Three Essays on Learning and Dynamic Coordination Games

Dengwei Qi

Graduate School of Economics

Kyoto University

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

December 2022

Acknowledgements

I am very grateful to Atsushi Kajii, Tadashi Sekiguchi, Saori Chiba, and Chia-Hui Chen for their support, guidance, and patience during my graduate studies at Kyoto University. Particularly, I am deeply indebted to Atsushi Kajii for taking the time to discuss and provide feedback on my research regularly, even when the projects were immature and the research ideas were rough. All four professors are my advisors, and they were willing to take me in when I had little experience in economic theory. Without their trust and support, I could never initiate the doctoral study, let alone graduation.

Also, I am thankful to Graduate School of Economics at Kyoto University, especially the East Asia Programme, for recommending me on the MEXT scholarship. And again, my thanks goes to Atsushi Kajii, who generously hired me as a research assistant, and this supports me for the last two years of my graduate studies.

Last but not least, I would like to thank my family: my parents and my brother, who always support me and encourage me to pursue what I like.

Abstract

This dissertation consists of three chapters about dynamic coordination games, learning, and their connections.

The first chapter presents a dynamic coordination game in the context of an investment crash. Agents decide whether and when to invest in a risky project while observing past activities over time. The optimal action timing of an agent is determined by her constant trading off the informational gain of delaying versus its opportunity cost. The chapter characterizes the optimal action timing of each agent and further shows the uniqueness of such a monotone equilibrium, shedding testable insights into analysis and policy guidance. Various comparative statics questions are answered, including the impact of learning on coordination success and behaviors. Furthermore, the analysis applies to all ranges of information precisions, resolving a long difficulty in the literature in which most existing studies can only tackle the vanishing noise situations. Additionally, I show that full learning about the state achieves in the limit, and give conditions on which observing actions reveals more accurate information about the state than directly observing it.

The second chapter is based on a classical market-based learning model in the presence of both private and public observations of market aggregates. Existing studies show that the learning speed is slow with only public learning, which undermines the value of information since market situations change. This chapter incorporates a private learning channel through observing market data with idiosyncratic noise. I demonstrate that now learning efficiency is

improved, in the sense that both public and private information become limit accurate and furthermore, the asymptotic learning rates are linear, higher than in the pure public learning case. Various intuitive features of the learning process have also been verified.

The third chapter combines the above two chapters by constructing a model in which agents interact and learn from each other prior to a coordination game. Learning still happens from both public and private observation of market aggregates. I show that learning, or equivalently the higher information precision, improves agents' expected payoffs and coordination success. In addition, I demonstrate that the incorporation of private learning, instead of dispersing agents' information, contributes to its conformity, and thus prompts multiple equilibria in the global game.

Table of contents

Chapter 1

Learning and Strategic Delay in a Dynamic Coordination Game

1.1. Introduction

Coordination games of incomplete information like currency crises or investment crashes impact the economy massively and draw much attention from economists. One prominent approach to analyze such problems is the global games model pioneered by Carlsson and Van Damme (1993) and Morris and Shin (1998). It introduces asymmetric information into the traditional coordination game framework and remarkably obtains a unique and analytically convenient equilibrium, shedding testable insights on policy guidance and welfare implications. However, the existing studies are mostly in static contexts, despite the economic activities are inherently dynamic. Budget-constrained agents delay their investment decisions to learn from their predecessors' behaviors, for example. That said, learning and delaying behaviors of agents are prevalent in practice and worth exploring, but the static models cannot provide predictions or analysis for those dynamic aspects.

Particularly in coordination games, learning and delay behaviors become more notable because agents face not only the payoff uncertainty about the economic fundamentals, but also the strategic uncertainty about their opponents' (past, current, and future) beliefs. Consequently, it is almost inevitable to extend the static models into multi-periods and consider learning and delaying behaviors of agents, to capture their intrinsic motivations to mitigate both sorts of uncertainties. And the investigation into dynamic environments is not a simple extension of the static model because of the strategic delay consideration of agents to try to select the optimal action timing. That is, delay provides informational gains through agents' observation of past activities, but is also costly due to discounting and shrinking opportunities, so agents must constantly trade off the benefit and the cost of delay to determine when to act, and this crucial trade-off cannot be captured in static frameworks.

Therefore in this paper, we construct an $N \in \mathbb{N}$ period model in an investment context to investigate the impact of learning and delay options on agents' behaviors, based on the static global game of Morris and Shin (2000). The prospect of an investment project, or the state, is deterministic but ex ante unknown, and a continuum of heterogeneously informed agents can undertake a fixed-size investment once. They independently select the investment timing (if at all) out of N periods, while observing a stream of noisy signals about past activities over time, which represents the learning behavior and is the informational gain of delay. To capture the coordination motive and the opportunity cost of delay, we let the payoff of the investment to an agent, paid at the end of the game, be positively correlated with the aggregate investment, while negatively with her investment timing, if ever invested. Hence agents with one-time investment opportunity need to trade off the informational gain of delay versus its opportunity cost to decide their optimal investment timing.

After constructing the model, we solve for its equilibrium and demonstrate the existence and the uniqueness of a monotone equilibrium, in which agents take a symmetric threshold strategy profile (i.e., an agent invests in one period if and only if her belief about the state exceeds some threshold prescribed for that period.) This monotone form of strategy is documented in almost all relevant literature and is as well intuitively appealing in this dynamic environment. To see it, agents select their investment timing by trading off the informational gain of delay against its opportunity cost, so if an agent believes the state is good enough in one period, she expects the investment is profitable and consequently, the expected opportunity cost to her is huge while the informational gain is little; thereby she invests immediately. Otherwise she delays to the next period, in which she updates her information by observing what others have done, and then make decisions by the same trade-off logic as before, and so on.

Noteworthy in equilibrium, the investment decisions of an agent only depend on her beliefs about the state, even though the payoff involves her fellow agents' behaviors. This is so because (i) the payoff depends on the aggregate investment and (ii) the aggregate investment is (shown to be) deterministic given the state. There two properties are standard in global games literature and essentially stem from the Law of Large Numbers. Recall that agents form a continuum, which allows us to characterize a one-to-one relation between the aggregate action and the state. Hence a belief about the state suffices to evaluate the corresponding aggregate investment. Also with this deterministic relation, the information learned from past actions is shown to be summarized in a closed-form statistic centered around the state, for all learning precision levels. This is one of the novelty of this paper because the past literature in dynamic environments only allows analysis in limit accurate learning situations (cf., Dasgupta (2007)). The comparison to the existing literature will be elaborated in the literature review section soon.

Our following analysis is thus focused on this unique monotone equilibrium and addresses two questions. The first probes the dynamics of agents' behaviors in equilibrium, and the second investigates comparative statics, particularly the impact of learning and delay options on coordination success and welfare.

First, we summarize agents' equilibrium behaviors. In period 1, the optimistic agents who observe favorable signals (that exceed the equilibrium threshold of period 1) invest immediately, since they believe the investment's prospect is already good and thus outweigh the opportunity cost of delay over its informational gain. In the subsequent intermediate periods, the remaining agents constantly revise their expectations about the investment through cumulative learning, and depending on learning efficiency, a large or small fraction of agents will switch into investing. Noteworthy, if the learning efficiency is modest (i.e., the accuracy of endogenous signals is low), in every period will a few agents newly invest, so the relatively inertia phenomenon documented in the literature (Angeletos et al. (2007)) is expected. Also by implication, had no learning effect existed, agents would only act in the first period and stay inactive till the last period. We indeed verify this conjecture and show that the mere delay option without learning opportunities has no impact on the game, relative to the one-shot game. In the last period, there is no stage to delay to, so another positive fraction of remaining agents will choose to invest.

Next we discuss comparative statics. To begin with, we contrast agents' behaviors with that in the one-shot game. Results show that agents are less aggressive (i.e., less likely to invest) in the intermediate periods than in the static game. Intuitively, agents are tempted by the information learned from delaying and hence choose to wait. On the other hand, agents behave more aggressively in the last period of the dynamic game, due to a higher expected total investment and the coordination motive.

We next investigate the values of learning opportunities and the consequent welfare implications. It is demonstrated that learning opportunities increase agents' expected continuation payoffs and thus improve coordination success and social welfare. Intuitively, coordination fails because agents, facing the uncertainty about whether others will cooperate, may choose not to invest, even if it is their collective interest to do so. Learning alleviates this problem by reducing the strategic uncertainty among agents, since it makes agents' signals more accurate and thereby better aligned. Also we show that agents more accurately infer the state in the presence of learning, indicating the payoff uncertainty is also mitigated.

Note that agents learn the state through observing past activities, and we are interested in how efficient such a learning mechanism is, relative to learning from directly observing the state. We find that as long as agents' initial information is precise enough, observing actions reveals more accurate information than directly observing the state. Intuitively, learning efficiency of observing actions depends on (i) how accurately agents' private information is about the state and (ii) how accurately endogenous signals reflect their actions (and hence their private information). The two channels are shown to be mutually reinforced and therefore, if one of them is accurate enough, it is possible that indirect learning delivers more accurate information about the state than direct learning.

Lastly, we discuss the equilibrium selection in the dynamic model. One of the remarkable results that static global game models provide is the uniqueness in equilibrium when information among agents is sufficiently diffused (see Morris and Shin (2003)), resolving the indeterminacy of equilibria problem in complete information coordination games. And we indeed obtain a unique monotone equilibrium in this dynamic environment. However, other forms of strategies than a threshold strategy cannot be excluded to constitute an equilibrium. This is because the dynamic environment provides other dimensions for coordination and thus multiplicity. For instance, if all agents believe their opponents will take some specific strategy form, so may they, and this mutual effect in turn justifies the usage of that strategy form. Aside from this, as Angeletos and Werning (2006) demonstrate, when learning is through public observation of actions, multiple equilibria can arise even when agents are endowed with limit accurate private information. The feature is also present in our model when we consider that learning is through public observation of actions in Section 4.

Also in section 4, we extend the game to infinite periods and show that the properties of the N-period game are still valid; furthermore, we find that agents completely learn the true state in the limit, avoiding the usual information cascade when learning is through observing past activities (Bikhchandani et al. (1992) and Banerjee (1992)). Indeed, in our model, pooled information of agents reveals the true state, so it is at least plausible for agents to fully learn the state. And the signal structures we consider are continuous due to normal noise; as Lee (1993) demonstrate, this continuity prevents information cascade, because any tiny variation in agents' behaviors will be, at least noisily, reflected by signals.

1.1.1 Related literature

This paper is most related to Angeletos et al. (2007) and Dasgupta (2007). Angeletos et al. (2007) investigate a dynamic regime change game in which short-lived agents (in the sense that agents are new and given a unit of perishable endowment every period) repeatedly decide whether to attack a regime, while observing the outcomes of the past attacks. By contrast, agents in our model are long-lived and have budget constraints in the sense that they can only act at most once, and thus face an active timing problem. Moreover, we consider that all past activities cumulatively affect the payoff of the investment, while they assume only the action of the present period affects agents' payoffs. In addition, a continuous payoff structure is assumed in our paper, as opposed to the discrete payoff structures of the regime change game (which pays either a lump sum or nothing, depending on whether the regime switches), so our result about the dynamics of agents' behaviors complements that of Angeletos et al. (2007): agents in our model respond continuously to information variations, while their agents have complete inertia unless receiving a large change of information.

It is worth stressing that though payoffs are continuous in parameters in our model, agents' strategies are not because of the feature of the threshold strategy. That said, agents' actions can change discontinuously and dramatically with a small perturbation of information (even given the state of the world); to see it, consider those with signals around the threshold. Consequently, volatile non-fundamental variations of actions exist in our model, which is one of the highlights of the global games approach to explain sudden changes of behaviors in crises phenomena; see Morris and Shin (2003).

Dasgupta (2007) considers a regime change game in a two-period span, with agents endowed with limit accurate private information as well as learning is of limit accurate, so learning is almost immaterial there. Our analysis instead spans N, and further infinite, periods and applies to all learning efficiencies. Furthermore, the almost fully informed agents in Dasgupta's work always benefit from the delay option, while we find that, when agents are not fully informed, what helps improve coordination success is the learning effects and that the delay option alone does not affect the outcomes, relative to the static game.

Some works focus exclusively on learning effects, especially the effects of public signals on equilibrium selection in global games. The pioneers are Angeletos and Werning (2006), who show the rise of multiple equilibria when learning is through public observation. Most distinctively, our paper differs from theirs because in that their game is essentially static, in the sense that one group of agents act in the first period in the financial market of Grossman and Stiglitz (1976), and then another group, observing price or activity in the market, act in a static global game; the two groups share no payoff transfers. Also connecting to the rational learning literature, our learning mechanism, particularly the Gaussian signal structure, has the similar updating rule as in Vives (1993).

There are works on global coordination games that study different aspects than this paper. For example, Hellwig et al. (2006) consider endogenous interest rates, Angeletos et al.

(2006) analyze the signaling effects, and Szkup and Trevino (2015) study costly information acquisition. See also Morris and Shin (1998) for currency crises, Goldstein and Pauzner (2005) for bank runs, and Edmond (2013) for sociopolitical revolutions.

In very different setups, the option value of delay has been examined by Chamley and Gale (1994) in a noncooperation environment with perfect observation of past activities. Gale (1995) studies strategic delay in a complete information coordination game.

The rest of the paper is structured as follows. Section 2 investigates the game comprising two periods and captures our core results. Section 3 considers multiple periods and confirms the validity of the results in the two-period model. Section 4 discusses the extension concerning infinite periods and public learning.

1.2. The Two-Period Model

In this section, we examine a two-period game with a linear payoff structure; it captures our core results. The stage game is based on Morris and Shin (2000). The general model that comprises $N \in \mathbb{N}$ periods and a general payoff will be explored in Section 3.

1.2.1 Setup

A measure-one continuum of agents, denoted i or j, independently decide whether to invest in a risky project at time $t = 1$ or $t = 2$, if at all. An agent can invest at most once irreversibly. Let $a_{ti} \in \{0,1\}$ denote agent i's action at time t, where 1 (or 0) refers to investing (or not investing); it is then required that $a_{1i} + a_{2i} \in \{0,1\}$. Moreover, let $a_t = \int_i a_{ti} di$ be the aggregate investment at time t, and $\hat{a}_t = \sum_1^t a_k$ the cumulative investment till t. The return of the project is determined after all investment decisions are completed, so payoffs are realized at the end of time 2. The payoff of an agent who does not invest is normalized to 0, and that to investing is the sum of two factors. The first is the total investment $\hat{a}_2$, and the second is the exogenous investment environment, which is driven by other economic fundamentals. We summarize the second factor by a single parameter $r \in \mathbb{R}$. In sum, the return to an agent who chooses $a_{ti} \in \{0,1\}$ equals

$$a_{ti}(r+\hat{a}_2). \tag{1.1}$$

In each $t = 1,2$, agent i chooses $a_{ti} \in \{0,1\}$ to maximize her aggregate expected payoff, $E[\sum_{t=1}^{2} \delta^{t-1} a_{ti}(r+\hat{a}_2)]$, given her available information at that time, where $\delta \in (0,1)$ is the timing cost on investment. Note that δ acts similarly as a discount factor, but since agents only receive payments at the end of the game, δ is interpreted as shrinking opportunities.

The state parameter r is deterministic but ex ante unknown, and is uniformly distributed over the entire real line, so agents hold an improper prior about it: $r \sim \text{Unif}(\mathbb{R})$. In period 1, agent i observes a private signal x_{1i} about the realization of r:

$$x_{1i} = r + \frac{1}{\sqrt{\tau_1}}\varepsilon_{1i}, \tag{1.2}$$

and in period 2, agent i additionally receives x_{2i} about the past activity a_1:

$$x_{2i} = \Phi^{-1}(a_1) + \frac{1}{\sqrt{\tau_2}}\varepsilon_{2i}, \tag{1.3}$$

where Φ is the CDF of the standard normal and $\tau_t > 0$, $t = 1,2$, measures the information quality, and ε_{ti} is a standard normal variable, independent across time and agents and of r (i.e., $\varepsilon_{ti}|_r = \varepsilon_{ti} \sim \mathcal{N}(0,1)$, i.i.d. for any t and i). Here we follow the literature (Dasgupta (2007) and Angeletos and Werning (2006)) to choose the analytically convenient information aggregation technology Φ^{-1}, but as we will see soon, the qualitative results of the paper are valid for other learning technologies. Furthermore, we impose the Law of Large Numbers

(LLN) convention through out the paper, namely, the proportion of agents who receive signals higher than some real number is equal to the probability of an individual agent receiving such signals. Consequently, no aggregate uncertainty about r exists since the idiosyncratic noise cancels out: $\int_i \varepsilon_{1i} di = 0$.

In summary, the game proceeds as follows.

0. Nature randomly draws r from $\mathbb{R}$.

1. In period 1, agent i privately observes x_{1i} about r and then makes an investment decision. The total investment a_1 is thus determined.

2. Subsequently in period 2, agent i privately observes x_{2i} about a_1 and then takes a feasible action. The aggregate investment $\hat{a}_2 (= a_1 + a_2)$ of the game is thus determined.

3. The payoffs to investment depending on r and $\hat{a}_2$ are realized at the end of period 2.

Recall that in period 2, the only feasible action to agents who have invested is action 0.

1.2.2 Threshold Strategies and Monotone Equilibria

In line with the literature, we consider that agents play a symmetric threshold strategy in each period - an agent invests iff her expectation of r at that period exceeds some threshold number. Specifically, a *threshold strategy* σ_1 *in period 1* for agent i who observes x_{1i} takes the form

$$\sigma_1(x_{1i}) = \begin{cases} 1, & \text{if } x_{1i} > x_1 \\ 0, & \text{otherwise,} \end{cases}$$

for some $x_1 \in \mathbb{R}$ (we differentiate signals and thresholds by subscript i). By construction, the agent selects not investing at a tie when $x_{1i} = x_1$. The expression of a threshold strategy in period 2 requires closer inspection because of endogenous learning. To see it, note that $a_1(r) = P(x_{1i} > x_1 \mid r) = \Phi(\sqrt{\tau_1}(r - x_1))$ for any realization of r, when all agents follow a threshold strategy with threshold x_1 in period 1. Hence endogenous signal x_{2i} becomes

$$x_{2i} = \sqrt{\tau_1}(r - x_1) + \frac{1}{\sqrt{\tau_2}}\varepsilon_{2i},$$

rearranging which we obtain

$$\frac{x_{2i}}{\sqrt{\tau_1}} + x_1 = r + \frac{1}{\sqrt{\tau_2 \tau_1}}\varepsilon_{2i}.$$

Therefore, if we define

$$x'_{2i} \equiv \frac{x_{2i}}{\sqrt{\tau_1}} + x_1,$$

then

$$x'_{2i} = r + \frac{1}{\sqrt{\tau'_2}}\varepsilon_{2i},$$

where $\tau'_2 \equiv \tau_1 \tau_2$. Note that x'_{2i} is informationally equivalent to x_{2i} with respect to r. Therefore by Bayes' rule, agent i's updated belief about r in period 2 can be summarized by a sufficient statistic $\hat{x}_{2i}(x_{1i}, x_{2i})$ with

$$\hat{x}_{2i}(x_{1i}, x_{2i}) = \hat{x}_{2i}(x_{1i}, x'_{2i}) = \frac{\tau_1 x_{1i} + \tau'_2 x'_{2i}}{\tau_1 + \tau'_2} = r + \frac{1}{\sqrt{\hat{\tau}_2}}\varepsilon_{2i}, \tag{1.4}$$

where $\hat{\tau}_2 \equiv \tau_1 + \tau'_2$.[1] A *threshold strategy* σ_2 *in period 2* is defined by the rule

$$\sigma_2(x_{1i}, x_{2i}) = \begin{cases} 1 - \sigma_1(x_{1i}), & \text{if } \hat{x}_{2i}(x_{1i}, x_{2i}) > x_2 \\ 0, & \text{otherwise,} \end{cases}$$

[1] Recall ε_{1i} and ε_{2i} are both standard normals, and by abusing notation, we let ε_{2i} in (1.4) denote a normal noise in agent i's belief towards r at $t = 2$.

for some $x_2 \in \mathbb{R}$, given agents follow σ_1 in period 1. This completes the definition of a threshold strategy profile in the dynamic game. [2] When no confusion might occur, we write $\hat{x}_{2i} \equiv \hat{x}_{2i}(x_{1i}, x_{2i})$ and $\hat{x}_{1i} \equiv x_{1i}$ (and hence define $\hat{\tau}_1 = \tau_1' \equiv \tau_1$), and denote a threshold strategy profile by its thresholds, say, (x_1, x_2).

It should be noted that when all the agents follow (x_1, x_2), the size of investment a_t at time $t = 1, 2$ is a deterministic function of r, such that

$$a_1(r) = \int_i P(\hat{x}_{1i} > x_1 \mid r) di = P(\hat{x}_{1i} > x_t \mid r), \quad a_2(r) = P(x_{1i} < x_1, \hat{x}_{2i} > x_2 | r),$$

by LLN.[3] That said, when all agents play a threshold strategy, the payoff to investment depends on $(r, \hat{a}_2(r))$, so that the estimation about r suffices to evaluate decisions even $\hat{a}_2$ enters the payoff function. In this paper, we consider symmetric perfect Bayesian equilibria in which all agents follow a threshold strategy profile, and call such equilibria monotone equilibria. In what follows, we refer to monotone equilibria as equilibria unless otherwise stated.

It is worth stressing that the key in obtaining an analytical form of posterior belief $\hat{x}_{2i}(x_{1i}, x_{2i})$ in period 2 is the transformation from x_{2i} centered around $\Phi^{-1}(a_1)$ to x_{2i}' centered around r. The transformation is plausible because no aggregate uncertainty exists in the model ($\int_i x_{1i} di = r$ indicates pooling the continuum's information reveals r), so that when all agents follow a threshold strategy in period 1, the aggregate activity a_1 is deterministic given r, and thus the observation of a monotone function of it (i.e., $\Phi^{-1}(a_1)$) leads to an estimation of r. This line of reasoning implies that the specific aggregation rule Φ^{-1} of x_{2i}

[2] Note that σ_2 is only well defined when agents take a threshold strategy in period 1; it suffices for our purpose since we restrict to agents playing a threshold strategy profile.

[3] Note that $a_2(r) = P(x_{1i} < x_1, \hat{x}_{2i} > x_2 | r) = P(\hat{x}_{2i} > x_2 | r, x_{1i} < x_1) P(x_{1i} < x_1 | r) = P(\hat{x}_{2i} > x_2 | r) P(x_{1i} < x_1 | r)$, since r suffices to estimate $\hat{x}_{2i}$ by (1.4).

is not qualitatively restrictive: any one-to-one aggregation rule results in an estimation of r from observing a_1; we choose Φ^{-1} to obtain the well-behaved transformed signal.

Moreover, since the estimation of r is derived from x_{2i}, the quality of the estimation depends on how precise (i.e., τ_2) x_{2i} reflects a_1 and how precise (i.e., τ_1) a_1 reflects r. Indeed, the induced precision level $\tau_2' = \tau_1 \tau_2$ of x_{2i}' verifies this. It also highlights that the endogenous information is generated by social learning, or from individuals' private information, so the more accurate information agents initially hold, the more accurate information their actions convey. Noteworthy, the precision level τ_2' is the same as that of endogenous signals obtained from rational expectations equilibrium price (Grossman and Stiglitz (1976)), underscoring that the specific learning rule Φ^{-1} provides results consistent with the literature.

We close this section by characterizing agent i's cross-period beliefs about one another, and show that a higher expectation of r leads to a higher expectation of $\hat{a}_2$. The results are useful in the equilibrium characterization later.

Lemma 1. *When agents follow a threshold strategy profile (x_1, x_2), for time $t \neq k \in \{1,2\}$ and any signal realization $\hat{x}_{ki}, \hat{x}_{tj}$, we have, for $i \neq j$,*

$$\hat{x}_{tj}|_{\hat{x}_{ki}} \sim \mathcal{N}\left(\hat{x}_{ki}, \frac{\hat{\tau}_k + \hat{\tau}_t}{\hat{\tau}_k \hat{\tau}_t}\right),$$

and for $i = j$

$$\hat{x}_{2i}|_{x_{1i}} \sim \mathcal{N}\left(x_{1i}, \frac{\tau_2'}{\hat{\tau}_2 \tau_1}\right). \tag{1.5}$$

Moreover, $E[\hat{a}_2|\hat{x}_{ki}]$ strictly increases in $\hat{x}_{ki}$.

Proof. For $i \neq j$, since $\hat{x}_{ki}(= r + \varepsilon_{ki}/\sqrt{\hat{\tau}_k})$, we have

$$\hat{x}_{tj} = r + \frac{1}{\sqrt{\hat{\tau}_t}}\varepsilon_{tj} = \hat{x}_{ki} - \frac{1}{\sqrt{\hat{\tau}_k}}\varepsilon_{ki} + \frac{1}{\sqrt{\hat{\tau}_t}}\varepsilon_{tj}.$$

For $i=j$, since $x'_{2i}=x_{1i}-\varepsilon_{1i}/\sqrt{\tau_1}+\varepsilon_{2i}/\sqrt{\tau'_2}$, we have

$$\hat{x}_{2i}=\frac{\tau_1 x_{1i}+\tau'_2 x'_{2i}}{\tau_1+\tau'_2}=x_{1i}+\frac{\tau'_2}{\tau_1+\tau'_2}(-\frac{1}{\sqrt{\tau_1}}\varepsilon_{1i}+\frac{1}{\sqrt{\tau'_2}}\varepsilon_{2i}),$$

so (1.5) holds.

For the second part, note that $(E[a_1|\hat{x}_{ki}])'=(P(x_{1j}>x_1|\hat{x}_{ki}))'>0$, and that $E[a_2|\hat{x}_{ki}]=P(x_{1j}<x_1,\hat{x}_{2j}>x_2|\hat{x}_{ki})=(1-E[a_1|\hat{x}_{ki}])P(\hat{x}_{2j}>x_2|\hat{x}_{ki})$, so

$$\begin{aligned}\frac{d}{d\hat{x}_{ki}}E[\hat{a}_2|\hat{x}_{ki}]&=\frac{d}{d\hat{x}_{ki}}E[(a_1+a_2)\mid \hat{x}_{ki}]\\&=(E[a_1\mid\hat{x}_{ki}])'(1-P(\hat{x}_{2j}>x_2|\hat{x}_{ki}))+(1-E[a_1\mid\hat{x}_{ki}])(\underbrace{P(\hat{x}_{2j}>x_2|\hat{x}_{ki})}_{=\Phi(\sqrt{\cdot}(\hat{x}_{ki}-x_2))})'>0.\end{aligned}$$

Q.E.D.

Note that, complying with our intuition, the point (iii) states that i makes more accurate inferences about her own belief than about others'. In what follows, we abbreviate $E[a_t|\hat{x}_{ki}]$ and $E[\hat{a}_t|\hat{x}_{ki}]$ to $a_t(\hat{x}_{ki})$ and $\hat{a}_t(\hat{x}_{ki})$, respectively.

1.2.3 Equilibrium Characterization

The Static Game

We first consider when the game only consists of the first period, namely a static game, to illustrate how to solve for the unique monotone equilibrium; it also serves as a benchmark for later comparative statics analysis. Since the one-shot game is a standard static global game, the equilibrium can be easily characterized as in the following proposition.

Proposition 1. *In the static game with signal structure $x_{1i} = r + \varepsilon_{1i}/\sqrt{\tau_1}$ and payoff structure $r + a_1$, there exists a unique equilibrium which is a monotone equilibrium characterized by a threshold strategy x^*_{st} with $x^*_{st} = -1/2$.*

A detailed proof can be found in Morris and Shin (2000), and here we sketch it. Think of a marginal agent with signal x_{1i} such that she is indifferent between investing or not, namely, $E[r + a_1|x_{1i}] = 0$. It is straightforward to verify it has a unique solution, $x_{1i} = -1/2$, and we claim it is x^*_{st}.[4] Indeed, following symmetric strategy $-1/2$ (investing iff $x_{1i} > -1/2$) is optimal, because $E[r + a_1|x_{1i}] > 0$ iff $x_{1i} > -1/2$. The global uniqueness is obtained by the standard iterated dominance argument so we omit its proof.

Note that in verifying the threshold strategy x^*_{st} constitutes an equilibrium, it suffices to consider the marginal agent who is indifferent between the two actions. This is due to payoff's monotonicity in r and a_1: a higher signal realization indicates higher r and a_1, resulting in higher expected payoffs from investing, so an agent with signals higher than the cutoff signal expects the payoff to investing exceeds 0 and thus invests. Since the monotonicity of the payoff holds in the dynamic game, one can expect that an equilibrium shall be readily found by identifying such a marginal agent. However, her role is more subtle there, since instead of balancing whether to invest or not, the agent trades off between acting now versus delaying and acting optimally later. A careful analysis is therefore required and conducted below.

The Two-Period Game

We now solve the two-period model. Let $R_1(x_{1i};(x_1,x_2))$ denote the expected continuation payoff for agent i who observes x_{1i} and delays in period 1, given all other agents follow some threshold strategy (x_1,x_2) in the game. To compute R_1, agent i infers her to be received

[4]Formally, $a_1(x_{1i}) = \Phi(\sqrt{\tau_1/2}(x_{1i} - x_1))$ given other agents follow some threshold strategy x_1. Thus the only symmetric solution to the equation is $x_1 = -1/2$.

signal $\hat{x}_{2i}$ from x_{1i}, since $\hat{x}_{2i}$ determines whether she will invest later and if so, her expected payoff. We claim

$$\begin{aligned} R_1(x_{1i};(x_1,x_2)) &= \delta E\left[E\left[r+\hat{a}_2|\hat{x}_{2i}>x_2\right]\big|x_{1i}\right] \\ &= \delta \int_{x_2}^{\infty} E[r+\hat{a}_2 \mid \hat{x}_{2i}] f(\hat{x}_{2i} \mid x_{1i}) d\hat{x}_{2i}, \end{aligned}$$

where $f(\hat{x}_{2i}|x_{1i})$ is the density of $\hat{x}_{2i}$ given x_{1i} and by (1.5) in Lemma 1, equals $(P(\cdot \leqslant \hat{x}_{2i}|x_{1i}))' = (\Phi(\sqrt{\tau_1 \hat{\tau}_2/\tau_2'}(\hat{x}_{2i}-x_{1i})))'$; also see Lemma 1 for the formula of $E[\hat{a}_2|\hat{x}_{2i}]$ given thresholds (x_1,x_2). Let $R_2(\hat{x}_{2i};(x_1,x_2)) \equiv 0$ for any $\hat{x}_{2i}$ and (x_1,x_2), meaning the continuation payoff at the last stage is 0.

To better understand the formula of R_1, suppose the agent who observes $\hat{x}_{2i}$ has reached period 2; then given others follow (x_1,x_2) in the game, her expected payoff to following x_2, denoted $\tilde{R}_1(\hat{x}_{2i};(x_1,x_2))$, is

$$\tilde{R}_1(\hat{x}_{2i};(x_1,x_2)) = \begin{cases} E[r+\hat{a}_2 \mid \hat{x}_{2i}], & \text{if } \hat{x}_{2i} > x_2 \\ 0, & \text{otherwise.} \end{cases}$$

In period 1, the agent forms an expectation of this value through her current signal x_{1i}, which is what she expects to obtain by delaying and thus is R_1, so

$$R_1(x_{1i};(x_1,x_2)) = \delta E[\tilde{R}_1(\hat{x}_{2i};(x_1,x_2)) \mid x_{1i}] = \delta \int_{x_2}^{\infty} E[r+\hat{a}_2 \mid \hat{x}_{2i}] f(\hat{x}_{2i} \mid x_{1i}) d\hat{x}_{2i}.$$

With this result, the following proposition establishes the existence and the uniqueness of a monotone equilibrium.

Proposition 2. *A unique monotone equilibrium characterized by a threshold strategy profile* (x_1^*, x_2^*) *exists in the two-period game, where* x_t^* *uniquely solves*

$$E[r + \hat{a}_2 | x_t^*] = R_t(x_t^*; (x_1^*, x_2^*)), \quad \textit{for } t = 1, 2.$$

Proof. Let (x_1, x_2) denote an arbitrary threshold strategy. In period 2, given the others follow (x_1, x_2) in the game, agent i with belief $\hat{x}_{2i}$ expects her investment payoff to be

$$G_2(\hat{x}_{2i}; (x_1, x_2)) \equiv E[r + \hat{a}_2 \mid \hat{x}_{2i}; (x_1, x_2)], \tag{1.6}$$

where we write (x_1, x_2) to emphasize it is used to compute $E[\hat{a}_2 | \hat{x}_{2i}]$, which by Lemma 1 increases in $\hat{x}_{2i}$, so (1.6) increases in $\hat{x}_{2i}$. If a threshold strategy $x_2' \in \mathbb{R}$ is an equilibrium strategy in period 2, agent i should invest (i.e., $G_2(\hat{x}_{2i}; (x_1, x_2')) > 0$) if $\hat{x}_{2i} > x_2'$ and should not if $\hat{x}_{2i} < x_2'$; therefore by the continuity of (1.6) in $\hat{x}_{2i}$, i observing $\hat{x}_{2i} = x_2'$ must be indifferent between investing or not, namely,

$$G_2(x_2'; (x_1, x_2')) = 0. \tag{1.7}$$

We show in Appendix that $G_2(x_2'; (x_1, x_2'))$ is continuous, strictly increasing in x_2' and converges to $-\infty$ (*resp.* ∞) as $x_2' \to -\infty$ (*resp.* ∞). Hence a unique solution, given any x_1, to (1.7) exists, and we call it $x_2^*(x_1)$. Note that $x_2^*(x_1)$ is the only candidate for an equilibrium threshold in period 2, given x_1. The increasing monotonicity of (1.6) thus verifies $x_2^*(x_1)$ constitutes an equilibrium in period 2, since $G_2(\hat{x}_{2i}; (x_1, x_2^*(x_1))) > 0$ if $\hat{x}_{2i} > x_2^*(x_1)$.

With the above result, we proceed to period 1. When all agents except i follow some threshold strategy x_1 in period 1 and all agents follow $x_2^*(x_1)$ in period 2 (note that the deviation of a measure-zero agent does not affect the optimal strategy in period 2), the investment payoff of agent i with x_{1i} equals

$$G_1(x_{1i};(x_1,x_2^*(x_1))) \equiv E[r+\hat{a}_2 \mid x_{1i};(x_1,x_2^*(x_1))],$$

which increases in x_{1i} by Lemma 1. Let $\Delta(x_{1i};(x_1,x_2^*(x_1)))$ denote the payoff difference of agent i between investing and delaying in period 1, namely,

$$\Delta(x_{1i};(x_1,x_2^*(x_1))) = G_1(x_{1i};(x_1,x_2^*(x_1))) - R_1(x_{1i};(x_1,x_2^*(x_1))).$$

A threshold strategy x_1' constitutes an equilibrium in period 1 only if agent i observing $x_{1i} = x_1'$ is indifferent between investing and delaying, that is, the payoff difference is zero:

$$\Delta(x_1';(x_1',x_2^*(x_1'))) = 0. \tag{1.8}$$

Likewise, we show in Appendix that $\Delta(x_1';(x_1',x_2^*(x_1')))$ is continuous, strictly increasing in x_1' and converges to $-\infty$ (*resp.* ∞) as x_1' converges to $-\infty$ (*resp.* ∞), so that a unique solution, denoted by x_1^*, to (1.8) exists such that $\Delta(x_1^*;(x_1^*,x_2^*(x_1^*))) = 0$. And x_1^* is the only candidate for equilibrium threshold strategies in period 1. Indeed, it is optimal because

$$\frac{\partial}{\partial x_{1i}}\Delta(x_{1i};(x_1^*,x_2^*(x_1^*))) > 0, \tag{1.9}$$

namely, investing in period 1 is optimal ($\Delta(x_{1i};(x_1^*,x_2^*(x_1^*))) > 0$) if $x_{1i} > x_1^*$. And (1.9) can be obtained by the same way in which we compute $\partial\Delta(x_1';(x_1',x_2^*(x_1')))/\partial x_1' > 0$ in Appendix. Setting $x_2^* = x_2^*(x_1^*)$, then (x_1^*,x_2^*) is the unique threshold equilibrium stated in the proposition.

Q.E.D.

We have focused on symmetric strategies and this is without loss of generality, as Remark 1 shows; the essence is that every agent is infinitesimally small and faces the same decision problem.

Remark 1 (Exclusion of Asymmetric Strategies). *There can only be symmetric threshold strategies in equilibrium. Suppose by contradiction that the unit agents are divided into groups and each group in equilibrium follow threshold* $(x^*_{11}, x^*_{12}, \cdots, x^*_{1N})$ *respectively in period 1 and* $(x^*_{21}, x^*_{22}, \cdots, x^*_{2M})$ *respectively in period 2, where* $N, M \in \mathbb{N}$*. Then in period 2, an agent in group* $i \in \{1, \cdots, M\}$ *who observes* x^*_{2i} *must be indifferent between investing or not:*

$$E[r + \hat{a}_2 | x^*_{2i}; (x^*_{11}, \cdots, x^*_{1N}, x^*_{21}, \cdots, x^*_{2i}, \cdots, x^*_{2M})] = 0. \tag{1.10}$$

Since agents are infinitesimally small, the value of $\hat{a}_2$ *only depends on the thresholds of the population and is invariant of individuals' actions. Hence we have by* (1.10) *that* x^*_{2i} *equals the negative* $\hat{a}_2$*. Similarly for a group* $j \neq i$ *agent, she solves*

$$E[r + \hat{a}_2 | x^*_{2j}; (x^*_{11}, \cdots, x^*_{1N}, x^*_{21}, \cdots, x^*_{2M})] = 0,$$

and thus x^*_{2j} *also equals the negative* $\hat{a}_2$ *and thus equals* x^*_{2i}*. The similar argument applies to period 1.*

Remark 2 (Investment of Variable Size). *If the action space is replaced by an interval* $[0, 1]$ *with* $\sum_{t=1}^{2} a_{ti} \in [0, 1]$*, the monotone equilibrium stays unchanged. That said, no agent will split their endowment even if they can. This is so because the monotone equilibrium* (x^*_1, x^*_2) *holds due to* $\{0, 1\} \subset [0, 1]$*, and its uniqueness gives the result. Intuitively, when agents expect positive returns and face delay costs, it is not wise for them to keep endowment unused, whereas when they expect negative returns, being infinitesimal means that investing has neither payoff nor information values.*

The proposition establishes the uniqueness in a monotone equilibrium, consistent with the static global games literature. Yet the general uniqueness, namely the exclusion of other strategy forms, does not obtain. Even though taking a threshold strategy is intuitively appealing because of payoff's increasing monotonicity in the state. However, due to coordination motives, if agents believe their opponents take some other specific strategy, they may follow that form of the strategy. More technically speaking, when iteratively eliminating strictly dominated strategies (which is the key to establish general uniqueness; see Proposition 1), we will encounter an open interval of signal realizations in which all strategy forms are plausible to constitute an equilibrium.

What is more severe here is that, with endogenous learning from past activity, since arbitrary strategy forms in period 1 need to be taken into account in solving for general equilibria, agents face arbitrary information structures about r from observing a_1, thereby leaving the room for other forms of equilibria.

Remark 3 (Complementarities in Action Timings)**.** *In the model, agents contemplate their own action timing but not those of others, because the payoff depends on the activities* $\hat{a}_2$ *throughout the game. If, instead, the payoff to investing at time* t *depends only on the current investment size* a_t*, multiplicity also can occur. This is so because agents now have coordination motives in action timing, and if an agent believes all others will act in one particular period, so will her; see Dasgupta et al. (2012) for further discussion on this line of reasoning.*

1.2.4 Equilibrium Analysis

In this section, we contrast agents' behaviors between the static and the dynamic games, and investigate the values of learning and the delay option. The consequent welfare implications are also discussed. All the conclusions apply to the general N-period model.

Changes in Behavior

The two-period game can be perceived as (i) adding a stage before the static game or (ii) adding a stage afterwards. We first consider case (i) and compare agents' behavior in period 2 of the dynamic game to that in the static game. Results show that agents in period 2 tend to invest more frequently than in the static game ($x_2^* < x_{st}^*$). Intuitively, in period 2, there is no delay option, which is the same as in the static game; meanwhile, agents know if the game were static, the same fraction of agents would invest, and adding an additional previous stage means weakly more agents invest. So the conclusion follows by the strategic complementarity. For case (ii), agents in period 1 of the dynamic game are tempted by the delay option and thus invest less frequently than in the static game (i.e., $x_{st}^* < x_1^*$), reflecting the informational value from learning. The following proposition establishes these results.

Proposition 3. *Comparing to in the static game, agents are more aggressive in period 2, and less in period 1:* $x_2^* < x_{st}^* < x_1^*$.

Proof. Note that $\hat{a}_2(x_2^*) = a_1(x_2^*) + (1 - a_1(x_2^*))P(\hat{x}_{2j} > x_2^*|x_2^*) > 1/2 = a_1(x_{st}^*)$, so if $x_2^* \geqslant x_{st}^*$, then $E[r + \hat{a}_2|x_2^*] > E[r + a_1|x_{st}^*] = 0$, contradicting the equilibrium condition of period 2 in the dynamic game.

For the second half, note that if $\delta = 0$, then $R_1(x_{1i}; (x_1^*, x_2^*)) = 0$ and $a_2(x_{1i}) = 0$ (recall that agents in period 2 being indifferent to invest or not choose action 0), for any x_{1i}. Hence x_1^* solves $E[r + a_1|x_1^*] = 0$ and thus equals x_{st}^*. As δ increase, R_1 also increases, so x_1^* must increase to balance the equilibrium condition (1.8) of period 1. Therefore $x_1^* > x_{st}^*$.

Q.E.D.

Noteworthy, within the two-period game, agents behave more aggressively in period 2 than in period 1 since $x_2^* < x_1^*$. Contrasting this phenomenon with case (i) earlier, in which that agents in period 2 invest more frequently than in the static game ($x_2^* < x_{st}^*$) is due to coordination motives. Here for the dynamic game, coordination motives do not play a role since agents enjoy the same payment $\hat{a}_2$ whichever period they invest. Instead, here is because of the decreasing continuation payoff that changes from a strictly positive value R_1 to 0 at the last stage. Following this line of logic, we obtain the effect of continuation payoffs on agents' behavior: the lower continuation payoffs to delaying to the next period, the more aggressively agents behave in the current period. Its proof follows the proof of the second half in Proposition 3.

The Value of Information and Welfare Analysis

Note that in equilibrium,

$$\delta E[r+\hat{a}_2 \mid x_{1i}] = R_1(x_{1i}) + \underbrace{\delta \int_{-\infty}^{x_2^*} E[r+\hat{a}_2 \mid \hat{x}_{2i}] f(\hat{x}_{2i} \mid x_{1i}) d\hat{x}_{2i}}_{<0} < R_1(x_{1i}),$$

where the negativity of the second term is by the definition of x_2^*. The term $\delta E[r+\hat{a}_2 \mid x_{1i}]$ is the expected payoff to delaying without learning (i.e., when agent i holds a constant signal x_{1i}), which is shown strictly lower than the continuation payoff with learning existed; hence the value of information is positive. To see the intuition of why learning improves agents' expected payoffs, note that learning makes agents' signals better aligned, alleviating their strategic uncertainty and thereby making them better coordinate. In addition, learning mitigates the payoff uncertainty, as is reflected in $\hat{\tau}_2 > \tau_1$, namely, agents better infer the state in the presence of learning.

We next compare the interim welfare of agents between the dynamic and the static games, after agents' signals are realized yet before the state is revealed. Results show that agent i expects a higher payoff in the dynamic game when (i) she invests in the first stage in the dynamic game ($x_{1i} > x_1^*(> x_{st}^*)$), or (ii) she invests in period 2 ($x_{1i} \leqslant x_1^*$ and $\hat{x}_{2i} > x_2^*$) and her belief is driven upward after learning ($\hat{x}_{2i} > x_{1i}$). The increased welfare in case (i) originates from the higher expected total investment in the dynamic game, and that in case (ii) is due to, by $\hat{x}_{2i} > x_{1i}$, both higher state and higher investment size.

Some computation gives the conclusion. For example, in period 1, the expected payoff for agent i with x_{1i} is

$$\begin{cases} x_{1i} + a_1(x_{1i}), & \text{if } x_{1i} > x_{st}^* \\ 0, & \text{otherwise,} \end{cases}$$

in the static game, and

$$\begin{cases} x_{1i} + a_1(x_{1i}) + a_2(x_{1i}), & \text{if } x_{1i} > x_1^* \\ R_1(x_{1i}; (x_1^*, x_2^*)), & \text{if } x_{1i} < x_1^*, \end{cases}$$

in the dynamic game. Therefore, i's welfare increases if $x_{1i} > x_1^*(> x_{st}^*)$. Otherwise if $x_{1i} < x_1^*$, the agent proceeds to period 2 and similar comparison can be made. Note that there are inconclusive situations in which the direction of the welfare change depends on cost parameter δ versus information precision τ_1 and τ_2. For example, when i invests in both games but learning drives her belief down ($x_2^* < \hat{x}_{2i} < x_{st}^* < x_{1i} < x_1^*$), then welfare comparison depends on $x_{1i} + a_1(x_{1i})$ versus $\hat{x}_{2i} + \hat{a}_2(\hat{x}_{2i})$.

The Value of Delay

This subsection explores the option value of delay in isolation from the learning effects. To this end, we consider the game in which agent i cannot observe x_{2i}; one can think of it as $\tau_2 \to 0$, so that x_{2i} is completely noisy and ignored.

Proposition 4. *When x_{2i} is not observable, the dynamic game is essentially static: $x_2^* \geqslant x_1^* = x_{st}^*$.*

Proof. Suppose that agent i holds a constant belief x_{1i}. Her expected payoff to investing in period 2 is $\delta E[r+\hat{a}_2 \mid x_{1i}]$, so that if she will invest, she will only invest in period 1 due to $\delta < 1$. That said, the agent in period 2 stays inactive for sure, so x_2^* can be any number larger than x_1^*. Since agents will not invest in period 2, the payoff to delaying to period 2 is 0 and also $\hat{a}_2 = a_1$, so x_1^* is such that $E[r+a_1|x_1^*] = 0$ and thus equals x_{st}^*. *Q.E.D.*

Intuitively, with no learning benefit but only cost from delaying, agents act (if at all) as soon as possible, as is indicated by that the continuation payoff to delaying to period 2 is at most $\delta E[r+\hat{a}_2 \mid x_{1i}]$, a mere discounted current payoff. Consequently $\hat{a}_2 = a_1$, so that threshold $x_1^* = x_{st}^*$ and the strategic stage ends there.

1.2.5 Learning Efficiency

In this subsection, we pay attention to the learning mechanism in our paper, and contrast it with learning through directly observing the state r. That is, instead of observing an endogenous signal about past activity as in (1.3), if agent i is directly endowed with an exogenous signal x_{2i} such that

$$x_{2i} = r + \frac{1}{\sqrt{\tau_2}}\varepsilon_{2i}, \tag{1.11}$$

will she infer the state r more accurately at $t = 2$? Surprisingly, the agent estimates r more accurately when she learns though observing action, as long as her initially information precision $\tau_1 > 1$. To see it, by Bayes' rule, there exists a sufficient signal $\hat{\hat{x}}_{2i}$ for agent i that summarizes her information about r contained in x_{1i} and x_{2i}, such that

$$\hat{\hat{x}}_{2i} = r + \frac{1}{\sqrt{\tau_1+\tau_2}}\varepsilon_{2i},$$

when learning is through directly observing r as in (1.11). Recall that the precision level of agent i's information by indirect learning is $\hat{\tau}_2 = \tau_1 + \tau_1\tau_2$ at $t = 2$. Therefore, as long as $\tau_1 > 1$, learning through observation of actions reveals more accurate information about r.

Intuitively, learning efficiency of direct observation on r is fixed (=$\tau_1 + \tau_2$), while its precision level depends on how accurate agents know about r (measured by τ_1) and how accurate the endogenous signal reflects their private information (measured by τ_2), when learning is through observing the past activity. The two channels are mutually reinforced, as is reflected by that indirect learning precision in period 2 is $\tau_2' = \tau_1\tau_2$. Therefore, indirect learning can be more accurate when one of the channels is accurate enough, and we confirm the condition is $\tau_1 > 1$.

1.3. The N-Period Model

We now augment the game to $N \in \mathbb{N}$ periods and consider a general payoff structure. In the game, the unit of agents decide the optimal timing of investment (if at all) between $t = 1, 2, \cdots, N$. The first two periods run identically as before, and notations a_{ti}, a_t, and $\hat{a}_t$ bear similar meanings. Agent i's payoff in period t to action $a_{ti} = 0$ is 0, and her payoff to $a_{ti} = 1$ is now summarized by an increasing and continuously differentiable function $U(r, \hat{a}_N)$, namely, the return of investment increases in state r and aggregate investment $\hat{a}_N$, indicating the coordination feature of the game. We assume that $U(r, \hat{a}_N)$ is concave in each component and that $\lim_{r\to\infty} U(r, \hat{a}_N) = \infty$ and $\lim_{r\to-\infty} U(r, \hat{a}_N) = -\infty$, for any $\hat{a}_N \in [0, 1]$. By implication, when the state is extremely good (or bad), investing strictly dominates (or is strictly dominated) regardless of others' actions.

We now describe the endogenous signals that agents receive in periods $t = 3, \cdots, N$. To maintain analyticity and similarity to Section 2, we let agent i observe, in $t = 3, \cdots, N$,

$$x_{ti} = \Phi^{-1}(\bar{a}_{t-1}) + \frac{1}{\sqrt{\tau_t}}\varepsilon_{ti}, \quad \tau_t > 0, \tag{1.12}$$

where $\bar{a}_{t-1} = (\hat{a}_{t-1} - \hat{a}_{t-2})/(1 - \hat{a}_{t-2})$ is average action in $t-1$ (note that $\hat{a}_{t-1} - \hat{a}_{t-2}$ denotes the new investment at $t-1$ and that $1 - \hat{a}_{t-2}$ the fraction of agents who reach $t-1$), and $\varepsilon_{ti} \sim \mathcal{N}(0, 1)$ is independent of all other variables. By LLN, the average action equals the likelihood of investment for an individual agent; therefore, when agents follow a threshold strategy profile denoted by $(x_1, x_2, \cdots, x_N)$ in period $1, 2, \cdots, N$, $\bar{a}_t(r) = P(\hat{x}_{ti} > x_t | r)$, where $\hat{x}_{ti}$ is i's expectation of r at time t. Note that the structure is consistent with x_{2i} defined in Section 2 since the average action $\bar{a}_1 = a_1$.

In each period t, agent i still chooses a_{ti} to maximize her conditional expected total payoff $E[\sum_{t=1}^{N} a_{ti}\delta^{t-1}U(r, \hat{a}_N) | x_{1i}, \cdots, x_{ti}]$, where $\delta \in (0, 1)$.

1.3.1 Learning Under a Threshold Strategy

As in the two-period setup, for $t = 3, \cdots, N$, x_{ti} can be transformed into an informationally equivalent (with respect to r) signal x'_{ti} centered around r, when agents follow a threshold strategy profile before time t. The definition of a threshold strategy profile is similar to that in Section 2 and thus omitted. Agent i's updated belief about r in period t can be summarized by a unidimensional statistic $\hat{x}_{ti}(x_{1i}, x_{2i}, \cdots, x_{ti})$ that is normally distributed given r. We still let $\hat{x}_{1i} = x'_{1i} = x_{1i}$ and $\hat{\tau}_1 = \tau'_1 = \tau_1$. It turns out that the precisions of x'_{ti} and $\hat{x}_{ti}$, denoted by τ'_t and $\hat{\tau}_t$ respectively, are such that $\tau'_t = \tau_t \hat{\tau}_{t-1}$ and $\hat{\tau}_t = \sum_{k=1}^{t} \tau'_k$, for all $t \geqslant 2$. Lemma 2 summarizes the results.

Lemma 2. *Suppose that agents follow a threshold strategy profile with respective thresholds* $\{x_1, x_2, \cdots, x_N\}$.

(i) Let $x'_{ti} \equiv x_{ti}/\sqrt{\hat{\tau}_{t-1}} + x_{t-1}$ *for any* i *and* $t \geqslant 2$*; then* x'_{ti} *is sufficient for* x_{ti} *with respect to* r *and*

$$x'_{ti} = r + \frac{1}{\sqrt{\tau'_t}}\varepsilon_{ti}.$$

(ii) $\hat{x}_{ti}$ *can be expressed by* $\hat{x}_{ti}(x'_{1i}, \cdots, x'_{ti}) = (\sum_{k=1}^{t} \tau'_k)^{-1}(\sum_{k=1}^{t} \tau'_k x'_{ki})$ *and particularly,*

$$\hat{x}_{ti} = r + \frac{1}{\sqrt{\hat{\tau}_t}}\varepsilon_{ti}.$$

(iii) For any $t, k \in \{1, 2, \cdots, N\}$, $\hat{x}_{tj}$ *is normally distributed given* $\hat{x}_{ki}$ *(when* $i = j$*, let* $t > k$*). And moreover,* $E[\hat{a}_N|\hat{x}_{ki}]$ *increases in* $\hat{x}_{ki}$.

Proof. The proofs are by indication on t. For (i), it holds at $t = 2$ by Section 2. Assume inductively that it holds until $t = N-1$. Then $\bar{a}_{N-1}(r) = \Phi(\sqrt{\hat{\tau}_{N-1}}(r - x_{N-1}))$. So at $t = N$,

$$x_{Ni} = \sqrt{\hat{\tau}_{N-1}}(r - x_{N-1}) + \frac{1}{\sqrt{\tau_N}}\varepsilon_{ti}.$$

Rearranging and comparing it with x'_{Ni} and τ'_N give the conclusion. Then (ii) follows by Bayes' rule.

For (iii), when $i \neq j$, the first part is similar to Lemma 1 (iii). When $i = j$ and let $t > k$, it follows

$$\hat{x}_{ti}|_{\hat{x}_{ki}} = \frac{\hat{\tau}_k \hat{x}_{ki} + \tau'_{k+1} x'_{(k+1)i} + \cdots \tau'_t x'_{ti}}{\hat{\tau}_k + \tau'_{k+1} + \cdots \tau'_t}|_{\hat{x}_{ki}},$$

and note that for $n \in \{k+1, \cdots, t\}$, $x'_{ni}|_{\hat{x}_{ki}} = r + \varepsilon_{ni}/\sqrt{\tau'_n} = \hat{x}_{ki} - \varepsilon_{ki}/\sqrt{\hat{\tau}_k} + \varepsilon_{ni}/\sqrt{\tau'_n}$ is normally distributed given $\hat{x}_{ki}$.

The monotonicity of $E[\hat{a}_2|\hat{x}_{ki}]$ holds by Lemma 1. Assume inductively that $\hat{a}'_{t-1}(\hat{x}_{ki}) > 0$ till $t = N-1$. Then at $t = N$,

$$\frac{d}{d\hat{x}_{ki}}\hat{a}_N(\hat{x}_{ki}) = \frac{d}{d\hat{x}_{ki}}[\hat{a}_{N-1} + (1-\hat{a}_{N-1})\bar{a}_N](\hat{x}_{ki}) = (1-\bar{a}_N)\hat{a}'_{N-1} + (1-\hat{a}_{N-1})\bar{a}'_N > 0.$$

Q.E.D.

These results are extensions to those in Section 2 and follow the discussions there.

1.3.2 Equilibrium Characterization

We still restrict to a monotone equilibrium and now solve for it. Provided that agents follow a threshold strategy profile denoted $(x_1, x_2, \cdots, x_N)$, the expected continuation payoff $R_t(\hat{x}_{ti}; \{x_t\}_{t=1}^N)$ of agent i with $\hat{x}_{ti}$ at time t is

$$\begin{aligned} R_t(\hat{x}_{ti}; \{x_t\}_{t=1}^N) = \delta \int_{x_{t+1}}^{\infty} & E[U(r, \hat{a}_N) \mid \hat{x}_{(t+1)i}] f(\hat{x}_{(t+1)i} \mid \hat{x}_{ti}) d\hat{x}_{(t+1)i} \\ + \delta \int_{-\infty}^{x_{t+1}} & R_{t+1}(\hat{x}_{(t+1)i}; \{x_t\}_{t=1}^N) f(\hat{x}_{(t+1)i} \mid \hat{x}_{ti}) d\hat{x}_{(t+1)i}, \end{aligned} \tag{1.13}$$

where $f(\hat{x}_{(t+1)i} \mid \hat{x}_{ti})$ is the conditional density of $\hat{x}_{(t+1)i}$ on $\hat{x}_{ti}$, whose value can be deduced by Lemma 2 (iii). Let $R_N(\hat{x}_{Ni}; \{x_t\}_{t=1}^N) \equiv 0$ for any $\hat{x}_{Ni} \in \mathbb{R}$. Note that we have let $R_t(\cdot)$ represent the face value at time t, instead of being discounted to time 1. The following proposition characterizes the unique monotone equilibrium.

Proposition 5. *There exists a unique monotone equilibrium characterized by* $(x_1^*, x_2^*, \cdots, x_N^*)$ *in the N-period game, where* x_t^* *is the unique solution to*

$$E[U(r, \hat{a}_N)|x_t^*] = R_t(x_t^*; \{x_t^*\}_{t=1}^N), \quad t = 1, 2, \cdots, N. \tag{1.14}$$

Similar to that in the two-period model, the proof starts from the last period N and takes as given that all agents in all previous periods play some threshold strategy profile, so as to characterize x_N^*. Next proceeding the argument backward and in each period $1 \leqslant t \leqslant N-1$, it is taken as given that agents play some threshold strategy profile before t and act optimally after t, and sequentially obtains $x_{N-1}^*, x_{N-2}^*, \cdots, x_1^*$. Recall that in checking that no agent wants to deviate at time t, the key is that $x_{t+1}^*, \cdots, x_N^*$ will not be disturbed by an infinitesimally small agent's deviation.

Proof. Fix an arbitrary threshold strategy profile $(x_1, x_2, \cdots, x_N)$ and an agent i. At $t = N$, given that all agents expect i follow $(x_1, x_2, \cdots, x_N)$ in the game, the payoff of agent i with belief $\hat{x}_{Ni}$ to investing is

$$E[U(r, \hat{a}_N) \mid \hat{x}_{Ni}; (x_1, \cdots, x_N)], \tag{1.15}$$

which increases in $\hat{x}_{Ni}$. An threshold strategy x'_N constitutes an equilibrium threshold at $t = N$ only if observing it makes agent i indifferent between investing or not, that is, it is such that

$$E[U(r, \hat{a}_N) \mid x'_N; (x_1, \cdots, x_{N-1}, x'_N)] = 0.$$

Similarly as in Proposition 2, the LHS is continuous, converges to infinity as x'_N converges to infinity, and strictly increases in x'_N, so there exists a unique such x'_N that solves the above equation. And the increasing monotonicity of (1.15) verifies that the solution indeed constitutes an equilibrium threshold at $t = N$. We denote it by $x_N^*(x_1, x_2, \cdots, x_{N-1})$ and shorthand it by x_N^*.

Proceeding to $t = N-1$, taken as given that all agents expect some agent i follow $(x_1, x_2, \cdots, x_{N-1}, x_N^*)$, the payoff of investing immediately to agent i with $\hat{x}_{(N-1)i}$ is

$$E[U(r, \hat{a}_N) \mid \hat{x}_{(N-1)i}; (x_1, \cdots, x_{N-1}, x_N^*)],$$

while delaying to the next period has an expectation value given by

$$R_{N-1}(\hat{x}_{(N-1)i};(x_1,\cdots,x_{N-1},x_N^*)).$$

Let $\Delta_{N-1}(\hat{x}_{(N-1)i};(x_1,\cdots,x_{N-1},x_N^*))$ denote the payoff difference for i investing at $N-1$ or delaying, that is,

$$\Delta_{N-1}(\hat{x}_{(N-1)i};(x_1,\cdots,x_{N-1},x_N*)) \equiv$$
$$E[U(r,\hat{a}_N) \mid \hat{x}_{(N-1)i};(x_1,\cdots,x_{N-1},x_N^*)] - R_{N-1}(\hat{x}_{(N-1)i};(x_1,\cdots,x_{N-1},x_N^*)).$$

An threshold strategy x'_{N-1} that constitutes an equilibrium threshold at $t=N-1$ must be such that

$$\Delta_{N-1}(x'_{N-1};(x_1,\cdots,x_{N-2},x'_{N-1},x_N^*)) = 0. \tag{1.16}$$

We demonstrate in Appendix the unique existence of such x'_{N-1} that solves (1.16), by showing the LHS is strictly increasing in x'_{N-1} and converges to infinity as x'_{N-1} goes infinity. Also, we confirm that the solution indeed constitutes an equilibrium at $t=N-1$ by showing in Lemma 3 in Appendix that $\Delta_{N-1}(\hat{x}_{(N-1)i};(x_1,\cdots,x_{N-1},x_N*))$ increases in $\hat{x}_{(N-1)i}$. Let the solution be denoted by $x^*_{N-1}(x_1,x_2,\cdots,x_{N-2})$, or for notational simplicity, by x^*_{N-1}. By backward induction and similarly, we can characterize $x^*_{N-2},x^*_{N-3},\cdots,x^*_1$. *Q.E.D.*

It is noteworthy that $\Delta_t \neq \Delta_k$ so that $x_t^* \neq x_k^*$ for $t \neq k \in \{1,\cdots,N\}$, indicating agents respond to information changes continuously and that a positive fraction of them move from not investing to investing every period. This observation is in contrast to the dynamic regime change games (*cf.* Angeletos et al. (2007)) in whose model agents stay inertia for a series of periods. The difference occurs because the payoff structure in this paper is continuous in r and $\hat{a}_N$, while it is discrete in their regime change game.

However, if learning precisions $\{\tau_t\}_{t\geqslant 2}$ are moderate (so $\hat{x}_{ti} \approx \hat{x}_{(t+1)i}$) and the delaying cost is not too severe (e.g., $\delta \to 1$ so $R_t \approx R_{t+1}$), the number of new active agents between periods should be small, since the differences between the continuations payoffs evaluated at t and $t+1$ are small. So x_t^* and x_{t+1}^* are near. In this situation, agents' behaviors experience relative inertia in intermediate periods (also documented in Angeletos et al. (2007)), and the dynamics of the game are now such that an active first stage followed by a relative tranquil phase, till the last stage at which less optimistic agents also invest, because the continuation value of delay in the last period drops discontinuously to 0 from some positive number R_{N-1}.

1.3.3 Equilibrium Analysis

In this section, we confirm our two-period results. The positive information value is easy to obtain, since

$$\begin{aligned}\delta E[U(r,\hat{a}_N) \mid \hat{x}_{ti}] = \delta \int_{x_{t+1}^*}^{\infty} E[U(r,\hat{a}_N) \mid \hat{x}_{(t+1)i})] f(\hat{x}_{(t+1)i} \mid \hat{x}_{ti}) d\hat{x}_{(t+1)i} \\ + \delta \int_{-\infty}^{x_{t+1}^*} \underbrace{E[U(r,\hat{a}_N) \mid \hat{x}_{(t+1)i})]}_{<R_{t+1}(\hat{x}_{(t+1)i})} f(\hat{x}_{(t+1)i} \mid \hat{x}_{ti}) d\hat{x}_{(t+1)i} < R_t(\hat{x}_{ti}),\end{aligned}$$

by the definition of $R_t(\hat{x}_{ti})$. Perceiving the first term $\delta E[U(r,\hat{a}_N) \mid \hat{x}_{ti}]$ as the expected payoff at time $t+1$ in the absence of learning, the strict inequality then shows the information is of positive value. Next, we verify (i) comparing to in the static game, agents are more aggressive in the last stage ($x_N^* < x_{st}^*$) and less ($x_{st}^* < x_t$) in earlier periods $t < N$, and (ii) the mere delay option has zero impact. The proofs are relegated to Appendix.

Proposition 6. *(i) $x_N^* < x_{st}^* < x_t^*$, for $t = 1, 2, \cdots, N-1$.*

(ii) When learning does not exist such that x_{ti} for any $t \geqslant 2$ is unobservable, the game is essentially static: $x_1^* = x_{st}^*$ and agents stay inactive after period 1.

1.4. Discussions

1.4.1 Learning Efficiency

We now investigate the learning efficiency of observing actions, by comparing it with learning through directly observing the state r. That is, if the signal structures of x_{ti}, for $t = 2, 3, \cdots, N$, are such that

$$x_{ti} = r + \frac{1}{\sqrt{\tau_t}} \varepsilon_{ti}, \tag{1.17}$$

will it improve the accuracy with which agents infer the state r, relative to the signal structures (1.12) in the paper?

Proposition 7. *If the initial information is precise $\tau_1 > 1$, observing the actions as in* (1.12) *reveals more accurate information about the state r than directly observing r as in* (1.17)*, for all periods $t \geqslant 2$.*

Recall that the information precision through observing actions is $\hat{\tau}_t = \sum \tau_k'$. Let $\hat{\hat{\tau}}_t$ denote the precision level of learning through observing r. We conclude by comparing them.

Proof. In period $t \geqslant 2$, when agents directly observe r as in (1.17), there exists a sufficient statistic, denoted $\hat{\hat{x}}_{ti}$, of $x_{1i}, x_{2i}, \cdots, x_{ti}$ with respect to r; by Bayes' rule,

$$\hat{\hat{x}}_{ti} = r + \varepsilon_{ti} / \sqrt{\hat{\hat{\tau}}_t}, \quad \text{with } \hat{\hat{\tau}}_t = \tau_1 + \tau_2 + \cdots + \tau_N.$$

Therefore, $\hat{\tau}_t > \hat{\hat{\tau}}_t$ for all $t \geqslant 2$, whenever $\tau_1 > 1$. *Q.E.D.*

1.4.2 Infinite Periods

Now we augment the game into infinite periods by setting $N \to \infty$, and demonstrate that the equilibrium properties are similar as when N is finite. Also, we find that agents fully learn the true state in the limit. Defined analogously, a threshold strategy is denoted by $\{x_t\}_{t=1}^{\infty}$ and continuation payoffs by $\{R_t\}_{t=1}^{\infty}$. We restrict to that agents follow a symmetric threshold strategy $\{x_t\}_{t=1}^{\infty}$ in the game, so we obtain the similar transformed endogenous signals x'_{ti} and cumulative signals $\hat{x}_{ti}$ as in Section 3, since the learning processes only depend on that agents play a threshold strategy.

The equilibrium concept we consider now is, however, an ε-symmetric monotone equilibrium which consists of a symmetric threshold strategy, such that no agent can expect to gain more than $\varepsilon > 0$ by deviating from the strategy, given others also follow it. This enables us to implement the previous backward induction argument in characterizing the equilibrium. In detail, for any $\varepsilon > 0$, due to $\delta \in (0,1)$, there exists $N^*_\varepsilon \in \mathbb{N}$ such that

$$R_t(\hat{x}_{ti}; \{x_t\}_t) < \varepsilon,$$

for every $t \geqslant N^*_\varepsilon$, signal $\hat{x}_{ti}$, and threshold strategy $\{x_t\}$. Henceforth fix a random $\varepsilon > 0$ and consequently an N^*_ε. We claim there exists an ε-monotone equilibrium with an identical equilibrium threshold after period N^*_ε. In what follows, we assume that agents follow some identical threshold after N^*_ε to solve for an equilibrium, and then verify it is indeed optimal for agents to follow such a constant threshold strategy after period N^*_ε.

For any $t \geqslant N^*_\varepsilon$, a threshold strategy $x^*_{N^*_\varepsilon}$ constitutes an equilibrium strategy in period t only if it solves

$$E[U(r, \hat{a}_\infty) \mid x^*_{N^*_\varepsilon}; (x_1, \cdots, x^*_{N^*_\varepsilon}, x^*_{N^*_\varepsilon}, \cdots)] = 0,$$

where $\hat{a}_\infty = \sum_t^\infty a_t \in (0,1)$. Such $x^*_{N^*_\varepsilon}$ exists; to see its monotonicity, as $x^*_{N^*_\varepsilon}$ increases, the state r increases and the expected fractions of investors in periods other than t increase while the expected fraction of investors in period t remains constant (which is $1/2$). To check $x^*_{N^*_\varepsilon}$ indeed is an ε-equilibrium strategy in period t, note that when observing a signal higher than

$x^*_{N^*_\varepsilon}$, deviating from investing (which gives a positive payoff) to delaying (which gives $R_t < \varepsilon$) increases the expected payoff by at most ε; when observing a signal lower than $x^*_{N^*_\varepsilon}$, deviating from not investing to investing clearly lowers the expected payoff. Since t is arbitrary as long as larger than N^*_ε, we have shown that for periods $t = N^*_\varepsilon, N^*_\varepsilon + 1, \cdots$, it is optimal for agents to follow a constant threshold strategy $x^*_{N^*_\varepsilon}$. Next, proceed to period $N^*_\varepsilon - 1$ and take as given that all agents in periods $t \geqslant N^*_\varepsilon$ follow $x^*_{N^*_\varepsilon}$; the argument goes exactly the same as in Proposition 5, so we obtain its unique equilibrium threshold $x^*_{N^*_\varepsilon - 1}$, and proceeding backward to obtain $x^*_{N^*_\varepsilon - 2}$,$x^*_{N^*_\varepsilon - 3}$ till x^*_1.

Noteworthy, agents fully learn the actual state in the limit, because their information precision $\sum_{t=1}^{\infty} \tau_t \to \infty$. Such a property holds even when the equilibrium thresholds are now constant after some certain periods. To see the reason, recall that equilibrium thresholds start to be constantly $x^*_{N^*_\varepsilon}$ from period N^*_ε; then at $t = N^*_\varepsilon$, a positive fraction of agents will move to invest because $x^*_{N^*_\varepsilon - 1} \neq x^*_{N^*_\varepsilon}$. This movement changes the total investment size in period N^*_ε (from that in $N^*_\varepsilon - 1$) and thereby makes agents in period $t = N^*_\varepsilon + 1$ learn new information and consequently, a further fraction of agents will move to invest in period $N^*_\varepsilon + 1$, and so on. Essentially, the fully learning of the state stems from that (i) there is no public learning and hence no crowding out effect as in the herding literature, and (ii) pooling everyone's information reveals the true state. Of course, that observational precisions τ_t being exogenously given and bounded away from zero is also a reason.

1.4.3 Proper Priors and Public Learning

The analysis till now is conducted with agents holding an improper prior, and we claim it is almost without loss of generality. Now we mention how to extend the model to a proper prior game. Let agents hold a common prior as follows:

$$r \sim \mathcal{N}(\alpha, 1/\beta),$$

where $\alpha \in \mathbb{R}$ and $\beta > 0$. Still restrict to agents taking a threshold strategy profile; agent i's belief about r at each stage is summarized by a unidimensional statistic $\hat{x}_{ti}$, by the same Gaussian updating process as in Section 3.1. Consequently, the equilibrium characterization is analogous, so is the analysis part *when* there exists a unique monotone equilibrium.

Moreover, consider the case where learning is from public observation of actions, so the signal structure of period $t \in \{2, 3, \cdots, N\}$ becomes

$$x_{ti} = \Phi^{-1}(\bar{a}_{t-1}) + \frac{1}{\sqrt{\tau_t}}\varepsilon_t,$$

where $\varepsilon_t \sim \mathcal{N}(0, 1)$ represents the market-wise noise, independent of all other variables. Let x_{1i} still be private.

We now elaborate on the potential arise of multiple equilibria in the presence of public learning. To this end, it suffices to consider period 2 and show there exist multiple optimal strategies. Results from Section 2 state that the informativeness of the public signal about r (assuming an improper prior) is $\tau_2' = \tau_1\tau_2$, which converges to infinity as $\tau_1 \to \infty$. Therefore, the ratio of the precision of the public information to the square root of that of private information, namely $\tau_2'/\sqrt{\tau_1}$, diverges to ∞ as $\tau_1 \to \infty$. Hence with public learning, multiplicity in monotone equilibria arises even if private information is infinitely diffused; see Morris and Shin (2004) and Angeletos and Werning (2006) for proofs on why the ratio determines multiplicity. Noteworthy, the proof shows that there are multiple optimal symmetric threshold strategies, to say nothing of optimal strategies in other forms. Intuitively, it is known that complete information coordination games admit multiple equilibria; when the ratio is large, indicating public information dominates private information, global games exhibit similarity to the complete information environment and thus have multiple equilibria. Note that adding a common prior only increases the ratio and thereby only contributes to the rise of multiplicity.

How about the learning property in the limit when learning is public? We have shown that full learning of the true state obtains in the limit with only private observation, and attributed it to the absence of the crowding-out effect from the public information. However, even with public learning as in this subsection, full learning is plausible when agents interact long enough, as long as they play a threshold strategy. To see it, the learning mechanism stays the same as Section 3.1 when agents follow a threshold strategy, so the information precision about the state is always increasing and due to observational precisions τ_t are exogenous, agents in the limit learn the true state. This result crucially depends on the continuous signal structures in this paper, which, shown generally by Lee (1993), avoids the information cascade.

1.5. Conclusions

This paper constructs a dynamic coordination game with learning and delay opportunities factored in. It tractably analyzes agents' optimal action timings, which are determined though constantly trading off the information gain of delay against its opportunity costs. A unique monotone equilibrium is characterized and in it, learning is shown to improve agents' expected payoff, while the mere delay option impose no impact on agents' behaviors, relative to the one-shot game. Additionally, the dynamics of agents' behaviors are characterized and depending on the learning efficiency, the tranquil intermediate periods documented in the literate obtain. Conditions of welfare enhancement, and the contrast to learning by directly observing the state, are also given. The analysis applies for all ranges of learning efficiencies, generalizing the existing studies that usually focus on the limit accurate signals. We illustrate the paper in an investment context; the applicability to other coordination scenarios including currency crises or bank runs is straightforward.

Chapter 2

The Rate of Learning with Public and Private Observations

2.1. Introduction

Aggregate market data accumulate and diffuse the dispersed information in the market and market participants benefit from those data. The data can be publicly known to everyone like the price system, or privately known like the information learned from local communications or at different timings. We in this paper investigate, in the presence of both public and private data, how efficient each of them is in collecting and revealing the market information. The efficiency encompasses two meanings - (i) whether the data can aggregate all the market information and (ii) how long does it take to achieve so, if ever. Vives (1993) demonstrates in a pure public learning framework that the price system indeed completely assembles the market information, but in a long run, due to the crowding-out effect of public data. The market situation varies over time, however, so the information revealed slow can be of little value at the time when agents learn and utilize it. As a result, we add a private learning channel to Vives' model and explore whether this increases the learning efficiency and to what extend.

To be more precise on why the inclusion of a private learning channel is conjectured to increase the learning efficiency, note that the price system is endogenously determined by the market demands versus supplies; therefore, over time, as price accumulates more information, agents respond more to this public signal and less to their private information, so price contains increasingly less private information. This is the well known negative externality of public information on learning processes, or the crowding-out effect. Therefore, adding a private learning channel in principle alleviates this crowding-out problem, as agents now respond to private endogenous signals as well.

In fact, the inclusion of a private learning channel is more than an ex-post tool to improve the learning efficiency, but rather a starting point to reflect the actual learning processes, because local interactions and observing on neighbours' behaviors are prevalent and inevitable in practice. Furthermore, even for the same piece of information (like trade data), market participants check t at different timings and thus obtain differently imperfect information; such behavior also represents private learning.

Aside from the practical significance, as is mentioned that market situation varies in the long run, it is also theoretically important to understand how much speed of the learning process can be improved (if at all) with the incorporation of a private learning channel. We will prove that allowing agents to observe private endogenous signals indeed improves the learning efficiency and obtain the exact learning rate in the limit (which is linear).

In the detail of our model, we incorporate a private learning channel to the pure public learning model of Vives (1993). A continuum of privately informed agents trade a risky asset repeatedly, and the return of the asset depends on a deterministic yet ex-ante unknown state. Our interest is on whether agents can learn this state by repeated interactions (i.e., without exogenous information expect the one endowed at the beginning of the game), how many rounds of interactions are needed, and the features of the learning process. The price of the asset varies among periods and in every period, it is endogenously determined by the market force through the demand versus supply relation, and hence the price contains agents' private information. To embed both private and public learning, we let agents observe, at the beginning of every period, the realized prices publicly and the past aggregate activities privately (i.e., observing is with idiosyncratic noise). After a certain period $T \in \mathbb{N}$, the state is revealed, and agents get paid. The game then ends.

We solve for bayesian equilibria of this model and obtain a unique equilibrium, in which an agent's action is a linear function of his private information and the public information. Noteworthy, in decision-making, the weight that an agent puts on her private information is positively correlated with the precision level of her private information, and negatively with the public information's precision. This result verifies the crowding out effect of public information, and (informally) supports our earlier conjecture that private learning improves the learning efficiency. In addition, note that, to make a decision, an agent must infer the state, while the observation is about past prices and activities; therefore, to tract the updating process of an agent's information and for later equilibrium analysis, we must transform those endogenous signals to some informationally equivalent statistics that are centered around the state. A well-behaved transformation is achieved, and we further show such transformation, or equivalently the way agents update their beliefs, is unique in equilibrium.

Since the equilibrium is unique and the equivalent signals are well-behaved, we obtain clean and analytical results concerning the learning process. Primarily, we find that both public and private information converges to limit accuracy when agents interact long enough, and the asymptotic rates are linear (the rate of learning is $t > 0$ if τ_t / t converges to a positive, less than infinity, constant, where τ_t is the information precision at time t). This finding states that both price and aggregate actions become precise with enough rounds of interactions, consistent with the pure public learning literature. Moreover, Vives (1993) finds that the learning rate in the limit is $t^{1/3}$, when learning is purely public. Therefore, we confirm that a private learning channel increases the learning efficiency. To be concrete, if it takes 1000 rounds for agents to reach some precision level when learning is through a public channel, it takes around 10 rounds for agents to reach the same level of precision when a private learning channel is incorporated.

Note that standard learning rate from observing i.i.d. exogenous signals that are centered around θ is also linear. Hence we have demonstrated that the incorporation of a private learning channel improves the learning efficiency to the extend that the learning rate is restored to the standard one. The essential reason for the linear rate is because both private and public signals become limit accurate, and consequently (i) agents equally respond to each type of signal and (ii) the limit accuracy indicates the asymptotic precisions of agents' private and public information are both constant (infinity), so that they are as if observe i.i.d. signals in the limit, which is indeed the case and will be verified in the paper.

Though adding a private learning channel increases the learning rate, the learning is still slow, by comparing to the rational expectations models of Grossman and Stiglitz (1976) and Grossman and Stiglitz (1980). In their model, price finishes accumulating information instantly or simultaneously when the agents act. In our model, the informativeness of either private or public signals can increase at most a positive constant, which is bounded by the observational errors of those endogenous signals. What private learning contributes in this regard is that it increases the per-period increments in informativeness of agents' information (in inferring the state). However, such slowness should be regarded as a success of our learning process, because as Grossman and Stiglitz argue, when information is reflected too fast, it is better for everyone to wait an instant so as to have others act first, and then to respond with better information. However, if everyone thinks so, no one would actually act first. Such situation is termed the paradox of efficient information markets by them, and is resolved in our model.

Related Literature Our paper is closet to Vives (1993) who investigates the learning rate in a pure public learning environment. Vives shows that agents ultimately learn the true state of the world in the limit at a rate of $t^{1/3}$. Our model introduces a private learning channel and implements a different context (that agents as traders trade a homogenous asset repeatedly and wants to infer its actual value, while Vives thinks of a firm uncertain about the demand tries to learn it from price). Hence our work complements his.

Essentially, both of our models are a (dynamic) extension to the traditional rational expectations models such as Feldman (1987) and Townsend (1978). While in their models, price does not depend on agents' private information but instead, directly depends on the economy state. We endogenize price such that its realization results from the market demand versus its supply, thereby containing agents' private information. Another important distinguish is that agents in rational expectations models make decisions using the information about the state that contains in the spot, unrealized price, while agents in our models instead can only use the information that contains in the past prices. Also, as discussed, Grossman and Stiglitz (1976) and Grossman and Stiglitz (1980) are also classical in static rational expectations models.

Another strand of literature that particularly focuses on the public learning process is the herding literature (e.g., Banerjee (1992) and Bikhchandani et al. (1992)). These studies document that public signals fail to gather all the information among agents, because agents tend to rely too much on public signals and thus stop utilizing their private information. That Vives' model and ours obtain full learning is partly because agents form a continuum (so that the pool information reveals the true state) and action spaces are continuous. As proved generally by Lee (1993), full learning can be achieved in such a situation.

The analysis concerning the learning process has also been extended to various directions. Burguet and Vives (2000) consider that agents can costly acquire private information and examine how the release of public information affects their information acquisition behaviors. In our setup, information acquisition is free and agents automatically observe both types of information in every period. This feature also ensures that agents in our model do not have free rider problems or motivations, which are one of the central issues in the strategic experimentation literature such as Bolton and Harris (1999) and Keller et al. (2005). Related to the experimentation literate, researchers recently start exploring the effects of an outside agency (such as a platform) in soliciting private information; see Kremer et al. (2014) and Che and Hörner (2018). In these models, agents act sequentially and the platform can decide which piece of information an agent can observe. The central question is to design the optimal way for the platform to optimally disclose information to agents, given its goal.

Our analysis is in a discrete time framework; Amador and Weill (2012) examine a continuous time environment with a focus on the welfare implications of public learning, with the co-existence of both public and private learning. They implement the pure prediction framework as in Vives (1997), and obtain their results by explicitly solving stochastic differential questions. We obtain our conclusions by tracking the changes in agents' information in each period.

The remaining of the paper is structured as follows. Section 2 constructs the model and Section 3 presents our core techniques and ideas by a careful analysis of period 1. The general model is solved in Section 4.

2.2. The Model

We introduce a private learning channel, in the form of private endogenous signals about past activities, to the market-based learning model of Vives (1993).

Time is discrete and lasts $T \in \mathbb{N}$ periods, indexed by $t \in \{1, \cdots, T\}$, and a measure-one continuum of agents, indexed by $i \in [0,1]$, independently decide how much of a risky asset to trade in a competitive market in every period. Think of the asset as a financial asset, and let $a_{ti} \in \mathbb{R}$ denote agent i's demand in period t; by construction, the action space is the entire real line and a negative action means a supply or short behavior. Denote by $A_t = \int_i a_{ti} di$ the aggregate action in period t. The payment of the asset is its liquidation value that is only revealed at the end of the game. The value is contingent on economic situations, which are summarized by a parameter θ; we also call θ the state. Trading the asset incurs a quadratic adjustment cost and of course, agents also pay for its spot price. Specifically, agent i pays adjustment cost $\lambda a_{ti}^2/2$ for trading a_{ti} units in period t, with $\lambda > 0$ a known constant. The price of the asset in period t, denoted p_t, is endogenously determined by the market supply versus its demand, with

$$p_t = A_t + \frac{1}{\sqrt{\beta_\varepsilon}} \varepsilon_t, \tag{2.1}$$

where $\varepsilon_t \sim N(0,1)$ is the standard normal, independent of all other parameters and represents the independent periodical demand shock, and $\beta_\varepsilon > 0$ measures the scale of the shock. Note that ε_t commonly affects all agents. One interpretation of ε_t is noise traders in the market who are also sensitive to prices. Therefore, agent i who demands a_{ti} in period t obtains a net payoff $\pi_t(a_{ti})$ with

$$\pi_t(a_{ti}) = (\theta - p_t)a_{ti} - \frac{\lambda}{2} a_{ti}^2. \tag{2.2}$$

The payoffs are paid only after the liquidation value θ is revealed at the end of the game. Hence only $p_t a_{ti} + \frac{\lambda}{2} a_{ti}^2$, instead of $\pi_t(a_{ti})$, is known to agent i after period t.

Several features are worth stressing before we describe agents' information. First, the quadratic adjustment cost indicates the increasing marginal cost of the asset, and this is for avoiding the full revelation of the market information within one period. Since agents form a continuum and are risk neutral, the one-shot game is enough to accumulate all the market information if the cost is with a constant return (Palfrey (1985)). Likewise, the noise traders measured by ε_t are also more than just reflecting the practical market situations; their existence prevents the price from being fully informative after one period.

For the information in the game, the state parameter θ is drawn by nature from a normal distribution with mean p_0 and variance $1/\beta_0$, so $\theta \sim N(p_0, 1/\beta_1)$, where $p_0, \beta_0 > 0$. For notational convenience, we without loss of generality assume that agents instead hold a common improper prior over θ such that $\theta \sim Unif(\mathbb{R})$, and that they observe a public signal p_0 about the realized θ in period 1, such that

$$p_0 = \theta + \frac{1}{\sqrt{\beta_0}}\varepsilon_0,$$

where $\varepsilon_0 \sim N(0,1)$ is the standard normal. [1] Furthermore in period 1, agent i observes a private signal x_{1i} about the realized θ:

$$x_{1i} = \theta + \frac{1}{\sqrt{\tau_1}}\varepsilon_{1i},$$

where $\tau_1 > 0$ is the signal precision and $\varepsilon_{1i} \sim N(0,1)$ is i's observational noise, independent of all other parameters. That said, when agent i makes a decision in period 1, her information set is $\{p_0, x_1\}$. This is the only information that i obtains from exogenous sources in the model. In each of the subsequent periods $t = 2, 3, \cdots, T$, agent i privately observes an

[1] Note that in the improper prior case, p_0 is a random variable (instead of a constant) whose realization depends on the realization of θ.

endogenous signal x_{ti} about the market demand A_{t-1} in the last period in the form of

$$x_{ti} = A_{t-1} + \frac{1}{\sqrt{\tau_\varepsilon}}\varepsilon_{ti}, \tag{2.3}$$

where $\tau_\varepsilon > 0$ is the observational precision and $\varepsilon_{ti} \sim N(0,1)$ are i.i.d. standard normals, representing the idiosyncratic observational noise that is independent of all other parameters and periods. Of course, agent i also (publicly) knows the realized market price p_{t-1} in the last period. Thus in period t, the information set for agent i is $\{x_{ki}, p_{k-1}\}_{k=1}^{t}$. The objective of agent i is to select $a_{ti} \in \mathbb{R}$, for every t, to maximize her expected sum of payoffs $\sum_{t=1}^{T} E[\pi_t(a_{ti})|\{x_{ki}, p_{k-1}\}_{k=1}^{t}]$. Noteworthy, since the action space is $\mathbb{R}$ in every period and an infinitesimally small agent cannot affect A_t and hence cannot affect payoffs or information of other periods, the optimization problems to an agent are essentially static in the sense that it is equivalent to consider agent i maximizes $E[\pi_t(a_{ti})|\{x_{ki}, p_{k-1}\}_{k=1}^{t}]$ separately for every t. This fact also suggests that the undiscounted environment we consider is without loss of generality.

In the paper, we impose the Law of Large Numbers (LLN) convention, so that with probability 1, the proportion of agents who receive signals higher than some number equals the probability of an individual agent receiving such signals. Consequently, the idiosyncratic noise in private signals cancels out: $\int_i \varepsilon_{ti} di = 0$ for all t, and hence we have $\int_i x_{1i} di = \theta$. That is, pooling everyone's information reveals the true state, which suggests it is potential that agents can learn the true state after enough rounds of interactions, even though no other external sources of information exist after period 1. We will verify this conjecture.

Also, it is worth stressing that agents rationally understand the information externality of their aggregate activities, though individually they cannot change its contents. That said, given agent i's information set $\{x_{ki}, p_{k-1}\}_{k=1}^{t}$ in period t, she takes an action using all the information concerning θ in this set. But since the endogenous signals are about actions or prices, we must, and will, transform the information structures to those centered around θ for analysis. We will achieve a well-behaved, and unique transformation, which allows us to analyze the model analytically.

We focus on agents playing a symmetric strategy in the game. With little abusing notation, let $a_t(\{x_{ki}, p_{k-1}\}^t) \in \mathbb{R}$ denote a symmetric strategy of an agent i with information set $\{x_{ki}, p_{k-1}\}^t$ in period t, which prescribes a trading quantity out of $\mathbb{R}$ from her information set. The equilibrium concept is the Bayesian equilibria for the dynamic game defined below.

Definition 1 (Equilibrium)**.** *An equilibrium is a sequence of strategies $\{a_t^*\}_{t=1}^T$, the endogenous signals $\{\{x_{ti}\}_i, p_t\}_{t=1}^T$, and the aggregate actions $\{A_t\}_{t=1}^T$ such that*
(i) a_t^* solves $E[\pi_t(\cdot)|\{x_{ki}, p_{k-1}\}_{k=1}^t]$,
(ii) x_{ti} and p_t are updated by Bayes' rule and follow (2.3) *and* (2.1) *respectively, and*
(iii) A_t is the aggregate action at date t,
for any $t = 1, \cdots, T$, for any realization of $\{\{x_{1i}\}_i, p_0\}$, and θ.

We will demonstrate the equilibrium is unique, which justifies the exclusion of asymmetric strategies. As is discussed, solving the dynamic game is to solve for a sequence of equilibria of the one-shot game; we thus implement an induction argument to solve the game, and this is the key to the uniqueness - when the uniqueness in every period is proved inductively, it also means that the way agents update their information about θ from observing the unique equilibrium actions and prices is also unique.

2.3. Equilibrium Analysis

We illustrate how to solve the dynamic game by carefully analyzing period 1. Knowing how agents will act in period 1, we can characterize the endogenous signals that agents observe in period 2; the updating way turns out to be unique in equilibrium. Then period 2 can be solved in a similar way as solving period 1, and so on.

2.3.1 Period One

Suppose that agent i has observed signal realizations p_0 and x_{1i} at the beginning of period 1. She needs to infer θ to make a decision, and her estimation of θ is a weighted average between his private signal and the common prior:

$$\begin{aligned} E[\theta \mid p_0, x_{1i}] &= \frac{\tau_1 x_{1i} + \beta_0 p_0}{\tau_1 + \beta_0} \\ &\equiv \delta_1 x_{1i} + (1 - \delta_1) p_0, \end{aligned}$$

where $\delta_1 = \tau_1/(\tau_1 + \beta_0)$ is the common weight on the private information. Note that such a convex-combination valuation for θ is common knowledge among agents, so we conjecture (and will verify) that the aggregate action A_1 is also linear and is function of θ and p_0. In turn, as the market price p_1 is determined by A_1 and the demand shock ε_1, we conjecture that there exists at least one equilibrium price in the form as a linear function of θ, p_0, and ε_1, such that

$$p_1(\theta, p_0) = m_1 \theta + n_1 p_0 + \frac{1}{\sqrt{\beta_\varepsilon}} \varepsilon_1, \tag{2.4}$$

where m_1, n_1 are to be determined coefficients. In what follows, we assume that p_1 is in such a linear form and solve for the corresponding equilibrium; the successful characterization pins down the unique m_1 and n_1 and furthermore justifies such p_1. After that we show the linear equilibrium is the only equilibrium.

Recall that agent i in period 1 chooses $a_{1i} \in \mathbb{R}$ to maximize

$$E[(\theta - p_1) a_{1i} - \frac{\lambda}{2} a_{1i}^2 \mid x_{1i}, p_0]. \tag{2.5}$$

Substituting p_1 in the form of (2.4) in (2.5) gives

$$(1-m_1)E[\theta \mid x_{1i}, p_0]a_{1i} - (n_1 p_0 + \frac{1}{\sqrt{\beta_\varepsilon}} E[\varepsilon_1 \mid x_{1i}, p_0])a_{1i} - \frac{\lambda}{2} a_{1i}^2.$$

Solving by the first order condition, we obtain the optimal action $a_1^*(x_{1i}, p_0)$ such that

$$\begin{aligned} a_1^*(x_{1i}, p_0) &= \frac{1}{\lambda} \{(1-m_1)E[\theta \mid x_{1i}, p_0] - n_1 p_0\} \\ &= \frac{1}{\lambda} \{(1-m_1)\delta_1 x_{1i} + [(1-m_1)(1-\delta_1) - n_1]p_0\} \end{aligned}$$

Therefore, the aggregate action A_1 is a function of θ and p_0, and equals

$$A_1(\theta, p_0) = \frac{1}{\lambda} \{(1-m_1)\delta_1 \theta + [(1-m_1)(1-\delta_1) - n_1]p_0\} \tag{2.6}$$

Since the market clearing condition is such that $p_1(\theta, p_0) = A_1(\theta, p_0) + \varepsilon_1/\sqrt{\beta_\varepsilon} = m_1\theta + n_1 p_0 + \varepsilon_1/\sqrt{\beta_\varepsilon}$, by comparing coefficients, we obtain

$$m_1 = \frac{\delta_1}{\lambda + \delta_1}, \quad n_1 = \frac{\lambda}{1+\lambda} \frac{1-\delta_1}{\lambda + \delta_1}.$$

Till now, we have solved for an equilibrium with a linear price function. It is the only equilibrium due to the concavity of the optimization problem (2.5) in a_{1i} (note that p_1 is not affected by a small agent's choice of a_{1i}).

Knowing the forms of m_1, n_1 in equilibrium, we can simplify $a_1^*(x_{1i}, p_0)$ to

$$a_1^*(x_{1i}, p_0) = m_1 x_{1i} + n_1 p_0.$$

Recall that $m_1 = \delta_1/(\lambda + \delta_1)$ increases in δ_1 and that δ_1 increases in τ_1, we confirm an intuitive result that the more precise private signal is, the more weight agents puts on it in making a decision. Consequently, their actions reveal more private information to the market. We can likewise obtain the crowding-out effect of the public information p_0.

Learning from Period 1

We now investigate the learning process in period 2 from observing A_1 and p_1. We have demonstrated that agent i makes a decision based on her estimation of θ, while the endogenous signals x_{2i} and p_1 are centered around A_1. Therefore, we first transform them, respectively, into an informational equivalent statistic centered around θ. We achieve this by demonstrating that, when agents follow the equilibrium strategy in period 1, there exist two statistics, denoted x'_{2i} and p'_1, that are centered around θ from observing A_1 and p_1, respectively, such that the information set $\{x_{1i}, x'_{2i}; p_0, p'_1\}$ has the same informational contents about θ as $\{x_{1i}, x_{2i}; p_0, p_1\}$ does, for any x_{ti} and p_{t-1}, for $t = 1, 2$. After this, we show that the transformation is unique in equilibrium, meaning that all agents in period 2 update their information symmetrically and uniquely. The arguments below are on the equilibrium path.

In equilibrium, $A_1(\theta, p_0) = m_1\theta + n_1 p_0$ by (2.6), so that

$$x_{2i} = m_1\theta + n_1 p_0 + \frac{1}{\sqrt{\tau_\varepsilon}}\varepsilon_{2i}.$$

Rearranging it yields

$$\frac{1}{m_1}(x_{2i} - n_1 p_0) = \theta + \frac{1}{m_1\sqrt{\tau_\varepsilon}}\varepsilon_{2i}.$$

If we define $x'_{2i}(x_{2i}, p_0)$ by

$$x'_{2i}(x_{2i}, p_0) = \frac{1}{m_1}(x_{2i} - n_1 p_0), \tag{2.7}$$

then

$$x'_{2i}(x_{2i}, p_0) = \theta + \frac{1}{\sqrt{\tau_2}}\varepsilon_{2i}, \tag{2.8}$$

where $\tau_2 = m_1^2 \tau_\varepsilon$. Note that θ is normally distributed given $x'_{2i}(x_{2i}, p_0)$ and by construction, x'_{2i} represents agent i's *new* private information (that is only started to be known in period 2 from observing x_{2i}, and no one else knows this information) about θ. More precisely, from observing x_{2i}, agent i newly and privately learns that a normal distribution with mean $x'_{2i}(x_{2i}, p_0)$ and variance $1/\tau_2$ can represent θ, and this is exactly what observing $x'_{2i}(x_{2i}, p_0)$ in the form of (2.8) reveals. Such information represents the new information learned through the private learning channel. Likewise, there is new public information that agent i, and every other agent, learns from observing p_1. Since $p_1 = m_1\theta + n_1 p_0 + \varepsilon_1/\sqrt{\beta_\varepsilon}$ in equilibrium, we have

$$\frac{1}{m_1}(p_1 - n_1 p_0) = \theta + \frac{1}{m_1\sqrt{\beta_\varepsilon}}\varepsilon_1.$$

Therefore, by defining

$$p'_1(p_0, p_1) = \frac{1}{m_1}(p_1 - n_1 p_0),$$

then θ is normally distributed given $p'_1(p_0, p_1)$ such that

$$p'_1(p_0, p_1) = \theta + \frac{1}{\sqrt{\beta_1}}\varepsilon_1,$$

where $\beta_1 = m_1^2 \beta_\varepsilon$. Similarly, $p'_1(p_0, p_1)$ is a statistic that represents the new information about θ to agent i from observing p_1. To economize on notation, we abbreviate them by x'_{2i} and p'_1, and also write $x'_{1i} = x_{1i}$ and $p'_0 = p_0$. Till now, we have transformed the endogenous signals centered around A_1 into signals centered around θ, such that $\{x'_{ti}, p'_{t-1}\}_{t=1,2}$ is informationally equivalent to $\{x_{ti}, p_{t-1}\}_{t=1,2}$ with respect to θ.

Next, we distinguish agent i's private information (x_{1i}, x_{2i}) with the public information (p_0, p_1), and show that each type of information can be respectively summarized by a unidimensional statistic. This procedure allows us to track the private learning channel and the public learning channel separately. To be specific, define

$$\hat{x}_{2i} = \frac{\tau_1 x_{1i} + \tau_2 x'_{2i}}{\tau_1 + \tau_2}; \tag{2.9}$$

then $\hat{x}_{2i}$ indicates i's updated private information about θ in period 2 and by Bayes' rule

$$\theta|_{\hat{x}_{2i}} \sim N(\hat{x}_{2i}, \frac{1}{\hat{\tau}_2}),$$

where $\hat{\tau}_2 = \tau_1 + \tau_2$ is the precision level when i uses only her private information to estimate θ. Likewise, by defining

$$\hat{p}_1 = \frac{\beta_0 p_0 + \beta_1 p'_1}{\beta_0 + \beta_1}, \tag{2.10}$$

then $\hat{p}_1$ is the updated public information about θ in period 2 and moreover,

$$\theta|_{\hat{p}_1} \sim N(\hat{p}_1, \frac{1}{\hat{\beta}_1}),$$

where $\hat{\beta}_1 = \beta_0 + \beta_1$.

Till now, we have shown that, when agents follow the equilibrium in period 1, there exist closed-form statistics x'_{2i} and p'_1 that are centered around θ and represent new information about θ from observing A_1 and p_1, respectively. Consequently, agent i's updated private and public information about θ is summarized in $\hat{x}_{2i}$ and $\hat{p}_2$, respectively, such that θ is normally distributed given each of them. Note that such representations are unique because activity A_1 and price p_1 are unique in equilibrium in period 1. We will show by induction that endogenous signals in later periods $t = 3, 4, \cdots, T$ also have a similar, unique Gaussian transformation. Before that, we note several important features of the learning process.

First, the information equivalence between x_{ti} and x'_{ti} is only for θ. For example, given period-one information p_0 and x_{1i}, observing x_{2i} (but not p_1) does not reveal the same information as observing x'_{2i}; see the definition (2.7) of x'_{2i}. Essentially, this is because A_1 is (at least partially) determined by p_1 and thereby x_{2i} contains the information about p_1, but x'_{2i} is only about θ. While in the set meaning, the complete information equivalence holds: $\{x_{ti}, p_{t-1}\}_{t=1,2}$ is information equivalent to $\{x'_{ti}, p'_{t-1}\}_{t=1,2}$.

Limiting to information about θ, the equivalence between x_{2i} and x'_{2i} allows us to characterize the increment in the informativeness of agent i's private information, which is τ_2. Recall that $\tau_2 = m_1^2 \tau_\varepsilon < \tau_\varepsilon$, suggesting that the indirect learning about θ from observing A_1 is lower than the accuracy of learning A_1 itself. And $m_1 = \delta_1/(\lambda + \delta_1)$ is increasing in δ_1, verifying our intuition that learning precision increases in agents' private information precision (and decreases in that of public information). How to interpret the role of λ in determining the learning accuracy? Recall that λ is the quadratic adjustment cost that is irrelevant to agents' information, and hence the higher it is, the more weight agents put on it and less on private information. For example, as mentioned, if we set $\lambda = 0$ (the constant return case), agents only use their private information and complete learning achieves after one period.

Also, since $\tau_2 < \tau_\varepsilon$ and $\beta_2 < \beta_\varepsilon$, informativeness of signals can increase at most τ_ε or β_ε regardless of whether agents are fully informed or not, and such maximum increments happen only when agents solely respond to their private information (and the quadratic cost). In this sense learning in our model is slow. Yet this slowness is also a success for the learning mechanism since as is discussed in Introduction, if information is reflected too fast, the paradox of efficient informational markets arises (Grossman and Stiglitz (1976)).

2.3.2 Period *t*

With the understanding of the equilibrium in period 1 and the learning process in period 2, we are ready to solve the dynamic game. In period $t = 2,3,\cdots,T$, let $\hat{x}_{ti}(x_{1i},x_{2i},\cdots,x_{ti})$ denote agent i's private information about θ, i.e., her private estimation of θ that is only known to herself, from observing $x_{1i},x_{2i},\cdots,x_{ti}$, and let $\hat{p}_{t-1}(p_0,p_1,\cdots,p_{t-1})$ denote the public information about θ from observing $p_0,p_1,\cdots,p_{t-1}$; we respectively shorthand them by $\hat{x}_{ti}$ and $\hat{p}_{t-1}$. We state our results of the equilibrium in two parts. The first part concerning (i) and (ii) in the proposition describes the updating rules of the learning processes, and the second part concerning (iii) dictates the equilibrium strategies and prices.

Proposition 8. *A unique equilibrium of the dynamic game exists such that the following three statements hold.*
(i) For any period $t \in \{2,3,\cdots,T\}$ and information set $\{x_{ki},p_{k-1}\}_{k=1}^{t}$, there exists a set $\{x'_{ki},p'_{k-1}\}_{k=1}^{t}$ that shares the same information contents about θ, where $x'_{1i} = x_{1i}$, $p'_0 = p_0$, and for other k,

$$x'_{ki} = \frac{1}{m_{k-1}}(x_{ki} - n_{k-1}\hat{p}_{k-1}), \quad p'_k = \frac{1}{m_k}(p_k - n_k\hat{p}_{k-1}),$$

with m_k, n_k are coefficient defined in (iii). Moreover, x'_{ki} and p'_t are both centered around θ in the sense that they can be expressed in the following forms:

$$x'_{ti} = \theta + \frac{1}{\sqrt{\tau_t}}\varepsilon_{ti}, \quad p'_t = \theta + \frac{1}{\sqrt{\beta_t}}\varepsilon_t,$$

where $\tau_t = m_{t-1}^2\tau_\varepsilon$, and $\beta_t = m_t^2\beta_\varepsilon$.

(ii) For any period $t \in \{1,2,\cdots,T\}$, agent i's private and public information about θ,

namely $\hat{x}_{ti}$ *and* $\hat{p}_{t-1}$*, can be expressed in the forms of* $\hat{x}_{1i} = x_{1i}$*,* $\hat{p}_0 = p_0$*, and for other t,*

$$\hat{x}_{ti} = \theta + \frac{1}{\sqrt{\hat{\tau}_t}}\varepsilon_{ti}, \quad \hat{p}_t = \theta + \frac{1}{\sqrt{\hat{\beta}_t}}\varepsilon_t,$$

where $\hat{\tau}_t = \sum_{k=1}^{t} \tau_k$ *and* $\hat{\beta}_t = \sum_{k=0}^{t} \beta_k$*.*

(iii) In any period $t \in \{1, 2, \cdots, T\}$*, equilibrium price* p_t *is a linear function of* θ*,* $\hat{p}_{t-1}$*, and* ε_t *such that*

$$p_t(\theta, \hat{p}_{t-1}) = m_t\theta + n_t\hat{p}_{t-1} + \frac{1}{\sqrt{\beta_\varepsilon}}\varepsilon_t,$$

where

$$m_t = \frac{\delta_t}{\lambda + \delta_t}, \quad n_t = \frac{\lambda}{1+\lambda}\frac{1-\delta_t}{\lambda + \delta_t}, \quad \textit{with} \quad \delta_t = \frac{\hat{\tau}_t}{\hat{\tau}_t + \hat{\beta}_{t-1}}.$$

And the unique optimal strategy $a_t^*(\hat{x}_{ti}, \hat{p}_{t-1})$ *in period t is a linear function of* $\hat{x}_{ti}$ *and* $\hat{p}_{t-1}$*, such that*

$$\begin{aligned} a_t^*(\hat{x}_{ti}, \hat{p}_{t-1}) &= \frac{1}{\lambda}\{(1-m_t)\delta_t\hat{x}_{ti} + [(1-m_t)(1-\delta_t) - n_t]\hat{p}_{t-1}\} \\ &= m_t\hat{x}_{ti} + n_t\hat{p}_{t-1}. \end{aligned}$$

The proof is by induction and is similarly as we solve for period 1. For example in period 2, note that the maximization problem of agent i is the same as in period 1, and θ is still normally distributed given her updated information $\hat{x}_{2i}$ and $\hat{p}_1$ (see their definitions in (2.9) and (2.10)), so the information structures are also the same between the two periods. What changes is just the precision levels of an agent's information. Therefore, an equilibrium in period 2 can be solved in a similar way as solving period 1 (i.e., posit a linear price function and solve for its corresponding equilibrium) and thereby in period 3, learning from period 2 is also similar as learning in period 2 from period 1. Hence we can solve period 3, and so on. The equilibria for each period together constitute the equilibrium of the dynamic game. Due to the uniqueness in equilibria in every period and the induction argument, the way agents update their information about θ upon observing actions and prices is also symmetric and unique in equilibrium, establishing the general uniqueness of the game.

Note that all agents are rational, so that they make decisions in every period taken as given that agents in all previous stages act optimally. Also they understand the impact of their cumulative actions on the game, so the transformed endogenous signals they infer are the equilibrium ones.

Proof. Consider an arbitrary period $t \in \{2,3,\cdots,T-1\}$ and an agent i with information set $\{x_{ki}, p_{k-1}\}_{k=1}^{t}$. By Section 3.1, the statements (i) and (ii) that concern the learning process hold for period 2, and (iii) holds for period 1. Assume inductively that they hold for all periods till t, t included. We in below prove that those statements are valid for period $t+1$, so that the proposition is established since t is arbitrary.

For (i), since $x_{ti} = A_{t-1} + \varepsilon_{ti}/\sqrt{\tau_\varepsilon}$ and in equilibrium, $A_{t-1}(\theta, \hat{p}_{t-1}) = m_{t-1}\theta + n_{t-1}\hat{p}_{t-1}$ by the induction assumption, we have

$$\frac{1}{m_{t-1}}(x_{ti} - n_{t-1}\hat{p}_{t-1}) = \theta + \frac{1}{m_{t-1}\sqrt{\tau_\varepsilon}}\varepsilon_{ti}.$$

Thus it suffices to define

$$x'_{ti} = \frac{1}{m_{t-1}}(x_{ti} - n_{t-1}\hat{p}_{t-1}), \quad \text{and } \tau_t = m_{t-1}^2\tau_\varepsilon.$$

Similarly, since $p_t(\theta, \hat{p}_{t-1}) = m_t\theta + n_t\hat{p}_{t-1} + \varepsilon_t/\sqrt{\beta_\varepsilon}$, let

$$p'_t = \frac{1}{m_t}(p_t - n_t\hat{p}_{t-1}), \quad \text{and } \beta_t = m_t^2\beta_\varepsilon.$$

Then $\{x'_{ki}, p'_{k-1}\}_{k=1}^{t}$ is informationally equivalent to $\{x_{ki}, p_{k-1}\}_{k=1}^{t}$ with respect to θ. Consequently (ii) holds by Bayes' rule if we set

$$\hat{x}_{ti} = \frac{\tau_1 x'_{1i} + \tau_2 x'_{2i} + \cdots + \tau_t x'_{ti}}{\tau_1 + \tau_2 + \cdots + \tau_t}, \quad \hat{p}_{t-1} = \frac{\beta_0 p'_0 + \beta_1 p'_1 + \cdots + \beta_{t-1} p'_{t-1}}{\beta_0 + \beta_1 + \cdots + \beta_{t-1}},$$

and $\hat{\tau}_t = \sum_{k=1}^{t} \tau_k$ and $\hat{\tau}_t = \sum_{k=1}^{t} \beta_k$. Noteworthy, the above argument also proves the uniqueness in the updating rule of agents' information, when agents follow the equilibrium strategy in all earlier periods.

We now establish (iii) in three steps. First, suppose the price function p_t is linear as stipulated, i.e.,

$$p_t(\theta, \hat{p}_{t-1}) = m_t \theta + n_t \hat{p}_{t-1} + \frac{\varepsilon_t}{\sqrt{\beta_\varepsilon}},$$

where m_t, n_t are coefficients to be determined.

Secondly, recall that the maximization problem of agent i in period t is

$$\max_{a_{ti} \in \mathbb{R}} E[(\theta - p_t) a_{ti} - \frac{\lambda}{2} a_{ti}^2 \mid (\hat{x}_{ti}, \hat{p}_{t-1})].$$

Substitute the linear p_t into it and solve by the first order condition; the unique maximizer $a_t^*(\hat{x}_{ti}, \hat{p}_{t-1})$ is then such that

$$\begin{aligned} a_t^*(\hat{x}_{ti}, \hat{p}_{t-1}) &= \frac{1}{\lambda}\left((1 - m_t) E[\theta \mid (\hat{x}_{ti}, \hat{p}_{t-1})] - n_t \hat{p}_{t-1}\right) \\ &= \frac{1}{\lambda}\left\{(1 - m_t)\delta_t \hat{x}_{ti} + [(1 - m_t)(1 - \delta_t) - n_t]\hat{p}_{t-1}\right\}, \end{aligned}$$

where $\delta_t = \hat{\tau}_t / (\hat{\tau}_t + \hat{\beta}_t)$. Therefore the aggregate action A_t is also a linear function of θ and $\hat{p}_{t-1}$:

$$A_t(\theta, \hat{p}_{t-1}) = \frac{1}{\lambda}\left((1 - m_t)\delta_t \theta + [(1 - m_t)(1 - \delta_t) - n_t]\hat{p}_{t-1}\right).$$

Thirdly, by the market clearing condition $p_t = A_t + \varepsilon_t/\sqrt{\beta_t}$, plug A_t as in the second step in and compare the coefficients; we have

$$m_t = \frac{\delta_t}{\lambda + \delta_t}, \quad n_t = \frac{\lambda}{1+\lambda}\frac{1-\delta_t}{\lambda + \delta_t}.$$

Due to the concavity of the optimization problem and the induction argument, the obtained equilibrium is the unique equilibrium in the game.

Q.E.D.

Remark 4. *Note that A_t is a transformation from p_t minus the demand shock. Therefore, the signal x_{ti} can be written as*

$$x_{ti} = p_t - \varepsilon_t/\sqrt{\beta_t} + \varepsilon_{ti}/\sqrt{\tau_t},$$

even off the equilibrium path. This reformulation leads to a new interpretation of endogenous signal x_{ti}. Now x_{ti} can represent the different timings of agents checking the market price within each period. To be clear, though price in literature is often regarded as a public signal, it, as when represented by x_{ti}, bears private features since agents observe it with idiosyncratic noise because of the fine timings in checking it (meanwhile, ε_t explains common components like noise traders which affect the market, or simply the measurement error of a government agency who complies the data.). This interpretation conveys the feature of modern markets in which prices are constantly changing.

2.4. Evolvements of Precisions

We now characterize the asymptotic feature of the learning processes, which is decomposed into two related properties:

(i) Both private and public information fully reveals θ in the limit.

(ii) Both types of information converge at the same rate, and the rate is linear.

The first feature states that agents will learn the true state if they interact long enough, and both $\hat{x}_{Ti}$ and $\hat{p}_T$ becomes limit accurate for large T. The similar conclusion is obtained in the (static) rational expectations model of Grossman and Stiglitz (1980) who demonstrate that the endogenous public signal, price, becomes fully informative, when agents are fully informed exogenously or the noise is vanishingly small. Our approach is dynamic and thus differs from them in the learning process, and additionally we can track the evolvements of agents' information precisions (i.e., due to closed-form $\hat{\tau}_t$ and $\hat{\beta}_t$, we know the changes of information precision between periods).

The second property deals with the rate of learning. It has been proved by Vives (1993) that the rate is $T^{1/3}$ if private learning is absent (e.g., $\tau_\varepsilon \to 0$ so that agents completely ignore her private endogenous signals x_{ti}). The rate is slower when there is only a public learning channel, because public signals becomes increasingly informative over time and agents consequently respond less to their private information. Hence the price system, which is determined by agents' actions, reflects less private information. We demonstrate that private learning can accelerate learning, and that the rate increases to T in the limit, which is the same as when agents receive exogenous i.i.d. signals centered around θ.

Now we prove the above arguments. Denoted by $\Delta\hat{\tau}_t \equiv \hat{\tau}_{t+1} - \hat{\tau}_t$ the change in the informativeness of private information from time t to time $t+1$, and analogously define $\Delta\hat{\beta}_t \equiv \hat{\beta}_{t+1} - \hat{\beta}_t$. Recall that by construction, $x'_{(t+1)i}$ and p'_t respectively represent the new private and public information that agent i learns in period $t+1$ about θ, so the changes in informativeness of each type of information in period $t+1$ is the precision of $x'_{(t+1)i}$ and p'_t, i.e., $\Delta\hat{\tau}_t = \tau_{t+1}$ and $\Delta\hat{\beta}_t = \beta_t$. In turn, we have the following proportion. Now we prove the above arguments. Denoted by $\Delta\hat{\tau}_t \equiv \hat{\tau}_{t+1} - \hat{\tau}_t$ the change in the informativeness of private information from time t to time $t+1$, and analogously define $\Delta\hat{\beta}_t \equiv \hat{\beta}_{t+1} - \hat{\beta}_t$. Recall that by construction, $x'_{(t+1)i}$ and p'_t respectively represent the new private and public information that agent i learns in period $t+1$ about θ, so the changes in informativeness of each type of information in period $t+1$ is the precision of $x'_{(t+1)i}$ and p'_t, i.e., $\Delta\hat{\tau}_t = \tau_{t+1}$ and $\Delta\hat{\beta}_t = \beta_t$. In turn, we have the following proportion.

Proposition 9. *(i) For all $t = 1, 2, \cdots T-1$, $\hat{\beta}_t$ and $\hat{\tau}_t$ increase at the same rate in the following sense:*

$$\hat{\beta}_t = \alpha + \frac{\beta_\varepsilon}{\tau_\varepsilon}\hat{\tau}_{t+1},$$

for some constant α.

(ii) $\hat{\tau}_T \to \infty$ and $\hat{\beta}_T \to \infty$, as $T \to \infty$.

(iii) The rate of convergence of both public and private signals is linear:

$$\hat{\tau}_T/T \to (\frac{1}{\lambda + 1 + \lambda\beta_\varepsilon/\tau_\varepsilon})^2\tau_\varepsilon, \quad \hat{\beta}_T/T \to (\frac{1}{\lambda + 1 + \lambda\beta_\varepsilon/\tau_\varepsilon})^2\beta_\varepsilon,$$

as $T \to \infty$, when $\beta_\varepsilon > 0, \tau_\varepsilon > 0$.

Proof. For (i), recall by (i) of Proposition 1 that $\tau_{t+1} = m_t^2\tau_\varepsilon$ and $\beta_t = m_t^2\beta_\varepsilon$, so the relative increment of precision levels is constant in the sense that

$$\frac{\Delta\hat{\beta}_t}{\Delta\hat{\tau}_t} = \frac{\beta_t}{\tau_{t+1}} = \frac{\beta_\varepsilon}{\tau_\varepsilon}.$$

We thus have $\hat{\beta}_t = \hat{\tau}_{t+1}\beta_t/\varepsilon_t$ for all $t = 1, 2, \cdots, T-1$. Therefore,

$$\hat{\beta}_t - \frac{\beta_\varepsilon}{\tau_\varepsilon}\hat{\tau}_{t+1} = \beta_0 - \frac{\beta_\varepsilon}{\tau_\varepsilon}\tau_1 \equiv \alpha.$$

And $\hat{\tau}_{t+1} = \hat{\tau}_t + \tau_{t+1}$, where τ_{t+1} is finite and hence does not affect rates.

For (ii), note that if the precision of one type of information goes to ∞, so does the other type because of (i). Now suppose by contradiction that none goes to ∞; then $\Delta\hat{\tau}_t = \tau_{t+1} = m_t^2\tau_\varepsilon > 0$ for any t since now $m_t = \delta_t/(\lambda+\delta_t) > 0$. In consequence, $\hat{\tau}_T = \sum_{t=1}^{T}\Delta\hat{\tau}_t + \tau_1$ diverges, causing a contradiction.

For (iii), the proof is by inspecting $\hat{\tau}_t$, which equals $\sum_{k=1}^{t-1}\Delta\hat{\tau}_k + \tau_1$, and we show that $\hat{\tau}_T/T$ converges to the stated constant. In below, we characterize the rate at which $\hat{\tau}_T \to \infty$, and the rate for $\hat{\beta}_T \to \infty$ equals that for $[\alpha + (\hat{\tau}_T\beta_\varepsilon)/\tau_\varepsilon] \to \infty$. For any $t = 1, \cdots, T$,

$$\begin{aligned}
\hat{\tau}_t = \sum_{k=1}^{t-1}\Delta\hat{\tau}_k + \tau_1 = \tau_\varepsilon\sum_{k=1}^{t-1}m_k^2 + \tau_1 &= \tau_\varepsilon\sum_{k=1}^{t-1}\left(\frac{\delta_k}{\lambda+\delta_k}\right)^2 + \tau_1 \\
&= \tau_\varepsilon\sum_{k=1}^{t-1}\left(\frac{\hat{\tau}_k}{(\lambda+1)\hat{\tau}_k + \lambda\hat{\beta}_k}\right)^2 + \tau_1 \quad \left(\because\ \delta_k = \frac{\hat{\tau}_k}{\hat{\tau}_k + \hat{\beta}_k}\right) \\
&= \tau_\varepsilon\sum_{k=1}^{t-1}\left(\frac{\hat{\tau}_k}{(\lambda+1)\hat{\tau}_k + \lambda(\alpha + \beta_\varepsilon\hat{\tau}_{k+1}/\tau_\varepsilon)}\right)^2 + \tau_1
\end{aligned}$$

Then by (ii), there exists $T_\varepsilon \in \mathbb{N}$ such that for all $t \geqslant T_\varepsilon$, $\hat{\tau}_t \to \infty$. Therefore, divide $\hat{\tau}_T$ in the above form by T, and let $T \to \infty$; we obtain

$$\begin{aligned}
\hat{\tau}_T/T = &\underbrace{\tau_\varepsilon\sum_{k=1}^{T_\varepsilon-1}\left(\frac{\hat{\tau}_k}{(\lambda+1)\hat{\tau}_k + \lambda\hat{\beta}_k}\right)^2/T}_{\to 0} + \tau_\varepsilon\sum_{k=T_\varepsilon}^{T-1}\left(\frac{\hat{\tau}_k}{(\lambda+1)\hat{\tau}_k + \lambda(\alpha+\beta_\varepsilon\hat{\tau}_{k+1}/\tau_\varepsilon)}\right)^2/T + \underbrace{\tau_1/T}_{\to 0} \\
&\to \left(\frac{1}{\lambda+1+\lambda\beta_\varepsilon/\tau_\varepsilon}\right)^2\tau_\varepsilon,
\end{aligned}$$

which is a non-zero constant, so $\hat{\tau}_T/T \to \infty$ at a linear rate. *Q.E.D.*

This statements (i) and (ii) confirm and extend the features of the learning processes discussed in Section 3 to multiple periods. In detail, the increase in informativeness in each period is bounded by τ_ε and β_ε, verifying that learning is slow. Also, the increment of informativeness of agents' information about θ, namely $\Delta\tau_t$ and $\Delta\beta_t$, is positively correlated with agents' private information and negatively with their public information. Thirdly, it states that each type of information becomes arbitrarily accurate after enough rounds of interaction, indicating that all agents' private information can be fully elicited by the market force.

We present a heuristic informal proof for (iii) here. Recall that, for any t, $\Delta\hat{\tau}_t = (\delta_k/(\lambda+\delta_k))^2\tau_\varepsilon$, and $\delta_k = \hat{\tau}_k/(\hat{\tau}_k+\hat{\beta}_k)$, and $\hat{\beta}_t = \alpha + \hat{\tau}_{t+1}\hat{\beta}_\varepsilon/\hat{\tau}_\varepsilon$, so that

$$\Delta\hat{\tau}_T = \left(\frac{\hat{\tau}_k}{(\lambda+1)\hat{\tau}_k+\lambda(\alpha+\beta_\varepsilon\hat{\tau}_{k+1}/\tau_\varepsilon)}\right)^2\tau_\varepsilon \to \left(\frac{1}{\lambda+1+\lambda\beta_\varepsilon/\tau_\varepsilon}\right)^2\tau_\varepsilon, \quad as \quad T\to\infty,$$

since $\hat{\tau}_T \to \infty$. Therefore, it is as if agents receive i.i.d. signals centered around θ with precision $(\lambda+1+\lambda\beta_\varepsilon/\tau_\varepsilon)^{-2}(\tau_\varepsilon+\beta_\varepsilon)$ in the limit in each period, so the asymptotic learning rate is expected to be the same as observing i.i.d. signals and thus is linear. The defining factor for the rate jumping back from $t^{1/3}$ in pure public learning case to linearity is that both $\hat{\tau}_t$ and $\hat{\beta}_t$ converge to infinity at the same rate. Consequently, contrasting with the sole public learning case in which $\hat{\tau}_T = \tau_1$ while $\hat{\beta}_T \to \infty$ (so that $\hat{\tau}_T/\hat{\beta}_T \to 0$ as $T\to\infty$), private information in our setting will not be crowded out.

It is noteworthy that the linear rate is guaranteed as long as private learning exists ($\tau_\varepsilon > 0$), though a smaller τ_ε decreases $\Delta\tau_T$ and $\Delta\beta_T$ and thus leads to a lower per-period increment in informativeness. Also, a higher τ_ε results in a higher asymptotic learning rate. Hence we say that the inclusion of a private learning channel improves the learning efficiency. Note that the initial precision level τ_1 impacts little on the learning efficiency, since τ_1 does not affect $\Delta\hat{\tau}_t$ or $\Delta\hat{\beta}_t$, nor the asymptotic learning rate (see Proposition 2. (iii)). The impact of τ_1 on the informativeness of agents' beliefs is completely characterized by its absolute value in this relationship: $\hat{\tau}_T = \sum_{t=1}^{T-1}\Delta\hat{\tau}_t + \tau_1$. Therefore, to improve the learning efficiency, the exogenous information precision is not of paramount importance, and one should think about ways to improve the observational accuracy of the endogenous signals.

Remark 5 (No Initial Public Information)**.** *Suppose agents in prior believe that θ is randomly drawn from the entire real line: $\theta \sim Unif(\mathbb{R})$ and that p_0 is unobservable, that is, agent i's information set is only $\{x_{1i}\}$ in period 1. Then learning accelerates since we can show that now p_1 is a function of θ and ε_1 only (not including p_0 as before). The proof is similar - start by positing a linear price function $p_1(\theta) = m'_1\theta + \varepsilon_1/\sqrt{\beta_1}$ for some coefficient m'_1, and solve agent i's optimization problem by the first order condition, which gives $m'_1 = 1/2$ in equilibrium. Note that p_1 now is centered around θ, and hence there is no need to transform it into some p'_1. As is discussed, the direct observation on θ is faster than learning from transformed signals; hence the learning efficiency improves. But the learning rate in the limit is still linear since in periods $t = 2, 3, \cdots, N$, agents' actions start to contain public information. One can envisage that from period 2, the structure of this new game is similar to the original game in Section 2.*

Chapter 3

Learning and Multiplicity in Global Games

3.1. Introduction

Coordination problems are prevalent in the economy, and though partially informed agents have aligned preferences in these games to coordinate on the same action, they often fail to do so because of payoff uncertainty about the market conditions as well as strategic uncertainty about their fellow agents' beliefs. Learning, whether through public observation of market price or private observation of nearby agents' moves, in principle mitigates both kinds of uncertainties and thus mitigates coordination failures. This article studies the impact of learning on a dynamic coordination game - agents interact and learn repeatedly from both public and private observation of past actions, and then they participate in a global coordination game.

Specifically, we construct the coordination game as a global game of Carlsson and Van Damme (1993) and Morris and Shin (2004). The advantage of the global game approach is its selection of a unique equilibrium, when agents' information is dispersed enough. As a result, it paves a safe way to conduct comparative statics analysis of the impact of learning on coordination incidence and welfare. Later studies such as Angeletos and Werning (2006), nonetheless, cast doubts on the uniqueness, remarking the rise of multiple equilibria when endogenous public information is factored in, even when agents' information is most diffused (namely, when agents are endowed with limit accurate private signals). Their work, among other existing studies, overlooks the role of private endogenous information, however. While private learning is a significant component: practically agents extensively learn through local observation and private talks, and theoretically uniqueness in global games is obtained by introducing private information to agents so as to form informational heterogeneity among them. Therefore, we are interested in, when private learning happens in tandem with public learning, whether the unique equilibrium can be established.

At the same time, though global games are initially treated as an equilibrium selection device as described above, it has gained popularity recently for modeling realistic coordination problems and delivering robust predictions, since it sustains the unique equilibrium. Therefore, as discussed, we will also investigate practical issues such as the learning's impact on agents' welfare.

In detail, we consider that a continuum of agents interact for T periods, in which the learning stage, based on Vives (1993, 1997), consists of the first $T-1$ periods, and the coordination stage, modeled as a global game of regime change, happens in the Tth period. In each of the first $T-1$ periods, every agent takes an action to minimize the quadratic distance to an unknown, payoff-relevant state parameter θ. Agents have private, partial information about θ initially, and can further privately and publicly observe the realized aggregate action of the previous period. After T-1 periods, they proceed to the coordination stage in which each of them, instead of minimizing the distance, decides whether to invest in a risky project which has the regime change feature (namely, it either succeeds, if the mass of agents investing exceeds some threshold depending on θ, or not). Then the game ends and the payoffs are realized.

Note that an implication of the model is the abstraction of the direct payoff linkage between periods by assuming agents' action space is independent over time. This is for a clean analysis of learning's effects on the equilibrium selection and is also standard in literature such that Angeletos and Werning (2006). Therefore, the two stages are only connected through the information linkage, because they share the same state θ and thus past actions reveal valuable information in inferring θ.

We now summarize our findings, which are two-folded. Theoretically, as for whether the co-existence of public and private learning relieves multiplicity, the answer is negative - multiplicity always arises after enough rounds of learning. Furthermore and notably, private learning facilitates multiplicity in some situations because it accelerates the learning rate of public signals. Consequently, agents quickly learn common information, and thus multiplicity appears sooner, compared to the sole public learning case.

Practically, we take advantage of the existence of the unique equilibrium to present safe comparative statics results. We show that, to the extent that uniqueness holds, the presence of learning improves agents' expected payoffs from the coordination game. The result is hardly surprising, because with endogenous signals, agents estimate θ more accurately and they know their opponents also better guess θ; hence both payoff uncertainty and strategic uncertainty are alleviated, so their payoffs increase.

Furthermore, since learning happens through both public and private channels, we attain several novel and intuitive features of the learning process. Primarily, we demonstrate that both types of endogenous information converge to limit accuracy about θ at a linear rate when agents interact long enough.[1] The conclusion holds as long as private learning is available, and it should be compared with Vives (1993) who demonstrates the convergence rate of endogenous public information is $t^{1/3}$ (t is the periods of learning) with only public learning available. Our finding thus stresses that private learning facilitates overall learning efficiency in the sense of a higher convergence rate. The result is conceivable because, as implied by the herding literature, social learning becomes slow with only public observation (for agents put increasing weights on public information along the time, lowering the informativeness of endogenous signals), while private learning guarantees that private information will not be crowded out en route. Noteworthy, since we show that learning fully reveals the state in the limit, our model justifies the usual assumption in the global games literature in which agents are assumed to be endowed with almost accurate exogenous information.

Related Literature This paper relates to two strands of literature: global games and learning. We describe them in order. The global games model is initiated by Carlsson and Van Damme (1993); see also Morris and Shin (2003) for a comprehensive survey. Both of them show uniqueness holds for sufficiently dispersed information among agents. Morris and Shin (2004) explicitly characterizes the necessary and sufficient condition, as a ratio of informativeness of public information to the square root of that of private information, that determines uniqueness. All information in these studies is exogenously given, however. Later studies from various aspects examine whether endogenous public information leads to multiplicity, and among them, Angeletos and Werning (2006) who explore that agents learn from rational expectations equilibrium price prior to a global game, is most similar to our work. Several aspects are distinct, though. First, we consider multi-period learning, while they examine a one-spot rational expectations equilibrium model of Grossman and Stiglitz (1976). Therefore, our work circumvents the paradox of the impossibility of efficient information markets.

[1] The rate is said to be t^n if τ_t/t^n converges to a positive constant, where t is time and τ_t is time-t's information precision. A linear rate is when $n = 1$.

Secondly, we also scrutinize the role of private learning and show that it in some situations makes multiplicity easier to occur. Actually, endogenous private learning is overlooked in other existing literature as well; previous studies tend to vary precisions of private information exogenously. Our work is thus novel in analyzing the role of private endogenous signals. Thirdly, we investigate the impact of learning, which lacks in the framework of Angeletos and Werning (2006). This point is also examined by Dasgupta (2007) in an almost completely different environment, and Szkup (2020) gives a general analysis for static global games.

There are other researches that, though methodologically less relevant to this article, examine the relationship between endogenous public information and multiplicity. Angeletos et al. (2006), for example, document the occurrence of multiplicity when an omniscient mechanism designer signals to agents before their decision-making, and Hellwig et al. (2006) show multiplicity when the market force is factored in. Also, Angeletos et al. (2007) verify the rise of multiplicity in dynamic games in which agents face a new global game repeatedly. On the other hand, Szkup and Trevino (2015) confirms uniqueness for agents who costly acquire precision levels of signals prior to a global game. Notably, we obtain a similar result to the last one in the demonstration of the positive value of information, when we show learning improves welfare.

For learning, we base our analysis on Vives (1993, 1997), both of which restrict to public learning. We incorporate an analogous private learning channel into the models. For the exclusive learning literature that does not consider anything about global games, Amador and Weill (2012) examine the co-existence of both types of learning for exploring welfare implications. Their work, however, distinguishes from ours in that it is a continuous-time model that delivers results by explicitly solving stochastic differential equations. We, on the other hand, consider a discrete-time model and derive our conclusions by tracking the changes in agents' information in each period.

3.2. The Model: The Coordination Stage

Our model is based on Vives (1997) for the learning stage and Morris and Shin (2004) for the coordination stage. We consider that agents interact repeatedly and observe past aggregate activities publicly as well as privately prior to a global coordination game. To familiarize readers who are new to the global games approach, we first investigate the coordination stage with exogenous information, and then proceed to incorporate the learning stage into the game.

The economy consists of a measure-one continuum of agents $i \in [0,1]$ who independently decide whether to invest or not in a risky investment project. Let $a_i \in \{0,1\}$ denote agent i's action, with $a_i = 1$ indicates investing and $a_i = 0$ means not investing, and $A = \int_i a_i di$ the mass of agents who choose to invest. The payoff to an agent who does not invest is normalized to zero, and the payoff to an agent who invests depends on whether the investment project succeeds or not. The success of the investment depends on an unknown state parameter θ as well as on the number of agents who invest in the project. Specifically, The project succeeds if and only if the mass of agents who invest exceeds some threshold: $A_T \geqslant 1 - \kappa\theta$, where $\kappa > 0$ is a known constant that captures the dependence of success on the state; we for now set $\kappa \equiv 1$. Therefore, agent i who takes action a_i obtains payoff

$$\begin{cases} a_i(1-c), & \text{if } A \geqslant 1-\kappa\theta \\ -a_i c, & \text{otherwise,} \end{cases} \tag{3.1}$$

where $c \in (0,1)$ represents the cost and is a known constant, and $\kappa = 1$. Note that only the payoff difference between the two actions matters in an agent's decision-making, so the normalization of payoff to $a_i = 0$ is innocuous.

It is also worth remarking that an agent finds it dominant to invest if $\theta \geqslant 1$, because the project will succeed even without coordination, and on the opposite, one should never invest if $\theta < 0$. The interesting case is when $\theta \in (0,1)$ - it is optimal for an agent to invest if and only if enough of her fellow agents invest. In this sense, the game exhibits a coordination feature. Noteworthy that a higher θ means fewer investments are required for the project's success, so we say θ indicates the state of the economy or the project's prospect. Furthermore, each agent must infer θ as well as his opponents' actions when making decisions, suggesting that higher order inferences are involved in agents' decision-making, but as we will show, there exists a simple form of strategies that governs agents' behaviors in equilibrium.

The state parameter θ is uniformly distributed over the entire real line, so agents hold an improper prior about it: $\theta \sim \text{Unif}(\mathbb{R})$. At the beginning of the game, each agent i observes two signals, one private x_{1i} and one public p_1, about the realization of θ, such that

$$x_{1i} = \theta + \frac{1}{\sqrt{\tau_1}}\varepsilon_{1i}, \quad p_1 = \theta + \frac{1}{\sqrt{\beta_1}}\varepsilon_1, \tag{3.2}$$

where τ_1, β_1, measuring the precision of each signal, are positive constants, and $\varepsilon_{1i}, \varepsilon_1$ are standard normals ($\varepsilon_{1i}, \varepsilon_1 \sim N(0,1)$) independent of each other and all other parameters. [2] Throughout, the Law of Large Numbers (LLN) convention is imposed, so that the proportion of agents who receive signals higher than some number is equal to the probability of an individual agent receiving such signals. An immediate consequence is that the idiosyncratic noise in private signals cancels out in the population: $\int_i \varepsilon_{1i} di = 0$, and hence no aggregate uncertainty prevails in our model.

The Threshold Strategy and Equilibrium As in the literature, we restrict to agents playing a symmetric threshold strategy; denoted by $a(x_{1i}, p_1) \in \{0,1\}$ a symmetric strategy, it assigns

[2] Note that, observing the public signal p_1 is equivalent to assuming that $\theta \sim N(p_1, 1/\beta)$ in prior, but p_1 is then a constant, instead of a random variable, in the proper prior case. We write p_1 as a public signal for later notation simplicity.

a probability to investment from one's information set. A threshold strategy $a(x_{1i}, p_1)$ of agent i takes the following form; there exists a threshold number $x(p_1) \in \mathbb{R}$,

$$a(x_{1i}, p_1) = \begin{cases} 1, & \text{if } x_{1i} > x(p_1) \\ 0, & \text{otherwise,} \end{cases}$$

given any realizations of x_{1i} and p_1. A symmetric equilibrium is defined by a strategy profile (which consists of a symmetric strategy), such that the action to agent i prescribed by that strategy maximizes her expected payoff, given i's information and others agents also follow the strategy. An equilibrium in which agents play a threshold strategy is referred to as a monotone equilibrium; we sometimes call a monotone equilibrium simply an equilibrium. Note that in a monotone equilibrium, an agent's strategy is completely characterized by the threshold number $x(p_1)$; hence we in what follows focus on finding the equilibrium threshold and denote it by $x^*(p_1)$.

We now solve this static global game as in the following proposition.

Proposition 10. *A unique monotone equilibrium characterized by $x^*(p_1)$, for any realization p_1, exists, and the monotone equilibrium is the only equilibrium form if and only if $\beta_1/\sqrt{\tau_1} \leqslant \sqrt{2\pi}$.*

The proof is based on Morris and Shin (2004), and the idea is to show the existence of a pair $(x^*(p_1), \theta^*(p_1))$, such that (i) it is optimal for agent i to invest iff $x_{1i} > x^*(p_1)$ and that (ii) the project succeeds iff $\theta \geqslant \theta^*(p_1)$, given any realization of p_1. The first condition is also referred to as the payoff indifference condition and is characterized by the following equation:

$$P(\theta \geqslant \theta^*(p_1) | x^*(p_1)) = c,$$

and the second is called the critical mass condition which states

$$P(x_{1i} > x^*(p_1) \mid \theta^*(p_1)) = 1 - \theta^*(p_1).$$

Rewriting the payoff indifference condition, we obtain

$$x^*(p_1) = \frac{\tau_1 + \beta_1}{\tau_1}\theta^*(p_1) - \frac{\beta_1}{\tau_1}p_1 + \frac{\sqrt{\tau_1 + \beta_1}}{\tau_1}\Phi^{-1}(1-c).$$

Plugging $x_1^*(p_1)$ of this form into the critical mass condition gives an equation only about $\theta^*(p_1)$

$$\Phi\left(\frac{\beta_1}{\sqrt{\tau_1}}\theta^*(p_1) - \frac{\beta_1}{\sqrt{\tau_1}}p_1 + \frac{\sqrt{\tau_1 + \beta_1}}{\sqrt{\tau_1}}\Phi^{-1}(1-c)\right) - \theta^*(p_1) = 0.$$

There exists a solution $\theta^*(p_1)$ since as $\theta^*(p_1) \to -\infty$ (resp. ∞), the L.H.S converges to ∞ (resp. $-\infty$). It is also standard to check that the solution $\theta^*(p_1)$ is unique if $\beta_1/\sqrt{\tau_1} < \sqrt{2\pi}$.

The implication of the equilibrium is straightforward - the better the state, agents on average have higher posteriors about θ (and they know their opponents also think so) and thus invest more frequently, and in turn the higher chance of success coordination (i.e., θ exceeds $\theta^*(p_1)$ are states in which the investment succeeds). In addition, the proposition claims that if the information is highly public among agents, measured by the ratio of informativeness of the public signal to the square root of that of the private signals, multiple equilibria arise. This is so because, with complete information, there exist two equilibria (all investing versus none investing) at states $\theta \in (0,1)$, and as a result when public information becomes dominant, the game is expected to behave similarly as in complete information environments and thus has multiple equilibria.

Also, it becomes evident that the driving force of uniqueness in global games is the level of heterogeneity in agents' information and so, it is conceivable that public endogenous signals studied in the literature like Angeletos and Werning (2006) cause the uniqueness result to fail, and also that private learning might rebuild it. We however will verify that adding private endogenous information can sometimes prompt the rise of multiple equilibria.

3.3. The Two-period Model with a Learning Stage

We now add a learning stage before the coordination game. For illustration, we in this section consider that the learning stage only lasts one period, so the game is a two-period dynamic game. In the next section, we allow agents to interact and learn for multiple periods before entering the coordination game stage. The two-period setting is enough to deliver our result concerning the information's impact on coordination.

3.3.1 The Setup

The state of the economy is still characterized by θ towards which the unit of agents hold an improper prior, and in period 1, every agent i observes two signals x_{1i} and p_1 in the structures of (3.2) about the realized θ. However now, before entering the coordination stage, which happens in period 2, agents first have a learning stage in period 1 in a form of a pure prediction game of Vives (1997).

Specifically, in period 1, the objective of every agent i is to minimize the quadratic distance between her action to the state parameter θ, given her information set (x_{1i}, p_1). Let $a_{1i} \in \mathbb{R}$ denote her action in period 1 (we use subscripted a_{ti} to indicate an action in the learning stage and a_i for the coordination stage). That is, agent i after observing (x_{1i}, p_1) chooses $a_{1i} \in \mathbb{R}$ to minimize

$$E[(a_{1i} - \theta)^2 \mid x_{1i}, p_1]. \tag{3.3}$$

Let A_1 denote the aggregation action in period 1:

$$A_1 \equiv \int_i a_{1i} di. \tag{3.4}$$

After everyone makes a decision, A_1 is determined and the game proceeds to period 2, which essentially is the same as before, except that agents' will observe the realized A_1. In detail, at the beginning of period 2, agent i observes a private signal x_{2i} and a public signal p_2 about the aggregate activity in period 1, in the form of

$$x_{2i} = A_1 + \frac{1}{\sqrt{\tau_\varepsilon}} \varepsilon_{2i}, \qquad p_2 = A_1 + \frac{1}{\sqrt{\beta_\varepsilon}} \varepsilon_2, \tag{3.5}$$

where ε_{2i} and ε_2 are both standard normals, independent of all other parameters and represent the individually specifical noise and the market wise noise, respectively; τ_ε and β_ε are known positive constants that capture observational errors. Therefore, the information set for agent i in period 2 is $\{x_{1i}, x_{2i}, p_1, p_2\}$. Agent i decides whether to invest in the investment project given this information set, the payoff of the project is still given by (3.1) (here, we still denote a_i and $A = \int_i a_i di$ the i's action and the aggregate action in the coordination stage). At the end of period 2, the state is revealed and agents get paid for their payoffs from both period 1 and period 2. The game then ends.

We still focus on agents playing a symmetric strategy in each period, and a strategy of an agent in period $t = 1, 2$ is a function that maps from her information set $\{x_{ki}, p_k\}_{k=1}^{t}$ to an action out of $\mathbb{R}$ in period 1, or out of $\{0, 1\}$ in period 2. Note that the game is essentially static since there is no payoff linkage between the two periods, because the action space is $\mathbb{R}$ in period 1 and all agents are infinitesimal. That said, the optimal myopic behavior is the optimal behavior of the game. Still, we consider that agents play a threshold strategy in period 2. An equilibrium of the two-period game consists of equilibria of each stage game, and agents in period 2 update their information by Bayes' rule.

3.3.2 Solving the Learning Stage

We now solve for the equilibrium in period 1. It is easy to see that there exists a unique optimal strategy for agent i with information (x_{1i}, p_1) to maximize (3.3), which is given by

$$a_1^*(x_{1i}, p_1) = E[\theta \mid x_{1i}, p_1] = \delta_1 x_{1i} + (1-\delta_1)p_1.$$

where $\delta_1 = \tau_1/(\tau_1+\beta_1)$ is the common weight on one's private information. Consequently $A_1(\theta, p_1)$ is a function of θ and p_1 such that

$$A_1(\theta, p_1) = \delta_1\theta + (1-\delta_1)p_1. \tag{3.6}$$

Knowing the form of A_1 in equilibrium, we are ready to characterize the endogenous signals upon observing A_1 in period 2. Note that agents rationally understand the information impact of their aggregate actions on the game, so the structures of endogenous signals are on the equilibrium path, that is, A_1 equals $A_1(\theta, p_1)$ in the form of (3.6). Substituting $A_1(\theta, p_1)$ into $x_{2i} = A_1 + \varepsilon_{2i}/\sqrt{\tau_\varepsilon}$ gives

$$x_{2i} = \delta_1\theta + (1-\delta_1)p_1 + \frac{1}{\sqrt{\tau_\varepsilon}}\varepsilon_{2i}.$$

Rearranging it gives

$$\frac{1}{\delta_1}(x_{2i} - (1-\delta_1)p_1) = \theta + \frac{1}{\delta_1\sqrt{\tau_\varepsilon}}\varepsilon_{2i}.$$

Therefore, if we define

$$x'_{2i} = \frac{1}{\delta_1}(x_{2i} - (1-\delta_1)p_1),$$

we have an informationally equivalent signal x'_{2i} to x_{2i} with respect to θ, such that

$$x'_{2i} = \theta + \frac{1}{\sqrt{\tau_2}}\varepsilon_{2i},$$

where $\tau_2 = \delta_1^2\tau_\varepsilon$. [3] Likewise, by plugging $A_1(\theta, p_1)$ into $p_2 = A_1 + \varepsilon_2/\sqrt{\beta_\varepsilon}$ and rearranging, we obtain an informationally equivalent statistic p'_2 to p_2 in regard to θ, such that

$$p'_2 \equiv \frac{1}{\delta_1}(p_2 - (1-\delta_1)p_1) = \theta + \frac{1}{\sqrt{\beta_2}}\varepsilon_2,$$

where $\beta_2 = \delta_1^2\beta_\varepsilon$. Note that x'_{2i} and p'_2 are centered around θ, and thus θ is normally distributed given each of them. Therefore we can summarize every agent i's private information (that is only known to agent i), denoted by $\hat{x}_{2i}(x_{1i}, x_{2i})$, from observing x_{1i}, x_{2i}, and the public information $\hat{p}_2(p_1, p_2)$ from observing p_0, p_1, in period 2. By Bayes' rule, they can be expressed by

$$\hat{x}_{2i}(x_{1i}, x_{2i}) = \hat{x}_{2i}(x_{1i}, x'_{2i}) = \frac{\tau_1 x_{1i} + \tau_2 x'_{2i}}{\tau_1 + \tau_2}, \quad \hat{p}_2(p_1, p_2) = \frac{\beta_1 p_1 + \beta_2 p'_2}{\beta_1 + \beta_2}.$$

Note that due to gaussian updating, θ is normally distributed given $\hat{x}_{2i}$ and $\hat{p}_2$ such that

$$\hat{x}_{2i}(x_{1i}, x_{2i}) = \theta + \frac{1}{\sqrt{\hat{\tau}_2}}\varepsilon_{2i}, \quad \hat{p}_2 = \theta + \frac{1}{\sqrt{\hat{\beta}_2}}\varepsilon_2,$$

[3] Note that with little abusing notation, we still let ε_{2i}, and later ε_2, be standard normals that may have different realizations to those in x_{2i} and p_2.

where $\hat{\tau}_2 = \tau_1 + \tau_2$ and $\hat{\beta}_2 = \beta_1 + \beta_2$. It is worth stressing the above characterization is unique in equilibrium due to unique a_1^* and $A_1(\theta, p_1)$, indicating that agents update their beliefs about θ upon observing A_1 in a symmetric and unique way. This result is key to establishing the uniqueness in equilibria of the whole dynamic game. In detail, the optimal strategy of an agent, given her information, is unique; therefore, once we show the updating rules of agents from learning are unique and symmetric, the equilibrium of the whole game is unique.

3.3.3 Solving the Coordination Stage

Note that the information set for agent i in period 2 is thus $\{\hat{x}_{2i}, \hat{p}_2\}$, which share the same structure as their information set $\{x_{1i}, p_1\}$ in period 1 (i.e., an additive structure centered around θ with a Gaussian noise), Therefore, we can readily solve the coordination stage as similar to that in Proposition 1. That is, we characterize a pair of thresholds $(x^*(\hat{p}_2), \theta^*(\hat{p}_2))$, for any realization of $\hat{p}_2$ by the following two equations:

$$P(\theta \geqslant \theta^* \mid x^*(\hat{p}_2)) = c, \text{ and } P(\hat{x}_{2i} > x^*(p_1) \mid \theta^*(\hat{p}_2)) = 1 - \theta^*(\hat{p}_2).$$

The first equation is the payoff indifference condition and the second is the critical mass condition. Solve them simultaneously as we do in Section 2 (i.e., obtain the expression of $x^*(\hat{p}_2)$ in terms of $\theta^*(\hat{p}_2)$ from the first equation, substitute it into the second to get a single equation that is only about $\theta^*(\hat{p}_2)$); we can conclude the existence of the pair, and furthermore obtain that the pair is unique if and only if $\hat{\beta}_2/\sqrt{\hat{\tau}_2} < 2\sqrt{\pi}$.

Therefore, we know the uniqueness in global games is indeed determined by the ratio of public informativeness to the square root of private informativeness, and since

$$\frac{\hat{\beta}_2}{\sqrt{\hat{\tau}_2}} < \frac{\hat{\beta}_2}{\sqrt{\tau_1}},$$

the introduction of a private learning channel, or the increasing precision in private signals, indeed increases the likelihood of a unique equilibrium in the global game. This is an intuitive result that, however, does not hold in a general setting when learning lasting multiple periods, as we will show later. The main reason is that though private learning initially disperses information among agents, it in the long run improves the learning efficiency with which agents learn from public signals, so that information conformity is actually easier to be established among agents in the presence of a private learning channel.

Another impact of learning is certainly on coordination behavior and outcomes, as is easy to notice that the equilibrium threshold pair has changed (from the one in Section 2). We now demonstrate that learning improves coordination success, as long as the equilibrium is unique, by varying the efficiency of the learning processes. Intuitively, with a higher τ_ε, agents' signals are better aligned with each other as well as with the state; hence both strategic and payoff uncertainty is alleviated so that they make more accurate decisions on average and expect higher payoffs. In below, let $D(x_{1i}, p_1)$ be the payoff difference between investing to not investing in the coordination stage, from the perspective of agent i with (x_{1i}, p_1) in period 1. That is,

$$D(x_{1i}, p_1) = P\left(\hat{x}_{2i} > x^*, \theta \geqslant \theta^* | x_{1i}, p_1\right)(1-c) - P(\hat{x}_{2i} > x^*, \theta < \theta^* | x_{1i}, p_1)c.$$

When there is no learning, then $\hat{x}_{2i} = x_{1i}$ so that this is the expected payoff of agent i following the threshold strategy. Consequently, if a higher τ_ε increases this value (i.e., $\partial D(x_{1i}, p_1)/\partial \tau_\varepsilon > 0$), we establish that learning improves agents' welfare.

Proposition 11. *The higher τ_ε, the higher expected payoffs from the coordination stage.*

Proof. Since τ_1 is given, it is equivalent to examining the impact of $\hat{\tau}_2$ on $D(x_{1i}, p_1)$. We compute that

$$\frac{\partial}{\partial \hat{\tau}_2} D(x_{1i}, p_1) > 0.$$

The detailed computation is in Appendix. *Q.E.D.*

By implication, agents make more accurate decisions when better informed, and the accuracy entails that agents invest more frequently when the actual state is high, and less when the actual state is low. Public signals also align agents' information and hence are expected to have the same positive impact on agents' behaviors. It is noteworthy that the above analysis is conducted away from the limit precise environment, which is a focus of the global games for the selection of a unique equilibrium. Since we treat the global game methodology more than just the equilibrium selection device but also a useful model producing practical insights, it is important to demonstrate that the approach applies to situations away from the limit.

3.4. The T-Period Model

In this section, we augment the game into $T \in \mathbb{N}$ periods in which learning (i.e., the pure prediction game) happens repeatedly in period 1 to $T-1$, and the coordination stage is in period T.

In detail, at the beginning of period 1, agents still hold an improper prior over θ such that $\theta \sim \text{Unif}(\mathbb{R})$, and observe two signals x_{1i} and p_1 about $\theta's$ realization in the forms of (3.2). In period $t = 1, 2, \cdots, T-1$, which together comprises the learning stage, agent i faces a problem to select an action a_{ti} out of $\mathbb{R}$ to minimize the mean square error in predicting θ:

$$\min_{a_{ti} \in \mathbb{R}} E[(\theta - a_{ti})^2],$$

given whatever information she has. Period T is the coordination stage and is essentially the same as in section 2. That is, agent i chooses an action $a_i \in \{0, 1\}$ to maximize (3.1), given her information. Let $A_t = \int_i a_{ti} di$ denote the aggregate activity in period $t = 1, 2, \cdots, T-1$, and $A = \int_i a_i di$ the total activity in period T.

Now we describe the information structures in the game. Agents in each period $t = 2, 3, \cdots, T$ observe the past activity A_{t-1} through both a private and a public channel. Specifically, agent i observes x_{ti} and p_t about A_{t-1} and we structure them by

$$x_{ti} = A_{t-1} + \frac{1}{\sqrt{\tau_\varepsilon}} \varepsilon_{ti}, \qquad p_t = A_{t-1} + \frac{1}{\sqrt{\beta_\varepsilon}} \varepsilon_t,$$

where ε_{ti} and ε_t are respectively periodical private and public shocks that are standard normals, independent of each other, periods and all other parameters; $\tau_\varepsilon > 0$ and $\beta_\varepsilon > 0$ are observational precisions.

The agent i's objective in the game is to maximize the (undiscounted) sum of expected payoffs, and as before, this is equivalent to maximizing the expected payoff in each period separately, due to action spaces being $\mathbb{R}$ and infinitesimal agents. We without loss of generality focus on agents playing symmetric strategies, denoted by $a_t(\{x_{ki}, p_k\}_{k=1}^t) \in \mathbb{R}$ for period t, which prescribe a quantity (out of $\mathbb{R}$ or $[0, 1]$) from an agent's information set. And a symmetric equilibrium consists of a symmetric strategy profile such that the action prescribed that the strategy maximizes the agent's expected payoff of the game, given her information and all other agents play that strategy, in every period; and agents update their information by Bayes' rule. We still restrict to agent play a threshold strategy at the coordination game stage in period T.

It is straightforward to note that there exists a unique optimal strategy, denoted $a_t^*(\{x_{ki}, p_k\}_{k=1}^t)$, for agent i in the learning stage in periods $t = 1, 2, \cdots, T-1$. It is such that

$$a_t^*(\{x_{ki}, p_k\}_{k=1}^t) = E[\theta \mid \{x_{ki}, p_k\}_{k=1}^t].$$

We now show a well-behaved expression exists for this unique optimizer.

3.4.1 Solving the Game

We first solve for the equilibrium in the learning stage. Let $\hat{x}_{ti}(\{x_{ki}\}_{k=1}^t)$ denote agent i's private estimation of θ in period t from her private signal observations $\{x_{ki}\}_{k=1}^t$. Likewise, let $\hat{p}_t(\{p_k\}_{k=1}^t)$ denote the public estimation of θ in period t from the public information $\{p_k\}_{k=1}^t$, for all t. We shorthand the two statistics by $\hat{x}_{ti}$ and $\hat{p}_t$, respectively, and will give an analytical form for each of them.

Proposition 12. *For any $t = 1, 2, \cdots, T-1$, there exists a unique equilibrium characterized by $a_t^*(\hat{x}_{ti}, \hat{p}_t)$ such that*

$$a_t^*(\hat{x}_{ti}, \hat{p}_t) = E[\theta \mid (\hat{x}_{ti}, \hat{p}_t)].$$

Our proof consists of three steps. First we posit functional forms for $\hat{x}_{ti}$ and $\hat{p}_t$, and second we take as given that estimations are in these forms to solve for an equilibrium (if such an equilibrium exists). Thirdly, we indeed successfully characterize such an equilibrium, justifying the initial posit on the forms of $\hat{x}_{ti}$ and $\hat{p}_t$.

Proof. Step 1. We begin with positing the following assumption 1.

Assumption 1: For every t, agent i's estimations $\hat{x}_{ti}$ and $\hat{p}_t$ about θ, respectively, can be expressed in the following forms:

$$\hat{x}_{ti} = \theta + \frac{1}{\sqrt{\hat{\tau}_t}}\varepsilon_{ti}, \quad \hat{p}_t = \theta + \frac{1}{\sqrt{\hat{\beta}_t}}\varepsilon_t,$$

where $\hat{\tau}_t$ and $\hat{\beta}_t$ are positive constants that will be given later. As a result, the state θ is normally distributed given $\hat{x}_{ti}$ and $\hat{p}_t$, respectively.

Step 2. Given Assumption 1, we obtain immediately that the unique optimizer $a_t^*(\hat{x}_{ti}, \hat{p}_t)$ for agent i in period $t = 1, 2, \cdots, T-1$ equals

$$\begin{aligned} a_t^*(\hat{x}_{ti}, \hat{p}_t) = E[\theta \mid \hat{x}_{ti}, \hat{p}_t] &= \frac{\hat{\tau}_t}{\hat{\tau}_t + \hat{\beta}_t}\hat{x}_{ti} + \frac{\hat{\beta}_t}{\hat{\tau}_t + \hat{\beta}_t}\hat{p}_t \\ &\equiv \delta_t \hat{x}_{ti} + (1-\delta_t)\hat{p}_t, \end{aligned}$$

where $\delta_t = \hat{\tau}_t / (\hat{\tau}_t + \hat{\beta}_t)$ is the common weight on one's private information.

Till now, we have obtained an equilibrium for the learning stage provided Assumption 1 holds. Since as discussed the learning stage admits only one equilibrium, if we can show that Assumption 1 holds given agents follow strategy $a_t^*(\hat{x}_{ti}, \hat{p}_t)$ in the above form, we finish our proof. Step 3 achieves this.

Step 3 Given a_t^*, we know that the aggregate activity A_t in the proposed equilibrium is a function of θ and $\hat{p}_t$, such that

$$A_t(\theta, \hat{p}_t) = \delta_t \theta + (1-\delta_t)\hat{p}_t. \tag{3.7}$$

Recall that $x_{ti} = A_{t-1} + \varepsilon_{ti}/\sqrt{\tau_\varepsilon}$ so that

$$x_{ti} = \delta_{t-1}\theta + (1-\delta_{t-1})\hat{p}_{t-1} + \frac{1}{\sqrt{\tau_\varepsilon}}\varepsilon_{ti}.$$

Rearranging it we obtain

$$\frac{1}{\delta_{t-1}}(x_{ti} - (1-\delta_{t-1})\hat{p}_t) = \theta + \frac{1}{\delta_{t-1}\sqrt{\tau_\varepsilon}}\varepsilon_{ti},$$

and consequently if we denote the L.H.S by x'_{ti}, such that

$$x'_{ti} \equiv \delta_{t-1}(x_{ti} - (1-\delta_{t-1})\hat{p}_t) = \theta + \frac{1}{\sqrt{\tau_t}}\varepsilon_{ti}, \tag{3.8}$$

where $\tau_t = \delta_{t-1}^2 \tau_\varepsilon$, then x'_{ti} is a sufficient statistic to x_{ti} with respect to information about θ. As a result, agent i's information (i.e., her estimation) of θ in period t is, by Bayes' rule

$$\sum_{k=1}^{t} \frac{\tau_1 x_{1i} + \tau_2 x'_{2i} + \cdots + \tau_t x'_{ti}}{\tau_1 + \tau_2 + \cdots + \tau_t}, \tag{3.9}$$

which is exactly what $\hat{x}_{ti}$ represents and hence we obtain an analytical form of $\hat{x}_{ti}$ (with $\hat{\tau}_t = \sum_{k=1}^{t} \tau_k$) that satisfies Assumption 1 (i.e., θ is normally distributed given a statistic defined by (3.9) and the expectation of θ from private information is also given by (3.9); this is exactly the definition of $\hat{x}_{ti}$ in Assumption 1). Similarly, we can obtain $\hat{p}_t$ (with precision $\hat{\beta}_t = \sum_k \beta_k$, where $\beta_k = \delta_{t-1}^2 \beta_\varepsilon$) and thus complete our proof. *Q.E.D.*

Since agents' updated information $\hat{x}_{ti}$ and $\hat{p}_t$ are both of an additive structure with Gaussian noise, the coordination game stage can be solved similarly as we show in previous sections. That is, an equilibrium in period T characterized by a threshold pair $(x^*(\hat{p}_T), \theta^*(\hat{p}_T))$

for any realization of $\hat{p}_T$, exists, and it is the only equilibrium if and only if

$$\hat{\beta}_T / \sqrt{\hat{\tau}_T} \leqslant \sqrt{2\pi}. \tag{3.10}$$

Also because of the similarity in information structures, the (positive) impact of learning precisions on agents' expected payoffs can be similarly obtained to that in Proposition 3.

3.4.2 The Rise of Multiplicity

In this section, we show that, in the multi-period case, private learning, unlike definitely dispersing agents' information in the two-period setting, sometimes facilitates information homogeneity and thus prompts multiplicity. To this end, we first tract the respective evolvements of agents' private information and public information, namely, how $\hat{\tau}_t$ and $\hat{\beta}_t$ vary among periods, so as to obtain their relative values at time T. This is because this determines multiplicity due to condition (3.10). The following proposition summarizes our results concerning the evolvements of $\hat{\tau}_t$ and $\hat{\beta}_t$.

Proposition 13. *(i) Precisions $\hat{\tau}_T$ and $\hat{\beta}_T$ both converge to ∞ as $T \to \infty$.*

(ii) The rates of convergence of both public and private information are the same and are linear, in the sense that

$$\hat{\tau}_T / T \to \frac{1}{(1+\beta_\varepsilon/\tau_\varepsilon)^2}\tau_\varepsilon, \quad \hat{\beta}_T / T \to \frac{1}{(1+\beta_\varepsilon/\tau_\varepsilon)^2}\beta_\varepsilon,$$

as $T \to \infty$.

The first property states that both private and public information becomes fully revealing of θ if learning lasts long enough, and the second characterizes their convergence rates. Intuitively, with private signals incorporated, the crowding-out effect from public information diminishes, and Vives (1997) shows that pure public learning case successfully reveals the true θ, so that information in our environment should also be fully revealing; hence (i) should hold. Consequently, agents respond to both types of information equally since they are both limit accurate and as a result, the convergence rate should be the same as if observing i.i.d. signals (here, identically limit accurate signals), which is linear; hence we expect (ii). It is noteworthy that the above results hold as long as private learning exists $\tau_\varepsilon > 0$, so our results are expected to be applicable to a variety of sceneries.

The results and the arguments are essentially the same as those in Chapter 2, so we only state what matters for later analysis and relegate other proofs to Appendix. Denote by $\Delta\hat{\tau}_t \equiv \hat{\tau}_{t+1} - \hat{\tau}_t$ the change in the informativeness of private information from time t to time $t+1$, and analogously define $\Delta\hat{\beta}_t \equiv \hat{\beta}_{t+1} - \hat{\beta}_t$. By the constructions of x'_{ti} and p'_t in Proposition 3, step 3, we know $\Delta\hat{\tau}_t = \tau_{t+1}$ and $\Delta\hat{\beta}_t = \beta_{t+1}$. Therefore, for all $t = 1, 2, \cdots T$,

$$\Delta\hat{\tau}_t = (\frac{\hat{\tau}_t}{\hat{\tau}_t + \hat{\beta}_t})^2 \tau_\varepsilon, \quad \Delta\hat{\beta}_t = (\frac{\hat{\tau}_t}{\hat{\tau}_t + \hat{\beta}_t})^2 \beta_\varepsilon,$$

and consequently,

$$\Delta\hat{\beta}_t = \frac{\beta_\varepsilon}{\tau_\varepsilon} \Delta\hat{\tau}_t.$$

The above equation indicates that the relative increment of precision levels is constant, and hence $\Delta[\hat{\beta}_t - \beta_\varepsilon/\tau_\varepsilon \hat{\tau}_t] = 0$, for all t; in turn we have

$$\hat{\beta}_t - \frac{\beta_\varepsilon}{\tau_\varepsilon}\hat{\tau}_t = \beta_1 - \frac{\beta_\varepsilon}{\tau_\varepsilon}\tau_1 \equiv \alpha,$$

denoted the constant by α. Therefore, we have the relation

$$\hat{\beta}_t = \alpha + \frac{\beta_\varepsilon}{\tau_\varepsilon}\hat{\tau}_t, \quad \text{for all } t.$$

Hence the two types of information converge at the same rate.

Now we inspect the impact of private learning on multiplicity. First, by statement (ii) of the proposition, we know that

$$\frac{\hat{\beta}_T}{\sqrt{\hat{\tau}_T}} = \frac{\alpha + \hat{\tau}_T \beta_\varepsilon / \tau_\varepsilon}{\sqrt{\hat{\tau}_T}} \to \infty, \quad as \quad T \to \infty. \tag{3.11}$$

An immediate consequence of (3.11) and (3.10) is that multiplicity inevitably arises after enough rounds of learning, regardless of whether private learning is introduced; hence our hope that private learning sufficiently disperses agents' information fails to bear fruit. Noteworthy, since both $\hat{\tau}_T$ and $\hat{\beta}_T$ become limit accurate, the first-order beliefs about state θ are sure to be the same. We can at most hope that higher-order beliefs (e.g., how one agent perceives the beliefs of others) are dispersed by private learning. However, since all signals are essentially generated by θ, a higher first-order precision means more accurate inferences about others' beliefs as well, and thus agents' higher-order beliefs are also better aligned.

Also noteworthy, the convergence rate of $\hat{\beta}_t/\sqrt{\hat{\tau}_t}$ in (3.11) is the same as $\sqrt{\hat{\tau}_T} \to \infty$, or equivalently $\sqrt{\hat{\beta}_T} \to \infty$. On the other hand, if there were no private observation, the convergence rate of (3.11) would be the same as $\hat{\beta}_T \to \infty$, seemingly suggesting that the existence of private learning disperses agents' information and slows down (from $\hat{\beta}_T$ to $\sqrt{\hat{\beta}_T}$) the rise of homogeneity in information. This implication, however, is not true because as Proposition 4 (ii) shows, the presence of private learning accelerates the rate of $\hat{\beta}_T \to \infty$ to T, from $T^{1/3}$ in the pure public learning case. Hence the convergence rate of (3.11) is actually increased to $T^{1/2}$ (from $T^{1/3}$) with the incorporation of private learning. For a numeric example, if it takes 1000 rounds for the common knowledge in the ratio form (3.11) to appear with only public learning, it only requires 100 rounds when private learning is also available.

As a result, judging by the convergence rate of $\hat{\beta}_T/\sqrt{\hat{\tau}_T}$, the presence of private learning may, contrary to our intuitions, contributes to the happening of multiple equilibria. However, note that the rate is merely an indicator since it is asymptotic in the limit, while multiplicity arises as soon as $\hat{\beta}_T/\sqrt{\hat{\tau}_T} > \sqrt{2\pi}$ (the ratio is a necessary and sufficient condition for multiplicity). As a result, we need to take a closer look at the relationship between the introduction of private learning and the value of $\hat{\beta}_T/\sqrt{\hat{\tau}_T}$. We find that private learning does facilitate multiplicity, in some situations.

In detail, multiplicity arises if and only if

$$\frac{\hat{\beta}_T}{\sqrt{\hat{\tau}_T}} = \frac{\alpha + \hat{\tau}_T \beta_\varepsilon/\tau_\varepsilon}{\sqrt{\hat{\tau}_T}} > \sqrt{2\pi},$$

rewriting which we obtain

$$(\alpha\frac{\beta_\varepsilon}{\tau_\varepsilon})^2\hat{\tau}_T^2 + 2(\alpha\frac{\beta_\varepsilon}{\tau_\varepsilon} - \pi)\hat{\tau}_T + \alpha^2 > 0. \tag{3.12}$$

Therefore, if its discriminant $4\pi(\pi - 2\alpha\beta_\varepsilon/\tau_\varepsilon) < 0$, namely,

$$(\beta_1 - \tau_1\frac{\beta_\varepsilon}{\tau_\varepsilon})\frac{\beta_\varepsilon}{\tau_\varepsilon} > \pi/2, \quad (\because \alpha = \beta_1 - \frac{\beta_\varepsilon}{\tau_\varepsilon}\tau_1)$$

then (3.12) holds for all realizations of $\hat{\tau}_T$, meaning multiplicity is certain to appear even after one round of learning. Such scenarios happen, confirming our intuitions, when public information dominates from the start: β_1/τ_1 is high. On the other hand, when the discriminant is negative, suggesting uniqueness holds for some periods of interactions (that is, there exists $\hat{\tau}_T$ that makes $\hat{\beta}_T/\sqrt{\hat{\tau}_T} \leqslant \sqrt{2\pi}$). Note that private learning first disperses agents' information and then assembles it, as we have seen in the two-period model (see Section 3.3). Therefore, it is after some periods that the private learning starts to create publicity of information (*i.e.,* enlarging $\hat{\beta}_T/\sqrt{\hat{\tau}_T}$) and contributes to multiplicity. Therefore we claim the following.

Claim: Private learning facilitates multiplicity if it takes many periods for $\hat{\beta}_T/\sqrt{\hat{\tau}_T}$ to reach $\sqrt{2\pi}$.

Nevertheless, without further defining values of the parameters $\tau_1, \beta_1, \tau_\varepsilon$, and β_ε, we cannot conclude the exact number of periods needed for multiplicity. The following proposition outlines two situations in which it takes long enough periods for $\hat{\beta}_T/\sqrt{\hat{\tau}_T}$ to pass $\sqrt{2\pi}$ and thus arises multiplicity.

Proposition 14. *Private learning facilitates multiplicity if*

(i) β_1/τ_1 and $\beta_\varepsilon/\tau_\varepsilon$ are small enough, or

(ii) the investment succeeds iff $A_T \geqslant 1-\kappa\theta$ with large enough κ.

Situation (i) states that if publicity of information is initially small and public information is very noisy, then it takes many rounds for agents to form information conformity. For situation (ii), recall that we have set $\kappa \equiv 1$ in the successful coordination condition $A \geqslant 1-\kappa\theta$ throughout (see (3.1)). If we relax it, an immediate result from Proposition 1 arises as in the following remark.

Remark 6. *Given that the investment succeeds whenever $A_T \geqslant 1-\kappa\theta$ with $\kappa > 0$, there is a unique equilibrium iff $\hat{\beta}_T/\sqrt{\hat{\tau}_T} \leqslant \kappa\sqrt{2\pi}$.*

Therefore we conclude situation (ii). The intuition of the remark is straightforward - a higher κ means the success of coordination depends more on states rather than agents' behaviors, and hence beliefs about opponents' actions matter less. Now we prove (i) in Proposition 5.

Proof. Note that $\hat{\beta}_t/\sqrt{\hat{\tau}_t} = (\alpha + \hat{\tau}_t \beta_\varepsilon/\tau_\varepsilon)/\sqrt{\hat{\tau}_t}$. Hence the increment of the ration in period $t+1$ from period t is

$$\begin{aligned}
\Delta \frac{\hat{\beta}_t}{\hat{\tau}_t} &\equiv \frac{\alpha + \hat{\tau}_{t+1}\beta_\varepsilon/\tau_\varepsilon}{\sqrt{\hat{\tau}_{t+1}}} - \frac{\alpha + \hat{\tau}_t \beta_\varepsilon/\tau_\varepsilon}{\sqrt{\hat{\tau}_t}} \\
&= \frac{\Delta \tau_t \beta_\varepsilon/\tau_\varepsilon}{\sqrt{\hat{\tau}_{t+1}\hat{\tau}_t}}.
\end{aligned}$$

And $\hat{\beta}_T/\sqrt{\hat{\tau}_T} = \sum_{t=1}^{T} \Delta(\hat{\beta}_t/\sqrt{\hat{\tau}_t}) + \beta_1/\tau_1$, so the ratio is small when β_1/τ_1 and $\beta_\varepsilon/\tau_\varepsilon$ are small. *Q.E.D.*

References

Amador, M. and Weill, P.-O. (2012). Learning from private and public observations of others actions. *Journal of Economic Theory*, 147(3):910–940.

Angeletos, G.-M., Hellwig, C., and Pavan, A. (2006). Signaling in a global game: Coordination and policy traps. *Journal of Political economy*, 114(3):452–484.

Angeletos, G.-M., Hellwig, C., and Pavan, A. (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica*, 75(3):711–756.

Angeletos, G.-M. and Werning, I. (2006). Crises and prices: Information aggregation, multiplicity, and volatility. *American Economic Review*, 96(5):1720–1736.

Banerjee, A. V. (1992). A simple model of herd behavior. *The quarterly journal of economics*, 107(3):797–817.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, 100(5):992–1026.

Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.

Burguet, R. and Vives, X. (2000). Social learning and costly information acquisition. *Economic theory*, 15(1):185–205.

Carlsson, H. and Van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, 61(5):989–1018.

Chamley, C. and Gale, D. (1994). Information revelation and strategic delay in a model of investment. *Econometrica*, 62(5):1065–1085.

Che, Y.-K. and Hörner, J. (2018). Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, 133(2):871–925.

Dasgupta, A. (2007). Coordination and delay in global games. *Journal of Economic Theory*, 134(1):195–225.

Dasgupta, A., Steiner, J., and Stewart, C. (2012). Dynamic coordination with individual learning. *Games and Economic Behavior*, 74(1):83–101.

Edmond, C. (2013). Information manipulation, coordination, and regime change. *Review of Economic studies*, 80(4):1422–1458.

Feldman, M. (1987). An example of convergence to rational expectations with heterogeneous beliefs. *International Economic Review*, 28(3):635–650.

Gale, D. (1995). Dynamic coordination games. *Economic theory*, 5(1):1–18.

Goldstein, I. and Pauzner, A. (2005). Demand–deposit contracts and the probability of bank runs. *Journal of Finance*, 60(3):1293–1327.

Grossman, S. J. and Stiglitz, J. E. (1976). Information and competitive price systems. *American Economic Review*, 66(2):246–253.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408.

Hellwig, C., Mukherji, A., and Tsyvinski, A. (2006). Self-fulfilling currency crises: The role of interest rates. *American Economic Review*, 96(5):1769–1787.

Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.

Kremer, I., Mansour, Y., and Perry, M. (2014). Implementing the "wisdom of the crowd". *Journal of Political Economy*, 122(5):988–1012.

Lee, I. H. (1993). On the convergence of informational cascades. *Journal of Economic theory*, 61(2):395–411.

Morris, S. and Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3):587–597.

Morris, S. and Shin, H. S. (2000). Rethinking multiple equilibria in macroeconomic modeling. *NBER macroeconomics Annual*, 15:139–161.

Morris, S. and Shin, H. S. (2003). Global games: Theory and applications. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume 1*, pages 56–114. Cambridge University Press.

Morris, S. and Shin, H. S. (2004). Coordination risk and the price of debt. *European Economic Review*, 48(1):133–153.

Palfrey, T. R. (1985). Uncertainty resolution, private information aggregation and the cournot competitive limit. *The Review of Rconomic Studies*, 52(1):69–83.

Szkup, M. (2020). Multiplier effect and comparative statics in global games of regime change. *Theoretical Economics*, 15(2):625–667.

Szkup, M. and Trevino, I. (2015). Information acquisition in global games of regime change. *Journal of Economic Theory*, 160:387–428.

Townsend, R. M. (1978). Market anticipations, rational expectations, and bayesian analysis. *International Economic Review*, 19(2):481–494.

Vives, X. (1993). How fast do rational agents learn? *The Review of Economic Studies*, 60(2):329–347.

Vives, X. (1997). Learning from others: A welfare analysis. *Games and Economic Behavior*, 20(2):177–200.

Appendix A

Appendix to Chapter 1

Computations in Proposition 2 The monotonicity of $G_2(x_2'; (x_1, x_2'))$ in x_2' follows from that, given x_1,

$$\frac{\partial}{\partial x_2'} G_2(x_2'; (x_1, x_2')) = \frac{\partial}{\partial x_2'} \{E[r|x_2'] + P(x_{1j} > x_1|x_2') + (1 - P(x_{1j} > x_1|x_2'))\underbrace{P(\hat{x}_{2j} > x_2'|x_2')}_{=1/2}\}$$
$$= \frac{\partial}{\partial x_2'} \left\{ x_2' + \frac{1}{2}\Phi(\sqrt{\cdot}(x_2' - x_1)) + \frac{1}{2} \right\} > 0.$$

For the boundary value of $G_2(x_2'; (x_1, x_2'))$, since $E[r|\hat{x}_{2i}] = \hat{x}_{2i}$ and $\hat{a}_2 \in [0, 1]$, when $x_2' \to \infty$,

$$G_2(x_2'; (x_1, x_2')) = E[r + \hat{a}_2 \mid x_2'; (x_1, x_2')] \to \infty.$$

Now we prove the monotonicity of $\Delta(x_1'; (x_1', x_2^*(x_1')))$ in x_1'. Note that

$$\frac{\partial}{\partial x_1'} G_1(x_1'; (x_1', x_2^*(x_1'))) = \frac{\partial}{\partial x_1'} \left\{ E[r \mid x_1'] + \underbrace{P(x_{1j} > x_1'|x_1')}_{=1/2} + (1 - P(x_{1j} > x_1'|x_1'))P(\hat{x}_{2j} > x_2^*(x_1')|x_1') \right\}$$
$$= \frac{\partial}{\partial x_1'} \left\{ x_1' + \frac{1}{2} + \frac{1}{2}\Phi(\sqrt{\cdot}(x_1' - x_2^*(x_1'))) \right\} > 0,$$

since $dx_2^*(x_1')/dx_1' \in [0,1]$ by taking the total derivative of (1.7) with respect to x_1'. Also,

$$\begin{aligned}\frac{\partial}{\partial x_1'} R_1(x_1';(x_1',x_2^*(x_1'))) =& \delta \frac{d}{dx_1'} \int_{-\infty}^{\infty} E[r+\hat{a}_2 \mid \hat{x}_{2i}] f(\hat{x}_{2i} \mid x_1') \mathbf{1}_{\hat{x}_{2i}>x_2^*(x_1')} d\hat{x}_{2i} \\ \leqslant& \delta \frac{d}{dx_1'} \int_{-\infty}^{\infty} \underbrace{E[r+\hat{a}_2 \mid \hat{x}_{2i}]}_{=E[r+\hat{a}_2|\hat{x}_{2i},x_1']} f(\hat{x}_{2i} \mid x_1') d\hat{x}_{2i} \\ =& \delta \frac{d}{dx_1'} E[r+\hat{a}_2 \mid x_1'],\end{aligned} \tag{A.1}$$

where the first equation follows from $R_1's$ definition with $\mathbf{1}$ being the indicator function, the inequality is due to $\mathbf{1} \in [0,1]$. Therefore,

$$\frac{\partial}{\partial x_1'} \Delta(x_1';(x_1',x_2^*(x_1'))) \geqslant (1-\delta) \frac{\partial}{\partial x_1'} G_1(x_1';(x_1',x_2^*(x_1'))) > 0.$$

Next for the boundary value, since $R_1 \geqslant 0$, we have $\Delta(x_1';(x_1',x_2^*(x_1'))) \to -\infty$ as $x_1' \to -\infty$. On the other hand, as $x_1' \to \infty$, it is similar as in (A.1)) to obtain $R_1(x_1';(x_1',x_2^*(x_1'))) \leqslant \delta E[r+\hat{a}_2 \mid x_1';(x_1';x_2^*(x_1'))]$, so we also have

$$\Delta(x_1';(x_1',x_2^*(x_1'))) \geqslant (1-\delta) E[r+\hat{a}_2 \mid x_1';(x_1';x_2^*(x_1'))] \to \infty.$$

Existence of a Unique Solution to (1.16) For a pedagogical purpose, we verify the general case by showing

$$\begin{aligned}&\Delta_t(x_t';(x_1,\cdots,x_t',x_{t+1}^*,\cdots,x_N^*)) \equiv \\ &E[U(r,\hat{a}_N) \mid x_t';(x_1,\cdots,x_t',\cdots,x_N^*)] - R_t(x_t';(x_1,\cdots,x_t',\cdots,x_N^*)) = 0\end{aligned}$$

admits a unique solution x_t'. First recall by Lemma 2 we have $E[U(r,\hat{a}_N)|\hat{x}_{ti}]$ strictly increasing in $\hat{x}_{ti}$, given any threshold strategy profile the population plays. Hence in checking the monotonicity of $E[U(r,\hat{a}_N)|x_t';(x_1,\cdots,x_t',\cdots,x_N^*)]$ in x_t', if we can show that the investment at time t will not decrease, then by the result of Lemma 2, the aggregate investment increases in x'_t. Indeed, at time t, the fraction of agents who invest (in the eyes of the agent observing x'_t) equals $P(\hat{x}_{tj} > x'_t | x'_t) = 1/2$, namely, it is invariant. So we have $E[U(r, \hat{a}_N) \mid x'_t]$ is strictly increasing in x'_t. Next we show that $\partial R_t(x'_t)/\partial x'_t < \delta E[U(r, \hat{a}_N) | x'_t]$ by backward induction.

For $t = N-1$, it is the same as in the two-period model to obtain that $\partial R_{N-1}(x'_{N-1})/\partial x'_{N-1} < \delta E[r + \hat{a}_{N-1} | x'_{N-1}]$. Assume backward inductively that $\partial R_k(x'_k)/\partial x'_k < \delta E[U(r, \hat{a}_N) | x'_k]$ for all $k = N-2, N-3, \cdots, t+1$, so at time t,

$$\begin{aligned}
&\frac{\partial}{\partial x'_t} \int_{-\infty}^{x^*_{t+1}} R_{t+1}(\hat{x}_{(t+1)i}; x_1, \cdots, x'_t, x^*_{t+1}, \cdots, x^*_N) f(\hat{x}_{(t+1)i} \mid x'_t) d\hat{x}_{(t+1)i} \\
\leqslant &\delta \frac{\partial}{\partial x'_t} \int_{-\infty}^{x^*_{t+1}} E[r + \hat{a}_t \mid \hat{x}_{(t+1)i}] f(\hat{x}_{(t+1)i} \mid x'_t) d\hat{x}_{(t+1)i},
\end{aligned}$$

by noting that x^*_{t+1} increases in x'_t due to strategic complementarity (a higher signal to an individual does not affect x^*_{t+1}, but a higher threshold means fewer agents invest, which causes agents less willing to invest and thus x^*_{t+1} decreases). Therefore,

$$\begin{aligned}
\frac{\partial}{\partial x'_1} R_t(x'_t; x_1, \cdots, x'_t, x^*_{t+1}, \cdots, x^*_N) \leqslant &\delta \frac{d}{dx'_t} \int_{x^*_{t+1}}^{\infty} E[r + \hat{a}_t \mid \hat{x}_{(t+1)i}] f(\hat{x}_{(t+1)i} \mid x'_t) d\hat{x}_{(t+1)i} \\
&+ \delta \frac{\partial}{\partial x'_t} \int_{-\infty}^{x^*_{t+1}} E[r + \hat{a}_t \mid \hat{x}_{(t+1)i}] f(\hat{x}_{(t+1)i} \mid x'_t) d\hat{x}_{(t+1)i} \\
= &\delta \frac{\partial}{\partial x'_t} \int_{-\infty}^{\infty} E[r + \hat{a}_t \mid \hat{x}_{(t+1)i}] f(\hat{x}_{(t+1)i} \mid x'_t) d\hat{x}_{(t+1)i} \\
= &\delta \frac{\partial}{\partial x'_t} E[r + \hat{a}_t \mid x'_t].
\end{aligned}$$

Hence we conclude that

$$\frac{\partial}{\partial x'_t} \Delta_t(x'_t; (x_1, \cdots, x'_t, x^*_{t+1}, \cdots, x^*_N)) \geqslant (1-\delta) \frac{\partial}{\partial x'_t} E[U(r, \hat{a}_N) | x'_t] > 0.$$

Computation in Proposition 5 The following lemma establishes the monotonicity of $\Delta_t(\hat{x}_{ti};(x_1,\cdots,x_N))$ in agent i's current belief $\hat{x}_{ti}$, given any $(x_1,\cdots,x_N)$ and t.

Lemma 3. *Let $(x_1,x_2,\cdots,x_N)$ be an arbitrary threshold strategy profile. The expected continuation payoff $R_t(\hat{x}_{ti};x_1,x_2,\cdots,x_N)$ for agent i with $\hat{x}_{ti}$ at time t satisfies*

$$\frac{d}{d\hat{x}_{ti}}R_t(\hat{x}_{ti};x_1,x_2,\cdots,x_N) < \delta_t\frac{d}{d\hat{x}_{ti}}E[U(r,\hat{a}_N)\mid\hat{x}_{ti}],$$

for $t=1,2,\cdots,N-1$ and any $\hat{x}_{ti}$.

Proof. Note that agents can at most invest once, so it suffices to show that the derivative of each integrand in $R_t(\hat{x}_{ti};\{x_t\}_t)$ satisfies the property stated in the Lemma. For example for its first term concerning $t+1$,

$$\begin{aligned}&\frac{d}{d\hat{x}_{ti}}\delta_{t+1}\int_{-\infty}^{\infty}E[U(r,\hat{a}_N)\mid\hat{x}_{(t+1)i}]f(\hat{x}_{(t+1)i}\mid\hat{x}_{ti})\mathbf{1}_{\hat{x}_{(t+1)i}\geqslant x_{t+1}}d\hat{x}_{(t+1)i}\\ \leqslant&\frac{d}{d\hat{x}_{ti}}\delta_{t+1}\int_{-\infty}^{\infty}E[U(r,\hat{a}_N)\mid\hat{x}_{(t+1)i}]f(\hat{x}_{(t+1)i}\mid\hat{x}_{ti})d\hat{x}_{(t+1)i}\\ =&\frac{d}{d\hat{x}_{ti}}\delta_{t+1}E[U(r,\hat{a}_N)\mid\hat{x}_{ti}].\end{aligned}$$

Therefore, $R_t'(\hat{x}_{ti};x_1,\cdots,x_N)\leqslant\max\{\delta_{t+1}(E[U(r,\hat{a}_N)\mid\hat{x}_{ti}])',\cdots,\delta_N(E[U(r,\hat{a}_N)\mid\hat{x}_{ti}])'\}<\delta_t(E[U(r,\hat{a}_N)\mid\hat{x}_{ti}])'$. *Q.E.D.*

Proof of Proposition 6 For (i), recall

$$E[U(r,\hat{a}_N)|x_N^*]=0,$$

$$E[U(r,a_1)|x_{st}^*]=0.$$

Since $\hat{a}_N(x_N^*)>a_1(x_N^*)$ and $U(x,\hat{a}_N)$ increase in both elements, by contradiction it can be proved that $x_N^*<x_{st}^*$ and $\hat{a}_N(\hat{x}_N^*)>a_1(x_{st}^*)$. The second half is as in the two-period game.

That is, if $\delta = 0$, then $R_t = 0$ and thus $x_t^* = x_{st}^*$ for all $t = 1, 2, \cdots, N$. Since R_t increases in δ, when $\delta > 0$, we must have x_t^* also increase to satisfy the equilibrium condition of period t.

For (ii), it suffices to confirm that $R_t(x_t^*; \{x_t^*\}_t) = 0$ at every $t = 1, 2, \cdots, N$ when learning lacks. Fix an arbitrary t and x_{1i}. Recall that

$$\begin{aligned}
R_t(x_{1i}; \{x_t^*\}_t) = \delta & \int_{x_{t+1}^*}^{\infty} E[U(r, \hat{a}_N) \mid \hat{x}_{(t+1)i}] f(\hat{x}_{(t+1)i} \mid x_{1i}) d\hat{x}_{(t+1)i} \\
& + \delta^2 \int_{x_{t+2}^*}^{\infty} \int_{-\infty}^{x_{t+1}^*} E[U(r, \hat{a}_N) \mid \hat{x}_{(t+2)i}] f(\hat{x}_{(t+2)i}, \hat{x}_{(t+1)i} \mid x_{1i}) d\hat{x}_{(t+1)i} d\hat{x}_{(t+2)i} \\
& + \cdots \\
& + \delta^{N-t} \int_{x_N^*}^{\infty} \int_{-\infty}^{x_{N-1}^*} \cdots \int_{-\infty}^{x_{t+1}^*} E[U(r, \hat{a}_N) \mid \hat{x}_{Ni}] f(\hat{x}_{Ni}, \cdots, \hat{x}_{(t+1)i} \mid x_{1i}) d\hat{x}_{(t+1)i} \cdots d\hat{x}_{Ni}.
\end{aligned}$$

Since agent i holds constant belief x_{1i}, each integrant is mutually exclusive; therefore, only one integrant remains, so

$$R_t(x_{1i}; \{x_t^*\}_t) = \delta^k E[U(r, \hat{a}_N) \mid x_{1i}],$$

for some $k \in \{1, 2, \cdots, N-t\}$. To pin down x_t^*, it is required that

$$E[U(r, \hat{a}_N) | x_t^*] = R_t(x_t^*; \{x_t^*\}_t) = \delta^k E[U(r, \hat{a}_N) \mid x_t^*].$$

If $E[U(r, \hat{a}_N) | x_t^*] \neq 0$, the two sides can never be equal, so $E[U(r, \hat{a}_N) | x_t^*] = 0$.

Appendix B

Appendix to Chapter 3

Computation in Proposition 2: Note that varying τ_ε does not affect public signals, so we in below fix an arbitrary $\hat{p}_2$ and write thresholds as (x^*, θ^*). Note that, for agent i with information set (x_{1i}, p_1) in period 1, her expected payoff at the coordination stage is

$$\begin{aligned} D(x_{1i}, p_1) &= P(\hat{x}_{2i} > x^*, \theta \geqslant \theta^* | x_{1i}, p_1)(1-c) - P(\hat{x}_{2i} > x^*, \theta < \theta^* | x_{1i}, p_1)c \\ &= (1-c)[\int_{\theta^*}^{\infty} g(\theta|x_{1i}, p_1)d\theta - \underbrace{\int_{\theta^*}^{\infty}\int_{-\infty}^{x^*} f(\hat{x}_{2i}|\theta)g(\theta|x_{1i}, p_1)\, d\hat{x}_{2i}d\theta}_{\text{Type II error}}] \\ &- c\underbrace{\int_{-\infty}^{\theta^*}\int_{x^*}^{\infty} f(\hat{x}_{2i}|\theta)g(\theta|x_{1i}, p_1)\, d\hat{x}_{2i}d\theta}_{\text{Type I error}}, \end{aligned}$$

where $g(\theta|x_{1i}, p_1) \equiv (G(\theta|x_{1i}, p_1))' = \sqrt{\tau_1 + \beta_1}\phi\left(\sqrt{\tau_1 + \beta_1}(\theta - (\tau_1 + \beta_1)^{-1}(\tau_1 x_{1i} + \beta_1 p_1))\right)$. Similarly $f(\hat{x}_{1i}|\theta) \equiv (F(\hat{x}_{1i}|\theta))' = \sqrt{\hat{\tau}_2}\phi(\sqrt{\hat{\tau}_2}(x^* - \theta))$. [1] Consider the Type I error; it equals

$$\int_{-\infty}^{\theta^*}\int_{x^*}^{\infty} g(\theta|x_{1i}, p_1)[1 - F(x^*|\theta)]d\theta.$$

[1]That is, we denote the CDFs and PDFs of nonstandard normals by (F, G) and (f, g), which can be obtained by transforming the standard normal's Φ or ϕ.

Differentiating Type I error w.r.t. $\hat{\tau}_2$ yields

$$
\begin{aligned}
&-\int_{-\infty}^{\theta^*} g(\theta|x_{1i},p_1)\frac{(x^*-\theta)}{2\hat{\tau}_2}f(x^*|\theta)d\theta = -\int_{-\infty}^{\theta^*}\tilde{g}(\theta|x^*)\frac{(x^*-\theta)}{2\hat{\tau}_2}\tilde{f}(x^*|x_{1i},p_1)\,d\theta \\
=&-\int_{-\infty}^{\theta^*}\tilde{g}(\theta|x^*)d\theta x^*\tilde{f}(x^*|x_{1i},p_1)/(2\hat{\tau}_2)+\int_{-\infty}^{\theta^*}\tilde{g}(\theta|x^*)\theta d\theta\tilde{f}(x^*|x_{1i},p_1)/(2\hat{\tau}_2),
\end{aligned}
$$

where we use the fact that $g(\theta|x_{1i},p_1)f(x^*|\theta) = g(\theta|x_{1i},p_1)f(x^*|\theta,x_{1i},p_1) = \tilde{g}(\theta|x^*,x_{1i},p_1)\tilde{f}(x^*|x_{1i},p_1) = \tilde{g}(\theta|x^*)\tilde{f}(x^*|x_{1i},p_1)$ in the first equation. And $\tilde{g}(\theta|x^*)$, $\tilde{f}(x^*|x_{1i},p_1)$ are PDFs with corresponding CDF denoted $\tilde{G}$ and $\tilde{F}$. Analogously, we differentiate Type II error w.r.t $\hat{\tau}_2$ and it equals

$$
\int_{\theta^*}^{\infty}\tilde{g}(\theta|x^*)d\theta x^*\tilde{f}(x^*|x_{1i},p_1)/(2\hat{\tau}_2)-\int_{\theta^*}^{\infty}\tilde{g}(\theta|x^*)\theta d\theta\tilde{f}(x^*|x_{1i},p_1)/(2\hat{\tau}_2).
$$

Therefore,

$$
\frac{\partial D(x_{1i},p_1)}{\partial\hat{\tau}_2} = \frac{\tilde{f}(x^*|x_{1i},p_1)}{2\hat{\tau}_2}\left[(1-c)\int_{\theta^*}^{\infty}\tilde{g}(\theta|x^*)\theta d\theta - c\int_{-\infty}^{\theta^*}\tilde{g}(\theta|x^*)\theta d\theta\right]. \tag{B.1}
$$

Here, we claim that the term within the square bracket can cancel out and equals $\tilde{g}(\theta^*|x^*)/\hat{\tau}_2$. To see this, note that

$$
\begin{aligned}
&\int_{\theta^*}^{\infty}\tilde{g}(\theta|x^*)\theta d\theta \\
=&\int_{\theta^*}^{\infty}\frac{\sqrt{\hat{\tau}_2}}{\sqrt{2\pi}}\exp\left(-\frac{\hat{\tau}_2(\theta-x^*)^2}{2}\right)(\theta-x^*)d\theta+\int_{\theta^*}^{\infty}\frac{\sqrt{\hat{\tau}_2}}{\sqrt{2\pi}}\exp\left(-\frac{\hat{\tau}_2(\theta-x^*)^2}{2}\right)x^*d\theta \\
=&\tilde{g}(\theta^*|x^*)/\hat{\tau}_2+x^*[1-\tilde{G}(\theta^*|x^*)],
\end{aligned}
$$

where the first term is obtained by the change of variables and the second term is by the definition of $\tilde{G}(\theta^*|x^*)$. Similarly, we obtain

$$\int_{-\infty}^{\theta^*} \tilde{g}(\theta|x^*)\theta d\theta = -\tilde{g}(\theta^*|x^*)/\hat{\tau}_2 + x^*\tilde{G}(\theta^*|x^*).$$

Note that in equilibrium $\tilde{G}(\theta^*|x^*) = 1-c$ by the payoff indifference condition. Therefore, (B.1) equals

$$\frac{\partial D(x_{1i},p_1)}{\partial \hat{\tau}_2} = \frac{1}{2\hat{\tau}_2^2}\tilde{f}(x^*|x_{1i},p_1)\tilde{g}(\theta^*|x^*) > 0$$

Proof of Proposition 4 Statement (i) First notice that if the precision of one type of information goes to ∞, so does the other type because of the above argument (the informativeness of both types of information increase at the same rate). And suppose none goes to ∞; then $\Delta\hat{\tau}_t > 0$ for any t, and thus $\hat{\tau}_T = \sum_{k=1}^{T-1}\Delta\hat{\tau}_t + \tau_1$ diverges, causing a contradiction.

Statement (ii) We characterize the rate for $\hat{\tau}_T \to \infty$, and the rate for $\hat{\beta}_T$ follows from $\hat{\beta}_T = \alpha + \hat{\tau}_T\beta_\varepsilon/\tau_\varepsilon$. From statement (i), there exists $T_\varepsilon \in \mathbb{N}$ such that for all $t > T_\varepsilon$, $\hat{\tau}_t \to \infty$. Next note that, for any $t = 1,\cdots,T$,

$$\begin{aligned}\hat{\tau}_t = \sum_{k=1}^{t-1}\Delta\hat{\tau}_k + \tau_1 = \tau_\varepsilon\sum_{k=1}^{t-1}\left(\frac{\hat{\tau}_k}{\hat{\tau}_k+\hat{\beta}_k}\right)^2 + \tau_1 &= \tau_\varepsilon\sum_{k=1}^{t-1}\left(\frac{\hat{\tau}_k}{(1+\beta_\varepsilon/\tau_\varepsilon)\hat{\tau}_k+\alpha}\right)^2 + \tau_1 \\ &= \tau_\varepsilon\sum_{k=1}^{t-1}\left(\frac{\hat{\tau}_k}{(1+\beta_\varepsilon/\tau_\varepsilon)\hat{\tau}_k+\alpha}\right)^2 + \tau_1.\end{aligned}$$

Divided both sides by T, and let $T \to \infty$; we obtain

$$\begin{aligned}\hat{\tau}_T/T = \tau_\varepsilon\sum_{k=0}^{T_\varepsilon}\left(\frac{\hat{\tau}_k}{(1+\beta_\varepsilon/\tau_\varepsilon)\hat{\tau}_k+\alpha}\right)^2/T + \tau_\varepsilon\sum_{k=T_\varepsilon+1}^{T-1}\left(\frac{\hat{\tau}_k}{(1+\beta_\varepsilon/\tau_\varepsilon)\hat{\tau}_k+\alpha}\right)^2/T + \tau_1/T \\ \to \frac{1}{(1+\beta_\varepsilon/\tau_\varepsilon)^2}\tau_\varepsilon,\end{aligned}$$

which is a non-zero constant, so $\hat{\tau}_T/T \to \infty$ at a linear rate.