to the length of information lag.

In the next section, we introduce a model of repeated product choice game. We consider a simple example with one period lag in Section 3. We present our main results in Section 4. We first derive an upper bound on the best eequilibrium payoff for the firm, then construct an eequilibrium to achieve this bound explicitly. Section 5 is devoted to discussions. Most of the proofs can be found in the appendix at the end of the paper.

2 Model

Stage Game

A firm (player 1) and a sequence of consumers (player 2) play the following product choice game in every period. The firm chooses the level of effort, which is either H ("high") or L ("low"). In each period, one consumer decides whether to buy the good that the firm produces or not. The stage game payoffs are summarized by the following matrix.

	B("Buy")	DB("Don't Buy")
H	$1, 1$	$-G, 0$
L	$1+G, -T$	$0, 0$

where $G, T > 0$ and $G - T < 1$ so that $(1, 1)$ is the most efficient outcome with transferable utility.[2] The firm's effort level is not observable. The firm's effort determines the distribution of the quality of its product, which is either g ("good") or b ("bad"). Let p and q be the probability that the quality is bad when the firm chooses a high effort or low effort respectively. We assume $0 < p < q$, so a product is more likely to be of a good quality when the firm chooses a high effort. The above stage game payoffs can be interpreted as follows. The good is sold at a fix price. Each consumer's payoff is the expected payoff, which is the expected quality minus the payment. It is worth purchasing a good only when the firm chose a high effort in producing it. The extra cost of exerting a high effort instead of a low

[1] Kobayashi and Ohta [8] studies a related model of multimarket contacts and shows that the eequilibrium in Abreu et al. [1] is in fact an optimal one for each n when δ is large enough.

[2] The detail of the payoffs do not matter for our results as long as (1) L is the dominant action and (2) "Buy" is the unique best response to H and "Don't Buy" is the unique best response to L.

effort is G.

Lag Structure

Let s_t be the quality of the good that is produced and consumed in period t, which would be never available if the good is not purchased in period t. The distinctive feature of our model is that s_t is not observed immediately (by any player) and observed with a random lag. We assume that the distribution of lag is invariant across the time and is given by $\lambda_n, n = 0, 1, ..., N$ (N can be ∞) independent of t, where λ_n is the probability that s_t becomes publicly available in the end of period $t + n$ to the firm and any consumer from period $t + n + 1$. The standard case of no information lag corresponds to $\lambda_0 = 1$. We call lag structure $\{\lambda_n\}$ N**-stochastic information lag** when N is the largest integer such that $\lambda_N > 0$.

The firm discounts future payoffs by discount factor $\delta \in (0, 1)$. We focus on pure strategy perfect Bayesian eequilibrium. Note that this means that each consumer knows the firm's action and best responds to it on the eequilibrium path. So only (H, B) and (L, DB) can be played on the eequilibrium path. Therefore the best eequilibrium payoff is bounded by 1 and every eequilibrium payoff can be ranked: the best eequilibrium payoff for the firm is the best eequilibrium payoff for the consumers in terms of the discounted average.

3 An Example with One Period Lag

In this section, we present the simplest example with deterministic 1 period lag. As a benchmark, let's first consider the standard case with no information lag, i.e. $\lambda_0 = 1$. In this case, we know that the following strategy achieves the best pure strategy eequilibrium: play (H, B) in the first period, if g is observed, then restart the game (hence continue with the best eequilibrium). If b is observed, then restart the game with probability $1 - \pi$ and use Nash reversion with probability π. Hence the best eequilibrium (discounted average) payoff v_0^* satisfies the following recursive equation.

$$v_0^* = (1 - \delta) + \delta \left[(1 - p)v_0^* + p(1 - \pi)v_0^*\right]$$

To maximize the payoff, we need to minimize π: the probability of punishment. So π is chosen to satisfy the incentive constraint with equality: $(1 - \delta)G = \delta(q - p)\pi v_0^*$. By solving these equations, we obtain $v_0^* = 1 - \frac{G}{\frac{q}{p} - 1}$ ([9], 7.6.2.). In particular, the firm cannot achieve the efficient outcome $(1, 1)$.

Now suppose that the quality of each product is observed with 1-period lag, i.e. $\lambda_1 = 1$. To derive the best eequilibrium payoff heuristically, we pretend that we can ignore the incentive constraint in the second period, then derive the best eequilibrium payoff v_1^*. Decompose the repeated game into the first two periods and the continuation game. Clearly it is optimal to have (H, B) played in the first two periods. We still need to provide an incentive for the firm to play H in the 1st period. The most efficient way to do so is to continue with the best eequilibrium from period 3 when g is observed in the end of the second period and impose some punishment when b is observed in the end of the second period. Hence v_1^* satisfies the following recursive equation.

$$v_1^* = (1-\delta^2) + \delta^2 \left[(1-p)v_1^* + p(v_1^* - \Delta)\right]$$

, where Δ is the punishment relative to the best eequilibrium payoff. We need to minimize Δ as before to obtain the best eequilibrium payoff that satisfies this recursive equation. So we choose Δ satisfying $(1-\delta) = \delta^2(q-p)\Delta$. Then we obtain $v_1^* = 1 - \frac{1}{1+\delta}\frac{G}{\frac{q}{p}-1}$. Note that information lag reduces the incentive cost, hence helps achieving a higher value.

Of course, v_1^* is still just an upper bound of the best eequilibrium payoff. We need to justify why the second period incentive constraint is not binding. In fact, we can construct an eequilibrium in such a way that Δ not only takes care of the incentive constraint in the 1st period, but also the incentive constraint in the 2nd period without any additional cost. The following strategy profile is such an eequilibrium.

- 1st period: Play (H, B);
- 2nd period: Play (H, B) (on and off path). If the realized quality s_1 (of the first period good) is g, then restart the game from the 1st period. If it is b, then proceed to the 3rd period.
- 3rd period: Play (H, B) (on and off path). If the realized quality s_2 is g, then restart the game from the 1st period. If it is b, then randomize between going to the punishment period and going to the above "2nd period" and treating the current period as the "1st period" (hence the signal associated with the current effort will used as s_1 above).
- Punishment period: Play (L, DB) and randomize between restarting the game from the 1st period and playing (L, DB) forever. The probability to restart the game depends on the realized signal s_3 and would be higher when $s_3 = g$ and lower when $s_3 = b$.

We set the randomization in the 3rd period so that the 1st period incentive constraint is binding. Then we can show that the second period incentive constraint is indeed automatically satisfied. Note that two consecutive bad signals (in the 2nd period and the 3rd period) would trigger a move to the punishment period, and the firm would have stronger incentive to play H in the second period because she is one period closer to the possible punishment. As for the incentive in the 3rd period, note that the firm does not know whether the game goes back to the 2nd period (hence the current period will be treated as the 1st period) or the punishment period. We set the randomization in the punishment period so that the firm is indifferent between H and L in the 3rd period conditional on the event that the game moves to the punishment period. Then the firm's incentive constraint conditional on the event that the game moves to the 2nd period is the same as the incentive constraint in the 1st period, hence is satisfied as well.

When the quality of each good is observed with deterministic N-period lag, we can use a similar heuristic derivation of the best eequilibrium payoff v_N^*, ignoring the incentive constraints from period 2 to period $N+1$. The best eequilibrium payoff v_N^* with deterministic N-period lag is $v_N^* = 1 - \frac{1}{1+\sum_{n=1}^{N}\delta^n}\frac{G}{\frac{q}{p}-1}$. Hence the best eequilibrium payoff v_N^*, would increase as N increases as long as the firm is patient enough. In the limit as $N \to \infty$ (while $\delta \to 1$ at the same time), the incentive cost vanishes and the first best payoff 1 is approximately achieved.

Next we try to illustrate how a more dispersion of lag would hurt the firm. Suppose that a quality is observed immediately or with 2-period lag with equal probability, i.e. $\lambda_0 = \lambda_2 = 0.5$. So the expected length of lag is still 1, but now the lag is random. This time we assume that the best eequilibrium is played as the continuation eequilibrium as soon as s_i is observed in the ith period for any $i = 1, 2, 3$ and ignore the incentive constraint of the second period and the third period conditional on the event that no signal is observed in the 1st period and the 2nd period respectively.

Then the best eequilibrium payoff v_λ^* would satisfy the following recursive equation.

$$\begin{aligned} v_\lambda^* &= 0.5\left[(1-\delta) + \delta\left\{(1-p)v_\lambda^* + p(v_\lambda^* - \Delta_0)\right\}\right] + 0.25\left[(1-\delta^2) + \delta^2 v_\lambda^*\right] \\ &+ 0.125\left[(1-\delta^3) + \delta^3 v_\lambda^*\right] + 0.125\left[(1-\delta^3) + \delta^3\left\{(1-p)v_\lambda^* + p(v_\lambda^* - \Delta_2)\right\}\right] \end{aligned}$$

where Δ_0 is the punishment when the quality s_1 of period-1 good is revealed to be bad immediately, and Δ_2 is the expected punishment conditional on s_1 being revealed to be bad in the end of period 3 and no other public signal is observed. Δ_0 and Δ_2 satisfy the following 1st period incentive constraint with equality: $(1-\delta)G = 0.5\delta(q-p)\Delta_0 + 0.125\delta^3(q-p)\Delta_2$. Then we obtain $v_\lambda^* = 1 - \frac{1}{1+0.5\delta+0.25\delta^2}\frac{G}{\frac{q}{p}-1}$. Note that this is strictly less than v_1^*, even when $\delta \to 1$.

Our theorem generalizes this observation and show that the best eequilibrium payoff with lag distribution $\{\lambda_n\}$ is given by

$$v_\lambda^* = 1 - \frac{1}{1+E(\lambda,\delta)} \frac{G}{\frac{q}{p}-1}$$

where $E(\lambda,\delta)$ the expected discounted waiting time before the first arrival of a public signal. In the above examples, $E(\lambda,\delta)=0$ with no lag, $E(\lambda,\delta)=\sum_{n=1}^{N}\delta^n$ with deterministic N-period lag, and $E(\lambda,\delta)=0.5\times 0+0.25\delta+0.25(\delta+\delta^2)=0.5\delta+0.25\delta^2$ with the stochastic lag. We use this formula to show that the best eequilibrium payoff would be higher when the lag distribution is higher in the order of first order stochastic dominance/second order stochastic dominance.

4 Main Results

4.1 Upper bound

We provide a semi-formal, but still a heuristic way to derive an upper bound over all the eequilibrium payoffs. The proof that it is indeed an upper bound is provided in the appendix. Then we actually construct an eequilibrium to achieve the bound in the next subsection. Let $\overline{v}_\lambda(\delta)$ be a least upper bound of all pure strategy eequilibrium payoffs given discount factor δ.

Take any N-stochastic information lag structure. Let $\gamma_n, n=0,1,...N$ be the probability that n periods pass before the first arrival of a public signal. γ_n is completely determined by the lag distribution, namely, $\gamma_0=\lambda_0$, $\gamma_1=\lambda_1+(1-\lambda_0-\lambda_1)\lambda_0$ etc. Let $E(\lambda,\delta)=\sum_{n=1}^{N}\gamma_n\left(\sum_{k=1}^{n}\delta^k\right)$ be the expected discounted waiting time before the first arrival of a public signal.

Suppose that we can ignore the incentive constraint of the firm before the arrival of a first public signal, except for the first period. So this is as if the firm can commit to choose high effort before the first public signal arrives except for the first period. We assume that the firm chooses high effort in every period until the first public signal arrives. When the first signal arrives, assign $\overline{v}_\lambda(\delta)$ as the continuation payoff. This is consistent with our assumption that we can ignore the incentive constraints except for the 1st period incentive constraint. When the first signal arrives in period n and the set of realized signals include s_1, assign $\overline{v}_\lambda$ as the continuation payoff when $s_1=g$ and $\overline{v}_\lambda(\delta)-\Delta^n$ when $s_1=b$, where Δ^n is used as a punishment to provide the incentive for the firm to exert a high effort in the 1st period.

Then the firm's payoff would be given by

$$\sum_{n=0}^{N}\gamma^n\left[(1-\delta^{n+1})+\delta^{n+1}\left(\frac{\gamma_n^1}{\gamma_n}\left\{(1-p)\overline{v}_\lambda(\delta)+p\overline{v}_\lambda(\delta)-\Delta^n\right\}+\frac{\gamma_n-\gamma_n^1}{\gamma_n}\overline{v}_\lambda(\delta)\right)\right]$$

where γ_n^1 is the probability that the first public signals arrive in period $n+1$ and they include s_1.

To minimize the expected cost associated with these punishments, we choose $\Delta^n, n = 0, 1, ..., N$ to satisfy the following binding incentive constraint in the 1st period.

$$(1-\delta)G=\sum_{n=0}^{N}\gamma_n^1\delta^{n+1}(q-p)\Delta^n$$

We can write this as

$$\frac{(1-\delta)G}{\frac{q}{p}-1}=\sum_{n=0}^{N}\gamma_n^1\delta^{n+1}p\Delta^n$$

Using this, we can eliminate $\Delta^n, n = 0, 1, ..., N$ from the above expression of the firm's payoff to obtain

$$\sum_{n=0}^{N}\gamma^n\left[(1-\delta^{n+1})+\delta^{n+1}\overline{v}_\lambda(\delta)\right]-\frac{(1-\delta)G}{\frac{q}{p}-1}$$

This value should be larger than any eequilibrium payoff, hence larger than $\overline{v}_\lambda(\delta)$ as follows.

$$\overline{v}_\lambda(\delta)\le\sum_{n=0}^{N}\gamma^n\left[(1-\delta^{n+1})+\delta^{n+1}\overline{v}_\lambda(\delta)\right]-\frac{(1-\delta)G}{\frac{q}{p}-1}$$

This gives us an upper bound of the set of pure strategy eequilibrium payoffs.

$$\overline{v}_\lambda(\delta)\le 1-\frac{1}{1+E(\lambda,\delta)}\frac{G}{\frac{q}{p}-1}$$

If this value is negative, then the best eequilibrium payoff is 0 and achieved by the repetition of the Nash eequilibrium. So we obtain the following proposition.

Proposition 1. *Given information lag structure $\{\lambda_n\}$ and discount factor $\delta \in [0, 1)$, every pure strategy perfect Bayesian eequilibrium payoff is bounded above by*

$$\max\left\{1-\frac{1}{1+E(\lambda,\delta)}\frac{G}{\frac{q}{p}-1},\ 0\right\}$$

The formal proof of this proposition is in the appendix. In the proof, we set up an auxiliary finitely repeated games with stochastic termination where the game is terminated and the firm receives a terminal payoff as soon as the first public signal arrives or the $N+1$st period is over, whichever comes first. We can assign any terminal payoff that is equal to or less than $\overline{v}_\lambda(\delta)$ as a function of the realized public signals and the firm's actions in all the periods except for the first period. We show that this auxiliary game can replicate any eequilibrium in the original game with N-stochastic information lag, hence we obtain a bound of eequilibrium payoffs by maximizing the best eequilibrium payoff over all possible selection terminal payoffs (for ∞-stochastic information lag, we obtain a bound by using a longer and longer horizon ($N \to \infty$)). Since this bound is larger than or equal to $\overline{v}_\lambda(\delta)$, we have a recursive expression regarding $\overline{v}_\lambda(\delta)$, from which we obtain the above expression that bounds $\overline{v}_\lambda(\delta)$ from the above.

4.2 Best Eequilibrium

We characterize the best eequilibrium payoff as a function of information lag structure and other parameters by explicitly constructing an eequilibrium that achieves the above upper bound when the firm is patient enough.

Theorem 1. *Suppose that the lag structure is given by N-stochastic information lag $\{\lambda_n\}$ for some finite $N < \infty$, then there exists $\underline{\delta}$ such that the best pure strategy eequilibrium payoff $v^*_\lambda(\delta)$ is given by*

$$v^*_\lambda(\delta) = 1 - \frac{1}{1+E(\lambda,\delta)} \frac{G}{\frac{q}{p}-1}$$

for any $\delta \in [\underline{\delta}, 1)$, assuming that the above expression is positive.

Proof. See the appendix. □

Remark: For $N = \infty$, we cannot achieve the bound $\overline{v}_\lambda(\delta)$ exactly, but we can construct a sequence of eequilibrium that approximates $\overline{v}_\lambda(\delta)$ as $\delta \to 1$.

Below we provide a description of eequilibrium strategy that achieves the above value when the firm is patient enough.[3] Our construction generalizes the construction for the deterministic 1-period lag case in the previous section. The eequilibrium consists of 3 phases, one transition period, and Nash reversion.

- **Phase 1:** Always play (H, B) (on and off the eequilibrium path). Reset the game and start a new Phase 1 as soon as any public signal is observed. If no public signal is observed for N periods, then go to the transition period.

[3] This description can be translated into a more formal expression using a finite state automaton.

- **Transition Period:** Play (H, B). Note that s_1 (the signal associated with the first period action in the preceding Phase 1) realizes with probability 1 conditional on the game reaching to the transition period. If only s_1 is observed and $s_1 = g$ or any other public signal is observed in this period, then reset the game and start a new Phase 1. If only s_1 is observed and $s_1 = b$, then go to Phase 2.

- **Phase 2:** Always play (H, B) (on and off the eequilibrium path). If any public signal associated with the firm's actions in current Phase 2 is observed, then reset the game and start a new Phase 1. Otherwise, if any public signal associated with the preceding Phase 1 and the preceding transition period realizes with less then N period lag or realizes with N period lag and it is a good signal g in the $n(\leq N)$th period of Phase 2, then go to the $n + 1$th period of a new Phase 1 (to the transition period if $n = N$) and treat the t th period of current Phase 2 for $t = 1, ..., n$as the t th period of new Phase 1 retroactively. Suppose that none of the above has occurred for N periods. Then N consecutive bad signals should have been observed where the nth bad signal is associated with the firm's action in the $n + 1$st period in the preceding Phase 1 (or the firm's action in the preceding transition period for $n = N$). Given this event, randomize between going to a new transition period (and treating the current Phase 2 as the Phase 1 preceding the transition period) and going to Phase 3.

- **Phase 3:** Play (L, NB) (on and off the eequilibrium path) for N periods. In the end of this phase, reset the game and start a new Phase 1 with probability $\beta(z)$ or start Nash reversion (play (L, NB) forever) with probability $1 - \beta(z)$, where $z \in \{g, b\}^N$ is an N-tuple of public signals associated with the actions in the preceding Phase 2.

Here is a high level intuition for why this is an eequilibrium. A punishment would occur with some probability when and only when $N + 1$ consecutive bad signals with longest lag (N-period lag) associated with the actions in Phase 1 and the transition period are observed. We treat the firm's actions from the the 1st period of Phase 1 to the transition period in a symmetric way. Hence if the incentive constraint for the 1st period action is satisfied, then the incentive constraint in any later period in Phase 1 and the transition period would be satisfied, because the 1st period is most distant from the timing of possible punishment. We choose the probability of punishment (probability to move to Phase 3 at the end of Phase 2) in such a way that the first period incentive constraint is binding to avoid any loss of efficiency. The incentive constraints in Phase 3 are trivially satisfied because the firm's action does not affect future payoffs at all and the firm plays a dominant action L there. In Phase 2, the firm is uncertain about whether the game would move to Phase 3 or the current Phase 2 would be treated as Phase 1. We choose the probability to go back to Phase 1 in the end of Phase 3 so that the firm is indifferent between high effort and low effort in Phase 2 conditional on moving to Phase 3 at the end. Hence we can check the firm's incentive in Phase 2 conditional on it being treated as Phase 1. Then such incentive constraints in Phase 2 are identical to the ones in in Phase 1, hence satisfied.

4.3 Comparative Statics in Lag Structure

Our formula shows that the best eequilibrium payoff is increasing in the expected waiting time for the first public signal (with δ suitably adjusted). So we can do comparative statics of the best eequilibrium payoff with respect to the change of lag structures very easily, we just need to check how $E(\lambda, \delta)$ depends on $\{\lambda_n\}$.

Remember that $E(\lambda, \delta) = \sum_{n=1}^{N} \gamma_n (\sum_{k=1}^{n} \delta^n)$. This is the expected value of increasing and concave function on integers with respect to the distribution of waiting time for the first public signal. Hence any change of lag structure that would increase this distribution in the order of second order stochastic dominance would increase the the best eequilibrium payoff.

Our main comparative statics result with respect to information lag is the following theorem.

Theorem 2. *Suppose that a lag structure $\{\lambda_n\}$ second order stochastically dominates a lag structure $\{\lambda'_n\}$. Then $E(\lambda, \delta) \geq E(\lambda', \delta)$ for any $\delta \in [0, 1)$, hence there exists $\underline{\delta}$ such that $v^*_\lambda(\delta) \geq v^*_{\lambda'}(\delta)$ for any $\delta \in [\underline{\delta}, 1)$.*

If a lag distribution changes in such a way that the new lag distribution FOSD the original distribution, each public signal would realize later with higher probability, so it is intuitively clear that the new distribution of the first public signal FOSD the corresponding original distribution as well.

What is more subtle is the case where a lag distribution changes so that the new lag distribution SOSD the original lag distribution. Suppose that the original lag variable is a new lag variable with 0 mean noise (the variance of the noise can depend on the realized lag of the new lag variable), this is the best case scenario for the old lag distribution. Take the timing of the first public signal for the new lag distribution, which is a random variable. Let's add the 0 mean noise to this timing, then the derived distribution of the timing is already second order stochastically dominated by the distribution with the new lag distribution. Furthermore, the actual timing associated with the original lag distribution is even earlier (thus FOSDed further), because a 0 mean noise is applied to every signal, not just the one that realizes earliest, hence the first public signal may change, especially the first public signal is hit by a positive lag shock (see the random lag example in Section 3).

5 Discussion

5.1 A Simpler Optimal Eequilibrium

Here we illustrate that we can find a simpler optimal eequilibrium for some special class of information lag.

Suppose that the lag structure is given by deterministic N-period lag (i.e. $\lambda_N = 1$) Divide the repeated game into two independent games. One game starts at period 1, and includes period 1 to $N+1$, then period $2(N+1)+1$ to $3(N+1)$ and so on. The other game starts at period $N+2$ then continues until $2(N+1)$, then $3(N+1)+1$ to $4(N+1)$ and so on. So each game consists of a sequence of $N+1$ -period blocks.

Then we can construct an eequilibrium that is similar to the one in [1]. The play is either in the cooperative phase or the Nash reversion phase for each game. The Nash reversion phase is an absorbing one, where the players play (L, NB) forever. In the cooperative phase, (H, B) is played in every period on and off the eequilibrium path. No signal associated with the actions in the current cooperative phase will be observed until the very end of the phase (the signal about the first period action will be observed in the end of the $N+1$ th period before moving to the other independent game). Since one block is separated from the next one by $N+1$ periods in each game, every signal associated with one block will have realized by when the next block of the same game starts. In the beginning of the next block, the cooperative phase would restart if there is an even one good signal g among $N+1$ public signals. If every signal is b, then the cooperative phase restarts with probability $1-\pi$ and the Nash reversion starts with probability π.

The only binding incentive constraint is the one in the first period:

$$(1-\delta)\,G = \delta^{2(N+1)}\pi(q-p)p^N V(\delta).$$

where $V(\delta)$ is the the eequilibrium payoff in the beginning of the cooperative phase for each component game , which satisfies the recursive equation:

$$V(\delta) = \left(1-\delta^{N+1}\right) + \delta^{2(N+1)}(1-\pi p^{N+1})V(\delta)$$

By solving this, we obtain

$$V(\delta) = \frac{1-\delta^{N+1}}{1-\delta^{2(N+1)}} - \frac{1-\delta}{1-\delta^{2(N+1)}}\frac{G}{\frac{q}{p}-1}$$

If we add the eequilibrium payoffs of the two component games, then we obtain

$$V(\delta) + \delta V(\delta) = 1 - \frac{1}{1 + \sum_{n=1}^{N} \delta^n} \frac{G}{\frac{q}{p} - 1},$$

which is exactly the best eequilibrium payoff $V^*(\delta)_N$ with deterministic N-information lag.

References

[1] Abreu, D., P. Milgrom, and D. Pearce, "Information and Timing in Repeated Partnerships," *Econometrica* 59 (1991) 1713-1733.

[2] Abreu, D., D. Pearce, and E. Stacchetti, "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58 (1990) 1041–1063.

[3] Fuchs, W., "Contracting with Repeated Moral hazard and Private Evaluations," *American Economic Review* 97, no. 4 (2007) 1432-1448.

[4] Fudenberg, D., Y. Ishii, and S. D. Kominers, "Delayed-Response Strategies in Repeated Games with Observation Lags," *Journal of Economic Theory* 150 (2014) 487-514.

[5] Fudenberg, D., and W. Olszewski, 'Repeated Games with Asynchronous Monitoring of an Imperfect Signal," *Games and Economic Behavior* 72, no. 1 (2011) 86-99.

[6] Green, E. J., and R. H. Porter, "Noncooperative Collusion under Imperfect Price Information," *Econometrica* 52 (1984) 87-100.

[7] Kandori, M., "The Use of Information in Repeated Games with Imperfect Monitoring," *The Review of Economic Studies* 59, no. 3 (1992) 581-593.

[8] Kobayashi, H., and K. Ohta, "Optimal Collusion Under Imperfect Monitoring in Multimarket Contact," *Games and Economic Behavior* 76, no. 2 (2012) 636-647.

[9] Mailath, G. and L. Samuelson, *Repeated Games and Reputations* (2006), Oxford University Press.

6 Appendix: Proofs

6.1 Proof of Proposition 1

TBA

6.2 Proof of Theorem 1

In the following, we verify that the strategy suggested in the main text is indeed an eequilibrium and achieves the bound for the firm.

First we check one-shot deviation constraints for the firm at every history (including off-path histories). Note that consumers' incentive constraints are trivially satisfied.

Incentive in Phase 3.

All the incentive constraints in Phase 3 are trivially satisfied because L is the strictly dominant action and all the public signals generated by the actions in this phase will not affect the future play of the game at all.

Incentive in Phase 1 and at the transition period

Consider the k th period of Phase 1 for $k = 1, ..., N$ and the transition period on the eequilibrium path. The firm's action in this period matters through a realization of s_k if and only if all $s_n, n = 1, ..., N + 1$ realizes with the longest lag (each s_n realizing in period $N + n$, which is the $n - 1$ th period in the Phase 2 (or the transition period for $n = 1$)), $s_n = b$ for $n = 1, ..., N + 1$ except for $n = k$, and any $s_n, n = N + 2, ..., 2N + 1$ does not realize by the end of Phase 2.

Hence the firm's one-shot deviation constraint in the k th period on the eequilibrium path for $k = 1, ..., N + 1$ is

$$(1 - \delta) G \leq \delta^{2n-k+2} \Pi_{t=1}^{k-1} \lambda_N(t) \lambda_N^{N-k+2} \gamma_N \pi (q - p) p^N \left(V^T(\delta) - V^3(\delta)\right) \text{ for } k = 1, ..., N + 1,$$

- $\lambda_N(t) = \frac{\lambda_N}{\sum_{n=t-1}^{N} \lambda_n}$ is the probability that a signal realizes with the longest lag conditional on it has not realized for $t - 1$ periods
- π is the probability to go to Phase 3 in the end of Phase 2 conditional on the above event (given which s_k plays a pivotal role) and s_k being a bad signal b
- $V^T(\delta)$ is the firm's continuation payoff at the transition period and V^3 is the firm's continuation payoff in the beginning of Phase 3.

For the one-shot deviation constraints in the k th period of Phase 1 and the transition period <u>off</u> the equilibrium path, some of p in this expression is replaced by $q > p$. So those one-shot deviation constraints off the eequilibrium path would be automatically satisfied if this one-shot deviation constraint on the eequilibrium path is satisfied.

Next note that the above condition for k would be automatically satisfied if it is satisfied for $k = 1$, because $\delta^{2n-k+2} \leq \delta^{2n+1}$ and $\lambda_N(t) \geq \lambda_N$ for any $t = 1, ..., N$. So every one-shot deviation constraint on and off the eequilibrium path in Phase 1 and the transition period would be satisfied if the following one-shot deviation constraint in the 1st period of Phase 1 is satisfied.

$$(1-\delta) G \leq \delta^{2n+1} \lambda_N^{N-k+2} \gamma_N \pi (q-p) p^N \left(V^T(\delta) - V^3(\delta)\right), \tag{1}$$

Incentive in Phase 2.

We define the probability $\beta(z)$ to go back to Phase 1 in the end Phase 3 as follows.

$$\beta(z) = \eta \sum_{n=1}^{N} \delta^{n-1} \mathbf{1}(z_n = g),$$

when we use the public signal associated the firm's action in the kth period of Phase 2 as an input for z_k. We choose small η so that this is indeed a probability given any $z \in \{g, b\}^N$

Suppose that the firm believes with probability 1 that the game would reach the end of Phase 2 and move to Phase 3, then the firm's one-shot deviation constraint in the k th period of Phase 2 (on and off the eequilibrium path) would be independent of k (by construction of β) and given by

$$(1-\delta) G \leq \delta^{2N} \eta (q-p) V^1(\delta). \tag{2}$$

where V^1 is the continuation payoff in the beginning of Phase 1. We choose η so that this constraint is binding.

Note that if the firm believes with probability 1 that the game would move to the transition period in the end of Phase 2 when it is reached, then the one-shot deviation constraint in the k th period of Phase 2 coincides with the one-shot deviation constraint in the k th period of Phase 2.

Since the actual one-shot deviation constraints in Phase 2 are mixtures of the above two type of one-shot deviation constraints conditional on two distinct events, they are satisfied.

To summarize, all the on-shot deviation constraints on and off the eequilibrium path are satisfied if we can choose π and η in such a way that (1) and (2) are satisfied.

Next we derive the firms' payoff V^1, assuming that (1) and (2) are binding. V^1 satisfy the following recursive equation.

$$V^1 = \sum_{n=0}^{N-1} \gamma^n \left[(1-\delta^{n+1}) + \delta^{n+1} V^1\right] + \gamma^N \left[(1-\delta^{N+1}) + \delta^{N+1}\left\{(1-p)V^1 + pV^2\right\}\right]$$

Note that the binding (1) is equivalent to the following condition where the continuation payoffs are evaluated just after the transition period.

$$(1-\delta)\, G = \gamma^N \delta^{N+1} (q-p) \left(V^1 - V^2\right).$$

Using the above two equations, we can eliminate V^2 to derive V^1 as follows

$$\begin{aligned} V^1 &= 1 - \frac{1}{1 - \sum_{n=0}^{N} \gamma_n \delta^{n+1}} \frac{1-\delta}{\frac{q}{p} - 1} \\ &= 1 - \frac{1}{1 + E(\lambda, \delta)} \frac{1-\delta}{\frac{q}{p} - 1} \end{aligned}$$

where $E(\lambda, \delta) = \sum_{n=1}^{N} \gamma_n \left(\sum_{k=1}^{n} \delta^k\right)$. So V^1 achieves the bound in Theorem 1.

If this bound is strictly positive as $\delta \to 1$, it can be verified that we can find $\pi \in (0,1)$ and small enough η such that $\beta(z) \in (0,1)$ so that the above strategy is well defined and (1) and (2) are satisfied with equality when δ is large enough. So there exists some lower bound of discount factor $\underline{\delta}$ in this case such that the constructed strategy is an eequilibrium and achieves the bound as the equilibrium payoff for any $\delta \geq \underline{\delta}$.

6.3 Proof of Theorem 2

TBA