London School of Economics and Political Science

Essays in Social Learning

Min Zhang

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy, London, June 2015.

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 29, 047 words.

London, 22 June 2015

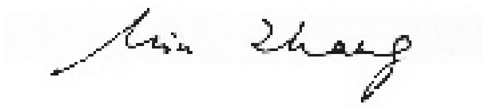

Min Zhang

Statement of Conjoint Work

I confirm that Chapter 3, "Learning and Price Dynamics in Durable Goods Markets", was jointly co-authored with Mr Francesco Palazzo. I contributed a minimum of 50% of this work.

General decisions about the directions of research, and the proofs of the main results, were made equally between the authors.

Francesco Palazzo

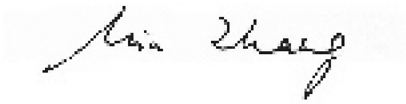

Min Zhang

Acknowledgements

I am deeply indebted to my supervisor Erik Eyster for his consistent inspiration and encouragement throughout the course of my PhD. I benefit so much from many long and fruitful discussions with Erik; he helps me transform from a graduate student to a proper economist. I would also like to thank Balazs Szentes as my advisor over the second year of my PhD; his suggestions and comments on my work are always sharp and helpful.

Many other researchers have provided feedback on several parts of the thesis: Andrew Ellis, Leonardo Felli, Gilat Levy, Matthew Levy, Francesco Nava, Andrea Prat, Matthew Rabin, Ronny Razin, as well as the seminar participants at LSE. I thank them for all the useful comments and discussions.

I want to thank Francesco Palazzo, whom I am fortunate to know at LSE as a good friend and a smart co-author. I have learned a lot studying and working with him. Thanks also to my fellow PhD students at LSE for making the PhD experience so much more enjoyable.

A special thanks to Pierpaolo Benigno for his encouragement and generous support on my PhD study.

Finally, I want to thank my family for their consistent support. I thank my parents for their financial support during my first years at LSE, and more importantly, for their care and love since I was born. I thank my wife Yanjie for coming with me to the UK despite a promising career in China, and for believing in me throughout the years.

Abstract

This thesis contains two theoretical essays built upon the canonical models of social learning, and one that applies social learning theory to durable goods markets.

The first chapter, "Non-Monotone Observational Learning", revisits the canonical social-learning model that rationalizes herding in the long run, to investigate the possibility of non-imitative behavior in the short run generated by non-monotone learning: *ceteris paribus*, when some predecessor(s) switch to actions revealing greater confidence in one state of the world, agents become less confident in that state. I characterize conditions on the underlying information structures that lead to non-monotone learning. In particular, in a general setting with continuous private signals, I provide a necessary condition for non-monotone learning with an argument for its plausibility, as well as two non-restrictive sufficient conditions that do not rely on parametrization.

The second chapter, "Does Public Information Disclosure Help Social Learning?", studies the effect of releasing exogenous public information in the canonical social-learning model that predicts incomplete learning. To improve social learning, I show that it is weakly better to postpone the disclosure of a public signal irrespective of its precision. However, such weak monotonicity no longer holds if the objective is to maximize the discounted sum of people's expected payoffs or if the model goes beyond the canonical binary setting. On the other hand, it is suboptimal to ever release a public signal less precise than people's private signals even if sophisticated releasing strategies are allowed.

The last chapter, "Learning and Price Dynamics in Durable Goods Markets", is joint work with Francesco Palazzo. We study how markets for durable goods with unobservable and time-varying aggregate market conditions determine price dynamics with market participants constantly learning from public observations. We set up a dynamic auction model with two key features: first, agents enjoy heterogeneous private use values and later resell the asset; second, prices do not incorporate all available information dispersed in the economy. Informational frictions slow down learning and affect price movements asymmetrically across high and low aggregate demand states. Learning and the resale motive are the predominant force for durable goods with short resale horizons, slow time-varying aggregate demand, and similar use values across agents.

Contents

List of Figures

1 Non-Monotone Observational Learning

Whereas rational observational learning generates herds in the long run, it can lead to behavior quite different from herding or imitation in the short run. This work revisits the canonical binary-state model of observational learning, in which agents sequentially choose a binary action and the history of actions is publicly observable, to investigate the possibility of non-imitative behavior generated by non-monotone learning: *ceteris paribus*, when some predecessor(s) switch to actions revealing greater confidence in one state of the world, agents become less confident in that state. In a special case with binary signal space, we show that most agents always form such non-monotone posterior beliefs with respect to the first agent's action. In a general setting with continuous signals, we provide a necessary condition for non-monotone learning, and show that it fails only for information structures that never generate public beliefs between $\frac{1}{3}$ and $\frac{2}{3}$ throughout the learning process. We also provide two non-restrictive sufficient conditions for non-monotone learning on information structures that are not explicitly parameterized.

1.1 Introduction

The theoretical literature on observational learning has demonstrated that, when rational agents with common preferences act sequentially, they tend to imitate their predecessors and eventually exhibit herd behavior. In the canonical example of restaurant choice due to Banerjee (1992), agents sequentially choose to eat at one of the two restaurants, A and B, which are equally likely to be the better one *a priori*, and each agent obtains a binary private signal indicating the better restaurant before making her choice. Assuming all signals have the same precision and agents follow their private signals when indifferent, when one restaurant has been chosen twice more than other, all future agents will ignore their private signals and choose that restaurant.

Rationalizing herd behavior and imitation is considered as one key contribution of the literature.[1] Yet we are keen to know what else rational observational learning may predict. Do rational agents always tend to imitate their predecessors throughout the learning process? When they imitate, do they imitate each and every predecessor? Answers to questions like these rely on a thorough understanding of rational learning models beyond the well-established long-run predictions. In addition, more studies into the short-run behavioral implications of rational observational learning serve to distinguish rationality from other potential explanations of social learning, both theoretically and empirically.

Let us start by looking into the learning dynamics in the canonical example above. We can see that agents' incentives to imitate their predecessors result from a "monotone" evolution of posterior beliefs. More specifically, each agent forms posterior beliefs that are monotone with respect to her observation, in the sense that she believes one restaurant, say A, is (weakly) more likely to be better than the other when more of her predecessors have switched to choose A. When she observes that A has been chosen at least twice more than B, her posterior belief after such history *dominates* her private information and she then chooses A regardless of her private signal.

Is such monotonicity a general feature of observational learning? It seems plausible at first glance. By choosing a restaurant, each agent reveals more confidence in that

[1] Quoting the preface of Chamley (2004) which provides a comprehensive overview of the literature, "Learning by individuals from the behavior of others and imitation pervade the social life... herds, fads, bubbles, crashes, and booms are cited as proofs of the irrationality of individuals. However, most of these colorful events will appear in the models of rational agents...".

restaurant than the other, so it seems natural that, if more agents have chosen A instead of B in the past, the history becomes stronger evidence supporting A. However, this is *not* necessarily the case as the interpretation of each action in the history closely depends on other actions, *e.g.*, how much confidence in A agent t expresses by choosing A may change dramatically if the action of agent $t-1$ changes. And these interpretations are particularly affected by the underlying information structure of the model. In a variety of settings, rational agents could indeed form non-monotone posterior beliefs with respect to their observations, which we refer as "non-monotone learning": *ceteris paribus*, when some predecessor(s) switch to actions revealing greater confidence in one state of the world, agents become less confident in that state. As a consequence, agents have *less* incentive to imitate their predecessors, a feature that is absent in the canonical example and also overlooked in models extended from it.

For instance, in the restaurant-choice problem, certain signal structures could lead agent 3 to believe B is more likely to be the better restaurant when she observes a history (A, B) than when she observes a history (B, B). Fixing agent 2's action, the *more* agent 1's action reveals confidence in A, the *less* confident agent 3 becomes in A. We refer this particular case of non-monotone learning, where agent 3's beliefs shift against the action of agent 1, as agent 3 forming posterior beliefs that are *anti-imitative* of agent 1.[2] It is worth mentioning that fixing agent 2's action makes the comparison more interesting, because agent 3's posterior belief is always consistent with agent 2's action, *i.e.*, the *Overturning Principle* in Smith and Sørensen (2000).[3]

To further elaborate such an example, we adopt a setting close to that of Callander and Hörner (2009) in which agents are heterogeneously informed. Some agents are *experts* who have private signals about whether A or B is better with precision $q > 0.5$. Others are just *amateurs* who know nothing at all. In Callander and Hörner (2009), agents can only observe the total number of agents having chosen each option, and they show that uninformed agents (*amateurs*) should follow the *minority* of their predecessors rather than the majority, when informed agents (*experts*) are rare.[4] Such behavior is clearly different

[2] The formal definition of anti-imitative beliefs is provided in Section 1.2.

[3] See the proof of Theorem 3 in Smith and Sørensen (2000) or Lemma 1.3.2 in this paper.

[4] Callander and Hörner (2009) assume that informed agents have perfect private information, *i.e.*, $q = 1$, hence they only focus on the rational behavior of uninformed agents.

from what might be expected based on the canonical example, and more fundamentally, the posterior beliefs of uninformed agents are clearly not monotone with respect to the summary statistic they observe. Here we show that similar but less extreme information structures can indeed lead to non-monotone learning even if we maintain the assumption that each agent observes the order of all predecessors' moves as in the canonical example. In particular, when the fraction of experts in the population, r, is sufficiently large or the precision q is sufficiently high, agent 3 will form beliefs that are anti-imitative of agent 1 as described in the last paragraph. To get the intuition, think about the comparison between history (A, B) and history (B, B). (A, B) reveals one weak piece of evidence against B from the first action A, as it could just be a random choice by an amateur. Meanwhile (A, B) reveals one strong piece of evidence supporting B from the second action B, as it must be an informative choice by an expert since an amateur, lacking private information, would have followed the first action A. On the other hand, (B, B) reveals two weak pieces of evidence supporting B as each action B could come from an amateur. When r (q) is very high which implies the weak (strong) evidence is fairly insignificant (significant), (A, B) may turn out to be an overall stronger piece of evidence supporting B than (B, B). Furthermore, we also provide conditions on q and r for massive instances of anti-imitative beliefs: each agent from agent 3 on will form posterior beliefs that are anti-imitative of agent 1 after any possible equilibrium history, even though they all share the same preference as agent 1 and there are no strategic effects at all.[5]

Knowing that different information structures can dramatically affect a rational agent's inference from her observations and possibly lead to non-monotone learning, we then turn to a general model with continuous private signals due to Smith and Sørensen (2000). A nice feature of this model is that it guarantees behavioral differences whenever learning is non-monotone. In particular, if one agent forms beliefs that are anti-imitative of a predecessor and private signals are continuous, she would indeed choose B with higher probability had that predecessor switched from B to A, *i.e.*, she *anti-imitates* that predecessor.[6] On the other hand, we want to emphasize that such anti-imitative behavior does not contradict but rather sit side by side with the long-run herd behavior that has been

[5] Agent 2's posterior belief cannot be anti-imitative according to the overturning principle.

[6] We show in Section 1.2 that, with binary private signals, anti-imitative beliefs do not necessarily induce such behavioral difference.

demonstrated in this model.[7] Hopefully our investigation could draw further attention to comprehensive and thorough studies of many existing learning models in which rational behavior during the learning process may be quite different from the long-run predictions already established.

We first characterize the learning dynamics along any equilibrium path by introducing two transition functions that describe the movement of public beliefs from one period to the next, one for each possible realization of the most recent action. In other words, these transition functions govern the learning process in term of beliefs along all possible histories. A necessary condition for non-monotone learning is that the transition functions are non-monotone in public beliefs. To see why non-monotone transition functions are necessary, let us consider a change of action from A to B by a number of agents. According to the overturning principle, the public belief right after the last agent who changes his action now supports B instead of A. If the transition functions are always monotone, then all the public beliefs generated thereafter will always shift toward B and hence learning cannot be non-monotone. It is probably too abstract to think about the shape of transition functions, but we make an interesting observation: if transition functions are monotone, then public beliefs cannot enter the interval $(\frac{1}{3}, \frac{2}{3})$ after any history. This observation casts doubt on the plausibility of monotone transition functions, as it is very hard to believe that we live in a world which never allows us to have moderate public beliefs about the unknown.[8] Therefore we must at least worry about non-monotone learning and anti-imitation most of the times.

We then provide a sufficient condition for non-monotone learning and anti-imitation based on our findings in the binary-signal setting. There we have shown that agent 3's posterior belief is anti-imitative of agent 1 when the fraction of amateurs is sufficiently large. It suggests that learning is probably non-monotone as well in the continuous model, which then leads to anti-imitation, when private signals are most likely uninformative. Hence we consider distributions of *private beliefs* that have high density around $\frac{1}{2}$, *i.e.*, most agents are almost uninformed, and provide a sufficient condition for each agent

[7]See for example Theorem 3(b) about action convergence in Smith and Sørensen (2000).

[8]For example, in medical research physicians are allowed to offer patients randomization to different treatments only if *clinical equipoise* exists, *i.e.*, there is genuine uncertainty in the expert medical community over whether a treatment will be beneficial. In other words, the community of physicians should regard the treatments as (roughly) equally preferable. See Freedman (1987).

$t+2$ to anti-imitate agent t.[9] Loosely speaking, it requires the distributions to have "thick" tails, in the sense that the conditional expectation of private beliefs larger than $\frac{1}{2}$ is sufficiently far away from $\frac{1}{2}$. The intuition behind this condition is analogous to what we have in the binary-signal setting: high density around $\frac{1}{2}$ corresponds to a large r and thick tails correspond to a large q.[10] If learning is conducted in a society where most people barely have any private knowledge but those who do have some knowledge are sufficiently knowledgeable on average, it is bound to exhibit anti-imitative behavior. Such information structures are reasonable in many contexts of social learning such as technology adoption or development of medical treatments, where most of us really know little but those technicians or physicians are usually well recognized for their expertise.[11]

Another interesting result is a boundary condition on the distributions of private beliefs that guarantees non-monotone learning. We show that the transition functions of public beliefs are decreasing around the boundaries when the distribution of private beliefs are sharply diminishing around the boundaries. As learning completes eventually, public beliefs are bound to be close to the boundaries after sufficiently long histories and then decreasing transition functions lead to non-monotone learning.[12] This condition can be satisfied by a variety of common distributions; being a boundary condition, it can also be "approximately" satisfied by essentially every distribution.[13] On the other hand, in the absence of explicit parametrization of the information structure, we find it hard and most likely intractable to get a necessary and sufficient condition for non-monotone learning in general.

---

[9] For technical convenience we focus on the distributions of *private beliefs* rather than private signals as in Smith and Sørensen (2000).

[10] In the binary-signal model, the conditional expectation of private beliefs larger than $\frac{1}{2}$ is exactly q.

[11] Learning and information aggregation are indeed more desirable when individuals have limited private knowledge. For example, Conley and Udry (2010) investigate the diffusion of a new agricultural technology in Ghana, and find evidence that farmers adjust their inputs to align with those who were surprisingly successful in previous periods, which indicates the presence of social learning. However the input choices for another crop of *known* technology indicate an absence of social learning effects.

[12] We assume unbounded private beliefs in the model to ensure complete learning. See Theorem 1(b) in Smith and Sørensen (2000).

[13] See further discussion in Section 1.3 and Appendix.

1.1.1 Related Literature

The theoretical literature on observational learning has conventionally focused on asymptotic properties and few papers have deliberately studied the dynamics of learning and its behavioral implications short of the limit. This work is closest to Callander and Hörner (2009) as both highlight the impact of information structures on the learning process, which leads to behavior quite different from the canonical predictions. We differ from Callander and Hörner (2009) by maintaining the full observation assumption and considering more general information structures, but as mentioned before, the information structure in their paper shares the intuition of non-monotone learning in our work. On the other hand, Eyster and Rabin (2014) study the impact of observation structures instead and question the rationality of imitative behavior.[14] They provide a necessary and sufficient condition on observation structures for rational anti-imitation, and the logic lies in the fact that rational agents need to take into account the redundancy of previous actions under those structures.[15]

Other work in the literature mainly studies the efficiency of information aggregation and the long-run herd behavior but has substantially extended the first models by Banerjee (1992), Bikhchandani *et al.* (1992), and Smith and Sørensen (2000) in different respects. The assumption of full observation has been relaxed by, for example, Çelen and Kariv (2004) where each agent is only allowed to observe her immediate predecessor. They show that beliefs and actions end up cycling and learning is never complete.[16] Perhaps the most comprehensive generalization in this respect is by Acemoglu *et al.* (2011), which allows the observation structure to be a network topology.[17] They show that asymptotic learning is achieved with unbounded private beliefs when the network topology has *expanding observations*: agents should not be confined to receive information from a bounded subset of other agents. Few recent papers such as Guarino *et al.* (2011) and

[14] An observation structure is a directed network and they consider general networks other than the canonical single-file setting. See also Jackson (2008) for comprehensive discussion on social networks.

[15] Note that their condition on observation structures is necessary and sufficient when the action space is continuous, hence it is still possible to have anti-imitation in the canonical model with binary actions even though their condition is violated.

[16] Incomplete learning and non-converging actions are not due to their observational assumption though. In fact, according to Acemoglu *et al.* (2011), learning is complete (with unbounded private beliefs) in the canonical model even if one can only observe her immediate predecessor.

[17] The network topology in Acemoglu *et al.* (2011) is similar to Eyster and Rabin (2014) but also allows stochastic sampling like in Banerjee and Fudenberg (2004) or Smith and Sørensen (2008).

Herrera and Hörner (2013) alter the observational assumptions in another way by assuming that only certain realization of actions is observable. Yet none of these papers talks about (non-)monotone properties of learning process or anti-imitation.

There is also a growing number of papers that relax the assumption of myopic preferences and introduce payoff interdependence among agents. It is less surprising to see anti-imitation or even contrarian behavior when there are negative externalities such as congestion costs in Eyster *et al.* (2014). But observational learning models with payoff externalities are generally hard to solve due to the existence of forward-looking incentives.[18] Dasgupta (2000) studies social learning in coordination problems and demonstrates that agents exhibit herd behavior as complete imitation under certain information structures. In the context of sequential elections, Ali and Kartik (2012) manage to characterize conditions on the payoff interdependence that will guarantee sincere behavior by agents and eventually a herd.

Most empirical and experimental work on observational learning follows the main focus in the theoretical literature and justifies imitative behavior and herding. For example, Moretti (2011) uses box-office data and finds that the sales of movies with positive and negative opening-weekend surprises in demand diverge over time: a movie that experiences larger sales in Week 1 will experience further increasing sales in subsequent weeks. He considers such imitative behavior by consumers as a result of social learning and further quantifies the effect of social learning on movie sales. Cai *et al.* (2009) conduct a field experiment to distinguish imitative observational learning from salience in which they tell diners either the recently popular dishes or the "feature dishes". They find that diners react more strongly to popularity than to salience, which convincingly suggests that diners imitate. Our work, on the other hand, will potentially raise the question whether the imitative behavior detected by these papers comes from *rational* observational learning.[19]

Herd behavior has also been extensively studied in the sequential trading model of financial market introduced by Glosten and Milgrom (1985), and Park and Sabourian (2011) further investigate the possibility of contrarianism: traders buy (sell) assets after

[18] Besides Eyster *et al.* (2014), a lot of work assumes only backward incentives such as queuing models by Debo *et al.* (2012) and Cripps and Thomas (2014).

[19] In Cai *et al.* (2009), the popularity of dishes is sorted by the actual number of plates sold in the previous week. So roughly speaking, they are really only looking at how agent 2 reacts to agent 1, and we know that there is no anti-imitation due to the overturning principle.

observing histories that reveal bad (good) information about the asset value. Their work builds on the seminal paper by Avery and Zemsky (1998) which shows that there are no informational cascades in such model with informationally efficient prices and it is unlikely to have herding behavior unless signals are "non-monotonic". Park and Sabourian (2011) instead argue that the monotonicity of signals defined by Avery and Zemsky (1998) is disputable and describe conditions on the underlying information structure that are necessary and sufficient for herding or contrarianism. Dasgupta and Prat (2008) bring the sequential trading model together with the reputational herding model established by Scharfstein and Stein (1990) and demonstrate herd behavior when agents have career concerns. However, due to the existence of a competitive market maker who consistently adjusts the bid (ask) prices based on the public histories, all these papers implicitly impose heterogeneous payoff functions of traders.

The remainder of the paper is structured as follows. We begin with the simple model with binary private signals in Section 1.2. Section 1.3 develops some general results for the model with continuous private signals. Section 1.4 concludes.

1.2 A Simple Setting with Binary Private Signals

1.2.1 Setup

We consider a simple variant of the setting by Callander and Hörner (2009). There is an underlying state of the world, $\theta \in \{A, B\}$, whose realization is unknown to the population, a countable set of agents. Agents hold a common prior of θ, $\Pr(\theta = A) = \Pr(\theta = B) = 0.5$. There is an infinite time horizon, $t \in \{1, 2, 3, \ldots\}$, and at each period t an agent is chosen to make a once-in-a-life-time binary decision, $a_t \in \{A, B\}$.[20] Unlike Callander and Hörner (2009), the history of past actions, $\boldsymbol{h}_t \equiv (a_1, a_2, \ldots, a_t)$, is publicly observable for all the future agents. Agents have common payoff functions $u(a_t, \theta) = 1_{\{a_t = \theta\}}$.

There are two informational types of agent. An agent can be an *expert*, who receives a private signal $\sigma_t \in \{A, B\}$, which matches the true state with probability $q \in (\frac{1}{2}, 1)$, before making her decision. Or she can be an *amateur*, who does not receive any private signals. Each agent's type is her private information but the probability of being an

[20] We simply refer the agent acting at period t as agent t throughout the paper.

amateur, $r \in (0,1)$, is the same across all agents and commonly known.

Let $l_{t+1} \equiv l(\boldsymbol{h}_t) \equiv \ln \frac{\Pr(\theta=B|\boldsymbol{h}_t)}{\Pr(\theta=A|\boldsymbol{h}_t)} = \ln \frac{\Pr(\boldsymbol{h}_t|\theta=B)}{\Pr(\boldsymbol{h}_t|\theta=A)}$ be the posterior log-likelihood ratio of agent $t+1$ after observing history $\boldsymbol{h}_t$ but before acquiring her private signal. We also call l_{t+1} the *public belief* as *history* $\boldsymbol{h}_t$ is publicly observed by all future agents, with $l_1 = 0$ as the prior.[21] Let us first solve the Bayesian decision problem of each agent.[22]

Lemma 1.2.1 *Agent* 1 *follows her private signal if she is an expert, and (by assumption) randomly chooses between A and B if she is an amateur. Agent $t \geq 2$ follows her immediate predecessor if she is an amateur or $|l_t| > \ln \frac{q}{1-q}$, and follows her private signal otherwise.*

Proof. See Appendix. ■

An amateur, absent of private signals, will just follow her immediate predecessor, as she knows no more than her immediate predecessor, who has the same preference and acted rationally. On the other hand, an expert will stick to her private signal as long as its precision outweighs the public belief, *i.e.*, $\ln \frac{q}{1-q} \geq |l_t|$, but follows what the public belief suggests otherwise. With individual decision problems solved, we can then characterize the dynamics of $\{l_t\}_{t=1}^{\infty}$ along every *equilibrium* history.[23]

Lemma 1.2.2 *Let $M \equiv \ln \frac{r\cdot\frac{1}{2}+(1-r)\cdot q}{r\cdot\frac{1}{2}+(1-r)\cdot(1-q)}$, $L \equiv \ln \frac{r+(1-r)\cdot q}{r+(1-r)\cdot(1-q)}$, and $H \equiv \ln \frac{q}{1-q}$. Along every equilibrium path $\boldsymbol{h}_\infty \equiv (a_1, a_2, \ldots, a_t, \ldots)$, the public beliefs evolve in the following*

[21] The term *public belief* has been broadly used in the literature, which is the posterior likelihood of a certain state after a history. We use log-likelihood ratio here for technical convenience.

[22] It is not hard to see that in this binary decision problem, each agent's decision rule is unique up to a tie-breaking rule. Here we simply assume an expert follows her private signal when indifferent and the first agent randomly chooses between A and B if she is an amateur. Tie-breaking assumptions are no longer important for the continuous-signal model in Section 1.3.

[23] As all agents are myopic here, we do not emphasize a particular notion of equilibrium, and an equilibrium history is simple a history that is consistent with the individual decision rule described in Lemma 1.2.1. Readers can nevertheless assume what we have in mind is the standard *Bayesian Nash Equilibrium* throughout the paper.

way:

$$
\begin{aligned}
l_1 &= 0;\, l_2 = \begin{cases} M & \textit{if } a_1 = B \\ -M & \textit{if } a_1 = A \end{cases}; \\
\forall t &\geq 2,\, l_{t+1} = \begin{cases} l_t + H & \textit{if } a_t = B \textit{ and } a_{t-1} = A \\ l_t - H & \textit{if } a_t = A \textit{ and } a_{t-1} = B \\ l_t + L & \textit{if } a_t = a_{t-1} = B \textit{ and } l_t \leq H \\ l_t - L & \textit{if } a_t = a_{t-1} = A \textit{ and } l_t \geq -H \\ l_t & \textit{otherwise} \end{cases}.
\end{aligned}
$$

Proof. See Appendix. ■

Except for the first action a_1, each action later in the history reveals a strong piece of evidence, a weak piece of evidence, or no evidence at all.[24] If a_{t+1} is different from a_t, it must come from an expert hence it is a strong piece of evidence, which shifts the public belief by $\pm H$. If a_{t+1} is the same as a_t and $|l_t| \leq H$, it is only meaningful if it comes from an expert, which happens with probability $1 - r$, hence it is a weak piece of evidence and only shifts the public belief by $\pm L$. If a_{t+1} is the same as a_t but $|l_t| > H$, it reveals no more information as both types will follow the immediate predecessor a_t anyway, hence the public belief remains unchanged thereafter, *i.e.*, *Informational Cascade*.[25]

1.2.2 Non-Monotone Learning

To introduce the (non-)monotone property we are interested in this paper, let us set a simple linear order $\precsim$ on $\{A, B\}$ such that $A \precsim B$ (and of course $B \precsim B$). Then we can induce a partial order on each Cartesian product $\{A, B\}^t$, $\forall t \in \mathbb{N}^+$:

$$\forall \boldsymbol{h}_t, \boldsymbol{h}'_t \in \{A, B\}^t,\ \boldsymbol{h}_t \precsim \boldsymbol{h}'_t \text{ if and only if } a_\tau \precsim a'_\tau \text{ for any } \tau \leq t.$$

[24] The first action a_1 comes from either an informative choice by an expert or a random choice by an amateur, and by Bayes rule it turns out to be a "mediocre" piece of evidence compared to later actions.

[25] See Bikhchandani *et al.* (1992).

Definition 1.1 *Learning is **monotone** if*

$$\forall t \in \mathbb{N}^+,\ \forall \text{ two equilibrium histories } \boldsymbol{h}_t \text{ and } \boldsymbol{h}'_t \text{ s.t. } \boldsymbol{h}_t \precsim \boldsymbol{h}'_t,\ l_{t+1} \equiv l(\boldsymbol{h}_t) \leq l(\boldsymbol{h}'_t) \equiv l'_{t+1}.$$

*Learning is **non-monotone** otherwise, or equivalently,*

$$\exists \text{ two equilibrium histories } \boldsymbol{h}_t \text{ and } \boldsymbol{h}'_t \text{ s.t. } \boldsymbol{h}_t \precsim \boldsymbol{h}'_t,\ \boldsymbol{h}_t \neq \boldsymbol{h}'_t \text{ and } l_{t+1} > l'_{t+1}.$$

The non-monotonicity given by Definition 1.1 is straightforward. Suppose some actions A(s) in an equilibrium history $\boldsymbol{h}_t$ are switched to B(s) while the other actions remain unchanged, and let us compare the posterior belief of agent $t+1$ after observing this new history $\boldsymbol{h}'_t$ with the original posterior belief after observing $\boldsymbol{h}_t$.[26] If the new posterior belief is smaller, *i.e.*, agent $t+1$ is less confident in state B although more of her predecessors showed confidence in state A, we say learning is non-monotone.

Clearly the canonical restaurant-choice example discussed in the introduction does not exhibit such non-monotonicity, however we will show that learning can be non-monotone under the setting introduced in Section 1.2. In fact, we focus on some particular cases where only *one* agent's action is altered in the history.

Definition 1.2 *Consider two equilibrium histories $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$ that differ only in one action:*

$$\exists \tau \leq t \text{ s.t. } a_\tau = A,\ a'_\tau = B;\ \forall \tau' \leq t \text{ and } \tau' \neq \tau,\ a_{\tau'} = a'_{\tau'}.$$

*We say the posterior belief of agent $t+1$ is **anti-imitative** of agent τ under the pair $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$ if $l_{t+1} > l'_{t+1}$.*

*We say the posterior belief of agent $t+1$ is **always anti-imitative** of agent τ if $l_{t+1} > l'_{t+1}$ for* every *such pair $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$.*

Anti-imitative posterior belief refers to a special case of non-monotone learning: fixing the actions of all the predecessors of agent $t+1$ other than agent τ, agent $t+1$ becomes less confident in state B though agent τ's action reveals more confidence in state B. In other words, it is as if the posterior belief of agent $t+1$ "anti-imitates" agent τ's action.

[26] Definition 1 requires the new history $\boldsymbol{h}'_t$ to be an equilibrium history as well, otherwise posterior beliefs are not well-defined after off-equilibrium histories and the comparison becomes meaningless.

Now let us go back to the model and investigate the possibility of non-monotone learning, particularly anti-imitative posterior belief. We consider the case with $\tau = 1$, *i.e.*, only agent 1's action is altered between two histories[27]. Obviously by Lemma 1.2.2 agent 2's posterior belief cannot be anti-imitative of agent 1, hence we start the analysis from agent 3 onwards. For convenience we present most results here using H, M, and L defined in Lemma 1.2.2.

Proposition 1.2.1 *1. The posterior belief of agent* 3 *is always anti-imitative of agent* 1 *if and only if*

$$H > 2M + L, \text{ or equivalently, } \frac{r^3}{(1-r)^2(1+r)} > 4q(1-q).$$

2. The posterior belief of each agent $t \geq 3$ *is always anti-imitative of agent* 1 *if and only if*

$$H > 2M + L \text{ and } H > 2M + (k^* - 2)L \text{ where } k^* \equiv \min\{k \in \mathbb{N}^+ | M + kL > H\}.$$

Proof. See Appendix. ■

The logic behind the first result in Proposition 1.2.1 is as follows. According to Definition 1.2, agent 3 needs to compare two histories, $\boldsymbol{h}_2 = (A, B)$ and $\boldsymbol{h}'_2 = (B, B)$.[28] History (A, B) reveals one *mediocre* evidence against state B from the first action A but one *strong* evidence in favor of state B from the second action B, and by Lemma 1.2.2 the posterior belief after (A, B) can be precisely calculated as $l_3 = -M + H$. On the other hand, history (B, B) reveals one *mediocre* and one *weak* evidence in favor of B from the first and second action B respectively, and by Lemma 1.2.2 the posterior belief after (B, B) is $l'_3 = M + L$. To make the posterior belief of agent 3 anti-imitative, we need $l_3 > l'_3$ or simply $H > 2M + L$. It is not hard to see that this inequality can be satisfied by some pair (r, q) such that r or q is sufficiently high, *i.e.*, either the weak/mediocre piece

[27] With two informational types (*expert* and *amateur*), anti-imitative beliefs can only appear with respect to the first agent's action. Other formats of anti-imitative beliefs and non-monotone learning can appear in either binary-signal settings with more informational types or the continuous-signal setting in Section 1.3. Nevertheless, the simple setting here suffices to capture the intuition about how heterogeneous informational types can lead to anti-imitative beliefs.

[28] The other comparison between (A, A) and (B, A) is symmetric.

of evidence is fairly insignificant or the strong piece of evidence is sufficiently dominant.[29]

The second result in Proposition 1.2.1 builds on the first one. Suppose the posterior belief of agent 3 is already anti-imitative, *i.e.*, $l((A,B)) > l((B,B))$. By Lemma 1.2.2, extending the two histories by one *same* action will shift the posterior belief toward the *same* direction and by the *same* amount, $\pm H$ or $\pm L$, therefore the inequality still holds. This logic works for further extensions as well, *until* the posterior belief after one history, say l_t, grows beyond the precision of private signals, *i.e.*, $|l_t| > \ln \frac{q}{1-q}$. Since then l_t no longer changes and l'_t will catch up until $|l'_t| > \ln \frac{q}{1-q}$ as well. The extra condition in the second result precisely takes account of this subtle difference and assures that l'_t will not exceed l_t in the end.

We have shown that, contrast to the canonical model and what people might have learned from the existing literature, learning is indeed non-monotone in this simple setting with heterogeneous informed agents. More surprisingly, all the successors of agent 1 except agent 2 will form beliefs that are anti-imitative of agent 1, even though they all have the same preference as agent 1. Nevertheless, we do want to point out that the second result of Proposition 1.2.1 does not contradict with the existing results on the long-run behavior of agents, *e.g.*, an informational cascade occurs eventually or a herd arises eventually.[30] In fact, anti-imitative *belief* does not necessarily lead to anti-imitative *behavior*, especially when the private signals are discrete.

Definition 1.3 *Consider two equilibrium histories $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$ that differ only in one action:*

$$\exists \tau \leq t \text{ s.t. } a_\tau = A,\ a'_\tau = B;\ \forall \tau' \leq t \text{ and } \tau' \neq \tau,\ a_{\tau'} = a'_{\tau'}.$$

*We say agent $t+1$ **anti-imitates** agent τ under the pair $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$ if*

$$\Pr(a_{t+1} = B|\boldsymbol{h}_t) > \Pr(a_{t+1} = B|\boldsymbol{h}'_t).$$

[29] In Appendix, we indeed show that agent 3's posterior belief is anti-imitative of agent 1 when $r \geq \frac{\sqrt{5}-1}{2}$, for any $q \in (\frac{1}{2}, 1)$.

[30] See Proposition 1 in Bikhchandani *et al.* (1992) or Theorem 3(a) in Smith and Sørensen (2000).

We say agent $t+1$ ***always anti-imitates*** *agent* τ *if*

$$\Pr(a_{t+1} = B|\boldsymbol{h}_t) > \Pr(a_{t+1} = B|\boldsymbol{h}'_t) \textit{ for } \text{every} \textit{ such pair } (\boldsymbol{h}_t, \boldsymbol{h}'_t).$$

This definition of anti-imitative behavior is in the same spirit of *anti-imitation* defined in Eyster and Rabin (2014), except that we use probabilistic criterion here on account of the binary action space.[31] To see the difference between anti-imitative belief and anti-imitative behavior, let us reconsider agent 3 after history (A, B) or (B, B). Suppose the posterior belief of agent 3 is already anti-imitative, *i.e.*, $l((A, B)) > l((B, B))$ or $H > 2M + L$. If agent 3 is an amateur, she will choose B after both histories according to Lemma 1.2.1. If agent 3 is an expert, she will follow his private signal after both histories according to Lemma 1.2.1, since $l((A, B)) = H - M < H$ and $l((B, B)) = M + L < H$. Therefore, although agent 3's posterior belief is anti-imitative, there is *no* probabilistic difference in her behavior and she does *not* anti-imitate agent 1.

Corollary 1.2.1 *Under any pair of equilibrium histories* $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$ *that differ only in the first action, at most* $(k^* - 3)$ *agents anti-imitate agent* 1*, where* k^* *has been defined in Proposition* 1.2.1*.*

Proof. See Appendix. ■

Since amateurs always follow their immediate predecessors, only the behavior of experts could be different after two histories that differ in the first action. And according to Lemma 1.2.2, that happens only when the posterior belief after one history, say l_t, has exceeded the precision of private signals while the posterior belief after the other, say l'_t, has not. In that situation an expert will follow her immediate predecessor after l_t but follow her private signal after l'_t. However, as we discussed earlier, l'_t will catch up along the history and eventually exceeds the threshold as well. Hence such behavioral difference can only exist for a finite number of future agents, and the upper bound of that number is precisely $(k^* - 3)$. It is worth noting that Corollary 1.2.1 is indeed consistent with the long-run behavior predicted by the existing literature.

[31] See Definition 5 in Eyster and Rabin (2014) for a comparison. They assume continuous action space as in Lee (1993) and hence define anti-imitation as one agent's action being decreasing in some action(s) she observes.

In the next section we will turn to a model with continuous private signals, in order to better understand the impact of underlying information structures on the learning process. Besides, as we will see, the continuous model gets rid of the subtle difference between anti-imitative belief and anti-imitative behavior and thus allows a clear focus.

1.3 A Model with Continuous Private Signals

1.3.1 Setup

We consider the observational learning model with continuous private signals due to Smith and Sørensen (2000). There is an underlying state of the world, $\theta \in \{A, B\}$, whose realization is unknown. Agents hold a common prior of θ, $\Pr(\theta = A) = \Pr(\theta = B) = 0.5$. Agents move sequentially over an infinite time horizon, $t \in \{1, 2, 3, \ldots\}$, and each agent t makes a once-in-a-life-time binary decision, $a_t \in \{A, B\}$, with payoff $u(a_t, \theta) = 1_{\{a_t=\theta\}}$.The history of past actions, $\boldsymbol{h}_t \equiv (a_1, a_2, \ldots, a_t)$, is publicly observable for all the future agents. Besides, each agent at period t will receive a private signal $\sigma_t \in [\underline{c}, \overline{c}]$ before making her choice. Conditional on θ, $\{\sigma_t\}_{t=1}^{\infty}$ are independently and identically distributed across t.

Following Smith and Sørensen (2000), we work directly with $q_t \equiv \Pr(\theta = B|\sigma_t) \in [0, 1]$, the *private belief* of agent t after observing her private signal σ_t. The reason for such normalization is that what matters for each agent is the information generated by her private signal rather than the realization of private signal itself. Let $G_\theta(x)$ be the cumulative distribution function of q_t conditional on θ. $G_A(x)$ and $G_B(x)$ capture the information structure, and if both are *differentiable* with density function $g_A(x)$ and $g_B(x)$ respectively, the unconditional density function of q_t, $g(x) \equiv \frac{g_A(x)+g_B(x)}{2}$ is already sufficient. In particular, Bayesian updating implies $\frac{g_A(x)}{g_B(x)} = \frac{x}{1-x}$ and thus $g_A(x) = 2xg(x)$ and $g_B(x) = 2(1 - x)g(x)$.[32] We will use $g(x)$ in most results when convenient.

We impose the following assumptions on $g(x)$:

1. *Full support*: $g(x)$ is strictly positive on $[0, 1]$.

[32] The private belief structure is commonly used in the literature. Curious readers can look at Appendix A in Smith and Sørensen (2000) and Section 3.A in Smith *et al.* (2012) for the justification of it.

2. *Differentiability*: $g(x)$ is continuously twice differentiable on $(0,1)$; both $\lim_{x\to 0+} g''(x)$ and $\lim_{x\to 1-} g''(x)$ exist.

3. *Symmetry*: $g(x) = g(1-x)$, $\forall x \in [0,1]$.

The assumption of full support implies that private beliefs are *unbounded*, which guarantees that no history of actions is off equilibrium, as each agent has strictly positive probability to follow her private signal no matter what the history was. Hence we are free to compare any pair of histories when investigating whether learning is monotone or not. The assumption of twice differentiability allows us to conduct certain mathematical analysis later.[33] The assumption of symmetry, like the binary-signal model in Section 1.2, reduce the number of comparisons (by half) we need to make as we can focus on the pattern of histories rather than every particular realization.

1.3.2 A Necessary Condition for Non-Monotone Learning

We start by solving the Bayesian decision problem of each agent as well as characterizing the dynamics of *public beliefs*, $\{l_t\}_{t=1}^{\infty}$. Recall again that public belief after a history $\boldsymbol{h}_t$ is exactly the posterior belief of agent $t+1$ before her getting her private belief. For convenience, here we use posterior probability that state B is the true state rather than posterior log-likelihood ratio used in Section 1.2, *i.e.*, $l_{t+1} \equiv \Pr(\theta = B|\boldsymbol{h}_t)$ with $l_1 = \frac{1}{2}$.

Lemma 1.3.1 *Agent t who forms public belief l_t from history and private belief q_t from her signal will choose $a_t = B$ if and only if $q_t \geq 1 - l_t$.*

Therefore the stochastic process $\{l_t\}_{t=1}^{\infty}$ in equilibrium is characterized by the following transition functions:

$$l_{t+1} = \begin{cases} m(l_t) \equiv \frac{l_t G_B(1-l_t)}{l_t G_B(1-l_t)+(1-l_t)G_A(1-l_t)} & \text{if } a_t = A \\ n(l_t) \equiv \frac{l_t[1-G_B(1-l_t)]}{l_t[1-G_B(1-l_t)]+(1-l_t)[1-G_A(1-l_t)]} & \text{if } a_t = B \end{cases}.$$

Moreover, $m(l_t) = 1 - n(1-l_t)$.

Proof. See Appendix. ■

[33] In fact $k(x)$ being continuous differentiable is sufficient for most of the analysis. We only need twice differentiability in the proof of Proposition 1.3.3 that needs higher order Taylor expansion, but the proposition itself does not explicitly involve $k''(x)$.

Similar to the binary-signal model in Section 1.2, what rational agents are really doing here is comparing the public belief generated from the observed history with the private belief generated from their private signals. For instance, agent t will choose B under one of the following three circumstances:

> both the public belief and her private belief supports B, *i.e.*, $q_t > \frac{1}{2}$ and $l_t > \frac{1}{2}$;
>
> her private belief supports B and is stronger than the public belief which supports A, *i.e.*, $q_t > \frac{1}{2}$, $l_t < \frac{1}{2}$ but $q_t \geq 1 - l_t$;
>
> the public belief supports B and is stronger than her private belief which supports A, *i.e.*, $q_t < \frac{1}{2}$, $l_t > \frac{1}{2}$ but $l_t \geq 1 - q_t$.

Clearly these can be summarized as just $q_t \geq 1 - l_t$. The dynamics of public beliefs then simply comes from the individual decision rule and Bayes rule. In particular, the transition function, $m(\cdot)$ or $n(\cdot)$, captures how the public beliefs evolve from one period to the next, when the most recent action is A or B.

We may well start to answer the main question: Can learning be non-monotone? Let us stick to Definition 1.1, 1.2, and 1.3, respectively, for *(non-)monotone learning*, *anti-imitative belief* and *anti-imitative behavior*.[34] Recall that there are no off-equilibrium histories with unbounded private beliefs, which can be easily seen from Lemma 1.3.1, hence these definitions apply to every possible pair of histories. We first provide a *necessary* condition for non-monotone learning.

Lemma 1.3.2 (*Overturning Principle*) *$l_{t+1} > \frac{1}{2}$ if $a_t = B$ and $l_{t+1} < \frac{1}{2}$ if $a_t = A$.*

Proof. See Appendix or the proof of Theorem 3 in Smith and Sørensen (2000). ■

Proposition 1.3.1 *Learning is non-monotone* only if *the transition function $m(\cdot)$ or $n(\cdot)$ is non-monotone. Or equivalently,*

$$\exists x \in (0,1) \ \ s.t. \ \frac{g_B(1-x)}{G_B(1-x)} - \frac{g_A(1-x)}{G_A(1-x)} > \frac{1}{(1-x)x} \ \ or \ \frac{g_B(x)}{G_B(x)} - \frac{g_A(x)}{G_A(x)} > \frac{1}{x(1-x)}.$$

[34] As posterior log-likelihood ratio is a strictly monotone transformation of posterior probability, there is no need to provide redundant definitions here.

Proof. See Appendix. ■

Quite intuitively, non-monotone learning requires non-monotone transitions of posterior beliefs. To see it clearly, let us think about two different histories $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$ such that $\boldsymbol{h}_t \precsim \boldsymbol{h}'_t$. If $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$ differ on the last action, *i.e.*, $a_t = A$ and $a'_t = B$, then clearly $l_{t+1} < l'_{t+1}$ by the overturning principle. If $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$ only differ up to some earlier period $\tau < t$, then $l_{\tau+1} < l'_{\tau+1}$, again by the overturning principle. Since after period τ actions are identical between the two histories, agents should update their beliefs using the same transition function from period $\tau + 1$ on along the two histories. If both $m(\cdot)$ and $n(\cdot)$ are monotone, then the inequality "$l \leq l'$" will be preserved from period $\tau + 1$ on as well, and hence non-monotone learning, *i.e.*, $l_{t+1} > l'_{t+1}$, is not possible.

We automatically have the following corollary that gives the necessary condition in terms of $g(x)$.

Corollary 1.3.1 *Learning is non-monotone* only if $g(x)$*, the unconditional density function of private beliefs, is such that*

$$\exists x \in (0,1) \quad s.t. \quad g(x)[\frac{(1-x)}{\int_x^1 (1-s)g(s)ds} - \frac{x}{\int_x^1 sg(s)ds}] > \frac{1}{x(1-x)}$$
$$or\ g(x)[\frac{x}{\int_0^x sg(s)ds} - \frac{1-x}{\int_0^x (1-s)g(s)ds}] > \frac{1}{x(1-x)}.$$

Proof. Simply apply $g_A(x) = 2xg(x)$ and $g_B(x) = 2(1-x)g(x)$. ■

Numerical calculations can easily provide some simple *symmetric* distributions on $[0,1]$ that satisfy this necessary condition.

1. *Linear Density*: $g(x) = (4 - 8\alpha)\left|x - \frac{1}{2}\right| + 2\alpha$, $\forall \alpha \in (\frac{5}{7}, 1)$;

2. *Quadratic Density*: $g(x) = \alpha(x - \frac{1}{2})^2 + (1 - \frac{\alpha}{12})$, $\forall \alpha \in (-6, -2)$;

3. *Beta Distribution*: $g(x) = \frac{x^{\alpha-1}(1-x)^{\alpha-1}}{B(\alpha,\alpha)}$ with $B(\alpha,\alpha) \equiv \int_0^1 s^{\alpha-1}(1-s)^{\alpha-1}ds$, for sufficiently large α.[35]

[35] We get $\alpha > 28$ roughly, using Kumaraswamy's distribution as in Jones (2009) to approximate Beta distribution.

With the complicated expression in Corollary 1.3.1, it is probably hard to see whether the necessary condition is likely to be satisfied or not. Nevertheless we make the following observation that the necessary condition is violated only by information structures that always generate sufficiently strong public beliefs along the learning process.

Claim 1.3.1 *If transition functions* $n(\cdot)$ *and* $m(\cdot)$ *are monotone, then in equilibrium*

$$l_t \notin (\frac{1}{3}, \frac{2}{3}), \ \forall t \geq 2.$$

Proof. See Appendix. ■

When transition functions are monotone, we have a "stronger" overturning principle (compared to Lemma 1.3.2): $l_{t+1} \geq \frac{2}{3}$ if $a_t = B$ and $l_{t+1} \leq \frac{1}{3}$ if $a_t = A$.[36] However it seems extraordinary that public beliefs cannot enter the interval $(\frac{1}{3}, \frac{2}{3})$ after any history. Imagine that agents later in the sequence observe a very long history of alternating A's and B's, $(A, B, A, B, \ldots, A, B)$. If the public belief is very close to $\frac{1}{2}$, which is what we would naturally expect after observing a large but equal number of A's and B's, then the transition functions must be non-monotone. In fact it is very hard to believe that the underlying information structure never allow agents to have moderate public beliefs throughout the learning process. Hence the plausibility of monotone transition functions is questionable in general, which suggests that we must at least worry about learning being non-monotone most of the times.

Yet the necessary condition does not guarantee non-monotone learning or anti-imitation. The distribution of private beliefs determines not only transition functions $m(\cdot)$ and $n(\cdot)$, but at the mean time also the possible posterior beliefs that can be generated by different histories. Hence, the selection of distributions is in general hard because we have to make sure not only that transition functions are non-monotone, but also that such non-monotonicity is relevant in equilibrium. If transition functions are decreasing only over a subset of $[0, 1]$ that is, roughly speaking, never "entered" or "passed through" by posterior beliefs generated by any history, learning is still monotone and agents do not

[36] These bounds depend on the technical assumptions on the information structure. For example, if $k(0) = 0$ and $k'(0) > 0$, the interval becomes $(\frac{2}{5}, \frac{3}{5})$. However what matters here is that the public beliefs are always sufficiently bounded away from $\frac{1}{2}$ when the transition functions are monotone.

anti-imitate. With this concern in mind, we make further observations and provide a few interesting *sufficient* conditions.

1.3.3 Two Sufficient Conditions for Non-Monotone Learning

We begin with a sufficient condition that is somewhat analogous to what we have seen in Section 1.2. Recall that in the simple model with binary private signals, the posterior belief of agent 3 is always anti-imitative of agent 1 when the fraction of uninformed agents is sufficiently large, but agent 3 never anti-imitates agent 1. Back to the current model with continuous private signals/beliefs, we first claim the equivalence between anti-imitative beliefs and anti-imitative behavior and then provide a similar condition on the continuous distribution of private beliefs that guarantees anti-imitation by agent 3.

Claim 1.3.2 *Let* $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$ *be a pair of histories such that*

$$\exists \tau \leq t \text{ s.t. } a_\tau = A,\ a'_\tau = B;\ \forall \tau' \leq t \text{ and } \tau' \neq \tau,\ a_{\tau'} = a'_{\tau'}.$$

Agent $t+1$ *anti-imitates agent* $\tau \leq t$ *under* $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$ if and only if *the posterior belief of agent* $t+1$ *is anti-imitative of agent* τ *under* $(\boldsymbol{h}_t, \boldsymbol{h}'_t)$. *In other words, Definition* 1.2 *and* 1.3 *are equivalent.*

Proof. See Appendix. ■

Anti-imitative beliefs implies that, according to Lemma 1.3.1, the threshold in the decision rule of agent $t+1$ is lower after $\boldsymbol{h}_t$ than after $\boldsymbol{h}'_t$, *i.e.*, $l_{t+1} > l'_{t+1} \Longrightarrow 1 - l_{t+1} < 1 - l'_{t+1}$. When private belief q_{t+1} is unbounded and continuously distributed, lower threshold always leads to higher probability of choosing B, hence agent $t+1$ also exhibits anti-imitative behavior.

Proposition 1.3.2 *Consider a sequence of continuously twice differentiable and symmetric density functions on* $[0,1]$, $\{g^s(\cdot)\}_{s=1}^{\infty}$. *Let* Z_s *be a random variable on* $[0,1]$ *that is distributed according to* g^s *and let* $\varepsilon_s \equiv E[Z_s | Z_s \geq \frac{1}{2}] - \frac{1}{2}$.[37]

[37] ε_s is essentially the *first absolute central moment* of the distribution g^s.

Let $X_{\frac{1}{2}}$ be an almost surely constant random variable with sole realization $\frac{1}{2}$. Suppose the sequence $\{g^s(\cdot)\}_{s=1}^{\infty}$ and $\{Z_s\}_{s=1}^{\infty}$ are such at

1. $Z_s \xrightarrow{d} X_{\frac{1}{2}}$ *as* $s \to \infty$*;*
2. $\exists S \in \mathbb{N}^+$ *s.t.* $\forall s > S$, $g^s(\dfrac{1}{2}) \cdot \varepsilon_s \geq \dfrac{1}{2}$.

Then $\forall t \geq 1$, $\exists S_t \in \mathbb{N}^+$ s.t. $\forall s > S_t$, agent $t+2$ always anti-imitates agent t in equilibrium when the unconditional density function of private beliefs is $g^s(\cdot)$.

Proof. See Appendix. ■

The conditions in Proposition 1.2.1 for the binary-signal model, particularly that r is sufficiently large, suggest that we might also have anti-imitation in the continuous model when the distribution of private beliefs is heavily centered around $\frac{1}{2}$, *i.e.*, most agents are more or less uninformed. Indeed, Proposition 1.3.2 says that agent $t+2$ always anti-imitates agent t under some heavily centered distributions that satisfies an extra condition. Roughly speaking, this extra condition requires "thick" tails of the probability density functions that converge to $\delta_{\frac{1}{2}}(\cdot)$, in the sense that the conditional expectation of private beliefs on the left(right) of $\frac{1}{2}$ should move to $\frac{1}{2}$ somehow "slower" than the increase of the density at $\frac{1}{2}$. In fact, to make an analogy, the binary-signal model does satisfy this condition in a particular way: the conditional expectation of private beliefs on the left or right of $\frac{1}{2}$ is $1-q$ or q, which is always bounded away from $\frac{1}{2}$.

Mathematically, what this condition really does is to make sure that the transition functions are decreasing around $\frac{1}{2}$. Think about the comparison between history $\boldsymbol{h}_2 = (A, B)$ and $\boldsymbol{h}'_2 = (B, B)$. After the first action, $l_2 < \frac{1}{2} < l'_2$ but both l_2 and l'_2 are close to $\frac{1}{2}$ when most agents are almost uninformed. If transition function $n(\cdot)$ is decreasing around $\frac{1}{2}$ while l_2 and l'_2 are so close to $\frac{1}{2}$ that both are within the decreasing region of $n(\cdot)$, we have $l_3 = n(l_2) > n(l'_2) = l'_3$. Hence agent 3 anti-imitates agent 1.

Let us give an example in order to better illustrate this condition. Transition function $n(\cdot)$ is not decreasing around $\frac{1}{2}$ with the following sequence of (truncated) Normal distributions,

$$g_N^s(x) \equiv \frac{s}{2\int_0^{s/2} e^{-t^2} dt} \cdot \exp[-s^2(x - \frac{1}{2})^2];$$

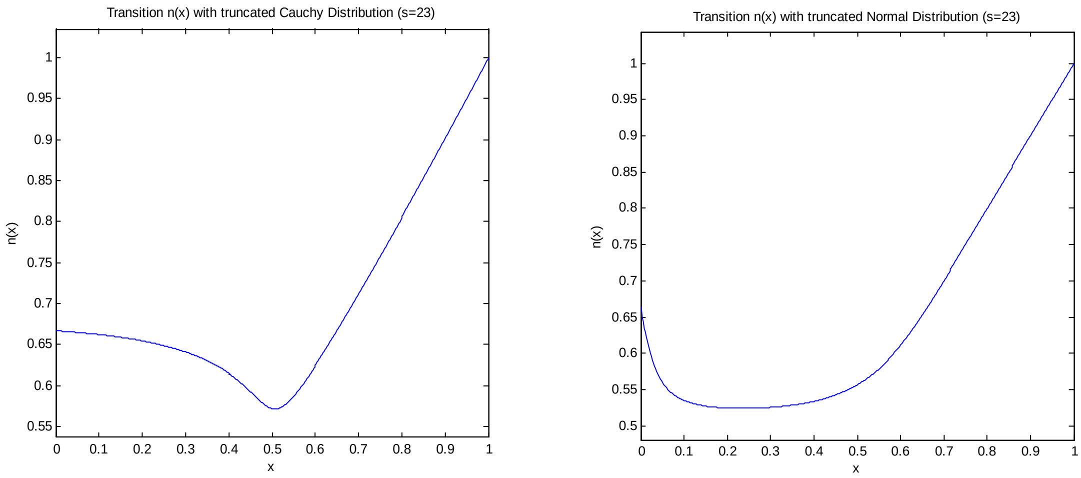

Figure 1: Examples of Proposition 1.3.2.

but it is with the following sequence of (truncated) Cauchy distributions

$$g_C^s(x) \equiv \frac{s}{2\arctan(\frac{s}{2})} \cdot \frac{1}{1+s^2(x-\frac{1}{2})^2},$$

which are known as heavy-tailed distributions.[38]

We like to point out that Proposition 1.3.2 does not require an explicit functional form of $g^s(\cdot)$, but it does not tell how close to $\delta_{\frac{1}{2}}(\cdot)$ the thick-tailed distribution needs to be either. If we are looking at a particular class of distributions with explicit parameters, we can always have a precise condition by using Lemma 1.3.1:

$$\text{agent 3 always anti-imitates agent 1 } \textit{if and only if } n(m(\frac{1}{2})) > n(n(\frac{1}{2})).$$

Take the sequence of (modified) Cauchy distributions $\{g_C^s(\cdot)\}_{s=1}^{\infty}$ for example. Numerical calculation yields that agent 3 always anti-imitate agent 1 when $s > 22.1$.

Careful readers probably have noticed in the figure above that, although the transition function $n(\cdot)$ associated with (truncated) Normal distributions is not decreasing around $\frac{1}{2}$, it is decreasing around 0. The fact that some distributions of private beliefs generate transitions functions that are decreasing around 0 or 1 is what drives the next sufficient

[38] See Johnson and Kotz (1982). Note that the distributions here are not exactly Normal and Cauchy distributions. Since the support is $[0, 1]$ rather than $\mathbb{R}$, we have to use truncated distribution.

condition. We first recall the famous result by Smith and Sørensen (2000) that learning is complete eventually with unbounded private beliefs.

Lemma 1.3.3 (*Complete Learning*) *With unbounded private beliefs,* $l_t \longrightarrow 1_{\{\theta=B\}}$ *almost surely as* $t \longrightarrow +\infty$.

Proof. See Theorem 1(b) in Smith and Sørensen (2000). ■

Complete learning implies that public beliefs will be sufficiently close to either 0 or 1 after a sufficiently long history. If the transition functions are decreasing around 0 and 1, we can be certain that eventually public beliefs will reach those decreasing regions, which then erases the "gap" between non-monotone learning and non-monotone transitions discussed earlier. Therefore we have the following sufficient condition for non-monotone learning.

Proposition 1.3.3 *Learning is non-monotone* if $\frac{\lim_{x\to 0+} g'(x)}{g(0)} > 3$.

Proof. See Appendix. ■

When the density function $g(\cdot)$ is diminishing sufficiently fast at the boundaries of $[0, 1]$, the transition functions are decreasing around 0 or 1. In particular, $m(\cdot)$ is decreasing around 1 and $n(\cdot)$ is decreasing around 0, *i.e.*, $n(\cdot)$ is decreasing over $(0, \varepsilon)$ for some small ε and symmetrically $m(\cdot)$ is decreasing over $(1 - \varepsilon, 1)$.

Let us compare the posterior beliefs after any two histories, $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$, such that $\boldsymbol{h}_t \precsim \boldsymbol{h}'_t$ and $\boldsymbol{h}_t \neq \boldsymbol{h}'_t$. Generically the beliefs are different as well, $l_{t+1} \neq l'_{t+1}$.[39] Complete learning implies that both l_{t+1} and l'_{t+1} are within either $(0, \varepsilon)$ or $(1 - \varepsilon, 1)$ when t is sufficiently large. Suppose the true state is A and hence the former is the case. If $l_{t+1} > l'_{t+1}$, learning is non-monotone by definition. If $l_{t+1} < l'_{t+1}$, we can simply extend $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$ by an action B, *i.e.*, $\boldsymbol{h}_{t+1} = (\boldsymbol{h}_t, B)$ and $\boldsymbol{h}'_{t+1} = (\boldsymbol{h}'_t, B)$. But now $l_{t+2} = n(l_{t+1}) > n(l'_{t+1}) = l'_{t+2}$, so learning is non-monotone. Notice that when $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$ only differ in one action, we effectively get anti-imitation by either agent $t + 1$ or agent $t + 2$.

The sufficient condition in Proposition 1.3.3 can be satisfied by a variety of (truncated) common distributions on $[0, 1]$:

[39] For a generic density function $k(\cdot)$, we can always construct two different histories that generate different posterior beliefs.

1. *Linear Density*: $g(x) = (4 - 8\alpha)\left|x - \frac{1}{2}\right| + 2\alpha$, $\forall \alpha \in (\frac{5}{7}, 1)$;

2. *Quadratic Density*: $g(x) = \alpha(x - \frac{1}{2})^2 + (1 - \frac{\alpha}{12})$, $\forall \alpha \in (-6, -2)$;

3. (*Truncated*) *Double Exponential Distribution*: $g(x) = \frac{\exp(-\frac{\left|x-\frac{1}{2}\right|}{\alpha})}{2\alpha(1-\exp(-\frac{1}{2\alpha}))}$, $\forall \alpha \in (0, \frac{1}{3})$;

4. (*Truncated*) *Normal Distribution*: $g(x) = \frac{\alpha}{2\int_0^{\alpha/2} e^{-t^2} dt} \cdot \exp[-\alpha^2(x - \frac{1}{2})^2]$, $\forall \alpha > 2\sqrt{3}$;

5. (*Truncated*) *Cauchy Distribution*: $g(x) = \frac{\alpha}{2\arctan(\frac{\alpha}{2})} \cdot \frac{1}{1+\alpha^2(x-\frac{1}{2})^2}$, $\forall \alpha > 2\sqrt{3}$.[40]

In fact, as a boundary condition, it can be "approximately" satisfied by essentially every symmetric and twice differentiable density function with full support $[0, 1]$, and we provide a claim about this in Appendix.

We have characterized two circumstances where non-monotone transition functions become sufficient for non-monotone learning: transitions function are decreasing around $\frac{1}{2}$ or around $0(1)$. Unfortunately it is not very clear to us what will happen if the transition functions are decreasing somewhere else in general. We do want to emphasize though, that Lemma 1.3.1 gives an explicit algorithm to calculate public beliefs and hence a precise condition for non-monotone learning is always achievable, at least numerically, once we restrict attention on certain classes of distributions with parameters. Nevertheless, it is hard to establish a precise condition that, like Proposition 1.3.2 and 1.3.3, applies to a general distribution without any parametrization.

1.4 Conclusion

In this paper we reconsider the standard observational learning models where agents act myopically and share common preferences. We study how underlying information structures affect the evolution of posterior beliefs as well as the behavior of rational agents. We show that learning is not always monotone in these models: rational agents often form posterior beliefs that are non-monotone with respect to the actions they observe. As a result, alongside the long-run herd behavior that has been well established by the

[40]For (truncated) Cauchy distribution with $\alpha > 14.3$, we could indeed have very large decreasing regions, *i.e.*, $n(\cdot)$ is decreasing over $(0, \frac{1}{2})$ and $m(\cdot)$ is decreasing over $(\frac{1}{2}, 1)$, which have also been captured by the earlier graph. In that case learning is non-monotone even under a pair of very short histories.

existing literature, it is rational for agents to anti-imitate some of the predecessors during the learning process.

We first look into a simple model with binary private signals and two informational types of agents, *i.e.*, agents are either uninformed or informed by a private signal with certain precision. We find that learning is non-monotone when the probability of being uninformed is sufficiently large or the private signals are sufficiently precise. In particular, under such information structures, the third agent as well as each agent after her always form posterior beliefs that are non-monotone with respect to the action of the first agent. Consequently, some of them anti-imitate the first agent: *ceteris paribus*, they are more likely to choose one action when the first agent has switched to the opposite, even in the absence of any strategic concern or preference heterogeneity.

Next we investigate the observational-learning model with continuous private signals that has been extensively studied in the literature since Smith and Sørensen (2000). We provide an intuitive necessary condition for non-monotone learning and anti-imitation: transitions of public beliefs need to be non-monotone. And we argue that the necessary condition is likely to be satisfied as any information structure violating it never generates moderate public beliefs during learning. Then we make further observations on when this necessary condition could become sufficient and obtain two sufficient conditions. We find that when the transition functions are non-monotone over certain subsets of the unit interval, such as around the middle or near the two boundaries, learning is non-monotone and hence some agents anti-imitate their predecessor(s). Though we are still in search for a general necessary and sufficient condition, we do have a complete characterization of the learning process so it is possible to get a precise condition on parameter values if we restrict attention to information structures captured by certain classes of distributions. However non-parametric results other than what we have presented are in general quite hard to get and further work out of this paper is most welcome.

We treat this paper as an interesting contribution to the literature of observational learning, where most work has been done on the efficiency of information aggregation and the long-run behavior of agents. Our work shows that, in the short run, rational agents may indeed act quite differently from what we might expect based on the asymptotic outcomes we already know. It is clearly important as well to understand the behavioral implications of rational learning short of the limit, a lot of which we are still not clear about or at least cannot simply induce from the existing results.

1.5 Appendix

1.5.1 Omitted Proofs

Proof of Lemma 1.2.1**.** Agent 1's decision rule is trivial. Now without loss of generality let us assume $a_t = A$ and consider agent $t+1$. If agent $t+1$ is an amateur, she will find A optimal as well since she holds no more information than agent t, who has the same preference and rationally chose A. In other words, agent $t+1$'s posterior, before receiving her private signal if she is an expert, should be (weakly) in favor of A:

$$l_{t+1} = \ln \frac{\Pr(\boldsymbol{h}_t = (a_1, a_2, \ldots, a_t = A)|\theta = B)}{\Pr(\boldsymbol{h}_t = (a_1, a_2, \ldots, a_t = A)|\theta = A)} \leq 0.$$

If agent $t+1$ is an expert with private signal σ_{t+1}, she will update her posterior by:

$$\begin{aligned} l'_{t+1} &= \ln \frac{\Pr(\boldsymbol{h}_t = (a_1, a_2, \ldots, a_t = A), \sigma_{t+1}|\theta = B)}{\Pr(\boldsymbol{h}_t = (a_1, a_2, \ldots, a_t = A), \sigma_{t+1}|\theta = A)} \\ &= \begin{cases} l_{t+1} + \ln \frac{1-q}{q} & \text{if } \sigma_{t+1} = A \\ l_{t+1} + \ln \frac{q}{1-q} & \text{if } \sigma_{t+1} = B \end{cases} \\ &= \begin{cases} l_{t+1} - \ln \frac{q}{1-q} < 0 & \text{if } \sigma_{t+1} = A \\ \ln \frac{q}{1-q} - |l_{t+1}| \geq 0 & \text{if } \sigma_{t+1} = B \text{ and } |l_{t+1}| \leq \ln \frac{q}{1-q} \\ \ln \frac{q}{1-q} - |l_{t+1}| < 0 & \text{if } \sigma_{t+1} = B \text{ and } |l_{t+1}| > \ln \frac{q}{1-q} \end{cases}. \end{aligned}$$

Assuming agent $t+1$ will follow her private signal when indifferent, *i.e.*, $l'_{t+1} = 0$, we can see that her decision rule is exactly what Lemma 1.2.1 describes. ■

Proof of Lemma 1.2.2**.** By Lemma 1.2.1, the first action a_1 is either a random choice if agent 1 is an amateur or the same as σ_1 if agent 1 is an expert. Hence by Bayes rule,

$$l_2 = \ln \frac{\Pr(a_1|\theta = B)}{\Pr(a_1|\theta = A)} = \begin{cases} \ln \frac{r \cdot \frac{1}{2} + (1-r) \cdot q}{r \cdot \frac{1}{2} + (1-r) \cdot (1-q)} = M & \text{if } a_1 = B \\ \ln \frac{r \cdot \frac{1}{2} + (1-r) \cdot (1-q)}{r \cdot \frac{1}{2} + (1-r) \cdot q} = -M & \text{if } a_1 = A \end{cases}.$$

For all $t \geq 2$, if $a_{t+1} \neq a_t$ then by Lemma 1.2.1 agent $t+1$ must be an expert with

$\sigma_{t+1} = a_{t+1}$.[41] Hence by Bayes rule,

$$\begin{aligned} l_{t+1} &= \ln \frac{\Pr(\boldsymbol{h}_t = (a_1, a_2, \dots, a_t, a_{t+1} \neq a_t) | \theta = B)}{\Pr(\boldsymbol{h}_t = (a_1, a_2, \dots, a_t, a_{t+1} \neq a_t) | \theta = A)} \\ &= l_t + \ln \frac{\Pr(a_{t+1} | \theta = B)}{\Pr(a_{t+1} | \theta = A)} = \begin{cases} l_t + \ln \frac{q}{1-q} = l_t + H & \text{if } a_{t+1} = B \\ l_t + \ln \frac{1-q}{q} = l_t - H & \text{if } a_{t+1} = A \end{cases}. \end{aligned}$$

If $a_{t+1} = a_t$ and $|l_t| \leq \ln \frac{q}{1-q} = H$, by Lemma 1.2.1 we know that agent $t+1$ is either an amateur or an expert with $\sigma_{t+1} = a_{t+1}$. Hence by Bayes rule,

$$\begin{aligned} l_{t+1} &= \ln \frac{\Pr(\boldsymbol{h}_t = (a_1, a_2, \dots, a_t, a_{t+1} = a_t) | \theta = B)}{\Pr(\boldsymbol{h}_t = (a_1, a_2, \dots, a_t, a_{t+1} = a_t) | \theta = A)} \\ &= l_t + \ln \frac{\Pr(a_{t+1} = a_t | \theta = B, |l_t| \leq H)}{\Pr(a_{t+1} = a_t | \theta = A, |l_t| \leq H)} = \begin{cases} l_t + \ln \frac{r+(1-r)q}{r+(1-r)(1-q)} = l_t + L & \text{if } a_{t+1} = B \\ l_t + \ln \frac{r+(1-r)(1-q)}{r+(1-r)q} = l_t - L & \text{if } a_{t+1} = A \end{cases}. \end{aligned}$$

If $a_{t+1} = a_t$ but $|l_t| > \ln \frac{q}{1-q} = H$, by Lemma 1.2.1 agent $t+1$ always follows agent t regardless of her type and private signal, hence the public belief remains unchanged. ■

Proof of Proposition 1.2.1. For the first result, let us consider agent 3's posterior belief after observing $\boldsymbol{h}_2 = (A, B)$ or $\boldsymbol{h}'_2 = (B, B)$. According to Lemma 1.2.2,

$$\begin{aligned} & l_3 = -M + H, \, l'_3 = M + L \\ \implies & l_3 > l'_3 \text{ iff } H > 2M + L. \end{aligned}$$

Due to the symmetric structure of this model, the other comparison between $\widetilde{\boldsymbol{h}}_2 = (A, A)$ and $\widetilde{\boldsymbol{h}}'_2 = (B, A)$ will yield the same inequality.[42]

$$\begin{aligned} & \widetilde{l}_3 = -M - L, \, \widetilde{l}'_3 = M - H \\ \implies & \widetilde{l}_3 > \widetilde{l}'_3 \text{ iff } H > 2M + L. \end{aligned}$$

Hence the posterior belief of agent 3 is always anti-imitative of agent 1 if and only if $H > 2M + L$. Plug in the definitions of H, M, and L in Lemma 1.2.2, we can rewrite

[41] Note that $a_{t+1} \neq a_t$ implies $|l_t| \leq \ln \frac{q}{1-q} = H$ according to Lemma 2.1. Hence a history such that $a_{t+1} \neq a_t$ and $|l_t| > H$ for some t is off equilibrium and not considered by Lemma 2.2.

[42] We will use the symmetric structure of private signals and equilibrium dynamics to simplify future proofs as well.

this inequality in terms of r and q:

$$
\begin{aligned}
& H > 2M + L \\
\iff\ & \ln \frac{q}{1-q} > 2\ln \frac{r\cdot\frac{1}{2} + (1-r)\cdot q}{r\cdot\frac{1}{2} + (1-r)\cdot(1-q)} + \ln \frac{r + (1-r)\cdot q}{r + (1-r)\cdot(1-q)} \\
\iff\ & \frac{q}{1-q} > [\frac{r\cdot\frac{1}{2} + (1-r)\cdot q}{r\cdot\frac{1}{2} + (1-r)\cdot(1-q)}]^2 \cdot \frac{r + (1-r)\cdot q}{r + (1-r)\cdot(1-q)} \\
\iff\ & \frac{r^3}{(1-r)^2(1+r)} > 4q(1-q).
\end{aligned}
$$

Note that $4q(1-q) < 1$ for $q \in (\frac{1}{2}, 1)$, so clearly we can find some pair (r, q) that satisfies this inequality. For example,

$$
\forall r \in [\frac{\sqrt{5}-1}{2}, 1),\ \forall q \in (\frac{1}{2}, 1),\ \frac{r^3}{(1-r)^2(1+r)} \geq 1 > 4q(1-q).
$$

For the second result, let us start with agent 4. By symmetry, we only focus on two comparisons: $\boldsymbol{h}_3 = (A, B, B)$ with $\boldsymbol{h}'_3 = (B, B, B)$, $\widetilde{\boldsymbol{h}}_3 = (A, B, A)$ with $\widetilde{\boldsymbol{h}}'_3 = (B, B, A)$.

$$
\begin{aligned}
& l_4 = l_3 + L,\ l'_4 = l'_3 + L \\
\implies\ & l_4 > l'_4 \text{ iff } l_3 > l'_3 \text{ iff } H > 2M + L; \\
& \widetilde{l}_4 = l_3 - H,\ \widetilde{l}'_4 = l'_3 - H \\
\implies\ & \widetilde{l}_4 > \widetilde{l}'_4 \text{ iff } l_3 > l'_3 \text{ iff } H > 2M + L.
\end{aligned}
$$

Hence $H > 2M + L$ guarantees that agent 4's posterior belief is also anti-imitative. In fact we can see that, as long as each new action added to the histories updates the posterior beliefs in exactly the same way, we do not need extra conditions for anti-imitative beliefs of future agents. However, according to Lemma 1.2.2, the update differs when one posterior belief has absolute value bigger than H but not the other, and we need extra conditions to take that into account.

Consider the earliest such instance: $\boldsymbol{h}_4 = (A, B, B, B)$ with $\boldsymbol{h}'_4 = (B, B, B, B)$. The posterior belief after $\boldsymbol{h}_4$ is

$$
l_5 = -M + H + 2L = H + (2L - M) > H,
$$

since it is easy to verify that $L < M < 2L$. Therefore the only *equilibrium* history extended from $\boldsymbol{h}_4$ is $\boldsymbol{h}_t = (A, B, B, B, \ldots, B)$ for $t \geq 4$ and the only valid equilibrium history extended from $\boldsymbol{h}'_4$ to compare is $\boldsymbol{h}'_t = (B, B, B, B, \ldots, B)$ for $t \geq 4$. By Lemma 1.2.2 it is not difficult to calculate the corresponding l_{t+1} and l'_{t+1}:

$$
\begin{aligned}
l_{t+1} &= l_5 = H + (2L - M), \; \forall t \geq 4; \\
l'_{t+1} &= \begin{cases} M + (t-1)L & \text{if } 4 \leq t \leq k^* + 1 \\ M + k^* L & \text{if } t > k^* + 1 \end{cases},
\end{aligned}
$$

where $k^* \equiv \min\{k \in \mathbb{N}^+ | M + kL > H\}$.[43] Clearly, $l_{t+1} > l'_{t+1}$ for every $t \geq 4$ iff

$$H + (2L - M) > M + k^* L \Longleftrightarrow H > 2M + (k^* - 2)L.$$

Note that $H > 2M + L$ and $M > L$ implies that $k^* \geq 3$, so this extra condition is not redundant in general.

Moreover, the same extra condition will be yielded if we compare two histories where the actions herd on A rather than B eventually. For example, let us consider $\widetilde{\boldsymbol{h}}_4 = (A, A, A, A)$ and $\widetilde{\boldsymbol{h}}'_4 = (B, A, A, A)$. By symmetry $\widetilde{l}'_5 = -l_5 < -H$, hence the only equilibrium history extend from $\widetilde{\boldsymbol{h}}'_4$ is $\widetilde{\boldsymbol{h}}'_t = (B, A, A, A, \ldots, A)$ for $t \geq 4$ and the only valid equilibrium history extended from $\widetilde{\boldsymbol{h}}_4$ to compare is $\widetilde{\boldsymbol{h}}_t = (A, A, A, A, \ldots, A)$ for $t \geq 4$. By symmetry again,

$$
\begin{aligned}
\widetilde{l}'_{t+1} &= -l_{t+1} = -H - (2L - M), \; \forall t \geq 4; \\
\widetilde{l}_{t+1} &= -l'_{t+1} = \begin{cases} -M - (t-1)L & \text{if } 4 \leq t \leq k^* + 1 \\ -M - k^* L & \text{if } t > k^* + 1 \end{cases},
\end{aligned}
$$

where k^* is the same as before. Clearly, $\widetilde{l}_{t+1} > \widetilde{l}'_{t+1}$ for every $t \geq 4$ iff

$$-M - k^* L > -H - (2L - M) \Longleftrightarrow H > 2M + (k^* - 2)L.$$

For a general pair of equilibrium histories that differ only in the first action, agents

[43] It takes $(k^* - 3)$ periods for l'_{t+1} to grow until beyond H.

might flip between A and B multiple times before a herd starts eventually. However, we do not need other extra conditions to take care of them because, as argued earlier, all those actions before the eventual herd update the posterior beliefs in exactly the same way. Therefore, the only condition required here, in addition to $H > 2M + L$, is $H > 2M + (k^* - 2)L$.

Lastly, to see the existence of (r, q) that satisfies $H > 2M + (k^* - 2)L$, we can rewrite this extra condition as

$$2L - M > M + k^*L - H \iff 2 - \frac{M}{L} > k^* - (\frac{H - M}{L}).$$

As $\frac{M}{L} \in (1, 2)$ and $\frac{H-M}{L} \in [k^* - 1, k^*)$,

$$2 - \frac{M}{L} > k^* - (\frac{H - M}{L}) \iff \psi(\frac{M}{L}) < \psi(\frac{H - M}{L}),$$

where $\psi(x) \equiv x - \max\{n \in \mathbb{N} | n \leq x\}$, $\forall x > 0$. It is easy to verify that $\frac{M}{L} \to 2$ as $r \to 1$, thus $\psi(\frac{M}{L}) > 0.5$ for sufficiently large r and

$$\text{if } \psi(\frac{M}{L}) > 0.5,\ \psi(\frac{M}{L}) < \psi(\frac{H - M}{L}) \iff \psi(\frac{H}{L}) \in (2\psi(\frac{M}{L}) - 1, \psi(\frac{M}{L})) \subset [0, 1).$$

Obviously $\frac{H}{L} \to +\infty$ as $r \to 1$, hence for sufficiently large r, $\psi(\frac{H}{L})$ will go through the whole interval $[0, 1)$ infinitely many times. Therefore the extra condition can be satisfied for some (but not every) sufficiently large r. ■

Proof of Corollary 1.2.1. Clearly agent 2 will not anti-imitate agent 1 so let us consider two equilibrium histories with length at least 2 that differ only in the first action. Without loss of generality let us also assume the last action of both histories is B, *i.e.*, $\boldsymbol{h}_t = (A, \ldots, B), \boldsymbol{h}'_t = (B, \ldots, B)$, $t \geq 2$.

If agent $t+1$ is an amateur, she always follows her immediate predecessor and chooses B, so an amateur cannot anti-imitate agent 1. Now suppose agent $t + 1$ is an expert. According to Lemma 1.2.1, the only way to induce anti-imitative behavior of her is that

$$l_{t+1} > H \text{ and } l'_{t+1} \leq H,$$

in which case agent $t+1$ always chooses B after $\boldsymbol{h}_t$ but still follows her private signal σ_{t+1} after l'_{t+1}. Then unconditional on her informational type, the probability of agent $t+1$ choosing B is,

$$\Pr(a_{t+1} = B|\boldsymbol{h}_t) = r \cdot 1 + (1-r) \cdot 1;$$
$$\Pr(a_{t+1} = B|\boldsymbol{h}'_t) = r \cdot 1 + (1-r) \cdot \Pr(\sigma_{t+1} = B|\boldsymbol{h}'_t) = \frac{1+r}{2} < 1.$$

However, l_{t+1} never changes after period $t+1$ while l'_{t+1} will gradually increase until it exceeds H according to Lemma 1.2.2, and there is no more behavioral difference for either type from then on. The number of periods it takes for l'_{t+1} to increase until beyond H has been calculated in the proof of Proposition 1.2.1, which is at most $k^* - 3$. ■

Proof of Lemma 1.3.1. Agent t choose $a_t = B$ if and only if $\Pr(\theta = B|l_t, \sigma_t) \geq 0.5$.[44] By Bayes rule,

$$\Pr(\theta = B|l_t, \sigma_t) = \frac{l_t f_B(\sigma_t)}{l_t f_B(\sigma_t) + (1-l_t) f_A(\sigma_t)} \text{ and } q_t \equiv \Pr(\theta = B|\sigma_t) = \frac{\frac{1}{2} f_B(\sigma_t)}{\frac{1}{2} f_B(\sigma_t) + \frac{1}{2} f_A(\sigma_t)},$$

therefore

$$\Pr(\theta = B|l_t, \sigma_t) \geq 0.5 \iff \frac{f_B(\sigma_t)}{f_A(\sigma_t)} \geq \frac{1-l_t}{l_t} \iff \frac{q_t}{1-q_t} \geq 1 - l_t \iff q_t \geq 1 - l_t.$$

By symmetry, we only show the transition function is $l_{t+1} = m(l_t)$ when $a_t = A$.

$$\begin{aligned} l_{t+1} &\equiv \Pr(\theta = B|\boldsymbol{h}_t) = \Pr(\theta = B|\boldsymbol{h}_{t-1}, a_t = A) \\ &= \frac{\Pr(a_t = A|\theta = B, \boldsymbol{h}_{t-1}) \Pr(\theta = B|\boldsymbol{h}_{t-1})}{\Pr(a_t = A|\theta = B, \boldsymbol{h}_{t-1}) \Pr(\theta = B|\boldsymbol{h}_{t-1}) + \Pr(a_t = A|\theta = A, \boldsymbol{h}_{t-1}) \Pr(\theta = A|\boldsymbol{h}_{t-1})} \\ &= \frac{l_t G_B(1-l_t)}{l_t G_B(1-l_t) + (1-l_t) G_A(1-l_t)} \text{ according to agent } t\text{'s decision rule.} \end{aligned}$$

The fact that $m(l_t) = 1 - n(1-l_t)$ comes directly from symmetry. ■

Proof of Lemma 1.3.2. By symmetry, we only show $l_{t+1} < \frac{1}{2}$ if $a_t = A$. By Lemma 1.3.1,

$$l_{t+1} = \frac{l_t G_B(1-l_t)}{l_t G_B(1-l_t) + (1-l_t) G_A(1-l_t)} \text{ if } a_t = A.$$

[44] Tie-breaking rule is not important given continuous distributions of private beliefs.

Note that

$$\forall x \in (0,1),\ \frac{g_B(x)}{g_A(x)} = \frac{x}{1-x} \Longrightarrow \frac{G_B(x)}{G_A(x)} < \frac{x}{1-x}$$
$$\Longrightarrow\ l_{t+1} = \frac{l_t G_B(1-l_t)}{l_t G_B(1-l_t) + (1-l_t)G_A(1-l_t)} < \frac{l_{t-1}(1-l_{t-1})}{2l_{t-1}(1-l_{t-1})} = \frac{1}{2}.$$

■

Proof of Proposition 1.3.1. Suppose that both $m(\cdot)$ and $n(\cdot)$ are monotone, *i.e.*,

$$\forall 0 < x < x' < 1,\ m(x') \geq m(x) \text{ and } n(x') \geq n(x).$$

Let $\mathcal{F}$ be the set of all finite-order compositions of $m(\cdot)$ and $n(\cdot)$, *i.e.*,

$$\mathcal{F} = \cup_{i \in \mathbb{N}^+} \mathcal{F}_i, \text{ with } \mathcal{F}_1 = \{m, n\} \text{ and } \mathcal{F}_{i+1} = \{m \circ f, n \circ f | \forall f \in \mathcal{F}_i\} \text{ for } i \geq 1.$$

Obviously all functions in $\mathcal{F}$ are monotone: $\forall f \in \mathcal{F}$, $\forall 0 < x < x' < 1$, $f(x') \geq f(x)$.

Now take any two histories $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$ such that $\boldsymbol{h}_t \precsim \boldsymbol{h}'_t$ and $\boldsymbol{h}_t \neq \boldsymbol{h}'_t$. Let $\overline{\tau} \equiv \max\{\tau \leq t | a_\tau = A \text{ and } a'_\tau = B\}$. If $\overline{\tau} = t$, $l_{t+1} < \frac{1}{2} < l'_{t+1}$ according to Lemma 1.3.2, which violates Definition 1.1. If $\overline{\tau} < t$,

$$l_{\overline{\tau}+1} < \frac{1}{2} < l'_{\overline{\tau}+1} \text{ by Lemma 1.3.2;}$$
$$a_\tau = a'_\tau,\ \forall \overline{\tau} < \tau \leq t \Longrightarrow \text{by Lemma 1.3.1, } \exists f \in \mathcal{F} \text{ s.t. } l_{t+1} = f(l_{\overline{\tau}+1}),\ l'_{t+1} = f(l'_{\overline{\tau}+1}).$$

But f is monotone, so $l_{t+1} = f(l_{\overline{\tau}+1}) \leq f(l'_{\overline{\tau}+1}) = l'_{t+1}$, which again violates Definition 1.1. Therefore learning cannot be non-monotone.

To get the explicit conditions on $g_\theta(\cdot)$ for non-monotone transition functions, let us first look at $m(\cdot)$.[45] $m(\cdot)$ is continuously differentiable by Lemma 1.3.1 and

$$sgn(\frac{dm(x)}{dx}) = sgn\ (\frac{d(\frac{m(x)}{1-m(x)})}{dx}) = sgn(\frac{d(\frac{xG_B(1-x)}{(1-x)G_A(1-x)})}{dx}) \text{ by Lemma 1.3.1.}$$

[45] Herrera and Hörner (2012) derived a necessary and sufficient condition for monotone transition functions, which is the increasing hazard ratio property of private signals. It can be verified that the explicit conditions in Proposition 1.3.1 indeed violate that property, and hence transition functions are non-monotone.

The explicit derivative is

$$\frac{d(\frac{xG_B(1-x)}{(1-x)G_A(1-x)})}{dx} = \frac{xG_B(1-x)}{(1-x)G_A(1-x)}\{\frac{1}{x(1-x)} - [\frac{g_B(1-x)}{G_B(1-x)} - \frac{g_A(1-x)}{G_A(1-x)}]\}.$$

Hence

$$sgn(\frac{dm(x)}{dx}) = sgn(\frac{1}{x(1-x)} - [\frac{g_B(1-x)}{G_B(1-x)} - \frac{g_A(1-x)}{G_A(1-x)}])$$
$$\Longrightarrow \quad m(\cdot) \text{ is non-monotone iff } \exists x \in (0,1) \text{ s.t. } \frac{g_B(1-x)}{G_B(1-x)} - \frac{g_A(1-x)}{G_A(1-x)} > \frac{1}{(1-x)x}.$$

Note that $m(x) = 1 - n(1-x)$ (from Lemma 1.3.1) implies that $\frac{dm(x)}{dx} < 0$ if and only if $\frac{dn(1-x)}{d(1-x)} < 0$. So clearly,

$$n(\cdot) \text{ is non-monotone iff } \exists x \in (0,1) \text{ s.t. } \frac{g_B(x)}{G_B(x)} - \frac{g_A(x)}{G_A(x)} > \frac{1}{x(1-x)}.$$

■

Proof of Claim 1.3.1. We first show that $\lim_{x\to 0+} n(x) = \frac{2}{3}$ and $\lim_{x\to 0+} n(x) = \frac{2}{3}$. Note that

$$\lim_{x\to 0+} \frac{n(x)}{1-n(x)} = \lim_{x\to 0+} \frac{x[1-G_B(1-x)]}{(1-x)[1-G_A(1-x)]}$$
$$= \lim_{x\to 0+} \frac{xG_A(x)}{(1-x)G_B(x)} \text{ by symmetry.}$$

Both G_A and G_B are continuously twice differentiable as $g(x)$ is continuously differentiable, so Taylor expansion yields

$$G_A(x) = 0 + g_A(0)x + O(x^2) = 2g(0)x + O(x^2),$$
$$G_B(x) = 0 + g_B(0)x + \frac{\lim_{x\to 0+} g'_B(x)}{2}x^2 + O(x^3) = g(0)x^2 + O(x^3).$$

Hence

$$\lim_{x\to 0+} \frac{n(x)}{1-n(x)} = \lim_{x\to 0+} \frac{2g(0)x^2 + O(x^3)}{g(0)x^2 + O(x^3)} = 2.$$

Therefore $\lim_{x\to 0+} n(x) = \frac{2}{3}$ and by symmetry $\lim_{x\to 1-} m(x) = \frac{1}{3}$.

By definition of $n(\cdot)$ and $m(\cdot)$ in Lemma 1.3.1, $\lim_{x\to 0+} m(x) = 0$ and $\lim_{x\to 1-} n(x) = 1$. If $n(\cdot)$ and $m(\cdot)$ are monotone,

$$\forall x \in [0,1],\ n(x) \in [\frac{2}{3}, 1] \text{ and } m(x) \in [0, \frac{1}{3}].$$

By Lemma 1.3.1, l_t ($t \geq 2$) must be within the image of $n(\cdot)$ or $m(\cdot)$, hence $l_t \notin (\frac{1}{3}, \frac{2}{3})$. ■

Proof of Claim 1.3.2**.** By Lemma 1.3.1,

$$\begin{aligned}
& \Pr(a_{t+1} = B|\boldsymbol{h}_t) - \Pr(a_{t+1} = B|\boldsymbol{h}'_t) > 0 \\
\iff & \Pr(q_{t+1} \geq 1 - l_{t+1}) - \Pr(q_{t+1} \geq 1 - l'_{t+1}) > 0 \\
\iff & \Pr(q_{t+1} \in [1 - l_{t+1}, 1 - l'_{t+1})) > 0 \\
\iff & l_{t+1} > l'_{t+1},
\end{aligned}$$

where the first equivalence comes from the fact that q_{t+1} is independent of $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$. Clearly Definition 1.2 and 1.3 are now equivalent. ■

Proof of Proposition 1.3.2**.** We start by showing if $g^s(\cdot)$ is the unconditional pdf of private beliefs and $g^s(\frac{1}{2}) \cdot \varepsilon_s \geq \frac{1}{2}$, then the transition functions satisfy that $m'(\frac{1}{2}) = n'(\frac{1}{2}) < 0$.

$m'(\frac{1}{2}) = n'(\frac{1}{2})$ comes directly from symmetry. Note that

$$\begin{aligned}
G_B(\frac{1}{2}) &= \int_0^{\frac{1}{2}} 2xg^s(x)dx = \frac{\int_0^{\frac{1}{2}} xg^s(x)dx}{\int_0^{\frac{1}{2}} g^s(x)dx} = E[Z_s|Z_s \leq \frac{1}{2}] = \frac{1}{2} - \varepsilon_s; \\
G_A(\frac{1}{2}) &= 1 - G_B(\frac{1}{2}) = \frac{1}{2} + \varepsilon_s.
\end{aligned}$$

On the other hand, $g^s(\frac{1}{2}) \cdot \varepsilon_s \geq \frac{1}{2} \Longrightarrow g^s(\frac{1}{2}) \geq \frac{1}{2\varepsilon_s} > \frac{1}{2\varepsilon_s} - 2\varepsilon_s = \frac{1-4\varepsilon_s^2}{2\varepsilon_s}$. Hence

$$\begin{aligned}
& \frac{g_B(\frac{1}{2})}{G_B(\frac{1}{2})} - \frac{g_A(\frac{1}{2})}{G_A(\frac{1}{2})} = \frac{g^s(\frac{1}{2})}{\frac{1}{2} - \varepsilon_s} - \frac{g^s(\frac{1}{2})}{\frac{1}{2} + \varepsilon_s} \\
= \ & g^s(\frac{1}{2}) \cdot \frac{8\varepsilon_s}{1 - 4\varepsilon_s^2} > \frac{1 - 4\varepsilon_s^2}{2\varepsilon_s} \cdot \frac{8\varepsilon_s}{1 - 4\varepsilon_s^2} = 4.
\end{aligned}$$

In the proof of Proposition 1.3.1 we have verified that

$$n'(\frac{1}{2}) < 0 \text{ iff } \frac{g_B(\frac{1}{2})}{G_B(\frac{1}{2})} - \frac{g_A(\frac{1}{2})}{G_A(\frac{1}{2})} > \frac{1}{\frac{1}{2}(1-\frac{1}{2})} = 4,$$

so clearly $m'(\frac{1}{2}) = n'(\frac{1}{2}) < 0$.

Take $t = 1$ and we want to show that agent 3 always anti-imitate agent 1. By symmetry we only compare history $\boldsymbol{h}_3 = (A, B)$ with $\boldsymbol{h}'_3 = (B, B)$. By Lemma 1.3.1,

$$l_3 = n(m(\frac{1}{2})) = n(G_B(\frac{1}{2})) \text{ and } l'_3 = n(n(\frac{1}{2})) = n(G_A(\frac{1}{2})).$$

Hence

$$l_3 - l'_3 = n(\frac{1}{2} - \varepsilon_s) - n(\frac{1}{2} + \varepsilon_s).$$

Clearly $\varepsilon_s \to 0$ as $s \to \infty$ since $Z_s \xrightarrow{d} X_{\frac{1}{2}}$ as $s \to \infty$. Both $n(\cdot)$ and $m(\cdot)$ are continuously differentiable, so by first-order Taylor expansion,

$$\exists \widetilde{S}_1 \text{ s.t. } \forall s > \widetilde{S}_1,\ l_3 - l'_3 = -2\varepsilon_s \cdot n'(\frac{1}{2}) + O(\varepsilon_s^2).$$

We know from above that $n'(\frac{1}{2}) < 0$ when $s > S$, therefore $\forall s > S_1 \equiv \max\{S, \widetilde{S}_1\}$, $l_3 - l'_3 > 0$ and agent 3 always anti-imitates agent 1.

Now take $t = 2$. First we compare $\boldsymbol{h}_3 = (A, A, A)$ with $\boldsymbol{h}'_3 = (A, B, A)$. Again by first-order Taylor expansion,

$$\begin{aligned}
l_4 &= m(m(m(\frac{1}{2}))) = m(m(\frac{1}{2} - \varepsilon_s)) \\
&= m(m(\frac{1}{2}) - m'(\frac{1}{2})\varepsilon_s + O(\varepsilon_s^2)) \\
&= m(\frac{1}{2} - \varepsilon_s - m'(\frac{1}{2})\varepsilon_s + O(\varepsilon_s^2)) \\
&= m(\frac{1}{2}) - [m'(\frac{1}{2}) + (m'(\frac{1}{2}))^2]\varepsilon_s + O(\varepsilon_s^2); \\
l'_4 &= m(n(m(\frac{1}{2}))) = m(n(\frac{1}{2} - \varepsilon_s)) \\
&= m(n(\frac{1}{2}) - n'(\frac{1}{2})\varepsilon_s + O(\varepsilon_s^2)) \\
&= m(\frac{1}{2} + \varepsilon_s - n'(\frac{1}{2})\varepsilon_s + O(\varepsilon_s^2)) \\
&= m(\frac{1}{2}) + [m'(\frac{1}{2}) - m'(\frac{1}{2})n'(\frac{1}{2})]\varepsilon_s + O(\varepsilon_s^2).
\end{aligned}$$

Note that $m'(\frac{1}{2}) = n'(\frac{1}{2})$ and thus

$$\exists \widetilde{S}_2 \text{ s.t. } \forall s > \widetilde{S}_2,\ l_3 - l'_3 = -2\varepsilon_s \cdot m'(\frac{1}{2}) + O(\varepsilon_s^2).$$

Then we compare $\boldsymbol{h}_3 = (A, A, B)$ with $\boldsymbol{h}'_3 = (A, B, B)$. Similarly we have

$$\begin{aligned}
l_4 &= n(m(m(\frac{1}{2}))) = n(m(\frac{1}{2} - \varepsilon_s)) \\
&= n(m(\frac{1}{2}) - m'(\frac{1}{2})\varepsilon_s + O(\varepsilon_s^2)) \\
&= n(\frac{1}{2} - \varepsilon_s - m'(\frac{1}{2})\varepsilon_s + O(\varepsilon_s^2)) \\
&= n(\frac{1}{2}) - [n'(\frac{1}{2}) + n'(\frac{1}{2})m'(\frac{1}{2})]\varepsilon_s + O(\varepsilon_s^2); \\
l'_4 &= n(n(m(\frac{1}{2}))) = n(n(\frac{1}{2} - \varepsilon_s)) \\
&= n(n(\frac{1}{2}) - n'(\frac{1}{2})\varepsilon_s + O(\varepsilon_s^2)) \\
&= n(\frac{1}{2} + \varepsilon_s - n'(\frac{1}{2})\varepsilon_s + O(\varepsilon_s^2)) \\
&= n(\frac{1}{2}) + [n'(\frac{1}{2}) - (n'(\frac{1}{2}))^2]\varepsilon_s + O(\varepsilon_s^2).
\end{aligned}$$

Note that $m'(\frac{1}{2}) = n'(\frac{1}{2})$ and thus

$$\exists \widehat{S}_2 \text{ s.t. } \forall s > \widehat{S}_2,\, l_3 - l_3' = -2\varepsilon_s \cdot n'(\frac{1}{2}) + O(\varepsilon_s^2).$$

Therefore, $\forall s > S_2 \equiv \max\{S, \widetilde{S}_2, \widehat{S}_2\}$, $l_4 - l_4' > 0$. By symmetry we don't need to compare other pairs so agent 4 always anti-imitates agent 2.

For $t \geq 3$, we want to compare $\boldsymbol{h}_{t+1}$ and $\boldsymbol{h}'_{t+1}$ that differ only in the t-th action ($a_t = A$ and $a_t' = B$). Using first-order Taylor expansion recursively like we did, it is easy to get that when s is sufficiently large,

$$l_{t+2} - l_{t+2}' = \begin{cases} -2\varepsilon_s \cdot m'(\frac{1}{2}) + O(\varepsilon_s^2) & \text{if } a_{t+1} \text{ is } A \\ -2\varepsilon_s \cdot n'(\frac{1}{2}) + O(\varepsilon_s^2) & \text{if } a_{t+1} \text{ is } B \end{cases}.$$

Hence $\exists S_t \geq S$ s.t. $\forall s > S_t$, $l_{t+2} - l_{t+2}' > 0$ and agent $t+2$ always anti-imitates agent t.

■

Proof of Proposition 1.3.3**.** We first show that,

$$\text{if } \lim_{x\to 0+} \frac{g'(x)}{g(0)} > 3,\ \lim_{x\to 0+} \frac{dm(x)}{dx} < 0.$$

Recall that $sgn[\frac{dm(x)}{dx}] = sgn\ [\frac{d(\frac{m(x)}{1-m(x)})}{dx}]$ and we have already calculated the $\frac{d(\frac{m(x)}{1-m(x)})}{dx}$ in the proof of Proposition 1.3.1:

$$\begin{aligned} & \lim_{x\to 1-} \frac{d(\frac{m(x)}{1-m(x)})}{dx} \\ = & \lim_{x\to 1-} \frac{xG_B(1-x)}{(1-x)G_A(1-x)}[\frac{1}{x(1-x)} - \frac{g_B(1-x)}{G_B(1-x)} + \frac{g_A(1-x)}{G_A(1-x)}] \\ = & \lim_{z\to 0+} \frac{G_A(z)G_B(z) + z(1-z)[g_A(z)G_B(z) - G_A(z)g_B(z)]}{z^2G_A^2(z)} \quad (z \equiv 1-x) \\ = & \lim_{z\to 0+} \frac{G_A(z)G_B(z) + 2z(1-z)g(z)[(1-z)G_B(z) - zG_A(z)]}{z^2G_A^2(z)}. \end{aligned}$$

To simplify the expressions later, let $R \equiv g(0) > 0$ and $W \equiv \lim_{x\to 0+} g'(x)$. Both G_A and G_B are continuously three times differentiable as $g(\cdot)$ is continuously twice differentiable, so Taylor expansion yields that

$$
\begin{aligned}
g(z) &= g(0) + \lim_{z\to 0+} g'(z)z + O(z^2) = R + Wz + O(z^2); \\
G_A(z) &= 0 + g_A(0)z + \frac{\lim_{z\to 0+} g'_A(z)}{2} z^2 + O(z^3) = 2Rz + (W-R)z^2 + O(z^3); \\
G_B(z) &= 0 + g_B(0)z + \frac{\lim_{z\to 0+} g'_B(z)}{2} z^2 + \frac{\lim_{z\to 0+} g''_B(z)}{6} z^3 + O(z^4) = Rz^2 + \frac{2}{3}Wz^3 + O(z^4).
\end{aligned}
$$

Plug these back to the expression above,

$$
\begin{aligned}
& \lim_{x\to 1-} \frac{d(\frac{m(x)}{1-m(x)})}{dx} \\
= & \lim_{z\to 0+} \frac{G_A(z)G_B(z) + 2z(1-z)g(z)[(1-z)G_B(z) - zG_A(z)]}{z^2 G_A^2(z)} \\
= & \lim_{z\to 0+} \frac{(R^2 - \frac{1}{3}RW)z^4 + O(z^5)}{4R^2z^4 + O(z^5)} \\
= & \frac{3R - W}{12R} < 0 \text{ since } \frac{W}{R} > 3 \text{ and } R > 0.
\end{aligned}
$$

Therefore $\lim_{x\to 1-} \frac{dm(x)}{dx} < 0$ and by symmetry $\lim_{x\to 0+} \frac{dn(x)}{dx} < 0$. By continuous differentiability of $m(\cdot)$ and $n(\cdot)$,

$$
\exists \varepsilon > 0 \text{ s.t. } \forall x \in (0, \varepsilon),\ \frac{dn(x)}{dx} < 0;\ \forall x \in (1-\varepsilon, 1),\ \frac{dm(x)}{dx} < 0.
$$

Without loss of generality, let us assume $\theta = A$. Consider two histories, $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$, such that

1. $\boldsymbol{h}_t \precsim \boldsymbol{h}'_t$ and $\boldsymbol{h}_t \neq \boldsymbol{h}'_t$;
2. $l_{t+1} \neq l'_{t+1}$ and $\{l_{t+1}, l'_{t+1}\} \subset (0, \varepsilon)$.

Such histories exist when t is sufficient large, because of Lemma 1.3.4 and the fact that, for a generic density function $g(\cdot)$, we can always construct two different histories that generate different posterior beliefs.[46] If $l_{t+1} > l'_{t+1}$, learning is monotone by definition. If

[46] Since we are allowed to increase the length of the two histories arbitrarily, a density function $k(\cdot)$ that always leads to $l_{t+1} = l'_{t+1}$ is non-generic.

$l_{t+1} < l'_{t+1}$, we can simply extend $\boldsymbol{h}_t$ and $\boldsymbol{h}'_t$ by an action B, *i.e.*,

$$\boldsymbol{h}_{t+1} = (\boldsymbol{h}_t, B) \text{ and } \boldsymbol{h}'_{t+1} = (\boldsymbol{h}'_t, B).$$

According to Lemma 1.3.1, $l_{t+2} = n(l_{t+1}) > n(l'_{t+1}) = l'_{t+2}$. So learning is again non-monotone. ∎

1.5.2 Additional Claim

Here we present a claim related to Proposition 1.3.3, and show that any information structure satisfying the primary assumptions imposed in Subsection 1.3.1 can be well approximated by another information structure that leads to non-monotone learning. It is not hard to understand this claim, given that Proposition 1.3.3 only imposes a boundary condition on the information structure.

Claim 1.5.1 *Take any density function $g(\cdot)$ on $[0,1]$ that satisfies the assumptions of full support, differentiability, and symmetry. There exists a density function on $[0,1]$, $g(\cdot;\gamma)$, such that*

$$g(\cdot;\gamma) \to g(\cdot) \text{ as } \gamma \to 0, \text{ and } \lim_{x\to 0+} \frac{dg(x;\gamma)}{dx} > 3 \cdot g(0;\gamma) \text{ for any } \gamma > 0.$$

Learning is then non-monotone with $g(0;\gamma)$ being the unconditional density function of private beliefs.

Proof. Construct $g(x;\gamma) \equiv g(x) - \gamma + \gamma b(x;\theta_\gamma)$, where $0 < \gamma < \min_{x\in[0,1]} g(x)$ and $b(x;\theta_\gamma) \equiv \frac{[x(1-x)]^{\theta_\gamma}}{B(1+\theta_\gamma, 1+\theta_\gamma)}$ is the pdf of Beta distribution.

Note that

$$\begin{aligned} &\lim_{\substack{x\to 0+ \\ \theta\to 0+}} \frac{db(x;\theta)}{dx} \to +\infty \\ \Longrightarrow\quad &\exists \theta_\gamma > 0 \text{ s.t. } \lim_{x\to 0+} \frac{db(x;\theta_\gamma)}{dx} > \frac{[3\cdot g(0) - \lim_{x\to 0+} g'(x)]}{\gamma} - 3. \end{aligned}$$

Clearly $g(\cdot;\gamma) \to g(\cdot)$ as $\gamma \to 0$, and

$$\begin{aligned}
& \lim_{x\to 0+} \frac{dg(x;\gamma)}{dx} \\
= \; & \lim_{x\to 0+} g'(x) + \gamma \lim_{x\to 0+} \frac{db(x;\theta_\gamma)}{dx} \\
> \; & \lim_{x\to 0+} g'(x) + 3\cdot g(0) - \lim_{x\to 0+} g'(x) - 3\gamma \\
= \; & 3\cdot g(0) - 3\gamma = 3\cdot g(0;\gamma).
\end{aligned}$$

■

2 Does Public Information Disclosure Help Social Learning?

This work studies the effect of releasing exogenous public information in the canonical social-learning model that predicts informational cascades and incomplete learning. In particular, we consider the simple setting with binary states, binary actions, and binary private signals. For the purpose of increasing the average expected payoff of the population, we show that it is weakly better to postpone the disclosure of a public signal irrespective of its precision. However, such weak monotonicity no longer holds if the objective is to maximize the discounted sum of people's expected payoffs or if the model goes beyond the binary setting. On the other hand, it is suboptimal to ever release a public signal that is less precise than people's private signals even if sophisticated releasing strategies are allowed as noisy public information crowds out private information and harms information aggregation.

2.1 Introduction

As important as the rationalization of herd behavior, one contribution of the theoretical literature on rational social learning is the prediction of *incomplete learning*, *i.e.*, inefficient information aggregation among the population, when people have boundedly accurate private information. In the canonical binary model due to Bikhchandani *et al.* (1992), agents make binary choices between A and B sequentially over an infinite time horizon, and before taking her action, each agent receives a private binary signal indicating which option is better with uniform precision and observes all the past actions. Eventually rational agents herd without fully learning the truth and *informational cascade* arises.[47]

Like every other economic model that predicts inefficient outcomes, we naturally ask ourselves of potential ways to improve efficiency in such environments. In fact, as Bikhchandani *et al.* (1992) pointed out, informational cascades are *fragile*: since information stops to aggregate, the cascades and hence the herds are vulnerable to new pieces of information. Therefore it is worthwhile to investigate whether and how disclosure of exogenous public information can improve social learning.

We introduce a social planner to the canonical binary model, who receives an extra signal about the truth with certain precision and decides whether and when to release it to the public to maximize the average expected payoff of the population.[48] In particular, the social planner is looking for the optimal timing (if any) of releasing that public signal, given its precision, that essentially maximizes the expected payoff of limiting agents.

We first provide an *anti-transparency* result: the social planner should never release a *noisy* public signal that is less precise than people's private signals. An informational cascade arises when one action, say A, has been chosen at least twice more than B, and agents start to herd on A.[49] Releasing a noisy public signal then cannot break down the informational cascade as agents will continue to herd on A even when the public signal

[47] Strictly speaking, incomplete learning does not necessarily imply informational cascades when private signals are continuous rather than discrete. See Herrera and Hörner (2012) for a discussion about a necessary and sufficient condition on the distribution of private signals for informational cascades.

[48] Due to the presence of herd behavior, the average expected payoff of the population is equivalent to the expected payoff of limiting agents. This is a common objective of interest in the literature of social learning, and we use it as a measure of social welfare for the social planner. In Section 2.3 we will discuss an alternative objective function of the social planner.

[49] We assume each agent follows her private signal when indifferent.

suggests B: combining the history and the public signal, future agents still find action A sufficiently attractive. Hence a noisy public signal has no effect on the limiting expected payoff if it is released after an informational cascade has arisen. On the other hand, when an informational cascade has not yet formed, releasing a noisy public signal may induce a wrong cascade in the future more likely than people's own private signals due to the lower precision, hence lowers the limiting expected payoff. Therefore overall to release a noisy public signal is a *bad* idea for the social planner. Moreover, this result is robust when sophisticated releasing strategies are allowed, *i.e.*, a noisy public signal should never be released even if the social planner can make the timing of disclosure contingent on the history of past actions.[50]

The other result, perhaps more interesting, is a *monotonicity* result: the expected payoff of limiting agents is *weakly increasing* in the period at which the public signal is released, regardless of its precision. In other words, the social planner should always *postpone* the disclosure of any public information.[51] The intuition behind this result is that the benefit of releasing a public signal is greater when an informational cascade has arisen than when it has not. Before a cascade starts the information aggregation of private signals is still going on, so a public signal released then may crowd out the next private signal(s) in terms of updating people's belief. Hence the "net" informational contribution of the public signal is lower than when it is released after a cascade has started. Meanwhile, the probability of entering an informational cascade is weakly increasing over time, thus implies that the benefit of releasing a public signal is also weakly increasing over time.

Nevertheless the monotonicity result seems not compelling especially for extremely precise public signals: suppose the social planner holds a public signal that perfectly reveals the truth, then she should naturally release it *as early as possible* so that everyone can learn the truth from it and choose the right action. This thought experiments casts doubt on whether the limiting expected payoff or the average expected payoff is a proper objective for the social planner, and we reconsider the whole problem assuming that

[50] In general numerous Bayesian Nash equilibria exist with contingent releasing strategies, so we focus on a selection of equilibria to make meaningful prediction. See Subsection 2.2.3 for details.

[51] Note that the monotonicity result is true even for noisy public signals, but does not contradict with the anti-transparency result: releasing a noisy public signal is bad, but if the social planner were to release one, she should postpone as much as possible.

the social planner wants to maximize the *discounted* sum of people's expected payoffs instead. Although the optimal timing of disclosure is not yet clear to us, we show that the monotonicity result is no longer true. If the social planner is indifferent between two periods to release the public signal before, now she strictly prefers the earlier period of the two due to the discount factor.

We also present an alternative setting with ternary states, actions and signals, and show that the monotonicity result does not hold either. In this setting, at some point in the history an action could be excluded by all the agents afterwards: *e.g.*, an agent observing history (A, B, A, B, A, B) would not choose C regardless of her private signal and neither would all her successors. We call this situation a *trap* away from action C, and unlike a herd, the informational "depth" of a trap can increase over time; hence a public signal could fail to break down a wrong trap if it is released too late, which is not what the social planner wants.

Related literature. This paper is related to a stream of papers on anti-transparency. Morris and Shin (2002) presented a model where every agent wants to minimize a loss function made up of two components: loss in the distance between her action and the underlying state, and loss in the distance between her action and the average action in the population, *i.e.*, a "beauty-contest" term.[52] With later comments by Svensson (2006) and Morris *et al.* (2006), it can be shown that in such a model the welfare with noisy public information could indeed be worse than the welfare without.[53] Demertzis and Hoeberichts (2007) further explored this anti-transparency result by introducing costly information acquisition to the model, where people might free-ride on public information and abandon private information acquisitions. In this paper we get an anti-transparency result as well in the canonical social-learning model, but without beauty-contest-like preference or costly information acquisition.[54]

This work is clearly related to the social learning literature initiated by Banerjee (1992), Bikhchandani *et al.* (1992), and Smith and Sørensen (2000). Nevertheless few papers talked about disclosure of public information in social-learning models. Bikhchan-

[52] See Keynes (1936).

[53] In that model, public information serves as a coordination device for the second loss term, and people could overlook their private signals when they put a sufficiently high weight on the second loss term.

[54] Compared to Morris and Shin (2002), noisy public information distorts social welfare in this model through informational externality rather than payoff interdependence.

dani *et al.* (1992) pointed out the fragility of informational cascades and only briefly discussed the effect of releasing extra information, while this work further looks into this issue and investigates the optimal timing of release. Gill and Sgroi (2008) also augmented the standard model to allow a principal to provide public information to the agents by subjecting herself to a test of certain toughness at the beginning.[55] On the other hand, as discussed before, the monotonicity result in this paper might question the plausibility of limiting efficiency, which is the common objective of interest in most of the literature, as a good measure of social welfare in social-learning models.

The remainder of the paper is structured as follows. Section 2.2 sets up the canonical binary model and provides the main result. Section 2.3 discusses settings with impatient social planner and with ternary states/actions. Section 2.4 concludes.

2.2 A Simple Setting with Binary Choice

2.2.1 Setup and Preliminaries

There is a population of countably infinite agents who are exogenously ordered to make a binary choice sequentially. Each agent is labelled by the period of her turn, $t \in T = \{1, 2, 3, ...\}$.

The state of the world θ is realized out of a binary state space $\Theta \equiv \{1, -1\}$ before anyone makes the choice, with $\Pr\{\theta = 1\} = 1/2$. After the realization of θ, every agent t receives a private signal $s_t \in \{1, -1\}$ and the private signals are conditionally i.i.d. with

$$\Pr\{s_t = 1|\theta = 1\} = \Pr\{s_t = -1|\theta = -1\} = q \in (\frac{1}{2}, 1),$$

where the precision q is common knowledge to the whole population.

Before exerting her action $a_t \in A = \{1, -1\}$, agent t is allowed to observe the history of all her predecessors' choices, $\boldsymbol{h}_t \in H_t \equiv \{\emptyset\} \cup A^{t-1}$, where $\boldsymbol{h}_1 \equiv \emptyset$ denotes the empty history at period 1. Agents have identical utility function

$$u(a_t; \theta) = 1_{\{\theta=a_t\}}$$

[55] Essentially the outcome of the test is like a public signal with certain precision (based on the toughness) that is released at the beginning. Note that the principal in Gill and Sgroi (2008) does not have the same objective of the social planner in this paper though.

and are assumed to follow their own signal when indifferent.

Similar to Chapter 1, we call $w_t \equiv \log_{q/(1-q)}[\frac{\Pr(\boldsymbol{h}_t|\theta=1)}{\Pr(\boldsymbol{h}_t|\theta=-1)}]$ the *public belief* after history $\boldsymbol{h}_t$.[56]

This is essentially the canonical model in Bikhchandani *et al.* (1992) and it is well known that the Bayesian Nash equilibrium exhibits herd behavior eventually. For the purpose of future analysis though, let us restate the existing results as lemmata.

Lemma 2.2.1 *The Bayesian Nash equilibrium strategy of each agent t is given by*

$$a_t^*(\boldsymbol{h}_t, s_t) = a^*(w_t, s_t) \equiv \begin{cases} s_t & \text{if } |w_t| \leq 1 \\ sgn(w_t) & \text{otherwise} \end{cases},$$

where $sgn(x) \equiv \begin{cases} x/|x| & \text{if } x \neq 0 \\ 0 & \text{otherwise} \end{cases}$.

Hence, along the equilibrium path $\boldsymbol{h}_t^* = (a_1^*, a_2^*, ..., a_{t-1}^*)$ *with* $\boldsymbol{h}_1^* \equiv \emptyset$*, the dynamic of public beliefs is given by*

$$w_1^* = 0;\ w_{t+1}^* = \begin{cases} w_t^* + a_t^* & \text{if } |w_t^*| \leq 1 \\ w_t^* & \text{otherwise} \end{cases}.$$

Proof. See Appendix. ■

Lemma 2.2.1 shows that public belief w_t serves as a sufficient statistic for agent t's decision problem and in equilibrium w_t^* stops to update once it leaves interval $[-1, 1]$, which is exactly when an informational cascade, or a herd, starts.

Definition 2.1 *We say a* ***herd*** *on action* $1(-1)$ *starts at period* T *if*

$$\forall t \geq T,\ a_t(\boldsymbol{h}_t, s_t) = a_t(\boldsymbol{h}_t, \cdot) = 1(-1).$$

Lemma 2.2.2 *Along the equilibrium path described by Lemma* 2.2.1*, a herd starts even-*

[56] Compared to the definition of public belief l_t on page 15 of Chapter 1, we only change the base of the logarithm from e to $\frac{q}{1-q}$ here, *i.e.*, a linear transformation. It is convenient to use w_t here as it only takes integer values in equilibrium (prior to the disclosure of public information) and in fact represents the *net* number of private signals revealed by the history.

tually with probability 1.

Proof. See Appendix. ■

Note that the eventual herd could be incorrect, as Smith and Sørensen (2000) argued, if agents have *bounded private beliefs*, which is exactly the case here. The probability of a correct herd eventually is nevertheless important for our analysis later on social welfare, as it determines expected payoffs for future agents in the long run. Hence we would like to calculate this probability here.

For convenience, let us assume the realization of θ is 1 without loss of generality for the remainder of this section.[57] Define

$$p(x) \equiv \Pr(\text{plim}_{t\to\infty} a_t^*(\boldsymbol{h}_t^*, s_t) = 1 | w_1 = x, \theta = 1), \forall x \in \mathbb{R},$$

the probability of a correct herd eventually conditional on some initial public belief $w_1 = x$, which can be explicitly calculated according to the following useful lemma.

Lemma 2.2.3 *$p(x)$ can only take the following 7 discrete values:*

$$\begin{aligned}
p(x) &= 0 \equiv \alpha_1, \ \forall x < -1; \\
p(-1) &= \frac{q^3}{1-2q(1-q)} \equiv \alpha_2; \\
p(x) &= \frac{q^2}{1-q(1-q)} \equiv \alpha_3, \ \forall x \in (-1,0); \\
p(0) &= \frac{q^2}{1-2q(1-q)} \equiv \alpha_4; \\
p(x) &= \frac{q}{1-q(1-q)} \equiv \alpha_5, \ \forall x \in (0,1); \\
p(1) &= q + \frac{(1-q)q^2}{1-2q(1-q)} \equiv \alpha_6; \\
p(x) &= 1 \equiv \alpha_7, \ \forall x > 1.
\end{aligned}$$

[57] It is without loss of generality from an *ex-ante* perspective due to the symmetric setting.

In fact, $\boldsymbol{\alpha} \equiv (\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6, \alpha_7)^\top$ *satisfies* $Q\boldsymbol{\alpha} = \boldsymbol{\alpha}$ *where*

$$Q \equiv \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1-q & 0 & 0 & q & 0 & 0 & 0 \\ 1-q & 0 & 0 & 0 & q & 0 & 0 \\ 0 & 1-q & 0 & 0 & 0 & q & 0 \\ 0 & 0 & 1-q & 0 & 0 & 0 & q \\ 0 & 0 & 0 & 1-q & 0 & 0 & q \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Proof. See Appendix. ■

Note that the subgame starting from a period T is identical to the original game (and the equilibrium strategy is stationary according to Lemma 2.2.1), hence Lemma 2.2.3 actually tells us how to calculate the probability of a correct herd eventually if the public belief at period T is w_T. On the other hand, the matrix Q introduced in Lemma 2.2.3 also helps us to characterize the equilibrium public beliefs as a *monotone* Markov chain.

Definition 2.2 *A transition matrix* $C = (c_{ij})_{n\times n}$ *is* ***monotone*** *if*

$$\forall 1 \le i < j \le n, \forall k \le n, \sum_{m=1}^{k} c_{mi} \ge \sum_{m=1}^{k} c_{mj}.$$

A Markov chain is ***monotone*** *if it has a monotone transition matrix.*[58]

Lemma 2.2.4 *Let* $\mathcal{P} = \{P_1, P_2, P_3, P_4, P_5, P_6, P_7\}$ *be a finite partition of* $\mathbb{R}$ *with:*

$$P_1 = (-\infty, -1),\ P_2 = \{-1\},\ P_3 = (-1, 0),\ P_4 = \{0\},$$
$$P_5 = (0, 1),\ P_6 = \{1\},\ P_7 = (1, +\infty)\}.$$

Define $\pi_i^t \equiv \Pr(w_t^* \in P_i)$ *and* $\boldsymbol{\pi}^t = (\pi_1^t, \pi_2^t, \pi_3^t, \pi_4^t, \pi_5^t, \pi_6^t, \pi_7^t)$ *is hence the probability vector of* w_t^* *over partition* $\mathcal{P}$. *We have*

$$\boldsymbol{\pi}^{t+1} = \boldsymbol{\pi}^t Q \text{ with } \boldsymbol{\pi}^1 = (0, 0, 0, 1, 0, 0, 0),$$

[58] These definitions come from Keilson and Kester (1977).

where the transition matrix Q is given in Lemma 2.2.3 *and it is monotone.*

Proof. See Appendix. ■

2.2.2 Exogenous Release of Public Information

As we have seen in the previous subsection, a herd starts eventually but it is possibly on the wrong action. Bikhchandani *et al.* (1992) referred to the eventual herd as an informational cascade and pointed out that it is vulnerable to public information disclosure. Here we look into this issue more specifically by introducing public information release into the model.

In addition to the population of agents as before, there is a social planner who also receives a signal $\widetilde{s} \in \{1, -1\}$ after the realization of θ and

$$\Pr\{\widetilde{s} = 1|\theta = 1\} = \Pr\{\widetilde{s} = -1|\theta = -1\} = \widetilde{q} \in (\frac{1}{2}, 1).$$

The precision $\widetilde{q}$ is common knowledge to the whole population and $\widetilde{s}$ is conditionally independent of any s_t.

The social planner can decide whether and when to release the signal $\widetilde{s}$ to the public. Once $\widetilde{s}$ is released at period $\tau \geq 1$ it becomes public information and every agent afterwards, $t \geq \tau$, can take it into account before she makes her decision. The social planner wants to maximize the expected average payoff of the whole population,

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} Eu(a_t^*; \theta).$$

Note that due to the existence of herd behavior, the social planner's objective is essentially to maximize the probability of a correct herd eventually, or say, the probability of learning the truth eventually.

In this subsection we particularly consider the situation where the releasing strategy is *exogenous*, namely she has to decide a period $\tau \in \{1, 2, 3...\}$ to release or not to release at all before anything happens and commits to that. Keep in mind that we still assume the realization of θ is 1 without loss of generality.

Note that private signals are equally precise, hence the public belief w_t can also be interpreted as the *net* number of correct private signals revealed by history $\boldsymbol{h}_t$. We would like to first have a similar interpretation of the public signal by measuring its precision with respect to private signals:

Definition 2.3 *The public signal* $\widetilde{s}$ *has* ***relative precision*** $\lambda \in \mathbb{R}^+$ *if*

$$\log_{q/(1-q)}[\frac{\widetilde{q}}{1-\widetilde{q}}] = \lambda, \text{ or equivalently, } \widetilde{q} = \frac{q^\lambda}{q^\lambda + (1-q)^\lambda}.$$

When the public signal has **relative precision** λ, we have

$$\begin{aligned} & \frac{\Pr\{\widetilde{s}|\theta = 1\}}{\Pr\{\widetilde{s}|\theta = -1\}} = (\frac{\widetilde{q}}{1-\widetilde{q}})^{\widetilde{s}} = [(\frac{q}{1-q})^\lambda]^{\widetilde{s}} \\ = \; & [(\frac{q}{1-q})^{s_t}]^\lambda = [\frac{\Pr(s_t|\theta = 1)}{\Pr(s_t|\theta = -1)}]^\lambda \text{ whenever } s_t = \widetilde{s}. \end{aligned}$$

That is, learning a public signal in favor of one state with relative precision λ is equivalent to learning λ *net* private signals in favor of that state.

Now suppose the social planner releases the public signal at period τ. Then the subgame after the release is equivalent to the original game without public information, which we discussed in the previous subsection, but with an initial public belief inferred from bother the history before period τ and the public signal. Hence, a herd still starts eventually and the expected average payoff of the population is just the probability of a correct herd eventually, which depends only on the initial public belief according to Lemma 2.2.3. Meanwhile, using the relative precision, we can linearly describe the effect of the public signal on the public belief. These observations are summarized in the following lemma.

Lemma 2.2.5 *Suppose the social planner releases* $\widetilde{s}$ *with relative precision* λ *at period* $\tau \geq 1$ *and the history before that has generated a public belief* w_τ*. Then the new public belief after release will be*

$$\widetilde{w}_\tau = w_\tau + \lambda \widetilde{s},$$

and (under the assumption that the realization of θ *is* 1*) the expected average payoff of* *the population conditional on $\widetilde{w}_\tau$ is simply given by*

$$\text{plim}_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}u(a_t^*;\theta=1|\widetilde{w}_\tau)=p(\widetilde{w}_\tau)=p(w_\tau+\lambda\widetilde{s}).$$

Furthermore, let us denote $v(\tau;\lambda)$ as the (unconditional) expected average payoff of the population when the social planner is to release the public signal with relative precision λ at period τ, then

$$v(\tau;\lambda)=E_{w_\tau,\widetilde{s}}[p(w_\tau+\lambda\widetilde{s})].$$

Proof. See Appendix. ■

Keep in mind that without no release at all, the expected average payoff of the population, which is just the probability of a correct herd eventually, is equal to α_4 given in Lemma 2.2.3. Now we are in a position to provide the main result of this section.

Proposition 2.2.1 *1. It is never optimal to release a public signal less precise than the private signals. That is,*

$$\forall\lambda\in(0,1),\ \forall\tau\geq 1,\ v(\tau;\lambda)<\alpha_4.$$

2. It is strictly better to release a public signal no less precise than the private signals than not to release at all. That is,

$$\forall\lambda\in[1,+\infty),\ \exists\tau<\infty \text{ such that } v(\tau;\lambda)>\alpha_4.$$

3. It is always (weakly) better to release a public signal later than sooner regardless of its precision. That is,

$$\forall\lambda\in\mathbb{R}^+,\ \forall\tau\geq 1,\ v(\tau+1;\lambda)\geq v(\tau;\lambda).$$

Proof. See Appendix. ■

Here we would like to talk about the third statement of Proposition 2.2.1 in particular. The weak monotonicity of $v(\tau;\lambda)$ in τ mathematically comes from the fact that equilibrium public beliefs evolve according to a *monotone* transition matrix until the public signal is released, regardless of the value of λ. However, the intuition for this weak monotonicity is not as universal as the property itself. For illustrative purposes, let us focus on two cases, $\lambda = 1$ and $\lambda < 1$.

We can think of releasing a public signal as an "additional" agent joining in the sequence who always follows his own private signal $\widetilde{s}$. When $\lambda = 1$, the release has no ex-ante effect if a herd has not started yet because every agent t just follows her own private signal s_t, which has the same precision as $\widetilde{s}$, before a herd starts. In this case the ex-ante benefit of releasing $\widetilde{s}$ arises after a herd starts, where $\widetilde{s}$ is more likely to break down a wrong herd than to break down a correct herd as $\widetilde{q} > 1/2$. So the benefit of releasing $\widetilde{s}$ is increasing in the probability of herding at the time of release. It is easy to verify that the probability of herding is weakly increasing in t, which explains the weak monotonicity of $v(\tau;\lambda)$ in τ.

When $\lambda < 1$, however, releasing $\widetilde{s}$ has no effect once a herd starts: $|\widetilde{w}_\tau| = |w^*_\tau + \lambda\widetilde{s}| > 1$ and $sgn(\widetilde{w}_\tau) = sgn(w^*_\tau)$ when $w^*_\tau = \pm 2$ and $\lambda < 1$.[59] But it brings ex-ante disadvantage before a herd starts since it is more likely to induce a wrong herd than what a normal agent does, due to the lower precision $\widetilde{q} < q$. Therefore the harm of release is decreasing in the probability of herding, which in turn is weakly decreasing over time and hence explains the weak monotonicity of $v(\tau;\lambda)$ in τ.

It is worth pointing out that when $\lambda > 3$, the weak monotonicity is actually uniformity. In that case, the public signal is so strong that people start to herd on the action same as the realization of $\widetilde{s}$ immediately after it is released, so releasing at different periods makes no difference.

2.2.3 Contingent Release of Noisy Public Information

In addition to the *monotonicity* result, Proposition 2.2.1 also makes another observation: it is better not to release the public signal at all when it is less precise than private signals. This can be interpreted as an *anti-transparency* result: more (but noisy) pub-

[59] Bikhchandani *et al.* (1992) argued that releasing a public signal less informative than the private signal can still be beneficial when there is an information cascade. This is true under their assumption that agents play mixed strategies when indifferent, but not under the tie-breaking rule here.

lic information can be bad for social welfare.[60] However, the social planner has so far been restricted to use exogenous releasing strategies, hence a natural question would be whether this suboptimality of release when $\lambda < 1$ still holds if *contingent* releasing strategies are allowed, namely the social planner can decide whether to release or not at period t based on the realization of $\widetilde{s}$ and w_t.

Let $g(\widetilde{s}, w_t) \in \{0, 1\}$ be the strategy of the social planner: $g(\widetilde{s}, w_t) = 1$ means the social planner releases the public signal at period t after seeing $\widetilde{s}$ and w_t; $g(\widetilde{s}, w_t) = 0$ means not. And $g_t \in \{0, 1\}$ denotes the corresponding action. We restriction attention on pure strategies by the social planner.

Note that g_t is now relevant information for agents $\tau \geq t$, because, given a releasing strategy by the social planner, agents can possibly infer the realization of $\widetilde{s}$ from g_t and w_t. A natural issue arises here, like in lots of games with incomplete information, that there could potentially exist undesired equilibria due to lack of restriction on *off-equilibrium* beliefs. So we want to impose the following refinement on certain off-equilibrium path.

Definition 2.4 *Given a releasing strategy $g(\widetilde{s}, w_t)$ by the social planner, let $\mu(w_t, g_t)$ denote agents' belief at period t about the realization of $\widetilde{s}$ after observing w_t and g_t. That is,*

$$\mu(w_t, g_t) \equiv \Pr(\widetilde{s} = 1 | w_t, g_t, g(\cdot, \cdot)).$$

We say $\mu(w_t, g_t)$ is ***non-excessive*** *if*

$$\mu(w_t, 0) = \frac{1}{2}, \ \forall w_t \ s.t. \ g(\widetilde{s}, w_t) = 1 \text{ for any } \widetilde{s} \in \{-1, 1\}.$$

Non-excessive belief requires that, on an off-equilibrium path where the social planner does not release $\widetilde{s}$ while she should have released it regardless of its realization, agents should not make excessive inference about the realization of $\widetilde{s}$ in this symmetric world. We think this is a reasonable refinement and it indeed helps us get rid of meaningless equilibria which do not serve for the purpose of our analysis here.[61]

[60] The seminal paper on anti-transparency, Morris and Shin (2002), also used the average payoff of the population to refer to social welfare.

[61] Without restriction on non-excessive beliefs, one can show that to release the public signal after any history could be an equilibrium. But these equilibria do not help improve the social welfare.

Note that the agents' behavior still follows what has been described in Lemma 2.2.1 but under a public belief generated from both the previous actions and their inference about $\widetilde{s}$. Hence we will give the main result here that focuses on the social planner's releasing strategy.

Proposition 2.2.2 *When $\lambda < 1$ and agents' belief about $\widetilde{s}$ is non-excessive, there are 3 Bayesian Nash equilibria where the social planner's contingent releasing strategies are respectively:*

$$g^1(\widetilde{s}, w_t) = 0;$$
$$g^2(\widetilde{s}, w_t) = 1_{\{w_t = -\widetilde{s}\}};$$
$$g^3(\widetilde{s}, w_t) = 1_{\{w_t = \pm 2\widetilde{s}\}}.$$

Proof. See Appendix. ■

The interesting addition compared to the case with exogenous releasing strategy is g^2. With g^2, the social planner will release the public signal once he saw an history that is not a herd yet but against the realization of $\widetilde{s}$, which is reasonable because he wants to prevent the agents from starting an herd against the public signal too early. Unfortunately, from an ex-ante perspective, social welfare is *not* improved under contingent releasing strategies.

Corollary 2.2.1 *g^1 generates the same ex-ante average payoff of the population in equilibrium as g^3 does, which is better than what g^2 does. And none of them can improve social welfare compared to exogenous releasing strategies.*

Proof. g^1 means no release at all, which is also the best the social planner can do under exogenous releasing strategies. g^3 means to disclose the public signal when a herd has already started but in that case disclosure makes no difference as the noisy public signal can never break down a herd, hence social welfare is the same as with no release at all. On the other hand, the "separating" strategy g^2 implies that the agents can perfectly infer the realization of $\widetilde{s}$ after period 1, hence social welfare is the same as with exogenous release at period 2, which is worse than with no release at all when $\lambda < 1$ as we saw in Proposition 2.2.1. ■

2.3 Postponing Disclosure Is Not Always Better

Recall that the benefit of public information disclosure is weakly increasing over time in the binary-choice model. In this section, however, we are about to introduce two alternative settings under which postponing disclosure of public information is not necessarily a good decision for the social planner.

2.3.1 Impatient Social Planner

So far we have assumed that the social planner cares about the expected average payoff of the population *without discounting*, hence she essentially cares only about whether people eventually herd on the correct action or not, *i.e.*, limiting efficiency. Although limiting efficiency is the common objective of interest in the literature of social learning, it might not be a plausible measure of social welfare for a social planner.

For example, Proposition 2.2.1 says that the social planner is indifferent among all periods to release a public signal that is sufficiently precise. Imagine that the social planner has a public signal with perfect precision. Then naturally she should release the public signal as early as possible, because any delay would hurt some earlier agents. However this natural observation is not captured by the non-discounted average payoff as the social planner only cares about people in the limit. Hence in this subsection we introduce a discount factor δ in the social planner's objective and reconsider the timing of information disclosure. In particular, we show that it is not always better to postpone the release of a public signal.

Formally, with all the other configurations identical to the benchmark model, we assume the social planner now wants to maximize the discounted sum of people's expected payoff,

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}Eu(a_t^*;\theta).$$

For simplicity we restrict attention on exogenous releasing strategies and assume the public signal has the same precision q as the private signals.

Claim 2.3.1 *Let $V(\tau)$ be the discounted sum of people's expected payoff when the public signal is released at period τ. Then $\forall q \in (\frac{1}{2}, 1)$, $V(3) > V(4)$.*

Proof. See Appendix. ■

In the benchmark without discounting, the social planner is indifferent between releasing at period 3 and at period 4.[62] Hence with discounting it is not surprising to see that the social planner now strictly prefers to release the public signal at periods 3 than period 4. The optimal timing of disclosure with discounting is yet to be explicitly characterized, nevertheless we present this example mainly to bring up the concern for the plausibility of treating limiting efficiency as the main objective in social-learning models.

2.3.2 Ternary Setting

In this subsection we expand the binary setting in the benchmark model and allow the state/action/signal space to have three elements.[63] Under this new setting, even for a patient social planner who only cares about limiting efficiency as in the benchmark model, it is not always better to postpone the disclosure of public information.

Formally, the state of the world ψ is realized out of $\{L, M, R\}$ with

$$\Pr(\psi = L) = \Pr(\psi = M) = \Pr(\psi = R) = 1/3.$$

After the realization of ψ, every agent t receives a private signal $s_t \in \{L, M, R\}$ and the private signals are conditionally i.i.d. with

$$\begin{aligned}
&\Pr\{s_t = L|\psi = L\} = \Pr\{s_t = M|\psi = M\} = \Pr\{s_t = R|\psi = R\} = q \in (\frac{1}{3}, 1);\\
&\Pr\{s_t = L|\psi = M\} = \Pr\{s_t = R|\psi = M\} = \Pr\{s_t = M|\psi = R\} =\\
&\Pr\{s_t = L|\psi = R\} = \Pr\{s_t = R|\psi = L\} = \Pr\{s_t = M|\psi = L\} = \frac{1-q}{2}.
\end{aligned}$$

Each agent chooses a_t from $\{L, M, R\}$ and observes the history of past actions $\boldsymbol{h}_t \equiv (a_1, a_2, \ldots, a_{t-1})$. They have identical utility function $u(a_t; \psi) = 1_{\{\psi = a_t\}}$, and note that this degenerate utility function implies that the three states cannot be linearly ordered,

[62] See the proof of Proposition 2.2.1 for details.

[63] Ternary spaces are sufficient to capture the intuition we want to describe, yet not too complicated for analysis.

unlike many other economic models with multiple states.[64]

With ternary spaces, we want to specify a tie-breaking rule:

$$a_t^*(\boldsymbol{h}_t, s_t) = s_t \text{ if } s_t \in \operatorname*{arg\,max}_{a\in\{L,M,R\}} E_\psi[u(a;\psi)|\boldsymbol{h}_t, s_t],$$
$$a_t^*(\boldsymbol{h}_t, s_t) = a_{t-1} \text{ if } \operatorname*{arg\,max}_{a\in\{L,M,R\}} E_\psi[u(a;\psi)|\boldsymbol{h}_t, s_t] = \Psi\backslash\{s_t\}.$$

Namely, agent t follows s_t if it is one of the maximizers and chooses to follow her immediate predecessor if the two actions different from s_t are both maximizers.[65]

Again there is a social planner who receives a signal $\widetilde{s} \in \{L, M, R\}$ and decide whether/when to release it to the public. Here for simplicity $\widetilde{s}$ is assumed to be *equally* precise as the private signals. The (patient) social planner's objective is still to maximize the ex-ante average payoff of the population and we restriction attention on *exogenous* releasing strategies only in this subsection.

Before making further observations, we would like to introduce an idea similar to herd behavior:

Definition 2.5 *$\forall a \in \{L, M, R\}$, a* ***trap*** *away from action a starts at period T if $a_t(\boldsymbol{h}_t, \cdot) \neq a$ for all $t \geq T$.*

It is easy to see that a trap away from one action is equivalent to a herd on the other action in the binary model. However, with ternary spaces, a herd on action $a' \neq a$ is a trap away from action a, but *not* vice versa. And the difference between a herd and a trap is exactly what drives the following result that the benefit of releasing the public signal is no longer weakly monotone over time.

Claim 2.3.2 *Let $\overline{G}$ be the ex-ante average payoff of the population with* no *release at all, and let $G(\tau)$ be the ex-ante average payoff of the population if the public signal is released at period $\tau \geq 1$. Then $\forall q \in (\frac{1}{3}, 1)$, $G(3) > \overline{G}$ and $G(4) > G(5)$.*

Proof. See Appendix. ∎

[64] Specifically, the degenerate utility function rules out the scenario where an agent believes one state, say M, is more likely after observing an action L and an action R. This setting, though complicates the analysis, is crucial for the result we will present in this subsection.

[65] The specification itself is not very important; we just want a tie-breaking rule to guarantee deterministic outcomes and hence tractability.

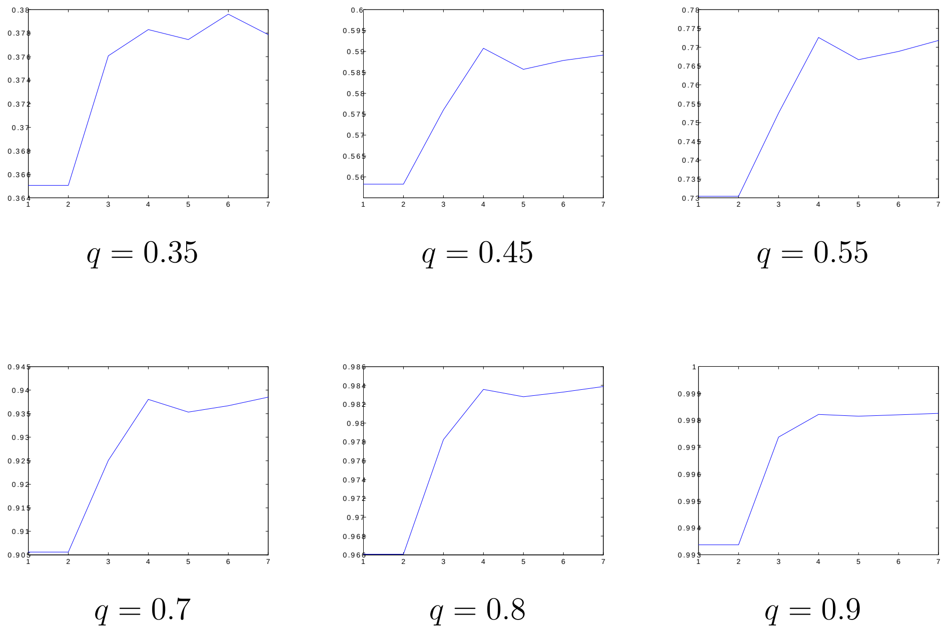

Figure 2: Non-monotone $G(\cdot)$ under different values of q.

$G(3) > \overline{G}$ is not a surprising: the public signal is equally precise as a private signal, so releasing it would not bring any harm but could possibly break down a wrong herd starting at period 3, if the first two actions are the same but wrong.[66] Meanwhile we lose weak monotonicity as $G(4) > G(5)$ for the following reason: releasing $\widetilde{s}$ at period 4 is possible to break down a trap away from true state ψ if $\widetilde{s} = \psi$, no matter what $\boldsymbol{h}_3$ is; however, if $\psi = R$ but $\boldsymbol{h}_4 = (L, M, L, L)$, the trap away from R could not be broken down even if $\widetilde{s} = R$ as long as it is released at period 5. In general, weak monotonicity fails here because the existence of traps rather than herds: a trap is not necessarily an informational cascade and information can still aggregate over time for the two "surviving" actions before a herd finally starts, hence the social planner could face the danger of not being able to break down a wrong trap if the public signal is released too late.

On the other hand, the optimal timing of release is unclear to us and in principal it shall depends on the value of q. See Figure 2 for some examples.

[66] It is not difficult to see that a herd will arise when, in the history, the number of one action is larger than the number of the other two actions by at least 2.

2.4 Conclusions

In this paper we look into the effect of public information disclosure on social learning. In the canonical binary model, if a social planner were to choose a certain period to release a public signal, she should release it as late as possible regardless of the precision of the public signal: a *monotonicity* result. Meanwhile, when the public signal is less precise than people's private signals, releasing it would do no good on social welfare even if the timing of release can be contingent on the history of actions: an *anti-transparency* result.

We present two alternative settings where the monotonicity result no longer holds. Postponing the information disclosure could be bad for a social planner, if her objective is the discounted sum of people's expected payoffs, or if the state/action spaces are richer. Solving the optimal timing of disclosure in these two settings is a challenging but interesting follow-up to this work.

As to the anti-transparency result, a relevant and interesting question is: what is the lower bound of the (relative) precision of a public signal that could improve social welfare once released in a more general setting, *e.g.*, agents have private signals of heterogeneous precision? Some preliminary work suggests that this lower bound is lower and could be substantially lower than the "average" precision of people's private signals.

We treat this work as a contribution to the literature on social learning, with a particular focus on exogenous information intervention. Perhaps more importantly, we hope this work can also raise attention on welfare control or optimal design of social learning under different settings.

2.5 Appendix

Proof of Lemma 2.2.1**.** By standard Bayesian Nash equilibrium definition and the tie-breaking rule,

$$
\begin{aligned}
& a_t^*(\boldsymbol{h}_t, s_t) \\
= \; & \arg\max_{a\in\{1,-1\}} E_\theta(1_{\{\theta=a\}}|\boldsymbol{h}_t, s_t) = \arg\max_{a\in\{1,-1\}} \Pr(\theta = a|\boldsymbol{h}_t, s_t) \\
= \; & \begin{cases} s_t & \text{if } \Pr(\theta = 1|\boldsymbol{h}_t, s_t) = \Pr(\theta = -1|\boldsymbol{h}_t, s_t) \\ sgn(\Pr(\theta = 1|\boldsymbol{h}_t, s_t) - \Pr(\theta = -1|\boldsymbol{h}_t, s_t)) & \text{otherwise} \end{cases}.
\end{aligned}
$$

By Bayes' Rule and uniform prior,

$$
\begin{aligned}
& \Pr(\theta = 1|\boldsymbol{h}_t, s_t) = \frac{\Pr(\boldsymbol{h}_t, s_t|\theta = 1)}{\Pr(\boldsymbol{h}_t, s_t|\theta = 1) + \Pr(\boldsymbol{h}_t, s_t|\theta = -1)} = 1 - \Pr(\theta = -1|\boldsymbol{h}_t, s_t) \\
\Rightarrow \; & sgn(\Pr(\theta = 1|\boldsymbol{h}_t, s_t) - \Pr(\theta = -1|\boldsymbol{h}_t, s_t)) = sgn(\Pr(\boldsymbol{h}_t, s_t|\theta = 1) - \Pr(\boldsymbol{h}_t, s_t|\theta = -1)) \\
\Rightarrow \; & a_t^*(\boldsymbol{h}_t, s_t) \\
= \; & \begin{cases} s_t & \text{if } \Pr(\boldsymbol{h}_t, s_t|\theta = 1) = \Pr(\boldsymbol{h}_t, s_t|\theta = -1) \\ sgn(\Pr(\boldsymbol{h}_t, s_t|\theta = 1) - \Pr(\boldsymbol{h}_t, s_t|\theta = -1)) & \text{otherwise} \end{cases}.
\end{aligned}
$$

By definition of w_t and independence between s_t and $\boldsymbol{h}_t$,

$$
\begin{aligned}
& \frac{\Pr(\boldsymbol{h}_t, s_t|\theta = 1)}{\Pr(\boldsymbol{h}_t, s_t|\theta = -1)} = \frac{\Pr(\boldsymbol{h}_t|\theta = 1)}{\Pr(\boldsymbol{h}_t|\theta = -1)} \frac{\Pr(s_t|\theta = 1)}{\Pr(s_t|\theta = -1)} \\
= \; & (\frac{q}{1-q})^{w_t} (\frac{q}{1-q})^{s_t} = (\frac{q}{1-q})^{w_t + s_t}, \text{ where } \frac{q}{1-q} > 1 \\
\Rightarrow \; & sgn(\Pr(\boldsymbol{h}_t, s_t|\theta = 1) - \Pr(\boldsymbol{h}_t, s_t|\theta = -1)) = sgn(w_t + s_t) = \begin{cases} s_t \text{ or } 0 & \text{if } |w_t| \leq 1 \\ sgn(w_t) & \text{otherwise} \end{cases},
\end{aligned}
$$

hence we get $a_t^*(\boldsymbol{h}_t, s_t)$ characterized in the Lemma.

On the other hand, apparently $w_1^* = 0$ and by definition of w_t^*,

$$
\begin{aligned}
(\frac{q}{1-q})^{w_{t+1}^*} &= \frac{\Pr(\boldsymbol{h}_{t+1}^*|\theta = 1)}{\Pr(\boldsymbol{h}_{t+1}^*|\theta = -1)} = \frac{\Pr(\boldsymbol{h}_t^*, a_t|\theta = 1)}{\Pr(\boldsymbol{h}_t^*, a_t|\theta = -1)} \\
&= \frac{\Pr(\boldsymbol{h}_t^*|\theta = 1)}{\Pr(\boldsymbol{h}_t^*|\theta = -1)} \frac{\Pr(a_t^*|\boldsymbol{h}_t^*, \theta = 1)}{\Pr(a_t^*|\boldsymbol{h}_t^*, \theta = -1)} \\
&= (\frac{q}{1-q})^{w_t^*} \frac{\Pr(a_t^*|\boldsymbol{h}_t^*, \theta = 1)}{\Pr(a_t^*|\boldsymbol{h}_t^*, \theta = -1)}.
\end{aligned}
$$

Meanwhile, if $|w_t^*| \leq 1$,

$$\begin{aligned} a_t^* = s_t \;\; & \Rightarrow \;\; \frac{\Pr(a_t^*|\boldsymbol{h}_t^*, \theta = 1)}{\Pr(a_t^*|\boldsymbol{h}_t^*, \theta = -1)} = \frac{\Pr(s_t|\boldsymbol{h}_t^*, \theta = 1)}{\Pr(s_t|\boldsymbol{h}_t^*, \theta = -1)} = \frac{\Pr(s_t|\theta = 1)}{\Pr(s_t|\theta = -1)} = (\frac{q}{1-q})^{s_t} \\ & \Rightarrow \;\; w_{t+1}^* = w_t^* + s_t = w_t^* + a_t^*; \end{aligned}$$

otherwise,

$$\begin{aligned} a_t^* = sgn(w_t^*) \;\; & \Rightarrow \;\; \frac{\Pr(a_t^*|\boldsymbol{h}_t^*, \theta = 1)}{\Pr(a_t^*|\boldsymbol{h}_t^*, \theta = -1)} = \frac{\Pr(sgn(w_t^*)|\boldsymbol{h}_t^*, \theta = 1)}{\Pr(sgn(w_t^*)|\boldsymbol{h}_t^*, \theta = -1)} = 1 \\ & \Rightarrow \;\; w_{t+1}^* = w_t^*. \end{aligned}$$

■

Proof of Lemma 2.2.2**.** According to Lemma 2.2.1 and Definition 2.1, a herd on action $sgn(w_t^*)$ starts at period t if and only if $|w_t^*| > 1$. Note that

$$\begin{aligned} \forall t \in \mathbb{N}^*, \; |w_t^*| \leq 1 \;\; & \Rightarrow \;\; \begin{cases} \forall k \in \mathbb{N}^*, \; a_{2k-1}^* + a_{2k}^* = 0 \\ \forall t \in \mathbb{N}^*, \; a_t^* = s_t \end{cases} \\ & \Rightarrow \;\; \forall k \in \mathbb{N}^*, \; s_{2k-1} + s_{2k} = 0. \end{aligned}$$

Hence a herd starts eventually unless $s_{2k-1} + s_{2k} = 0$, $\forall k \in \mathbb{N}^*$. However,

$$\begin{aligned} & \Pr(\forall k \in \mathbb{N}^*, \; s_{2k-1} = -s_{2k}) \leq 1 - \Pr(\exists k' \in \mathbb{N}^* \text{ s.t. } s_{k'} = s_{k'+1} = s_{k'+2}) \\ = \;\; & 1 - E_\theta[\Pr(\exists k' \in \mathbb{N}^* \text{ s.t. } s_{k'} = s_{k'+1} = s_{k'+2}|\theta)] = 1 - 1 = 0, \end{aligned}$$

where $\forall \theta \in \Theta$, $\Pr(\exists k' \in \mathbb{N}^* \text{ s.t. } s_{k'} = s_{k'+1} = s_{k'+2}|\theta) = 1$ due to $\{s_t\}_{t=1}^{\infty}$ being conditional i.i.d. and Law of Large Numbers. ■

Proof of Lemma 2.2.3**.** By Lemma 2.2.1, a herd starts in equilibrium when $|w_t^*| > 1$ and the herd is correct(wrong) if $w_t^* > 1(< -1)$. Then we can immediately see that $\alpha_1 = 0$ and $\alpha_7 = 1$. For the remaining cases, let us look into the transition of w_t^θ. (Recall that we have assumed that the realization of θ is 1 without loss of generality)

If $w_1 = 0$,

$$\begin{aligned} & a_1^* = s_1 \text{ and } w_2^* = s_1 \\ \Rightarrow \;& \Pr(w_2^* = 1 | w_1 = 0) = \Pr(s_1 = 1) = q, \\ & \Pr(w_2^* = -1 | w_1 = 0) = \Pr(s_1 = -1) = 1 - q. \end{aligned}$$

Note that $\forall T < \infty$ (especially $T = 2$ here),

$$\begin{aligned} p(x) &\equiv \Pr(\text{plim}_{t\to\infty} a_t^*(\boldsymbol{h}_t^*, s_t) = \theta | w_1 = x) \\ &= \Pr(\text{plim}_{t\to\infty} a_t^*(\boldsymbol{h}_t^*, s_t) = \theta | w_T = x) \end{aligned}$$

since the subgame starting from agent T is identical to the original game, thus we have $\alpha_4 = (1-q)\alpha_2 + q\alpha_6$. Through similar arguments,

$$\begin{aligned} & \Pr(w_2^* = 2 | w_1 = 1) = q \text{ and } \Pr(w_2^* = 0 | w_1 = 1) = 1 - q \\ \Rightarrow \;& \alpha_6 = q\alpha_7 + (1-q)\alpha_4 = q + (1-q)\alpha_4; \\ & \Pr(w_2^* = 0 | w_1 = -1) = q \text{ and } \Pr(w_2^* = -2 | w_1 = -1) = 1 - q \\ \Rightarrow \;& \alpha_2 = q\alpha_4 + (1-q)\alpha_1 = q\alpha_4. \end{aligned}$$

Solve the three linear equations together to get α_2, α_4 and α_6 as stated in the Lemma.

If $w_1 = x \in (-1, 0)$,

$$\begin{aligned} & a_1^* = s_1 \text{ and } w_2^* = x + s_1 \\ \Rightarrow \;& \Pr(w_2^* = x + 1 \in (0,1) | w_1 = x) = q, \\ & \Pr(w_2^* = x - 1 < -1 | w_1 = x) = 1 - q \\ \Rightarrow \;& \alpha_3 = q\alpha_5 + (1-q)\alpha_1 = q\alpha_5; \end{aligned}$$

through similar argument,

$$\begin{aligned}
&\Pr(w_2^* = x' + 1 > 1 | w_1 = x' \in (0,1)) = q, \\
&\Pr(w_2^* = x' - 1 \in (-1,0) | w_1 = x' \in (0,1)) = 1 - q \\
\Rightarrow\ &\alpha_5 = q\alpha_7 + (1-q)\alpha_3 = q + (1-q)\alpha_3.
\end{aligned}$$

Solve the two linear equations together to get α_3 and α_5 as stated in the Lemma.

Combing all these linear equations together, we have exactly $Q\boldsymbol{\alpha} = \boldsymbol{\alpha}$. In other words, $\boldsymbol{\alpha}$ is an eigenvector of P associated with eigenvalue 1, with restriction that $\alpha_7 = 1$ and $\alpha_1 = 0$. ■

Proof of Lemma 2.2.4. $\boldsymbol{\pi}^1 = (0, 0, 0, 1, 0, 0, 0)$ simply because $w_1^* = 0$. As to the transition between w_t^* to w_{t+1}^*, it is identical to the transition between w_1^* and w_2^* illustrated in the proof of Lemma 2.2.3:

$$\begin{aligned}
&\Pr(w_{t+1}^* > 1 | w_t > 1) = 1,\ \Pr(w_{t+1}^* < -1 | w_t^* < -1) = 1; \\
&\Pr(w_{t+1}^* = 1 | w_t^* = 0) = q,\ \Pr(w_{t+1}^* = -1 | w_t^* = 0) = 1 - q; \\
&\Pr(w_{t+1}^* = 2 | w_t^* = 1) = q,\ \Pr(w_{t+1}^* = 0 | w_t^* = 1) = 1 - q; \\
&\Pr(w_{t+1}^* = 0 | w_t^* = -1) = q,\ \Pr(w_{t+1}^* = -2 | w_t^* = 0) = 1 - q; \\
&\Pr(w_{t+1}^* \in (0,1) | w_t^* \in (-1,0)) = q,\ \Pr(w_{t+1}^* < -1 | w_t^* \in (-1,0)) = 1 - q; \\
&\Pr(w_{t+1}^* > 1 | w_t^* \in (0,1)) = q,\ \Pr(w_{t+1}^* \in (-1,0) | w_t^* \in (0,1)) = 1 - q.
\end{aligned}$$

Therefore the transition matrix is exactly matrix Q, which is indeed monotone according to Definition 2.2.

Note that by Lemma 2.2.1, w_t^* can only take values ± 2, ± 1 and 0, hence $\pi_3^t = \pi_5^t = 0$ for any t and $w_t^* > 1(< -1)$ indicates $w_t^* = 2(-2)$. ■

Proof of Lemma 2.2.5. By definition of public beliefs,

$$(\frac{q}{1-q})^{w_\tau} = \frac{\Pr(\boldsymbol{h}_\tau | \theta = 1)}{\Pr(\boldsymbol{h}_\tau | \theta = -1)} \text{ and } (\frac{q}{1-q})^{\widetilde{w}_\tau} = \frac{\Pr(\boldsymbol{h}_\tau, \widetilde{s} | \theta = 1)}{\Pr(\boldsymbol{h}_\tau, \widetilde{s} | \theta = -1)}.$$

Since $\boldsymbol{h}_t$ and $\widetilde{s}$ are independent,

$$\begin{aligned}(\frac{q}{1-q})^{\widetilde{w}_\tau} &= \frac{\Pr(\boldsymbol{h}_\tau, \widetilde{s}|\theta=1)}{\Pr(\boldsymbol{h}_\tau, \widetilde{s}|\theta=-1)} \\ &= \frac{\Pr(\boldsymbol{h}_\tau|\theta=1)}{\Pr(\boldsymbol{h}_\tau|\theta=-1)} \cdot \frac{\Pr(\widetilde{s}|\theta=1)}{\Pr(\widetilde{s}|\theta=-1)} \\ &= (\frac{q}{1-q})^{w_\tau}(\frac{\widetilde{q}}{1-\widetilde{q}})^{\widetilde{s}} = (\frac{q}{1-q})^{w_\tau}(\frac{q}{1-q})^{\lambda\widetilde{s}}.\end{aligned}$$

Hence we have $\widetilde{w}_\tau = w_\tau + \lambda\widetilde{s}$.

On the other hand,

$$\begin{aligned}& \text{plim}_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} u(a_t^*; \theta=1|\widetilde{w}_\tau) \\ = \; & \text{plim}_{T'\to\infty} \frac{1}{T'}\sum_{t=\tau}^{\tau+T'} u(a_t^*; \theta=1|\widetilde{w}_\tau) \\ = \; & \Pr(\text{plim}_{t\to\infty} a_t^* = 1|w_1 = \widetilde{w}_\tau) = p(\widetilde{w}_\tau),\end{aligned}$$

where the last equation comes from Lemma 2.2.3. Finally, the (unconditional) expected average payoff is just

$$v(\tau;\lambda) = E_{\widetilde{w}_\tau} p(\widetilde{w}_\tau) = E_{w_\tau,\widetilde{s}}[p(w_\tau + \lambda\widetilde{s})].$$

■

Proof of Proposition 2.2.1. For $\tau \geq 1$, let $\widetilde{\boldsymbol{\pi}}^\tau = (\widetilde{\pi}_1^\tau, \widetilde{\pi}_2^\tau, \widetilde{\pi}_3^\tau, \widetilde{\pi}_4^\tau, \widetilde{\pi}_5^\tau, \widetilde{\pi}_6^\tau, \widetilde{\pi}_7^\tau)$ be the probability vector of $\widetilde{w}_\tau$ on the partition $\mathcal{P}$ introduced in Lemma 2.2.3. Then by Lemma 2.2.3 we have

$$E_{\widetilde{w}_\tau}(p(\widetilde{w}_\tau)) = \widetilde{\boldsymbol{\pi}}^\tau \cdot \boldsymbol{\alpha}.$$

Note that the public information is irrelevant for agents before period τ so in equilibrium w_t^* for $t \leq \tau$ still evolves according to Lemma 2.2.4. Bearing in mind as well that $\widetilde{s}$ is independent of w_t^* and distributed according to

$$\Pr\{\widetilde{s}=1|\theta=1\}=\Pr\{\widetilde{s}=-1|\theta=-1\}=\widetilde{q},$$

we can derive $\widetilde{\boldsymbol{\pi}}^\tau$ explicitly and prove the proposition case by case on λ as follows: (we then omit the argument λ in $v(\cdot;\cdot)$ in each case)

Case I ($0<\lambda<1 \Leftrightarrow \frac{1}{2}<\widetilde{q}<q$)

$$\begin{aligned}
&\widetilde{w}_1 = w_1^* + \lambda\widetilde{s} \text{ and } \boldsymbol{\pi}^1 = (0,0,0,1,0,0,0)\\
\Rightarrow\ &\widetilde{\boldsymbol{\pi}}^1 = (0,0,1-\widetilde{q},0,\widetilde{q},0,0)\\
\Rightarrow\ &v(1) = \widetilde{q}\alpha_5 + (1-\widetilde{q})\alpha_3 = \frac{\widetilde{q}q + q^2(1-\widetilde{q})}{1-q(1-q)} < \frac{2q^2-q^3}{1-q(1-q)} < \alpha_4;\\
&\widetilde{w}_2 = w_2^* + \lambda\widetilde{s} \text{ and } \pi^2 = (0,1-q,0,0,0,0,q,0)\\
\Rightarrow\ &\widetilde{\boldsymbol{\pi}}^1 = ((1-q)(1-\widetilde{q}),0,(1-q)\widetilde{q},0,q(1-\widetilde{q}),0,q\widetilde{q})\\
\Rightarrow\ &v(2) = (1-q)\widetilde{q}\alpha_3 + q(1-\widetilde{q})\alpha_5 + q\widetilde{q} = \frac{q\widetilde{q} + q^2(1-\widetilde{q})}{1-q(1-q)} = v(1) < \alpha_4;\\
&\widetilde{w}_\tau = w_\tau^* + \lambda\widetilde{s} \text{ and } \boldsymbol{\pi}^\tau = 2q(1-q)\boldsymbol{\pi}^{\tau-2} + ((1-q)^2,0,0,0,0,0,q^2) \text{ for } \tau\geq 3\\
\Rightarrow\ &\widetilde{\boldsymbol{\pi}}^\tau = 2q(1-q)\widetilde{\boldsymbol{\pi}}^{\tau-2} + ((1-q)^2,0,0,0,0,0,q^2);\\
\Rightarrow\ &v(\tau) = 2q(1-q)v(\tau-2) + q^2\\
\Rightarrow\ &sgn(v(\tau+1)-v(\tau)) = sgn(v(\tau-1)-v(\tau-2)).\\
&\text{If } v(\tau-2) < \alpha_4 = \frac{q^2}{1-2q(1-q)}\\
\Rightarrow\ &v(\tau-2) < 2q(1-q)v(\tau-2) + q^2 = v(\tau) < 2q(1-q)\alpha_4 + q^2 = \alpha_4\\
\Rightarrow\ &v(1) = v(2) < v(3) = v(4) < \alpha_4.
\end{aligned}$$

Recursively we have $v(1)=v(2)<v(3)=v(4)<v(5)=v(6)<\cdots$ and $v(\tau)<\alpha_4$ for any $\tau\geq 1$.

Case II ($\lambda = 1 \Leftrightarrow \widetilde{q} = q$)

$$\widetilde{w}_1 = w_1^* + \lambda \widetilde{s} \text{ and } \boldsymbol{\pi}^1 = (0,0,0,1,0,0,0)$$

$$\Rightarrow \widetilde{\boldsymbol{\pi}}^1 = (0, 1-\widetilde{q}, 0, 0, 0, \widetilde{q}, 0)$$

$$\Rightarrow v(1) = \widetilde{q}\alpha_6 + (1-\widetilde{q})\alpha_2 = q(q + \frac{(1-q)q^2}{1-2q(1-q)}) + \frac{(1-q)q^3}{1-2q(1-q)} = \alpha_4;$$

$$\widetilde{w}_2 = w_2^* + \lambda \widetilde{s} \text{ and } \pi^2 = (0, 1-q, 0, 0, 0, 0, q, 0)$$

$$\Rightarrow \widetilde{\boldsymbol{\pi}}^2 = ((1-q)(1-\widetilde{q}), 0, 0, (1-q)\widetilde{q} + q(1-\widetilde{q}), 0, , 0, q\widetilde{q})$$

$$\Rightarrow v(2) = [(1-q)\widetilde{q} + q(1-\widetilde{q})]\alpha_4 + q\widetilde{q} = 2q(1-q)\alpha_4 + q^2 = \alpha_4;$$

$$\widetilde{w}_\tau = w_\tau^* + \lambda \widetilde{s} \text{ and } \boldsymbol{\pi}^\tau = 2q(1-q)\boldsymbol{\pi}^{\tau-2} + ((1-q)^2, 0, 0, 0, 0, 0, q^2) \text{ for } \tau \geq 3$$

$$\Rightarrow \widetilde{\boldsymbol{\pi}}^\tau = 2q(1-q)\widetilde{\boldsymbol{\pi}}^{\tau-2} + ((1-\widetilde{q})(1-q)^2, \widetilde{q}(1-q)^2, 0, 0, 0, (1-\widetilde{q})q^2, \widetilde{q}q^2)$$

$$\Rightarrow v(\tau) = 2q(1-q)v(\tau-2) + \widetilde{q}(1-q)^2\alpha_2 + (1-\widetilde{q})q^2\alpha_6 + \widetilde{q}q^2 =$$

$$2q(1-q)v(\tau-2) + (1-q)q\alpha_4 + q^3 > 2q(1-q)v(\tau-2) + q^2$$

$$\Rightarrow \; sgn(v(\tau+1) - v(\tau)) = \; sgn(v(\tau-1) - v(\tau-2)).$$

$$\text{If } v(\tau-2) \geq \alpha_4 \Rightarrow v(\tau) > 2q(1-q)v(\tau-2) + q^2 \geq \alpha_4$$

$$\Rightarrow v(4) = v(3) > v(2) = v(1) = \alpha_4.$$

Recursively we have $\alpha_4 = v(1) = v(2) < v(3) = v(4) < v(5) = v(6) < \cdots$.

Case III ($1 < \lambda < 2 \Leftrightarrow q < \widetilde{q} < \frac{q^2}{q^2+(1-q)^2} = \alpha_4$)

$$
\begin{aligned}
& \widetilde{w}_1 = w_1^* + \lambda\widetilde{s} \text{ and } \boldsymbol{\pi}^1 = (0,0,0,1,0,0,0) \\
\Rightarrow\ & \widetilde{\boldsymbol{\pi}}^1 = (1-\widetilde{q},0,0,0,0,0,\widetilde{q}) \\
\Rightarrow\ & v(1) = \widetilde{q} < \alpha_4; \\
& \widetilde{w}_2 = w_2^* + \lambda\widetilde{s} \text{ and } \pi^2 = (0,1-q,0,0,0,0,q,0) \\
\Rightarrow\ & \widetilde{\boldsymbol{\pi}}^2 = ((1-q)(1-\widetilde{q}),0,q(1-\widetilde{q}),0,(1-q)\widetilde{q},0,q\widetilde{q}) \\
\Rightarrow\ & v(2) = (1-q)\widetilde{q}\alpha_5 + q(1-\widetilde{q})\alpha_3 + q\widetilde{q} = \frac{2q(1-q)\widetilde{q}+q^3}{1-q(1-q)} \\
\Rightarrow\ & v(2)-v(1) = \frac{q^3(1-\widetilde{q})-(1-q)^3\widetilde{q}}{1-q(1-q)} > 0 \text{ as } \widetilde{q} < \alpha_4 < \frac{q^3}{q^3+(1-q)^3}; \\
& \widetilde{w}_\tau = w_\tau^* + \lambda\widetilde{s} \text{ and } \boldsymbol{\pi}^\tau = 2q(1-q)\boldsymbol{\pi}^{\tau-2} + ((1-q)^2,0,0,0,0,0,q^2) \text{ for } \tau \geq 3 \\
\Rightarrow\ & \widetilde{\boldsymbol{\pi}}^\tau = 2q(1-q)\widetilde{\boldsymbol{\pi}}^{\tau-2} + ((1-\widetilde{q})(1-q)^2,0,\widetilde{q}(1-q)^2,0,(1-\widetilde{q})q^2,0,\widetilde{q}q^2) \\
\Rightarrow\ & v(\tau) = 2q(1-q)v(\tau-2) + \widetilde{q}(1-q)^2\alpha_3 + (1-\widetilde{q})q^2\alpha_5 + \widetilde{q}q^2 \\
\Rightarrow\ & sgn(v(\tau+1)-v(\tau)) = sgn(v(\tau-1)-v(\tau-2)). \\
& v(3) = 2q(1-q)v(1) + \widetilde{q}(1-q)^2\alpha_3 + (1-\widetilde{q})q^2\alpha_5 + \widetilde{q}q^2 \text{ with } v(1) = \widetilde{q} \\
\Rightarrow\ & v(3)-v(2) = q(1-q)\widetilde{q} - \widetilde{q}(1-q)q = 0 \Rightarrow v(3) = v(2) > v(1).
\end{aligned}
$$

Recursively we have $v(1) < v(2) = v(3) < v(4) = v(5) < v(6) < \cdots$. As the sequence $\{v(\tau)\}_{\tau=1}^{\infty}$ is weakly monotonic and bounded between 0 and 1, $\lim_{\tau\to\infty} v(\tau)$ exists and it satisfies

$$
\begin{aligned}
& \lim_{\tau\to\infty} v(\tau) = 2q(1-q)\lim_{\tau\to\infty} v(\tau) + \widetilde{q}(1-q)^2\alpha_3 + (1-\widetilde{q})q^2\alpha_5 + \widetilde{q}q^2 \\
\Rightarrow\ & \lim_{\tau\to\infty} v(\tau) = \frac{\widetilde{q}(1-q)^2\alpha_3 + (1-\widetilde{q})q^2\alpha_5 + \widetilde{q}q^2}{1-2q(1-q)} > \\
& \frac{q(1-q)^2\alpha_3 + (1-q)q^2\alpha_5 + q^3}{1-2q(1-q)} > \frac{q^2}{1-2q(1-q)} = \alpha_4.
\end{aligned}
$$

Thus $\exists T < \infty$ s.t. $v(T) > \alpha_4$.

Case IV ($\lambda = 2 \Leftrightarrow \widetilde{q} = \frac{q^2}{q^2+(1-q)^2} = \alpha_4$)

Similar to Case III, we have $v(1) = \widetilde{q} = \alpha_4$;

$\widetilde{w}_2 = w_2^* + \lambda\widetilde{s}$ and $\pi^2 = (0, 1-q, 0, 0, 0, 0, q, 0)$

$\Rightarrow$ $\widetilde{\boldsymbol{\pi}}^2 = ((1-q)(1-\widetilde{q}), q(1-\widetilde{q}), 0, 0, 0, (1-q)\widetilde{q}, q\widetilde{q})$

$\Rightarrow$ $v(2) = (1-q)\widetilde{q}\alpha_6 + q(1-\widetilde{q})\alpha_2 + q\widetilde{q} = (1-q)\alpha_4\alpha_6 + q(1-\alpha_4)\alpha_2 + q\alpha_4 =$

$\alpha_4[\alpha_4 + 2q(1-\alpha_4)] > \alpha_4$;

$\widetilde{w}_\tau = w_\tau^* + \lambda\widetilde{s}$ and $\boldsymbol{\pi}^\tau = 2q(1-q)\boldsymbol{\pi}^{\tau-2} + ((1-q)^2, 0, 0, 0, 0, 0, q^2)$ for $\tau \geq 3$

$\Rightarrow$ $\widetilde{\boldsymbol{\pi}}^\tau = 2q(1-q)\widetilde{\boldsymbol{\pi}}^{\tau-2} + ((1-\widetilde{q})(1-q)^2, 0, 0, \widetilde{q}(1-q)^2 + (1-\widetilde{q})q^2, 0, 0, \widetilde{q}q^2)$

$\Rightarrow$ $v(\tau) = 2q(1-q)v(\tau-2) + [\widetilde{q}(1-q)^2 + (1-\widetilde{q})q^2]\alpha_4 + \widetilde{q}q^2$

$\Rightarrow$ $sgn(v(\tau+1) - v(\tau)) = sgn(v(\tau-1) - v(\tau-2))$.

$v(3) = 2q(1-q)v(1) + [\widetilde{q}(1-q)^2 + (1-\widetilde{q})q^2]\alpha_4 + \widetilde{q}q^2$ with $v(1) = \widetilde{q} = \alpha_4$

$\Rightarrow$ $v(3) = [2q(1-q) + \alpha_4(1-q)^2 + (1-\alpha_4)q^2 + q^2]\alpha_4 = v(2) > v(1)$.

Recursively we have $\alpha_4 = v(1) < v(2) = v(3) < v(4) = v(5) < v(6) = \cdots$.

Case V ($2 < \lambda < 3 \Leftrightarrow \alpha_4 = \frac{q^2}{q^2+(1-q)^2} < \widetilde{q} < \frac{q^3}{q^3+(1-q)^3}$)

Similar to Case III, we have $v(1) = \widetilde{q} > \alpha_4$;

$\widetilde{w}_2 = w_2^* + \lambda\widetilde{s}$ and $\pi^2 = (0, 1-q, 0, 0, 0, 0, q, 0)$

$\Rightarrow$ $\widetilde{\boldsymbol{\pi}}^2 = ((1-\widetilde{q}), 0, 0, 0, 0, 0, \widetilde{q})$

$\Rightarrow$ $v(2) = \widetilde{q} = v(1)$;

$\widetilde{w}_\tau = w_\tau^* + \lambda\widetilde{s}$ and $\boldsymbol{\pi}^\tau = 2q(1-q)\boldsymbol{\pi}^{\tau-2} + ((1-q)^2, 0, 0, 0, 0, 0, q^2)$ for $\tau \geq 3$

$\Rightarrow$ $\widetilde{\boldsymbol{\pi}}^\tau = 2q(1-q)\widetilde{\boldsymbol{\pi}}^{\tau-2} + ((1-\widetilde{q})(1-q)^2, 0, (1-\widetilde{q})q^2, 0, \widetilde{q}(1-q)^2, 0, \widetilde{q}q^2)$

$\Rightarrow$ $v(\tau) = 2q(1-q)U(\tau-2) + \widetilde{q}(1-q)^2\alpha_5 + (1-\widetilde{q})q^2\alpha_3 + \widetilde{q}q^2$

$\Rightarrow$ $sgn(v(\tau+1) - v(\tau)) = sgn(v(\tau-1) - v(\tau-2))$.

$v(3) = 2q(1-q)v(1) + \widetilde{q}(1-q)^2\alpha_5 + (1-\widetilde{q})q^2\alpha_3 + \widetilde{q}q^2$ with $v(1) = \widetilde{q}$

$\Rightarrow$ $v(3) - \widetilde{q} = -(1-q)^2\widetilde{q} + \alpha_5(1-q)^2\widetilde{q} + (1-\widetilde{q})\alpha_3 q^2 = \dfrac{q^4(1-\widetilde{q}) - (1-q)^4\widetilde{q}}{1-q(1-q)} > 0$

$\Rightarrow$ $v(3) > \widetilde{q} = v(2) = v(1)$.

Recursively we have $\alpha_4 < v(1) = v(2) < v(3) = v(4) < v(5) = v(6) < \cdots$.

Case VI ($\lambda = 3 \Leftrightarrow \widetilde{q} = \frac{q^3}{q^3+(1-q)^3}$)

$$
\begin{aligned}
& \text{Similar to Case V, we have } v(1) = v(2) = \widetilde{q} > \alpha_4; \\
& \widetilde{w}_\tau = w^*_\tau + \lambda \widetilde{s} \text{ and } \boldsymbol{\pi}^\tau = 2q(1-q)\boldsymbol{\pi}^{\tau-2} + ((1-q)^2, 0, 0, 0, 0, 0, q^2) \text{ for } \tau \geq 3 \\
\Rightarrow\ & \widetilde{\boldsymbol{\pi}}^\tau = 2q(1-q)\widetilde{\boldsymbol{\pi}}^{\tau-2} + ((1-\widetilde{q})(1-q)^2, (1-\widetilde{q})q^2, 0, 0, 0, \widetilde{q}(1-q)^2, \widetilde{q}q^2) \\
\Rightarrow\ & v(\tau) = 2q(1-q)v(\tau-2) + \widetilde{q}(1-q)^2\alpha_6 + (1-\widetilde{q})q^2\alpha_2 + \widetilde{q}q^2 \\
\Rightarrow\ & sgn(v(\tau+1) - v(\tau)) = sgn(v(\tau-1) - v(\tau-2)). \\
& v(3) = 2q(1-q)v(1) + \widetilde{q}(1-q)^2\alpha_6 + (1-\widetilde{q})q^2\alpha_2 + \widetilde{q}q^2 \text{ with } v(1) = \widetilde{q} \\
\Rightarrow\ & v(3) - \widetilde{q} = -(1-q)^2\widetilde{q} + \alpha_6(1-q)^2\widetilde{q} + (1-\widetilde{q})\alpha_2 q^2 = \frac{q^5(1-\widetilde{q}) - (1-q)^5\widetilde{q}}{1-2q(1-q)} > 0 \\
\Rightarrow\ & v(3) > \widetilde{q} = v(2) = v(1).
\end{aligned}
$$

Recursively we have $\alpha_4 < v(1) = v(2) < v(3) = v(4) < v(5) = v(6) < \cdots$.

Case VII ($\lambda > 3 \Leftrightarrow \widetilde{q} > \frac{q^3}{q^3+(1-q)^3}$)

$$
\begin{aligned}
\forall \tau \geq\ & 1, \widetilde{w}_\tau = w^*_\tau + \lambda \widetilde{s} \text{ with } |w^*_\tau| \leq 2 \text{ by Lemma 2.2.1} \\
\Rightarrow\ & |\widetilde{w}_\tau| > 1 \text{ and } sgn(\widetilde{w}_\tau) = sgn(\widetilde{s}) \text{ as } \lambda > 3 \\
\Rightarrow\ & \widetilde{\boldsymbol{\pi}}^\tau = ((1-\widetilde{q}), 0, 0, 0, 0, 0, \widetilde{q}) \\
\Rightarrow\ & v(\tau) = \widetilde{q} > \alpha_4.
\end{aligned}
$$

Thus we have $\alpha_4 < v(1) = v(2) = v(3) = v(4) = \cdots$. ■

Proof of Proposition 2.2.2**.** Firstly, note that *before* the period when the social planner would release the signal according to his releasing strategy, the equilibrium public belief w^*_t still evolves according to Lemma 2.2.1 and $w^*_t \in \{-2, -1, 0, 1, 2\}$. Thus for the social planner, whether to release the public signal or not depends on just five scenarios:

$$w_t = 2\widetilde{s},\ w_t = \widetilde{s},\ w_t = 0,\ w_t = -\widetilde{s},\ w_t = -2\widetilde{s}.$$

Note also that once the public signal is released or *fully inferred* by the agents, the subgame after that is just the standard case without public information but with an initial public belief $\widehat{w}_t = w_t^* + \lambda\widetilde{s}$, and social welfare is just $p(\widehat{w}_t)$ according to Lemma 2.2.3 as a herd starts eventually.

Suppose agents believe $g(\widetilde{s}, w_t) = 0$ is the releasing strategy of the social planner:

$$
\begin{aligned}
& \text{If } w_t^* = \pm 2\widetilde{s} \\
\Rightarrow\ & \text{releasing } \widetilde{s} \text{ would not break down the herd since } 0 < \lambda < 1 \\
\Rightarrow\ & \text{makes no difference.} \\
& \text{If } w_t^* = 0 \\
\Rightarrow\ & \text{releasing } \widetilde{s} \text{ makes } \widehat{w}_t = \lambda\widetilde{s} \\
\Rightarrow\ & g(\lambda\widetilde{s}) = \widetilde{q}\alpha_5 + (1-\widetilde{q})\alpha_3 < \alpha_4 \text{ as } \widetilde{q} < q; \\
& \text{without release, } \widehat{w}_t = w_t^* = 0 \\
\Rightarrow\ & g(0) = \alpha_4 \Rightarrow \text{not a profitable deviation.}
\end{aligned}
$$

$$
\begin{aligned}
& \text{If } w_t^* = \widetilde{s} \\
\Rightarrow\ & \text{releasing } \widetilde{s} \text{ makes } \widehat{w}_t = (\lambda+1)\widetilde{s} \\
\Rightarrow\ & p(\widehat{w}_t) = \frac{\widetilde{q}q}{\widetilde{q}q + (1-q)(1-\widetilde{q})}\alpha_7 + \frac{(1-\widetilde{q})(1-q)}{\widetilde{q}q + (1-q)(1-\widetilde{q})}\alpha_1 = \frac{\widetilde{q}q}{\widetilde{q}q + (1-q)(1-\widetilde{q})}; \\
& \text{without release, } \widehat{w}_t = w_t^* = \widetilde{s} \\
\Rightarrow\ & p(\widehat{w}_t) = \frac{\widetilde{q}q}{\widetilde{q}q + (1-q)(1-\widetilde{q})}\alpha_6 + \frac{(1-\widetilde{q})(1-q)}{\widetilde{q}q + (1-q)(1-\widetilde{q})}\alpha_2 \\
>\ & \frac{\widetilde{q}q}{\widetilde{q}q + (1-q)(1-\widetilde{q})} \text{ as } \widetilde{q} < q \\
\Rightarrow\ & \text{not a profitable deviation.}
\end{aligned}
$$

$$
\begin{aligned}
& \text{If } w_t^* = -\widetilde{s} \\
\Rightarrow\ & \text{releasing } \widetilde{s} \text{ makes } \widehat{w}_t = (\lambda - 1)\widetilde{s} \\
\Rightarrow\ & p(\widehat{w}_t) = \frac{\widetilde{q}(1-q)}{\widetilde{q}(1-q) + q(1-\widetilde{q})}\alpha_3 + \frac{q(1-\widetilde{q})}{\widetilde{q}(1-q) + q(1-\widetilde{q})}\alpha_5; \\
& \text{without release, } \widehat{w}_t = w_t^* = -\widetilde{s} \\
\Rightarrow\ & p(\widehat{w}_t) = \frac{\widetilde{q}(1-q)}{\widetilde{q}(1-q) + q(1-\widetilde{q})}\alpha_2 + \frac{q(1-\widetilde{q})}{\widetilde{q}(1-q) + q(1-\widetilde{q})}\alpha_6; \\
& \frac{\widetilde{q}(1-q)}{\widetilde{q}(1-q) + q(1-\widetilde{q})}(\alpha_2 - \alpha_3) - \frac{q(1-\widetilde{q})}{\widetilde{q}(1-q) + q(1-\widetilde{q})}(\alpha_5 - \alpha_6) \\
=\ & \frac{q}{\widetilde{q}(1-q) + q(1-\widetilde{q})} \frac{q}{1 - q(1-q)} \frac{(1-q)^2}{1 - 2q(1-q)}[q^2(1-\widetilde{q}) - \widetilde{q}(1-q)^2] \\
>\ & 0 \text{ as } \widetilde{q} < \frac{q^2}{q^2 + (1-q)^2} \\
\Rightarrow\ & \text{ not a profitable deviation.}
\end{aligned}
$$

Therefore $g^1(\widetilde{s}, w_t) = 0$ is indeed an equilibrium strategy of the social planner.

Suppose agents believe $g(\widetilde{s}, w_t) = 1_{\{w_t = \pm 2\widetilde{s}\}}$ is the releasing strategy of the social planner:

$$
\begin{aligned}
& \text{If releasing, } \widehat{w}_t = w_t^* + \lambda\widetilde{s} \text{ and } |\widehat{w}_t| > 1 \\
\Rightarrow\ & p(\widehat{w}_t) = sgn(\widehat{w}_t) = sgn(w_t^*); \\
& \text{if no release} \\
\Rightarrow\ & \text{ by Bayes Rule, } \widehat{w}_t = w_t^* - \lambda\widetilde{s} \text{ and } |\widehat{w}_t| > 1 \\
\Rightarrow\ & p(\widehat{w}_t) = sgn(\widehat{w}_t) = sgn(w_t^*) \\
\Rightarrow\ & \text{makes no difference;} \\
& \text{if not to release when } w_t^* = \pm 2\widetilde{s} \text{ but releasing later at } t' > t \\
\Rightarrow\ & \text{ by Bayes Rule, } \widehat{w}_k = w_t^* - \lambda\widetilde{s} \text{ and } |\widehat{w}_k| > 1 \text{ for } k = t, t+1, \ldots, t'-1, \\
& \widehat{w}_{t'} = \widehat{w}_{t'-1} + 2\lambda\widetilde{s} = w_t^* + \lambda\widetilde{s} \text{ and } |\widehat{w}_{t'}| > 1 \\
\Rightarrow\ & p(\widehat{w}_{t'}) = sgn(\widehat{w}_{t'}) = sgn(w_t^*) \\
\Rightarrow\ & \text{makes no difference.}
\end{aligned}
$$

Note that to release earlier at $t'' < t$ when $|w_{t''}| \leq 1$ is also not profitable because the original strategy at t'' is not to release until $w_t = \pm 2\widetilde{s}$ later, which is equivalent to not to release at all since $\lambda < 1$ and has been shown above to give better outcome. Therefore $g^3(\widetilde{s}, w_t) = 1_{\{w_t = \pm 2\widetilde{s}\}}$ is indeed an equilibrium strategy of the social planner.

Suppose agents believe $g(\widetilde{s}, w_t) \equiv 1_{\{w_t = 0\}}$ is the releasing strategy of the social planner:

$$\begin{aligned}
& \text{If releasing, } \widehat{w}_t = \lambda \widetilde{s} \\
\Rightarrow\ & p(\widehat{w}_t) = \widetilde{q}\alpha_5 + (1 - \widetilde{q})\alpha_3 < \alpha_4; \\
& \text{if no release at all} \\
\Rightarrow\ & \text{by non-excessive belief, } \widehat{w}_t = 0 \text{ and } p(0) = \alpha_4; \\
\Rightarrow\ & \text{it is a profitable deviation.}
\end{aligned}$$

Therefore $g(\widetilde{s}, w_t) \equiv 1_{\{w_t = 0\}}$ is not an equilibrium strategy of the social planner.

Suppose agents believe $g(\widetilde{s}, w_t) = 1_{\{w_t = \widetilde{s}\}}$ is the releasing strategy of the social planner:

$$\begin{aligned}
& \text{If releasing, } \widehat{w}_t = (1 + \lambda)\widetilde{s} \\
\Rightarrow\ & p(\widehat{w}_t) = \frac{\widetilde{q}q}{\widetilde{q}q + (1 - q)(1 - \widetilde{q})}; \\
& \text{if no release} \\
\Rightarrow\ & \text{by Bayes Rule, } \widehat{w}_t = (1 - \lambda)\widetilde{s} \\
\Rightarrow\ & p(\widehat{w}_t) = \frac{\widetilde{q}q}{\widetilde{q}q + (1 - q)(1 - \widetilde{q})}\alpha_5 + \frac{(1 - \widetilde{q})(1 - q)}{\widetilde{q}q + (1 - q)(1 - \widetilde{q})}\alpha_3 \\
>\ & \frac{\widetilde{q}q}{\widetilde{q}q + (1 - q)(1 - \widetilde{q})} \text{ as } \widetilde{q} < q \\
\Rightarrow\ & \text{it is a profitable deviation.}
\end{aligned}$$

Therefore $g(\widetilde{s}, w_t) = 1_{\{w_t = \widetilde{s}\}}$ is not an equilibrium strategy of the social planner.

Suppose agents believe $g(\widetilde{s}, w_t) = 1_{\{w_t = -\widetilde{s}\}}$ is the releasing strategy of the social planner:

$$
\begin{aligned}
& \text{If releasing, } \widehat{w}_t = (\lambda - 1)\widetilde{s} \\
\Rightarrow \; & p(\widehat{w}_t) = \frac{\widetilde{q}(1-q)}{\widetilde{q}(1-q)+q(1-\widetilde{q})}\alpha_3 + \frac{q(1-\widetilde{q})}{\widetilde{q}(1-q)+q(1-\widetilde{q})}\alpha_5; \\
& \text{if no release} \\
\Rightarrow \; & \text{by Bayes Rule, } \widehat{w}_t = -(1+\lambda)\widetilde{s} \\
\Rightarrow \; & p(\widehat{w}_t) = \frac{\widetilde{q}(1-q)}{\widetilde{q}(1-q)+q(1-\widetilde{q})}\alpha_1 + \frac{q(1-\widetilde{q})}{\widetilde{q}(1-q)+q(1-\widetilde{q})}\alpha_7 \\
= \; & \frac{q(1-\widetilde{q})}{\widetilde{q}(1-q)+q(1-\widetilde{q})} < \frac{\widetilde{q}(1-q)}{\widetilde{q}(1-q)+q(1-\widetilde{q})}\alpha_3 + \frac{q(1-\widetilde{q})}{\widetilde{q}(1-q)+q(1-\widetilde{q})}\alpha_5 \text{ as } \widetilde{q} < q \\
\Rightarrow \; & \text{not a profitable deviation;} \\
& \text{if not to release when } w_t = -\widetilde{s} \text{ but releasing later at } t' > t \\
\Rightarrow \; & \text{by Bayes Rule, } \widehat{w}_k = -(1+\lambda)\widetilde{s} \text{ for } k = t, t+1, \ldots, t'-1 \text{ and} \\
& \widehat{w}_{t'} = \widehat{w}_{t'-1} + 2\lambda\widetilde{s} = (\lambda-1)\widetilde{s} \\
\Rightarrow \; & p(\widehat{w}_{t'}) = \frac{\widetilde{q}(1-q)}{\widetilde{q}(1-q)+q(1-\widetilde{q})}\alpha_3 + \frac{q(1-\widetilde{q})}{\widetilde{q}(1-q)+q(1-\widetilde{q})}\alpha_5 \Rightarrow \text{makes no difference;}
\end{aligned}
$$

$$
\begin{aligned}
& \text{if instead releasing earlier at } t'' < t \text{ when } w_{t''} = 0 \\
\Rightarrow \; & \widehat{w}_{t''} = \lambda\widetilde{s} \text{ and } p(\widehat{w}_{t''}) = \widetilde{q}\alpha_5 + (1-\widetilde{q})\alpha_3; \\
& \text{at } t'' \text{ the original strategy is not to release until } w_t = -\widetilde{s} \text{ later} \\
\Rightarrow \; & \text{agents can perfectly infer } \widetilde{s} \text{ at } t''+1 \text{ since } w_{t''+1} = \pm 1 \\
\Rightarrow \; & \widehat{w}_{t''+1} = w_{t''+1} + \lambda\widetilde{s} \text{ and} \\
& p(\widehat{w}_{t''+1}) = \widetilde{q}[q\alpha_7 + (1-q)\alpha_3] + (1-\widetilde{q})[q\alpha_5 + (1-q)\alpha_1] = \widetilde{q}\alpha_5 + (1-\widetilde{q})\alpha_3 \\
\Rightarrow \; & \text{makes no difference;} \\
& \text{if instead releasing earlier at } t''' < t \text{ when } w_{t'''} = \widetilde{s} \\
\Rightarrow \; & \widehat{w}_{t'''} = (1+\lambda)\widetilde{s} \text{ and } p(\widehat{w}_{t'''}) = \frac{\widetilde{q}q}{\widetilde{q}q + (1-q)(1-\widetilde{q})}; \\
& \text{at } t''' \text{ the original strategy is not to release but agents can infer } \widetilde{s} = w_{t'''} \\
\Rightarrow \; & \text{makes no difference.}
\end{aligned}
$$

Therefore $g(\widetilde{s}, w_t) = 1_{\{w_t = -\widetilde{s}\}}$ is indeed an equilibrium strategy of the social planner. ■

Proof of Claim 2.3.1. We start by calculating the discounted sum of people's expected payoffs without any public information, $\overline{V}$. Without loss of generality, we assume $\theta = 1$ and use $EU_t \equiv Eu(a_t^*; \theta = 1)$ to simplify the notation. Then $\overline{V} = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} EU_t$.

Recursive calculation similar to the proof of Proposition 2.2.1 yields

$$\begin{aligned} EU_1 &= EU_2 = q, \\ EU_3 &= EU_4 = q^2 + 2q(1-q)EU_1, \\ &\dots \\ EU_{2k+1} &= EU_{2k+2} = q^2 + 2q(1-q)EU_{2k-1}. \end{aligned}$$

Hence

$$EU_{2k+1} = [2q(1-q)]^k [q - \frac{q^2}{(1-q)^2 + q^2}] + \frac{q^2}{(1-q)^2 + q^2};$$
$$\overline{V} = (1-\delta)\sum_{k=0}^{\infty} (\delta^{2k} + \delta^{2k+1}) EU_{2k+1} = \frac{\delta^2 q^2 + (1-\delta^2)q}{1 - \delta^2 2q(1-q)} > q.$$

If the public signal is released at period 1,

$$\begin{aligned} EU_1 &= q, \\ EU_2 &= EU_3 = q^2 + 2q(1-q)EU_1, \\ &\dots \\ EU_{2k} &= EU_{2k+1} = q^2 + 2q(1-q)EU_{2k-1}. \end{aligned}$$

Compared to the case without public information, we have

$$V(1) = \frac{\overline{V} - (1-\delta)q}{\delta} > \overline{V} \text{ as } \overline{V} > q.$$

Clearly there is no difference between releasing at period 1 and at period 2, so $V(2) = V(1) > \overline{V}$.

If the public signal is released at period 3,

$EU_1 = EU_2 = q$;

with probability q^3, a correct herd starts after the release;

with probability $(1-q)^3$, a wrong herd starts after the release;

with probability $3q(1-q)$, it is as if only the public signal is released for agents $t \geq 3$.

Hence

$$\begin{aligned} V(3) &= (1-\delta)(q+\delta q) + (1-\delta)q^3\frac{\delta^2}{1-\delta} + 3q(1-q)\delta^2 V(1) \\ &= (1-\delta^2)q + q^3\delta^2 + 3q(1-q)\delta^2 V(1). \end{aligned}$$

If the public signal is released at period 4,

$EU_1 = EU_2 = q$; $EU_3 = q^2 + 2q(1-q)EU_1$;

with prob. q^3, a correct herd starts after the release;

with prob. $(1-q)^3$, a wrong herd starts after the release;

with prob. $q(1-q)$, it is as if only the public signal is released for agents $t \geq 4$.

with prob. $2q(1-q)$, the public signal is as if released after one action for agents $t \geq 4$.

Hence

$$\begin{aligned} V(4) &= (1-\delta)(q+\delta q+\delta^2 q^2) + (1-\delta)q^3\frac{\delta^3}{1-\delta} + q(1-q)\delta^3 V(1) + 2q(1-q)\delta^2 V(2) \\ &= (1-\delta^2)q + (1-\delta)\delta^2 q^2 + q^3\delta^3 + 2q(1-q)(\delta^2 + \frac{\delta^3}{2})V(1). \end{aligned}$$

Therefore,

$$\begin{aligned} V(3) - V(4) &= q^3\delta^2(1-\delta) - (1-\delta)\delta^2 q^2 + q(1-q)\delta^2(1-\delta)V(1) \\ &= q(1-q)\delta^2(1-\delta)[V(1) - q] > 0 \text{ as } V(1) > \overline{V} > q. \end{aligned}$$

■

Proof of Claim 2.3.2**.** Note that there would still be a herd eventually due to *bounded* private beliefs, so the ex-ante average payoff of the population is again the probability of a correct herd eventually.

Let us first calculate $\overline{G}$, the probability of a correct herd eventually without any public information. Using similar recursive arguments as in the proof of Proposition 2.2.1 but with more tedious algebra, we have

$$\overline{G} = 6q(\frac{1-q}{2})^2\overline{G} + q^2 + 4q\frac{1-q}{2}q\frac{1+q}{2} + 4q^2(\frac{1-q}{2})^2(A+B),$$

where A and B are given by

$$\begin{cases} A = q\frac{1+q}{2} + q\frac{1-q}{2}A + (1-q)qB \\ B = (\frac{1+q}{2})^2 + \frac{1+q}{2}\frac{1-q}{2}A + \frac{1-q}{2}qB \end{cases}.$$

A and B are in fact the probabilities of a correct herd conditional on the event that a trap has started corresponding to two different tie-breaking situations.

It is easy to see that $G(1) = G(2) = \overline{G}$, because without the public signal the first two agents always follow their own private signals and releasing $\widetilde{s}$ is just "adding" an agent who always follow her signal, which does not affect social welfare from an ex-ante perspective. By exploring all possible situations of the first two actions, the value of $G(3)$ can be calculated as follows:

$$\begin{aligned} G(3) &= q^2[q + 2\frac{1-q}{2}(\frac{1+q}{2} + \frac{1-q}{2}A)] + 2(\frac{1-q}{2})^2q^2B \\ &\quad +2(\frac{1-q}{2})^2q\overline{G} + 4q\frac{1-q}{2}[q(\frac{1+q}{2} + \frac{1-q}{2}A) + \frac{1-q}{2}\overline{G} + \frac{1-q}{2}qB] \\ &\Rightarrow G(3) - \overline{G} = q^2(1-q)[\frac{1+q}{2} + \frac{1-q}{2}(A+B) - 1] > 0 \text{ as } A+B > 1. \end{aligned}$$

For $G(4)$ and $G(5)$, if the first three actions cancel each other then it is as if the public signal were released three periods earlier; otherwise either a trap or a herd starts and calculations similar to those above can be applied. Indeed we have

$$\begin{aligned} G(4) &= 6q(\frac{1-q}{2})^2G(1)+2q^2(\frac{1-q}{2})^2B+q^2[q+2\frac{1-q}{2}(\frac{1+q}{2}+\frac{1-q}{2}A)] \\ &+4q^3\frac{1-q}{2}+4q^2(\frac{1-q}{2})^2A \\ &+4q^2(\frac{1-q}{2})^2[q+2\frac{1-q}{2}(\frac{1-q}{2}\overline{G}+\frac{1-q}{2}qB+q\frac{1+q}{2}+q\frac{1-q}{2}A)] \\ &+4(\frac{1-q}{2})^3q[\frac{1-q}{2}q\overline{G}+q(\frac{1-q}{2}\overline{G}+\frac{1-q}{2}qB+q\frac{1+q}{2}+q\frac{1-q}{2}A)] \\ &+4(\frac{1-q}{2})^2q\{qB+\frac{1-q}{2}[\frac{1-q}{2}q\overline{G}+q(\frac{1-q}{2}\overline{G}+\frac{1-q}{2}qB+q\frac{1+q}{2}+q\frac{1-q}{2}A)]\}; \end{aligned}$$

$$\begin{aligned} G(5) &= 6q(\frac{1-q}{2})^2G(2)+2q^2(\frac{1-q}{2})^2B+q^2[q+2\frac{1-q}{2}(\frac{1+q}{2}+\frac{1-q}{2}A)] \\ &+4(\frac{1-q}{2})^3\frac{1-q}{2}q^2\overline{G} \\ &+4(\frac{1-q}{2})q^2\frac{1-q}{2}(q\frac{1+q}{2}+q\frac{1-q}{2}A+q\frac{1-q}{2}B+\frac{1-q}{2}C) \\ &+4(\frac{1-q}{2})q^2\frac{1+q}{2}[q+\frac{1-q}{2}+\frac{1-q}{2}(\frac{1-q}{2}A+\frac{1-q}{2})] \\ &+4(\frac{1-q}{2})^2q[q(q\frac{1+q}{2}+q\frac{1-q}{2}A+\frac{1-q}{2}qB+\frac{1-q}{2}C)+(1-q)q^2B], \end{aligned}$$

where $C \equiv \frac{1-q}{2}\overline{G}+q\frac{1+q}{2}+q\frac{1-q}{2}A+\frac{1-q}{2}qB$.

It can be verified that

$$sgn(G(4)-G(5)) = sgn(q^2(A+B-1)+3\frac{1-q}{2}(1-\overline{G})-\frac{1+q}{2}).$$

When $q=1$ or $\frac{1}{3}$,

$$\begin{aligned} & A=B=\overline{G}=1 \text{ or } \frac{1}{3} \\ \Rightarrow\ & q^2(A+B-1)+3\frac{1-q}{2}(1-\overline{G})-\frac{1+q}{2}=0; \end{aligned}$$

Meanwhile, $q^2(A+B-1)+3\frac{1-q}{2}(1-\overline{G})-\frac{1+q}{2}$ is in fact convex in q on $(\frac{1}{3},1)$, therefore

$$q^2(A+B-1)+3\frac{1-q}{2}(1-\overline{G})-\frac{1+q}{2} > 0 \text{ and } G(4) > G(5).$$

■

3 Learning and Price Dynamics in Durable Goods Markets

A durable good provides a private use value to its user, and it is eventually resold in a secondary market. This paper analyzes what determines different learning and price dynamics in durable goods markets. Our model includes three main features: (i) buyers have heterogeneous private use values and a common expected resale horizon; (ii) an unobservable and time-varying aggregate state determines the distribution of use values in the population; and (iii) trade takes place in markets with a limited number of buyers. Informational frictions slow down learning and affect price movements asymmetrically in high and low aggregate states. We disentangle two sources of price variability. Idiosyncratic volatility is prevalent in markets with very heterogenous use values, a long resale horizon and a small number of buyers. Aggregate volatility mirrors the sensitivity of prices to new price information, and it weights more when the resale motive dominates, *i.e.*, for goods with short resale horizons, significant persistence of the aggregate state, and similar use values.

3.1 Introduction

Since the initial contribution of Hayek (1945), a vast literature in economic theory has been studying how the price system aggregates dispersed private information. No social planner has access to all available information, and in a market-based economy prices have the fundamental of influencing decisions by consumers, firms and governments. How information is *incorporated* into asset prices is the main focus of the Rational Expectation literature, and one of the most debated topics in finance.

Although there is a vast asset pricing literature on financial securities, less attention has been devoted to price patterns in durable goods markets. Notable examples are real-estate, machineries, automotive, but also artwork, collectibles and musical instruments. These goods provide a private use value to users, but they are often resold on the market after some time. Durable goods represent a sizeable portion of household and corporate balance sheets, and as such they play a central role in the economy as consumption goods, production inputs or pledgeable collateral. Many papers focus on a specific market—especially real-estate and vehicles—and try to match a few empirical facts, either with a rather specific model, or with a slight adaptation of a workhorse asset pricing model. In the former case, the results cannot be applied *sic et simpliciter* to other markets that share few similarities; in the second case, models overlook some specific—but potentially relevant—market features.[67]

In this paper we broadly focus on durable goods—a sufficiently large class of assets—and we study how a few common characteristics affect learning and information aggregation. Our model does not pretend to match precise price patterns for a specific market, but it rather aims to highlight a few economic mechanisms common to all durable goods markets.[68]

We develop a dynamic trading model with time-varying and unobservable aggregate demand conditions. Our framework explicitly considers two peculiar characteristics of durable goods. First, they trade in decentralized markets where sellers enter into private

[67] An example of this dihcotomic approach is the real estate literature. Some authors use Lucas-type models and derive estimates for risk and liquidity premia, other papers set up search and matching models including a rental sector, geographic dispersion, and private use values.

[68] We do not deal with any specific price puzzle, and we actually exclude *a priori* the existence of risk premia by assuming agents' risk neutrality. Our main focus is on information.

negotiations with a limited number of potential buyers. Second, they provide utility as consumption goods until re-sold to a different user at a future point in time.[69] There exist great variation within each characteristic. On the one hand, trade decentralization admits a large variety of trade protocols. On the other hand, the consumption vs. resale trade-off depends on several intrinsic characteristics of the market.

Learning patterns depend on prices if the latter provide useful information on the underlying aggregate demand. Heterogeneity in trading protocols leads to different ways in which agents update their beliefs. These informational frictions may have different origins: the absence of an organized trading platform, legal restrictions on information disclosure, or bidders' incentives to manipulate prices. We abstract from any single source of friction and focus directly on the relationship between disclosed information and learning dynamics. We present two main results. First, trading games revealing *coarser* information sets lead to a *slower* learning process. Second, different trading protocols may affect beliefs *asymmetrically* between high and low aggregate demand states. In particular, when only winning bids are disclosed, beliefs tend to adjust more rapidly when the aggregate state is low.

If the trading protocol determines which information is revealed to agents, other intrinsic characteristics of the durable good influence its price *sensitivity* to new information. We consider three main dimensions: the expected resale horizon, the persistence of aggregate demand states, and the degree of heterogeneity in private use values. To explicitly solve the model, we assume sellers trade via second-price auctions. Thanks to an analytic solution for the bidding strategy, we obtain several comparative statics results. First, prices respond more to new information when buyers have more similar private use values. Second, a *longer* expected resale horizon increases the relative importance of private use values *vis-á-vis* future resale prices. Similarly, price sensitivity is larger when aggregate states are more persistent. Lastly, price volatility can be decomposed into two factors: idiosyncratic and aggregate. The former depends on the heterogeneity in buyers' use values, and it is driven by the consumption motive. The latter captures price sensitiveness to current information, and it depends on the interest in forecasting

[69] Other products may share the same two features. We explicitly refer to durable goods just to focus our attention on a relevant set of markets which possess these broad characteristics.

future prices.

Despite theoretical in nature, we believe our paper points out a few general ideas with a broad range of potential applications. For example, suppose a credit officer has to decide on the loan terms applied to two otherwise identical customers with different collateral goods: one has an classic car, and the other one a modern corporate car. Which car is the less risky collateral? To answer this question, it might be a good idea to understand who participates in these markets, and for what purpose. Classic cars are mostly bought for their subjective use value, and usually resold after a long time. On the contrary, buyers of corporate cars have similar use values and a fast car turnover, and they significantly care about the future resale price. Our model provides a framework to explain how these different characteristics affect price volatility.

Overview of the model and results. We briefly sketch our model setup to discuss our results in more detail. Trade takes place through a sequence of trading rounds with N bidders. Aggregate market conditions in period t depend on the distribution of private values from which individual bidders are sampled. In particular, their per-period use value in period t come from one of two possible distribution functions F_{θ_t}, $\theta_t \in \{H, L\}$. The state of the world θ_t is never publicly revealed, and it varies overtime according to a Markov process with state persistence ρ_j, $j = H, L$. Unless θ_t realizations are independent overtime, the observable public history provides information on the likelihood of future states of the world. Private use values have a double role: (i) they measure individual benefits from enjoying the good; and (ii) they provide information on the underlying state of the world. A winning bidder resells his good at a future random time: he faces an $\alpha \leq 1$ probability to sell his good in the next period. Higher values of α denote shorter resale horizons.[70] An owner enjoys his individual use value until resale. For simplicity, losing bidders and sellers go out of the market with no future possibility of re-entering.

The aggregate state θ_t may be considered as a reduced form to capture all those elements such as fashion, business and credit cycles that affect, at a given point in time, the willingness to purchase the good among agents in the population. It is often difficult to directly observe this state and we assume buyers only observe previous transaction

[70] α could be interpreted as the likelihood of being hit by a liquidity shock that forces the owner to sell the object.

prices. For example, a real-estate buyer may collect information on past prices in a local market but he may not have (or be able to process) information on unsuccessful bids, or on the real-estate market at large.

Our setup captures a few characteristics of market demand that widely vary across durable goods. The parameter α is a reduced form to capture the expected resale horizon for the good. The state persistence parameter ρ measures how likely an aggregate state will persist in future periods; lower values of ρ_H and ρ_L denote a more volatile aggregate environment. The distributions F_θ describe a more or less dispersed distribution of private use values among agents in the economy. Finally, the number of bidders N provides a measure of market competition, but also, to a certain extent, market liquidity. Thanks to an explicit characterization of the bidding function it is possible to derive analytically some general comparative statics results, and it would be straightforward to simulate other statistical properties for specific functional forms F_θ.

In Section 3.2 we discuss how differences in the information revealed through prices affect learning dynamics. The more information is disclosed by a trade protocol, the faster beliefs converge to the true state. In this respect, durable goods markets may exhibit a more sluggish price adjustment process relative to a centralized market.[71] The second result is less intuitive. In general, the speed of learning differ between high and low aggregate states. This asymmetry depends on the information revealed by the trade protocol. For example, a first-price auction reveals the highest valuation among the N bidders. In this case, learning is faster in the low state because low prices are more informative in revealing the underlying aggregate demand state.[72]

In Section 3.3 we assume—for reasons of analytical tractability—that the object is sold in a sequence of second-price auctions.[73] Prices are more sensitive to new information in markets in which: (i) the resale horizon is shorter ($\alpha \uparrow$); (ii) the expected demand between high and low aggregate states is larger; (iii) the current state of market demand is more likely to last longer ($\rho_j \uparrow$). Under these circumstances, sellers weight more the

[71] For example, in a modern stock exchange dealers' price quotes and traders' limit orders can be freely observed in real time by all market participants. Without strategic price manipulation markets disclose all private information.

[72] An analogous, but opposite, logic would hold if the lowest valuation among N bidders were revealed.

[73] In a second-price auction, players do not engage in strategic price manipulation: the optimal bidding strategy truthfully reveals private signals. In other auction formats—such as the first-price auction—manipulative incentives may arise, and the equilibrium analysis becomes analytically intractable.

informational content of recent prices, which is more effective in predicting future resale values. Price variability can be decomposed into two different components. The first one reflects the heterogeneity in private use values, and it has a purely idiosyncratic nature. The second type of uncertainty is over future market conditions. A decrease in the resale horizon ($\alpha \uparrow$) increases the aggregate variability component, decreasing the idiosyncratic one; thus, the overall effect is ambiguous. An increase in state persistence ($\rho_j \uparrow$) does not affect the idiosyncratic variance, and it increases the variability due to the future resale component.

Related literature. This work is closely related to the literature on learning in asset pricing models; see Timmermann (1993) for an early reference. This strand of literature argues that learning may explain some classic asset pricing puzzles (equity premium, risk-free rate and excess-volatility). Weitzman (2007) considers a Bayesian framework with risk-averse preferences, while Ju and Miao (2012) introduce ambiguity aversion as an additional explanatory factor. Compared to this literature, we focus more on the different microeconomic determinants of price dynamics, and we abstract from any discussion on risk premia by assuming risk neutral preferences.

This work is also related to the literature on auctions with resale. A small number of papers study this topic in a two-period setting. Gupta and Lebrun (1999) consider a setup in which private values are publicly revealed in the second period. Haile (2001, 2003) study the revenue performance of different auction formats in a *symmetric* environment. In his model, bidders' initial types come from the same initial distribution but they are not publicly announced in the second period. Within a similar symmetric environment, Zheng (2002) and Lebrun (2012) provide conditions to obtain the optimal auction outcomes first derived in Myerson (1981).[74] Differently from this literature, we do not assume that the same set of bidders re-trades in future periods. The latter case is important for industries in which market players rarely change overtime, and manipulative incentives may arise when the same goods are re-traded among the same set of bidders.[75] However, in many durable goods markets this type of strategic interaction seems less relevant. For example,

[74] Lebrun (2012) extend the results in Zheng (2002) for the symmetric environment, and it provides results for a specific class of asymmetric environments. Garratt and Tröger (2006), Halafir and Krishna (2008) and Virag (2013) provide additional results for asymmetric environments.

[75] For example, government concessions in telecommunications, oil, electricity.

in the real-estate market, buyers and sellers often do not have any previous information on the identity of their counterpart.[76]

The next section discusses some general results on public/social learning. Section 3.3 presents and solves the dynamic auction model, and provides some comparative statics analysis. Section 3.4 concludes. All proofs are in the Appendix.

3.2 Information Revelation and Learning

3.2.1 Model Setup

We consider a sequential market for a durable object. Time is discrete $t \in \{0, 1, 2, \ldots\}$. In each period t, there is an underlying state $\theta_t \in \{H, L\}$. The stochastic process $\{\theta_t\}_{t=0}^{\infty}$ is a homogenous Markov process with transition matrix:

$$P = \begin{bmatrix} \rho_H & 1 - \rho_H \\ 1 - \rho_L & \rho_L \end{bmatrix}$$

with $0 \leq \rho_H, \rho_L \leq 1$. The prior on θ_t is denoted by $\boldsymbol{\pi}_t = (\pi_t, 1 - \pi_t)$ with $\pi_t \equiv \mathbb{P}(\theta_t = H)$.

There is a population of infinitely many agents interested in the object. When an object is offered on sale, $N \geq 2$ agents are randomly drawn from the population to enter the market. Each buyer attaches a private use value to the object. The private values generated in each period t, $\{v_{it}\}_{i=1}^{N}$, are i.i.d. distributed according to a cumulative density function (cdf) F_{θ_t} across the N agents. The realizations of $\{\theta_t\}_{t=0}^{\infty}$ are not known to the agents, but both P and $\boldsymbol{\pi}_0$ are common knowledge.

Both F_H and F_L are continuously differentiable on the common support $[0, 1]$. Moreover, the corresponding probability density function (pdf) f_H and f_L are strictly positive everywhere on $[0, 1]$, and satisfy the monotone likelihood ratio (MLR) property: $\frac{f_H(\cdot)}{f_L(\cdot)}$ is strictly monotone on $[0, 1]$.

[76]Similar to our model, in some Australian cities a significant portion of home sales take place via auctions: `www.bloomberg.com/news/2013-04-23/australia-turns-to-auctions-as-housing-revives-mortgages.html`

3.2.2 Public Beliefs Dynamics

The information revealed in a trading round depends on the trade protocol. For example, there is a substantial difference between auctions and centralized exchanges, but there are also significant differences among auction formats. In this subsection, we abstract from a specific trade protocol, and we directly consider the information revealed after a trading round.[77]

Consider a vector $X_k = \{v_{1,k}, ..., v_{N,k}\}$ of private signals dispersed among N traders. We assume the trade protocol leads to publicly observe a statistic $T(X_k)$. $T(\cdot)$ is assumed not to depend on the previous history of the game, hence it is invariant in all periods k. In other words, the statistic T captures which information in the vector of private valuations X_k possessed by the N bidders in period k is publicly revealed after trade.[78] It is an equilibrium object because it depends on the trade protocol and players' strategies. It implicitly incorporates both informational constraints, due to the market organization, and informational frictions, due to players' strategic behaviour.

We explore two different issues related to learning. First, we provide a sufficient condition that ranks which statistic leads to a faster public belief convergence towards the true state. Second, we analyze whether a particular T leads to a more rapid price adjustment in one of the two states of the world. For these purposes, we restrict attention to the full persistence case $\rho_H = \rho_L = 1$.

Consider a probability space $\langle \mathbb{R}^N, \mathcal{B}, \mu \rangle$ endowed with the standard Borel σ-algebra and Lebesgue probability measure. Let a measurable function $T_i : \mathbb{R}^N \to \mathbb{R}^M$, $M \leq N$ be an observable statistic of the underlying $X_k = \{v_{1,k}, ...v_{N,k}\}$ and let $\sigma(T_i)$ be the σ-algebra generated by T_i. We denote with S_X the support of X.

Definition 3.1 *T_j is coarser than T_i if $\sigma(T_j) \subset \sigma(T_i)$ and $\exists A \in \sigma(T_i)$ s.t. $A \notin \sigma(T_j)$ and $\mu(A) > 0$.*

For a statistic T let S_T denote its support and $\mathcal{C}_T(A) \equiv \bigcup_{y \in A} \{X \in \mathbb{R}^N : T(X) = y\}$ the set of counter images of $A \subseteq S_T$.

[77] In Section 3.3, we solve a specific model where agents participate in second-price auctions. In this section, we adopt a more general approach to point out a few general properties of learning dynamics.

[78] For example, in Section 3.3 the second-highest price is publicly revealed and in equilibrium buyers infer the corresponding private value, hence $T(X_k) = v_k^{(2,N)}$.

We use this general notation to express public belief dynamics under different trade protocols. Let $f_\theta^T(y)$ be the probability density function of statistic T under state $\theta \in \{H, L\}$. Formally,

$$f_\theta^T(y) = \int_{\mathcal{C}_T(y)} f_\theta^X(x)\,\mathrm{d}\mu(x)$$

It is easier to describe the evolution of public beliefs with the log-likelihood ratio:

$$l_{k+1}(l_k, y_k) = \ln\frac{\pi_{\kappa+1}}{1-\pi_{\kappa+1}} = \ln\frac{\pi_\kappa}{1-\pi_\kappa}\frac{f_H^T(y_k)}{f_L^T(y_k)} = l_k + \ln\frac{f_H^T(y_k)}{f_L^T(y_k)} \tag{1}$$

where $y_k = T(x_k)$ is the value of statistic T when the vector of private use values for the N bidders in period k is $x_k \in \mathbb{R}^N$. To stress that the log-likelihood l_k depends on T, we add a superscript T. Let $\Delta l_{k+1}^T(y_k) = l_{k+1}^T(l_k^T, y_k) - l_k^T = \ln\frac{f_H^T(y_k)}{f_L^T(y_k)}$ denote the change in the log-likelihood ratio from period k to $k+1$ under statistic T. Assume there exists $M > 0$ such that $|\Delta l^T(x)| < M$ for every $x \in S_X$.

For $q \geq 1$ equation (1) generalizes into:

$$l_{k+q}^T = l_k^T + \sum_{m=1}^{q} \Delta l_{k+m}^T(y_{k+m-1})$$

Taking the expected value:

$$\mathbb{E}_{k,\theta}^X[l_{k+q}^T] = l_k^T + \sum_{m=1}^{q} \mathbb{E}_{k,\theta}^X\left[\Delta l_{k+m}^T(T(x_{k+m-1}))\right] = l_k^T + q\mathbb{E}_\theta^X\left[\Delta l^T(T(x))\right]$$

The last equation exploits the fact that—conditional on θ—samples are i.i.d. in all periods. In the remainder of the paper, we simply use $\mathbb{E}_\theta\left[\Delta l^T\right]$ rather than $\mathbb{E}_\theta^X\left[\Delta l^T(T(x))\right]$.

Beliefs converge to the true state as $\mathbb{E}_H\left[\Delta l^T\right] > 0 > \mathbb{E}_L\left[\Delta l^T\right]$.[79] Moreover, for two different statistics T_1 and T_2, public beliefs are expected to converge more rapidly to the true state θ under statistic T_1 if

$$\left|\mathbb{E}_\theta\left[\Delta l^{T_1}\right]\right| > \left|\mathbb{E}_\theta\left[\Delta l^{T_2}\right]\right| \tag{2}$$

[79] These inequalities follow from a simple application of Gibbs' inequality. They are strict inequalities because of the MLR assumption earlier.

The following claim provides an intuitive but still insightful result. If two statistics can be ranked according to Definition 3.1, it is possible to conclude that convergence is slower for the coarser one.

Claim 3.2.1 *If T_2 is coarser than T_1 then equation (2) holds.*

Although the result in Claim 3.2.1 is not surprising, it highlights an important property of markets in which information is only partially revealed. More severe informational frictions lead to more sluggish trade dynamics, and, possibly, a slower price adjustment.

A less intuitive result is that trade protocols may create differences—between high and low states—in the speed of convergence of public beliefs. In turn, more rapid learning is likely to be positively correlated with a more rapid price adjustment.

Claim 3.2.2 *Consider a statistic $T(\cdot)$, and let $\pi_0 = 1/2$. Define:*

$$\tau_H \equiv \inf\{k \geq 0 : \pi_k \geq 1 - \epsilon\} \qquad \tau_L \equiv \inf\{k \geq 0 : \pi_k \leq \epsilon\}$$

Then:

$$\lim_{\varepsilon \to 0} \tfrac{\mathbb{E}_H[\tau_H]}{\mathbb{E}_L[\tau_L]} \geq \left| \tfrac{\mathbb{E}_L[\Delta l^T]}{\mathbb{E}_H[\Delta l^T]} \right| > 1 \quad \textit{if} \;\; \mathbb{E}_H[\Delta l^T] + \mathbb{E}_L[\Delta l^T] < 0;$$
$$\lim_{\varepsilon \to 0} \tfrac{\mathbb{E}_L[\tau_L]}{\mathbb{E}_H[\tau_H]} \geq \left| \tfrac{\mathbb{E}_H[\Delta l^T]}{\mathbb{E}_L[\Delta l^T]} \right| > 1 \quad \textit{if} \;\; \mathbb{E}_H[\Delta l^T] + \mathbb{E}_L[\Delta l^T] > 0.$$

Claim 3.2.2 points out a learning story based on the nature of the information revealed in previous trading rounds. Compared to the 'rockets and feathers' story, our mechanism is likely to run the opposite way. If a trade protocol only reveals winning bids, a more rapid adjustment should be observed downward. We discuss the intuition in the context of an example.

Example of Claim 3.2.2. Consider pdfs $f_H(x) = 2x$ and $f_L(x) = 2(1-x)$. Suppose the trade protocol reveals, in equilibrium, the j-th order statistic out of N bidders. The next table summarizes the numerical values of the critical expression $\mathbb{E}_H[\Delta l^T] + \mathbb{E}_L[\Delta l^T]$ in Claim 3.2.2:[80]

[80] We compute it numerically as explicit integrals cannot be obtained.

$N \mid j$	1	2	3	4	5	6	7	8	9	10
1	0									
2	−0.09	0.09								
3	−0.22	0	0.22							
4	−0.37	−0.13	0.13	0.37						
5	−0.52	−0.28	0	0.28	0.52					
6	−0.68	−0.44	−0.14	0.14	0.44	0.68				
7	−0.83	−0.60	−0.30	0	0.30	0.60	0.83			
8	−0.99	−0.77	−0.46	−0.15	0.15	0.46	0.77	0.99		
9	−1.14	−0.94	−0.63	−0.31	0	0.31	0.63	0.94	1.14	
10	−1.29	−1.12	−0.80	−0.48	−0.16	0.16	0.48	0.80	1.12	1.29

The value is negative (positive) when beliefs move more rapidly toward state L (H), as the expecting hitting time under state L (H) is shorter. The table shows that:

1. If $j < \frac{N+1}{2}$ convergence is faster towards state L.

2. If $j > \frac{N+1}{2}$ convergence is faster towards state H.

3. If $j = \frac{N+1}{2}$, (N odd), there is no difference.

The table captures an intuitive result. If the trade protocol reveals a sequence of higher order statistic ($j < \frac{N+1}{2}$), low-value observations are more informative than high-value ones, and learning in more rapid in state L.[81] Fixing j and increasing the sample size N, there is more and more asymmetry toward state L.[82] Increasing the sample size N, the j-th order statistic is relatively 'higher', low-value observations become more informative, and there is a greater asymmetry towards state L. For example, if only the winning bid is revealed, low demand states are learnt more rapidly in a large market.

In conclusion, prices are not equally informative on both aggregate states. This phenomenon depends on the original distribution functions F_θ, but also on the trading protocol, and the number of market participants.

[81] On the contrary, if a sequence of lower order statistic ($j > \frac{N+1}{2}$) is revealed, high-value observations are more informative than low-value ones, and learning is more rapid in state H.

[82] Alternatively, there is less asymmetry towards H as values get smaller downwards in each column.

3.3 Dynamic Auction Model

3.3.1 Trading Protocol

Consider the model setup in Subsection 3.2.1. Now assume agents trade in second-price auctions according to the following protocol.

1. Consider a stochastic sequence $\{t_k\}_{k=0}^{\infty}$, with t_0 normalized to 0. At each t_k, N new agents enter the market, and participate in a sealed-bid second-price auction. The winner of the auction at t_k is the seller in the next available auction at t_{k+1}.

2. The waiting time between t_{k+1} and t_k is a random variable, which is i.i.d distributed across k according to a geometric distribution with parameter $\alpha \in (0, 1]$. That is,

$$\mathbb{P}(\triangle_k \equiv t_{k+1} - t_k = x) = \alpha(1-\alpha)^{x-1}, \forall x \in \mathbb{N}^+, \forall k \in \mathbb{N}.$$

   Due to the i.i.d. feature of the waiting time, we simply call the auction at t_k as "auction k". We also label each bidder in auction k by ik, $i \in \{1, 2, \ldots, N\}$. Denote his private value and bid as $w_{ik} \equiv v_{i,t_k}$ and b_{ik}, respectively.

3. The winner of auction k resells the object at the next available auction $k+1$. The revenues from resale are discounted at rate δ per period. Meanwhile, he enjoys his private value of the object, w_{ik}, in every period before auction $k+1$, and he discounts his utility at rate δ per period.

4. The trading price $p_k \equiv b_k^{(2)}$, the second highest bid in auction k, is publicly observed by the whole population before the next auction starts. There is no information generated between two adjacent auctions, other than the realization of the waiting time in between. Hence, the information set for each bidder ik is $\mathcal{I}_{ik} = \{w_{ik}, \{p_\tau\}_{\tau<k}, \{\Delta_\tau\}_{\tau<k}\}$.

3.3.2 Equilibrium characterization

Let $\mathbf{b} \equiv \{b_{ik}\}_{i \leq N, k \in \mathbb{N}}$ denote the action profile of all market entrants and every bidder ik's payoff is given by

$$u_{ik}(\mathbf{b}; w_{ik}) = 1_{\{b_{ik}=b_k^{(1)}\}} \cdot \{-b_k^{(2)} + \sum_{s=0}^{\triangle_{k-1}} \delta^s w_{ik} + \delta^{\Delta_k} b_{k+1}^{(2)}\}$$

Denote the public belief about the underlying state θ before auction k by $\boldsymbol{\pi}_k = (\pi_k, 1-\pi_k)$ with $\pi_k \equiv \mathbb{P}(\theta_{t_k} = H | \{p_\tau\}_{\tau<k})$.

We consider a Perfect Bayesian Equilibrium with *symmetric, time-invariant and monotone* strategies. In turn, every bidder ik can restrict his attention to the information set $\{v_{ik}, \boldsymbol{\pi}_k\}$.

Definition 3.2 *A pure strategy profile $\mathbf{b}^* \equiv \{b_{ik}^*(\mathcal{I}_{ik})\}_{i\leq N, k\in\mathbb{N}}$ is a perfect Bayesian equilibrium with symmetric, time-invariant and monotone strategies if*

1. $b_{ik}^*(\mathcal{I}_{ik}) = b^*(w_{ik}; \boldsymbol{\pi}_k)$, $\forall i, t$;

2. $\frac{\partial b^*(w_{ik};\boldsymbol{\pi}_k)}{\partial w_{ik}} > 0$, $\forall w_{ik} \in [0,1]$;

3. $b^*(w_{ik}; \boldsymbol{\pi}_k) = \arg\max_b \mathbb{E}[u_{ik}(b, b_{-ik}^*; w_{ik}) | w_{ik}, \boldsymbol{\pi}_k]$, $\forall i, k$;

4. $\boldsymbol{\pi}_{k+1} = \left(\frac{\pi_k g_H(b^{*-1}(p_k;\boldsymbol{\pi}_k))}{\pi_k g_H(b^{*-1}(p_k;\boldsymbol{\pi}_k)) + (1-\pi_k) g_L(b^{*-1}(p_k;\boldsymbol{\pi}_k))}, \frac{(1-\pi_k) g_L(b^{*-1}(p_k;\boldsymbol{\pi}_k))}{\pi_k g_H(b^{*-1}(p_k;\boldsymbol{\pi}_k)) + (1-\pi_k) g_L(b^{*-1}(p_k;\boldsymbol{\pi}_k))} \right) P^{\Delta_k}$

   where $g_\theta(\cdot)$ is the pdf of the 2^{nd} order statistic among N i.i.d random variables distributed according to F_θ, $\forall \theta \in \{H, L\}$.

The next proposition provides an explicit characterization of the equilibrium.

Proposition 3.3.1 *Let $\rho_H + \rho_L \geq 1$. There is a unique perfect Bayesian equilibrium with symmetric, time-invariant and monotone strategies:*

$$b^*(w_{ik}; \pi_k) = \frac{1}{1-\delta+\alpha\delta} \left\{ w_{ik} + \frac{\alpha\delta}{1-\delta} \left[c_L + \frac{1-\rho_L}{1-\delta(\rho_H+\rho_L-1)} \Delta c \right] + \gamma_{ik} \frac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)} \Delta c \right\}$$

with

1. $$\gamma_{ik} \equiv \frac{\pi_k f_H(w_{ik}) h_H(w_{ik})}{\pi_k f_H(w_{ik}) h_H(w_{ik}) + (1-\pi_k) f_L(w_{ik}) h_L(w_{ik})}$$

where $h_\theta(\cdot)$ is the pdf of the 1^{st} order statistic among $N-1$ i.i.d random variables distributed according to F_θ, $\forall\theta \in \{H, L\}$;

2. $c_\theta \equiv \int_0^1 x g_\theta(x)\,\mathrm{d}x$, *$\forall\theta \in \{H, L\}$, and $\Delta c \equiv c_H - c_L$.*

The equilibrium bidding function $b^*(w_{ik}, \pi_k)$ can be decomposed in a private value (PV) and a resale value (RV) component.

$$\begin{aligned} \text{PV} &= \tfrac{w_{ik}}{1-\delta+\alpha\delta} \\ \text{RV} &= \tfrac{\alpha}{1-\delta+\alpha\delta}\left[\tfrac{\delta}{1-\delta}\left(c_L + \tfrac{1-\rho_L}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right) + \gamma_{ik}\tfrac{\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right] \end{aligned}$$

The private value component is the expected discounted use value of the good until resale takes place. An increase in the expected resale horizon ($\alpha \downarrow$) increases the private value component but decreases the resale value component. Bidders expect to enjoy the good for a longer time, so their use value gains importance relative to the expected future resale price. The resale value component includes a constant term, and another term which depends on belief γ_{ik}. The latter depends on the public belief π_k, and on the private use value w_{ik}. The random variable w_{ik} enters in two distinct updating. First, w_{ik} is a signal on the current state of the world because it comes from the common distribution F_θ. Second, in equilibrium a winning bidder realizes that all other $N-1$ bidders had lower private use values. This last updating is analogous to the inference carried out by a winning bidder in a static common value auction. In this respect, our model may offer a dynamic micro-foundation of a static common value auction. The future resale price is at the root of the interdependence among bidders' valuations.

3.3.3 Comparative statics

In this subsection we carry out a few comparative statics exercises to highlight the main determinants of different price dynamics.

The sensitiveness of $b^*(w_{ik}, \pi_k)$ with respect to γ can be measured through a simple elasticity measure:

$$\eta_\gamma^b := \frac{\gamma}{b}\frac{\partial b}{\partial \gamma} = \frac{\gamma \frac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Delta c}{w_{ik} + \frac{\alpha\delta}{1-\delta}\left[c_L + \frac{1-\rho_L}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right] + \gamma\frac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Delta c}$$

A higher value of η_γ^b denotes a greater sensitivity of bidders' strategies to their present beliefs about the state of the world. It is easy to show that η_γ^b is increasing in α, Δc and ρ_j, $j = H, L$. In words, the bidding strategy is more sensitive to new information if: (i) the resale horizon is shorter ($\alpha \uparrow$); (ii) the expected difference between aggregates states is larger ($\Delta c \uparrow$);[83] or (iii) each state is more persistent ($\rho_j \uparrow$). The intuition for each variable is pretty straightforward. When the resale horizon is shorter, present information is more accurate to predict the state of the world at the future time of resale. Similarly, when states of the world are more persistent, current beliefs are more precise in predicting future states. As a result, prices respond more to new information (Fig. 3). Finally, a greater difference Δc increases the variability in the possible resale values between the two aggregate states, and agents adjust their bids more sharply.

Lastly, we derive a statistical measure of dispersion for *realized* prices. Our variance measure is derived assuming a deterministic resale horizon, say, q periods long, and a future state of the world θ_{k+1}.[84] Public beliefs move between any two trading periods according to the law of motion in Definition 3.2, and—for a given π_k and a fixed resale horizon q— it is immediate to get the value of π_{k+1}.[85]

Three different factors contribute to price variability:

$$\begin{aligned}\mathrm{Var}_\theta\Big(b^*\Big(w_{k+1}^{(2)}, \pi_{k+1}\Big)\Big|\pi_{k+1}\Big) = \Big(\tfrac{1}{1-\delta+\alpha\delta}\Big)^2 \Big[\mathrm{Var}_\theta\big(w^{(2)}\big) + \Big(\tfrac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Big)^2 \mathrm{Var}_\theta\Big(\gamma\left(w_{k+1}^{(2)}, \pi_{k+1}\right)\Big|\pi_{k+1}\Big) \\ + 2\tfrac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\ \mathrm{cov}_\theta\left(w_{k+1}^{(2)}, \gamma\left(w_{k+1}^{(2)}, \pi_{k+1}\right)\Big|\pi_{k+1}\right)\Big]\end{aligned} \tag{3}$$

[83] Specifically it is the difference in the expected second highest use value out of N bidders between the high and low state of the world.

[84] Notice the difference with the variance computed according to the subjective belief of a bidders in auction k. In this case, bidders do not know neither the present nor the future state.

[85] If we did not condition on a fixed resale horizon, we could have alternatively computed a measure of *expected* variance using as weights the probability to resale in a given future period.

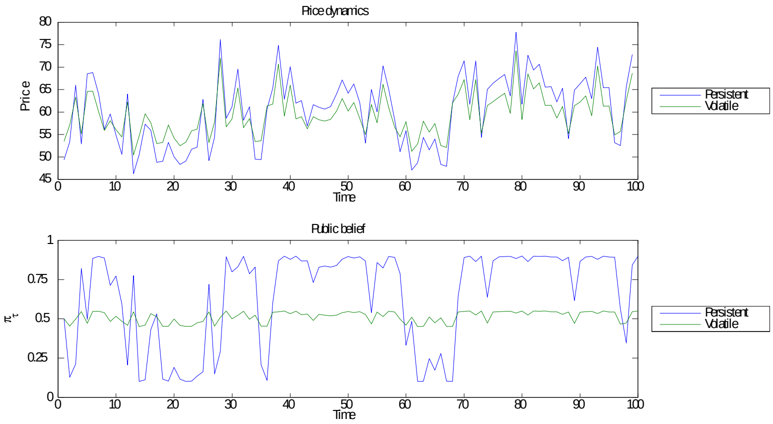

Simulation with $N = 4$, $\delta = 0.9$, $\alpha = 0.6$, $\pi_0 = 0.5$. The persistent line refers to $\rho_H = \rho_L = 0.9$ while the volatile one to $\rho_H = \rho_L = 0.55$. Private values w_{ik} are randomly sampled from lognormal distributions with mean $\mu_H = 8$, $\mu_L = 4$ and variance $\sigma = 2$.

Figure 3: Difference in state persistence.

The first term $\text{Var}_\theta\left(w^{(2)}\right)$ captures the heterogeneity in private use values. This idiosyncratic component depends on the initial distribution F_θ, and on the number of bidders N. A higher dispersion in subjective use values increases this quantity. The effect of an increase in N is not obvious, and it depends on the specific $F_\theta(\cdot)$ (see Papadatos (1995)). The second term in equation (3) reflects the uncertainty over the future beliefs held by the second highest bidder in auction $k+1$. It is a product of two quantities: a multiplicative constant, and the variance of γ_{k+1} conditional on π_{k+1}. The former is increasing in α and ρ_j; the latter is a complex quantity to analyze without additional assumptions on the functional forms for the pdfs. Lastly, the third term captures bidders' updating of γ_{ik} with the private use value $w_{i,k+1}$. The latter is used as an informative signal on the underlying aggregate state. The covariance term is always positive and it further increases price variability. The last two terms in equation (3) represent the volatility due to the uncertainty over future market conditions.

A decrease in the resale horizon ($\alpha \uparrow$) increases aggregate variability, decreasing the idiosyncratic one. The overall effect is ambiguous. A increase in the state persistence ($\rho_j \uparrow$) does not affect idiosyncratic variance, but it increases the aggregate one. Unfortunately, it is difficult to derive additional comparative statics results without assuming a specific distribution. Nonetheless, thanks to Proposition 3.3.1, it is straightforward to simulate any quantity of interest once we assume a specific form for F_θ.

3.4 Conclusion

This paper proposes a model for durable goods markets. We explicitly consider the possibility to re-sell an object, and we discuss what the potential implications are for learning and price dynamics.

We first present two results on the dynamics of public beliefs. First, the finer is the information publicly revealed in equilibrium, the faster is the convergence of public beliefs to the true state of the world. Second, trade protocols may lead public beliefs to move upward or downward at different rates. In particular, if only winning bids are disclosed, beliefs tend to adjust more rapidly when aggregate demand is low.

In the second part of the paper, we consider a dynamic auction model. Thanks to an analytic characterization of the bidding strategy, we provide some comparative statics results. A longer expected resale horizon increases the importance of private use values, and prices are less sensitive to current information. In this case, price volatility is mainly driven by the idiosyncratic tastes of users. If states of the world tend to last longer, prices respond more to current information. This is also the case when the difference in market conditions between high and low states is large.

This paper assumes an exogenous resale decision which is independent from previous price dynamics. This is clearly a strong assumption. Endogenous resale decisions play a decisive role in shaping market dynamics. For example, there is strong empirical evidence on the positive correlation between volume and prices in the real-estate market. Solving a dynamic auction model with endogenous entry is a challenging direction of extension, and we hope to address it in the future.

3.5 Appendix

Proof of Claim 3.2.1. We prove the statement only for $\theta = H$ as an analogous arguments holds for $\theta = L$.[86] Equation (2) can be easily rewritten as

$$\mathbb{E}_H^X\left[\ln \frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}\right] < 0.$$

Jensen's inequality implies

$$\mathbb{E}_H^X\left[\ln \frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}\right] \leq \ln \mathbb{E}_H^X\left[\frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}\right]$$

Notice

$$\begin{aligned}\mathbb{E}_H^X\left[\frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}\right] &= \int\limits_{S_X} \frac{f_L^{T_1}(T_1(x)) f_H^{T_2}(T_2(x))}{f_H^{T_1}(T_1(x)) f_L^{T_2}(T_2(x))} f_H^X(x)\,\mathrm{d}\mu(x) \\ &= \int\limits_{S_{T_1}} \int\limits_{\mathcal{C}_{T_1}(y)} \frac{f_L^{T_1}(T_1(x)) f_H^{T_2}(T_2(x))}{f_H^{T_1}(T_1(x)) f_L^{T_2}(T_2(x))} f_H^X(x)\,\mathrm{d}\mu(x)\mathrm{d}\mu(y)\end{aligned}$$

For every $y \in S_{T_1}$ the function $\frac{f_L^{T_1}(T_1(x)) f_H^{T_2}(T_2(x))}{f_H^{T_1}(T_1(x)) f_L^{T_2}(T_2(x))}$ is constant for every element in $\mathcal{C}_{T_1}(y)$. For $\frac{f_L^{T_1}}{f_H^{T_1}}$ this is true by definition of $\mathcal{C}_{T_1}$, while it follows from coarseness for $\frac{f_H^{T_2}}{f_L^{T_2}}$. Then,

$$\begin{aligned}&\int\limits_{S_{T_1}} \int\limits_{\mathcal{C}_{T_1}(y)} \frac{f_L^{T_1}(T_1(x)) f_H^{T_2}(T_2(x))}{f_H^{T_1}(T_1(x)) f_L^{T_2}(T_2(x))} f_H^X(x)\,\mathrm{d}\mu(x)\mathrm{d}\mu(y) \\ =& \int\limits_{S_{T_1}} \frac{f_L^{T_1}(y) f_H^{T_2}(y)}{f_H^{T_1}(y) f_L^{T_2}(y)} f_H^{T_1}(y)\,\mathrm{d}\mu(y) = \int\limits_{S_{T_1}} \frac{f_H^{T_2}(y)}{f_L^{T_2}(y)} f_L^{T_1}(y)\,\mathrm{d}\mu(y)\end{aligned}$$

[86] This proof uses coarseness in order to reduce the expression to a standard Gibbs' inequality.

As $\sigma(T_2) \subset \sigma(T_1)$ we can rewrite $f_\theta^{T_2}(z) = \int\limits_{y \in S_{T_1}: T_2(\mathcal{C}_{T_1}(y))=z} f_\theta^{T_1}(y) \mathrm{d}\mu(y)$. Therefore:

$$\begin{aligned}
\int\limits_{S_{T_1}} \frac{f_H^{T_2}(y)}{f_L^{T_2}(y)} f_L^{T_1}(y)\, \mathrm{d}\mu(y) &= \int\limits_{S_{T_2}} \frac{f_H^{T_2}(z)}{f_L^{T_2}(z)} \int\limits_{y \in S_{T_1}: T_2(\mathcal{C}_{T_1}(y))=z} f_L^{T_1}(y) \mathrm{d}\mu(y) \mathrm{d}\mu(z) \\
&= \int\limits_{S_{T_2}} \frac{f_H^{T_2}(z)}{f_L^{T_2}(z)} f_L^{T_2}(z) \mathrm{d}\mu(z) = 1
\end{aligned}$$

As a result $\ln \mathbb{E}_H^X \left[\frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}} \right] = \ln 1 = 0$.

Finally observe that Jensen's inequality holds strictly. In fact, ln is a strictly concave function and $\frac{f_L^{T_1} f_H^{T_2}}{f_H^{T_1} f_L^{T_2}}$ is not constant almost everywhere because coarseness implies the existence at least two sets A, B s.t. $A \subset B$, $A \notin \sigma(T_2)$ and $\mu(A) > 0$ where $T_2(x)$ is constant $\forall x \in B$ while $T_1(x) \neq T_1(x')$ for $x \in A$ and $x' \in B \backslash A$. ■

Proof of Claim 3.2.2. As $\pi_0 = \frac{1}{2}$ we have $l_k^T = \sum_{i=0}^{k-1} \Delta l_{i+1}^T$ where $\Delta l_{i+1}^T = l_{i+1}^T - l_i^T = \ln \frac{f_H^T(y_i)}{f_L^T(y_i)}$ $i = 0, 1, ..., k$ is a sequence of i.i.d. random variables.

Hitting times τ_H and τ_L can be equivalently stated in terms l_k:

$$\tau_H^l := \inf \left\{ k > 0 : l_k^T \geq \ln \frac{1-\epsilon}{\epsilon} \right\} \qquad \tau_L^l := \inf \left\{ k > 0 : l_k^T \leq \ln \frac{\epsilon}{1-\epsilon} \right\}$$

Applying Wald (1944)'s lemma to the sequence of i.i.d random variables Δl_i:

$$\mathbb{E}_\theta[l_{\tau_\theta^l}^T] = \mathbb{E}_\theta[\tau_\theta^l] \mathbb{E}_\theta[\Delta l^T] \qquad \forall \theta \in \{H, L\} \tag{4}$$

By Gibbs' inequality $\mathbb{E}_H[\Delta l^T] > 0$ and $\mathbb{E}_L[\Delta l^T] < 0$. If $\mathbb{E}_H[\Delta l^T] + \mathbb{E}_L[\Delta l^T] < 0$ then

$$\mathbb{E}_L[\Delta l^T] = -\left(\mathbb{E}_H[\Delta l^T] + c \right)$$

where $c \equiv -\int_{S_T} \ln \frac{f_H^T(y)}{f_L^T(y)} \left(f_H^T(y) + f_L^T(y) \right) \mathrm{d}\mu(y) > 0$. Note that c only depends on the primitives and it is independent of ε.

Substituting in equation (4):

$$\begin{aligned}\mathbb{E}_H[\tau^l_H]\mathbb{E}_H[\Delta l^T] = & \ \mathbb{E}_H[l^T_{\tau^l_H}] \\ \mathbb{E}_L[\tau^l_L](\mathbb{E}_H[\Delta l^T] + c) = & \ -\mathbb{E}_L[l^T_{\tau^l_L}]\end{aligned}$$

Hence

$$\begin{aligned}\mathbb{E}_H[\Delta l^T](\mathbb{E}_H[\tau^l_H] - \mathbb{E}_L[\tau^l_L]) = & \ \mathbb{E}_H[l^T_{\tau^l_H}] + \mathbb{E}_L[l^T_{\tau^l_L}] + \mathbb{E}_L[\tau^l_L]c \\ \implies \quad \mathbb{E}_H[\Delta l^T]\left(\tfrac{\mathbb{E}_H[\tau^l_H]}{\mathbb{E}_L[\tau^l_L]} - 1\right) = & \ \tfrac{\mathbb{E}_H[l^T_{\tau^l_H}] + \mathbb{E}_L[l^T_{\tau^l_L}]}{\mathbb{E}_L[\tau^l_L]} + c\end{aligned}$$

Note that $|\Delta l^T| < M$ implies $\mathbb{E}_L[l^T_{\tau^l_L}] \geq \ln \frac{\varepsilon}{1-\varepsilon} - M$ and by definition $\mathbb{E}_H[l^T_{\tau^l_H}] \geq \ln \frac{1-\varepsilon}{\varepsilon}$, hence

$$\mathbb{E}_H[\Delta l^T](\frac{\mathbb{E}_H[\tau^l_H]}{\mathbb{E}_L[\tau^l_L]} - 1) \geq \frac{-M}{\mathbb{E}_L[\tau^l_L]} + c$$

On the other hand, by Wald's lemma,

$$\mathbb{E}_L[\tau^l_L] = \frac{\mathbb{E}_L[l^T_{\tau^l_L}]}{\mathbb{E}_L[\Delta l^T]} \geq \frac{\ln \frac{1-\varepsilon}{\varepsilon}}{M} > 0$$

so

$$\mathbb{E}_H[\Delta l^T](\frac{\mathbb{E}_H[\tau^l_H]}{\mathbb{E}_L[\tau^l_L]} - 1) \geq -\frac{M^2}{\ln \frac{1-\varepsilon}{\varepsilon}} + c$$

Since $0 < \mathbb{E}_H[\Delta l^T] < M$,

$$\begin{aligned}\tfrac{\mathbb{E}_H[\tau^l_H]}{\mathbb{E}_L[\tau^l_L]} - 1 > & \ -\tfrac{M}{\ln \frac{1-\varepsilon}{\varepsilon}} + \tfrac{c}{\mathbb{E}_H[\Delta l^T]} \\ = & \ -\tfrac{M}{\ln \frac{1-\varepsilon}{\varepsilon}} + \tfrac{-\mathbb{E}_H[\Delta l^T] - \mathbb{E}_L[\Delta l^T]}{\mathbb{E}_H[\Delta l^T]} \\ = & \ -\tfrac{M}{\ln \frac{1-\varepsilon}{\varepsilon}} - 1 + \left|\tfrac{\mathbb{E}_L[\Delta l^T]}{\mathbb{E}_H[\Delta l^T]}\right| \text{ as } \mathbb{E}_L[\Delta l^T] < 0.\end{aligned}$$

Note that $\ln \frac{1-\varepsilon}{\varepsilon} \to \infty$ as $\varepsilon \to 0$, hence $\forall \delta > 0$, $\exists \overline{\varepsilon} > 0$ such that $\forall \varepsilon < \overline{\varepsilon}$, $\frac{M}{\ln \frac{1-\varepsilon}{\varepsilon}} < \delta$. The proof for the other case is symmetric. ■

Proof of Proposition 3.3.1. Consider the subgame starting from auction k.

Note that Δ_k is statistically independent of the underlying state and all private values, hence we can integrate it out when calculating the expected payoff of bidder ik:

$$
\begin{aligned}
& \mathbb{E}[u_{ik}(b, b^*_{-ik}; w_{ik})|w_{ik}, \boldsymbol{\pi}_k] \\
= \quad & \mathbb{P}(b^{*(2)}_k < b|w_{ik}, \boldsymbol{\pi}_k)\Big\{-\mathbb{E}(b^{*(2)}_k|w_{ik}, \boldsymbol{\pi}_k) \\
+ \quad & \sum_{x=1}^{\infty}\alpha(1-\alpha)^{x-1}\left[\sum_{s=0}^{x-1}\delta^s w_{ik} + \delta^x \mathbb{E}(b^{*(2)}_{k+1}|w_{ik}, \boldsymbol{\pi}_k, \Delta_k = x)\right]\Big\} \\
= \quad & \mathbb{P}(b^{*(2)}_k < b|w_{ik}, \boldsymbol{\pi}_k)\Bigg\{-\mathbb{E}(b^{*(2)}_k|w_{ik}, \boldsymbol{\pi}_k) + \sum_{x=1}^{\infty}\alpha(1-\alpha)^{x-1}\frac{1-\delta^x}{1-\delta}w_{ik} \\
+ \quad & \sum_{x=1}^{\infty}\left[\frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x\ \mathbb{E}(b^{*(2)}_{k+1}|w_{ik}, \boldsymbol{\pi}_k, \Delta_k = x)\right]\Bigg\} \\
= \quad & \mathbb{P}(\theta_{t_k} = H|w_{ik}, \boldsymbol{\pi}_k)\mathbb{P}(b^{*(2)}_k < b|\theta_{t_k} = H)\Big\{-\mathbb{E}(b^{*(2)}_k|\theta_{t_k} = H) + \sigma w_{ik} \\
+ \quad & \sum_{x=1}^{\infty}\left[\frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x E(b^{*(2)}_{k+1}|\theta_{t_k} = H, \Delta_k = x)\right]\Big\} \\
+ \quad & \mathbb{P}(\theta_{t_k} = L|w_{ik}, \boldsymbol{\pi}_k)\mathbb{P}(b^{*(2)}_k < b|\theta_{t_k} = L)\Big\{-\mathbb{E}(b^{*(2)}_k|\theta_{t_k} = L) + \sigma w_{ik} \\
+ \quad & \sum_{x=1}^{\infty}\left[\frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x\ \mathbb{E}(b^{*(2)}_{k+1}|\theta_{t_k} = L, \Delta_k = x)\right]\Big\}
\end{aligned}
$$

For convenience, let us introduce the following notation:

$z_k \equiv w^{(2)}_k \equiv v^{(2)}_{t_k}$:	2^{nd} highest realization of private values in auction k
$\boldsymbol{\rho}^x_j \ \equiv ((P^x)_{j1}, (P^x)_{j2})$	the j-th row of matrix P^x, $j \in \{1, 2\}$, $\forall x \in \mathbb{N}^+$
$\boldsymbol{e}_{k+1} \equiv (e^H_{k+1}, e^L_{k+1})^\intercal$ $e^\theta_{k+1} \equiv \mathbb{E}(b^{*(2)}_{k+1}|\theta_{t_{k+1}} = \theta)$	the expectation of equilibrium resale revenue conditional on the state of next auction
$\widetilde{\boldsymbol{\gamma}}_k = (\widetilde{\gamma}_k, 1 - \widetilde{\gamma}_k)$ $\widetilde{\gamma}_k \equiv \frac{\pi_k f_H(z_k)h_H(z_k)}{\pi_k f_H(z_k)h_H(z_k) + (1-\pi_t)f_L(z_k)h_L(z_k)}$	the belief used by the 2nd-highest-value bidder of auction k in his bidding function.

Using the notations above, we have

$$
\mathbb{E}(b^{*(2)}_{k+1}|\theta_{t_k} = H, \Delta_k = x) = \boldsymbol{\rho}^x_1\boldsymbol{e}_{k+1};
$$
$$
\mathbb{E}(b^{*(2)}_{k+1}|\theta_{t_k} = L, \Delta_k = x) = \boldsymbol{\rho}^x_2\boldsymbol{e}_{k+1}.
$$

Assuming monotone and symmetric bidding strategy, we can rewrite bidder ik's problem as

$$\begin{aligned}\max_b \quad & \pi_{ik} \int_0^{b^{*-1}(b;\boldsymbol{\pi}_k)} \left[-b^*(y;\boldsymbol{\pi}_k) + \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x \boldsymbol{\rho}_1^x \boldsymbol{e}_{k+1} \right] h_H(y)\,\mathrm{d}y \\ + \quad & (1-\pi_{ik}) \int_0^{b^{*-1}(b;\boldsymbol{\pi}_k)} \left[-b^*(y;\boldsymbol{\pi}_k) + \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x \boldsymbol{\rho}_2^x \boldsymbol{e}_{k+1} \right] h_L(y)\,\mathrm{d}y\end{aligned}$$

where $\pi_{ik} \equiv \mathbb{P}(\theta_{t_k} = H | w_{ik}, \boldsymbol{\pi}_k) = \frac{\pi_k f_H(w_{ik})}{\pi_k f_H(w_{ik}) + (1-\pi_k) f_L(w_{ik})}$, bidder ik's posterior about θ_{t_k}. Note that $\boldsymbol{e}_{k+1}$, the expected equilibrium resale revenue, will depend on x, the realization of Δ_k, and y, the realization of z_k, through public belief $\boldsymbol{\pi}_{k+1}$, therefore it cannot be taken out of the integral.

FOC yields

$$\begin{aligned}0 \quad = \quad & \pi_{ik} \left[-b^*(w_{ik};\boldsymbol{\pi}_k) + \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x \boldsymbol{\rho}_1^x \boldsymbol{e}_{k+1} \right] h_H(w_{ik}) \\ + \quad & (1-\pi_{ik}) \left[-b^*(w_{ik};\boldsymbol{\pi}_k) + \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x \boldsymbol{\rho}_2^x \boldsymbol{e}_{k+1} \right] h_L(w_{ik})\end{aligned}$$

Using $\boldsymbol{\gamma}_{it}$ and B defined in the proposition we can rewrite the FOC as

$$\begin{aligned}b^*(w_{ik};\boldsymbol{\pi}_k) \quad &= \quad \sigma w_{ik} + \boldsymbol{\gamma}_{ik} \left[\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x P^x \boldsymbol{e}_{k+1} \right] \\ &= \quad \sigma w_{ik} + \left[\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^x (\boldsymbol{\gamma}_{ik} P^x) \boldsymbol{e}_{k+1} \right]\end{aligned}$$

Now we need to solve for the equilibrium object $\boldsymbol{e}_{k+1}$.

Since the bidding strategy is time-invariant,

$$\begin{aligned} e^{\theta}_{k+1} &\equiv \mathbb{E}(b^{*(2)}_{k+1}|\theta_{t_{k+1}} = \theta) \\ &= \mathbb{E}\left[b^*(z_{k+1}; \boldsymbol{\pi}_{k+1}(\boldsymbol{\pi}_k, y = w_{ik}, x; P))\big|\theta_{t_{k+1}} = \theta\right] \\ &= \mathbb{E}\left[\sigma z_{k+1} + \sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x'}(\widetilde{\gamma}_{k+1}P^{x'})\boldsymbol{e}_{k+2}\middle|\theta_{t_{k+1}} = \theta\right] \end{aligned}$$

Put it back into the bidding function above,

$$\begin{aligned} &b^*(w_{ik}; \boldsymbol{\pi}_k) \\ = \quad & \sigma w_{ik} + \left[\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x}(\boldsymbol{\gamma}_{ik}P^{x})\boldsymbol{e}_{k+1}\right] \\ = \quad & \sigma w_{ik} + \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x}(\boldsymbol{\gamma}_{ik}P^{x})\sigma\begin{pmatrix} E[z_{k+1}|\theta_{t_{k+1}} = H] \\ E[z_{k+1}|\theta_{t_{k+1}} = L] \end{pmatrix} \\ + \quad & \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x}(\boldsymbol{\gamma}_{ik}P^{x})\begin{pmatrix} \mathbb{E}\left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x'}(\widetilde{\gamma}_{k+1}P^{x'})\boldsymbol{e}_{k+2}\middle|\theta_{t_{k+1}} = H\right] \\ \mathbb{E}\left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x'}(\widetilde{\gamma}_{k+1}P^{x'})\boldsymbol{e}_{k+2}\middle|\theta_{t_{k+1}} = L\right] \end{pmatrix} \\ = \quad & \sigma w_{ik} + \sigma\sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x}\mathbb{E}_{ik}[z_{k+1}] \\ + \quad & \sum_{x=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x}\mathbb{E}_{ik}\left[\sum_{x'=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x'}(\widetilde{\gamma}_{k+1}P^{x'})\boldsymbol{e}_{k+2}\right] \end{aligned}$$

where $\mathbb{E}_{ik}[\cdot]$ denote the expectation of bidder ik conditional on the event that he wins auction k and the highest value among others is exactly equal to his value, and his waiting time for resale is x.

Now consider $\boldsymbol{e}_{k+2}$. Let us label the 2nd highest bidder in auction k as bidder $\widetilde{k}$, $\forall k \in \mathbb{N}$.

$$\begin{aligned} e^{\theta}_{k+2} &\equiv \mathbb{E}\left[b^{*(2)}_{k+2}|\theta_{t_{k+2}} = \theta\right] \\ &= \mathbb{E}\left[b^*(z_{k+2}; \boldsymbol{\pi}_{k+2}(\boldsymbol{\pi}_{k+1}, y = z_{k+1}, x; P))|\theta_{t_{k+2}} = \theta\right] \\ &= \mathbb{E}\left\{\sigma z_{k+2} + \sum_{x''=1}^{\infty} \frac{\alpha}{1-\alpha}(\delta - \alpha\delta)^{x''}(\widetilde{\gamma}_{k+2}P^{x''})\boldsymbol{e}_{+3}\middle|\theta_{t_{k+2}} = \theta\right\} \end{aligned}$$

Hence

$$\left(\widetilde{\boldsymbol{\gamma}}_{k+1}P^{x'}\right)\boldsymbol{e}_{k+2} = \sigma\mathbb{E}_{\widetilde{k+1}}[z_{k+2}] + \mathbb{E}_{\widetilde{k+1}}\left[\sum_{x''=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x''}\left(\widetilde{\boldsymbol{\gamma}}_{k+2}P^{x''}\right)\boldsymbol{e}_{k+3}\right]$$

where $\mathbb{E}_{\widetilde{k+1}}[\cdot]$ denote the expectation of bidder $\widetilde{k+1}$ conditional on the event that he wins auction $k+1$ and the highest value among others is exactly equal to his value, and his waiting time for resale is x'.

Plugging back this value into the bidding function of bidder ik,

$$\begin{aligned}
& b^*(w_{ik};\boldsymbol{\pi}_k) \\
= \; & \sigma w_{ik} + \sigma\sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x\mathbb{E}_{ik}[z_{k+1}] \\
+ \; & \sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x\mathbb{E}_{ik}\left[\sum_{x'=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x'}\sigma\mathbb{E}_{\widetilde{k+1}}[z_{k+2}]\right] + r_{ik} \\
= \; & \sigma w_{ik} + \sigma\sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x\mathbb{E}_{ik}[z_{k+1}] \\
+ \; & \sigma\sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x\left[\sum_{x'=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x'}E_{ik}[z_{k+2}]\right] + r_{ik} \\
= \; & \sigma w_{ik} + \sigma\sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x\boldsymbol{\gamma}_{ik}P^x\boldsymbol{c} \\
+ \; & \sigma\sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x\left[\sum_{x'=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x'}\boldsymbol{\gamma}_{ik}P^{x+x'}\boldsymbol{c}\right] + r_{ik} \\
= \; & \sigma w_{ik} + \sigma\boldsymbol{\gamma}_{ik}\left[\sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^xP^x\right]\boldsymbol{c} + \sigma\boldsymbol{\gamma}_{ik}\left[\sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^xP^x\right]^2\boldsymbol{c} + r_{ik} \\
= \; & \sigma w_{ik} + \sigma\boldsymbol{\gamma}_{ik}B\boldsymbol{c} + \sigma\boldsymbol{\gamma}_{ik}B^2\boldsymbol{c} + r_{ik}
\end{aligned}$$

The second equation comes from law of iterated expectation and the fourth equation comes from the fact that waiting time is i.i.d. across auctions.

The residual term r_{ik} is:

$$\sum_{x=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^x\mathbb{E}_{ik}\left[\sum_{x'=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x'}\sigma\mathbb{E}_{\widetilde{k+1}}\left[\sum_{x''=1}^{\infty}\frac{\alpha}{1-\alpha}(\delta-\alpha\delta)^{x''}(\widetilde{\boldsymbol{\gamma}}_{k+2}P^{x''})\boldsymbol{e}_{k+3}\right]\right]$$

Note that the expected present value of the resale revenue from auction $k+m$ for bidder ik goes to 0 as m goes to infinity, due to the existence of discount rate δ. Hence we can recursively solve for the bidding function following the argument above, and finally get

$$\begin{aligned} b^*(w_{ik};\boldsymbol{\pi}_k) &= \sigma w_{ik} + \sigma\boldsymbol{\gamma}_{ik}\left(\sum_{s=1}^{\infty} B^s\right)\boldsymbol{c} \\ &= \sigma[w_{ik} + \boldsymbol{\gamma}_{ik} B(I-B)^{-1}\boldsymbol{c}] \end{aligned}$$

To complete the proof, we need to verify that this is indeed a monotone bidding function.

Let $D = \begin{bmatrix} d_{11} & d_{12} \\ d_{21} & d_{22} \end{bmatrix} \equiv B(I-B)^{-1}$ and rewrite $b^*(w_{ik};\boldsymbol{\pi}_k)$ as

$$\begin{aligned} b^*(w_{ik};\boldsymbol{\pi}_k) &= \sigma(w_{ik} + \boldsymbol{\gamma}_{ik} D\boldsymbol{c}) \\ &= \sigma\left\{w_{ik} + d_{21}c_H + d_{22}c_L + \gamma_{ik}\left[(d_{11} - d_{22} + d_{12} - d_{21})c_L + (d_{11} - d_{21})\Delta c\right]\right\} \end{aligned}$$

where $\Delta c \equiv c_H - c_L > 0$ and d_{ij}, $i,j = 1,2$, to be determined.

The matrix $B = \alpha\delta P[I - \delta(1-\alpha)P]^{-1}$ is equal to:

$$\begin{aligned} B &= \alpha\delta \begin{bmatrix} \rho_H & 1-\rho_H \\ 1-\rho_L & \rho_L \end{bmatrix} \begin{bmatrix} 1-\delta(1-\alpha)\rho_H & \delta(1-\alpha)(1-\rho_H) \\ -\delta(1-\alpha)(1-\rho_L) & 1-\delta(1-\alpha)\rho_L \end{bmatrix}^{-1} \\ &= \frac{\alpha\delta}{(1-\delta(1-\alpha))(1-\delta(1-\alpha)(\rho_H+\rho_L-1))} \cdot \\ & \quad \begin{bmatrix} \rho_H - \delta(1-\alpha)(\rho_H+\rho_L-1) & 1-\rho_H \\ 1-\rho_L & \rho_L - \delta(1-\alpha)(\rho_H+\rho_L-1) \end{bmatrix} \equiv \kappa \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \end{aligned}$$

Therefore

$$
\begin{aligned}
& D \\
= \quad & B(I-B)^{-1} = \kappa \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \begin{bmatrix} 1-\kappa b_{11} & -\kappa b_{12} \\ -\kappa b_{21} & 1-\kappa b_{22} \end{bmatrix}^{-1} \\
= \quad & \frac{1}{(1-\kappa b_{11})(1-\kappa b_{22})-\kappa^2 b_{12} b_{21}} \cdot \begin{bmatrix} \kappa b_{11}(1-\kappa b_{22})+\kappa^2 b_{12} b_{21} & \kappa b_{12} \\ \kappa b_{21} & \kappa b_{22}(1-\kappa b_{11})+\kappa^2 b_{12} b_{21} \end{bmatrix}
\end{aligned}
$$

Plugging back the value of κ and $b_{ij}, i, j = 1, 2$ we have

$$
\begin{aligned}
(1-\kappa b_{11})(1-\kappa b_{22})-\kappa^2 b_{12} b_{21} &= \tfrac{(1-\delta)(1-\delta(\rho_H+\rho_L-1))}{(1-\delta+\alpha\delta)(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}; \\
\kappa b_{11}(1-\kappa b_{22})+\kappa^2 b_{12} b_{21} &= \tfrac{\alpha\delta}{(1-\delta+\alpha\delta))(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}(\rho_H-\delta(\rho_H+\rho_L-1)); \\
\kappa b_{22}(1-\kappa b_{11})+\kappa^2 b_{12} b_{21} &= \tfrac{\alpha\delta}{(1-\delta+\alpha\delta)(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}(\rho_L-\delta(\rho_H+\rho_L-1)); \\
\kappa b_{12} &= \tfrac{\alpha\delta}{(1-\delta+\alpha\delta)(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}(1-\rho_H); \\
\kappa b_{21} &= \tfrac{\alpha\delta}{(1-\delta+\alpha\delta)(1-\delta(1-\alpha)(\rho_H+\rho_L-1))}(1-\rho_L)
\end{aligned}
$$

Therefore

$$
D = \frac{\alpha\delta}{(1-\delta)(1-\delta(\rho_H+\rho_L-1))} \begin{bmatrix} \rho_H-\delta(\rho_H+\rho_L-1) & 1-\rho_H \\ 1-\rho_L & \rho_L-\delta(\rho_H+\rho_L-1) \end{bmatrix}
$$

If we plug back the elements of D into the bidding function we get

$$
b^*(w_{ik};\boldsymbol{\pi}_k) = \sigma\left\{w_{ik} + \frac{\alpha\delta}{1-\delta}\left[c_L + \frac{1-\rho_L}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right] + \gamma_{ik}\frac{\alpha\delta(\rho_H+\rho_L-1)}{1-\delta(\rho_H+\rho_L-1)}\Delta c\right\}
$$

Since $\rho_H + \rho_L \in [1, 2]$ and γ_{ik} is strictly monotone in w_{ik}, $b^*(w_{ik};\boldsymbol{\pi}_k)$ is clearly strictly monotone in w_{ik} as well.

Lastly, notice that the $b^*(w_{ik};\boldsymbol{\pi}_k)$ is indeed the unique solution to the FOC, which implies that the expected payoff of each bidder would be a single-peaked function of her bid, hence FOC is sufficient for optimality. ■

References

ACEMOGLU, D., M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR (2011): "Bayesian Learning in Social Networks," *Review of Economic Studies*, 78(4), 1201-1236.

ALI, S. N., AND N. KARTIK (2012): "Herding with Collective Preferences," *Economics Theory*, 51(3), 601-626

AVERY, C., AND P. ZEMSKY (1998): "Multidimensional Uncertainty and Herd Behavior in Financial Markets," *American Economic Review*, 88(4), 724-748.

BANERJEE, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107(3), 797-817.

BANERJEE, A. V., AND D. FUDENBERG (2004): "Word-of-mouth Learning," *Games and Economic Behavior*, 46(1), 1-22.

BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades," *Journal of Political Economy*, 100(5), 992-1026.

BIKCHANDANI, S., AND C. HUANG (1989): "Auctions with Resale Market: An Exploratory Model of Treasury Bill Markets," *The Review of Financial Studies*, 2(3), 311-339.

CAI, H., Y. CHEN, AND H. FANG (2009): "Observational Learning: Evidence from a Randomized Natural Field Experiment," *American Economic Review*, 99(3), 864-882.

CALLANDER, S., AND J. HORNER (2009): "The Wisdom of the Minority," *Journal of Economic Theory*, 144(4), 1421–39.

CELEN, B., AND S. KARIV (2004): "Observational Learning under Imperfect Information," Games and Economics Behavior, 47(1), 72-86.

CHAMLEY, C. P. (2004): *Rational Herds: Economic Models of Social Learning*, Cambridge University Press.

Conley, T. G., and C. R. Udry (2010): "Learning About a New Technology: Pineapple in Ghana," *American Economic Review*, 100(1), 35-69.

Cripps, M., and C. Thomas (2014): "Strategic Experimentation in Queues," working paper, University of Texas at Austin.

Dasgupta, A. (2000): "Social Learning with Payoff Complementarities," mimeo, London School of Economics.

Dasgupta, A., and A. Prat (2008): "Information Aggregation in Financial Markets with Career Concerns," *Journal of Economic Theory*, 143(1), 83-113.

Debo, L. G., C. Parlour, and R. Uday (2012): "Signaling Quality via Queues," *Management Science*, 58(5), 876-891.

Demertzis, M., and M. Hoeberichts (2007): "The Cost of Increasing Transparency," *Open Economies Review*, 18(3): 263-280.

Eyster, E., and M. Rabin (2014): "Extensive Imitation Is Irrational and Harmful," *Quarterly Journal of Economics*, 129(4), 1861-1898.

Eyster, E., A. Galeotti, N. Kartik, and M. Rabin (2014): "Congested Observational Learning," *Games and Economic Behavior*, 87, 519-538.

Freedman, B. (1987): "Equipoise and The Ethics of Clinical Research," *The New England Journal of Medicine*, 317(3):141–145.

Garratt, R., and T. Tröger (2006): "Speculation in Standard Auctions with Resale," *Econometrica*, 74(3), 753-769.

Garratt, R., T. Tröger, and C. Zheng (2009): "Collusion via Resale," *Econometrica*, 77(4), 1095-1136.

Gill, D., and D. Sgroi (2008): "Sequential Decisions with Tests," *Games and Economic Behavior*, 63(2), 663-678.

Glosten, L., and P. Milgrom (1985): "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14(1), 71-100.

Guarino, A., H. Harmgart, and S. Huck (2011): "Aggregate Information Cascades," *Games and Economic Behavior*, 73(1), 167-185.

Gupta, M., and B. Lebrun (1999): "First Price Auctions with Resale," *Economics Letters*, 64(2),181-185.

Haile, P. A. (2000): "Partial Pooling at The Reserve Price in Auctions with Resale Opportunities," *Games and Economic Behaviour*, 33, 231-248.

Haile, P. A. (2001): "Auctions with Resale Markets: An Application to U.S. Forest Service Timber Sales," *American Economic Review*, 91(3), 399-427.

Haile, P. A. (2003): "Auctions with Private Uncertainty and Resale Opportunities," *Journal of Economic Theory*, 108(1), 72-110.

Halafir, I. E., and V. Krishna (2008): "Asymmetric Auctions with Resale," *American Economic Review*, 98(1), 87-112.

Hayek, F. A. (1945): "The Use of Knowledge in Society," *American Economic Review* 35(4), 519-530.

Herrera, H., and J. Horner (2012): "A Necessary and Sufficient Condition for Informational Cascades," mimeo, Yale University.

Herrera, H., and J. Horner (2013): "Biased Social Learning," *Games and Economic Behavior*, 80(C), 131-146.

Jackson, M. O. (2008): *Social and Economic Networks*, Princeton University Press.

Johnson, N. L., and S. Kotz (2006): *Encyclopedia of Statistical Sciences*, 2nd Edition, John Wiley & Sons.

Jones, M. C. (2009): "Kumaraswamy's Distribution: A Beta-type Distribution with Some Tractability Advantages," *Statistical Methodology*, 6(1), 70-81.

Ju, N., and J. Miao (2012): “Ambiguity, Learning and Asset Returns,” *Econometrica*, 80(2), 559-591.

Keilson, J., and A. Kester (1977): "Monotone Matrices and Montone Markovian Chains," *Stochastic Processes and Their Applications*, 5:231-241.

Keynes, J.M. (1936): “The General Theory of Employment, Interest, and Money,” *Macmillan Cambridge University Press.*

Lebrun, B. (2012): “Optimality and The English and Second-Price Auctions with Resale,” *Games and Economic Behaviour*, 75(2), 731-751.

Lee, I. H. (1993): “On the Convergence of Informational Cascades,” *Journal of Economic Theory*, 61(2), 395–411.

Milgrom, P. R. (1981): “Good News and Bad News: Representation Theorems and Applications,” *The Bell Journal of Economics*, 12, 380–391.

Milgrom, P. R., and R. J. Weber (1982): “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50(5), 1089-1122.

Moretti, E. (2011): “Social Learning and Peer Effects in Consumption: Evidence from Movie Sales,” *Review of Economics Studies*, 78(1), 356-393.

Morris, S., and H. S. Shin (2002): “Social Value of Public Information,” *American Economic Review*, 92(5): 1521-34.

Morris, S., H.S. Shin, and H. Tong (2006): “Social Value of Public Information: Morris and Shin (2002) Is Actually Pro Transparency, Not Con: Reply,” *American Economic Review*, 96(1): 453-455.

Myerson, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6(1), 58-73.

Papadatos, N. (1995): “Maximum Variance of Order Statistics,” *Annals of The Institute of Statistical Mathematics*, 47(1), 185-193.

Park, A., and H. Sabourian (2011): "Herding and Contrarian Behavior in Financial Markets," *Econometrica*, 79(4), 973-1026.

Scharfstein, D. S., and J. C. Stein (1990): "Herd Behavior and Investment," *American Economic Review*, 80(3), 465-479

Smith, L., and P. Sorensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68(2), 371-398.

Smith, L., and P. Sorensen (2008): "Rational Social Learning with Random Sampling," mimeo.

Smith, L., P. Sorensen, and J. Tian (2012): "Informational Herding, Optimal Experimentation, and Contrarianism," working paper, University of Wisconsin.

Svensson, L. (2006): "Social Value of Public Information: Comment: Morris and Shin (2002) Is Actually Pro Transparency, Not Con," *American Economic Review*, 96(1): 448-452.

Timmermann, A. (1993): "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices," *Quarterly Journal of Economics*, 108(4), 1135-1145.

Virag, G. (2013): "First-Price Auctions with Resale: The Case of Many Bidders," *Economic Theory*, 52(1), 129-163.

Wald, A. (1944): "On Cumulative Sums of Random Variables, " *The Annals of Mathematical Statistics*, 15(3), 283-296.

Weitzman, M. L. (2007): "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, 97(4), 1102-1130.

Zheng, C. Z. (2002): "Optimal Auction with Resale," *Econometrica*, 70(6), 2197-2224.