Topics In Game Theory: Information, etc.
Part 3 - Social Learning

Daniel N. Hauser

In many economic settings, we learn not only by observing information directly, but also by observing information filtered through the decision of others. Before deciding what restaurant to go to, what school to send my kids to, what phone to buy, I look and see what decision others have made in similar situations, and try to incorporate the information contained in their decision into my own decision. This sort of social learning is an important part of economic decision making. In these three lectures, I'm going to provide a brief overview of some of the main concepts, tools, and models of social/observational learning.

Unlike in most of the learning models we've seen in previous lectures, in these models decision makers do not observe all information directly. Instead, they observe decisions that are the result of private information held by other individuals. Throughout these lectures, we focus on settings where individuals are short-lived or myopic, and thus face a severe informational externality. As we've already seen, in many settings optimal learning requires decision makers to carefully balance trade-offs between *exploration* and *exploitation*. Here this trade-off manifests as a trade-off between *imitation* and *communication*. In order to most effectively communicate additional private information, decision makers may need to make suboptimal choices today, to facilitate better choices in the future. But, decision makers naturally do not internalize this externality; they always have the option to follow the lead of others, diminishing the connection between their own private information and their action. We'll show in simple settings how this imitation can lead to relatively extreme failures of information aggregation.

Given these forces, a natural first question to ask is what features of the environment enable or prevent long-run learning. In the first lecture, we'll study the canonical observational learning model (Bikhchandani, Hirshleifer, and Welch 1992; Banerjee 1992; Smith and Sørensen 2000), and characterize if and when decision makers learn to adopt the efficient action in the long run. The focus of this lecture is a characterization of what information structures that enable information aggregation look like, and how failures of aggregation can manifest.

In the second lecture, we'll discuss the speed of learning. How much slower to do agents learn if information is filtered through the actions of others. In this section, we'll go through some basics of large deviation theory and study the social learning model of Harel, Mossel, Strack, and Tamuz (2021).

Finally, we'll think about how these information aggregation results depend on our agents social networks. If decision makers only observe subsets of the population, will this help or hurt learning. In this section, we focus on two specific observation structures which provide insights into the forces driving the more general version of results in the literature. Acemoglu, Dahleh, Lobel, and Ozdaglar (2011) provides a much more general treatment of this than the examples we study here.

This is a large literature, that encompasses a number distinct models, in both Bayesian and non-Bayesian settings. In addition to the models I'll cover in this course, I've also included a non-exhaustive list of other work in this area.

1 Sequential Observational Learning

1.1 A Simple Example

Consider an urn filled with Red and Blue balls.

- There are two states of the world. In the Red state, the proportion of red balls is $\gamma > 1/2$. In the Blue state it is $1 - \gamma$.
- Prior probability $p \in (1/2, \gamma)$ that the state is red.
- Agents sequentially guess whether urn has more red or blue balls.
- Before guessing:
  - Draw a ball with replacement.
  - See all past guesses.

In the long run, will agents learn if there are more red or blue balls?

First agent

- Guess red iff you draw a red ball since

$$Pr(\text{Red State}|1\text{ red}) = \frac{p\gamma}{p\gamma + (1-p)(1-\gamma)} > p > 1/2$$

and

$$Pr(\text{Red State}|1\text{ blue}) = \frac{p(1-\gamma)}{p(1-\gamma)+(1-p)\gamma} < p < 1/2.$$

Next agent

- Inverts first guess, effectively sees two draws w/ replacement
- But now, if the first guess was red, then even after drawing an blue ball

$$Pr(\text{Red State}|1\text{ red, }1\text{ blue}) = \frac{p\gamma(1-\gamma)}{p\gamma(1-\gamma)+(1-p)\gamma(1-\gamma)} = p > 1/2.$$

  and after drawing a red ball

$$Pr(\text{Red State}|2\text{ red}) = \frac{p\gamma^2}{p\gamma^2+(1-p)(1-\gamma)^2} > p > 1/2.$$

- So guesses red no matter what
- Her guess contains no new info.

Everyone else is in the same position, so learning effectively shuts down. We have **herding.** Regardless of an individuals private information, their decision determined solely by the public history. It's easy to see that, in this simple model, if #Blue guesses-#Red guesses≥ 2 or if #Red guesses-#Blue guesses≥ 1, we are in a herd.

1.2 The Baseline Model

- Time is discrete $t = 1, 2, \ldots$.
- There are two states of the world $\theta \in \{B, R\}$.[1] Let $p = Pr(R)$.
- Signal process: $s_t \in \mathbb{R}$ from distribution F_θ with pdf/pmf f_θ, iid conditional on θ.
- Assume signals are:
  - Not perfectly revealing: $\operatorname{supp} F_B = \operatorname{supp} F_R$.
  - Informative: $F_B \neq F_R$.
- At each time t, there is a single agent who chooses an action $a_t \in \{B, R\}$.
- Time t agent observes $(a_1, a_2, \ldots a_{t-1}, s_t)$.

[1] This assumption is consequential. Specifically, generalizing our condition for information aggregation mutatis mutandis to the finite but > 2 state model misses some situations where information aggregates. See Kartik, Lee, and Rappoport (2021).

- Agents have preferences $u(a, \theta) = 1_{\{\theta = a\}}$.[2]

We make the following normalization. Normalize signals so that s satisfies

$$s = \log \frac{f_R(s)}{f_B(s)}.$$

Under this normalization, higher signals are now stronger evidence of state R, and the likelihood ratio satisfies

$$\frac{f_R(s)}{f_B(s)} = e^s.$$

Given a prior p and a signal s, Bayes rule is straightforward

$$\log \frac{Pr(\theta = R|s)}{Pr(\theta = B|s)} = s + \log \frac{p}{1-p}.$$

Throughout these notes, we adopt the convention that $0/0 = 1$. Somewhat sloppily, let $[\underline{L}, \overline{L}] = conv(\text{supp}\, F_R)$. It's easy to see that $\underline{L} < 0 < \overline{L}$.

Definition 1. *We say that signals are* bounded *if* $-\infty < \underline{L} < 0 < \overline{L} < \infty$ *and unbounded if* $\underline{L} = -\infty$ *and* $\overline{L} = \infty$*.*

1.3 Analysis

The belief process that matters here is the public belief,

$$L_t = \log \frac{Pr(\theta = R|a_1, a_2, \ldots a_t)}{Pr(\theta = B|a_1, a_2, \ldots a_t)}.$$

This is a stochastic process. Each agent's decision depends entirely on the realization of their signal s_t and L_{t-1}, as the private belief satisfies

$$L_t^{priv} = \log \frac{Pr(\theta = R|a_1, a_2, \ldots a_{t-1}, s_t)}{Pr(\theta = B|a_1, a_2, \ldots a_{t-1}, s_t)} = L_{t-1} + s_t.$$

Given a public belief L, an agent follows decision rule

$$a(L, s) = \begin{cases} R \text{ if } L + s > 0 \\ B \text{ if } L + s < 0. \end{cases}$$

Thus, the L_t process is Markovian. In each period, it either jumps up or down, and the size and likelihood of updates are fully pinned down by the state dependent probabilities of each action.

[2] This is relatively innocuous as long as the action space is finite, including if we allow for unobserved heterogeneity across agents. Lee (1993) and more recently Ali (2018b) discuss what happens with richer action spaces. Smith and Sørensen (2000) provide a brief treatment of another possible failure of learning that happens if agents disagree on whether B or R is optimal in state R.

Unfortunately, describing the short-run properties of this process, and the corresponding action process, are difficult. Almost all work in this literature focuses on asymptotics. The following result is immediate from what you've seen in part 1

Theorem 1. *Suppose the true state of the world is B. Then there exists a random variable L^* with values contained in $[-\infty, \infty)$ s.t. $L_t \to L^*$ F_B-almost surely. An analogous result holds in state R.*

Proof. Martingale convergence theorem applied to the random variable e^{L_t}. □

1.3.1 The Belief Process

Let's look more carefully at the belief process. Given L_{t-1} we know that if $a_t = R$ then

$$L_t = L_{t-1} + \log \frac{1 - F_R(-L_{t-1})}{1 - F_B(-L_{t-1})}$$

and if $a_t = B$ is

$$L_t = L_{t-1} + \log \frac{F_R(-L_{t-1})}{F_B(-L_{t-1})}.$$

It's easy to see that the log likelihood always moves down after a B action and up after an R action.[3] This is immediate from

$$F_R(L) - F_B(L) = \int_{-\infty}^{L} (e^s - 1) f_B(s)\, ds.$$

The right-hand side is decreasing below $L = 0$, increasing above $L = 0$ and is equal to 0 at $\overline{L}$ and $\underline{L}$, so $F_R(L) - F_B(L) \leq 0$ for all L, strictly so for $L \in (\underline{L}, \overline{L})$.

1.4 Bounded Signals

Theorem 2. *If signals are bounded, then agents never learn the state of the world, i.e. there exists a bounded random variable L s.t. $L_t \to L$ a.s. Moreover $\lim_{t\to\infty} Pr(a_t \neq \theta) > 0$.*

Proof. Suppose that $L_t > -\underline{L}$, then $L_t > 0$ and

$$L_t + s_t > L_t + \underline{L} > 0$$

so the optimal action is R, regardless of the signal, and $L_{t+1} = L_t$. Similarly, if $L_t < -\overline{L}$, then B is optimal after every signal, so $L_{t+1} = L_t$.

Let $U = \sup_L \log \frac{1-F_R(-L)}{1-F_B(-L)}$ and $D = \inf_L \log \frac{F_R(-L)}{F_B(-L)}$. These are both finite (in fact, they are bounded by $\overline{L}$ and $\underline{L}$ respectively), and $D < 0 < U$.

[3] Unsurprisingly given our normalization, F_R and F_B are ordered by stochastic dominance.

Now suppose $L_{t-1} \leq -\underline{L}$. Then by construction $L_t \leq -\underline{L} + U$. A similar argument establishes that $L_t \geq -\overline{L} + D$. So $L_t \leq \max\{L_0, -\underline{L} + U\}$ and $L_t \geq \min\{L_0, -\overline{L} + D\}$.

Therefore, $Pr(\theta = R|a_1, a_2, \ldots a_t)$ converges a.s. to some random variable L^* with bounded support contained (and can't converge a.s. to 1/2). By the dominated and martingale convergence theorems

$$E_R\left(\lim_{t\to\infty} \frac{Pr(\theta = B|a_1, a_2, \ldots a_t)}{Pr(\theta = R|a_1, a_2, \ldots a_t)}\right) = E_R\left(\frac{Pr(\theta = B)}{Pr(\theta = R)}\right)$$

and similarly

$$E_B\left(\lim_{t\to\infty} \frac{Pr(\theta = R|a_1, a_2, \ldots a_t)}{Pr(\theta = B|a_1, a_2, \ldots a_t)}\right) = E_B\left(\frac{Pr(\theta = R)}{Pr(\theta = B)}\right)$$

This immediately implies that with probability bounded away from 0 agents are asymptotically wrong in at least one state. □

When we look at the action and belief processes, there are a number distinct but closely related phenomena here. The language in the literature is a bit imprecise, I draw the following distinction, which roughly follows Smith and Sørensen (2000)

- *A [Information] Cascade:* An event where beliefs converge in finite time
- *A Limit Cascade:* An event where beliefs converge to an interior belief. At the limiting belief an information cascade occurs.
- *A Herd:* An event where at all future times, agents choose the same action.
- *Action Convergence:* An event where the frequency of an action converges to 1

Our theorem establishes that action convergence obtain a.s. It is relatively straightforward to see that limit cascades also obtain. In the finite signal case, information cascades also must arise, but this is not necessarily true in rich signal spaces. Somewhat more subtly, a herd also arises in finite time

Theorem 3. *A herd on some action arises in finite time a.s.*

Proof. Note that if $a_t = R$ then $L_t \geq 0$ and if $a_t = B$ then $L_t \leq 0$. Intuitively, since the time t agent has more information than the time $t + 1$ agent does before the signal is realized, it must be optimal for them to imitate the time t agent before seeing any additional information. Moreover, an information cascade occurs at beliefs above $-\underline{L}$ or below $-\overline{L}$.

Beliefs can only converge to points above $-\underline{L}$ or below $-\overline{L}$ a.s. To see this, note that if any point in $L \in (0, -\underline{L})$ was in the support of our limiting random variable, note that if L_t lies in a $\min(L, -\underline{L})/2$-ball around that point, with probability 1 eventually it either exits that ball or a B action is realized. So convergence to L is impossible. Clearly convergence to 0 is also impossible.

So, beliefs must converge to points in $(-\infty, -\overline{L}] \cup [-\underline{L}, \infty)$. Thus, along almost all sample paths, there exists a time T where for all $t > T$ either $L_t > 0$ or $L_t < 0$. This implies that all actions from time t on are the same. $\square$

1.5 Unbounded Signals

With unbounded signals, this breakdown no longer occurs. Just looking at the belief process, we see that after a B action

$$L_t = L_{t-1} + \log \frac{F_R(-L_{t-1})}{F_B(-L_{t-1})}.$$

Given our relationship between f_B and f_R and the fact that signals are unbounded $\log \frac{F_R(-L_{t-1})}{F_B(-L_{t-1})} < 0$, and in any fixed neighborhood of an interior likelihood, this is uniformly bounded away from 0. The update after an R action is similarly uniformly bounded away from 0 in any neighborhood of any interior likelihood ratio. Thus, beliefs cannot converge to any such neighborhood.[4] So we can conclude that agents learn the state.

Theorem 4. *Suppose $\theta = B$, then if signals are unbounded then almost surely $L_t \to -\infty$. Similarly, if $\theta = R$, $L_t \to \infty$.*

Note that this also implies that agents herd on the correct action in finite time a.s. So decision makers are eventually correct.

Many continuous distributions we use to model information are unbounded (e.g. normals with different means). At least in terms of asymptotics, these results suggest that this assumption is to some extent not innocuous. Do you think this assumption of unbounded signals makes sense? Are there other aspects of this model that contribute to the analysis we've done that, depending on the context, may not make much sense?

2 Rates of Learning

How fast do agents learn? Given our focus on asymptotics, it would be nice to know the rate that we approach the limit, not simply the limit itself. Even in the cases where agents learn the state of the world asymptotically, it seems like a great deal of information is being lost, and information about rates would be a first step in quantifying the magnitude of that loss.

[4] Alternatively the imitation argument in the previous theorem, combined with the observation that both actions always occur with positive probability uniformly bounded away from 0 across a neighborhood of the belief implies this.

2.1 A Brief Introduction to Large Deviations

We've said very little about rates of learning in any setting. Before we return to the more complicated social learning settings, it's going to be useful to think about what, if anything, we can say in the much simpler iid world. Consider the following

- Time is discrete $t = 1, 2, \ldots$
- Two equally likely states of the world $\theta \in \{\theta_0, \theta_1\}$
- Sequence of real valued random variables X_t, drawn iid conditional on the state from distribution F_0 or F_1 (with corresponding pdfs/pmfs).
- In each period, choose an action $a_t \in \{0, 1\}$, with utility $u(a, \theta) = 1_{\{\theta_a = \theta\}}$.

What is the probability that you guess wrong at time t? You've (hopefully) already spent a bunch of time thinking about problems like this. When? What if we rewrote this as

$$H_0 : \theta_0$$
$$H_1 : \theta_1.$$

The test statistic that Bayes rule induces should seem familiar. Each player finds $a_t = 0$ optimal iff

$$Pr(\theta = \theta_0 | X_1, X_2 \ldots X_t) \geq 1/2$$

which, is equivalent to

$$\log \frac{Pr(\theta = \theta_1 | X_1, X_2 \ldots X_t)}{Pr(\theta = \theta_0 | X_1, X_2, \ldots X_t)} \leq 0.$$

Finally, a now familiar consequence of Bayes rule lets us write this as

$$\sum_{i=1}^{t} \log \frac{f_1(X_t)}{f_0(X_t)} \leq 0.$$

What is the probability you're wrong at time t. It's the average of the type 1 and type 2 errors under this specific likelihood ratio test. Our goal is to characterize the rate that this error probability goes to 0. Large deviation theory is a set of results and tools that focus on characterizing the rate that the probability of atypical things goes to 0 in large samples. This is important; even though we know that eventually mistakes disappear, it would be worrying if agents were very confident of something that's very wrong for a very long time. Our goal is going to be, for a given sequence of iid random variables

$X_1, X_2, \ldots$ to characterize a function $\lambda^*(k) : (E(x), \infty) \to \mathbb{R}_+$ s.t.

$$\lim_{t\to\infty} \frac{1}{t} \log Pr\left(\sum_{i=1}^{t} X_i \geq k\right) = -\lambda^*(k).$$

Then, applying this result to our likelihood ratio process will tell us the speed that log-likelihoods concentrate around their mean, giving us an idea of how likely it is that strongly held misconceptions persist.

2.1.1 Chernoff Bounds

The object we're interested in is

$$Pr_0\left(\sum_{i=1}^{t} \log \frac{f_1(X_t)}{f_0(X_t)} > 0\right),$$

that is, the probability of being wrong when the null hypothesis is true. Let $L = \log \frac{f_1(X)}{f_0(X)}$ be the random variable given by X's log-likelihood. We already know one way to bound this. In math camp, we saw that Chebyshev's inequality

$$Pr_0\left(\left|\frac{1}{t}\sum_{i=1}^{t} \log \frac{f_1(X_t)}{f_0(X_t)} - E_0(L)\right| > k\right) \leq \frac{\sigma^2}{tk^2}$$

So if we set $k = -E_0(L)$ then this includes all events where the sum of the log-likelihoods are positive. So, our error probability goes to 0 at least linearly.

At the same time, this is not a great bound. Can we be more clever? Recall Markov's Inequality. For any positive random variable

$$Pr(X \geq k) \leq \frac{E(X)}{k}.$$

We can't apply this directly to our test statistic, since it's not positive. But, for all positive z, we can rewrite things. For any random variable X

$$Pr(X \geq k) = Pr(e^{zX} \geq e^{zk}) \leq \frac{E(e^{zX})}{e^{zk}}.$$

This is a bound on the probability in terms of the moment generating function. So, if the moment generating function is finite, we can construct a pretty tight bound by taking the inf over positive z's

$$Pr(e^{zX} \geq e^{zk}) \leq \inf_{z\geq 0} \frac{E(e^{zX})}{e^{zk}}.$$

This is called a Chernoff bound. It gives us a tool for controlling the probability of tail events using an object we know a lot about.[5] In our specific case

$$Pr\left(\frac{1}{t}\sum_{i=1}^{t}\log\frac{f_1(X_t)}{f_0(X_t)} \geq k\right) = Pr\left(exp\left(z\cdot\left(\sum_{i=1}^{t}\log\frac{f_1(X_t)}{f_0(X_t)}\right)\right) \geq \exp(z\cdot k)\right)$$
$$\leq \left(E(e^{z\log\frac{f_1(X)}{f_0(X)}}e^{-zk})\right)^t$$

So we get the bound

$$\frac{1}{t}\log Pr\left(\frac{1}{t}\sum_{i=1}^{t}\log\frac{f_1(X_t)}{f_0(X_t)} \geq k\right) \leq \inf_{z\geq 0}\log E(e^{z\log\frac{f_1(X)}{f_0(X)}}) - zk.$$

Let $\lambda(z) = \log E(e^{z\log\frac{f_1(X)}{f_0(X)}})$. This is the log of the moment generating function of the likelihood ratio, which is called the cumulant generating function. The object we care about is

$$\lambda^*(k) = -\inf_{z\geq 0}\lambda(z) - zk = \sup_{z\geq 0} zk - \lambda(z).$$

This is the Legendre transform of λ. For now, simply observe that:

- λ^* is strictly convex.
- λ^* is well defined and finite above $E(L)$.
- For each k, the optimal z exists and is the unique solution to $k = \lambda'(z)$ (or 0).
- λ^* is increasing and is 0 at $E(L)$. It is strictly increasing on the set of k's where there exists a z s.t. $\lambda'(z) = k$

To summarize, for any $k > E(L)$

$$Pr\left(\frac{1}{t}\sum\log\frac{f_1(X_t)}{f_0(X_t)} \geq k\right) \leq exp(-\lambda^*(k)t)$$

So our probability of errors decays at least at an exponential rate, which we can characterize using the cumulant generating function.

[5]For another application, a slightly subtle concavity argument gives us Hoeffding's Inequality

$$Pr(|\sum X_n - E(X_n)| \geq k) \leq 2\exp(-2k^2/(\sum_{i=1}^{n}(b_i - a_i)))$$

where b_i and a_i are numbers such that $a_i \leq X_i \leq b_i$ a.s. This and the Borel-Cantelli lemma that we'll see next class immediately give us a strong law of large numbers for bounded random variables, in a way that doesn't involve any 4th order polynomials. If you're desperate to see how to finish the proof for unbounded random variables, I think Steele (2015) provides a very nice, straightforward argument.

2.1.2 Cramer's Theorem

In fact, this form of the Chernoff bound is also in some sense a lower bound, in that we can show the following result (stated here in the context of our problem)

Theorem 5 (Cramer's Theorem)**.** *Fix any $k > L$. Then*

$$\lim_{t\to\infty} \frac{1}{t} \log Pr\left(\frac{1}{t}\sum_{i=1}^{t} \log \frac{f_1(X_t)}{f_0(X_t)} \geq k\right) = -\lambda^*(k)$$

Proof. The Chernoff bound we constructed already tells us that this is bounded above by something that goes to 0 exponentially at rate $\lambda^*(k)$.

So to show this result, we need to construct a lower bound on the error probability. Fix an arbitrary random variable Y and consider the random variable Y^η with density

$$f_\eta(y) = e^{\eta y - \lambda(\eta)} f_Y(y),$$

where $\lambda(\eta) = \log E(e^{\eta Y})$ is the cumulant generating function for Y. This is called the tilted distribution.[6] Note that

$$E(Y^\eta) = \frac{E(Y exp(\eta Y))}{E(exp(\eta Y))} = \lambda'(\eta),$$

So, by choosing η we can choose the mean of this random variable, in a way that relates it closely to our desired bound.

Finally, if we let $\overline{Y}_t = \sum_{i=1}^t Y_i$ and let $\overline{Y}_t^\eta = \sum_{i=1}^t Y_i^\eta$. Then, if we let η be the

[6] Geometrically, we can view the distribution we're constructing as in some sense as looking for the closest random variable that is has expectation above k, in the sense that, for the appropriate choice η, f_η solves the problem

$$\min_{g\in\Delta(\mathbb{R})} E\left(\log \frac{f_Y(s)}{g(s)}\right)$$
$$\text{s.t. } E_g(s) \geq k.$$

It's straightforward for finite random variables, and slightly less straightforward for arbitrary random variables to see that this is solved by the unique tilted distribution with mean k, and $\lambda^*(k)$ gives the value of this program. We're selecting the random variable that's most similar to Y with the desired mean, where similarity is measured using relative entropy. So this is in sort of the minimal cost way of adjusting our random variable so that its mean is above k.

solution to $k = \lambda'(\eta)$ then [7]

$$Pr(\overline{Y}_t \geq tk) = \int_{tk}^{\infty} f_{\overline{Y}_t}(s)\, ds = \int_{tk}^{\infty} e^{t\lambda(\eta) - \eta \sum_{i=1}^{t} y_i} f_{\overline{Y}_t^\eta}(s)\, ds \geq e^{-t\lambda^*(k)} Pr(\overline{Y}_t^\eta \geq tk).$$

This gives us

$$Pr\left(\frac{1}{t}\sum Y_t \geq k\right) \geq e^{-t\lambda^*(k)} Pr\left(\frac{1}{t}\sum Y_t^\eta \geq k\right).$$

The law of large numbers tells us $Pr(\frac{1}{t}\sum Y_t^\eta \geq k)$ converges to 1. So, applying this to our test statistic, for any $k \geq E(L)$

$$\frac{1}{t}\log Pr\left(\frac{1}{t}\sum \log \frac{f_1(X_t)}{f_0(X_t)} \geq k\right) \geq -\lambda^*(k) + o(t)$$

and from our Chernoff bound

$$\frac{1}{t}\log Pr\left(\frac{1}{t}\sum \log \frac{f_1(X_t)}{f_0(X_t)} \geq k\right) \leq -\lambda^*(k) + o(t)$$

So

$$\lim_{t\to\infty} \frac{1}{t}\log Pr\left(\frac{1}{t}\sum \log \frac{f_1(X_t)}{f_0(X_t)} \geq k\right) \to -\lambda^*(k).$$

□

So, not only do our probabilities of errors go to 0, they go to 0 screamingly fast.

It is going to be useful to generalize this bound to sample paths. The following theorem, presented here without proof, follows from theorem 7 of Harel et al. (2021), elegantly reproduced here as theorem 6.

Theorem 6. *Fix a sequence of iid random variables $X_1, X_2, \ldots$. For every k such that $k > E(X_1)$ and any deterministic sequence $(x_t)_{t=1}^{\infty}$ s.t. $\limsup x_t/t = k$ and $P(X_t \geq x_t) > 0$ for all t, it holds that*

$$\lim_{t\to\infty} \frac{1}{t}\log P\left(\bigcap_{\tau=1}^{t}\{X_t \geq x_t\}\right) = -\lambda^*(k).$$

[7] That this relationship holds for the sum is perhaps not immediate. But, recall that the distribution of the sum of two iid random variables is given by

$$f_{Y_1+Y_2}(y) = \int_{-\infty}^{\infty} f(t)f(y-t)\, dt.$$

How the densities of Y_t and Y_t^η relate is pretty direct from there :).

2.2 Social Learning

Can we get similar bounds for social learning. At least in the model we've seen this question only makes sense in the unbounded signal case. I'm going to focus on a different model where we get information aggregation in the long-run under all signal distributions, but before I do, let me briefly summarize some of what we know in the canonical setting with unbounded signals:

- Hann-Caruthers, Martynov, and Tamuz (2018) shows that $\frac{1}{t}\log Pr(a_t \neq \theta) \to 0$, so it goes to 0 sub-exponentially. In the Gaussian case, for any $\varepsilon > 0$ there exists a k s.t. $Pr(a_t \neq \theta) \geq \frac{k}{t^{1+\varepsilon}}$.
- They also show that $L_t/t \to 0$ where L_t is the log-likelihood, and for any r_t s.t. $r_t/t \to 0$ we can find a distribution where $\lim L_t/r_t > 0$. So the log-likelihood grows sublinearly (while it converges linearly with iid signals), but can grow arbitrarily close to linearly for some distribution. In state R, this rate r_t s.t. $L_t/r_t \to 1$ is characterized by the solution to the differential equation $\dot{r}_t = F_B(-r_t)$, so the tail of the distribution of log-likelihoods determines the rate.[8]
- Rosenberg and Vieille (2019) show that in state B the expect time until the first correct choice is finite iff the expected number of mistakes is finite. Both of these conditions are equivalent to the condition that $\int_0^1 1/F(p)\,dp < \infty$, where $F(s)$ is the distribution of $Pr(\theta = R|s)$.

All these conditions effectively say (i) social learning is slower, (ii) it isn't much slower if the signal distribution has fat tails. This should in some sense match our intuition. If a lot of really precise signals are realized, then actions are also going to be pretty informative. We're still losing a lot of information through actions, but if decision makers adopt an incorrect action, the fatter the tails of the signal distribution the faster that mistake will be corrected. So incorrect "herds" break faster with fat tails, while correct herds aren't going to be erroneously abandoned, because very precise signals that go against the state must be incredibly unlikely.

[8]One may be curious how a differential equation shows up here. In state R, we know that eventually

$$L_{t+1} = L_t + \log \frac{1 - F_R(-L_t)}{1 - F_B(-L_t)}.$$

A standard trick for solving such recurrences relations is to approximate them with the corresponding differential equation, moreover for high enough L_t, the adjustment term is approximately $F_B(-L_t)$. So, there's reason to hope that the solution to $\dot{r}_t = F_B(-r_t)$ tells us approximately how beliefs behave in the long-run. These sorts of continuous time approximations lie at the heart of all three of these results.

2.3 Another Model

Let's move away from the unbounded/bounded signal dichotomy and think about social learning in a setting where agents always eventually learn the state.

- Time is discrete, $t = 1, 2, \dots$
- Two states $\theta \in \{B, R\}$, $Pr(\theta = R) = 1/2$.
- There are N myopic decision makers, indexed by $i \in \{1, 2, \dots N\}$.
- In each period each agent draws a signal s_t^i, conditionally iid, normalized to be the log-likelihood, informative but not perfectly, etc.
- Agents then choose an action $a_t^i \in \{B, R\}$.
- Agents observe the history of all actions, and their past signals.

So now, as a decision maker, I have to decide whether to go with my own private information, or to follow the lead of the group. It's immediate that asymptotically our players will learn the state of the world, but a natural question to ask is how much slower do they learn relative to if everyone in the group just disclosed their private information. How much information is lost through the social aspect of this learning problem?

It turns out, quite a bit. We already know that for each i

$$\lim_{t \to \infty} \frac{1}{t} \log Pr(a_t^i \neq \theta) = -n\lambda^*(0)$$

if decision makers could observe every signal. This gives us a natural benchmark. How much slower is the rate of learning in our model than this rate?

Learning in this model is a bit less straightforward, as after $t = 1$ the probabilities of different actions for different agents depend on the entire history of private signal realizations for that agent, and the action vector gives an imperfect signal of the signal vector and to interpret that signal we need to understand how the each other agent is interpreting the action vector given their private signal realization and...

So, let's try to find another approach. Suppose the true state is R. Let

$$G_t = \bigcap_{i=1}^{n} \bigcap_{\tau=1}^{t} \{a_\tau^i = B\},$$

the event that agents have taken the wrong action in every period. Denote i's private likelihood ratio by $L_t^{i,priv} = \sum_{\tau=1}^{t} s_t$. There clearly exists some threshold q_t s.t.

$$G_t = G_{t-1} \cap \bigcap_{i=1}^{n} \{L_t^{i,priv} \leq q_t\}.$$

This threshold q_t is deterministic, and can be written recursively as

$$q_t = -(n-1)\log\frac{P_R(W_{t-1}^1)}{P_B(W_{t-1}^1)}$$

where $W_t^i = \bigcap_{1\le\tau\le t}\{L_t^{i,priv} \le q_t\}$.

So now everything has become a function of the private beliefs. A very complicated function of the private beliefs, but a function of the private beliefs nonetheless. First order stochastic dominance immediately tells us that these q_t's are all positive. Moreover, we know that

$$\lim_{t\to\infty}\frac{1}{t}\log P_R(W_{t-1}^1) = -\lambda_R^*(\liminf(-q_t/t)),$$

where λ_R^* is associated with $-s_t$. So, not only does $\lim q_t/t$ exist, it satisfies the equation

$$\lim_{t\to\infty}\frac{q_t}{t} = \lim_{t\to\infty} -(n-1)\frac{1}{t}\log P_R(W_{t-1}^1) + (n-1)\frac{1}{t}\log P_B(W_{t-1}^1) = (n-1)\lambda_R^*\left(-\lim_{t\to\infty}\frac{q_t}{t}\right)$$

Since in state B this probability goes to 1 and in state R it goes to 0. Thus β must be the unique solution to

$$\beta = (n-1)\lambda_R^*(-\beta).$$

So we're almost done, this β gives the rate

$$\lim_{t\to\infty}\frac{1}{t}\log P_R(G_t) = \lim_{t\to\infty}\frac{n}{t}\log P_R(W_{t-1}^1) = -n\lambda_R^*(-\beta) = -\frac{n}{n-1}\beta.$$

Note that β depends on n. Given our goal, we'd like a lower bound on this rate that was independent in n. Since this rate when agents observe all private information explodes, this would tell us that in large groups we're losing a ton of information. Finally, note that $\frac{n}{n-1}\beta \le E_R(s)$.[9] So we finally get our result:

$$\lim_{t\to\infty}\frac{1}{t}\log P_R(a_t \ne \theta) \ge -E_R(s).$$

[9]Since we know that λ_R^* is strictly convex and equal to 0 at the expected likelihood ratio $-L = E_R(-s)$, and $0 > -\beta > -L$ since λ^* is strictly increasing and 0 at $-L$. Looking at convex combinations, we get

$$\begin{aligned}\lambda_R^*(-\beta) &< \frac{\beta}{L}\cdot 0 + \frac{L-\beta}{L}\lambda^*(0)\\ \frac{1}{n-1}\beta &< \frac{L-\beta}{L}\lambda^*(0)\\ \frac{n}{n-1}\beta &< E_R(s)\end{aligned}$$

where the last line follows from the observation that $\lambda^*(0) < L$.

In contrast, if signals were observable

$$\lim_{t\to\infty} \frac{1}{t} \log P_R(a_t \neq \theta) = -n\lambda_R^*(0).$$

So as the group size grows large, mistakes would vanish much faster if decision makers saw all information. The highly correlated signals we get from social learning in large groups convey orders of magnitude less information, at least as t grows large. This failure is especially interesting because, for any $\varepsilon > 0$, conditional on the even that all agents have always taken the wrong action, there exists an agent who's private belief is with ε of placing probability 1 on the true state. So, conditional on this event, the social information swamps the agents much more precise private information (Of course, probability that every agent has always taken the wrong action is very unlikely as t grows large). Finally, if signals are Gaussian, they show that in each period t, the probability that agents all adopt the action the majority took in period 1 converges to 1 as the number of players go to ∞. So we can always find a group large enough that arbitrarily far in the future, agents prefer to herd on the consensus.

3 Observability

For the final part of the course, let's return to our canonical social learning setting. A natural question to ask here is what would happen if we relax the assumption that each agent observes all their predecessors actions. Acemoglu et al. (2011) provides a fairly general answer to that question. To avoid a bunch of technical complications, I'll focus on a few examples that highlight what can change.

3.1 A little more probability

Before we start, we need a bit more probability. The following theorem is a somewhat subtle, but powerful, implication of the countable additivity of probability.

Theorem 7 (First Borel-Cantelli Lemma)**.** *Let $E_1, E_2, \ldots$ be a collection of events. If*

$$\sum_{i=1}^{\infty} Pr(E_i) < \infty$$

then

$$Pr(\bigcap_{n=1}^{\infty} \bigcup_{k=n}^{\infty} E_k) = 0$$

This says that if we have a set of infinitely many events whose probabilities sum to something finite, then with probability 1 only finitely many of these events occur.[10] This is a very powerful result. Almost all the tools we have for establishing convergence involve constructing bounds on the probabilities of events. These are great for establishing convergence in probability, but clearly are going to give us trouble if we want to show something like almost sure convergence. This gives us a way to turn these statements from statements about convergence in probability to statements about almost sure convergence. If we can show that these probabilities go to 0 really fast, so fast that the sum converges, then we can turn our bounds into statements about probability 0 and 1 events. We can also establish a converse to this theorem

Theorem 8 (Second Borel-Cantelli Lemma)**.** *If $\sum Pr(E_i) = \infty$ and the events are independent, then*

$$Pr(\bigcap_{n=1}^{\infty} \bigcup_{k=n}^{\infty} E_k) = 1$$

So, with the addition of independence, we can also show that if the sum diverges then our sequence of events must occur infinitely often. We've informally made statements along these lines a bunch, statements like if we flip a coin infinitely many times it will eventually generate any finite sequence of heads and tails is basically what this lemma shows.[11]

3.2 Bounded Signals

The fundamental issue we encountered with bounded signals was that the information contained in the history very quickly swamped the information contained in any single signal, so information transmission shut down. What would change if we introduced a small chance that our decision makers just didn't observe any of the previous actions.

[10]The proof of this is pretty brief.

$$\bigcup_{n=1}^{\infty} E_n \supseteq \bigcup_{n=2}^{\infty} E_n \ldots \supseteq \bigcap_{n=1}^{\infty} \bigcup_{k=n}^{\infty} E_k.$$

The continuity property of probability tells us

$$Pr(\bigcap_{n=1}^{\infty} \bigcup_{k=n}^{\infty} E_k) = \lim_{N\to\infty} Pr(\bigcup_{n=N}^{\infty} E_N)$$

and subadditivity tells us that

$$Pr(\bigcup_{n=N}^{\infty} E_N) \leq \sum_{i=N}^{\infty} Pr(E_i)$$

since the sum starting at index 1 is finite and all terms are positive, the RHS must converge to 0.

[11]There are other versions of this lemma that relax the independence assumption in different ways, for instance by showing the appropriate sequence of conditional probabilities is infinite. A reasonably sophisticated probability textbook like Durrett "Probability: Theory and Examples" or the aforementioned Williams "Probability with Martingales" will have a lot on these results.

Assume this was unobserved by others. Specifically, suppose with probability $Q(t)$ a decision maker only observes their private signal, otherwise everything is the same. What can we say about

$$\lim_{t\to\infty} Pr(a_t \neq \theta)?$$

For simplicity, let's assume there are two signals $s_t \in \{\underline{s}, \overline{s}\}$, $\underline{s} < 0 < \overline{s}$ and prior likelihood that lies between them. The fundamental logic here would be unchanged for any bounded signal distribution, but this makes things a bit more straightforward. Given our Borel-Cantelli theorems and our intuition about herds, the following seems like a reasonable conjecture:

$$\lim_{t\to\infty} Pr(a_t \neq \theta) = 0 \iff \sum_{t=1}^{\infty} Q(t) = \infty \text{ and } Q(t) \to 0$$

Turns out, this conjecture is true

Theorem 9.

$$\lim_{t\to\infty} Pr(a_t \neq \theta) = 0 \iff \sum_{t=1}^{\infty} Q(t) = \infty \textit{ and } Q(t) \to 0$$

Proof. Clearly, we need $Q(t) \to 0$ to have any hope here. The Borel-Cantelli lemma tells us that intuitively we "observe" infinitely many signals if the sum condition holds, while we only observe finitely many if it does not. So in one case we should be able to deduce the state, and in the other we shouldn't.

Suppose $\sum Q(t) = \infty$. We know that the history observing types likelihood ratio a.s. converges. We want to show that it can't converge to any interior belief. At any belief we are in one of two situations. Either (i) the history observing types are in a cascade, suppose they are all playing R. Then they know a B action comes from an agent who doesn't observe the history, so their beliefs must update by $\underline{s}$. Therefore, for any belief in the cascade set, beliefs exit a ball around that belief of radius $\min(-\underline{s}, \overline{s})$ with probability 1 (with some minor adjustments if that ball includes beliefs outside of the cascade set). If we aren't in the cascade set, then both types are acting based on their private signal, so beliefs update by $\underline{s}$ after a B action and the same logic applies. Since beliefs converge, the public likelihood ratio must converge to the truth, which means that a.s. in finite time all types that observe the history play the correct action. So $Pr(a_t \neq \theta) \to 0$.

In the other case, note that there exists an N s.t. for any $T > N$ if all actions have been B and we are at time $t > T$ then all players who observe the history find B optimal. Moreover, by the Borel-Cantelli lemma, there exists a T^* s.t. $Pr(\bigcup_{t=T^*}^{\infty}\{\text{Player } t \text{ doesn't observe the history}\}) < 1$. This is independent for the signal process, so with positive probability in state R after time T^* everyone observes the history and before time T^* all realized signals are $\underline{s}$, so all players played $a_t = B$ for all $t \leq T^*$. Thus, players herd on the wrong action with positive probability. $\square$

3.3 Unbounded Signals

It seems intuitive that with unbounded signals decision makers can learn the state of the world even if they observe a fairly small subset of the history. On the other hand, it's also clear that for learning to happen in any of these models, agents need to consistently observe actions that depend on a lot of the history. If, for instance, every decision maker only observed their private signal and the first realized action, there's no way we can get asymptotic learning. No subset of agents can be too influential, as learning requires both contrarian behavior and that behavior to be observed by future agents.

We begin with an example to show how the unbounded signal result can generalize to fairly sparse networks. Suppose that each decision maker only observe their own private signal and the action of the previous agent. It turns out, in this setting, decision makers eventually take the correct action with probability 1. Like in much of what we've done in the last two lectures, beliefs here are a bit of a nightmare. So let's take try an alternative plan of attack

Observation. *It must be that the ex-ante probability of taking the correct action is increasing,* $Pr(a_t = \theta) \geq Pr(a_{t-1} = \theta) \geq 1/2$.

I can always choose to copy the previous agent's action, so I must do weakly better than them in expectation. Note that this tells us that ex-ante the probably of being right converges, but we need more if we want to show it converges to 1. A "natural" path forward would be to try to find a continuous function $Z : [1/2, 1] \to [1/2, 1]$ such that $Z(x) > x$ for any $x < 1$ and $Pr(a_t = \theta) \geq Z(Pr(a_{t-1} = \theta))$. Let's call such a Z an *improvement function*.

Intuitively, it seems like such a function must exist. No matter how far along we get in time, we know that, because of unbounded signals, there's a set of signals where we'd prefer to follow the signal rather than imitate the previous decision maker. Specifically let $U_t = \log \frac{Pr(a_{t-1}=B|\theta=R)}{Pr(a_{t-1}=B|\theta=B)}$ and $D_t = \log \frac{Pr(a_{t-1}=R|\theta=R)}{Pr(a_{t-1}=R|\theta=B)}$. The agents optimal decision satisfies

$$a_t = \begin{cases} R \text{ if } s_t > -U_t \\ B \text{ if } s_t < -D_t \\ a_{t-1} \text{ otherwise.} \end{cases}$$

Each player isn't mimicking the previous agent in one of the first two regions, so if we can show the probability of each of these regions is bounded away from 0 for any interior action probability, then we're in good shape.

Theorem 10. *Define the function* $Z : [1/2, 1] \to [1/2, 1]$ *as*

$$Z(\alpha) = \alpha + \frac{1}{2}\min\{(1 - F_B(\log((2-\alpha)/(1-\alpha)))), F_R(-\log((2-\alpha)/(1-\alpha)))\}.$$

This works.

Proof. The probability that our time t agent is right can be written in terms of the time $t-1$ probabilities

$$\begin{aligned}
Pr(a_t = \theta) &= Pr(\theta = R)(Pr(s_t > -U_t|R) + Pr(a_{t-1} = R|R)(Pr(s_t \in (-D_t, -U_t)|R))) \\
&\quad + Pr(\theta = B)(Pr(s_t < -D_t|B) + Pr(a_{t-1} = B|B)(Pr(s_t \in (-D_t, -U_t)|B))) \\
&= \frac{1}{2}[(1 - Pr(a_{t-1} = R|R))(1 - F_R(-U_t)) + (1 - Pr(a_{t-1} = B|B))(F_B(-D_t)) \\
&\quad + (1 - F_R(-D_t))Pr(a_{t-1} = R|R) + (F_B(-U_t))Pr(a_{t-1} = B|B)].
\end{aligned}$$

The first two terms capture the probability of the previous agent making a mistake and the current agent correcting it. Collect the terms where the $t-1$ agent was choosing B, how much do we improve in that situation

$$\begin{aligned}
&(1 - Pr(a_{t-1} = R|R))(1 - F_R(-U_t)) + (F_B(-U_t))Pr(a_{t-1} = B|B) - Pr(a_{t-1} = B|B) \\
&= Pr(a_{t-1} = B|R)\int_{-U_t}^{\infty} e^s \, dF_B(s) - (1 - F_B(-U_t))Pr(a_{t-1} = B|B) \\
&= Pr(a_{t-1} = B|R)\int_{-U_t}^{\infty} e^s - e^{-U_t} \, dF_B(s) \\
&\geq Pr(a_{t-1} = B|R)\int_{-U_t+\log(k+1)}^{\infty} e^s - e^{-U_t} \, dF_B(s) \\
&\geq kPr(a_{t-1} = B|B)(1 - F_B(-U_t + \log(k+1))).
\end{aligned}$$

for any $k > 0$, where the second line follows from our signal normalization, and the third line follows from the definition of U. Similarly, for any $k > 0$

$$\begin{aligned}
&(1 - Pr(a_{t-1} = B|B))(F_B(-D_t)) + (1 - F_R(-D_t))Pr(a_{t-1} = R|R) \\
&\geq Pr(a_{t-1} = R|R)(1 + kF_R(-D_t - \log(k+1))).
\end{aligned}$$

So

$$\begin{aligned}
&Pr(a_t = \theta) \\
&\geq Pr(a_{t-1} = \theta) + \frac{1}{2}(k_1 Pr(a_{t-1} = B|B)(1 - F_B(-U_t + \log(k_1+1))) \\
&+ k_2 Pr(a_{t-1} = R|R)F_R(-D_t - \log(k_2+1))).
\end{aligned}$$

Let $\alpha = Pr(a_{t-1} = \theta)$. First suppose $Pr(a_{t-1} = B|B) \geq \alpha$. Then we get a lower bound by letting $k_1 = e^{U_t} + 1$, $k_2 = 1 + e^{-D_t}$ of

$$\begin{aligned}
&Pr(a_t = \theta) \\
&\geq \alpha + Pr(a_{t-1} = B)(1 - F_B(\log(1 + 2e^{-U_t}))) + Pr(a_{t-1} = R)F_R(-\log(2e^{D_t} + 1)) \\
&\geq \alpha + Pr(a_{t-1} = B)(1 - F_B(\log((2 - \alpha)/(1 - \alpha)))) \\
&+ Pr(a_{t-1} = R)F_R(-\log((2 - \alpha)/(1 - \alpha))).
\end{aligned}$$

Since one of either $Pr(a_{t-1} = B)$ or $Pr(a_{t-1} = R)$ must be greater or equal to 1/2 we get

$$Pr(a_t = \theta) \geq \alpha + \frac{1}{2}\min\{(1 - F_B(\log((2 - \alpha)/(1 - \alpha)))), F_R(-\log((2 - \alpha)/(1 - \alpha)))\}.$$

The right hand side is our desired function $Z(\cdot)$ (with $Z(1) = 1$, which is the limit). This is continuous for continuous densities.[12] Therefore, since signals are unbounded, unless $Pr(a_{t-1} = \theta) = 1$

$$Pr(a_t = \theta) \geq Z(Pr(a_{t-1} = \theta)) > Pr(a_{t-1} = \theta)$$

and since $Pr(a_t = \theta)$ converges, it must converge to 1. □

What's really going on here? The important ingredients are (i) people further in the future observe decisions made by agents who are more likely to be correct, (ii) agents can, with positive probability, do better than imitating the people they observe, no matter how much information the people they observe have. (ii) is what we need unbounded signals for. (i) comes from our observability assumption, the farther in the future we get, the agents we observe are making decisions that are indirectly informed by the actions of all previous agents. A general version of this property is the following. Let Q_t be a distribution over subsets of $\{1, 2, \ldots t-1\}$. Assume agent t observes the actions taken by agents in the set H drawn from Q_t. This gives us a pretty rich set of possible observation networks to look at.

[12] In fact, we can make it continuous for any density by taking the function

$$\bar{Z}(\alpha) = \alpha + \sup_{[1/2,\alpha]} \frac{1}{2}\min\{(1 - F_B(\log((2 - \alpha)/(1 - \alpha)))), F_R(-\log((2 - \alpha)/(1 - \alpha)))\}$$

and making the observations that if $\bar{Z}(Pr(a_{t-1} = \theta)) > Pr(a_t = \theta)$ then we can find an α' between 0 and 1/2 where $Z(\alpha') > Pr(a_t = \theta)$ and note that our agent can always add noise to the previous observation to make it as-if they observed an agent who was right with probability α'.

Definition 2. $(Q_t)_{t=1}^{\infty}$ *has* expanding observations *if for all* $K \in \mathbb{N}$ *we have*

$$\lim_{t\to\infty} Q_t(\{H : \max_{i\in H} i < K\}) = 0.$$

In a network with expanding observations, the likelihood that agents only observe early actions is vanishing in the long-run. This rules out, for instance, networks where there's a set of agents C whose actions are all the most players observe. Clearly, in both the bounded and unbounded signal case, we need expanding observations to get asymptotic learning. It turns out, using a similar argument to the one we used to for the only previous action model, with unbounded signals expanding observations is all we need.

Theorem 11. *Suppose private beliefs are unbounded and the* Q_t *has expanding observations. Than asymptotic learning occurs in every equilibrium.*

4 More Social Learning Papers

- Networks/other observability structures:
  - Acemoglu et al. (2011)
  - Herrera and Hörner (2013)
  - Monzón and Rapp (2014)
  - Lobel and Sadler (2015)
  - Mossel, Sly, and Tamuz (2015)
  - Arieli and Mueller-Frank (2019)
  - Dasaratha and He (2021)
- Speed of social learning
  - Hann-Caruthers et al. (2018)
  - Rosenberg and Vieille (2019)
  - Harel et al. (2021)
- Changing States
  - Frongillo, Schoenebeck, and Tamuz (2011)
  - Dasaratha, Golub, and Hak (2018)
- Behavioral
  - Demarzo, Vayanos, and Zwiebel (2003)
  - Eyster and Rabin (2010)
  - Guarino and Jehiel (2013)
  - Molavi, Tahbaz-Salehi, and Jadbabaie (2018)
  - Bohren (2016)
  - Frick, Iijima, and Ishii (2020)
  - Bohren and Hauser (2021)
  - Frick, Iijima, and Ishii (2022)
- Search/Info Acquisition
  - Burguet and Vives (2000)
  - Mueller-Frank and Pai (2016)

  - Song (2016)
  - Ali (2018a)
  - Bobkova and Mass (2020)
- Market Microstructure/Financial Markets
  - Avery and Zemsky (1998)
  - Chari and Kehoe (2004)
  - Park and Sabourian (2011)
- Other Twists
  - Mossel, Mueller-Frank, Sly, and Tamuz (2020) (Static Equilibrium concept)
  - Wolitzky (2018) (Observing outcomes and confounded learning)
  - Callander and Hörner (2009) (uncertainty about precisions)
  - Ottaviani and Sørensen (2001) (how agents should be ordered)
  - Eyster, Galeotti, Kartik, and Rabin (2014) (negative congestion externality)
  - Arieli, Koren, and Smorodinsky (Forthcoming) (Pricing)
- Some Design Questions
  - Kremer, Mansour, and Perry (2014)
  - Che and Hörner (2017)

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011): "Bayesian learning in social networks," *Review of Economic Studies*, 78, 1201–1236.

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

——— (2018b): "On the Role of Responsiveness in Rational Herds," *Economic Letters*, 163, 79–82.

Arieli, I., M. Koren, and R. Smorodinsky (Forthcoming): "The Implications of Pricing on Social Learning," *Theoretical Economics*.

Arieli, I. and M. Mueller-Frank (2019): "Multidimensional social learning," *The Review of Economic Studies*, 86, 913–940.

Avery, C. and P. Zemsky (1998): "Multidimensional uncertainty and herd behavior in financial markets," *American economic review*, 724–748.

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

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *The Journal of Political Economy*, 100, 992–1026.

Bobkova, N. and H. Mass (2020): "Learning What Unites or Divides Us: Information Acquisition in Social Learning," *Available at SSRN 3639622*.

Bohren, A. (2016): "Informational Herding with Model Misspecification," *Journal of Economic Theory*, 222–247.

Bohren, J. A. and D. N. Hauser (2021): "Learning with heterogeneous misspecified models: Characterization and robustness," *Econometrica*, 89, 3025–3077.

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

Callander, S. and J. Hörner (2009): "The wisdom of the minority," *Journal of Economic theory*, 144, 1421–1439.

Chari, V. V. and P. J. Kehoe (2004): "Financial crises as herds: overturning the critiques," *Journal of Economic Theory*, 119, 128–150.

Che, Y.-K. and J. Hörner (2017): "Recommender systems as incentives for social learning," *The Quarterly Journal of Economics*.

Dasaratha, K., B. Golub, and N. Hak (2018): "Social learning in a dynamic environment," *arXiv preprint arXiv:1801.02042*.

Dasaratha, K. and K. He (2021): "Aggregative Efficiency of Bayesian Learning in Networks," .

Demarzo, P. M., D. Vayanos, and J. Zwiebel (2003): "Persuasion Bias, Social Influence, and Unidimensional Opinions," *Quarterly Journal of Economics*, 118, 909–968.

Eyster, E., A. Galeotti, N. Kartik, and M. Rabin (2014): "Congested observational learning," *Games and Economic Behavior*, 87, 519–538.

Eyster, E. and M. Rabin (2010): "Naive Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, 2, 221–243.

Frick, M., R. Iijima, and Y. Ishii (2020): "Misinterpreting others and the fragility of social learning," *Econometrica*, 88, 2281–2328.

——— (2022): "Stability and Robustness in Misspecified Learning Models," .

Frongillo, R. M., G. Schoenebeck, and O. Tamuz (2011): "Social learning in a changing world," in *International Workshop on Internet and Network Economics*, Springer, 146–157.

Guarino, A. and P. Jehiel (2013): "Social Learning with Coarse Inference," *Amer-*

ican Economic Journal: Microeconomics, 5, 147–174.

Hann-Caruthers, W., V. V. Martynov, and O. Tamuz (2018): "The speed of sequential asymptotic learning," *Journal of Economic Theory*, 173, 383–409.

Harel, M., E. Mossel, P. Strack, and O. Tamuz (2021): "Rational groupthink," *The Quarterly Journal of Economics*, 136, 621–668.

Herrera, H. and J. Hörner (2013): "Biased social learning," *Games and Economic Behavior*, 80, 131–146.

Kartik, N., S. Lee, and D. Rappoport (2021): "Observational Learning with Ordered States," *arXiv preprint arXiv:2103.02754*.

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

Lee, I. H. (1993): "On the Convergence of Infomational Cascades," *Journal of Economic Theory*, 61, 395–411.

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

Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie (2018): "A Theory of Non-Bayesian Social Learning," *Econometrica*, 86, 445–490.

Monzón, I. and M. Rapp (2014): "Observational learning with position uncertainty," *Journal of Economic Theory*, 154, 375–402.

Mossel, E., M. Mueller-Frank, A. Sly, and O. Tamuz (2020): "Social learning equilibria," *Econometrica*, 88, 1235–1267.

Mossel, E., A. Sly, and O. Tamuz (2015): "Strategic learning and the topology of social networks," *Econometrica*, 83, 1755–1794.

Mueller-Frank, M. and M. M. Pai (2016): "Social learning with costly search," *American Economic Journal: Microeconomics*, 8, 83–109.

Ottaviani, M. and P. Sørensen (2001): "Information aggregation in debate: who should speak first?" *Journal of Public Economics*, 81, 393–421.

Park, A. and H. Sabourian (2011): "Herding and contrarian behavior in financial markets," *Econometrica*, 79, 973–1026.

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

Smith, L. and P. N. Sørensen (2000): "Pathological Outcomes Observational Learning," *Econometrica*, 68, 371–398.

Song, Y. (2016): "Social learning with endogenous observation," *Journal of economic theory*, 166, 324–333.

Steele, J. M. (2015): "Explaining a mysterious maximal inequality—and a path to the law of large numbers," *The American Mathematical Monthly*, 122, 490–494.

Wolitzky, A. (2018): "Learning from Others' Outcomes," *American Economic Review*, 108, 2763–2801.