NBER WORKING PAPER SERIES

INFORMATION CASCADES AND SOCIAL LEARNING

Sushil Bikhchandani
David Hirshleifer
Omer Tamuz
Ivo Welch

Working Paper 28887
http://www.nber.org/papers/w28887

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021

We thank the editors, Steven Durlauf and David Romer, four anonymous referees, Nageeb Ali, Antonio Guarino, Navin Kartik, Eric Rasmusen, Rann Smorodinsky, and Calvin Wright for very helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Information Cascades and Social Learning
Sushil Bikhchandani, David Hirshleifer, Omer Tamuz, and Ivo Welch
NBER Working Paper No. 28887
June 2021
JEL No. D7,D8,D9

ABSTRACT

Social learning is the updating of beliefs based on the observation of others. It can lead to efficient aggregation of information, but also to inaccurate decisions, fragility of mass behaviors, and, in the case of information cascades, to complete blockage of learning. We review the theory of information cascades and social learning, and discuss important themes, insights and applications of this literature as it has developed over the last thirty years. We also highlight open questions and promising directions for further theoretical and empirical exploration.

Sushil Bikhchandani
Anderson School if Management
UCLA
110 Westwood Plaza
Los Angeles, CA 90095
sushil.bikhchandani@anderson.ucla.edu

David Hirshleifer
Marshall School of Business University
of Southern California 3670 Trousdale
Pkwy
Los Angeles, CA 90089
and NBER
hirshlei@marshall.usc.edu

Omer Tamuz
Caltech
1200 E California Blvd
Pasadena, CA 91125
omertamuz@gmail.com

Ivo Welch
Anderson School at UCLA (C519)
110 Westwood Place (951481) Los
Angeles, CA 90095-1482 and
NBER
ivo.welch@anderson.ucla.edu

A data appendix is available at http://www.nber.org/data-appendix/w28887

1 Introduction

People rely heavily on the information of others in forming their opinions and selecting actions. The updating of beliefs based on observation of the actions of others, observation of the consequences of these actions, or conversation with others, is called *social learning*. The effect of social learning on decisions is very well-documented.

The theory of social learning in economics starts with the basic insight that if many individuals each have conditionally independent signals about the state of the word, then if they were to share these signals they would have almost perfect knowledge of the state. If these agents are lined up in sequence and choose between two actions, and each agent passes all the information that the agent possesses to the next agent, then sufficiently late agents obtain almost perfect knowledge of state and almost surely decide correctly.

A key further insight starting with the models of Banerjee (1992*a*) and Bikhchandani, Hirshleifer and Welch (1992) (henceforth BHW) comes from the fact that agents may observe the history of actions but not the history of private signals. Filtering past signals through past actions has a drawback, as actions do not always vary across all possible signal values. In the terminology of Ali (2018*b*), there can be limited responsiveness of actions to signals. A lack of responsiveness limits the information that later agents can glean from past actions.

Even worse, in using their private signals and observations of predecessors, agents choose actions for their own purposes rather than to convey helpful information to later agents. This *information externality* in information transmission hinders social learning. For example, if an agent's action is independent of the agent's private signal, there is an information cascade (defined formally in §2.2).[1] An information cascade

[1] An alternative term, "herding," of Banerjee (1992*a*) has essentially the same meaning as "information cascades" in BHW. However, "herding" has several different meanings in economics, and even within the social learning literature. We therefore use the term "information cascade" for the concept introduced by these two papers.

blocks social learning by the next agent in the queue, and perhaps by further agents as well. In consequence, even with many agents, decisions can often be incorrect.

Indeed, a key implication of these models is that, under appropriate conditions, there will always be an information cascade, because the information available via social observation grows until it overwhelms a single agent's private signal. This makes the agent's action uninformative to later agents. As a result, if agents are ex ante identical, then all later agents face the same decision problem, are also in a cascade, and take exactly the same choice—one that may easily be incorrect.

Nevertheless, we emphasize to readers of this survey the distinction between an information cascade, defined as a situation where an agent ignores the agent's private signal, and a possible consequence of information cascades—that many consecutive agents always take the same action. That consequence (as developed in early cascades models) is interesting and important, but does *not* hold in all models and interesting applications of the concept of information cascades.

The model of BHW has a second key implication, that social outcomes are fragile. Agents in a cascade are somewhat close to indifferent between two action alternatives. So owing to cascades, there is a systematic tendency for later agents to reach a resting point at which behavior is sensitive to small shocks. It follows that a shock, such as the arrival of even low-precision public news can be enough to dislodge a long-standing cascade. So social learning theory offers a possible explanation for why behavior in groups is often idiosyncratic and volatile.[2] This is the case even though all individuals are fully rational, and enough information to make the correct decision is available to society.

Moreover, in a setting in which there is a probability of the state shifting, which changes the full-information optimal action, there can be "fads," wherein the

[2] A possible example is the rise and fall of surgical fashions and quack medical treatments on the part of physicians who rely on their colleagues' practices (Taylor (1979), Robin (1984)). Preference or pressure for conformity would instead tend to stabilize practices.

probability of behavior shifting is much higher than the probability of the state shifting. A further line of research has explored the implications of changing states on information aggregation.

Models of social learning can help explain volatile aggregate behaviors and dysfunctional social outcomes in a range of social and economic domains. In consequence, models of information externalities and cascades have been applied in such fields as anthropology, computer science, economics, law, political science, psychology, sociology, and zoology.

In this survey we provide an overview of social learning theory and information cascades in particular. We consider models with imperfectly rational updating, but with a focus on agents who engage in either Bayesian or quasi-Bayesian updating, and who are trying to make good choices (i.e., to optimize).

The problems of information externalities and of unresponsive actions in social learning have led economists to several directions of inquiry. When people make decisions in sequence and can observe each other, do they eventually make correct choices? In other words, is dispersed information aggregated effectively? How quickly do action choices improve over time? Can psychological bias improve social outcomes? What are the effects of observation of payoff outcomes, of choice in how much private information to acquire, and of observation of predecessors in general social networks?

As a benchmark case for consideration of these and other issues, in § 2 we present what we call the *Simple Binary Model* (SBM). This model provides a simple illustration of the intuition of information cascades. This setting helps isolate the effects of alternative assumptions about signals, actions, and social observation learning process. Under what conditions do information externalities lead to complete blockage of learning—either temporarily or permanently—to fragility with respect to shocks, and to better or worse social outcomes? In § 3 and § 4 we explore the effects on these outcomes of varying the action space and the signal distribution, respectively. Another variation is to allow agents to choose whether to act immediately or to delay, which offers new insight about boom and bust dynamics in investment and market entry contexts (see § 5).

In § 6 we relax the assumption that agents observe all of their predecessors and that agents do not observe predecessors' payoffs. We explain there the surprising conclusion that greater observation of predecessors can produce inferior social outcomes.

In models in which agents can acquire costly private signals there is a second type of information externality, since acquiring private information can confer a benefit on later observers. The availability of social information tends to reduce the incentive to acquire private information, potentially leading to severe impairment or blockage of social learning (see § 7).

§ 8 considers how endogenous price determination in market exchange affects costs of adoption and thereby the social learning process.

People observe and learn from those whom they are socially connected to, so the network of observation and communication affects social outcomes. A topic of study has been whether the information of some subset of individuals unduly influences society in settings with different network structures and in which agents can act either once or repeatedly (see § 9).

Another key direction is the study of social learning with boundedly rational agents (see especially § 6.2). Our coverage of limited rationality focuses on quasi-Bayesian decision makers rather than mechanistic agents.

Finally, in § 10 we cover models of information cascades and social learning in a variety of applied domains, including politics, law, product markets, financial markets, and organizational structure. The social learning approach, by focusing on the role of information externalities, offers new perspectives about a wide range of human behaviors.

The formal modeling of social influence began primarily outside of economics, usually with mechanistic assumptions about updating rules.[3] Interactions via social networks were studied heavily in sociology, as with the influential DeGroot model (DeGroot (1974)), and cultural evolution was studied in anthropology and other fields (Boyd and Richerson (1985)). In the 1990s, economists started to focus upon social learning itself, rather than just as an aspect of market or strategic interactions between small numbers of agents. A newer element offered by the social learning literature is the possibility that people observe or talk to others even when the target of observation has no incentive or intent to influence others. Also integral to this literature is that many agents act and observe *sequentially*, so that the learning process iterates, as is often the case in practice.

Even without learning (the focus of this survey), social interaction can cause behaviors to converge, owing, for example, to payoff externalities or preference interactions (Arthur (1989)), reputation effects (Scharfstein and Stein (1990), Ottaviani and Sørensen (2000)) or a preference for conformity (Bernheim (1994)). Some models also study the dynamics of how social interaction induces convergent or divergent behaviors (Kirman (1993)). Our focus is on convergence that derives from social learning and information cascades. However, most applications involve the interaction of several possible factors, including information, rewards and punishments, and payoff externalities. The integration of social learning and cascades with other factors provides a richer palette of models to describe how society chooses technologies, ideas, governments, organizational choices, conventions, legal precedents, and market outcomes.

For reasons of space, in this survey we omit models of social learning and information cascades that incorporate payoff externalities or utility interactions. Other surveys of social learning include , Gale (1996), Bikhchandani, Hirshleifer and Welch

[3] A highly interdisciplinary literature in mathematics, physics, and sociology studies heuristic models of social learning with repeated actions on networks. Overviews of the mathematics and physics literatures are provided by Castellano, Fortunato and Loreto (2009) and Boccaletti et al. (2006). In sociology, a quantitative literature studies opinion exchanges on networks, partially motivated by the question of how to measure network centrality (see, e.g., Katz (1953) and Bonacich (1987)).

(1998), Chamley (2004*b*), Vives (2010), Acemoglu and Ozdaglar (2011), and Golub and Sadler (2016)). Our review includes coverage of relatively new developments in the study of social networks, repeated moves, and psychological bias in this rapidly evolving field. Complementing our theoretical focus are several reviews with a primarily empirical focus, such as Hirshleifer and Teoh (2009), Anderson and Holt (2008), Sacerdote (2014), and Blume et al. (2011). These surveys cover tests of cascades theory in the experimental laboratory (Anderson and Holt (2008)), in field experiments (Duan, Gu and Whinston (2009)), and with archival data (Tucker, Zhang and Zhu (2013), Amihud, Hauser and Kirsh (2003)).

2 The Simple Binary Model: A Motivating Example

We illustrate several key intuitions about social learning in a setting with binary actions, states and signals, which we call the *Simple Binary Model*, hereafter, SBM. The SBM is a modified version of the binary example of Bikhchandani, Hirshleifer and Welch (1992), hereafter BHW.[4] The SBM illustrates how information externalities cause social learning to be severely impaired. In this setting, information cascades block social learning. The SBM can be extended to illustrate many further concepts. Nevertheless, some of the conclusions of the SBM are more robust than others to changes in the model. We recurrently discuss in this survey when these effects do or do not arise in various other social learning settings.

2.1 Basic Setup: Binary Actions, Signals, and States

Individuals $I_1, I_2, I_3, \ldots$ make choices in sequence. Each agent I_n chooses one of two actions, High ($a_n = H$) or Low ($a_n = L$). The underlying state θ, which is not observed,

[4]The simple binary model is the special case of the BHW model with symmetric binary signals. The model of Banerjee (1992*b*) differs in several substantive ways, as described in § 3.

Table 1: Notation Guide

	Notation	Base case assumption
State	θ	$\theta \in \{L, H\}$
Signal	s_n for agent I_n	$s_n \in \{\ell, h\}$
Action	a_n for agent I_n	$a_n \in \{L, H\}$
Utility	$u(\theta, a)$	$u(\theta,a) = \begin{cases} 1 & \text{if } \theta = a \\ 0 & \text{otherwise} \end{cases}$

takes one of two possible values, H or L. The two states are equally likely ex ante. We will often view the H state as one in which there is high payoff to some activity and the L state in which there is low payoff to that activity. In such cases we call action H "Adopt," and action L "Reject." Table 1 summarizes the notation used in this survey.

Each agent I_n receives a binary symmetric private information signal $s_n = h$ or $s_n = \ell$, with probability $p := \mathbb{P}[s_n = h|\theta = H] > 1/2$ and $\mathbb{P}[s_n = h|\theta = L] = 1 - p$. Signals are independent conditional on θ.

We refer to p as the *signal precision*. Let the *belief precision* of a possible belief q about the state θ be $|q - 0.5|$. The greater the signal precision, the closer the belief is to 0 or 1, implying more certainty about the state.

By Bayes' rule, the posterior probability $\mathbb{P}[\theta = H|s_n = h] = p$, and $\mathbb{P}[\theta = H|s_n = \ell] = 1 - p$. By the symmetry of the SBM, h and ℓ signals have offsetting effects on posterior beliefs, so $\mathbb{P}[\theta = H|s_1 = h, s_2 = \ell] = 1/2$, and if an agent sees or infers certain numbers of predecessor h and ℓ signals, the updated belief depends only on the difference between these numbers.

All agents have the same utility function $u(\theta, a)$, which is equal to 1 if $a = \theta$ and to 0 otherwise. Therefore, each agent chooses the action that is more likely to match the state given her information. If I_n assigns equal probabilities to both states, she is indifferent between the two actions, in which case we assume that she follows her private signal, $a_n = s_n$.[5]

We refer to *social information* in this survey as information derived from observing others, the *social belief* at any step in the sequence as the belief that an outside observer would have based only on the social information of agent I_n, and an agent's *private belief* as the belief implied solely by the agent's private signal. Each agent's action choice is of course based on the agent's full information set. It is common knowledge that each I_n knows the decision model and information structure of predecessors.

We contrast two regimes for social observation:

1. **The Observable Signals Regime:** Each agent observes both the signals and the actions of all predecessors.[6]

2. **The Only-Actions-Observable Regime:** Each agents observes the actions but not the private signals of all predecessors.

In the Observable Signals Regime, the pool of social information always expands, and by the law of large numbers the social belief becomes arbitrarily close to certainty about the correct action—Adopt if $\theta = H$, Reject if $\theta = L$—and thus eventually behave alike. An outcome in which all agents behave alike forever is called a *herd*. We define herds more formally in § 4.2.3.

In the Only-Actions-Observable Regime, each agent's action depends on the publicly observable action history and the agent's own private signal. In this regime, the

[5]The qualitative properties of this model are robust to changes in the tie-breaking rule. For the purpose of providing minimal examples, we sometimes employ different tie-breaking conventions for the behavior of indifferent agents. Although pedagogically convenient, such ties could be eliminated, with similar results, by slightly perturbations of the model parameters.

[6]Actions do not convey any additional information as a predecessor's signal is a sufficient statistic for her action.

precision of the social belief is weakly increasing with n, so it is tempting to conjecture that highly accurate outcomes will again be achieved. But in fact, the precision of the social belief hits a finite ceiling, as we will discuss. It is easy to show that the choices of a few early predecessors determine the actions of all later agents. In consequence, agents still herd — but often upon the wrong action, the choice that yields a lower payoff.

Throughout this review, our default premise will be the Only-Actions-Observable regime; we only mention assumptions when they deviate from the base set of assumptions of the SBM. The SBM, and variations thereof, are rich sources of insight into social learning.

2.2 Why Information Stops Accumulating: Information Cascades

To see why outcomes are inefficient in the Only-Actions-Observable Regime, consider each agent in sequence. Agent I_1, Ann, adopts if her private signal is h and rejects if it is ℓ. Agent I_2, Bob, and all successors can infer Ann's signal perfectly from her decision. So if Ann adopts, and Bob's private signal is h, he also adopts. Bob knows that there were two h signals: he infers one from Ann's actions and has observed one privately. If Bob's signal is ℓ, it exactly offsets Ann's h, making him indifferent between adopting and rejecting. By our tie-breaking convention, Bob follows his own signal and rejects. In this setting, regardless of which signals Ann and Bob receive, each chooses an action that matches his or her signal, and thus their actions reveal their signals to later agents.

Agent I_3, Carol, now faces one of three possible situations: (1) Ann and Bob both adopted (their actions were HH), (2) both rejected (LL), or (3) one adopted and the other rejected (HL or HL). In the HH case, Carol adopts regardless of which signal she received, since the majority of observed or inferred signals is h even if she received an ℓ. In other words, Carol's action is independent of her private information signal—she is in an information cascade, defined as follows:

Definition 1. *An agent is said to be in an* information cascade *if it is optimal for her to* *choose an action which is independent of her private signal.*

We reiterate that the definition of an information cascades does *not* assert that a long string of consecutive agents take the same action. In some interesting settings cascading by one agent does implies such strings, and in some it does not. Furthermore, in some settings without cascades there can still be realizations with identical actions by strings of consecutive agents.[7]

In the model here, crucially, in the HH case, Carol's action provides no information about her signal to her successors. Her action has not improved the public pool of information; the social belief remains unchanged. Hence all later agents face the same decision problem that she did. Once Carol is in an Adopt cascade, all her successors also adopt, ultimately based only on the observed actions of Ann and Bob. Similarly, in the second case, two Rejects (LL) put Carol and all subsequent agents into a Reject cascade.

In the third case, Ann has adopted and Bob has rejected, or vice versa. Carol infers that Ann and Bob observed opposite signals, so the social belief is that the two states are equally likely. Her decision problem is therefore identical to that of Ann, so Carol's decision is to follow her private signal. In turn, this makes the decision problem of agent I_4, Dan, isomorphic to Bob's—he need only pay attention to his latest predecessor.

More generally, taking agents pairwise starting from (I_1, I_2), the only way to avoid an information cascade is for each pair to contain one Adopt and one Reject. If this occurs through any even number of agents, the next agent can infer that the number of past h and ℓ signals was equal, so that past actions can be ignored. Following any such history of paired opposing actions, the next two agents may both receive the same signal (which occurs with probability $p^2 + (1-p)^2$). When the next two agents observe the same signals, they take the same action, so the next agent is in a cascade. Once an agent

[7] In those settings where cascading continues without a break through a string of consecutive agents, for convenience we sometimes refer to the entire string as a cascade.

is in a cascade, so are all successors.[8] It follows that a herd also occurs. Furthermore, since each pair has a fresh chance to start a cascade, a cascade happens eventually, almost surely.

Cascades tend to start early—based on a small preponderance of evidence in the social belief. Even in the least cascade-favorable scenario of very noisy signals ($p \approx 1/2$), by the time the 20$^{\text{th}}$ agent has to decide, the probability of not being in a cascade is already under 0.1%, and the expected number of agents who act on private information before a cascade starts is under four. So the private information of all but a few agents is lost to the group.

Once a cascade starts, no new information about the state becomes public; the accuracy of the social belief plateaus. An early preponderance of either Adopts or Rejects causes all subsequent agents to rationally ignore their private signals. These signals thus never join the public pool of knowledge. Agents disregard their private signals before the social belief becomes very accurate. Specifically, as soon as the social belief is more precise than the signal of just a single agent, the next agent falls into a cascade. We will call this *the logic of information cascades*.[9] The logic of information cascades implies that cascades are often incorrect—they are *idiosyncratic*.

The accuracy improvement from social observation in the Only-Actions-Observable Regime falls far short of the outcome in the Observable Signals Regime, in which eventually everyone makes the correct decision. Indeed, when private signals are very noisy (e.g., $p = 1/2 + \epsilon$), the social outcome is almost pure noise. This falls far short of perfect information aggregation, wherein decisions for later agents would become almost perfectly accurate. Specifically, the increase in accuracy that agents obtain from being able

[8] In general, however, a cascade can occur for just one agent or can continue for any given number of agents.

[9] In a setting with more than two possible signal values, two identical successive actions do not necessarily start a cascade. The general point is that a fairly low-precision social belief can be enough to trigger a cascade.

to observe the actions of predecessors is negligible.[10]

The social outcome in the Only-Actions-Observable Regime depends heavily on the order in which signals arrive. If signals arrive in the order $hh\ell\ell\ldots$, then a cascade starts and everyone adopts. If, instead, the same set of signals arrives in the order $\ell\ell hh\ldots$, everyone rejects. So cascades and welfare are *path dependent*.

2.3 Lessons of the Binary Model

The SBM illustrates a number of key social learning concepts that are empirically testable, and which also obtain in some more general settings.

Conformity: People end up following the behavior of others.

The conclusion that agents conform to the actions of others—even incorrect actions—holds in a wide array of social learning models. The SBM provides a useful benchmark for assessing which assumptions are crucial for this conclusion. In later sections we will see that a similar conclusion can hold with more general action spaces and signal distributions, and under different assumptions about the timing of moves, observability of others, and the decision of whether to acquire information. However, some variations of modeling assumptions such as unbounded private signals, certain psychological biases, and negative payoff externalities can break this conclusion, even when information cascades still occur.

Idiosyncrasy and Path Dependence: Despite a wealth of private information, which in the aggregate would assure the correct action if it could be made fully available for use, the behavior of most agents often ends up being incorrect—idiosyncrasy.

[10] In the Only-Actions-Observable Regime, the probability of a correct cascade (Adopt if and only if $\theta = H$) can be shown to be $\frac{p^2}{p^2+(1-p)^2}$. For a noisy signal as above, this is approximately $1/2 + 2\epsilon$. This is not much better than an agent could do based solely on private information, which results in a probability of choosing correctly of $p = 1/2 + \epsilon$.

Later in this survey we examine the robustness of this conclusion. Specifically, when does society converge to correct behavior? And under what conditions is there persistent idiosyncrasy?

A key problem in the SBM and some other social learning settings is that the actions of a few early agents tend to be decisive in determining the actions and success of a large numbers of successors. Social outcomes have low predictability. In other words, outcomes are path-dependent. This problem raises the question of whether policy can improve outcomes. As such, there has been extensive interest in path dependence even in settings without social learning (Arthur (1989)).

Information Externality: In almost all social learning models, each agent takes an action that is individually optimal, without consideration of the informational benefits to later agents. A greater benefit could be conferred upon later agents if early agents were instead to act in ways that reveal their own signals. The wastage of private information delays or even entirely blocks learning.

In a range of social learning settings, owing to the information externality, agents conform surprisingly quickly, blocking information aggregation. In other settings, the information externality does not lead to complete blockage of learning. But even in such settings, learning is much slower than in the Observable Signals Regime (see §4.3).

Fragility: Fragility is a recurring property of many of the social learning models that we explore in this survey. In the SBM, cascades start at a social belief that is not much more precise than a single private signal. Hence, any comparably informative additional (public or private) signal added to the model could dislodge the cascade—cascades are fragile (we define fragility formally in § 6). In other words, society spontaneously wanders to a position that is highly sensitive to small shocks. Even with more general signal distributions than the SBM, once the social belief is more informative than the most informative signal value, a cascade starts, resulting in fragility.

> This contrasts with settings in which social outcomes tend to be insensitive to small shocks (outside particular parameter values that happen to put the system close to a knife edge; Kuran (1989)).

The SBM and the cascades model of BHW provide some surprising outcomes: complete blockage of learning along with identical action choices. The broader intuition is that information externality limits learning, resulting in a similar but milder possibility that sometimes occurs in other settings: that learning becomes slow, and that actions (even incorrect ones) become very similar.

Due to its simplicity, the SBM serves as an intuition pump for more complex scenarios. It illustrates which assumptions are crucial and allows us to isolate the effects of changes in assumptions. For example, using a different tie-breaking rule, or allowing for finitely many signal values instead of two qualitatively makes no difference, as shown by BHW. We will likewise see that allowing for endogenous signal acquisition or for endogenous timing of decisions only strengthens the key punchlines. Various other settings also maintain information blockage as in the SBM.

However, in other settings the key implications of the SBM do not hold. Furthermore, varying model assumptions produces interesting new phenomena, such as the endogenous sudden onset of avalanches of activity by many agents triggered by the action of a single agent, eventual learning, or shattering of cascades after their formation. Two insights of the SBM, information externality and herding, are preserved in a very wide spectrum of social learning models, resulting in slow learning.[11]

These extensions sometimes involve relaxing the SBM's common knowledge assumptions, for example by allowing for privately known preferences, observations and signal precision. In the SBM, agents observe all predecessors and no one else. In more general models of learning in social networks, different agents may observe different subsets of predecessors, and agents may not always know what their predecessors

[11] However, in the model of Vives (1997), there is an information externality but no cascades or herding, since actions are taken in a continuum with full responsiveness to signals.

have observed. In models with imperfect rationality, the true economic environment is not common knowledge as agents may systematically misestimate the information sets of predecessors. Exploring different directions for generalization, and isolating the effects of different features of the social learning environment, is a focus of much of the remainder of this survey.

3 The Action Space

A key insight provided by the social learning literature is that sparsity of the action space can severely limit learning. This occurs not just for the mechanical reason that more choices can potentially reveal more signal values. The problem is that sparsity can further reduce the incentive of agents to take action choices that convey useful information to others.

In the SBM, agents choose from one of two actions, and receive a utility that is one if the action matches the state, and zero otherwise. More generally, agents choose an action from some action set, and receive a utility that depends on the action and the state. Cascades tend to start quickly, so that agents take actions that are independent of their signals.

At the opposite extreme from the two-action assumption, if the set of available actions is continuous, under reasonable assumptions even a small variation in an agent's belief causes a small shift in action. If so, agents' actions always depend on their private signals, and therefore reveal their private signals. Hence there is no cascade, and late agents eventually learn the state with near perfection.

Between the binary and continuous extremes, the action sets that agents choose from are large but finite. For example, firms may choose one among several technologies to adopt. Moreover, continuous choice sets are often truncated. For example, a scooter can be rented for different time periods, but the length is bounded below at zero.

Furthermore, there is evidence that people sometimes discretize their choice sets, such as restricting their investment orders (price, quantity, or market value) to round numbers (Harris (1991), Ikenberry and Weston (2008), Hirshleifer et al. (2019)).

When agents have more available choices, they can potentially attune their actions more finely to their private signals. This can allow later agents to learn more from their predecessors' actions. This in turn suggests that the problem of incorrect cascades may diminish as the action space is enriched. Indeed, one might think that a sufficiently rich action space will make the problems of social learning completely vanish. We will see that this is only sometimes true.

As an example of a simple setting in which agents eventually learn the state, suppose that the two possible states are $\theta = 0$ or 1. Agent I_n chooses an action a_n from the interval $[0, 1]$ and chooses the payoff that minimizes the mean-squared error, i.e., the utility function is $u(\theta, a_n) = -(\theta - a_n)^2$. In this setting, I_n's optimal action is strictly increasing in her beliefs and therefore perfectly reveals her private information. More generally, Ali (2018*b*) defines a notion called *responsiveness*, which loosely speaking means that actions are sensitive even to small differences in beliefs. He shows that responsiveness is a sufficient condition for eventual learning of the true state, for any private signal distribution. Intuitively, responsiveness causes learning predecessors' signals, as discussed earlier.

This reasoning may seem to suggest that a continuous action space always eliminates cascades, and results in eventual learning of the state. However, even with a continuous action space, an agent's optimal action could be unresponsive to the signal. This occurs if there is a particular action that is optimal at more than one belief, in which case it is optimal at an interval of beliefs. For example, this occurs if the action space is $[0, 3/4]$, and the utility is, as above, $u(\theta, a_n) = -(\theta - a_n)^2$. In this case $a_n = 3/4$ is optimal for all beliefs above $3/4$. Cascades can then form if the social belief rises enough above $3/4$ so that the next agent will take action $a_n = 3/4$ regardless of her private signal.

More generally, when there is an interval of beliefs supporting an action, it may be that no private signal can draw the posterior belief outside the interval. Hence, no information is revealed about private signals once this region is entered, and a cascade starts. When the set of actions is finite this always happens, and so a cascade starts (and never ends) almost surely (in the spirit of Proposition 1 in BHW). With positive probability, this cascade is incorrect.

Restricting to mean squared error preferences and binary signals, and allowing for multiple states, Lee (1993) obtains necessary and sufficient conditions for fully revealing information cascades. A sufficient condition is that the action space be finite. In this case, choices cannot be responsive; at least one action has the property of being optimal for a range of probability beliefs. This makes actions less informative about agents' private signals. In consequence, incorrect cascades can arise, i.e., there is idiosyncrasy.

Models with a continuum of actions usually have continuous utility functions. An early and prominent exception is the model of Banerjee (1992*b*), in which incorrect cascades occur despite a continuous action space. The unknown state θ is uniformly distributed on $[0, 1]$, and agents choose an action $a \in [0, 1]$. Each agent obtains a payoff of 1 if she chooses action $a = \theta$ and a payoff of zero if $a \neq \theta$. Each agent receives either no signal (and is uninformed) or one signal. An informed agent receives a signal about θ that is either fully revealing or is pure noise (in which case it is uniform on $[0, 1]$). An informed agent does not know which of these two possibilities is the case. In this model, the optimal action is not responsive to small changes in an agent's signal. Thus, early agents may fix upon an incorrect action—an incorrect cascade.

An interesting question is whether a larger set of available action choices improves social learning by making agents' decisions more informative to later agents. For example, does this help prevent incorrect cascades, or reduce their adverse effects by inducing incorrect cascades that are closer to the correct action. Certainly, sometimes adding even a few actions can increase welfare substantially (Talley (1999)), and when the set of actions is large enough to reveal the private belief to a high degree of accuracy, cascades will take a long time to form, and will be likely to be correct. However, the benefit from adding actions is non-monotonic; adding an action sometimes hastens the onset of cascades and reduces welfare. This can occur, for example, when an attractive intermediate action is added which agents fix upon prematurely instead of using their private signals. This phenomenon deserves explicit exploration.

4 General Signal Distributions

Private information is often finer-grained than the binary signals of the SBM. This raise several questions about social learning outcomes. What is the effect of the private signal distribution on the probability of agents settling on the correct action? When do agents eventually herd upon identical actions? For which signal distributions do information cascades still arise? Understanding these questions requires technical specifics, so this section is more formal than much of this survey.

As in the SBM, we assume a binary state and binary actions, and that I_n takes an action a_n after observing I_1 through I_{n-1}, with the goal of matching the state. We also still assume that private signals s_n are i.i.d. conditional on the state θ, but we now allow a general signal distribution. We still assume that private signals are inconclusive about the state, i.e., the private belief as defined in § 2, $b_n = \mathbb{P}[\theta = H|s_n]$, is in $(0, 1)$. Let the *social belief* be defined as

$$P_n = \mathbb{P}[\theta = H|a_1, \ldots, a_n],$$

the belief held by an outsider who can observe the agents' actions and has no private information. It is convenient to define the *social log-likelihood ratio* $R_n = \log \frac{P_n}{1-P_n}$.

The social belief P_n (or, equivalently, the public log-likelihood ratio R_n) converges to some probability almost surely as the number of agents becomes large, though not necessarily to the truth. This is a general property of any sequence of beliefs of

Bayesian agents who collect more information over time. This convergence follows from the Martingale Convergence Theorem (MCT), and the fact that the sequence $P_1, P_2, \ldots$ is a *bounded martingale*. More generally, the MCT is a useful tool for analyzing asymptotic outcomes in social learning settings, as it ensures convergence of beliefs under rather mild assumptions.

4.1 Bounded vs. Unbounded Private Signals

The distribution of private signals, and hence of the private beliefs b_n, is key to understanding whether agents eventually learn the state, whether agents converge upon the same action, and whether cascades occur. If agents may receive arbitrarily accurate signals, then it is intuitive that agents will end up making very good decisions. In contrast, bad outcomes are possible when signals are not arbitrarily accurate, as in the SBM.

The concept of arbitrarily accurate signals is formalized with the notion of unboundedness. We say that private signals are *bounded* if the resulting private belief b_n is not arbitrarily extreme, i.e., there is some $\varepsilon > 0$ such that the belief is supported between ε and $1 - \varepsilon$:

$$b_n > \varepsilon \text{ and } b_n < 1 - \varepsilon \quad \text{almost surely}.$$

We say that a private signal is *unbounded* if, for every $\varepsilon > 0$, the probabilities $\mathbb{P}[b_n < \varepsilon]$ and $\mathbb{P}[b_n > 1 - \varepsilon]$ are both non-zero. With an unbounded signal distribution, signal realization sometimes have arbitrarily high informativeness. Boundedness is a property of the possible beliefs induced by a signal, rather than the range of the signal values per se.[12]

[12] An example of a bounded signal on $[0, 1]$ is one which has density $f_L(s) = 3/2 - s$ in state L and $f_H(s) = 1/2 + s$ in state H. If, instead, the conditional densities in the two states are $f_L(s) = 2 - 2s$ and $f_H(s) = 2s$, then the signal is unbounded. The log-likelihood ratio $\log \frac{f_L(s)}{f_H(s)}$ (and therefore the belief) is bounded in the first case, but can take any value in the real line in the second case.A signal can be

As shown in Smith and Sørensen (2000), agents eventually learn the state (in a sense to be made precise later) if and only if signals are unbounded. Kartik et al. (2022) point out that when there are three or more states, unbounded signals (appropriately defined for the case of multiple states) are ruled out if the signal distribution has full support and is monotone. The authors obtain sufficient conditions on signals and utility functions that guarantee eventual learning of the state.

4.2 Asymptotic Learning, Cascades, and Limit Cascades

Agents sometimes make incorrect decisions even in the long run, as we have seen in the case of information cascades. To put this in a broader perspective, we now define asymptotic learning—a situation in which the social belief becomes arbitrarily accurate. We then consider the conditions under which asymptotic learning occurs. We also define limit cascades, which are a variant of cascades.

4.2.1 Asymptotic Learning

We say that there is *asymptotic learning* if the social belief P_n almost surely tends to 1 with n when the state is H, and to 0 when the state is L. Two other possible definitions for asymptotic learning are the following:

1. The sequence of actions $a_1, a_2, \ldots$ converges almost surely to θ, i.e., all agents from some I_n on take the action that matches the state.
2. The probability that agent I_n takes an action that matches the state tends to one with n.

In this setting both of these are equivalent to asymptotic learning. Intuitively, if beliefs are almost perfectly accurate, so are actions. And for actions to be almost always accurate

neither bounded nor unbounded; for example, it could be the case that the private beliefs can take values arbitrarily close to 1, but are bounded away from 0.

under the infinite range of possible signal realization sequences, beliefs must also almost always be highly accurate. When asymptotic learning fails, there is idiosyncrasy, as defined informally in § 2.

4.2.2 Information Cascades and Limit Cascades

Our definition of information cascade in § 2, that the agent takes the same action regardless of her private signal, can be rephrased as follows. An information cascade occurs when, for some agent I_n, the social log-likelihood ratio R_{n-1} is either so high or so low that no private signal can influence the action, i.e., cause I_n's action to depend on I_n's private signal. This occurs if and only if the social belief remains unchanged after observing the action.

A *limit cascade* occurs if the agents' actions converge, the limiting action is sometimes incorrect, and each agent chooses either action with positive probability. As with asymptotic learning, the probability that an agent chooses differently from her predecessor decreases quickly enough that, with probability one, from some point on, all agents choose the same action. As with cascades, this action is often incorrect, i.e., the social outcome is idiosyncratic. More formally, in a limit cascade the social log-likelihood ratio R_n tends to a limit that is not $+\infty$ or $-\infty$, but (unlike cascades proper) does not reach that limit in finite time (Smith and Sørensen (2000)); Gale (1996) provides an example of essentially the same phenomenon. Thus, the outside observer's belief converges to an interior point in $[0, 1]$. Empirically, the implication is essentially identical to that of information cascades: society may fix upon an incorrect action forever.

How much social information is aggregated under social learning? The *Two-Signal Principle* states that under appropriate conditions, very little—less than two reinforcing occurrences of the most informative possible private signal value (or, with unbounded signals, the supremum of that informativeness).

This is obviously true with unbounded signals, where the upper bound is infinite. For the case of bounded signals, recall from § 2.1 that when the state is binary, the precision of a possible belief r about the state θ is defined as $|r - 0.5|$. Consider for simplicity a setting in which the prior belief is symmetric.

To understand the Two-Signal Principle, observe that there is no way for an agent I_n's action to reflect more than two maximally informative private signal realizations unless at least one earlier agent's action reflects more than one such realization. So consider any point in the sequence in which more than one maximally informative private signal realization is incorporated into the social belief as observed by I_n. Then I_n will follow the action implied by the social belief regardless of I_n's private signal, i.e., I_n is in a cascade. It follows that I_n's action is not informative, which means that it does not increase the precision of the social belief. So I_{n+1} and all later agents are also in a cascade, and also do not increase the precision of the social belief. In other words, the zone of social beliefs that incorporates between one and two private signals is impassable.

4.2.3 Conditions for the Three Possible Social Learning Outcomes

There are three possibilities for the asymptotic outcome of the process: (i) an information cascade, (ii) a limit cascade, and (iii) neither of the two, in which agents herd on the correct action. These three cases can be equivalently characterized by the limiting behavior of the social belief P_n, which must converge by the Martingale Convergence Theorem to some $P_\infty \in [0, 1]$. In case (i) P_n converges in finite time to an interior point, $P_\infty \in (0, 1)$. In case (ii) it again converges to an interior P_∞, but not in finite time. And in case (iii) it becomes decisive, i.e., it converges to $P_\infty \in \{0, 1\}$. So, in this latter case the agents converge to the correct action. Smith and Sørensen (2000) describe the relation between signal structures and these three possible outcomes.

In the SBM, the only possible outcome is an information cascade. As shown by BHW, this holds more generally when the set of possible signal values is finite. When signals are bounded but not finite, either cascades or limit cascades can occur. For example, there are always limit cascades when private signals are distributed uniformly on $[0, 1]$ in state L and have density $f(s) = 1/2 + s$ on $[0, 1]$ in state H (Smith and Sørensen (2000)).

When signals are unbounded, there is asymptotic learning: Almost surely P_n converges to to either 0 or 1 and the actions converge to the state. Asymptotic learning fails when cascades or limit cascades occur with positive probability. In this case, the social belief converges to an interior $P_\infty \in (0, 1)$, and the probability that agent I_n chooses correctly tends to $\max\{P_\infty, 1 - P_\infty\} < 1$.[13]

Why do bounded signals block asymptotic learning? With bounded signals, once the social belief P_n is close enough to either 0 or 1, the private signal of agent n cannot shift the belief enough to change the action.

Unbounded signals imply that information cascades and limit cascades are impossible, and asymptotic learning is obtained. To see why, suppose that R_∞, the limit of R_n, is finite, so that a limit cascade or a cascade occurs. Let q_n be the probability that I_n chooses H conditioned only on social information. This probability depends only on R_n, and is strictly between 0 and 1 for any R_n, since signals are unbounded. In the long run since R_n converges to some finite R_∞, q_n will approach some probability $0 < q_\infty < 1$. So an observer who sees the action sequence is essentially receiving an infinite stream of binary signals of approximately constant precision. This observer's beliefs become perfectly accurate, which contradicts the premise that the social log likelihood ratio is bounded.

To sum up, in this setting, there is asymptotic learning if and only if signals are unbounded. When signals are unbounded, no matter how long past agents have been following a mistaken action, there will eventually be an agent with a strong enough

[13]For example, when a cascade starts, the social belief reaches P_∞. If $P_\infty > 0.5$, the agent adopts, and this is correct with probability P_∞. Similarly, if $P_\infty < 0.5$, the agent rejects, and this is correct with probability $1 - P_\infty$.

signal who will overturn that action. Conversely, when signals are bounded, then the social belief cannot become too accurate, since a cascade would be triggered, blocking further information aggregation.

When there is a cascade, resulting in poor information aggregation, an additional exogenous shock can easily dislodge cascading on the action that was most popular before the shock. So, as in the SBM, outcomes in the setting with general bounded signal distributions can be fragile, a concept developed explicitly in § 6.

A largely unexplored question is the extent to which unbounded signals still succeed in bringing about asymptotic learning when agents are not expected utility maximizers. An exception is Chen (2021), in which agents know the precisions of their own private signals but not of other agents' signals, and are ambiguity averse with respect to this parameter. This can deter agents from breaking a cascade, as agents may fear the worst-case scenario (from the perspective of deviating from the cascade) that predecessors in the cascade are very well-informed. In contrast, when contemplating following one's own signal, an agent is not very fearful of joining the cascade, as the worst that happens is that the agent loses the benefit of the agent's noisy signal. In consequence, even when signals are unbounded, there can be information cascades and no asymptotic learning.

A *herd* is a realization in which all agents behave alike from some point on. Formally, a herd starts from agent I_n if all later agents take the same action, that is, if $a_m = a_n$ for all $m \geq n$. BHW show that when private signals are finitely supported, herding occurs with probability one. In fact this is the case with an unbounded signal structure as well (Smith and Sørensen (2000)).

Crucially, a herd can occur even in the absence of information cascades. Herding is an ex post realization in which all agents from some point on take the same action. In contrast, an information cascade is an ex ante concept: for all possible signal realizations, an agent takes the same action. This said, explicit modeling shows cascades and herding to be connected; see, e.g., the SBM. Furthermore, in settings with limit cascades there are herds with substantial probability (indeed, probability one), but there are no cascades.

The crucial difference between the bounded and the unbounded signal setting is that with unbounded signals there is asymptotic learning: from some point on, all agents take the correct action.[14] This is somewhat surprising, since each agent in the herd has a positive probability of taking either action, i.e., there are no cascades.

4.3 Information Externalities and Speed of Learning

As mentioned in the Introduction, an information externality is a situation where an agent makes a decision in her own interest that affects the information available to some other agent. Information externalities reduce the rate of social learning relative to the socially optimal one, wherein agents adjust their actions to optimally balance the immediate benefits to themselves of choosing the best action with the benefit of conveying information to later agents (according to some social objective).

Under some conditions, we have seen that this wastage of private information completely blocks learning. Under others, there is asymptotic learning. Nevertheless, since agents do not choose their actions for the purpose of conveying information to others, a general principle is that asymptotic learning tends to be inefficiently slow relative to the socially optimal rate.

So a natural question is: Under asymptotic learning, how long does it take for a correct herd to form, and how do information externalities affect the speed of learning? One measure of efficiency is whether the expected time until a correct herd forms (i.e.,

[14]It is largely a matter of convenience whether to model signals as unbounded, so that incorrect cascades never occur, or bounded, so that either cascades or limit cascades occur. There is no way to empirically distinguish a signal distribution that includes values that are extremely rare and highly informative, versus one where such values do not exist at all. So for applications, either modeling approach is equally acceptable. For a similar perspective, see Gale (1996) and Chamley (2004*b*). In many applied contexts, it reasonable to model the signal space as finite. For example, people often obtain information signals through casual conversation with limited nuance.

until the last time the wrong action is taken) is infinite (Smith and Sørensen (2000)). Whether this expectation is finite or infinite depends upon the tail of the distribution of the private beliefs (Hann-Caruthers, Martynov and Tamuz (2018)). Rosenberg and Vieille (2019) focus on the expectation of the *first time* that a *correct* action is taken; this is (perhaps not obviously) very closely related to the time at which a correct herd starts. Their very elegant main result is that this is finite conditioned on a given state if and only if $\int \frac{1}{1-F(q)} \mathrm{d}q$ is finite, where F is the cumulative distribution function of the private belief in that state. Thus, when private beliefs have very thin tails on the "correct" end—i.e., very low probabilities of extremely informative correct-direction signals—the expected time can be long.

Even with continuous action spaces, which allow agents attune their actions finely to their unbounded private signals, owing to information externalities, the learning process may still be very slow (Vives (1993; 1997)). In these models, agents only observe a noisy signal about the average actions of predecessors, which relaxes the assumption that the actions of others is common knowledge. Similarly, Chamley (2004*a*), Acemoglu et al. (2011) and Dasaratha and He (2019) provide models with slow convergence. The review of Gale (1996) points out that very slow asymptotic learning, as occurs in several models, may be observationally indistinguishable from the learning stoppages that occur in settings with incorrect cascades.

It is also notable that that the *steadiness* of the rate of social learning is very different in different settings. In some learning as measured by increase in belief precision is steady and deterministic (e.g., Vives (1997)), whereas in others learning proceeds in stochastic fits and starts. In settings with unsteady rates of learning, there can be either temporary or permanent learning blockages.

4.4 Heterogeneous Precision and Influencers

Social psychologists report that people imitate the actions of experts. When a sports star uses a particular brand of equipment, it is an expert product endorsement likely to sway observers, as with Roger Federer's endorsement of Wilson tennis rackets and gear. Of course, sports gear often has a non-utilitarian "fashion" element. However, it is plausible that a tennis star knows how to choose an effective racket. Such endorsements may be less compelling for products that are unrelated to the endorser's primary domain of expertise. However, even outside this domain, observers may still view the decisions of an exceptionally successful individual as reflecting knowledge about what choices are most effective.

As the tennis example illustrates, some agents may have systematically more accurate signals than others. This raises the questions of whether small differences in accuracy can make big differences for social outcomes, whether increasing the accuracy of some agents necessarily improves social outcomes, and whether agents with higher accuracy should be placed earlier or later in the decision queue (Ottaviani and Sørensen (2001) study the latter question in a reputational learning setting).

Consider a variation of the SBM in which some agents, whom we call *influencers*, have more precise private signals than other agents. Depending on their location in the decision queue, influencers can trigger immediate information cascades. So even a small advantage in signal precision can have a large effect. For example, if we alter the SBM so that Ann has a signal precision $p' > p$, then even if p' is close to p, Bob will defer to her, as will all later agents. The resulting cascade has precision p'. In contrast, if $p' = p$, a cascade occurs only after two identical actions, which makes use of at least two signals, and therefore is more accurate.

The outcome depends crucially on the order of moves. If the high-precision agent, Ann, were second instead of first, there would be no immediate cascade. Placing an influencer later allows the actions of early agents to remain informative, which improves the accuracy of the cascade that ultimately ensues (see also § 10.1). These effects of influencers apply also in settings with with more general signal distributions, even though an influencer may not immediately trigger a cascade.

The drawback of leading off with the better-informed has not been lost on designers of judicial institutions. According to the Talmud, judges in the Sanhedrin (the ancient Hebrew high court) voted on cases in inverse order of seniority. Similar voting orders continue in some of today's courts (e.g., those in the U.S. Navy). Such strategic ordering can reduce the undue influence of older (and presumably wiser) judges.

4.5 Partial Cascades

As mentioned at the start of §3, a general feature of social learning is that a coarse action space reduces the informativeness of an agent's action, for both mechanical reasons and owing to information externalities. This is seen most starkly in the logic of information cascades: that with a coarse action space, an agent may choose an action independently of her private signal. In consequence, her action does not add to the pool of social information, and asymptotic learning fails. Furthermore, even in settings where there is asymptotic learning, the improvement of social information tends to be delayed.

The mechanical effect of having a coarse action space is present whenever the number of signal values exceeds the number of possible actions. More generally, information loss is exacerbated by information externality; self-interested agents have no incentive to convey information to later agents.[15]

To capture this more general point, Lee (1998) defines a weaker notion of information cascade. A *partial cascade* is a situation where an agent takes the same action for multiple signal values. (This terminology is due to Brunnermeier (2001);

[15] We mainly focus on coarseness in conjunction with information externalities. Coarseness is much less of a problem with altruistic agents, as they could choose their actions strategically to convey their private information to others (see Cheng, Hann-Caruthers and Tamuz (2018)).

Lee uses the term “cascades” for this concept.) An information cascade proper is the special case in which an agent takes one action for all the possible signal values. Partial cascades occur trivially when there are more possible signal values than actions. But even when the action space is rich, partial cascades (and cascades proper) can occur.

Lee applies this notion to a model of information blockage and stock market crashes (as discussed in § 10). In the model of Hirshleifer and Welch (2002), partial cascades result in either excessive maintenance of early actions (‘inertia”) or excessive action shifts (“impulsiveness”). Overall, there is a wide array of settings in which partial cascades hinder information aggregation.

5 Endogenous Timing of Actions

When faced with an irreversible decision, people often have an option to either act immediately, or to defer their decision to a time in which more information will be at their disposal. Examples include purchasing versus deferring the purchase of a product, or undertaking versus deferring an investment project. In social learning settings with the option to delay, there is the possibility of sudden booms and busts in action choices.

Consider a modification of the SBM in which, at any given date, any agent is free to choose among three options: adopt, reject, or delay. Adopting or rejecting is irrevocable. In this setting, there is no exogenous sequencing in the order of moves. The benefit to delay is that an agent can glean information by observing the actions of others. Thus, delay generates option value. The cost of delay could take the form of deferral of project benefits. Since acting early confers a positive information externality upon other agents, in equilibrium there can be excessive delay.

This equilibrium outcome when agents have an incentive to delay is seen most simply when agents have heterogeneous signal precision. We consider this case in the next subsection. We then consider a wider set of models.

5.1 Delay with Heterogeneous Signal Precision

Consider a setting as in the SBM except that there are agent-specific values of p (either commonly or privately known), so that agents differ in the precisions of their binary private signals. At a given point in time, an agent can act by choosing project H or L; or alternatively, can delay. As in the SBM, project H is optimal in state H, and L is optimal in state L. As discussed in BHW (p. 1002), high-precision agents have less to gain from waiting to see the actions of others—in the extreme, a perfectly-informed I_i ($p_i = 1$) has nothing to gain from waiting. So (without conclusively ruling out other possibilities) we focus on equilibria in which, among agents with h signals, those with higher precision adopt earlier than those with lower precision.[16]

Suppose first that each agent's signal accuracy is known to all. Once the agent with highest precision has acted (H or L), all remaining agents are in a cascade on the selected action (as in § 4.4). Intuitively, the agent with second-highest precision acts immediately rather than waiting to learn from others. Since this action is uninformative, so does the agent with next lower precision, and so forth.

Now suppose instead that precisions are only privately known. In equilibrium, agents can infer the signal accuracy of other agents from time elapsed with no action. In the continuous-time model of Zhang (1997), this results in an equilibrium in which delay fully reveals precision. Each agent has a critical maximum delay period, after which, if there are no actions by others, she becomes the first to act. The higher the precision, the shorter the agent's maximum delay period. All agents wait until the highest-precision agent acts. At that point, all other agents immediately act, choosing the same project. The incentive to free ride on the information of others by delaying is also present when there are continuous signals and actions, and when information

[16] To avoid technical issues in timing decisions, we suppose that agents have precisions drawn without replacement from a discrete distribution, such that no two agents have the same precision. This ensures that, if time periods are short enough, different agents act at different time periods in order of their precisions.

acquisition is endogenous (Aghamolla and Hashimoto (2020)).

In this model, cascades are *explosive* in the sense that there is an initial period during which all agents delay, and then, once the highest-precision agent acts, others immediately follow. Furthermore, since the cascade is based solely on one agent's signal, actions are highly idiosyncratic.

In the real options model of Grenadier (1999), the underlying asset value evolves as a geometric Brownian motion with an unknown drift. Investors differ in the accuracy of their private signals about the value of the drift. If two high-accuracy investors with positive signals invest, all other agents are in an information cascade, and also invest. Thus the broad insight from Zhang (1997) holds in Grenadier's setting as well.

The fact that less accurate agents have a greater benefit to delay suggests that when decisions are observable to others, there can be a strategic advantage to acquiring less information. Consistent with this, even if information acquisition is costless, agents may choose to remain imperfectly informed, as this encourages other agents to act earlier (Aghamolla (2018)). Moreover, even in nonstrategic settings, agents may acquire too little information, since there are externalities in information acquisition (§ 7).

5.2 Adopt versus Delay as an Indicator of Degree of Optimism

In the setting we have discussed, delay is informative about agents' precisions, but not about whether their signals favor project H or L. Suppose now that precisions are identical, that there is only a single project, and that the decision each period is whether to adopt it or to delay. As before, adopting the project is irrevocable, but delay is not. Then delay can be an indicator that the project is unattractive.

Symmetric pure strategy equilibria do not occur in such a setting, because when everyone moves early there is an incentive to delay and learn from their choices, and when everyone moves late there is an incentive for an agent with a favorable signal to invest earlier to advance the rewards from investing. Instead, there can be asymmetric pure strategy equilibria in which agents delay different amounts of time (Wang (2017)). Such an equilibrium endogenously generates sequential choice, resulting in an outcome that is similar to having an exogenous order of moves as in the SBM, and likewise gives rise to cascades. Alternatively, there can be mixed strategy equilibria in which agents randomize as to whether to delay. This can lead to idiosyncratic booms or voids in investment. Randomization in delay also occurs in the setting of Chamley and Gale (1994), in which a subset of agents randomly receives the option to invest—an option that can be exercised at any time. More agents receive the option to invest in better states, so receiving the option is a favorable indicator about state.[17] Furthermore, the incentive to wait to see whether others invest generates excessive deferral (see however Chamley (2004*a*)).

Overall, these models of timing decision reveal that there can be either inefficient delay, or a rush to invest even in unprofitable projects. This suggests that social learning may generate shifts in investment activity that are reminiscent of observed industry-wide booms and busts, or macroeconomic fluctuations.[18]

6 Observability Assumptions

What an agent learns from others, and the extent to which social learning is impaired by information externalities, depends on what the agent can observe about the past history of actions, payoffs, or even private signals. Observations of past actions can also be noisy

[17]The idea that information externalities result in stochastic delay is in the spirit of Hendricks and Kovenock (1989), who examine an experimentation setting in which two firms with private information decide how soon to drill for oil, where drilling causes the arrival of public information about the payoff outcome.

[18]Models of delay in which agents can observe the payoffs of predecessors (Caplin and Leahy (1993), Wagner (2018), Aghamolla and Guttman (2021)) have also been applied to such phenomena. Caplin and Leahy (1994) analyze booms and busts when firms can repeatedly adopt and terminate projects.

or limited, as with observation of only a specified subset of agents, a random sample, or a count of adoptions or other aggregate statistics. Observability can be asymmetric (as with greater observation of adoptions or of higher payoffs). Furthermore, agents can also be uncertain about whom their predecessors have observed.

Some key questions, under alternative observability regimes, are whether there is asymptotic learning, how quickly and how steadily information is aggregated, whether incorrect information cascades form, and how sensitive social outcomes are to the arrival of new public information. Another key question is whether greater observability increases welfare. We consider such questions while maintaining our convention of defaulting to SBM assumptions except where otherwise stated. We first consider rational settings, and then turn to imperfect rationality.

6.1 Rational Models

Many different observability assumptions have been studied in the literature. Rather than systematically covering all possible combinations, we organize the discussion in terms of general observations.

A herd at agent I_i, as defined in § 4, is the event that all later agents do the same thing as agent I_i.

Observation 1. Herds occur when there is sufficient observation of past social information.

As a benchmark, with no social observation, agents act based upon their own private signals, and there is no herd. In the other extreme, when all private signals are observable, there is eventually a herd on the correct action. A recurring theme of the literature is that enough social observation of actions and perhaps payoffs likewise produces a herd, but not necessarily a correct one. Mossel et al. (2020) give sufficient conditions for the formation of herds across a large class of social learning models.

There are many models with sufficient social observation to generate herding, including those of Banerjee (1992*b*), BHW, Smith and Sørensen (2000), Cao, Han and Hirshleifer (2011), models with costly information acquisition discussed in § 7, and models with social networks discussed in § 9. On the other hand, insufficient social observation precludes herding, as in Celen and Kariv (2004*a*) (which we discuss again in § 9), where each agent observes the action of the immediate predecessor only.

Observation 2. The ability to observe others' actions or payoffs tends to cause insufficient *exploration*.

A fundamental trade-off in an individual decision-making setting is between taking actions that generate new information that is relevant for an agent's future actions, versus taking the myopically best action based on current beliefs. This is known as the *explore/exploit trade-off*. This trade-off can also appear in social learning settings.

Consider now a setting as in the SBM except that previous payoffs are observable and are stochastic given the state. Observing these payoffs provides additional information about the state, with some actions generating more useful payoff information than other actions.

In consequence, there are two types of information externalities. The first, just as in the SBM, is that agents do not care that their actions convey information about their private signal to later agents. This externality affects the *aggregation* of private signals. The second type of information externality is in the *generation* of new information. Agents do not take into account the benefit to later agents of observing the payoffs derived from the chosen action. This is an externality in exploration (see, e.g., Rob (1991)). Owing to this externality, in a social multi-arm bandit settings with no private information, asymptotic learning fails (Bolton and Harris (1999)). An analysis with quasi-Bayesian agents is provided by Bala and Goyal (1998).

One could imagine that the additional information contained in the payoff observations would preclude information cascades and result in asymptotic learning.

However, this is not always the case, as incorrect herds can form on actions that generate limited payoff information.

To see this, consider a setting in which all past payoffs as well as actions are observable (Cao, Han and Hirshleifer (2011)). Suppose that the payoffs to action a are either 1 or -1, and to action b are either 2 or -2. There are four equally likely states: uu, ud, du and dd, where the first entry indicates a high (u) or low (d) payoff to action a and the second entry indicates a high or low payoff to action b.

Once an action is taken, its payoff is known to all, an assumption that is highly favorable to effective social learning. Nevertheless, a problem of inadequate experimentation remains. If private signals and payoff information about action a are initially favorable, whereas the prior beliefs about b are not very favorable, society can lock into a, for an expected payoff close to 1 under the belief that ud is likely, without ever trying b, whose payoff in state uu of 2 is even higher. So the ability to socially acquire both types of information (about private signals and about payoffs) does not solve the information externality problem. More generally, if payoffs are stochastic even conditional upon the state (or observation of payoffs is noisy), there is still a strictly positively probability that a given agent will cascade upon an incorrect action.

A possible interpretation of the payoff signal is that it is an online review posted by an agent who has adopted. In this application, Le, Subramanian and Berry (2016) show that the probability of an incorrect cascade can increase with the precision of the review. A more accurate review that is favorable to one option can prematurely deter useful exploration of the other option (see also Acemoglu et al. (2019)). When customer heterogeneity is sufficiently high, Ifrach et al. (2019) find that there is asymptotic learning, as there is always a chance of purchase strictly between zero and one, regardless of the social belief.

Observation 3. Social outcomes tend to be sensitive to the arrival of modest new public information.

Owing to information externalities, social learning is inefficiently slow, so on average, at any given date, the action taken by agent I_n is less accurately attuned to state than would be the case under learning that is efficient according to a standard social objective. If agents are rational, then they understand that on average they possess less information than under efficient learning, which makes them more willing to shift their actions in response to shocks. We call this tendency *sensitivity*.

Definition 2. *A social learning outcome is* sensitive *to shocks if a hypothetical one-time public disclosure of a signal with a distribution that is identical to that of the private signal possessed by a single agent is more likely to shift the agent's action than would be the case if learning from predecessors were efficient according to a social objective such as maximizing average utility.*

This is seen in starkest form in settings in which there are information cascades. The occurrence of even a small shock can easily dislodge a cascade. Each agent knows that a cascade is based upon information that is only slightly more accurate than the agent's own private signal. Thus, as emphasized by BHW, a key implication is that even long-standing cascades are fragile with respect to small shocks.

Definition 3. *A cascade is* fragile *if a hypothetical one-time public disclosure of a signal with a distribution that is identical to that of the private signal possessed by a single agent would, with positive probability, break the cascade, i.e., causes the next agent's action to depend on that agent's signal.*

For example, in the SBM, once a cascade starts, it remains fragile for all remaining agents.

Sensitivity and fragility are general concepts that could be defined in terms of different kinds of shocks to the system. For example, instead of the arrival of a public signal, the shock could be the arrival of a better-informed agent, or the of an agent whose preferences are known to differ from predecessors. In each of these cases, a long-standing cascade can easily be dislodged.

Owing to information cascades, there is a systematic, spontaneous tendency for the system to move to a position of precarious stability. In contrast, equilibrium is much more stable in models in which there are sanctions upon deviants or disutility from nonconformity (Kuran (1987)).

Observation 4. The Principle of Countervailing Adjustment: Seemingly favorable shifts in information availability can cause agents to take less informative actions, and do not necessarily improve average decisions or welfare.

The direct positive effect of greater availability of information to an agent tends to be opposed by the tendency of the next agent to take actions that are less responsive to their own private signals, to the detriment of later agents. We call this the *Principle of Countervailing Adjustment.* Indeed, the greater incentive of agents to mimic the actions of predecessors can ,make agents completely unresponsive to own-signals. An example is the presence of an agent with even slightly better private information (an "influencer"). As discussed in § 4.4, increasing the precision of one agent, I_1, can reduce average welfare by causing later agents to fall more readily into an incorrect cascade. Wu (2021) applies the principle to the formation of cascades in two decision queues.

One consequence of the principle of countervailing adjustment is that disclosure can reduce average welfare (see Result 2 of BHW). This contrasts with settings with no social interaction, where the free disposal property property implies that extra information never has negative value.

To see this, modify the influencer model discussed in § 4.4, so that all private signals have identical precisions, and instead have the slightly more accurate signal be a *public* disclosure made at date 0. Now agent I_1 is the first in the cascade, and all agents have lower expected utility than in the basic setting, as all now effectively act based upon just a single signal (the public disclosure). Of course, a sufficiently accurate early signal or public disclosure, such as conclusive information, can improve the social outcome.

More generally, a shift in information regime that might seem to make agents better informed (such as higher signal precision of early agents, greater observability of others, or greater publicly available information) can reduce average decision accuracy in the long run and can reduce average welfare. An example of this is the possible deleterious effect of increasing the precision of publicly posted information discussed by Le, Subramanian and Berry (2016) above.[19]

Observation 5. When each agent observes only a random sample of past actions, incorrect information cascades can occur, and may last forever.

To understand this observation, consider a sequential setting with random sampling of past actions and with no information about the order of past actions (see Smith and Sørensen (2020)). With bounded signals, information aggregation tends to be self-limiting, because as actions become more informative, an agent becomes more likely to cascade upon the preponderant action in the agent's observation sample. Whenever such cascading occurs, the agent's signal is not incorporated into the action history.

Suppose, for example, that past agents' actions were to become so accurate that even a single sampled H action would overwhelm the most extreme possible opposing private signal value. Then an observer will sometimes be in a cascade upon the predominant action in the agent's sample (e.g., in a sample of size k in which all observations in the sample are of the same action).

This reasoning suggests that owing to the possibility of information cascades, learning with random sampling may be quite slow. Indeed, a stronger claim is true: so long as all agents observe a sample size of at least 1, asymptotic learning fails (Smith and Sørensen (2020)).

[19] However, in general a shift in model structure can have different types of effects on the quality of information aggregation. For example, within the SBM, if each agent's private signal became more precise, cascades will tend to be more informative.

To see why, consider the case of a sample size of $N = 1$, where private signals are symmetric and binary. Consider a setting as in the SBM, except with the departure that each agent observes the action of one randomly selected predecessor instead of all predecessors. Suppose that a point is reached where for some agent I_n this random observation is more informative than a single private signal. (If this never happens, of course asymptotic learning fails.) Then agent I_n would be in an information cascade, so I_n's action would be exactly as informative as a sample of one action from among I_n's predecessors. In consequence, the sample observed by I_{n+1} is also more informative than I_{n+1}'s signal, so I_{n+1} would also be in a cascade. A similar argument holds for all later agents, so information stops accumulating. Consequently, asymptotic learning does not occur. This failure is similar to the fashion leader version of the SBM in §4.4. A similar intuition also applies to the random sampling model of Banerjee and Fudenberg (2004), in which it is possible that past payoffs as well as actions are observed.[20]

A further interesting implication of the sequential sampling setting of (Smith and Sørensen (2020)) is that assuming each agent does not observe too many predecessors, once different agents take different actions, herds never occur. This is because there is always a chance that an agent observes a set of predecessors who did not follow the currently-predominant action (Smith and Sørensen (2020)). The chance of observing such a deviant action does not decline rapidly enough with n to bring about herding. Similarly, if payoff information is also observed and if an early agent adopted a popular action and experienced low payoffs from doing so, there is always a chance that this agent is later observed, causing a later agent to deviate from the popular action.

If agents observe samples of payoff outcomes but not the past actions that led to these outcomes, it is again possible that agents do not converge to the same action. The need to simultaneously draw inferences about what actions predecessors have taken and the performance of those actions can confound inferences. Wolitzky (2018) considers a

[20] Monzón and Rapp (2014) consider a sampling setting in which agents also do not know their own positions in the decision queue. In this setting, under a stationarity assumption on sampling rules, incorrect cascades can last forever.

setting with two actions: action R(isky) has a probability of success that depends on the state, and action S(afe) has a fixed probability of success that is state independent. If action R always has lower probability of success (but has lower cost) than action S, then even for very large samples, there is not asymptotic learning.

Observation 6. Agents can incorrectly herd on rejection even when agents observe the aggregate number of adopts, but do not observe rejects.

Consider a scenario in which each agent observes only the aggregate count of the number of past adopts. In this setting, sequencing information is lost. In general this induces loss of two types of information. First, an agent does not know *the order* in which past actions were taken. Second, an agent does not know *how many* predecessors have acted—i.e., agents do not know their own positions in the queue. This occurs when an agent does not observe all past actions, one example being when an agent observes only adopts, not rejects; it is not hard to obtain information about how many Teslas have been sold, but we do not see how many people considered Teslas but opted not to buy.

In the setting of Guarino, Harmgart and Huck (2011), observation is asymmetric (we will refer to this as only observing past adopts, not rejects), there is a finite number of agents, agents cannot see the order of predecessors' actions, and they have no direct information about their own positions in the queue (though they can draw inferences about this from their observations of others). Since all that an agent observes is how many predecessors adopted, there is no way for a cascade on reject to get started. (If it could, then even I_1 would reject, since I_1 does not see any past adopts nor does I_1 know that no agent preceded her. In consequence, all agents would always reject, which is not consistent with equilibrium.) So the possibility of cascading is limited to just one action, and indeed, with a large (finite) population, such a cascade occurs with a probability that approaches one regardless of state.

In sharp contrast, when agents do have some idea about their own positions in the queue (based, for example, on observation of own-arrival-time), cascades on either action can occur. This is because an agent who is probably late in the queue and who observes few adopts infers that others probably arrived earlier and chose to reject (Herrera and Horner (2013)). Since inferences are noisy, cascade outcomes can be incorrect, resulting in incorrect herding on either adopt or reject.

Observation 7. Contrarian actions can reveal that an agent has high precision.

Consider a setting like the influencer model of § 4.4, in which each agents has a high or low precision binary signal, and this precision is private information. Then the decision of an agent to deviate from a cascade indicates that the agent has high precision. This can potentially cause subsequent agents to follow contrarians.

In this setting an observer knows that the minority choice was made in opposition to predecessors, which is indicative of strong private information. What is more surprising, as shown by Callander and Hörner (2009), is that agents who only observe a count of past adopts and rejects sometimes act in opposition to the majority of the actions they observe: Intuitively, the very presence of a minority action can be an indicator of a late, high precision deviator, as early adoption of an action tends to cause it to not be in the minority.

Observation 8. Even with partial observation of past signals, asymptotic learning can fail and incorrect cascades can occur.

The additional social information that agents obtain may be about the private signals of predecessors through conversation. Indeed, it is sometimes argued that owing to such communication, in the long run agents will learn the state, and that even in the short run that incorrect cascades will never occur. However, in practice people often do not pass on the full set of reasons for their actions—especially reasons that they have acquired from others. In settings with limited communication of past private signals, incorrect cascades may form and, with positive probability, last forever. We provide an example of this that extends the SBM in the Online Appendix, § A.1.3. Intuitively, being able to observe a predecessor's signal is much like having an extra endowed private signal, which in turn is much like having a more precise endowed signal. But (as in the SBM), increasing the precision of private signals does not eliminate incorrect cascades.

Observation 9. Reduced social network connectivity can improve social learning

BHW and Banerjee (1992*b*) point out in information cascades settings that there is a benefit to quarantining early agents so that some make decisions without observing others. We call such agents *sacrificial lambs*. The advantage of having nonsocial agents in social learning settings is that their actions are sensitive to their own signals. This makes their actions more informative to later observers. More generally, weaker social network connectivity can result in better social learning outcomes (see, e.g., the references in Wu, O'Connor and Smaldino (2023)).

A similar effect can derive from psychological biases that cause agents to use their own signals more heavily instead of imitating others. By bringing more private information to bear, such bias can improve learning and welfare.

6.2 Models with Imperfect Rationality

In models of individual decision making, irrationality makes an agent worse off. In contrast, in a social learning setting, irrationality can make most agents better off, because mistaken actions are sometimes more informative to later observers than correct ones. So psychological bias can help remedy information externalities, resulting in more accurate social beliefs. This is an example of the general phenomenon that irrationality can make interacting agents better off (Kreps et al. (1982)).

6.2.1 Overconfidence

Even if an agent is not a sacrificial lamb, the agent may place greater weight upon her own private information because she is overconfident about its precision. Such overconfident agents can break incorrect cascades, improving long-run learning. In

Bernardo and Welch (2001), occasional overconfident "entrepreneurs" overestimate the precisions of their own signals. This can cause them to make use of their own signals instead of following the actions of predecessors in an information cascade. So overconfidence can improve learning and outcomes for later agents. Arieli et al. (2023) similarly show that "condescending" agents who underestimate the precision of their peers' signals can improve aggregate outcomes.

To illustrate overconfidence, suppose, for example, that everyone's private signal has the same precision, known to all, except that I_1 to I_{10} each substantially overestimates their own precision. Then each acts based solely upon the agent's own signal, so the first 10 signals are revealed through their actions. In consequence the expected welfare of all later agents is improved.

A possible direction for future research is understanding the conditions under which overconfidence harms social learning instead of helping. For example, if overconfidence were growing rapidly with later agents, all agents would act based only on their private signals, and there would be no information aggregation.

6.2.2 Neglect of Social Observation by Predecessors

Other psychological biases can also influence social learning. An important one is that agents may neglect the fact that others are making social observations. As with overconfidence, such neglect can cause an agent to underweigh predecessors and rely excessively on her own signal instead of cascading, thereby improving welfare.

Alternatively, an agent may neglect the fact that the actions of two predecessors derive from a common source. This induces *correlation neglect* (also known as *persuasion bias*), the phenomenon that people sometimes treat the information of others as independent even when it is not—a type of double-counting (Enke and Zimmermann (2019)). These forms of neglect follow naturally from limited attention and cognitive processing power, since inferences about observation of others who in turn observe others can require extensive computation.

However, neglect of observation by predecessors can also make agents more prone to cascading. To see why, consider a setting in which an agent observes multiple predecessors who chose the same action, and mistakenly believes that these predecessors acted independently. Some of these predecessors may be in a cascade, making their actions uninformative. The mistaken belief that these actions are all informative can cause the agent to imitate predecessors and cascade instead of following the agent's own signal. This can happen even if the agent has higher private signal precision than do predecessors, so that if the agent were rational the agent would not be in a cascade.

In a social network, heavily connected agents provide correlated information to many observers, who are in turn observed by others. In consequence, correlation neglect increases the influence of agents who are more heavily connected in the social network (see DeMarzo, Vayanos and Zwiebel (2003) and the review of Golub and Sadler (2016)).

Neglect of predecessor's observation of others is captured in a sequential quasi-Bayesian setting in Bohren (2016). In Bohren's model, states are equiprobable, signals are bounded, and there is a given probability that each agent is *social* (observes the actions of predecessors) or *nonsocial* (does not observe any predecessor actions). Whether an agent is social is unknown to others. Agents may either underestimate this probability (resulting in correlation neglect) or overestimate it (which could be called correlation overestimation).

In the case of rational agents, this is a model of fragility in cascades and social learning. When enough agents take the same action, a cascade forms, in the sense that a social agent follows the social information implicit in the actions of predecessors. However, when a nonsocial agent arrives, this agent provides an information shock by taking an independent action. This is informative to later social agents, so that initial cascading may be broken. In the rational case, eventually social agents make correct decisions, since there is an infinite stream of nonsocial sacrificial lambs.

To understand outcomes in imperfectly rational cases, let q be the probability that any agent observes the actions of predecessors, and let $\hat{q}$ be agents' perception of that probability. If the possible values of $\hat{q}$ are in an intermediate interval (an interval which includes q), then Bohren shows that, just as in the rational case, with probability one the social agents eventually make correct choices.

When $\hat{q}$ is below this intermediate interval, agents severely underestimate the probability of agents being social, and therefore overestimate how informative are these actions. When, by chance, a strong enough preponderance of agents favors one of the available actions, agents fall into a cascade. So the preponderance of one action tends to grow over time. Since agents think this growing preponderance is coming largely from independent private signals, social agents grow increasingly confident in the accuracy of the latest cascade. In the limit agents become sure of either the wrong state or the correct one.

This possibility of strongly held faith in the wrong state provides an interesting contrast with the information cascades setting of BHW, in which there is also a failure of asymptotic learning, but agents have low confidence about the state, making cascades fragile. It also contrasts with a rational benchmark with continual arrival of nonsocial agents, in which asymptotically the beliefs of the social agents become arbitrarily strong, but always converge to the correct state.

When $\hat{q}$ is large (i.e., to the right of the above-mentioned interval), beliefs fluctuate forever, so again there is not asymptotic learning. Even if, at some date, there were a very strong preponderance of action H, for example, agents would believe that this derives almost entirely from cascading by predecessors. This makes the system fragile. When by chance (as must eventually happen) even a few nonsocial agents take the opposite action, the next social agent will no longer be in a cascade, and will sometimes choose L.

In the social learning model of Eyster and Rabin (2010), correlation neglect takes a more extreme form—observers think that each predecessor decided independently based *only* upon that agent's private information signal.[21] In their model, state and actions are continuous. Beliefs about others are analogous to $q = 1$ and $\hat{q} = 0$ in the Bohren (2016) model. In consequence, the views of early agents are very heavily overweighted by late agents, convergence to the correct belief is blocked (even with sharing of continuous beliefs or actions), and agents become highly confident about their mistaken beliefs.

A lesson that comes from analyses of imperfect rationality and social learning is that biases that cause agents to be more aggressive in using their own signals, such as overconfidence, or such as overestimation of how heavily others have observed their predecessors, tend to promote the use of private signals. Under appropriate conditions, this improves social information aggregation. In contrast, correlation neglect tends to have an opposite effect, causing agents to defer too much to history.

6.2.3 Other Heuristics and Psychological Biases

So far in this section we have discussed models that explicitly analyze the effects of psychological biases such as correlation neglect and overconfidence on social learning. Such models fully endogenize beliefs and behaviors. Another approach is to make exogenous assumptions about the agent's mapping from observed actions and payoffs into the agent's actions. Ellison and Fudenberg (1993, 1995) provide pioneering analyses using this heuristic agent approach (see the Online Appendix, § A.1 for details).

Bohren and Hauser (2019) examine a setting that allows for a variety of types of possible psychological biases in social learning, including correlation neglect. In this model, signals are continuous (and may be unbounded). They focus on settings in which enough information arrives so that if agents were rational there would be asymptotic learning (via the arrival of either public signals or nonsocial agents). However, owing to psychological bias, asymptotic learning can fail, which can take the form of

[21] Hirshleifer and Teoh (2003) and Eyster, Rabin and Vayanos (2018) apply such neglect of the signal-dependent behavior of others to financial markets.

convergence to a mistaken action, permanent disagreement over action, or infinite cycling. For example, when agents overreact to private signals, and where there is a positive probability of nonsocial types, there can be infinite cycling between actions. When agents underreact, there can be fixation upon a mistaken action. Furthermore, when there are incorrect herds, beliefs converge almost surely to the incorrect state. So, consistent with Bohren (2016), and in contrast with BHW, longer herds become increasingly stable.

7 Costly information acquisition

People often have a choice of whether or not to acquire information. We next examine the effects on social learning of costly acquisition of either direct private information signals about the state (§ 7.1) or about predecessors' actions (§ 7.2).

When agents can acquire private signals, it is unprofitable to do so if the signal will not (or is unlikely to) affect the agent's action. So social learning is self-limiting. Once social information becomes sufficiently informative, it crowds out costly information acquisition about the state. In particular, even when private signals are unbounded, social learning can become completely blocked.

This learning blockage effect has a close parallel with settings with exogenous private signals and long-continuing information cascades, such as the SBM. In such cascades settings, the signals of late agents do not contribute to social knowledge, because once a cascade forms, such signals do not affect actions. This results in a similar conclusion, that late agents do not contribute to social knowledge.

If private signals are costless, then asymptotic learning occurs when private signals are unbounded (as noted in § 4), and may occur when the action space is continuous (as described in § 3). Relative to these scenarios, a positive cost of observing private signals degrades learning. As we will discuss, even in settings with unbounded private signals or continuous action spaces, asymptotic learning does not occur if there is even a small cost of investigating.

On the other hand, if private signals are costless, introducing a cost of *observing predecessor's actions* can *improve* social learning. In such a setting, an agent with a very informative signal realization may choose not to observe others' actions. Thus, her action conveys greater incremental information, which benefits later agents.

7.1 Costly Acquisition of Direct Private Signals about State

Costs of acquiring private information introduce another information externality. In deciding whether to buy a signal, agents do not take into account the indirect benefit that accurate decisions confer upon later observers.

The effects of this externality are illustrated simply in a setting as in the SBM of § 2 with the addition of a choice of whether to acquire a private signal about the state. Each agent can pay some cost $c > 0$ and observe a binary signal with given precision p, or can pay nothing and observe no signal. In this setting, agents I_n, $n > 1$ will not acquire a signal, no matter how small the cost. To see this, observe that in this binary setting, even if I_2 were to observe a private signal, following I_1 would still be weakly optimal. Agent I_3 understands this, so by the same reasoning, I_3 does not acquire the signal; nor does any subsequent agent. So the social outcome incorporates even less information than in the SBM.

In more general settings as well, when there is at least some minimal cost of acquiring a private signal about the state, agents stop acquiring private signals, resulting in complete blockage in the growth of social information. So only a few individuals end up acquiring private signals. This is a version of the "Law of the Few" (see Galeotti and Goyal (2010)). For example, in a model with a continuous state and continuous actions studied by Burguet and Vives (2000), asymptotic learning is obtained if and only if the marginal cost of obtaining a signal tends to zero as signals become arbitrarily weak.

A similar outcome applies in Mueller-Frank and Pai (2016) in a setting in which agents acquire information about finite samples of predecessors' actions and payoffs. Agents differ in the realizations of their costs of sampling, which is private information, and do not see what samples were drawn by predecessors. Asymptotic learning occurs if and only if sampling costs are not bounded away from zero in the sense that costs can be arbitrarily close to zero for an unlimited number of agents.

As discussed in § 3, Ali (2018*a*) introduces a notion of responsiveness which, loosely speaking, requires that any change in an agent's beliefs changes the optimal action. Ali shows that even with responsiveness, there may not be asymptotic learning if information is costly to acquire, since the benefit of greater accuracy may not outweigh the cost of information. Responsiveness implies asymptotic learning if and only if the minimum across agents of the costs of gathering information is arbitrarily close to zero. The consistent message from these papers is that asymptotic learning is not robust to introducing non-negligible costs of acquiring (or processing) private information.

Alternatively, instead of costs of acquiring private information, there could be costs of acquiring private information signals by discussion with predecessors. Consider a modification to the SBM such that for a small fixed cost an agent can talk to a predecessor to find out the rationale behind her action choice. In other words, the agent can learn the predecessor's belief (which may reflect information that she has acquired in conversations with her predecessors). Nevertheless, as long as there is even a small cost of such conversations, incorrect cascades occur with positive probability. Intuitively, as beliefs become increasingly informative, at some point it pays for an agent to simply follow the action of the agent's immediate predecessor rather than paying to learn the predecessor's belief.

7.2 Costly or Noisy Observation of Past Actions

It is often costly to observe the actions of others. For instance, in evaluating the decision to invest in a startup firm, a venture capitalist can devote time and effort to gathering information about the decisions of earlier potential investors.

To illustrate the effects of costly observation, consider a setting as in the SBM except that there is a fixed cost for each agent of observing all predecessors' actions. If the cost is high enough, early agents will not incur it, and therefore will act solely on the basis of their own private signals. However, for this very reason, at some point the action history may become so informative that an agent finds it worthwhile to learn the previous choices. Once this point is reached, all subsequent agents will also find observation worthwhile. So observation costs can turn early agents into sacrificial lambs, as defined in § 6.2, to the benefit of many later agents.

Based on this, from the viewpoint of improving the accuracy of decisions (and perhaps welfare as well), typically the observation cost should be positive but not too large. With a zero cost, as in the SBM, actions tend to be very inaccurate. With too high a cost of observing predecessors, no agent will ever incur it. This insight is developed in the model of Song (2016), who derives conditions on signals and observation costs for asymptotic learning to occur.

7.3 Costly Information Acquisition, Limited Observation and Groupthink

Can social observation lead to decisions that are even worse than the decisions that agents would make under informational autarky? This might seem impossible, since any information gleaned by an agent via social observation is incremental to her own private information. However, psychologists have emphasized (Janis and Mann (1977)) that "groupthink" in group deliberations causes disastrous decision failures, as if interaction with others were harming instead of improving decisions. There is also evidence suggesting that observation of others sometimes result in degradation in decision quality (a zoological example is provided by Gibson, Bradbury and Vehrencamp (1991)).

Analytically, when there are investigation costs and noisy observation of past actions, agents in groups can come to decisions that are on average worse than if there were no social observation. The mechanism behind this is free-riding, which causes agents to underinvest in information.

To see this, first consider the SBM, in which others are observed without noise, with the modification that there is a small cost of acquiring private signals. As discussed at the start of § 7.1, starting with I_2 all agents follow I_1, so the social belief reflects only a single signal. This is no more accurate than if agents decided independently (though welfare is higher as agents save on investigation costs).

Suppose instead that observation of predecessors is noisy, where each agent observes binary signals about the actions of all predecessors. Suppose further that all agents observe the same binary noisy signal about the action of any given predecessor.

If the noise is sufficiently small relative to the cost of the signal, the net gain to I_2 of investigating is still negative, so she still does not investigate. But now, owing to observation noise, her action is less accurate than if she were to decide on her own. So observation of others reduces decision quality relative to informational autarky. (Nevertheless, I_2's welfare is higher than under autarky, as observation of others economizes on observation costs.)

What about later agents? Agent I_3 also just follows I_3's signal about I_1's action.[22] The same applies to all later agents, so everyone's action is less accurate than if they had decided independently.

[22]Agent I_3 ignores her signal about I_2's action, because she knows that I_2 imitated I_1 based on the same signal realization about I_1's action that I_3 observes. So if I_3's signal about I_2's action differs from I_3's signal about I_1's action, I_3 knows that this discrepancy *must* be caused by error in observation of I_2's action.

Decisions in groups can become even less accurate if agents observe only the latest predecessor. In this case noise can compound repeatedly. However, a point is eventually reached at which an agent again pays to acquire a private signal (Cao and Hirshleifer (1997)), so there are cyclical shifts in accuracy.

8 Social Learning in Markets

In a market for a product or financial asset of uncertain value, the decision to buy depends on the price, the agent's (buyer's) private information signal, and the decisions of predecessors that the agent has observed. This raises several questions. Does the price setting process promote or prevent cascades, including incorrect ones? How should a seller manage the social learning process? How does social learning affect market efficiency? What are the welfare consequences of social learning and cascades? We first discuss the case of monopoly pricing, in which the seller chooses prices to maximize expected profits, and then turn to competitive price-setting.

8.1 Monopoly

A monopolist may have an incentive to set price low enough to induce a cascade of buying. The dynamics of prices and buying depend on whether the monopolist must commit to a single price or can adjust prices in response to observation of early purchase decisions. We first discuss the fixed price case, which can also apply to products with menu costs (costs of changing prices; Sheshinski and Weiss (1977)). It also applies to the sale of equity shares of a firm in an Initial Public Offering (IPO), since a fixed price per share is mandated by U.S. law. Much of this literature focuses on the case of a seller who has no private information about the state. In § 8.2 we consider informed traders in competitive markets.

8.1.1 Fixed Price Case

Consider a setting that modifies the SBM so that, as in Welch (1992), an uninformed risk-neutral monopolist who offers to sell one unit of a product to each agent in a sequence at a fixed price until the monopolist's supply of the product, n units, is exhausted. The monopolist's cost of production is zero. As in the SBM, each agent receives a binary private signal about the state $\theta \in \{0, 1\}$, which is the unknown value of the product, and can observe the choices of all predecessors. The net gain to adopting (buying the product) is the difference between the state and the price. This contrasts with the SBM, in which the net value of adopting is exogenous. The monopolist is risk neutral and does not discount the future. We assume that when indifferent the customer buys.

If the price is sufficiently low, the seller induces a buying cascade, selling the entire inventory, but at a low price. If the price is sufficiently high, a no-buying cascade ensues. In the intermediate range, the initial agent's choice depends on her private signal, in which case her choice provides information about the state to later agents. Eventually, a cascade will ensue. This cascade can be either a buying or a non-buying cascade.

From the monopolist's perspective, demand is fragile. Just a few early agents with negative signals would cause buying to collapse. In a binary setting, for a sufficiently low signal precision, the profit-maximizing price is set just low enough that all agents always buy.

Intuitively, with low precision, even a price just slightly below the unconditional expected value of the product is enough to induce even an agent with an adverse signal to buy. It is therefore not worth it to the monopolist to risk the collapse of demand for a slightly higher price. This implication is consistent with the empirical finding of underpricing in IPO markets (Ritter and Welch (2002)). However, for higher signal precision, such underpricing does not occur, since selling to just a few buyers at a high price is worth more than selling to all at a low price.

In a setting with a uniform prior on θ, Welch (1992) shows that when the seller is also informed, a high-signal seller sets a higher price (with higher failure probability) to separate from a low-signal seller. Welch (1992) focused on cascades in the IPO market. Empirically, Amihud, Hauser and Kirsh (2003) find that IPO opportunities for investors tend to be either heavily oversubscribed or undersubscribed, with almost no IPOs in between. This is consistent with information cascades, in which there is positive feedback from early investor decisions to later ones.

8.1.2 Flexible Price Case

A monopolist may not be constrained to a fixed price. In Bose et al. (2008), the seller is risk neutral, uninformed, and can modify the price after observing each buyer's decision.[23] When buyers have binary signals, the relevant prices to consider offering to I_n are a low price that leads to a certain sale and a higher price that results in a sale only for a high signal. For each buyer, the values of the low and high prices depend on the actions of preceding buyers.

Bose et al. show that the seller starts with a high price that induces the first buyer to reveal her private signal. Once enough information is revealed, the value of further information to the seller is low. Eventually, the seller fixes a low price that induces a buying cascade. As the seller's discount factor increases, the value discovery phase becomes longer, and more information is revealed. In the limit, if the seller has no time discounting, there is complete value discovery, and the seller earns $\mathbb{E}[\theta]$ per buyer. This is the best conceivable asymptotic outcome for the seller, as rational buyers will never, on average, pay more than their ex ante expected valuation.

[23] Caminal and Vives (1996) and Caminal and Vives (1999) study social learning about product quality via market share in a duopoly.

8.2 Competitive Markets

In contrast with the monopolistic case, cascades do not occur under competitive price setting when agents have common values. In securities markets, when an agent buys or sells based on her private information, market prices should change to reflect at least some of the agent's private information. This makes it less attractive for an observer to imitate the trade, which opposes the formation of a buying or selling cascade.

To see this, consider a setting in which each trader receives a signal about an object with common value $\theta = 0$ or 1. The trader can buy one unit at the market makers' ask price A, sell one unit to the market makers at bid price $B \leq A$, or not trade. If market makers are uninformed, each must set A and B such that no trade occurs, as otherwise they would lose money in expectation. This follows from standard no-trade results for securities markets with no noise traders (Glosten and Milgrom (1985)). Thus, trivially, there is no trading cascade.

However, in settings with noise traders (Glosten and Milgrom (1985)), market makers profit at the noise traders' expense, and bid-ask spreads are set to accommodate trading. The adjustment of the market price to reflect private information discourages the occurrence of buy or sell cascades by making it optimal for traders to use their private information (see Avery and Zemsky (1998)).

The reason for this is that cascading would cause market makers on average to lose money, inconsistent with the proposed equilibrium. Specifically, a competitive market maker must set bid and ask prices equal to the expected values of the asset conditional on an observed buy order or sell order, respectively. In a buying information cascade (for example), the buy of an informed trader is uninformative. A noise trader's action is always uninformative, and a sell order can come only from a noise trader. So the competitive market maker must set the bid equal to the ask. The market maker breaks even on sells to noise traders, but on average loses money to the informed.[24]

Avery and Zemsky define a concept in the spirit of cascades which they call "herds" that is adapted to financial markets. To distinguish from our usage of "herd," we refer to this concept as a momentum herd. In a momentum herd, agents do not ignore their private signals, but make less use of their private signals than they would have in the absence of social observation. Specifically, an informed trader is in a *momentum herd* if the trader's optimal action is contrary to the optimal action the trader would have taken had she moved first, i.e., if she had the same private signal realization, no social information, and faced the initial bid and ask prices.

Momentum herds require going beyond a setting with binary signals and states. Avery and Zemsky present an example with three states in which momentum herds are possible. In a more general setting, Park and Sabourian (2011) provide necessary and sufficient conditions for momentum herds.

Cipriani and Guarino (2014) provide and test a model of momentum herds using stock market data. They modify the Avery-Zemsky model by dividing time into days, where each day consists of a finite number of trading periods. There is an asset whose fundamental value remains fixed during the trading day, and receives an independent shock at the end of each day. Agents are exogenously ordered and act once. They observe the history of prices and actions, and in addition each receives a private signal regarding the value on the day at which they trade. Cipriani and Guarino estimate the prevalence of momentum herds in their setting. They calibrate their model to NYSE stock market data and estimate that inefficiencies deriving from incorrect momentum herds constitute 4% of asset value.[25]

If there are transactions costs, then in contrast with Avery and Zemsky's finding,

[24]For the informed with an ℓ signal to be willing to buy, the conditional expected profit must be zero, which implies that the informed with an h signal on average profits at the expense of the market maker.

[25]There is also laboratory evidence that financial professionals from the Chicago Board of Trade tend to follow the incorrect actions of others instead of their own signals (Alevy, Haigh and List (2007)).

there can be cascades of no-trade (see Romano (2007) and Cipriani and Guarino (2008*a*)). Information asymmetry about the asset value decreases as successive traders buy or sell. Ultimately, a no-trade cascade starts when the value of an informed trader's private information does not justify the expected transaction cost to be incurred.

Lee (1998) considers a model in which each trader incurs a one-time transaction cost (perhaps cognitive) that enables her to trade repeatedly based upon a single private signal. Information blockage can occur, and may be temporary. Each agent I_n enters in period n, and after entering, can buy or sell any amount of a risky asset. Owing to the transaction cost, private information can be sidelined during several periods with no trading. This quiescent interval is shattered if a later agent trades upon observing a sufficiently extreme signal. Since multiple signal values can result in the same action the equilibrium at a given date is an example of a partial cascade as discussed in § 4. When agents suddenly start to trade, there is a sudden drop or jump in price. Lee calls this phenomenon an information avalanche.

In contrast with settings with pure common values, when agents have a private component to their valuations of assets, there can be information cascades of buying or selling. Even though prices adjust to reflect previous trades, there are cascades wherein, regardless of private signal realizations, informed traders with sufficiently low private values sell and those with sufficiently high private values buy. Private values are common for illiquid assets, such as real estate and private equity. Even for liquid assets, owing to risk aversion, an agent that is endowed with substantial holdings in a firm places less marginal value on a share than an agent with no holdings. Furthermore, even under risk-neutrality, investors who value control rights can place different values on the shares of a firm.

There are several models in which private values induce information cascades in asset market trading. In Cipriani and Guarino (2008*b*), information cascades, both incorrect and correct, may occur, with no asymptotic learning. A similar result occurs in Décamps and Lovo (2006), who consider trader heterogeneity in the value of an asset deriving from differences in risk aversion and initial endowments.[26]

9 Learning in Social Networks

Social networks—from word of mouth networks in iron-age villages to modern online social media—play an important role in the spread of information in human society. In general, what an agent learns by observing others depends on the agent's position in the social observation network. And the overall structure of the network affects aggregate social outcomes, such as whether there is asymptotic learning. So network structure provides empirical implications for behavior and outcomes. Accordingly, a large recent literature models social learning in networks.

Networks increase the complexity of the inferences that agents need to make. When agents do not observe the whole action history, they potentially learn about the actions of those that they do not observe from the actions of those that they do observe. Thus, calculating expected utilities may require taking into account the structure of the entire social network (Mossel, Sly and Tamuz (2015)). As this places heavy computational demands on agents, the rationality assumption becomes less plausible.

So, for tractability, network economists often make strong assumptions about the geometry of the network, and for both tractability and realism, often focus on non-Bayesian agents. Nevertheless, models with rational agents provide valuable benchmarks for evaluating the effects of different heuristic behaviors or beliefs.

A key question is how the geometry of the social network affects learning outcomes. As we will discuss, a general lesson from both rational and boundedly-rational models is that egalitarian networks—loosely speaking, networks in which no agent is much more important than others in the geometry of the network—tend to

[26] Also Chari and Kehoe (2004) find in a common value setting that if agents own investment projects and have a choice as to when to trade them, information cascades can occur.

facilitate social learning (Bala and Goyal (1998), Golub and Jackson (2010), Acemoglu et al. (2011), and Mossel, Sly and Tamuz (2015)).

9.1 Sequential Actions

In this section, we consider models of rational social learning with sequential actions on networks. Banerjee (1992*b*) and BHW assume that every agent observes the actions of all predecessors in the queue. In this simple network structure, each agent I_n observes the actions of all of her predecessors I_m, where $m < n$.

A subsequent literature retains the exogenous ordering of actions, but relaxes the complete observation structure, so that each agent observes only a subset of her predecessors. We discussed models with this feature in § 6.

To further consider such settings, let I_n's *neighborhood*, N_n, be the set of agents whose actions I_n observes before acting. Çelen and Kariv (2004*b*) study a model in which $N_n = \{I_{n-1}\}$: each agent observes only her immediate predecessor. In their model, the state is equal to the sum of the agents' private signals. With this state and network structure, neither herding nor information cascades occurs, but the probability that later agents mimic their immediate predecessors tends to one.

Acemoglu et al. (2011) introduce a general network structure: the neighborhood N_n of agent I_n can be any subset of $\{I_1, \ldots, I_{n-1}\}$, and, moreover, can be chosen at random, exogenously and independently. They study asymptotic learning (as defined in § 4.2) and ask the following question: Under what conditions does I_n take the correct action with a probability that tends to 1 as n becomes large? As BHW show, asymptotic learning does not occur when agents observe all of their predecessors and there are finitely many signals. Nevertheless, Acemoglu et al. (2011) show that asymptotic learning is possible even with bounded signals in some networks with incomplete observation structures.

For example, the presence of sacrificial lambs as defined in § 7—agents who are unable to observe others—can induce asymptotic learning. To see this, suppose that $N_n = \{\}$ with probability $1/n$, and that $N_n = \{I_1, \ldots, I_{n-1}\}$ with the remaining probability $(n-1)/n$. Sacrificial lambs act according to their private signals only. The sacrificial lambs choose the wrong action with a constant probability that does not tend to zero with n. But these mistaken actions become exceedingly rare, as the frequency of sacrificial lambs tends to zero. The actions of the sacrificial lambs reveal independent information items to their successors. Because the probability of arrival of a lamb decays slowly enough, there are infinitely many sacrificial lambs, and the rest eventually choose correctly with probability one. Acemoglu et al. (2011) provide a more general condition that ensures asymptotic learning. As in the sacrificial lambs example, this condition applies to stochastic networks only. Indeed, deterministically placed sacrificial lambs cannot produce asymptotic learning, since asymptotic learning requires that *all* late players choose correctly with high probability.

Acemoglu et al. also show, conversely, that asymptotic learning cannot hold when some agents play too important a role in the geometry of the network. This happens when there is a set of important agents $\{I_1, \ldots, I_M\}$ that are the sole source of social information for an infinite group of agents. When the signals of the important agents happen to indicate the wrong action—which occurs with some positive probability—infinitely many agents follow suit. Along these lines, Acemoglu et al. say that a network has *non-expanding observations* if there is some M and $\varepsilon > 0$ such that, for infinitely many agents I_n, the probability that N_n is contained in $\{I_1, \ldots, I_M\}$ is at least ε. In this case, asymptotic learning does not occur.

Turning to short-term dynamics, Acemoglu et al. (2011) show that social learning can sometimes induce beliefs that are contrarian with respect to a subset of predecessors, and as a result, anti-imitation (see also Eyster and Rabin (2014) discussed in § 6.2). Intuitively, suppose that both I_3 and I_2 observe I_1 only, and that I_4 observes I_1, I_2 and I_3. Then I_4 should place positive weight on I_3 and I_2, and negative weight on I_1 to offset double-counting. So, within a broad stream of imitation, there can be eddies of contrarian behavior.[27]

Until now we have considered network models in which the neighborhoods N_n are either deterministic, or drawn independently. Less is known about the case in which neighborhoods N_n are not independent (as analyzed by Lobel and Sadler (2015)). For example, a given agent may have a chance of being observed by either everyone or by no one.

In Arieli and Mueller-Frank (2019), agents are placed on a two or higher dimensional grid, and the timing of their actions is given by their distance from the origin. The observation structure is chosen at random according to a parameter p; each agent is independently, with probability p, *connected* to each of her grid neighbors who are closer to the origin. An agent I_n observes an agent I_m if there is a path of connected agents starting from I_n and ending in I_m. This implies that if I_n is not observed by any neighbor, then she would not be observed by any other agent, and so these events cannot be independent. Therefore this setting is a special case of Lobel and Sadler (2015), but not of the Acemoglu et al. (2011) setting, since the realized neighborhoods are not independent.

The authors define a weaker notion of asymptotic learning, which roughly speaking, is that a large fraction of agents in the limit choose the correct action. Their main result is that if $p < 1$ sufficiently large, such learning does occur. As in Acemoglu et al. (2011), this conclusion is based upon sacrificial lambs. For any $p < 1$ there is a constant fraction of agents who observe no other action. The actions of these agents provide independent information to observers. As $p \to 1$ these agents become more rare, but the network becomes more connected, delivering this information to a larger and larger fraction of the population. Thus, learning is achieved as long as there are some sacrificial lambs, regardless of how few.

An interesting open question is whether there exists a *deterministic* network structure in which asymptotic learning is attained for some bounded private signal

[27] For a formal example, see “Nonmonotonicity of Social Beliefs” in Acemoglu et al. (2011, Appendix B).

distribution. The sacrificial lamb mechanism by which asymptotic learning is achieved in *stochastic* networks does not seem to have an obvious analogue in deterministic networks. Thus, learning in deterministic networks would have to result from a different mechanism; it is not known whether such mechanisms exist.

9.2 Repeated Actions

Often in practice people take actions more than once while observing and learning from others bidirectionally. For example, in online social networks people observe product choices and lifestyle choices repeatedly over time. As a major deviation from the sequential learning setting, models of repeated action require different techniques and generate new insights. We focus on models with rational agents.

Under highly connected observation structures (such as one where there is a path of observation between any pair of agents), the concepts of herd behavior and information cascades need to be generalized to apply formally. Following the literature, we consider more general notions of agreement that capture phenomena that are similar to herding. Here *agreement* is the situation in which all agents eventually agree about what action is optimal. Likewise, we consider notions of *learning* that capture repeated action analogues of information cascades. Indeed, herding in sequential models is a form of agreement, since it occurs precisely when all agents (except a finite number, out of infinitely many) agree on the optimal action. In an analysis covering a wide class of social learning models, Mossel et al. (2020) find that agreement and asymptotic learning are closely related, and provide conditions under which each of these phenomena occur.

The study of agreement in social learning goes back at least to Aumann (1976). This seminal paper showed that two agents who have a common prior, who receive a signal regarding a binary state, and who have common knowledge of their posteriors, must have equal posteriors; agents cannot "agree to disagree."

Aumann's model does not consider the dynamics of how agents arrive at common knowledge of posteriors. Geanakoplos and Polemarchakis (1982) study how agreement is reached via social learning. They consider two agents who observe private signals about a state, and then repeatedly tell each other their posteriors. The authors show that the agents reach common knowledge of posteriors, and hence their posteriors will be identical.

Parikh and Krasucki (1990) extend this conclusion to a network setting. They consider a finite set of agents connected by a network. In each period, each agent I_n learns the posteriors of the members of her neighborhood N_n. Under the assumption that the network is *strongly connected*,[28] Parikh and Krasucki show that the agents reach common knowledge of posteriors, and hence agree.

This result was extended by Gale and Kariv (2003) to a setting in which each agent receives a signal at the initial date and then takes an action in each of infinitely many periods. Agents face a common decision problem at each period: they choose an action which results in a stage utility that depends on their action and the unknown state.[29] At each period, they observe their neighbors' actions, but not the stage payoffs. The agents are assumed to be myopic, so that at each period they choose an action that maximizes their stage utility. Neighboring agents do not necessarily eventually agree on actions, even at the limit. But any disagreement is due to indifference. This result was extended by Rosenberg, Solan and Vieille (2009) to forward-looking agents who maximize their discounted expected utility.

The mechanism at work here is the *imitation principle* (also known as the improvement principle; see Golub and Sadler (2016)), which asserts that in the long run an agent will be able to do at least as well as an agent that she observes, i.e., a neighbor. So if an agent disagreed with her neighbor and was not indifferent, she would believe that this neighbor could do better by imitating her, in contradiction with the imitation principle.

[28]The network is strongly connected if there is a chain of neighbors connecting every pair of agents.
[29]A *stage utility* is a payoff received in a particular period of a dynamic game.

While neighboring agents who disagree must be indifferent, non-neighbors can disagree without indifference. Consider a network of four agents I_1, I_2, I_3, I_4, who are connected along a chain, with each observing both the agent's predecessor and successor (if there is one). It is possible for I_1 and I_2 to converge to action L, while I_3 and I_4 converge to H. In this case I_2 and I_3 must be indifferent between L and H, but I_1 and I_4 need not be indifferent. This happens in the case of binary signals and actions, when I_1 and I_2 receive an ℓ signal, I_3 and I_4 receive an h signal, and the agents' tie breaking rule is to stick to their preceding-period action. In this case the agents' actions all immediately converge. After seeing that agent I_3 does not change her action, I_2 concludes that I_4 got an h, and thus I_2 becomes indifferent. Likewise, I_3 becomes indifferent. But I_1 does not know that I_2 is indifferent: from I_1's point of view, it may be that I_3 is also taking action L. Thus I_1 is not indifferent, and neither is I_4, and yet they disagree. In settings with sufficiently rich actions sets, indifference will not occur, but in practical settings action sets are not always rich.

Even when the agents do converge to the same action, it may be an incorrect one, so that the agents do not learn the state. In a quasi-Bayesian setting, Bala and Goyal (1998) show that the learning outcome depends on the network geometry. An important example of a strongly connected network in which agents may not converge to the correct action is the *royal family network*, in which all agents directly observe a small group, but that small group does not directly observe all others.

A similar phenomenon occurs in a setting with forward-looking, Bayesian agents who each receive a signal at period 0, and thereafter choose in each period a binary action with the objective of matching a binary state (Mossel, Sly and Tamuz (2015)). In a royal family network, incorrect signals received by the royal family can cause the entire population to eventually adopt the incorrect action. When this happens, the early period actions of agents in the population are still dependent upon their own private signals, but this information does not propagate through the network. The outcome is closely related to information cascades in models with a single action, in that social information can *eventually* cause the agents to disregard their own private signals.

Conversely, Mossel, Sly and Tamuz (2015) show that in infinite networks that are egalitarian, even with bounded signals, agents all converge to the correct action. In their terms, a network is said to be *egalitarian* if there are integers d and L such that (i) each agent observes at most d others and (ii) if agent I_n observes I_m, then there is a path of length at most L from I_m back to I_n. The first condition excludes agents who obtain large amounts of social information. The second excludes royal families, who are observed by many of their "subjects" but do not reciprocate by observing their subjects directly, or even indirectly through short paths.

Thus, both Bala and Goyal (1998) and Mossel, Sly and Tamuz (2015) conclude that networks in which a subset of agents plays too important a role can hamper the flow of information, resulting in failures of learning, and that asymptotic learning occurs in networks in which all agents play a similar role. The mechanism underlying the failure of asymptotic learning in royal family networks networks is similar to the cause of cascades in the SBM: many agents choose actions that are heavily influenced by social information, and thus do not reveal their private signals. In the SBM, this social information comes from the early agents, who are akin to a royal family in a repeated actions setting.

The association of egalitarianism with information aggregation also holds in models with heuristic updating of beliefs. For instance, Golub and Jackson (2010) use the framework of DeGroot (1974) to study the aggregation of information. In their model, agents follow DeGroot's update rule in which agents calculate their posteriors as the average of those of their neighbors. Golub and Jackson show that over time the agents' beliefs will converge arbitrarily close to being correct (for sufficiently large networks) as long as no agent has too high a degree. This condition can be viewed as another notion of egalitarianism.

When agents choose actions from a set that is rich enough to reveal their beliefs, information is fully aggregated regardless of the network structure (Mueller-Frank (2014)) for reasons similar to those discussed in § 3. Similarly, information is perfectly aggregate in the benchmark case of the model of DeMarzo, Vayanos and Zwiebel (2003) in which agents in a connected social network are interested in a state $\theta \in \mathbb{R}$ for which all have a uniform (improper) prior. Each agent initially observes a Gaussian signal with expectation equal to θ. In each period each agent observes her neighbors' posterior expectations of θ, and updates her posterior using Bayes' Law. Each agent's posterior is Gaussian, and thus is completely characterized by the agent's posterior expectation; the variance depends on the network structure and does not depend upon signal realizations. The agents converge to the same belief they would have if they were to share their private signals.[30]

Frongillo, Schoenebeck and Tamuz (2011) and Dasaratha, Golub and Hak (2019) consider a similar setting except that the state exogenously changes over time, agents receive a new signal at each date, and signals have heterogeneous precisions. Since the state changes, agents cannot learn it exactly. The efficiency of information aggregation depends on the social network and signal structures. In Dasaratha, Golub and Hak (2019), information is better aggregated when there is large heterogeneity in signal precisions, which helps agents filter out stale information that is entangled with information that is more relevant for the current state.

10 Applications and Extensions

We now turn to several extensions of the basic social learning setting that are tailored to specific applications. In these applications, theories of social learning offer new insights into behavior in both market and non-market settings.

[30] The main result of DeMarzo, Vayanos and Zwiebel (2003) describes the effects of persuasion bias, wherein agents update beliefs under the mistaken premise that other agents do not observe others.

10.1 Team Decisions, Optimal Contracting, and Organizational Design

Teams often need to make joint decisions that can be improved by aggregating the information of self-interested team members who may observe others' actions. Information externalities may hamper information aggregation. Such problems can be addressed through the design of a communication network and an incentive scheme for the team. We first discuss settings with rational agents.

As illustrated in the "influencers" example of § 4.4, when better-informed agents decide first, information aggregation can be especially poor. In such situations, anti-seniority voting systems achieve better information aggregation, when seniority is associated with precision.[31]

Khanna and Slezak (2000) study the choice of network structure—teams versus hierarchy—for a firm that seeks to aggregate the signals of multiple employees.[32] Their analysis shows that it can be desirable to assign agents to make decisions independently. This prevents information cascades, and allows a later observer (or more generally, observers) to obtain better social information. This is similar to the discussion of sacrificial lambs in § 6.

In Khanna and Slezak (2000), agents can acquire a private binary signal about the state, which gives information about the profitability of whether or not to adopt a project. The precision of an agent's private signal is increasing in costly private effort. Ex ante identical agents make their investigation decisions before any social observation. In equilibrium they all choose the same effort level. Each agent's action consists of a recommendation (e.g., to adopt or reject). The optimal compensation contract consists of a wage, a bonus if an agent's recommendation is correct ex post, and, if agents are organized into teams, a possible team bonus for a correct recommendation by the team.

[31] However, Ottaviani and Sørensen (2001) show in a reputational model related to Scharfstein and Stein (1990) that anti-seniority voting systems can perform poorly.

[32] We considered more generally the effects of networks on asymptotic learning in § 9.1.

In teams, agent recommendations are announced sequentially. Owing to social observation, in equilibrium agents tend to free ride in generating private signals, a decision that is made ex ante. Furthermore, unless the compensation scheme is heavily weighted toward individual accuracy, agents tend to cascade on the recommendations of earlier agents. The compensation scheme may optimally accommodate such cascading; first-best information generation and aggregation is not attained.

The alternative to teams is hierarchical organization, wherein each agent reports a recommendation to a higher manager without benefit of social observation. Hierarchies reduce free-riding in information acquisition, and prevent the formation of incorrect cascades. This benefit is achieved so long as direct communication among agents can be cheaply suppressed.[33]

In a multi-level hierarchy, an alternative to forcing early agents to make recommendations without observing each other is to artificially coarsen their recommendations. For example, the recommendations of 3 agents can be aggregated into a single overall adopt/reject recommendation on a project. This makes it unclear to an observer whether the preponderance in favor of an action was strong (3 to 0) or weak (2 to 1). Such coarsening can cause a supervisor to whom the agents report to use the supervisor's own private information instead of cascading, potentially improving the firm's decisions (Arya, Glover and Mittendorf (2006)).

10.2 Stigma, Prestige, and Related Phenomena

Job seekers with gaps in their resumes may be stigmatized, as gaps can be indicative that previous potential employers have rejected their applications. Kübler and Weizsäcker (2003) model unemployment stigma in a setting in which employers receive noisy signals about the quality of potential hires. Unlike the basic information cascades setting, both

[33] In a similar spirit, Sgroi (2002) analyzes the benefits to a firm of forcing some agents (such as a subset of customers) to make their decisions early.

employers and workers can take costly unilateral actions that modify signal precision. Kübler and Weizsäcker identify an equilibrium in which only cascades of rejection can occur. There is indeed field and other evidence that employers view gaps in resumes as indicating previous rejections (Oberholzer-Gee (2008), Kroft, Lange and Notowidigdo (2013)).

A similar mechanism can result in refusal of kidneys by patients who need transplants, with kidneys rather than workers being effectively stigmatized. In the U.S. kidney allocation mechanism, patients are sorted into a queue, in order of severity of their conditions. When a kidney becomes available, patients are offered it in sequence. Each patient is offered the kidney only if all predecessors reject it. In the model of Zhang (2010), patients have heterogeneous preferences as they differ in compatibility with an available organ and in the urgency of a transplant. There also is a common component to value; all patients prefer a high-quality kidney. Each patient's physician provides a medical judgment about the suitability of an available kidney, i.e., a private signal. If a patient is offered a kidney after refusals by earlier patients in the queue, she may be in a reject cascade.

Using data from the United States Renal Data System, Zhang (2010) finds strong evidence of social learning. Patients draw negative quality inferences from earlier refusals, causing further refusals. This leads to poor kidney utilization despite a severe shortage of available organs. There is also strong evidence of social learning in other health related decisions, such as choice of health plans (Sorensen (2006)).

Research on cultural evolution has hypothesized that the human mind has evolved to confer prestige upon successful individuals (Henrich and Gil-White (2001)). In this theory, people who defer to an admired individual benefit from being granted access to the information possessed by that individual about how to succeed in their shared environment. If different observers observe different payoff realizations, they will have different information about who is a good decision maker. So observers can acquire useful information by observing whom others defer to. This perspective on prestige suggests that there can be information cascades in conferring prestige. If so, prestige may be a very noisy indicator of decision ability. This is an interesting topic for future research.

10.3 Social Information Use by Animals

Zoologists have developed and applied social learning models, including models of information cascades, to understand the acquisition of social information by animals for decisions about mating, navigation, predator-avoidance, foraging, and habitat selection (see the discussion in BHW IV.B). Several empirical studies of imitation in animals conclude that social learning and cascades are important mechanisms, as distinct, e.g., from payoff externalities. Dall et al. (2005) discuss several proposed examples of animal information cascades, and Giraldeau, Valone and Templeton (2002) apply the logic of information cascades to the conditions under which animals optimally acquire social information.

An extensive empirical literature documents copying in mate choices in many species, including humans (Witte, Kniel and Kureck (2015)). In an experiment on mating decisions, female guppies switched their choice of mate to match the choices made by other females (Dugatkin and Godin (1992)). This is suggestive of overriding of private signals based on social information, as in information cascades, and of fragility of behavior in response to an information shock. In a study of mate choice among sage grouse, Gibson, Bradbury and Vehrencamp (1991) find evidence of imitation and idiosyncrasy: choices tend to be close to unanimous yet poorly correlated with observable characteristics of the so-called "lek displays." In an experiment with human subjects, seeing someone show interest in a member of the opposite sex was found to cause observers to rate the object of attention as more appealing (see Bowers et al. (2012) and the related evidence of Little et al. (2008)). This is consistent with social learning.

10.4 Sequential Voting

In primary elections for nominating party candidates for U.S. presidential elections, early primary election wins cause later voters to support the winner. In consequence, states that go early, such as Iowa and New Hampshire, exert a disproportionate influence on final outcomes. This effect is often called "momentum."

A leading explanation for political momentum is that later voters learn from the behavior of earlier ones. To measure beliefs, studies have used "thermometer ratings," the degree to which survey respondents view candidates favorably. After controlling for other determinants, more favorable poll results about a candidate are associated with subjects evaluating the candidate more positively (see Bartels (1988)). This finding is not explained by strategic voting motives, wherein a voter's willingness to support a candidate depends on the voter's perception of the candidate's prospects, nor by rational conformity preference, wherein the voter desires to vote for a popular candidate.

Social learning effects in sequential voting raise the possibility of the occurrence of some voters sometimes being in information cascades. Of course, sequential voting in practice involves settings that are more complex than the SBM, so the SBM conclusion that all agents start to behave identically may not hold. (As Definition 1 makes clear, information cascades do not necessarily induce unlimited strings of identical behavior.)

While the finding that voters are influenced by poll results is intuitive, equilibria in settings with sequential voting can be subtle. In a sequential setting in which voters with private information seek to elect the best candidate, Dekel and Piccioni (2000) show that there exist equilibria in which voters are not influenced by predecessors' votes. However, Wit (1997), Fey (2000), and Ali and Kartik (2012) show that there also exist equilibria in sequential voting in which voters are influenced by predecessors, resulting in information cascades.

In the cascade equilibrium, however, agents disregard the actions of predecessors who take an off-the-equilibrium-path action. As Wit (1997) and Fey (2000) note, under the reasonable assumption that a voter who breaks a cascade (i.e., off-the-equilibrium-path action) must have voted her signal, the cascade equilibrium is broken. However, the cascade equilibrium becomes plausible if we think of the game as one with even a small chance of observation error.[34]

Such an equilibrium is consistent with the evidence of voting momentum mentioned at the start of this section. With sequential voting, early success of a candidate seems to promote later success owing to favorable inferences drawn by later voters. Knight and Schiff (2010) document such political momentum. They estimate that voters in early states had far greater influence than voters in later states on the outcome of the 2004 U.S. presidential primaries, and that campaign advertising choices took into account such momentum effects. Moreover, Ali et al. (2008) provide laboratory evidence that suggests possible cascading behavior in sequential voting.

Battaglini (2005) modifies Dekel-Piccione's model to allow for a third choice – to abstain, and a cost of voting. With even a small cost of voting, Dekel and Piccione's conclusion does not hold: equilibrium outcomes in the simultaneous voting game are not supported as equilibria of the sequential voting game. So the argument provided by Dekel and Piccione for the existence of a non-cascades equilibrium in the sequential voting game does not carry through to a setting with voting costs. Furthermore, Battaglini shows that information aggregation is worse under sequential voting than under simultaneous voting.

Alternatively, if voting is costless, but voters have a direct preference for voting for the winner (not just for the best candidate), social learning results in information cascades (Callander (2007)). This illustrates a more general point: that information cascades and preference effects can reinforce each other in producing behavioral convergence. In Callander's model the preference is for supporting a winner, which in that

[34]Intuitively, when an agent sees a deviation from the cascade action, the observer may conclude that this is a false observation. Alternatively, deviations may be ignored if the observer believes that occasionally players may have unobservable deviant preferences.

setting is equivalent to a direct preference for conformity.

10.5 Changing States and Fads

The world is continually changing, so the best action to take also fluctuates over time. For example, competing web browsers sometimes leapfrog each other in functionality for users. Suppose that the state of the world determines which action is better, and that this state shifts over time. In such a setting, the mere *possibility* that a shock to the system (that is, a change in the true value of an action) could occur can be enough to dislodge a cascade, even if the shock does not actually occur. In consequence, the probability of actions shifting can be higher than the probability of the state shifting.

To see this, consider a setting as in the SBM except that just before I_{101}'s decision, the state is newly redrawn with probability 0.10, and remains fixed thereafter. So the state changes with probability 0.05. The possibility of a value shift breaks the cascade; I_{101} optimally follows her own signal. So the action often shifts even when the state does not. Eventually the system must settle into a new cascade—one that need not match the old one. It is easy to show that the probability of a cascade reversal is a little over $0.0935 \gg 0.05$ (see BHW). So the probability that the cascade shifts is 87% higher than would be the case under full information. BHW refer to such volatile outcomes as *fads*.

What if there is a probability *each period* that the state shifts? Then if the probability is not too large, there are still cascades. However, these cascades are transient; over time the social information in any cascade grows stale and some agent eventually returns to using her own signals (Moscarini, Ottaviani and Smith (1998)). Huang (2022) shows that the changes in action can be more frequent than the changes in state. In settings with repeated actions, it has also been found that the probability of change in the preponderant action can be either less than or greater than the probability of a change in state (inertia or impulsiveness respectively; see Hirshleifer and Welch

(2002)).

A fruitful further direction is to study whether something akin to fads occurs in social learning settings in which information cascades do not occur. A further interesting direction would be to study the determinants of fads and the stability of conformist social outcomes. For example, it might seem that less stable environments would be more fad-prone, but this is not obvious. A greater probability of state shift need not imply a greater excess probability of action shifts *relative to* the probability of state shift.

10.6 Law, Politics and Revolutions

A common interpretation of the legal norm for respect for precedent (*stare decisis*), at least since Oliver Wendell Holmes, is that judges acquire information from past decisions. Several legal scholars and economists have modelled respect for precedent as a form of social learning involving information cascades (Talley (1999) and Daughety and Reinganum (1999); see also Vermeule (2012)). Daughety and Reinganum (1999) also offer possible examples of actual precedential cascades among appellate courts.[35] These authors discuss the likely prevalence and stability of precedential cascades in different contexts. We discuss legal precedent more extensively in the Online Appendix, §A.3.

Turning to politics and revolutions, when citizens publicly protest or revolt against their government, the actions of early individuals convey information about the prevalence of dissatisfaction, and of the risks of government sanctions. So there can be positive feedback in submission or resistance to the regime.

In Kuran (1987, 1989), there is a pressure toward conformity, and agents have

[35]To the extent that *stare decisis* is mandated, a court may feel pressured to follow precedent rather than doing so for purely informational reasons. However, courts do sometimes deviate from precedent, perhaps because there can be gray areas in the applicability of a precedent. So this pressure is not absolute. The pressure of *stare decisis* as a social norm presents a challenge for empirically identifying informational effects. At a deeper level, social learning may explain why the norm of *stare decisis* originally emerged in common law legal systems.

publicly-declared preferences that differ from their actual preferences. This can support the survival of oppressive regimes. Lohmann (1994) models belief updating to analyze the maintenance and collapse of political regimes as information cascades.

In Lohmann's model, agents have different preferred policies, and derive disutilities when the government's policy differs from their the preferred ones. All agents know the probability distribution of preferred policies. Each agent receives a binary signal about the level of the future policy that will be selected by the current regime. Each agent in sequence has the opportunity to protest to promote an alternative regime that will implement policies that are potentially closer to the agent's preferred one. Each agent also has a disutility component from protesting which decreases with the number of fellow protestors (a "safety in numbers" effect). The existing regime collapses if enough information is revealed to show that more than some critical number of individuals support an alternative regime.

Lohmann applies the model to the Leipzig protests against the communist regime in East Germany during 1989-91. In her analysis, even a small protest, when larger than expected, can reveal strong opposition to the regime, causing the size of the protest to grow explosively. By the same token, when a large protest is expected, if the protest falls modestly short of expectations, the movement can suddenly collapse. Lohmann argues that a dynamic information cascades model helps explain the successful popular uprising against the communist regime in East Germany over five cycles of protests.

In a model of democratization, repression and regime change of Acemoglu and Robinson (2006), a repressive regime run by the rich can to stave off revolution by making concessions to the poor. Some fraction of the wealth is destroyed in a revolution, which limits the incentive of the poor to rebel.

Ellis and Fender (2011) combine features of Lohmann's model with those of Acemoglu and Robinson (2006). In the Ellis and Fender model, the poor have private information about regime strength. In what we call the L state, a revolution would destroy a higher fraction of the wealth to be appropriated by the poor than in the H state, making it less attractive for the poor to revolt in state L. So the regime is stronger in the H state than in the L state.

In the case of interest, the rich set the tax rate such that the poor would like to revolt if they knew that the state was H and not revolt if they knew the state was L. If the poor are not granted the franchise, based on their private binary signals, poor agents (hereafter, "agents") decide in sequence whether to rebel against the regime. All actions are observable. By assumption, a revolution occurs only if every agent rebels. At first, an agent with an h signal rebels, whereas an agent with an ℓ signal does not, thereby vetoing the rebellion. However, after a sufficiently long sequence of rebels (i.e., h signals), even an agent with an ℓ signal rebels. So revolution occurs as an information cascade.

10.7 Group Adoption of Conventions

Social conventions are sometimes transmitted from one group to the other. When there is limited observation across groups, distinct cascades of individual decisions can temporarily form. In Fisher and Wooders (2017), two SBMs (groups) run in parallel. A single agent, is common to both groups, and observes the actions of all predecessors in each group. Every later agent is in only one of the two groups. In the model, even very limited overlap in membership between groups can result in an information cascade jumping from one group to another.

10.8 Belief Cascades

Several authors have suggested that the dynamics of public opinions, such as the rise and fall of ideologies, and conspiracy theories, and the spread of fake news, are information cascades (Kuran and Sunstein (2000), Sunstein (2019, p.50)). In the context of social media, a possible pathway could be the binary decisions of whether or not to like or to forward a news item, making information cascades possible.

However, people often express opinions going beyond a binary choice to like or forward. This raises an important question which that we have not considered so far: are there information cascades in publicly expressed *beliefs*? In Online Appendix A.2, we provide a simple extension of the SBM in which belief cascades occur. In this setting, the action taken by agents is to make binary public *reports* about their beliefs, choosing the report that aligns most closely with their beliefs. Agents observe the reports of all predecessors before making their decisions. We define a *belief cascade* as a situation in which, having received the reported beliefs of some set of predecessors, an agent's reported belief is independent of the agent's private signal. In the model here, belief cascades occur, for the same reason that cascades occur in the SBM or in §4.

The key motivation for such a model is that people often communicate in coarse categories, such as binary partitions. When asked about a model of car, people often say that they like or dislike it, as opposed to reporting intensities of liking on a continuum, detailed reasons for their opinions, or probability distributions over which model is the best. Such digitized communication is inevitable owing to limited time and cognitive resources of speakers and listeners.

Whether belief cascades occur in practice is an empirical question. A possible application is to the spread of scientific claims. Since it is costly to read original source articles in detail, scholars often rely on descriptions of the original findings in later publications. Greenberg (2009) provides bibliometric evidence from the medical literature about the spread via chains of citations of unfounded scientific claims. The broader message is that a rich direction for future research is the application of social learning theory to the decision of which beliefs to publicly espouse.

11 Further Directions

We next discuss avenues for future theoretical research on information cascades and social learning.

11.1 Information Design

A fruitful direction for further exploration is the problem of information design with sequential social learning. One possible purpose could be to improve social welfare. We have seen that this can sometimes be achieved by assigning some agents to be sacrificial lambs. Alternatively, a manager or a seller can design or influence the signal structure or the observation network to achieve private objectives. This amounts to a form of sequential Bayesian persuasion (Kamenica and Gentzkow (2011)).

Several of the papers we have discussed have some form of information design, as when a monopolist sets prices to influence the decisions of early buyers who in turn influence later buyers; or when a principal designs an organization's observation structure to motivate agents appropriately (Welch (1992)), Bose et al. (2008), and Khanna and Slezak (2000). In these papers, the principal does not directly provide signals to agents. More generally, a principal may direct information to agents, perhaps dynamically as a function of past actions and outcomes (see, e.g., (Kremer, Mansour and Perry (2014); Che and Hörner (2018)). This is a promising path to pursue.

As discussed later in § 11.2, public discourse can be modeled as a part of a social learning process in which beliefs are transmitted to others. The internet has given firms much more flexibility to design social media platforms to influence who observes whom, and to curate what (mis)information an agent sees. This includes users' verbal opinions, recommendations, and anecdotes, as well as visual, aural, and even physiological traces. So the design of institutions to shape social learning is central to modern value creation and public policy. This has led to much debate about the design of social media algorithms.

Several questions about information design are being actively investigated. How does the introduction of bots and fake postings affect social learning? When do principals benefit from the ability to commit to fix their own actions and/or platform structures? How do the interests of social media platform owners diverge from maximizing social welfare, and can regulatory policies address such divergence?

11.2 Imperfect Rationality

Some models of imperfect rationality in social learning allow for overconfidence or for agents misestimating the extent to which others are observing predecessors. In models with repeated actions on networks, the common assumption of myopia can be viewed as weakening the assumption of forward-looking rationality (Gale and Kariv (2003)).

As discussed in § 6.2, biases that cause agents to be more aggressive in using their own signals tend to improve social information aggregation, whereas biases in the opposite direction worsen it. This suggests that the relative success of individuals and groups may depend upon the group's distribution of psychological agent types (e.g., Bernardo and Welch (2001)).

The distribution of psychological types affects both the speed of social learning and whether there is asymptotic learning. In evolutionary settings, the success of individual types depends on a balance between the direct benefit of having a psychological trait versus the benefits to being in a group in which others have that trait. It would be valuable to explore how this tension affects the extent to which imperfectly rational types survive. The survival of different psychological types (e.g., the degree of neurodiversity) can affect welfare in many ways. For example, types that heavily use their own private signals are informationally helpful to others.

Owing in part to different assumptions about the form of imperfect rationality, existing models (Bohren (2016), Bohren and Hauser (2019), Frick, Iijima and Ishii

(2020)) obtain dissonant conclusions about whether small biases have large effects. In models with large effect, biases can be socially amplified, generating even larger and more persistent effects on social outcomes, such as failures of asymptotic learning. Such social amplification suggests that policy interventions may also have large multiplier effects.

Another question is how imperfectly rational social learning influences boom/bust dynamics in activities such as investments or mergers and acquisitions. We have discussed rational models of booms and busts (Chamley and Gale (1994)), including some within market settings (Lee (1998)). However, many observers have argued that imperfect rationality contributes to booms and busts, as in the models of Bohren and Hauser (2019) and Hirshleifer and Plotkin (2021). It will be valuable to perform empirical testing with a view toward developing models that reflect the institutional facts of applied settings in a more granular way.

An especially rich direction is to analyze how social learning and information cascades influence public discourse. This may provide insight into the spread of spread of political and religious ideologies, folk economic thinking, and conspiracy theories. The wide availability of textual and social media databases in turn offers rich data to test hypotheses about social learning and misinformation. The model of belief cascades in § 10.8 (in which belief reports play the roles of actions) is a possible initial step. A further direction will be to incorporate a wider set of psychological biases that influence human communication. These include neglect of selection in what others choose to assert, and confirmation bias. Such analysis can provide insight into why in some domains discourse grows ever more sophisticated (as in the scientific scholarly community), whereas in others discourse remains persistently naïve (as with much political discourse).

11.3 Repeated Actions

We discussed the nascent literature on repeated actions by Bayesian agents in § 9.2. It will be useful to explore further the speed of learning in such settings (see Harel et al. (2021)) and the question of which social networks promote more efficient information aggregation deserves (Golub and Jackson (2010); Mossel, Sly and Tamuz (2015)).

A new strategic incentive that arises when agents act repeatedly is to induce others to convey useful information. For example, an agent whose private information opposes the agent's social information can benefit from following private information, even if that reduces the agent's current payoff, if that can induce others to take more informative actions. The topic of strategic information incentives is almost completely unexplored (but see the "mad king" example in Mossel, Sly and Tamuz (2015)).

Only a handful of papers consider social learning with repeated actions and a changing state (see Frongillo, Schoenebeck and Tamuz (2011) and § 10.5 on fads); even fewer do so in a repeated action setting (Dasaratha, Golub and Hak (2019)). Many interesting questions remain unexplored. In a steady state, how well attuned are actions to state? How does this concordance depend on the private signal structure, or on deviations from rationality? How quickly does the preponderant action change after a change in state?

11.4 The Speed of Learning

Most existing research emphasizes whether or not asymptotic learning occurs, rather than how quickly learning occurs (but see the recent exceptions noted in § 4.3). Under a social welfare function with discounting, it becomes vital that learning occur quickly, not just asymptotically. Information externalities impede that objective by delaying learning. Research is starting to explore how closely outcomes can approach an efficient benchmark (Chamley (2004*a*); Smith, Sørensen and Tian (2021)).

An issue for further exploration is the extent to which the principle of countervailing adjustment (§ 6.1) reduces the sensitivity of the speed of learning to the parameters of the learning process. Intuitively, parameter shifts that have a direct effect, at a given point in time, of making social information more accurate will tend to make the given agent's actions depend more weakly on her own private signal. This can oppose the direct increase in social information induced by the exogenous shift.

11.5 Payoff Externalities

In addition to the information externalities considered so far, there are sometimes also externalities wherein an agent's action directly affects the payoff of another agent. A topic that merits further exploration is the interaction between direct payoff externalities, social learning, and information cascades. Veeraraghavan and Debo (2011) model congestion (a negative externality) as the cost of waiting in a queue with random service times. The queue length conveys favorable information about the value of the service provided, so under some conditions there can be cascades on joining the longer queue. Eyster et al. (2014) describe how negative externalities prevent the convergence of agents to one action and improve social learning despite the continued occurrence of incorrect information cascades. Other analyses of social learning with payoff interactions are discussed extensively in Chamley (2004*b*). Since queues and congestion are very important in practice, future research is likely to explore these issues with assumptions that are more fully attuned to various applications.

12 Conclusion and Empirical Testing

Social learning and information cascades have offered a new understanding of three phenomena. First, individuals who learn by observing others often converge to the same action. Second, individuals who in the aggregate possess extensive information converge surprisingly often on an incorrect action, or only very slowly on the right one. Relatedly, there is path-dependence in social outcomes. Third, people are often attached to popular actions loosely, resulting in fragility of mass outcomes, and fads.

The analysis of social learning uncovers effects that differ from most of traditional information economics. Information economics focuses mainly on strategic motives as the actions of agents directly affect the payoffs of others. In contrast, in much of the social learning literature, externalities are purely informational. Thus, central topics of much of information economics, such as adverse selection and moral hazard, are not the main drivers of most social learning models.

The social learning literature provides a rich set of insights using different modeling assumptions, including the possibility that agents choose when to act, whether to acquire private signals, or making trades in a marketplace. Recent work has opened new vistas by analyzing cases in which agents take repeated actions, are imperfectly rational, or observe others in complex social networks. Several rich directions for future research are discussed in §11.

Although this review has focused on theoretical research, we now briefly discuss empirical considerations. A general challenge for empirical testing of social learning theories is that conformity can derive from many possible sources. These mechanisms include independent observation of the same information, a direct preference for conforming to the behavior of others, sanctions upon deviants, positive payoff externalities, and social learning.

Much research on testing for social influence (often called "peer effects") examines conditions for identification in the absence of information about the sequencing of actions. Manski (1993) describes the "reflection problem," which is that in certain simple social influence settings, empiricists cannot infer from the average behavior alone whether this average has influenced the decisions of group members. A further literature shows, however, that relatively modest restrictions on the structure of social influence allow identification of peer effects based on observations of individual behaviors, or even from average behaviors of groups (Blume et al. (2015)).

Data with information about the order in which agents act provides valuable further restrictions, since early agents influence later ones rather than vice versa. However, sequencing information alone does not necessarily fully identify social learning effects, since there can be unobserved external causes. Further possible partial identification derives from examining the effects of aggregate shocks, such as the prediction of the information cascades setting that outcomes are fragile. This contrasts sharply with settings in which conformity derives from positive payoff externalities, positive preference interactions, or from sanctions upon deviants; in such models, small shocks typically have small effects.

There are other possible means of identifying social learning effects. Often there are *negative* payoff externalities from following the same action, as with people joining a queue for the same restaurant or takeover bidders vying for the same target firm. In such contexts, observation of homogeneous behavior rules out payoff externalities as the dominant force, though it does not uniquely pinpoint social learning as the cause of conformity.

Some studies identify social learning using data on the structure of private information signals, or on the observation network, or design such elements within a laboratory or field experiment. Early laboratory experiments about social learning by Anderson and Holt (1997) provide evidence that information cascades occur, but also provide evidence of deviations from rational behavior.[36]

Several developments have created exciting opportunities for empirical testing of social learning effects. First, social media and other electronic databases now provide information about the network of social observation, such as the Facebook Social

[36]Many studies find that, inconsistent with rationality, subjects tend to overuse their private signals, which can break information cascades (see, e.g., Weizsäcker (2010)), consistent, e.g., with Bernardo and Welch (2004)). When agents over-rely on their private signals, we still expect cascades to form, even if there tends to be greater information aggregation.

Connectedness Index (Bailey et al. (2018)).

Second, social learning theories have traditionally been challenging to test because of lack of data on agents' private information signals. However, in recent years, textual and in some cases social-network data provide proxies for the information signals available to different agents or the beliefs that such agents have formed. Such proxies have been used to study the effects of language, images, and video on economic outcomes (see the review of Nekrasov, Teoh and Wu (2023)).

Third, field experimentation promises to address the external validity of laboratory tests. Salganik, Dodds and Watts (2006) provide evidence of social influence in song downloads in a field experiment on music marketplaces. Other research tests for social learning in the context of economic development (see, e.g., Banerjee et al. (2013)).

For all these reasons, empirical testing is playing an increasingly important role both for testing existing models and stimulating new theorizing. Testing the predictions of information cascades and social learning models outside the lab remains an important research challenge. An interesting direction for further exploration is whether conformists or deviants from the predominant action make better decisions. In some empirical settings, agents' actions and payoffs are indeed observable to the researcher. The answer to this question differs across models, so addressing it empirically, may help distinguish alternative mechanisms of social influence.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. 2019. "Learning From Reviews: The Selection Effect and the Speed of Learning." Working Paper, Massachusetts Institute of Technology.

Acemoglu, Daron, and Asuman E. Ozdaglar. 2011. "Opinion Dynamics and Learning in Social Networks." *Dynamic Games and Applications*, 1(1): 3–49.

Acemoglu, Daron, and James A. Robinson. 2006. *Economic Origins of Dictatorship and Democracy.* New York, NY:Cambridge University Press.

Acemoglu, Daron, Munther Dahleh, Ilan Lobel, and Asuman E. Ozdaglar. 2011. "Bayesian Learning in Social Networks." *Review of Economic Studies*, 78(4): 1201–1236.

Aghamolla, Cyrus. 2018. "Observational Learning, Endogenous Timing, and Information Acquisition." Working Paper, University of Minnesota.

Aghamolla, Cyrus, and Ilan Guttman. 2021. "Strategic timing of IPOs and disclosure: A dynamic model of multiple firms." *Accounting Review*, 96(3): 27–57.

Aghamolla, Cyrus, and Tadashi Hashimoto. 2020. "Information arrival, delay, and clustering in financial markets with dynamic freeriding." *Journal of Financial Economics*, 138(1): 27–52.

Alevy, Jonathan E., Michael S. Haigh, and John A. List. 2007. "Information Cascades: Evidence from a Field Experiment with Financial Market Professionals." *Journal of Finance*, 62(1): 151–180.

Ali, S. Nageeb. 2018*a*. "Herding with costly information." *Journal of Economic Theory*, 175: 713–729.

Ali, S. Nageeb. 2018*b*. "On the Role of Responsiveness in Rational Herds." *Economics Letters*, 163: 79–82.

Ali, S. Nageeb, and Navin Kartik. 2012. "Herding with collective preferences." *Economic Theory*, 51(3): 601–626.

Ali, S Nageeb, Jacob K Goeree, Navin Kartik, and Thomas R Palfrey. 2008. "Information aggregation in standing and ad hoc committees." *American Economic Review*, 98(2): 181–86.

Amihud, Yakov, Shmuel Hauser, and Amir Kirsh. 2003. "Allocations, adverse selection, and cascades in IPOs: Evidence from the Tel Aviv Stock Exchange." *Journal of Financial Economics*, 68(1): 137–158.

Anderson, Lisa R., and Charles A. Holt. 1997. "Information Cascades in the Laboratory." *American Economic Review*, 87(5): 847–862.

Anderson, Lisa R., and Charles A. Holt. 2008. "Information Cascade Experiments." In *Handbook of Experimental Economics Results*. Vol. 1, , ed. Charles R. Plott and Vernon L. Smith, Chapter 39, 335–343. Elsevier.

Arieli, Itai, and Manuel Mueller-Frank. 2019. "Multidimensional Social Learning." *Review of Economic Studies*, 86(3): 913–940.

Arieli, Itai, Yakov Babichenko, Stephan Müller, Farzad Pourbabaee, and Omer Tamuz. 2023. "The Hazards and Benefits of Condescension in Social Learning." *arXiv preprint arXiv:2301.11237*.

Arthur, W. Brian. 1989. "Competing Technologies, Increasing Returns, and Lock-in by Historical Events." *Economic Journal*, 99(394): 116–131.

Arya, Anil, Jonathan Glover, and Brian Mittendorf. 2006. "Hierarchical Reporting, Aggregation, and Information Cascades." *Managerial and Decision Economics*, 27(5): 355–362.

Aumann, Robert J. 1976. "Agreeing to Disagree." *Annals of Statistics*, 4(6): 1236–1239.

Avery, Christopher, and Peter Zemsky. 1998. "Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets." *American Economic Review*, 88(4): 724–48.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. 2018. "The Economic Effects of Social Networks: Evidence from the Housing Market." *Journal of Political Economy*, 126(6): 2224–2276.

Bala, Venkatesh, and Sanjeev Goyal. 1998. "Learning from Neighbours." *Review of Economic Studies*, 65(3): 595–621.

Banerjee, Abhijit. 1992*a*. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107(3): 797–817.

Banerjee, Abhijit. 1992*b*. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107(3): 797–817.

Banerjee, Abhijit, and Drew Fudenberg. 2004. "Word of Mouth learning." *Games and Economic Behavior*, 46(1): 1–22.

Banerjee, Abhijit V., Arun Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. “The Diffusion of Microfinance.” *Science*, 341(6144): 1236498.

Bartels, Larry M. 1988. *Presidential Primaries and the Dynamics of Public Choice.* Princeton, New Jersey:Princeton University Press.

Battaglini, Marco. 2005. “Sequential voting with abstention.” *Games and Economic Behavior*, 51(2): 445–463.

Bernardo, Antonio E., and Ivo Welch. 2001. “On the Evolution of Overconfidence and Entrepreneurs.” *Journal of Economics and Management Strategy*, 10(3): 301–330.

Bernardo, Antonio E., and Ivo Welch. 2004. “Liquidity and Financial Market Runs.” *Quarterly Journal of Economics*, 119(1): 135–158.

Bernheim, B. Douglas. 1994. “A Theory of Conformity.” *Journal of Political Economy*, 102(5): 841–877.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades.” *Journal of Political Economy*, 100(5): 992–1026.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1998. “Learning from the behavior of others: Conformity, fads, and informational cascades.” *Journal of Economic Perspectives*, 12(3): 151–170.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Rajshri Jayaraman. 2015. “Linear social interactions models.” *Journal of Political Economy*, 123(2): 444–496.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides. 2011. “Identification of Social Interactions.” In *Handbook of Social Economics*. Vol. 1, Chapter 18, 853–964. Amsterdam: North Holland.

Boccaletti, Stefano, Vito Latora, Yamir Moreno, Martin Chavez, and D-U Hwang. 2006. “Complex networks: Structure and Dynamics.” *Physics Reports*, 424(4-5): 175–308.

Bohren, J. Aislinn. 2016. “Informational herding with model misspecification.” *Journal of Economic Theory*, 163: 222–247.

Bohren, J. Aislinn, and Daniel N. Hauser. 2019. “Social Learning with Model Misspecification: A Framework and a Robustness Result.” Working paper, University of Pennsylvania.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica*, 67(2): 349–374.

Bonacich, Phillip. 1987. "Power and centrality: A family of measures." *American Journal of Sociology*, 92(5): 1170–1182.

Bose, Subir, Gerhard Orosel, Marco Ottaviani, and Lise Vesterlund. 2008. "Monopoly pricing in the binary herding model." *Economic Theory*, 37(2): 203–241.

Bowers, Robert I., Skyler S. Place, Peter M. Todd, Lars Penke, and Jens B. Asendorpf. 2012. "Generalization in mate-choice copying in humans." *Behavioral Ecology*, 23(1): 112–124.

Boyd, Robert, and Peter J. Richerson. 1985. *Culture and the Evolutionary Process.* Chicago, IL:University of Chicago Press.

Brunnermeier, Markus. 2001. *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis and Herding.* Oxford, UK:Oxford University Press.

Burguet, Roberto, and Xavier Vives. 2000. "Social Learning and Costly Information Acquisition." *Journal of Economic Theory*, 15(1): 185–205.

Callander, Steven. 2007. "Bandwagons and Momentum in Sequential Voting." *Review of Economic Studies*, 74: 653–684.

Callander, Steven, and Johannes Hörner. 2009. "The wisdom of the minority." *Journal of Economic Theory*, 144(4): 1421–1439.

Caminal, Ramon, and Xavier Vives. 1996. "Why Market Shares Matter: An Information-Based Theory." *The RAND Journal of Economics*, 27(2): 221–239.

Caminal, Ramon, and Xavier Vives. 1999. "Price Dynamics and Consumer Learning." *Journal of Economics and Management Strategy*, 8(1): 95–131.

Cao, H. Henry, and David Hirshleifer. 1997. "Limited observability, reporting biases, and informational cascades." Working paper, Fisher College of Business, The Ohio State University.

Cao, H. Henry, Bing Han, and David Hirshleifer. 2011. "Taking the Road Less Traveled by: Does Conversation Eradicate Pernicious Cascades?" *Journal of Economic Theory*, 146(4): 1418–1436.

Caplin, Andrew, and John Leahy. 1993. "Sectoral Shocks, Learning and Aggregate Fluctuations." *Review of Economic Studies*, 60: 777–794.

Caplin, Andrew, and John Leahy. 1994. "Business as Usual, Market Crashes, and Wisdom after the fact." *American Economic Review*, 84(3): 548–565.

Castellano, Claudio, Santo Fortunato, and Vittorio Loreto. 2009. "Statistical physics of social dynamics." *Reviews of Modern Physics*, 81(2): 591.

Celen, Bogachan, and Shachar Kariv. 2004*a*. "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." *American Economic Review*, 94(3): 484–497.

Çelen, Boğaçhan, and Shachar Kariv. 2004*b*. "Observational learning under imperfect information." *Games and Economic Behavior*, 47(1): 72–86.

Chamley, Christophe, and Douglas Gale. 1994. "Information Revelation and Strategic Delay in Irreversible Decisions." *Econometrica*, 62(5): 1065–1085.

Chamley, Christophe P. 2004*a*. "Delays and equilibria with large and small information in social learning." *European Economic Review*, 48(3): 477–501.

Chamley, Christophe P. 2004*b*. *Rational Herds: Economic Models of Social Learning.* Cambridge, UK:Cambridge University Press.

Chari, V.V., and Patrick J. Kehoe. 2004. "Financial Crises as Herds: Overturning the critiques." *Journal of Economic Theory*, 119(1): 128–150.

Cheng, Yu, Wade Hann-Caruthers, and Omer Tamuz. 2018. "A deterministic protocol for sequential asymptotic learning." 1735–1738, IEEE.

Chen, Jaden Yang. 2021. "Sequential Learning under Informational Ambiguity." *Available at SSRN 3480231*.

Che, Yeon-Koo, and Johannes Hörner. 2018. "Recommender systems as mechanisms for social learning." *Quarterly Journal of Economics*, 133(2): 871–925.

Cipriani, Marco, and Antonio Guarino. 2008*a*. "Herd Behavior and Contagion in Financial Markets." *B.E. Journal of Theoretical Economics*, 8(1): Article 24.

Cipriani, Marco, and Antonio Guarino. 2008*b*. “Transaction Costs and Informational Cascades in Financial Markets.” *Journal of Economic Behavior and Organization*, 68: 581–592.

Cipriani, Marco, and Antonio Guarino. 2014. “Estimating a Structural Model of Herd Behavior in Financial Markets.” *American Economic Review*, 104(1): 224–251.

Crawford, Vincent P., and Joel Sobel. 1982. “Strategic Information Transmission.” *Econometrica*, 50(6): 1431–1451.

Dall, Sasha R. X., Luc-Alain Giraldeau, Ola Olsson, John M. McNamara, and David W. Stephens. 2005. “Information and its use by animals in evolutionary ecology.” *Trends in Ecology & Evolution*, 20(4): 187–193.

Dasaratha, Krishna, and Kevin He. 2019. “Aggregative Efficiency of Bayesian Learning in Networks.” *arXiv preprint arXiv:1911.10116*.

Dasaratha, Krishna, Benjamin Golub, and Nir Hak. 2019. “Social Learning in a Dynamic Environment.” Working paper, Harvard University.

Daughety, Andrew F., and Jennifer F. Reinganum. 1999. “Stampede to Judgment: Persuasive Influence and Herding Behavior by Courts.” *American Law and Economics Review*, 1(1 & 2).

Décamps, Jean-Paul, and Stefano Lovo. 2006. “Informational cascades with endogenous prices: The role of risk aversion.” *Journal of Mathematical Economics*, 42(1): 109–120.

DeGroot, Morris H. 1974. “Reaching a Consensus.” *Journal of the American Statistical Association*, 69(345): 118–121.

Dekel, Eddie, and Michele Piccioni. 2000. “Sequential Voting Procedures in Symmetric Binary Elections.” *Journal of Political Economy*, 108(1): 34–44.

DeMarzo, Peter, Dimitri Vayanos, and Jeff Zwiebel. 2003. “Persuasion Bias, Social Influence, and Uni-Dimensional Opinions.” *Quarterly Journal of Economics*, 118: 909–968.

Duan, Wenjing, Bin Gu, and Andrew B. Whinston. 2009. “Informational Cascades and Software Adoption on the Internet: An Empirical Investigation.” *MIS Quarterly*, 33(1): 23–48.

Dugatkin, Lee A., and Jean-Guy J. Godin. 1992. “Reversal of female mate choice by copying in the guppy (*Poecilia reticulata*).” *Proceedings of the Royal Society of London, Series B*, 249: 179–184.

Ellis, Christopher J., and John Fender. 2011. “Information Cascades and Revolutionary Regime Transitions.” *Economic Journal*, 121(1): 763–792.

Ellison, Glenn, and Drew Fudenberg. 1993. “Rules of thumb for social learning.” *Journal of Political Economy*, 101(4): 612–643.

Ellison, Glenn, and Drew Fudenberg. 1995. “Word-of-Mouth Communication and Social Learning.” *Quarterly Journal of Economics*, 110(1): 93–125.

Enke, Benjamin, and Florian Zimmermann. 2019. “Correlation Neglect in Belief Formation.” *Review of Economic Studies*, 86(1): 313–332.

Eyster, Erik, and Matthew Rabin. 2010. “Naïve Herding in Rich-Information Settings.” *American Economic Journal: Microeconomics*, 2(4): 221–243.

Eyster, Erik, and Matthew Rabin. 2014. “Extensive Imitation is Irrational and Harmful.” *Quarterly Journal of Economics*, 129(4): 1861–1898.

Eyster, Erik, Andrea Galeotti, Navin Kartik, and Matthew Rabin. 2014. “Congested observational learning.” *Games and Economic Behavior*, 87: 519–538.

Eyster, Erik, Matthew Rabin, and Dimitri Vayanos. 2018. “Financial markets where traders neglect the informational content of prices.” *Journal of Finance*, 74(1): 371–399.

Fey, Mark. 2000. “Informational Cascades and Sequential Voting.” *University of Rochester. Mimeo*.

Fisher, James C.D., and John Wooders. 2017. “Interacting Information Cascades: On the Movement of Conventions Between Groups.” *Economic Theory*, 63(1): 211–231.

Frick, Mira, Ryota Iijima, and Yuhta Ishii. 2020. “Misinterpreting Others and the Fragility of Social Learning.” *Econometrica*, 88(6): 2281–2328.

Frongillo, Rafael M, Grant Schoenebeck, and Omer Tamuz. 2011. “Social learning in a changing world.” 146–157, Springer.

Gale, Douglas. 1996. “What have we learned from social learning?” *European Economic Review*, 40: 617–628.

Gale, Douglas, and Shachar Kariv. 2003. “Bayesian learning in social networks.” *Games and Economic Behavior*, 45(2): 329–346.

Galeotti, Andrea, and Sanjeev Goyal. 2010. "The Law of the Few." *American Economic Review*, 100(4): 1468–1492.

Geanakoplos, John D., and Herakles M. Polemarchakis. 1982. "We can't disagree forever." *Journal of Economic Theory*, 192–200.

Gibson, Robert M., Jack W. Bradbury, and Sandra L. Vehrencamp. 1991. "Mate Choice in Lekking Sage Grouse Revisited: The Roles of Vocal Display, Female Site Fidelity, and Copying." *Behavioral Ecology*, 2(2): 165–180.

Giraldeau, Luc-Alain, Thomas J. Valone, and Jennifer J. Templeton. 2002. "Potential disadvantages of using socially acquired information." *Philosophical Transactions: Biological Sciences*, 357(1427): 1559–1566.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics*, 14(1): 71–100.

Golub, Benjamin, and Evan Sadler. 2016. "Learning in Social Networks." In *The Oxford Handbook of the Economics of Networks*. , ed. Yann Bramoullé, Andrea Galeotti and Brian W. Rogers, Chapter 19, 504–542. New York, NY:Oxford University Press.

Golub, Benjamin, and Matthew O Jackson. 2010. "Naive learning in social networks and the wisdom of crowds." *American Economic Journal: Microeconomics*, 2(1): 112–49.

Greenberg, Steven A. 2009. "How citation distortions create unfounded authority: Analysis of a citation network." *British Medical Journal*, 339: b2680.

Grenadier, Steven R. 1999. "Information revelation through option exercise." *Review of Financial Studies*, 12(1): 95–129.

Guarino, Antonio, Heike Harmgart, and Steffen Huck. 2011. "Aggregate Information Cascades." *Games and Economic Behavior*, 73(1): 167–185.

Hann-Caruthers, Wade, Vadim V. Martynov, and Omer Tamuz. 2018. "The speed of sequential asymptotic learning." *Journal of Economic Theory*, 173: 383–409.

Harel, Matan, Elchanan Mossel, Philipp Strack, and Omer Tamuz. 2021. "Rational Groupthink." *Quarterly Journal of Economics*, 136(1): 621–668.

Harris, Lawrence. 1991. “Stock price clustering and discreteness.” *Review of Financial Studies*, 4(3): 389–415.

Hendricks, Kenneth, and Dan Kovenock. 1989. “Asymmetric Information, Information Externalities, and Efficiency: The Case of Oil Exploration.” *Rand Journal of Economics*, 20(2): 164–182.

Henrich, Joseph, and Francisco Gil-White. 2001. “The Evolution of Prestige: Freely conferred status as a mechanism for enhancing the benefits of cultural transmission.” *Evolution and Human Behavior*, 22(3): 165–196.

Herrera, Helios, and Johannes Horner. 2013. “Biased Social Learning.” *Games and Economic Behavior*, 80: 131–146.

Hirshleifer, David, and Ivo Welch. 2002. “An Economic Approach to the Psychology of Change: Amnesia, Inertia, and Impulsiveness.” *Journal of Economics and Management Strategy*, 11: 379–421.

Hirshleifer, David, and Joshua Plotkin. 2021. “Moonshots, Investment Booms, and Selection Bias in the Transmission of Cultural Traits.” *Proceedings of the National Academy of Sciences*, 118(e2015571118).

Hirshleifer, David, and Siew Hong Teoh. 2003. “Herd Behavior and Cascading in Capital Markets: A Review and Synthesis.” *European Financial Management*, 9(1): 25–66.

Hirshleifer, David, and Siew Hong Teoh. 2009. “Thought and Behavior Contagion in Capital Markets.” In *Handbook of Financial Markets: Dynamics and Evolution*. *Handbooks in Finance*, , ed. Thorsten Hens and Klaus Schenk-Hoppe, Chapter 1, 1–46. Amsterdam, The Netherlands:North-Holland.

Hirshleifer, David, Yaron Levi, Ben Lourie, and Siew Hong Teoh. 2019. “Decision Fatigue and Heuristic Analyst Forecasts.” *Journal of Financial Economics*, 133(1): 83–98.

Huang, Wanying. 2022. “The Emergence of Fads in a Changing World.” Working Paper, California Institute of Technology.

Ifrach, Bar, Costis Maglaras, Marco Scarsini, and Anna Zseleva. 2019. “Bayesian social learning from consumer reviews.” *Operations Research*, 67(5): 1209–1221.

Ikenberry, David L., and James Peter Weston. 2008. “Clustering in US Stock Prices after Decimalisation.” *European Financial Management*, 14(1): 30–54.

Janis, Irving L., and Leon Mann. 1977. *Decision Making: A Psychological Analysis of Conflict, Choice, and Commitment.* New York:Free Press.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow.* New York, NY:Farrar, Straus and Giroux.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590–2615.

Kartik, Navin, SangMok Lee, Tianhao Liu, and Daniel Rappoport. 2022. "Observational Learning with Ordered States." Columbia.

Katz, Leo. 1953. "A new status index derived from sociometric analysis." *Psychometrika*, 18(1): 39–43.

Khanna, Naveen, and Steve L. Slezak. 2000. "The Effect of Organizational Form on Information Aggregation and Project Choice: The Problem of Informational Cascades." *Journal of Economics and Management Strategy*, 9(1): 115–156.

Kirman, Alan. 1993. "Ants, Rationality, and Recruitment." *Quarterly Journal of Economics*, 108(1): 137–156.

Knight, Brian, and Nathan Schiff. 2010. "Momentum and Social Learning in Presidential Primaries." *Journal of Political Economy*, 118(6): 1110–1150.

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

Kreps, David M, Paul Milgrom, John Roberts, and Robert Wilson. 1982. "Rational cooperation in the finitely repeated prisoners' dilemma." *Journal of Economic theory*, 27(2): 245–252.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo. 2013. "Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment." *Quarterly Journal of Economics*, 128(3): 1123–1167.

Kübler, Dorothea, and Georg Weizsäcker. 2003. "Information Cascades on the Labor Market." *Journal of Economics*, 80(3): 211–229.

Kuran, Timur. 1987. "Preference Falsification, Policy Continuity and Collective Conservatism." *Economic Journal*, 97(387): 642–665.

Kuran, Timur. 1989. "Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution." *Public Choice*, 61(1): 41–74.

Kuran, Timur, and Cass R. Sunstein. 2000. "Controlling availability cascades." In *Behavioral Law and Economics*. , ed. Cass R. Sunstein, 374–397. Cambridge, UK:Cambridge University Press.

Lee, In Ho. 1993. "On the convergence of informational cascades." *Journal of Economic Theory*, 61: 395–411.

Lee, In Ho. 1998. "Market Crashes and Informational Avalanches." *Review of Economic Studies*, 65: 741–759.

Le, Tho Ngoc, Vijay Subramanian, and Randall Berry. 2016. "Quantifying the Utility of Noisy Reviews in Stopping Information Cascades." Submitted to the 55rd IEEE Conference on Decision and Control.

Little, Anthony C., Robert P. Burriss, Benedict C. Jones, Lisa M.DeBruine, and Christine A. Caldwell. 2008. "Social influence in human face preference: men and women are influenced more for long-term than short-term attractiveness decisions." *Evolution and Human Behavior*, 29(2): 140–146.

Lobel, Ilan, and Evan Sadler. 2015. "Information diffusion in networks through social learning." *Theoretical Economics*, 10(3): 807–851.

Lohmann, Susanne. 1994. "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91." *World Politics*, 47: 42–101.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies*, 60(3): 531–542.

Monzón, Ignacio, and Michael Rapp. 2014. "Observational learning with position uncertainty." *Journal of Economic Theory*, 154: 375–402.

Moscarini, Giuseppe, Marco Ottaviani, and Lones Smith. 1998. "Social Learning in a Changing World." *Economic Theory*, 11(3): 657–665.

Mossel, Elchanan, Allan Sly, and Omer Tamuz. 2015. "Strategic Learning and the Topology of Social Networks." *Econometrica*, 83(5): 1755–1794.

Mossel, Elchanan, Manuel Mueller-Frank, Allan Sly, and Omer Tamuz. 2020. "Social learning equilibria." *Econometrica*, 88(3): 1235–1267.

Mueller-Frank, Manuel. 2014. "Does one Bayesian make a Difference?" *Journal of Economic Theory*, 154: 423–452.

Mueller-Frank, Manuel, and Mallesh M. Pai. 2016. "Social Learning with Costly Search." *American Economic Journal: Microeconomics*, 8(1): 83–109.

Nekrasov, Alexander, Siew Hong Teoh, and Shijia Wu. 2023. "Limited Attention and Financial Decision Making." In *Handbook of Financial Decision Making*. , ed. Gilles Hilary and David McLean, forthcoming. Edward Elgar Publishing.

Oberholzer-Gee, Felix. 2008. "Nonemployment stigma as rational herding: A field experiment." *Journal of Economic Behavior & Organization*, 65(1): 30–40.

Ottaviani, Marco, and Peter Norman Sørensen. 2001. "Information aggregation in debate: Who should speak first?" *Journal of Public Economy*, 81(3): 393–421.

Ottaviani, Marco, and Peter N. Sørensen. 2000. "Herd behavior and investment: Comment." *American Economic Review*, 90(3): 695–704.

Parikh, Rohit, and Paul Krasucki. 1990. "Communication, consensus, and knowledge." *Journal of Economic Theory*, 52(1): 178–189.

Park, Andreas, and Hamid Sabourian. 2011. "Herding and Contrarian Behavior in Financial Markets." *Econometrica*, 79(4): 973–1026.

Ritter, Jay R., and Ivo Welch. 2002. "A Review of IPO Activity, Pricing, and Allocations." *Journal of Finance*, 57(4): 1795–1828.

Robin, Eugene D. 1984. *Matters of Life and Death: Risks vs. Benefits of Medical Care.* New York:Freeman and Co.

Rob, Rafael. 1991. "Learning and Capacity Expansion under Demand Uncertainty." *Review of Economic Studies*, 58(4): 655–675.

Romano, Maria Grazia. 2007. "Learning, Cascades and Transaction Costs." *Review of Finance*, 11(3): 527–560.

Rosenberg, Dinah, and Nicolas Vieille. 2019. "On the efficiency of social learning." *Econometrica*, 87(6): 2141–2168.

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille. 2009. "Informational externalities and emergence of consensus." *Games and Economic Behavior*, 66(2): 979–994.

Sacerdote, Bruce. 2014. “Experimental and Quasi-Experimental Analysis of Peer Effects: Two Steps Forward?” *Annual Review of Economics*, 6: 253–272.

Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. “Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market.” *Science*, 311(5762): 854–856.

Scharfstein, David S., and Jeremy C. Stein. 1990. “Herd behavior and Investment.” *American Economic Review*, 80: 465–479.

Sgroi, Daniel. 2002. “Optimizing Information in the Herd: Guinea Pigs, Profits, and Welfare.” *Games and Economic Behavior*, 39(1): 137–166.

Sheshinski, Eytan, and Yoram Weiss. 1977. “Inflation and Costs of Price Adjustment.” *Review of Economic Studies*, 44(2): 287–303.

Smith, Lones, and Peter Norman Sørensen. 2000. “Pathological Outcomes of Observational Learning.” *Econometrica*, 68(2): 371–398.

Smith, Lones, and Peter N. Sørensen. 2020. “Rational Social Learning with Random Sampling.” Working paper, University of Wisconsin.

Smith, Lones, Peter Norman Sørensen, and Jianrong Tian. 2021. “Informational herding, optimal experimentation, and contrarianism.” *The Review of Economic Studies*, 88(5): 2527–2554.

Song, Yangbo. 2016. “Social learning with endogenous observation.” *Journal of Economic Theory*, 166: 324–333.

Sorensen, Alan T. 2006. “Social learning and health plan choice.” *RAND Journal of Economics*, 37(4): 929–945.

Sunstein, Cass R. 2009. *A Constitution of Many Minds: Why the Founding Document Doesn’t Mean What It Meant Before.* Princeton, NJ:Princeton University Press.

Sunstein, Cass R. 2019. *Conformity: The Power of Social Influences.* Princeton, NJ:NYU Press.

Talley, Eric. 1999. “Precedential Cascades: An Appraisal.” *Southern California Law Review*, 73: 87–137.

Taylor, Richard. 1979. *Medicine out of Control.* Melbourne:Sun Books.

Tucker, Catherine, Juanjuan Zhang, and Ting Zhu. 2013. “Days on market and home sales.” *RAND Journal of Economics*, 44(2): 337–360.

Veeraraghavan, Senthil K., and Laurens G. Debo. 2011. “Herding in Queues with Waiting Costs: Rationality and Regret.” *Manufacturing & Service Operations Management*, 13(3): 329–346.

Vermeule, Adrian. 2012. *Law and the Limits of Reason.* . Reprint ed., Oxford University Press.

Vives, Xavier. 1993. “How fast do rational agents learn?” *Review of Economic Studies*, 60: 329–347.

Vives, Xavier. 1997. “Learning from others: a welfare analysis.” *Games and Economic Behavior*, 20: 177–200.

Vives, Xavier. 2010. *Information and Learning in Markets: The Impact of Market Microstructure.* Princeton, NJ:Princeton University Press.

Wagner, Peter A. 2018. “Who goes first? Strategic delay under information asymmetry.” *Theoretical Economics*, 13(1): 341–375.

Wang, Tao. 2017. “Information revelation through bunching.” *Games and Economic Behavior*, 102: 568–582.

Weizsäcker, Georg. 2010. “Do we follow others when we should? A simple test of rational expectations.” *American Economic Review*, 100(5): 2340–2360.

Welch, Ivo. 1992. “Sequential Sales, Learning, and Cascades.” *Journal of Finance*, 47(2): 695–732.

Wit, Jorgen. 1997. “Herding behavior in a roll-call voting game.” *Department of Economics, University of Amsterdam*.

Witte, Klaudia, Nina Kniel, and Ilka M. Kureck. 2015. “Mate-choice copying: Status quo and where to go.” *Current Zoology*, 61(6): 1073–1081.

Wolitzky, Alexander. 2018. “Learning from Others’ Outcomes.” *American Economic Review*, 108(10): 2763–2801.

Wu, Jiemai. 2021. “Correlated Information Cascades.” *Available at SSRN 3951879*.

Wu, Jingyi, Cailin O’Connor, and Paul E. Smaldino. 2023. “The Cultural Evolution of Science.” In *The Oxford Handbook of Cultural Evolution.* , ed. Gilles Hilary and David McLean, C78S1–C78N7. Oxford University Press.

Zhang, Jianbo. 1997. “Strategic Delay and the Onset of Investment Cascades.” *RAND Journal of Economics*, 28(1): 188–205.

Zhang, Juanjuan. 2010. "The Sound of Silence: Observational Learning in the U.S. Kidney Market." *Marketing Science*, 29(2): 315–335.

A Online Appendix

§A.1 provides discussion related to Section 6 on limited communication or observation of past actions. §A.2 describes a simple model of belief cascades, relating to the discussion in §10.8. §A.3 discusses legal precedent and information cascades, relating to the discussion in §10.6.

A.1 Further discussion of §6 on limited communication or observation of past actions

In the next two subsections, we discuss imperfectly rational models, as referred to in § 6.2.3, that make direct heuristic assumptions about the agent's mapping from observed actions and payoffs into the agent's actions. Then, in § A.1.3, we discuss a rational setting with partial communication of past signals.

A.1.1 Heuristic action rules based on observation of averages of past payoffs and of action frequency

If rational agents could observe the average payoffs experienced from possible actions undertaken by many in a large population, they would always choose the right action. However, it is possible that agents fall short of rationality. In the model of Ellison and Fudenberg (1993), there is a continuum of agents, where each period some given fraction of them has the opportunity to update their choices. Agents observe the average payoff of the two actions in the population from the previous period, and their popularities. Agents use rules of thumb for updating their actions, which take the form of specified probability of switching action based on these two variables. The paper explores conditions on the weighting of these considerations that promote or hinder long-run correct decisions.

A.1.2 Heuristic action rules based on sampling of past actions and payoffs

In Ellison and Fudenberg (1995), there is a large number of agents who take simultaneous actions at each discrete period. Each period, some fraction of agents exogenously stick to their current action, and the remainder observes an unbiased sample of the latest actions and payoffs of N and choose action based on this and based on their own latest payoff. Agents follow the following heuristic decision rule. If all the agents in an agent's sample takes the same action, the agent follows that action. If both actions are selected by at least one agent in the sample, the agent chooses whichever action has higher average payoff based on observed reports and the agent's own latest experienced payoff.

When the sample size N is small, learning from the experience of others causes the system to evolve to universal adoption of the correct action. In contrast, when the sample size is large, there is strong pressure toward diversity of behavior.[37] So the system never fixes on the correct action. This analysis is valuable in illustrating how non-obvious interesting conclusions about efficiency derive from reasonably plausible heuristic assumptions. On the other hand, when there is a 'split decision' in an agent's sample, it would be reasonable for an agent to take into account that a preponderance of 99 Adopts to 1 Reject (for example) might suggest almost conclusively that adopt is superior.[38] This would tend to oppose the diversity effect discussed above.

A.1.3 Recent past signals observable

The SBM and the BHW model are based on the premise that agents do not directly communicate their private signals. If the signals of agents are fully communicated to their followers, learning would be efficient, and there would be asymptotic learning.

[37] This is because if one action is very unpopular, with large N many adherents of the more popular action "hear about" the unpopular one and potentially convert to it.

[38] Ellison and Fudenberg consider a related effect which they call 'popularity weighting.'

In many practical contexts, more information is indeed transmitted than just action choices. For example, people are often free to talk about their private information, though owing to time and cognitive constraints, such communication may be imperfect. This raises the question of whether incorrect cascades still occur, and whether there is asymptotic learning.

We consider here the possibility of limited communication of private signals. Every private signal is directly communicated, which in principle could bring about asymptotic learning. However, we consider a case in which each signal is passed on to just one other agent, and provide an example in which incorrect cascades still form and last forever.

Consider the SBM of Section 2, except that in addition to observing all past actions, each agent observes the private signal of the agent's immediate predecessor. In contrast with the SBM, we assume that when indifferent, each agent always follows her own signal. This assumption is convenient but not crucial.

After actions HH or LL, I_3 infers that the first two signals were hh or $\ell\ell$ respectively, and falls into a cascade which could easily be incorrect. This cascade can be broken. To see this, suppose that the first two actions are HH, and that I_3 observes ℓ. So I_3 chooses H. If I_4 observes ℓ, this combines with observation of I_3's ℓ to make I_4 indifferent, so I_4 rejects, breaking the cascade.

Consider next the case in which I_3 observes h as well. Now, I_4 adopts, because I_4 knows that the first 3 signals were all h. At this point the cascade of adoption continues forever. I_5 knows that the first two signals were hh, and that the third and fourth signals were not $\ell\ell$. (I.e., they were either ℓh, hh, or $h\ell$.) So the net evidence in favor of state H is at least slightly stronger than two h signals. So even in the worst case in which I_5 directly observes two ℓ signals (I_5's and I_4's signals), I_5 strictly prefers H. Similar reasoning shows that all subsequent agents adopt.

The general point that inefficient cascades can form and last forever does not require this tie-breaking convention.[39] Intuitively, the ability to observe the predecessor's signal is a double-edged blade. The immediate effect is to increase an agent's information, which increases the probability that an agent decides correctly. But to the extent that this is the case, it later makes the action history more informative. That eventually encourages later agents to fall into a cascade, which blocks asymptotic learning.

In particular, the action sequence becomes informative enough that agents fall into cascades despite their access to an extra signal from the predecessor. The actions history eventually overwhelms an agent's information (even inclusive of observation of a predecessor's signal), at which point the accuracy of the social belief stops growing.

Another way to think of this is that being able to observe predecessors' signals is somewhat like being able to observe multiple private signals instead of one. This does not fundamentally change the argument for why incorrect cascades form—that the accumulation of information implicit in past actions must eventually overwhelm the signal(s) of a single agent.

A.2 A Model of Belief Cascades

We provide here a simple model of belief cascades (cascades of reported beliefs), as described informally in §10.8.

In the SBM, agent I_n's belief $\hat{p}_n$ about the state is a continuous variable. Suppose that we extend the SBM by having each agent directly and truthfully communicate a report of her belief (probability assessment) to the next agent. Then each agent's private signal always influences the next agent's belief, so no agent is ever in a cascade of reported beliefs. However, in alternative settings there can be cascades—including

[39]The following modification illustrates the point in a setting with no ties. Suppose that agents differ very slightly in their precisions, and that agents act in inverse order of their precisions, from low to high. Now agents are never indifferent. Agents who are close to indifference strictly prefer to follow their own signals, and the same analysis holds.

incorrect cascades—in reported beliefs. We focus on one possible source of belief cascades, based upon discreteness in the communication channel.

People often communicate in coarse categories, such as binary partitions.[40] Similarly, when asked for a suggestion about a movie or restaurant, or when discussing political topics, people often just name the preferred option, or say that an option is "hot" versus "sucks;" or "cool" versus "bogus."

To capture such limited bandwidth in the communication of beliefs, we consider a setting in which agents make decisions in a commonly known deterministic sequence, and each agent I_m observes the actions of some subset of the agent's predecessors.[41] Suppose that the action of each agent I_n is to report a binary indicator of state, $a_n \in \{L, H\}$.

We interpret the report a_n as the *reported belief* about the state. We will view reported belief L as a claim that the state is probably L, and reported belief H that the state is probably H.[42] This could take the form of expressing that "Candidate H will win the election" or "Candidate H will lose the election." Specifically, I_n *reports* belief L if her true belief $\hat{p}_n < 0.5$, reports H if $\hat{p}_n > 0.5$, and follows some indifference rule, such as flipping a coin between reporting L or H, if $\hat{p}_i = 0.5$. This reporting rule can be endogenized if agents desire a reputation with receivers for making good reports.

As defined in the main text, a *belief cascade* is a situation in which, having received the reported beliefs of some set of predecessors, an agent's reported belief is independent of the agent's private signal. A belief cascade is also an information

[40]As the well-known hedge fund manager and self-improvement author Ray Dalio put it, "It is common for conversations to consist of people sharing their conclusions rather than exploring the reasoning that led to those conclusions."

[41]Alternatively, each agent might also recall and pass on the history of all beliefs that have been reported to that agent. If signals are discrete, such a setting would effectively be equivalent to the model of BHW.

[42]A possible behavioral extension of the model would be to have the reports sometimes understood naively by receivers as indicating that the state is L or H for sure. If, with some probability, receivers make this mistake, information cascades can start very quickly.

cascade, as defined earlier, since the reported belief a_n can be interpreted as the agent's action.

In this model, if signals are bounded, then as in § 4, asymptotic learning can fail; even in the limit, reported and actual beliefs can be incorrect. Furthermore, if signals are also discrete, then for reasons similar to those of the model of BHW, there are incorrect belief cascades (see footnote 41).

In the modified SBM example above, belief cascades derive from communication bandwidth constraints. However, belief cascades may also derive from other possible mechanisms. In models of message sending with payoff interactions, message senders sometimes strategically report only coarsened versions of their beliefs (Crawford and Sobel (1982)). This raises the possibility that sequential reporting might result in cascades, which is an interesting topic for future research.[43]

A.3 Legal Precedent and Information Cascades

A common interpretation of the norm for respect for legal precedent (*stare decisis*), at least since Oliver Wendell Holmes, is that judges acquire information from past decisions. In the models of Talley (1999) and Daughety and Reinganum (1999), this takes the form of imitation of decisions across courts (see also Vermeule (2012)). Daughety and Reinganum (1999) model imitation in which courts at the same level observe

[43]Alternatively, in behavioral models with dual cognitive processing, an agent consists of two selves (sometimes called the "planner" and the "doer") who face distinct decision problems. This can correspond, for example, to the System 2 versus System 1 thinking distinction of Kahneman (2011). Suppose that the planner rationally updates and assesses probabilities, and, when $\hat{p}_i > 0.5$ instructs the doer that the state is H, and, when $\hat{p}_i < 0.5$ instructs the doer that the state is L. Such simplified instructions may be needed owing to limited cognitive processing power of the doer, who may face problems of distraction and time constraint. If people are typically in "doer" mode when engaged in casual conversation with others, then only the coarsened information is reported. This would again result in cascades in reported beliefs. Furthermore, these reported beliefs correspond to the genuine beliefs of people when in doer mode.

private signals about a binary state. This state indicates the single correct decision for a set of related cases to be considered by these courts. (The context is appellate courts that receive related cases.) The courts act in exogenous sequence. If signals are bounded, imitation across the courts results in information cascades. Even if signals are unbounded, there is a high probability of extensive imitation, resulting in poor information aggregation. The authors offer possible examples of actual precedential cascades among appellate courts.[44]

Talley argues that the conditions for cascades to occur in the precedential context are highly restrictive, in part because judges can relay granular information via written opinions. However, owing to limited time and attention, later courts may place heavier weight upon early decisions than upon the nuances of written opinions.

Daughety and Reinganum emphasize that incorrect cascades can be persistent, despite the finding of BHW that cascades are fragile, because of the rarity of shocks that might dislodge a judicial cascade. As discussed in § 2, informative public disclosures can dislodge cascades. However, the authors argue that in the U.S. judicial setting, such shocks take the form of cases being brought to the Supreme Court despite harmonious decisions of lower courts. Such review is rare, because the Supreme Court needs to wait for cases to be appealed, and because decisions that are harmonious across courts are rarely reviewed.

Sunstein (2009) argues that the adoption of legislation sometimes takes the form of information cascades across nations. Sunstein also argues that legal resolution of constitutional questions, such as the once-common, now rejected, view that the U.S. Constitution permits racial segregation, may be the product of information cascades.

[44]To the extent that *stare decisis* is mandated rather than just a description of court behavior, a court may feel pressured to follow precedent rather than doing so for purely informational reasons. However, courts do also deviate from precedent, perhaps because there can be gray areas in the applicability of a precedent. The pressure of *stare decisis* as a social norm presents a challenge for empirical testing for informational effects. At a deeper level, however, social learning may explain why the norm of *stare decisis* originally emerged in common law legal systems.