Social Learning with Heterogeneous Preferences

Pedro Brandão Solti*

November 14, 2022

Abstract

People learn from observing each others' actions, but their interests are not necessarily aligned. This paper analyzes a model of social learning allowing for heterogeneous preferences. There is a sequence of agents, each one observing an arbitrarily informative private signal and the action of their immediate predecessor. In a homogeneous preference environment, this game always leads to agents eventually taking the best action. With heterogeneous preferences, the value extracted from observational learning can be arbitrarily small. I develop an example in which information aggregation fails because agents' actions do not reflect the information of their predecessors. I show that the heterogeneity of behavior when agents are perfectly informed must vanish for there to be asymptotic learning for some signal structure. Similarly, the amount of heterogeneity of behavior under uncertainty must be asymptotically limited by that of a fictitious player in order for there to be asymptotic learning for all unbounded signal structures.

*University of Pennsylvania. Email: solti@sas.upenn.edu. I am deeply indebted to Andrew Postlewaite, George Mailath and Kevin He for their support at every stage of this paper. I also thank Cuimin Ba, Aislinn Bohren, Hanming Fang, Wayne Gao, Alice Gindin, Nawaaz Khalfan, Changhwa Lee, Annie Liang, Guillermo Ordoñez, João Thereze, Rakesh Vohra and audiences at UPenn and PETCO for helpful comments and suggestions.

Index

1. Introduction
2. The Anti-Herding Model
   (a) Homogeneous Preferences Version (Acemoglu et al. (2011))
   (b) Heterogeneous Preferences Version
3. The Model
   (a) Discussion on the Modeling Choices
4. The Gaussian World Example
   (a) Discussion on the Modeling of the Gaussian Example
5. Asymptotic Learning for Some Signal Structure
   (a) Discussion on Theorem 1
6. Asymptotic Learning for All Signal Structures
   (a) Robust Adequate Learning
   (b) Theorem 2: Sufficiency
   (c) Sketch of the Proof for Theorem 2
   (d) Theorem 3: Necessity
   (e) Sketch of the Proof for Theorem 3
   (f) Discussion
7. Prior-Specific Sufficient Conditions
8. Extensions
   (a) Break in Monotonicity
   (b) Heterogeneous Priors
   (c) Observing More than the Immediate Predecessor
9. Conclusion
10. Appendix 1: Proofs for the Gaussian World Example
11. Appendix 2: Proof of Theorem 1
12. Appendix 3: Proof of Theorem 2

13. Appendix 4: Proof of Theorem 3

14. Appendix 5: Proof of Theorem 4

15. Appendix 6: Proof of Proposition 6.

16. Appendix 7: Proofs of Heterogeneous Priors Extension.

17. Appendix 8: Proofs of Networks Extension.

.

1 Introduction

Most economic phenomena involve some sort of information transmission throughout society. Examples include the spreading of technologies in rural parts of the world (such as the Green Revolution), misinformation campaigns in social media or word-of-mouth advertisement of new products. This paper aims to show that the mismatch of preferences among the relevant agents has serious implications for the way in which information spreads.

In particular, the central concern of this paper relates to the process of observational learning, which is the phenomenon of people inferring what information led others to take the action they did, and learn from that exercise. The process of people successfully learning through observations, and transmitting information through their actions to future observes is known as social learning.

There is evidence that observational learning can be a very powerful tool for information to spread, sometimes even more powerful than direct access to information (e.g.: Krishnan and Patnam (2014)). Furthermore, classic positive results in the theoretical literature (Smith and Sørensen (2000), Acemoglu et al. (2011)) reinforce that strength.

Nevertheless, there is also evidence that different forms of heterogeneity in society can hinder that process of information accumulation[1]. For example, Tjernstrom (2017) studies the introduction of new maize technologies in rural Kenya through a randomized control trial, and finds that the variance of farmer's posteriors on the impacts of the new technology increases with the village soil heterogeneity, and that this weakens the strength of social learning.

This paper provides a new explanation for such phenomena. The combination of population heterogeneity and non-observation of the whole set of agents previous decisions can lead to the actions taken no longer reflecting information collected in the past.

This stands in contrast to the most common form of failure of information aggregation found in the literature, herding (Bikhchandani et al. (1992), Banerjee (1992)). In the herding models, failure happens because agents' actions stop reflecting their own private information, and are perfectly predictable given the actions of the previous players. In the anti-herding example, introduced in Section 2 failure happens because players' actions stop revealing past play, and only reflect their private signal.

That example is highly stylized and intends to portray the anti-herding effect in a very sharp way. A natural question that arises is how generalizable this phenomenon is. In particular, Acemoglu et al. (2011) studies a similar version of that model, with homogeneous preferences, and shows that information aggregates well. The reader may wonder whether small amounts of heterogeneity will not change that positive result.

More specifically, what are the conditions on the sequence of preferences such that information aggregates well for *some* (non-trivial) signal structure? And what conditions are sufficient or necessary for information to aggregate well for every (non-trivial, unbounded) signal structure?

Section 5 tackles the first question. Theorem 1 establishes a necessary and sufficient condi-

[1] Oyetunde-usman (2022) provides a good survey of the role of different sources of heterogeneity in technology adoption in East and West Africa.

tion for asymptotic learning to happen for some non-trivial (that is, not fully-revealing) signal structure. That condition states that the sequence of preferences must be such that, in the limit, perfectly informed players would act the same way (up to a relabeling).

Following that, Section 6 asks the analogous question, for *every* non-trivial, unbounded signal structure. Theorem 2 and Theorem 3, roughly speaking, find that it must be the case that there is a fictional agent for which, in the limit, player's preferences becomes arbitrarily close to being a refinement of that fiction players' preferences. This is a more stringent condition than the one for Theorem 1, since it induces similar behavior at uncertainty, whereas the previous condition only imposes that requirement for certainty beliefs.

Not having asymptotic learning does not mean that information aggregation is bad. It can be the case that agents are only ε-away from learning the true state of the world through observational learning, and the quantitative negative impact of heterogeneity of preferences is very small. Theorem 3 shows that this is not the case for at least one information structure. It states that there will be an infinite sequence of agents that will extract arbitrarily little value from observational learning. Their payoffs are very close to the ones they would have if they only observed their private signals. This means that heterogeneity will completely breakdown the social learning process infinitely many times.

The strength of that result depends on the geometry of the preferences. It is possible that some partition of the state space is learnable, but not necessarily the finest possible partition. Section 7 discusses how the geometric position of the prior beliefs plays a role in that process.

This paper also considers some extensions to the baseline model. Section 8.1 highlights another feature that distinguishes heterogeneous and homogeneous preferences: in the latter, agents always benefit from having observing strictly better informed people, but that needs not be true in the former cases. Section 8.2 shows how heterogeneity of priors can be easily incorporated into the baseline model. Finally, Section 8.3 shows that the results are qualitatively robust to more generic networks of observation.

2 The Anti-Herding Model

Suppose there are $\mathbb{N}$ rational agents that play sequentially. Each agent is identified by her position in line.

There are two states of the world, $\Theta = \{L, R\}$, and each agent has two actions available to her, $\{\ell, r\}$. They would like to match the state of the world.

Everyone shares a common prior of $P(\theta = R) = 1/2 + \epsilon$, for a very small $\epsilon > 0$. Furthermore, each agent gets a private, identically distributed, conditionally independent signal. Suppose that with a small probability ϕ, this signal is perfectly revealing, and with the remaining probability $1 - \phi$, the signal is perfectly uninformative.

All agents (except for the first one) get to observe the actions taken by their immediate predecessors, but not by any other player. This assumption is distinct from most of the literature, and its role will be more thoroughly discussed on Section 8.3.

We will consider different specifications for the sequence of preferences.

2.1 Homogeneous Preferences (Acemoglu et al)

Suppose all players share the same payoff function:

$$u_i(a_i, \theta) = \begin{cases} 1, \text{ if} & (a_i, \theta) \in \{(\ell, L), (r, R)\}; \\ 0, \text{ otherwise.} \end{cases}$$

That is, they get a payoff of 1 if they match the state of the world and 0 otherwise. This is a classical payoff structure that was already studied in the seminal papers in this literature, such as Bikhchandani et al. (1992), Banerjee (1992), or Smith and Sørensen (2000). This model has the exact same specification as Acemoglu et al. (2011), and the results in this subsection are present in that paper for a more general signal structure.

Given the payoff structure, agents will take action ℓ if their posterior[2] is lower than 1/2 and r otherwise.

The first player (P1) has only access to her private signal. If she receives the signal that reveals the state to be L, she will take action L. If she receives either the uninformative signal (which leads to a posterior of $1/2 + \epsilon$), or the signal that reveals the state to be R (which leads to a posterior of 1), she will take action r.

The second player (P2) observes the action of the first player. If he sees that P1 took action ℓ, he infers that she was informed that the state is L, and he updates his beliefs to become 0. He will also take action ℓ.

If P2 observes that P1 took action r, he cannot distinguish if she got an uninformative signal or an informative signal saying the state is R. He therefore updates his posteriors towards the state being R: $P(\theta = R \mid a_1 = r) > 1/2 + \epsilon$. In that case, if he get a perfectly revealing signal, he will match the state, otherwise he will take action r.

[2] Posterior refers to the probability of the state being R.

Making an induction argument, it can be concluded that, in equilibrium, if a player received a signal informing the state to be L, then all players in the future will correctly learn the state and take the optimal action. Otherwise, they will take action R and keep updating their posteriors towards 1. Since we know that, with probability one, eventually one agent will get a perfectly informative signal, we know that they will take the correct action in the limit.

In this stylized model, information aggregation is perfect. By observing only their predecessor, an agent will take exactly the same action as if she observed all the past private signals. There is no loss of payoff associated with not observing information directly, or with not observing the whole string of past players.

2.2 Heterogeneous Preferences

Now, suppose that players no longer share the same payoff function. Some players get large penalties for making mistakes when the state of the world is L, and small penalties for mistakes when the state of the world is R, and other players are the opposite. More specifically, for an $\varepsilon \in (0,1/2)$, suppose that for odd player $i \in \{1,3,5,...\}$,

$$u_i(a_i,\theta) = \begin{cases} 1, \text{ if} & (a_i,\theta) \in \{(\ell,L),(r,R)\}; \\ 1-\frac{\gamma}{1-\gamma}, \text{ if} & (a_i,\theta) = (r,L), \\ 0, \text{ if} & (a_i,\theta) = (\ell,R). \end{cases}$$

For even players $i \in \{2,4,6,8,...\}$,

$$u_i(a_i,\theta) = \begin{cases} 1, \text{ if} & (a_i,\theta) \in \{(\ell,L),(r,R)\}; \\ 1-\frac{\gamma}{1-\gamma}, \text{ if} & (a_i,\theta) = (\ell,R), \\ 0, \text{ if} & (a_i,\theta) = (r,L). \end{cases}$$

Those preferences are represented in Figure 1 below. In this case, odd players will optimally take action r when their posteriors are larger than γ, and ℓ otherwise, whereas even players will take action r when their posteriors are larger than $1-\gamma$, and ℓ otherwise.

Just like in the homogeneous preferences case, the first player will take action ℓ if she gets a perfectly informative signal telling her the state of the world is L, and r otherwise.

If P2 observes P1 took action ℓ, he will also take action ℓ. If he observes that she took action r, he will update his belief towards 1. For small values of ϕ or γ[3], this new posterior will be below $1-\gamma$. Therefore, if on top of observing P1 taking action r, P2 gets an uninformative signal, he will take action ℓ. If he observes a perfectly informative signal, he will match the state.

In other words, the only scenario in which P2 takes action r is when he observes a perfectly informative signal revealing the state to be R[4]. This mirrors the behavior of P1, that only takes

[3]More specifically, if $\gamma < \frac{1-2\epsilon(1-\phi)-\phi}{2-(1-2\epsilon)\phi}$.

[4]Note that it cannot be the case that P1 takes action ℓ and P2 receives a signal revealing the state to be R. This happens because P1 only takes action ℓ when the state of the world is L, and in those cases, P2 never receives the signal in question.

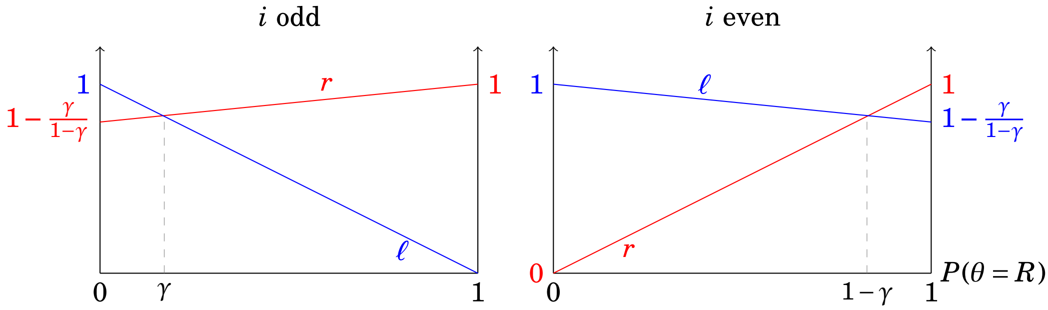

Figure 1: Graphical representation of the payoff functions in the anti-herding game.

action ℓ upon receiving a perfectly revealing signal that confirms the state to be L.

An induction argument therefore can be made that, in equilibrium, odd players take action ℓ upon seeing $s_i = L$ and r otherwise, whereas even players take action r upon seeing $s_i = R$, and ℓ otherwise. Their actions, therefore, only reflect their private signals. They do not reflect any information about the action of their predecessors. Information fails to accumulate, and people are as well off as they would be if they only observed their private signal. Observational learning has zero value.

This phenomenon stands in stark contrast to the herding models, started in Bikhchandani et al. (1992) and Banerjee (1992). In those models, the action of the past players are so informative about the state of the world that all of successors will copy them. Action stop reflection their private signals, which is the exact opposite of what happens in the model in this subsection. This is why this model is dubbed the *Anti-Herding Model*.

The fact that players only observe their immediate predecessor plays an important role here. Since the obstruction to social learning comes from actions not reflecting past play, if agents could observe all previous actions, there would be no obstruction. Eventually agents would take the full-information action. The failure stems from the combination of heterogeneity of preferences and not observing the entirety of the history of actions taken.

A similar model to this was introduced in Lobel and Sadler (2016), but with significant differences. The signal structure generates a full support of posteriors, which allows for some value to observational learning, even if agents' equilibrium payoffs do not converge to the full information one. That paper does not focus on understanding sufficient and necessary conditions for asymptotic learning of a general game, but rather on comparative statics related to the sparsity and homophily of the observational structure of this stylized game.

Even though the payoffs portrayed in Figure 1 are significantly different for odd and even players, there could be a version in which γ becomes arbitrarily close to 1/2, and the payoffs become arbitrarily close to the one in the homogeneous preferences case. As long as ϕ and ϵ are small enough, the analysis and the conclusions would be identical.

This is, of course, a highly stylized model. A natural question is how generalizable this insight is. The next session defines a much more general version of this model, which is analyzed in the subsequent sessions.

3 Model

At each (discrete) time period, a different player, indexed by i, makes a single decision. i denotes both the player and the time period. Player i has a finite set of actions A_i available to her, and a payoff function $u_i : A_i \times \Theta \to \mathbb{R}$, that depends on the action chosen, a_i, and on the true state of the world, $\theta \in \Theta$. It is assumed that there are finitely many states of the world.

Agents can have different utility functions, but they are all uniformly bounded by a certain M. Unlike in Smith and Sørensen (2000), payoff functions are common knowledge, which rules out the confounding effect present in that paper. It is assumed that for all states of the world, there is a uniformly strict best response, such as stated in Assumption 1.

Assumption 1. *For every state of the world $\theta \in \Theta$, there is a uniform strict best response. In other words, there exists an $\varepsilon > 0$, such that for all agents i, and all states θ, there exists an action $a_{i\theta} \in A_i$ for which*

$$u_i(a_{i\theta}, \theta) > u_i(a', \theta) + \varepsilon$$

for all $a' \neq a$.

The true state of the world θ is unknown. All players share a common prior $\mu \in \Delta(\Theta)$[5]. It is assumed that each agent$i$ receives a signal $s_i \in S$, which generically depends on the distribution of states of the world, and on the action of the predecessor.

Each signal is identified as the posterior that it induces on a Bayesian player with prior μ. For example, if $s_i = \delta_\theta$, where δ_θ is the Dirac-distribution that assigns weight 1 to the state θ, then the signal that i received perfectly reveals the true state of the world to be θ. It is also assumed that signals can be arbitrarily precise, which is formalized on Assumption 2.

Assumption 2. *Signals can be arbitrarily precise. That is, for all $\theta \in \Theta$, $\delta_\theta \in supp \mathbb{F}_\theta$.*

Furthermore, for every $\delta > 0$, there exists an $\varepsilon_\delta > 0$ such that, for all $i > I_\delta$, and for all θ,

$$\liminf F_i \left(\{\tilde{\mu} : ||\tilde{\mu} - \delta_\theta|| < \delta\} \right) > \varepsilon_\delta$$

The first condition in Assumption 2 says that signal structures are unbounded. This condition will play the same role as in Smith and Sørensen (2000), ruling out herdings. The second condition ensures that there are no neighborhoods around precise beliefs whose probabilities converge to zero. For example, a scenario in which player i gets a perfectly informative signal with probability $1 - 1/i$ and a perfectly informative signal with probability $1/i$ is ruled out.

Assumption 3 states that no player is perfectly informed just by observing their own signal. This assumption rules out some situations in which the analysis is trivial.

Assumption 3. *The signal structure is non-trivial, that is, it is never perfectly informative.*

On top of the private signals, each agent i observes the action chosen by her predecessor $i - 1$, and no one else's (with the obvious exception of the first player $i = 1$, who does not observe

[5] Section 8.2 shows that the model is easily extendable to a heterogeneous priors model.

anyone's actions). This modeling choice, even if present in Acemoglu et al. (2011), is dissonant from most of the literature, but is not necessary for the results. It is made to significantly simplify notation. A version of the results without this assumption is available on Section 8.3, and they have qualitatively similar ideas to the main versions of the results.

Definition 1. *The vector composed by the space of uncertainty Θ, the sequence of action spaces $\{A_i\}_{i\in\mathbb{N}}$, the sequence of utility functions $\{u_i\}_{i\in\mathbb{N}}$ and a sequence of signal structures F is a* ***game*** *$\mathcal{G}$:*

$$\mathcal{G} =< \Theta, \{A_i\}_{i\in\mathbb{N}}, \{u_i\}_{i\in\mathbb{N}}, \{F\}_{i\in\mathbb{N}} >$$

All players are Bayesian, and in equilibrium they choose their actions in order to maximize their expected payoff given the information that is available to them. Player i's strategy is therefore a mapping of their information into actions, $\sigma_i : S_i \times A_{i-1} \to A_i$, and a strategy profile is a collection of all players' strategies: $\sigma = \times_{i\in\mathbb{N}}\sigma_i$. The equilibrium notion is rationality.

If players could observe the precise information of their predecessors, the accumulation of information would be trivial. It would be akin to the situation in which player i observes all signals revealed thus far. As players get an arbitrarily large number of signals, they should eventually learn the state of the world, and therefore take the correct action in the limit. Information accumulation throughout generations would be seamless.

Nonetheless, observing actions instead of information acts as a filter. Player i cannot distinguish between the pieces of information that would all lead to the same action of $i-1$. The central question in the social learning literature is how much that filtering hinders information accumulation. To make that analysis, there needs to be a benchmark for information accumulation, which is the concept of *asymptotic learning*.

Definition 2. *The model features* adequate learning *if, in equilibrium, agents get arbitrarily close to the perfect information payoff. That is, for every $\theta \in \Theta$,*

$$|u_i(\sigma_i, \theta) - \max_{a\in A_i} u_i(a, \theta)| \to_p 0,$$

where, with a slight abuse of notation, $u_i(\sigma_i, \theta)$ represents the expected payoff induced by strategy σ_i in equilibrium.

3.1 Discussion

Using an argument based on the Martingale Convergence Theorem, Smith and Sørensen (2000) uses the unbounded signal structure assumption is sufficient for adequate learning to be achieved in the case of the whole string of predecessors' actions is observed by each players, when they all share the same two-states, two-actions payoff function.

Acemoglu et al. (2011) takes a step further and proves the same positive result for the case in which only a subset of the predecessors is observed. Under this structure, the Martingale Convergence Theorem is no longer applicable, and they use instead the Improvement Principle. This principle states that all players have a strategy available which is to copy their predecessor.

If they chose it, they would be doing at least as well as their predecessors. Nevertheless, they choose not to, which means that they must be doing better.

It is not surprising that the results from Acemoglu et al. (2011) can be extended to a more general homogeneous preference setting.

Proposition 1. *Suppose preferences are homogeneous, that is, $u_i = u_j \forall i,j$. Then, the model features adequate learning.*

Proof. This is a direct corollary of Theorem 2. □

Even though this result is not surprising, it is an important benchmark to which the heterogeneous preference case should be compared. This paper is dedicated to understanding the role of preference heterogeneity as an obstacle for information aggregation. The next section, the Gaussian World Example, illustrates these channels.

4 Gaussian World Example

The space of uncertainty has two dimensions, $\Theta = \Theta_1 \times \Theta_2 = \mathbb{R}^2$. Player i takes an action $a_i \in \mathbb{R}$.

Agent i's payoff function is quadratic loss onto dimension w_i, that is,

$$u_i(a_i, \theta) = -(a_i - w_i \cdot \theta)^2,$$

where $\cdot$ is the inner product operator, and w_i is a unitary vector that represents the weights that player i can assign to the different dimensions of the state of the world. Throughout this section, vector w_i will be called i's *preference vector*. It is without loss to assume that the first element of w_i is positive: $w_i \cdot (1,0) \geq 0$, and that $w_i \neq (0,-1)$[6], and that $||w_i|| = 1$.

One can think of many plausible stories that would fit this stylized model. For example, suppose there is a portfolio of investments available for agents to buy or sell. There are two dimensions that define the profitability of that portfolio: public health (if there is a raging pandemic), and intensity of climate change (droughts, floods...). In this case, Θ_1 represents the current state of the COVID pandemic, and Θ_2 represents climate conditions.

Positive values for a_i means that agent i decided to be in a long position, whereas agents with negative actions take on a short position.

Agents can have different relative degrees of exposition to COVID and the climate. For instance, some agents are soy bean farmers and some agents work for a virtual conferencing software. Even though both these categories of agents are exposed to both COVID and weather conditions, their relative level of exposition is different. This is capture by each agent's payoff vector w_i.

The state follows a multivariate Gaussian distribution:

$$\theta \sim N(0, I),$$

and each player receives a signal

$$s_i = \theta + \varepsilon_i,$$

where ε is a white noise:

$$\varepsilon_i \sim N(0, \Sigma).$$

It is assumed, without loss, that Σ is a diagonal matrix[7].

All that was described so far is common knowledge to all players, including their payoff vectors.

It is well-known that the best response of player i is to take the action equal to the expected value of $w_i \cdot \theta$. For instance, Figure 2 graphically represents the best response of two different

[6]This is because it is always possible to relable $\tilde{w}_i = -w_i$ and actions $\tilde{a}_i = -a_i$.

[7]It is without loss because it is always possible to rewrite $\Sigma = Q \Lambda Q'$, where Q is an orthonormal matrix and Λ a diagonal matrix. By changing the basis of the problem from the canonical to the one represented by Q, the variance-covariance matrix would thus be Λ, diagonal.

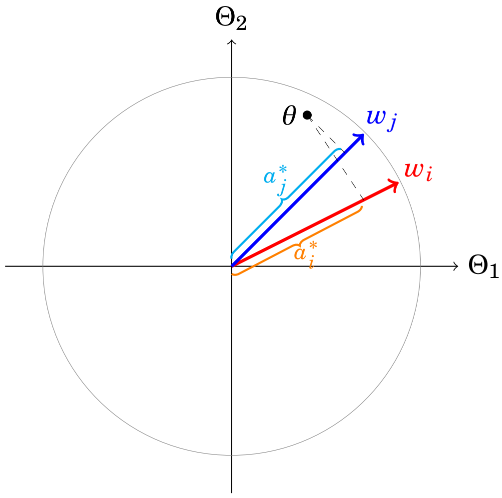

Figure 2: Best Response at Certainty in the Gaussian World.

players when they know that the state of the world is θ. Player i would like to project the point θ onto the dimension defined by the vector w_i, and take the action a_i that reflects that projection. Since players i and j have different preference vectors, they will choose different actions under certainty.

In this model, any sort of lingering heterogeneity in preference vectors will hinder adequate learning, as is made clear by Proposition 2.

Proposition 2. *The Gaussian World Example features asymptotic learning if and only if the preference vectors converge to being collinear, that is,*

$$\lim_{i\to\infty} |w_i \cdot w_{i+1}| = 1$$

Whereas the formal proof can be found in Appendix 1, the intuition for the result can be graphically conveyed. In Figure 3, the red vector represents the preference of the predecessor, and the blue one, that of the successor.

Both players $i-1$ and i care about public health and about the climate, but in different ways. The predecessor assigns a higher weigh to dimension Θ_1 than the successor. That is, she relatively cares more about COVID than the predecessor. Maybe $i-1$ works for an airline company, whereas i is a soybean farmer.

First, let's look at the only if direction. Suppose that adequate learning is achievable, and player $i-1$ is arbitrarily close to full information payoff. For the sake of the argument, suppose he is at the full information payoff. Knowing that, player i will observe $i-1$'s action, and learn that the state of the world lies somewhere along the dark dashed line in Figure 3. He will not be able to distinguish between θ and θ', for example.

We can do a change of basis from reflecting "Covid and Weather" to a new preference-based basis. That is, it is possible to pick a new basis (ϕ_1, ϕ_2) such that the vector w_{i-1} lies on the ϕ_1 axis. In this case, player $i-1$'s action will perfectly reveal the ϕ_1 dimension of the state of the world, but nothing about the ϕ_2 dimension.

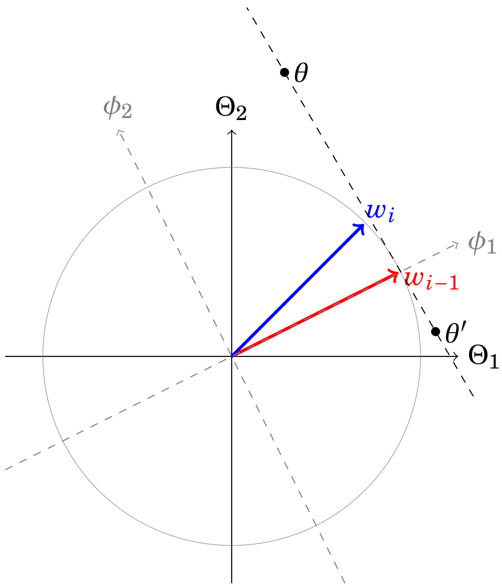

Figure 3: Visual Interpretation of Proposition 2.

Nevertheless, player i cares about both dimensions ϕ_1 and ϕ_2. The variance of the projection of the interim belief of θ onto w_i will not be zero, and will remain bounded away from zero after incorporating the private signal i receives. In other words, player i will not get a payoff close to zero in his best-case scenario (his successor being perfectly informed), and then there cannot be adequate learning as long as the preference vectors do not converge.

This means that, even though the process of information aggregation worked very well until agent $i-1$, part of that information will be lost between periods $i-1$ and i. The heterogeneity of information is associated with that loss of information accumulated throughout generations.

Notice that, were the vector of the successor $\tilde{w}_i = -w_{i-1}$, that problem would not have happened. In that case, player i also would not care about dimension ϕ_2, so that loss of information would not be associated with a lower payoff for agent i. What matters is not that the two players have the exact same preferences, just that they convey the same information.

Also, if i observed the whole string of agents (instead of only the predecessor), he would have access to a much more information. i would then have at least as good information about the ϕ_2 dimension as $i-1$. The "deletion of information" would not happen.

The if direction, that is, proving that there is asymptotic learning if all players share the same preference vector w_i, is simpler. Since in this model all the variables are Gaussian, the posteriors of the players will always be Gaussian. Since there is no uncertainty about the variance, and since the best response of a player will be the expectation of $w \cdot \theta$, a successor can perfectly back out the posterior of the predecessor, which is a sufficient statistic for all the signals received by all players so far. There is not only adequate learning, but the aggregation of information is as good in this model as in a situation in which a player could perfectly observe all signals received by all players before him.

4.1 Discussion on Technical Assumptions of the Gaussian World Example

It is important to remark that the Gaussian World Example has a couple of technical differences to some of the other models considered in this paper.

First of all, it is not the case that players have arbitrarily precise signals. That is, Assumption 2 does not hold. To see this, notice that the variance of the posterior for the players are deterministic in the Gaussian World Example, since all variables follow Normal distributions with known variances. Not only that, they will always be strictly greater than zero (even if they can approach zero in case of asymptotic learning). The posterior variance for any individual player is bounded away from zero.

This is not a problem, though. The only role that Assumption 2 plays in the main results of this paper is to prevent herdings (as an extension of the results in Smith and Sørensen (2000)). Because of the deterministic property of the Gaussian updating of beliefs, herdings are impossible in this scenario (see Appendix 1 for more details), therefore Assumption 2 is dispensable.

Also, for Theorem 2 onwards, it will be assumed that the set of states of the world and that the action space is finite. This assumption will be necessary for the algorithm used in Theorem 2 to converge. The theorems they refer to talk about robust learning, that is, the existence of asymptotic learning for all non-trivial, non-herding signal structures. Since the Gaussian World Example only considers a very specific subset of all possible signal structures, those results do not apply to this case, and therefore the finiteness assumptions are not necessary.

5 Asymptotic Learning for Some Signal Structure

The Gaussian World Example illustrates the impacts of heterogeneity of preferences on information aggregation. Informally, if even in the case of a perfectly informed predecessor, which should be eventually approximated if there is adequate learning, player i is not perfectly able to acquire all information necessary to inform his decision by observing player $i-1$'s action.

The goal of this section is to make the notion of "information necessary to inform his decision" precise in order to get a necessary condition for adequate learning. In other words, it is to find a condition on the sequence of preferences for which we can find a non-trivial (that is, never perfectly informative) unbounded signal structure that leads to asymptotic learning.

Definition 3. *A best-response correspondece for player i is a mapping $BR_i : \Delta(\Theta) \rightrightarrows A_i$ such that*

$$BR_i(\tilde{\mu}) = \arg\max_{a \in A_i} \mathbb{E}_{\tilde{\mu}}[u_i(a,\theta)]$$

Because of Assumption 1, the best response correspondence is always unitary for beliefs that assign probability 1 to a specific state of the world. That is, $|BR_i(\delta_\theta)| = 1$ for all i, θ.

Definition 4. *Given u_i and Θ, player i's* informational content of behavior at certainty *is a partition X_i of Θ such that*

$$BR_i(\delta_\theta) = BR_i(\delta_{\theta'}) \iff \exists x \in X_i \text{such that } \theta, \theta' \in x,$$

where δ_θ represents the Dirac-belief on state θ.

A player's informational content of behavior at certainty is a family of sets. It represents the inverse of the optimal behavior of i if she was perfectly informed. If player i would take action a^* both if she assigned probability 1 to the state of the world being θ and to being θ', then θ and θ' belong to the same element$x$ of X_i.

To make it more concrete, Figure 4 illustrates what the informational content of certainty looks like in the sec:gaussex]Gaussian World Example. If agent i is perfectly informed, she takes the same action for all states that share the same projection onto w_i. The loci of those states are elements of X_i, the set of all lines that are perpendicular to w_i. In Figure 4, each dashed black line is an element of X_i.

The concept of informational content of behavior at certainty represents that intuition of what aspects of the world a player care about. If two states belong to the same $x \in X_i$, then the player that assigns probability 1 to event x has no use to any further information. On the other hand, their behavior cannot be informative about distinctions within x.

These partitions can be more or less similar to each other, and this similarity will impact the usefulness of information transmission of perfectly informed agents. In order to make that intuition clear, we need a function that measures the distance between two informational contents at certainty. That is, a metric on the space of partitions of Θ.

May $\mathscr{P}(\Theta) = 2^\Theta \backslash \emptyset$, that is, the power set of Θ excluding the empty set.

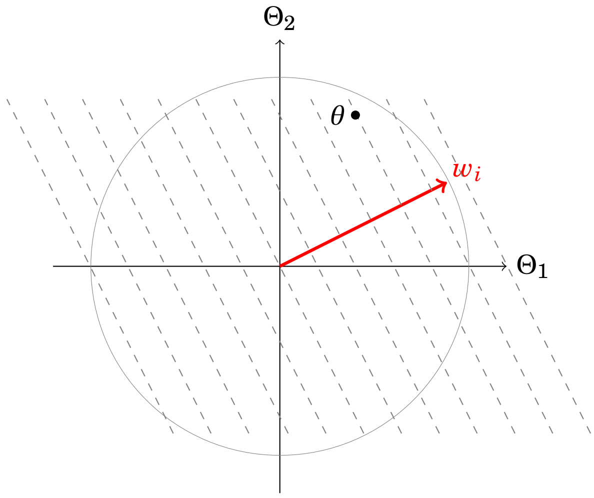

Figure 4: An illustration of the informational content of a player from Gaussian World Example.

Definition 5. *The entropy of a partition X of Θ is a function $H : \mathscr{P}(\Theta) \to \mathbb{R}$ such that*

$$H(X) = -\int_{x \in X} \mu(x) \log \mu(x) dx$$

Definition 6. *The mutual information of two partitions X and Y is a function $I : \mathscr{P}(\Theta)^2 \to \mathbb{R}$ such that*

$$I(X,Y) = \int_{x \in X} \int_{y \in Y} \mu(x,y) \log \frac{\mu(x,y)}{\mu(x)\mu(y)} dy dx$$

Definition 7. *The variation of information between two partitions X and Y is a function VI : $(\mathscr{P}^{\Theta})^2 \to \mathbb{R}$ such that*

$$VI(X,Y) = H(X) + H(Y) - 2I(X,Y)$$

Because the action space and the state space are finite sets, entropy has to be finite, and therefore the Variation of Information is well-defined. For the results in Theorem 1, finiteness of A_i and Θ are not necessary. It is possible to establish the same result, as long as it is assumed that $H(\Theta)$ is finite.

The definitions of entropy and mutual information are standard in the Economics literature, but Variation of Information is a less common concept. It is useful because it is a metric, as proved by Meila (2007).

Proposition 3. *Variation of Information is a metric of the space $\mathscr{P}(\Theta)$.*

Proof. See Meila (2007). □

(Shannon) Entropy is a measure of the amount of uncertainty in a random object, and mutual information measures how much about one random object one can learn from observing another random object. If two random objects are independent of each other, their mutual information is zero. The mutual information between a random object and itself is its entropy.

Variation of Information, as the name suggests, is an information-based measure. It compares the information revealed by two different random objects. This is a natural intuition for

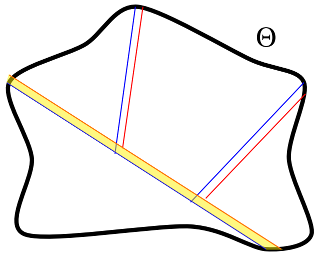

Figure 5: Graphical intuition for the proof of Theorem 1

a social learning setting, where the interactions between agents are purely through information externalities of their own actions.

Variation of Information is not necessarily the only metric that could be used in this setting.

At first glance, a reader could expect that differences in informational content of behavior at certainty are not a barrier to asymptotic learning as long as the signal structure is sufficiently informative. Since Proposition 1 states that, in the homogeneous preferences case, there is asymptotic learning, maybe there is some sort of continuity, and small heterogeneities in preferences can be overcome by more informative private signals, so that asymptotic learning would still be achievable. As Theorem 1 below states, this intuition is false.

Theorem 1. *Fix a sequence of preferences* $\{u_i\}$. *There will be some non-trivial signal structure* $\{F_i\}$ *for which the game* $<\Theta,\{A_i\},\{u_i\},F>$ *has asymptotic learning if and only if the sequence of informational contents at certainty converge, that is,*

$$VI(X_{i-1,}X_i) \to 0.$$

Proof. View Appendix 2. □

The intuition behind the proof is reflected on Figure **??**. In that figure, the amoeba-shaped set represents the set of all states of the world, Θ. The red lines represents the partition of the predecessor, X_{i-1}, and the blue lines represent the partition of the successor, X_i.

Suppose that asymptotic learning has been reached, such that player $i-1$ will always take the correct action. For the sake of the exercise, suppose the successor observes that the predecessor took the action compatible with the lower-left partition. This reveals that the state of the world is in that partition.

If the intersection between the lower-left partition of the successor and the lower-left partition of the predecessor are large enough, then there is a very high probability the predecessor is going to receive a signal that will lead him to take the action compatible with it. But there is always a chance that the true state of the world is in the yellow zone, the "discrepancy area", in which case agent i would be failing to maximize his payoff.

Since this mistake happens with strictly positive probability in the best-case scenario, it cannot be the case that the model features adequate learning.

The intuition for the other direction is trivial. Suppose there is a signal structure for which the induced posterior support is very close to being a full-information support. This signal structure will lead to asymptotic learning.

5.1 Discussion

Theorem 1 establishes a very strong negative result. If there is any lingering heterogeneity in behavior when agents are perfectly informed, then the observational process will not be able to lead to asymptotic learning if the players are not perfectly informed.

This suggests that the traditional results of information aggregation are knife-edge. If one adds a little bit of heterogeneity of preferences at certainty, the result breaks.

Nevertheless, the last statement is only made about the benchmark case of asymptotic learning. It can still be the case that, in the limit, agents' payoffs are going to be very close to the perfect information case, even though it will not converge to it.

Section 6 shows a result that says that, as long as there is some heterogeneity, and it does not have to be just at certainty, it is possible to construct a signal structure for which infinitely many players are going to be arbitrarily close to the payoff they would get if they observed no other players. That is, heterogeneity will completely obstruct social learning for some unbounded signal structure.

6 Asymptotic Learning for All Signal Structures

Theorem 1 gives a necessary and sufficient condition for asymptotic learning to be possible for *some* non-trivial signal structure. A natural question that comes from this statement is whether that condition is necessary and sufficient for asymptotic learning for *every* signal structure.

The answer is no, and the Anti-Herding Example is a counter-example. In that version of the model, all agents would match the state of the world if they were perfectly informed. That is, they all share the same informational content at certainty, $X_i = \{\{L\},\{R\}\}$. Nevertheless, as shown in Section 2, there is no asymptotic learning.

Figure 6 below gives an intuition why information aggregation fails for that case. For all players i, the support of distribution of interim is given by two points: $\{P(\theta = R \mid a_{i-1} = \ell), P(\theta = R \mid a_{i-1} = r)\}$, represented by the black dots in the graph. For all players, these two beliefs belong to the same belief basin, that is, they would lead to the same action. In the contingency that player i gets an uninformative signal, his action will "merge" these two points and thus obfuscate past play.

Those dynamics happen not at the border of the belief space, but rather at the interior. The conditions in Theorem 1 only impose restrictions on optimal behavior at the border of the belief space, and therefore that theorem fails to account for the kind of phenomena present in the Anti-Herding Example.

In other words, agents' actions fail to convey the level of certainty that they have about the state of the world. Since different agents require different levels of certainty to take action ℓ vis-

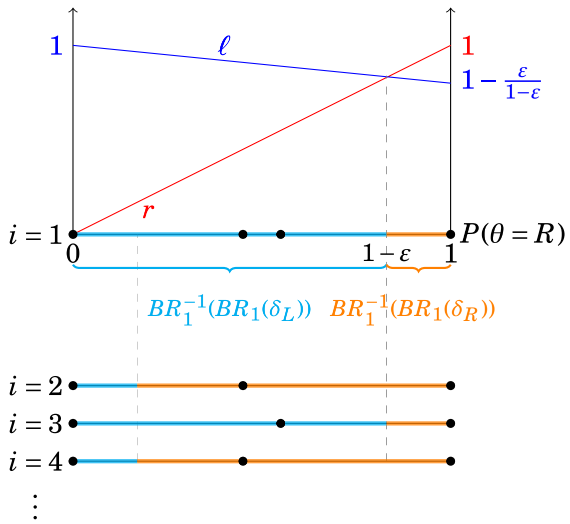

Figure 6: Representation of the Anti-Herding Example.

a-vis action r, observational learning fails to convey an important piece of information to inform decision-making.

6.1 Robust Adequate Learning

The Anti-Herding Example suggests that the level of precision that is conveyed by a player's action plays a big role in ensuring adequate learning. When observing the action of the predecessor, a player can infer what are the possible beliefs that would have led the predecessor to take that action. These beliefs form the *belief basin* of an action.

Definition 8. *The* belief basin of action a for player i *is* $B_i(a) \subseteq \delta(\Theta)$ *such that*

$$B_i(a) = BR_i^{-1}(\{a\}),$$

with BR_i^{-1} *being the upper inverse of the Best Response Correspondence.*

A certainty belief basin *is a belief basin for an* a *such that there exists a* θ *for which* $a = BR_i(\delta_\theta)$*.*

For example, in the Anti-Herding Example, the belief basin for ℓ is $[0, 1-\varepsilon]$ for odd players and $[0, \varepsilon]$ for even players.

When the belief basins do not line up perfectly, there is always chance of "contagion" as in the Anti Herding Example. By assigning more mass on one side of the misalignment or on the other, the signal distribution can induce different actions to the successor that gets an uninformative signal.

In other words, that example suggests that for certain sequences of preferences with some level of heterogeneity that does not vanish, it will be possible to construct information structures for which adequate learning will not be possible.

6.2 Theorem 2: Sufficiency

Definition 9. *A sequence of payoff structures $\{u_i\}$ has* ***robust adequate learning*** *if the game $<\Theta, \{A_i\}_i, \{u_i\}_i, \{F_i\}_i>$ features adequate learning for any signal structure $\{F_i\}_i$ that is unbounded and that does not have vanishing informativeness (as in Assumption 3).*

Theorem 2 below establishes a necessary and sufficient condition on the sequence of payoff functions for robust adequate learning. Before it is stated, it is important to formalize some definitions.

Definition 10. *The family of subsets of $\Delta(\Theta)$, C_i, is a coarsening of i's preferences if:*

- *the closure of its union is the whole belief space:*

$$\overline{\cup_{c \in C_i} c} = \Delta(\Theta), \text{ and}$$

- *for each $c \in C_i$, there exists a subset $A_i^c \subseteq A_i$ for which*
  - *c is formed by agglutinating the belief basins of the elements of A_i^c:*

$$c = \cup_{a \in A_i^c} B_i(a), \text{ and}$$

  - *there are different actions associated with each $c \in C_i$:*

$$c \neq c' \Rightarrow A_i^c \cap A_i^c = \emptyset$$

A coarsening of player i's preferences is convex *if all of its elements are convex sets.*

A coarsening of player i's preferences limited to $E \subseteq \Theta$ is a family of sets $\tilde{C}_i \subseteq \Delta(E)$ for which there exists a coarsening of player i's preferences C_i such that

$$\tilde{c} \in \tilde{C}_i \iff \exists c \in C_i \text{ such that } \tilde{c} = \bar{c} \cap \{\tilde{\mu} \in \Delta(\Theta) : \tilde{\mu}(E) = 1\}.$$

In other words, a coarsening if i's preferences are the preferences of a player $\tilde{i}$ constructed by agglutinating i's belief basins. To better understand the definition, Figure 7 illustrates it.

In the example in that figure, there are three states of the world, $\Theta = \{\theta, \theta', \theta''\}$. The first triangle represent the belief basins of a player i, and the second and third triangles represent two different coarsenings. The gray coarsening, consisting of the gray diamond and the white triangle, is achieved by agglutinating all belief basins except the top one, and the blue coarsening, consisting of the blue and the white sets, is formed by agglutinating all beliefs except the bottom left one.

The fourth triangle is an example of something that is not a coarsening. This is because the bottom middle belief basin intersects both the white and the pink shapes. For a set to be a coarsening, all belief basins of the player must be subsets of an element of the coarsening.

Furthermore, the gray coarsening is also convex, whereas the blue coarsening isn't. This is because the blue set is not convex.

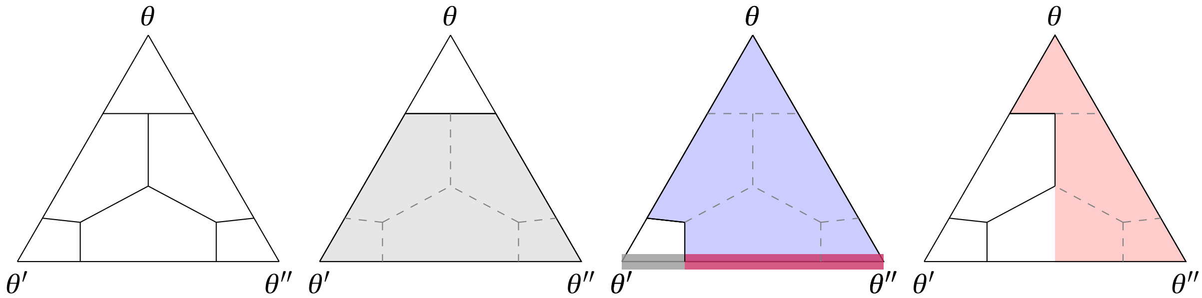

Figure 7: An example of the concept of preference coarsening.

Finally, the green and the orange segments represent the coarsening (c) limited to $E = \{\theta', \theta''\}$. That is, they are both subsets of $\Delta(\{\theta', \theta''\})$. The gray set is obtained by intersecting the closure of the white set with $\Delta(E)$, and the purple set, by intersecting the closure of the orange set with $\Delta(E)$.

An interpretation for a preference coarsening is to create categories for actions. For instance, suppose that there are two restaurants people in a village can go to, an Argentinean (A) restaurant and a Brazilian (B) restaurant. An action could be defined by the pair restaurant-dish ordered. For instance, ordering a *parillada* at the Argentinean restaurant or a *feijoada* at the Brazilian restaurant. The action space would be, for instance, {Argentinean parillada, Argentinean empanada, Brazilian feijoada, Brazilian pão de queijo}. A coarsening would be to aggregate all the dishes from the same restaurant, creating a new, coarser, action space defined by {Argentinean, Brazilian}. In order for there to be some preferences that rationalize the agglutination, it needs to be convex[8].

The convexity requirement can be easier to be met in lower dimensions. For example, in Figure 7, the blue aggregation is not convex. But if agents assign probability zero to state θ, their beliefs will necessarily lie on the bottom edge of the triangle. Restricted to the edge $\{\theta', \theta''\}$, the coarsening becomes the gray and the purple segments, which are convex. So the blue coarsening restricted to $\{\theta', \theta''\}$ is convex, even though the (irrestricted) blue coarsening is not.

Theorem 2 considers the case in which there is a sequence of coarsenings that converges, and the element it converges to is convex. That limit element can be identified with the preferences of a fictitious player. That situation approximates the homogeneous case in which all players share the same preferences as that fictitious limit player. Therefore, the informational content at certainty of that player can be learned, and there will be asymptotic learning.

That process can happen through an iterative fashion though. Consider the stylized example illustrated in Figure 8. There are three states of the world, $\{\theta, \theta', \theta''\}$. Preferences of the even-numbered players are represented by the belief basins drawn on the left triangle, and those of the odd-numbered players, by the ones on the right triangle.

First, it is possible to define a sequence of coarsenings given by { top triangle, gray diamond}

[8]It is not true that all convex partitions of a belief space can be rationalized by a payoff function. Nevertheless, it is true if that convex partition was obtained through a coarsening process of pre-existing preferences. This point is made in more details in Appendix 3.

in Figure 8. These coarsenings agglutinate all belief basins except for the top triangle for all players. Since the sequence is {{ top triangle, gray diamond},{ top triangle, gray diamond},{ top triangle, gray diamond},...}, it is just a repetition of that same element. It trivially converges. The limit of that sequence is that same coarsening { top triangle, gray diamond}. Both the top triangle and the gray diamond are convex sets, so the sequence converges to a convex coarsening. This convex coarsening represents the preferences of a fictitious player that only has access to two actions, and that would optimally take one action if his belief on θ was above a threshold, and the other action otherwise. Theorem 2 says that eventually players are going to be able to learn if the event of the world is $\{\theta\}$ or $\{\theta', \theta''\}$.

Second, notice that there is no sequence of coarsenings that converges to a convex coarsening that has $\delta_{\theta'}$ and $\delta_{\theta''}$ belonging to different elements. If the limit element includes δ_θ, it must include the whole extreme belief basin that it belongs to for all players. Since the limit element is convex, it must include the two diamond-shaped basins that border the aforementioned extreme belief basin for the even players. But that means that it must also include the large central basin for the odd player, and therefore the third diamond-shaped basin of the even player. Again, because of convexity, it must include the extreme belief basin that contains $\delta_{\theta''}$.

That means that one cannot directly use the same argument that we used to say that $\{\theta\}$ can be learned to say the same for $\{\theta'\}$. Nonetheless, it is still the case that θ' will be distinguished from θ''. We know that the state $\{\theta', \theta''\}$ will eventually be learned, which means that, if the true state lies in that event, eventually the beliefs of the players will belong to the bottom edge of the triangles.

Figure 8 projects that bottom edge onto the two colorful segments below the triangles. The segments represent the belief space $\Delta(\{\theta', \theta''\})$, and the colors represent the intersection of the closure of each belief basin with $\Delta(\theta', \theta'')$. We can now use the same procedure from the first part and generate a coarsening that combines the two orange areas and the pink area of the odd players, that is, {{blue}, {orange, pink, orange}, {green}}. For the even players, the coarsening is just the preferences restricted at $\{\theta', \theta''\}$, that is, {{blue}, {pink}, {green}}. These two coarsenings are the same and convex, so there is convergence, and $\delta_{\theta'}$ and $\delta_{\theta''}$ belong to different elements of the limit coarsening. Therefore, players will be able to tell θ' and θ'' apart, and perfectly learn the state of the world.

That intuition is formalized below.

Definition 11. *A sequence of coarsening of preferences $\{C_i\}$ converges to a family of belief basins C if, for each belief basin $b \in C$, there exists a sequence of elements $c_i^b \in C_i$ such that $c_i^b \to b$* [9].

Definition 12. *An event $e \subseteq E \subseteq \Delta(\Theta)$ is separated at E is there exists a preference C and a subsequence of preference coarsenings $\tilde{C}_j$ limited to E such that:*

- *the complement of the subsequence is sparse. That is, may $K \subset \mathbb{N}$ be the set of elements of the original sequence not in the subsequence ($K = \mathbb{N} \backslash Im(j)$). Then, $\lim_{k\to\infty} \frac{|\{\ell \in K : \ell \leq k\}|}{k} = 0$.*
- *$\tilde{C}_i$ converges to C;*

[9] The notion of convergence used here is the one based on the Hausdorff metric.

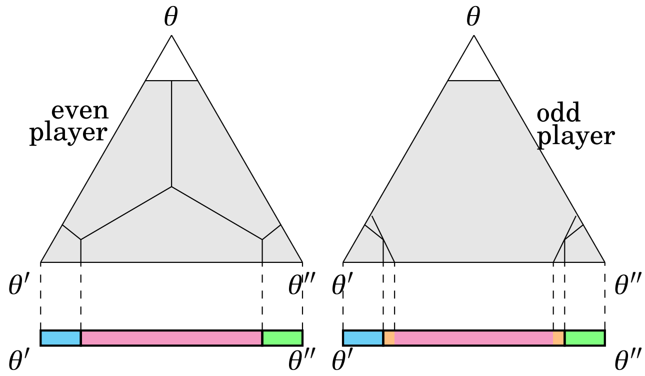

Figure 8: Example of iterative separation.

- *C is convex;*
- *there exists an element $\tilde{c_e} \in C$ such that $\{\tilde{\mu} \in \Delta(\Theta) : \tilde{\mu}(e \cap E) = 1\} \subseteq \tilde{c_e}$ and $\{\tilde{\mu} \in \Delta(\Theta) : \tilde{\mu}(e \cap E) = 0\} \cap \tilde{c_e} = \emptyset$.*

Let's break down the definition above. It starts by assigning probability 1 to event $E \subseteq \Delta(\Theta)$, and finding a sequence of coarsenings $\tilde{C}_i$. This sequence must be equipped with a subsequence that is non-sparse (that is, the share of people that do not appear in that subsequence converges to zero) and that is convergent to a certain convex family of sets C. Furthermore, there must be an element in C that contains all the beliefs that assign probability 1 to state e and none of the beliefs that assign probability 0 to e.

In the first round form example in Figure 8, E is the entire set Θ, the first event to be separated is $e = \{\theta', \theta''\}$, the sequence of coarsenings is {{ white triangle}, {gray diamond}}, that converges to C = {{ white triangle}, {gray diamond}}. All the beliefs that assign probability 1 to event $\{\theta', \theta''\}$, that is, the bottom edge, are contained in the gray diamond; and all the beliefs that assign probability zero to $\{\theta', \theta''\}$, that is, $\{\delta_\theta\}$ belong to the white triangle. Therefore, $\{\theta', \theta''\}$ is separated at Θ. By the same argument, we can say that its complement, $\{\theta\}$, is also separated at Θ.

For the second round, $E = \{\theta', \theta''\}$, $e = \{\theta'\}$, the sequence is {{blue}, {pink}, {green}, {blue}, {orange, pink, orange}, {green}} infinitely-many times, that converges to a limit C. All the beliefs that assign probability 1 to $\{\theta'\}$ are in the blue set, that belongs to C, and all the beliefs that assign probability 0 to $\{\theta'\}$ are in the green set, that belongs to C as well. Therefore, $\{\theta'\}$ is separated at $\{\theta', \theta''\}$.

Proposition 4. *Suppose e is separated at E. Then, there exists a partition P of E such that:*

- *e belongs to P, and*
- *all elements of P can be separated at E.*

Proof. See Appendix 3. □

In the example discussed, all elements of the partition $\{\{\theta\}, \{\theta', \theta''\}\}$ can be separated at $\{\theta, \theta', \theta''\}$, and all elements of the partition $\{\{\theta'\}, \{\theta''\}\}$ can be separated at $\{\theta', \theta''\}$.

Definition 13. *Partition Y of Θ is the* learnable partition *for the game $\mathscr{G}$ if it is the outcome of the* separation algorithm*, which is described below:*

1. *set $t = 0$ and $S_0 = \{\Theta\}$;*

2. *Consider each element $E \in S_t$, and find the finest partition of E such that all of its elements can be separated at E. Call this partition P_E^t.*

3. *Set $S_{t+1} = \cup_{E \in S_t} P_E^t$. ($S_{t+1}$ is itself a partition of Θ)*

4. *Stop if $S_{t+1} = S_t$. Otherwise, advance one step on t and return to step 2.*

Since Θ is finite, the algorithm above has to stop after at most $2^{|\Theta|}$ steps, and therefore the learnable partition is well-defined. Also, E is always separable at E, since we can always take the coarsening $\{\Delta(\Theta)\}$ for all players.

To reinforce those definitions, the separating algorithm would work as following for the discussed example from Figure 8:

1. $t = 0$ and $S_0 = \{\Theta\}$.

2. $P_\Theta^0 = \{\{\theta\}, \{\theta', \theta''\}\}$, since each element in that partition is separable at Θ.

3. $S_1 = \{\{\theta\}, \{\theta', \theta''\}\}$.

4. $S_0 \neq S_1$, so we advance one period and now $t = 1$.

5. The only partition possible for $\{\theta\}$ is itself, so $P_{\{\theta\}}^1 = \{\theta\}$. We know that all elements of $\{\{\theta'\}, \{\theta''\}$ are separable at $\{\theta', \theta''\}$, therefore $P_{\{\theta', \theta''\}}^1 = \{\{\theta'\}, \{\theta''\}\}$.

6. $S_2 = P_{\{\theta\}}^1 \cup P_{\{\theta', \theta''\}}^1 = \{\{\theta\}\} \cup \{\{\theta'\}, \{\theta''\}\} = \{\{\theta\}, \{\theta'\}, \{\theta''\}\}$.

7. $S_2 \neq S_1$, therefore we advance to $t = 2$.

8. The only partition possible for each element of S_2 is itself, therefore $P_{\{\theta\}}^2 = \{\theta\}$, $P_{\{\theta'\}}^2 = \{\theta'\}$, and $P_{\{\theta''\}}^2 = \{\theta''\}$.

9. $S_3 = P_{\{\{\theta\}\}}^2 \cup P_{\{\{\theta'\}\}}^2 \cup P_{\{\{\theta''\}\}}^2 = \{\{\theta\}, \{\theta'\}, \{\theta''\}\}$.

10. $S_3 = S_2$, therefore the algorithm stops. The learnable partition is $\{\{\theta\}, \{\theta'\}, \{\theta''\}\}$.

Even though the specific path the algorithm runs is not unique, the learnable partition that is its outcome is.

Proposition 5. *The learnable partition is unique.*

Proof. See Appendix 3. □

Theorem 2 below states a sufficient condition for robust asymptotic learning. First of all, the sequence of informational contents at certainty must converge: $X_i \to X$; otherwise Theorem 1 says that there will be no asymptotic learning. In addition to that, X being the learnable partition is a sufficient condition for robust asymptotic learning.

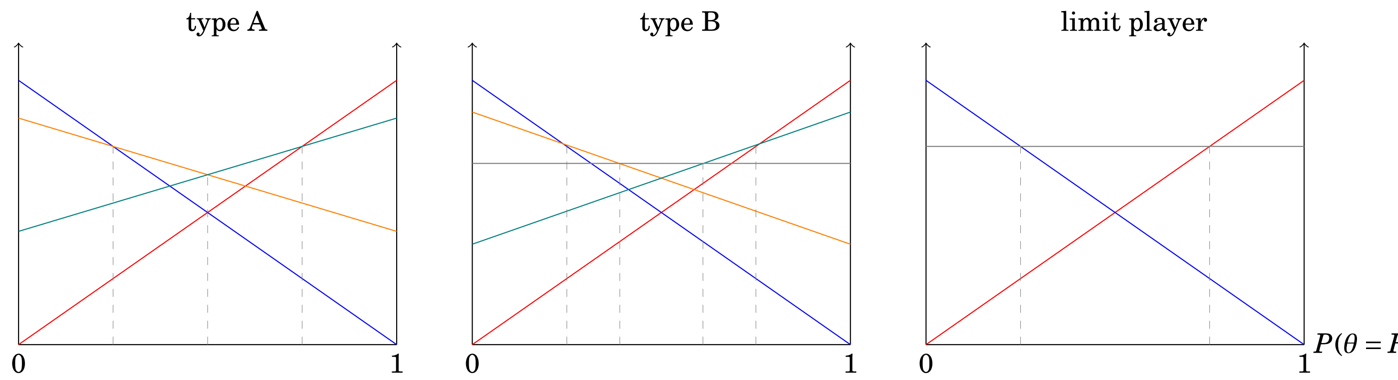

Figure 9: Example for the construction of the proof.

Theorem 2. *Fix the sequence of preferences* $\{u_i\}_i$. *Suppose that the sequence of informational contents at certainty converges:* $X_n \to X$.

If the learnable partition is X*, then there is adequate learning for all unbounded information structures.*

The proof is in Appendix 3, and a sketch of the proof using an example in Section 6.3.

An immediate corollary is that this result implies robust adequate learning for the homogeneous preferences case, generalizing the result in Acemoglu et al. (2011).

Corollary 1. *If all players share the same preferences, and signal structure is unbounded, then there will be adequate learning.*

Proof. Take the sequence of coarsenings of preferences to be exactly the sequence of preferences. Then, it trivially converges and it separates all states that are in different belief basins. Therefore, it must produce X as the learnable partition. □

6.3 Sketch of the proof of Theorem 2

This subsection discusses an example that illustrates the sufficiency direction of Theorem 2: if X is the learnable partition and the information structure is unbounded, then there is adequate learning. The full proof in Appendix 3 formalizes and generalizes the intuition conveyed in this subsection.

Consider a game in which there are two states of the world, and two types of players, with their preferences described in Figure 9. Type A has four actions available to her: blue, orange, green and red; whereas type B has 5: blue, orange, green, red and black.

The relevant feature of this example is that it has the two extreme belief basins (corresponding to the blue and the red actions) coinciding, but the remaining ones do not coincide. In fact, type A has 4 belief basins and type B has 5.

Consider the coarsening in which we agglutinate all non-extreme belief basins. We can construct a fictitious player, called "limit player", that has preferences described by that coarsening,

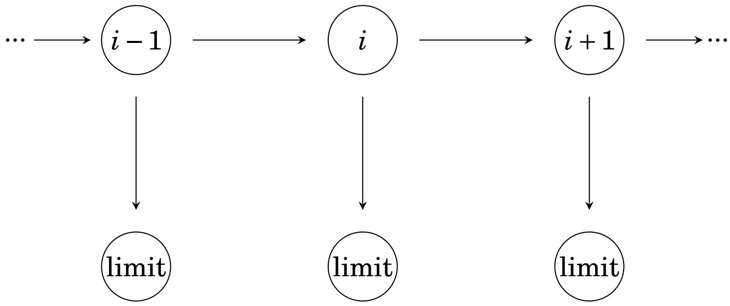

Figure 10: Observational structure of the altered version of the game.

just as in Figure 9[10].

We now construct a slightly altered version of the original game. There is the original sequence of players, each one observing its predecessor and receiving a private signal just like before. On top of that, there is a sequence of homogeneous players, all sharing the same preferences as the limit player. Each one of those limit players observes a censored version of only one of the original players, and they do not get any extra information besides that. This structure is made explicit in Figure10, where $(x \to y)$ represents that player y observes some signal measurable on player x's action.

The goal of this proof is to show that the sequence of payoffs of the limit players will be converging to that of a fully informed player. If that happens, it is because the actions of the original players must be getting arbitrarily informative about the state of the world, and therefore they must also be getting payoffs close to the full information one.

The limit players do not fully observe the action of the original player they share a period with, but rather a censored version of it. If in the sequence of coarsenings used in the separation algorithm, some belief basins of the original player were merged, then the limit player cannot distinguish between the corresponding actions. In this example, the limit player observes if the agents took the blue (extreme left) action, the red (extreme right) action, or one of the intermediate actions, but is unable to distinguish between any of the intermediate actions.

The limit players are sophisticated. They understand the structure of the game, and they not only know the utility functions of all agents, but they also know they are common knowledge. They understand the equilibrium structure, and they can derive the joint distribution of the posteriors of the agent he is observing and of the state of the world.

Using Bayes' rule, once a limit player observes an action, he updates his beliefs to become the average posterior of the player observed conditional on the belief basin of that action. Since he doesn't observe anything else, that will be his posterior. Figure 11 represents a limit player observing player $i-1$. The segment on the top represents the coarsening of the preferences of $i-1$, and the dots the average belief on each of those sets. The limit player "inherits" those beliefs as his posteriors and takes his preferred action. It is possible to calculate that limit player's expected payoff based on that.

[10]There are many preferences compatible with a single coarsening. Any one of those preferences will serve the purpose of the proof.

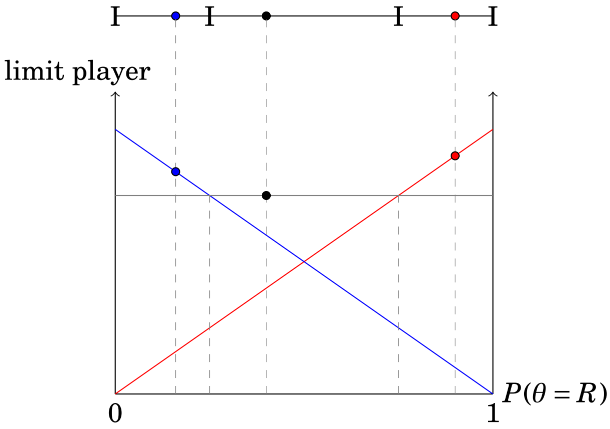

Figure 11: A limit player observing player $i-1$.

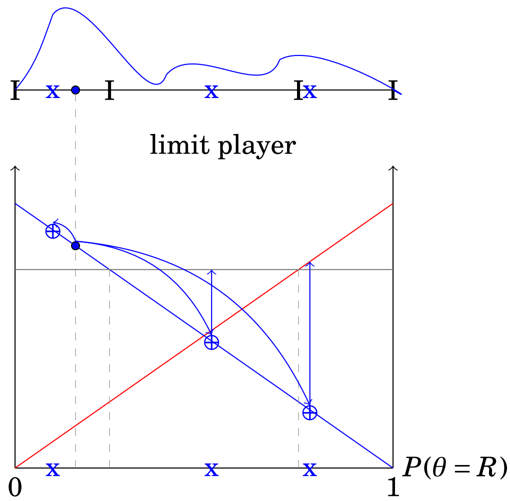

Figure 12: Limit player's payoff from observing player i.

Now, if i observes $i-1$, she will form her interim beliefs the same way the limit player formed his (bu being able to distinguish between the intermediary actions). But on top of that, she is receiving an (unbounded) signal. Therefore, it is possible to calculate a distribution of posteriors for i conditional on the action of $i-1$. This is represented in Figure 12. More specifically, the segment on the top represents the distribution of posteriors for i conditional on $i-1$ having chosen the extreme left action. The mean of that distribution is the blue point. The blue "x" symbols represent the average posterior of agent i conditional on both the action of player i (either extreme left, the average between the intermediary actions, or the extreme right), and the action of $i-1$ being extreme left: $P(\theta \mid a_i, a_{i-1} = \text{blue})$.

Now, suppose that the previous limit player, who knows that $i-1$ took the left action, is able to also see i's action. Then, his beliefs move from its original point (the blue circle) to one of the blue "x" points. If he is not allowed to change his actions after updating the beliefs, then his payoff cannot change. Graphically, the payoff given by the blue dot is the average of the three payoffs given by the blue $\oplus$ signs.

But if the limit player is allowed to switch his actions, then he would match the (censored) action of player i, and there would be an associated gain of payoff given by the vertical blue arrow.

We can make the same argument for the limit player that observes any of the other actions of $i-1$. Notice also that the average improved payoff is the same payoff of the limit player that only observes i, since they always repeat the action of player i. In other words, a limit player is better off observing i than observing $i-1$.

By creating this sequence of homogeneous players, we establish a version of the "improvement principle"[11] from the homogeneous preferences literature.

Furthermore, since the signals are unbounded, whenever the beliefs after observing $i-1$ do not assign probability 1 to a specific state, there is the possibility to receive an arbitrarily informative signal that will lead i (and therefore the limit player that observes i) to a different belief basin. Those switches are associated with a strictly positive increase in expected payoff. Based on that, an argument can be made that the sequence of expected payoffs of the limit players is monotonically increasing, and it converges to the highest possible payoff, the full information one.

We are not interested in the expected payoff of that fictitious limit player, but rather in the expected payoff of the original sequence of players. But notice that, if the sequence of limit players can get arbitrarily close to taking the correct action by observing a censored version of player i, so can player $i+1$, who is strictly better informed than the limit player. In other words, if there is asymptotic learning for the fictitious limit players, there will also be for the original players, establishing therefore the result.

6.4 Theorem 3: Necessity

Theorem 2 establishes a necessary condition for robust asymptotic learning. Theorem 3 below not only states that that same condition os necessary, but also quantifies how far away from asymptotic learning players could be. It says that there will be an unbounded signal structure for which infinitely-many players are going to be as close to the payoff they would get from observing only their private signal as possible.

Definition 14. *A subsequence $\{j\}$ of $\{i\}$ is non-sparse if the fraction of elements missing vanishes. That is, if*

$$\lim_{K\to\infty}\frac{1}{K}|\{k\in\{i\}:k\not\in\{i\},k<K\}|=0$$

Theorem 3. *Pick any $\varepsilon>0$. Pick a sequence of payoff structures $\{u_i\}$, such that the sequence informational contents at certainty induced by $\{u_i\}$ converges: $X_i\to X$. Suppose also that there is no non-sparse subsequence of $\{u_i\}$ that induces a learnable partition of X.*

Then, there will be a prior and a non-trivial, unbounded signal structure for which the game associated with it has a subsequence of agents whose equilibrium expected payoffs are ε-away from the expected payoff they would get had they only observed their private signal, and a signal perfectly revealing the events in the learnable partition.

[11] See Golub and Sadler (2016).

In other words, social learning will breakdown infinitely often. There will always be a player in the future that will extract almost zero value from observing his predecessor. Heterogeneity of preferences, in the sense that it does not allow the separation algorithm to produce a learnable partition of X, can be arbitrarily costly.

Of course, a consequence of Theorem 3 is that the learnable partition being X is a necessary condition for robust asymptotic learning. Combining this with Theorem 2, we get Theorem 2.

Corollary 2. *Suppose $X_i \to X$. There is robust asymptotic learning if and only if the learnable partition is X.*

The non-sparsity assumption rules out the situation in which heterogeneity disappears, except for increasingly rare "crazy" players. The periods of (almost-)homogeneous agents get arbitrarily long, and within those periods, asymptotic learning will eventually be approximated. Since the "crazy" types agree with the other at certainty (otherwise we would not have the convergence of X_i), it must be the case that they eventually will "match" the action of the non-crazy types, and information will be transmitted.

The role of the prior will be made more explicit in Section 7.

6.5 Sketch of the Proof of Theorem 3

There can be two reasons why the separation algorithm fails to produce a learnable partition of X. It can be because the belief basins fail to converge, as is the case in the anti-herding example; or it can be the case that they do converge, but to a coarsening that is not convex.

The way to construct an information structure that leads to failure can explore either of the two types of failure. The Anti-Herding Example example illustrates how the first type acts: by creating a signal structure that leads two different interim posteriors of player i to be in the same belief basin, it is possible to "merge" the beliefs of the predecessor, leading to a loss in information accumulated.

This section will show how the second type of failure, the non-convexity, can result in a failure of information accumulation.

In this example, there are three states of the world, $\{\theta, \theta', \theta''\}$. The goal is to create a cycle where the distribution of posteriors of infinitely many players is going to be the same as the distribution of posteriors of the first player.

Suppose the first player has the preferences describes in Figure 13. Her prior is the middle point μ. The support of her posterior has six points. With a very small probability, she get a perfectly informed signal, leading her belief to one of the vertices of the simplex. With the remaining probability, she gets a signal that leads her to one of the other three points signaled in Figure 13. The important feature of those points is that they lie on the "beak" of the belief space.

Suppose that Player 2, represented in Figure 14, observes that Player 1 took the "top" action. He can infer that she received either the perfectly informative signal that revealed the state to be θ, or the other signal that imperfectly moves her posterior towards θ. Since the latter is much more likely than the former, the average interim belief of Player 2, signaled as $\tilde{\mu}$.

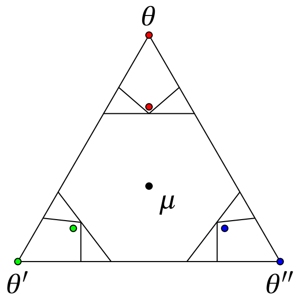

Figure 13: Payoff and Information Structure of the first player.

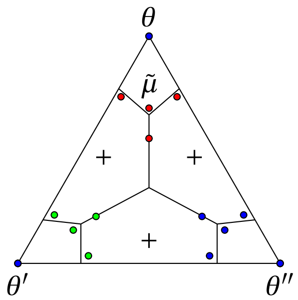

Figure 14: Payoff and Information Structure of the Second Player.

He then gets his own signal, that can only be perfectly informative (with a very small probability) or can take her posterior to be one of the other three red dots in the figure. An important feature of that example is that the interim belief is in an extreme belief basin, and also in the convex hull of those three belief points, but none of the three belief points are in the extreme belief basin. This is possible because the interim belief is in the convex hull of the complement of the extreme belief basin.

A similar structure happens in case Player 2 observes Player 1 taking any of the other actions he takes in equilibrium, which are represented by the colors green and blue.

Three of the belief points in the support lie on the boundary between the belief basins. For these points, Player 2 is indifferent between two actions, and he takes each with probability 50%.

If Player 3 observes that Player 2 took one of the extreme actions, she knows that the only belief on-path that would lead him to do so are the beliefs that assign probability 1 to a state, and therefore she would learn the state. If instead Player 3 observes Player 2 taking one of the other three actions, she knows that there are four beliefs on-path compatible with that action, so she averages between them. The average belief is represented by the "+″ sign on Figure 14.

Player 3's preferences and possible interim beliefs are represented in Figure 15. She can receive a perfectly informative signal with a very low probability, and a perfectly uninformative one with the remaining probability. In this case, when Player 4 observers her, he will know that extreme actions are only taken when Player 3 knows the state of the world, and the intermediate action is taken when Player 3 has one of the beliefs marked as "+″. The average of these three beliefs is exactly the original prior μ, so that if Player 4 observes the intermediate action, his interim belief will be identical to his prior. In other words, with a small probability Player 4 is

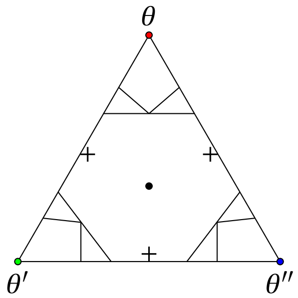

Figure 15: Payoff and Information Structure of the Third Player.

perfectly informed, and with a large probability, his posterior is his prior.

It is easy then to construct an information structure for Player 5 that is similar to the structure of Player 1, but with an even smaller probability of perfectly informative signals, so that the joint distribution of actions and states of Players 1 and 5 are identical. This way, the informativeness of Player 1's actions and of Player 5 are the same, thus establishing a cycle.

Of course this is a highly-stylized example in which a cycle is possible to be generated within a small number of periods. The formal proof for the result shows how to construct such an information structure in a more general setting in which the learnable partition is coarser than X.

The proof hinges on the possibility of increasing the probability of an uninformative signal to "save" some probability mass for future periods that can be used to construct signals that will lead the interim posteriors to either be in the "beak" (as in the example above) or in the "discrepancy" regions (as in the Anti-Herding Example), and thus be able to "merge" interim beliefs.

Furthermore, the proof leverages the non-sparsity assumption to show that there exists a finite number T such that those cycles can always be constructed within less than T periods. The last player of the cycle will be the one that profits close to nothing from observational learning.

6.6 Discussion

One limitation in the structure of the proof ofTheorem 3 is that it relies on constructing information structures that are not conditionally independent, let alone identical. It is not clear if it is possible to construct a similar result by imposing those assumptions.

Besides that, the results says that there is some informational structure for which there is no asymptotic learning. One can construct a sequence of games for which the limit of the limit of the payoff structures converges (for example, in the anti-herding game, if $\gamma \to 0$). The limit game contains robust asymptotic learning, but none of the games in the sequence do. There is a discontinuity.

The same not necessarily cannot be said if we keep the signal structure fixed. That is, if the limit of the limit of payoff structures converge, it can be the case that the set of informational structures that do not lead to asymptotic learning converges to the empty set. This is left for future research.

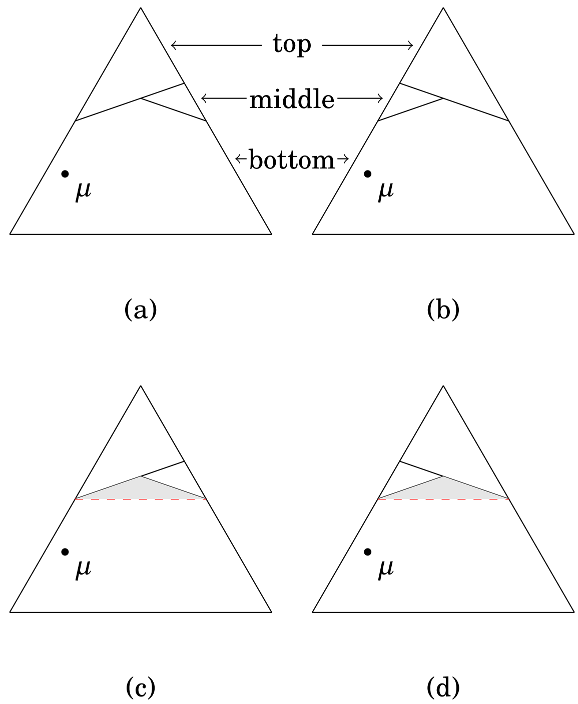

Figure 16: Example of the role of the priors.

7 Prior-Specific Sufficient Conditions

The conditions specified in Theorem 1 and in Theorem 2 refer solely to the payoff structure of the players. They are prior-free. Nevertheless, Theorem 3 is not prior-free. It says that there will be one prior and one signal structure associated with failure of information aggregation, but not that there will be such a signal structure for every prior. The logic of that condition is exemplified in Figure 16.

There are three states of the world. Triangles (a) and (b) represent the belief basins of odd and even players, respectively. Each player has access to three actions, “top”, “middle” and “bottom”, where “top” and “bottom” are extreme basins. The common prior is given by μ.

Notice that there is no coarsening of the belief basins that is convex. The sufficient conditions for Theorem 1 do not apply here.

Nevertheless, in this specific example, the positioning of the prior is such that, for every unbounded signal structure, there will be asymptotic learning.

Triangles (c) and (d) reproduce (a) and (b), but adding a red dashed line. This dashed line is a hyperplane that separates the prior μ from the belief basins middle and bottom, and it is the same line for all player.

From the martingale condition, we know that the interim beliefs (the beliefs of a player after observing the action of the predecessor but before observing his own signal) must average out to the prior. In other words, the prior must be in the convex hull of the support of the interim beliefs. That implies that the interim belief upon seeing the predecessor take the bottom action must lie below the red line. It will never be the case that the interim belief of a player will lie in the gray triangle.

Because of that, it is possible to construct an auxiliary game in which players have preferences described by four belief basins: “top”, “middle”, the gray triangle and the bottom diamond, and the information structure is such that no player ends up with a belief in the gray diamond area in equilibrium, and the outcomes are the same as in the original game.

This alternative game satisfies the conditions of Theorem 2, since it is possible to merge the top, middle and gray triangle basins, to form a convex coarsening. There must be asymptotic learning in this alternative game, and therefore there must also be in the original game.

Theorem 4. *If there exists an integer I and an event $x \subseteq \Theta$ such that, for all players $i > I$:*

- *δ_x lies in the same belief basin as the prior;*
- *for all $\theta \notin x$, δ_θ is not in the same belief basin as the prior, and*
- *there is a hyperplane separating the prior μ to all belief basins μ is not in,*

then x belongs to the learnable partition.

Proof. Proof in Appendix 5. □

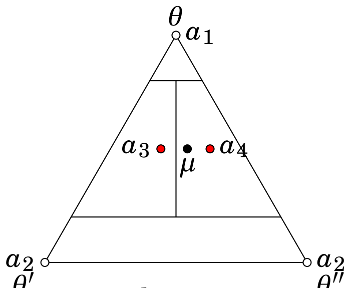

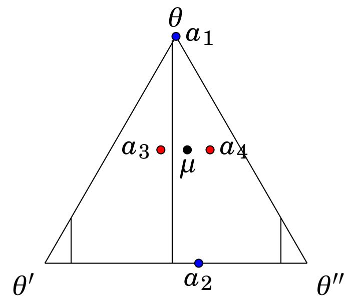

Figure 17: Example of Break of Monotonicity

8 Extensions

8.1 Break in Monotonicity

A natural question to be asked in social learning settings is whether a player would like his predecessor to be very well-informed.

At first glance, it may seem that the answer is "yes". The action of a very well-informed agents is going to reflect more information than that of a poorly-informed agent. Even though that intuition is valid for the homogeneous-preferences case, it is not applicable in the heterogeneous-preferences case.

Consider the following stylized example. There are three states of the world, $\{\theta,\theta',\theta''\}$. Figure 17 below represents the belief basins, and thus, the preferences of the players. The triangle on top stands for the simplex of player 1's beliefs, and the one below, that of player 2. The partition of the triangle represents the belief basins for each player. Both players start with a common prior μ.

Let's consider two alternative signal structures. First, suppose that player 1 has access to a perfectly revealing signal structure, and player 2 to a perfectly uninformative one. This can be seen as a reduced-form way to represent a social learning game where player 1 actually comes after a long line of players that share her exact preferences, and all players receive a full-support, but very uninformative, signal structure.

In this case, player 1 will be aware of the correct state of the world, and will take action a_1 if the state is θ, and action a_2 if the state is either θ' or θ''. The white dots on the top triangle of Figure 17 represent the support of her posterior, and is labeled with the associated optimal action.

Upon observing a_1, player 2 infers that the state of the world is θ and his belief becomes the blue dot denoted by a_2 on the bottom triangle. Upon observing a_2, player 2 learns that the state of the world is either θ' or θ'', but cannot distinguish between them, leading his posterior to be the blue dot labeled a_2. The prior μ, the posterior after observing a_1 and the posterior after observing a_2 all belong to the same belief basin. Observing his perfectly informed predecessor does not have any value to player 2, since it does not change his choice of action compared to not receiving any signals.

Now, consider an alternative, Blackwell-inferior signal structure for player 1 denoted by the red dots. Player 1 can now receive two signals, one of which would lead her to take action a_3 and the other one, action a_4. Player 1's actions are perfectly revealing of the signal she got, so upon seeing them, player 2 inherits her posteriors. Those two points belong to different belief basins, so different actions of player 1 would induce player 2 to also take different actions. In this case, observational learning has strictly positive value to player 2.

This example shows that in a scenario with heterogeneous preferences, a player may prefer his predecessor to be poorly informed. This stands in contrast with the homogeneous preferences case, in which Blackwell-improvements to the predecessor's information leads to an increase in payoff of all successors.

Proposition 6. *If preferences are homogeneous, in equilibrium, Blackwell improvements in the predecessor's information generate weakly higher payoffs for the successor.*

Proof. This is a direct implication of the improvement principle. □

In other words, a player being Blackwell more informed does not translate into that player being Blackwell more informative.

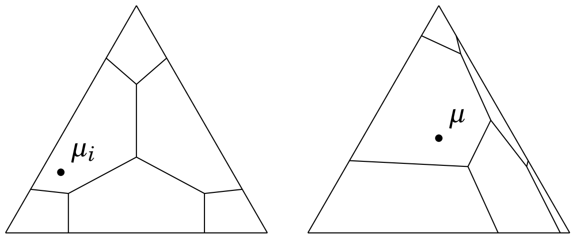

Figure 18: Example of the prior homogenization process.

8.2 Heterogeneous Priors

The results in this paper are able to accommodate a situation in which players not only have heterogeneous preferences, but also heterogeneous priors. That is, the description of a game becomes $\mathscr{G} =< \Theta, \{A_i\}_i, \{u_i\}_i, \{\mu_i\}_i, \{F^i\}_i >$. All priors, preferences and signal structures are common knowledge.

The key to this analysis is to construct a game of homogeneous priors, and show that there is adequate learning in the original game if and only if there is one in this new auxiliary game.

Definition 15. *The a-inducing signals for i is a set of signals $S_{ia} \subseteq S^i$ such that $s \in S_{ia}$ if and only if a is in the best response correspondence of a player i that observed s and nothing more.*

We can use the action inducing signals to construct the new preferences for each player in a way that they all share the same prior and their actions remain equally reflective of their private information. The new preferences rationalize a new player $\hat{\imath}$, that has a different prior, but takes the same action upon seeing the same evidence as i. Therefore, $\hat{\imath}$'s actions reflect the same information as i's actions.

Definition 16. *May $\mathscr{P}_i$ be the space of all bounded preferences for player i (that is, all functions from $\Theta \times A_i$ into $[0, M]$), and may $\beta : \Delta(\Theta) \times S^i \to \Delta(\Theta)$ be the Bayesian updating function ($\beta(\mu, s)$ is the posterior of a player with prior μ who observes signal s).*

The prior-homogenizing transformation for player i is a function $T : \mathscr{P}_i \times (\Delta(\Theta))^2 \to \mathscr{P}_i$ such that $BR_{u_i}(\beta(\mu_i, s)) = BR_{T(u_i,\mu_i,\mu)}(\beta(\mu, s))$.

The prior-homogenizing transformation is a function $T(u_i, \mu_i, \mu)$ that maps preferences and a pair of beliefs into a new set of preferences. It has the property that a player with preferences u_i and prior μ_i upon seeing any signal s will take the same action as a player with preferences $T(u_i, \mu_i, \mu)$ and prior μ that observes that same signal. Figure 18 exemplifies a prior-homogenizing transformation. Geometrically, the transformation "stretches" and "compresses" the belief basins for them to represent the same information received by the player.

Definition 17. *For a full-support belief $\mu \in \Delta(\Theta)$, a μ-homogenization of $\mathscr{G}$ is a homogeneous-prior game $\hat{\mathscr{G}}_\mu =< \Theta, \{A_i\}_i, \{T(u_i, \mu_i, \mu)\}_i, \{\mu_i\}_i, \{F^i\}_i >$, where all players share the same prior μ.*

Proposition 7. *The prior-homogenizing transformation is well-defined. That is, for any u_i, μ_i and $\mu \in \mathscr{P}_i \times (\Delta(\Theta))^2$, $T(u_i, \mu_i, \mu)$ exists and is a preference.*

For any heterogeneous prior game $\mathscr{G}$, and any prior $\mu \in \Delta(\Theta)$, there exists a μ-homogenization of $\mathscr{G}$.

Proposition 8. *There is adequate learning in $\mathscr{G}$ iff there is adequate learning in $\hat{\mathscr{G}}_\mu$ (for any μ).*

Proof. In equilibrium, the joint distribution of states of the world, signals, and actions are identical by construction in both games $\mathscr{G}$ and $\hat{\mathscr{G}}_\mu$, for any μ. Therefore, if the joint distribution of actions and states of the world in game $\hat{\mathscr{G}}_\mu$ converges to the one of perfectly informed agents, then it must be the case that the same happens in game $\mathscr{G}$. □

The natural conclusion of Proposition 8 is that all the analysis made so far can be made in the heterogeneous-prior case by recasting the problem as a homogeneous-prior problem and applying the results from this paper.

8.3 Observing More than the Immediate Predecessor

One of the modeling choices in this paper that represent a significant departure in respect to the canonical papers in the literature (such as Bikhchandani et al. (1992), Banerjee (1992), Smith and Sørensen (2000)). This modeling choice significantly simplifies notation, but the main results qualitatively go through when a more complex network of observation is allowed.

Acemoglu et al. (2011) was the first paper to analyze the canonical social learning in a network setting. The results and definitions and notation introduced in this section closely follow the ideas present in that paper.

In this section, player i can observe the action taken by a subset of agents $N_i \subseteq \{1,...,i\}$. This subset is called i's *neighborhood*.

Acemoglu et al. (2011) notes that, as long as players do not only observe some of the first K players (for some finite K), the two-states, two-actions homogeneous preferences model with unbounded signal structure has asymptotic learning. This is captured in the definition and proposition below.

Definition 18 (Acemoglu et al. (2011))**.** *The network has* expanding observations *if for all* $K \in \mathbb{N}$*, we have*

$$lim_{i\to\infty}\mathbb{1}\left(\max_{j\in N_i} j < K\right) = 0. \tag{1}$$

If the network does not satisfy this property, then we say it has nonexpanding observations.

An example of a network with nonexpanding is one in which all players only observe the action taken by the first player, $i = 1$. The main model in this paper, in which all agents observe their immediate predecessor, is an example of a network with expanding observations.

Definition 19. *A network contains* long-run sparsity *if, for every* $K > 0$*, the network obtained by deleting the first* K *nodes is connected.*

An example of a network that does not contain long-run sparsity is the following. Consider the same environment as the anti-herding example. There are two states of the world, $\{L,R\}$, and two actions available to all players, $\{\ell,r\}$.

There are two types of agents. Type A takes action ℓ if and only if his posterior is $P(\theta = R) < \gamma$, and type B, in and only if $P(\theta = R) < 1-\gamma$. That is, type A corresponds to the odd players in theanti-herding example, and type B, the even players. The information structure is also the same: with a small probability ϕ, the private signal is perfectly informative, and with a small probability $1-\phi$, perfectly uninformative.

Suppose that the line of agents is always of the format $\{A,A,B,A,A,B,A,A,B,A...\}$, that is, the type of the agent is B if and only if his position in the queue is a multiple of three.

Furthermore, suppose that the observation network is the following:

$$1 \to 4 \to 7 \to 11 \to 14 \to 17 \to ...$$

$$2 \to 3 \to 5 \to 6 \to 8 \to 9 \to 11 \to ...$$

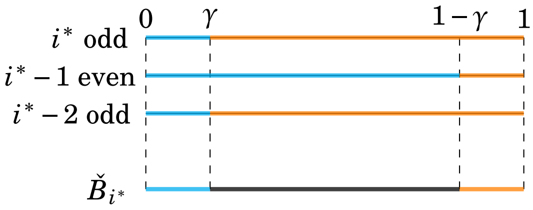

Figure 19: Neighborhood join for i^*.

This amounts to two disconnected networks. One of them is populated only by type As (and thus has homogeneous preferences), while the other one exactly matches the anti-herding example. Given the results seen before, the first network has features asymptotic learning whereas the second does not.

To avoid this sort of complication, it is without loss to analyze networks that satisfy long-run sparsity, otherwise the results of this paper can be applied to each branch of the network separately.

Definition 20. *Player i's* neighborhood join $\check{B}_i$ *is a family of subsets of* $\Delta(\Theta)$ *generated by the join of the set of belief basins of the players in i's neighborhood:*

$$\check{B}_i = \wedge_{j \in N_i} B_j.$$

Player i's neighborhood informational content at certainty $\check{X}_i$ *is a family of subsets of* Θ *generated by the join of the informational contents at certainty of the players in i's neighborhood:*

$$\check{X}_i = \wedge_{j \in N_i} X_j.$$

For example, in the anti-herding model, odd players' belief basins are $\{[0,\gamma),(\gamma,1]\}$; and even players' belief basins are $\{[1,1-\gamma),(1-\gamma,1]\}$. If a player i^* were to observe his two immediate predecessors, then his neighborhood join would be

$$\check{B}_{i^*} = \{[0,\gamma),(\gamma,1]\} \wedge \{[1,1-\gamma),(1-\gamma,1]\} = \{[0,\gamma),(\gamma,1-\gamma),(1-\gamma,1]\}.$$

This is graphically exemplified in Figure 19.

Intuitively, if i's neighborhood join is fine enough, she can get a lot of information from observational learning. This allows her to compare the information of all players that she observed and have a better grasp on the uncertainty level that they had.

The two conditions that need to be imposed to generalize the main results are expanding observations, and neighborhood joins being fine enough.

Theorem 1^***.** *Fix a sequence of preferences* $\{u_i\}$*. There will be some non-trivial signal structure* $\{F_i\}$ *for which the game has asymptotic learning if and only if the sequence of neighborhood informational contents at certainty is finer than i's own informational content at certainty. That is, if:*

$$VI(\check{X}_i, \check{X}_i \wedge X_i) \to 0.$$

Proof. See Appendix 8.

The intuition for this proof is very similar to the intuition for the proof of Theorem 1. If the condition in the theorem is satisfied, it means that players eventually could figure out all information they need by crossing the observations of their perfectly-informed neighbors.

Definition 21. *An event $e \subseteq E \subseteq \Delta(\Theta)$ is* separated* *at E is there exists a preference C and a subsequence of sets of coarsenings of neighborhood joins $\tilde{\check{X}}_j$ limited to E such that:*

- *the complement of the subsequence is sparse. That is, may $K \subset \mathbb{N}$ be the set of elements of the original sequence not in the subsequence ($K = \mathbb{N} \backslash Im(j)$). Then,* $\lim_{k\to\infty} \frac{|\{\ell \in K : \ell \leq k\}|}{k} = 0$.
- *$\tilde{\check{X}}_i$ converges to C;*
- *C is convex;*
- *there exists an element $\tilde{c_e} \in C$ such that $\{\tilde{\mu} \in \Delta(\Theta) : \tilde{\mu}(e \cap E) = 1\} \subseteq \tilde{c_e}$ and $\{\tilde{\mu} \in \Delta(\Theta) : \tilde{\mu}(e \cap E) = 0\} \cap \tilde{c_e} = \emptyset$.*

The original separation process looks at coarsenings of preferences. This process looks at coarsenings of neighborhood joins instead. This reflects the extra information that agents can get by exploring the joint distribution of the actions that are available to them.

Definition 22. *Partition Y of Θ is the* learnable partition* *if it is the outcome of the* separation algorithm*, *which is described below:*

1. *set $t = 0$ and $S_0 = \{\Theta\}$;*
2. *Consider each element $E \in S_t$, and find the finest partition of E such that all of its elements can be separated* at E. Call this partition P_E^t.*
3. *Set $S_{t+1} = \cup_{E \in S_t} P_E^t$. ($S_{t+1}$ is itself a partition of Θ)*
4. *Stop if $S_{t+1} = S_t$. Otherwise, advance one step on t and return to step 2.*

The learnable partition* is the analogue of the learnable partition when considering the networks game.

Theorem $2^*/3^*$. *Suppose $VI(\check{X}_i, \check{X}_i \wedge X_i) \to 0$, that there are expanding observations, and a non-sparse network. There is robust asymptotic learning if and only if the learnable partition is $\lim_{i\to\infty} \check{X}_i$.*

Proof. See Appendix 8.

This result is the analogue of Corollary 2, combining results that extend Theorem 2 and Theorem 3. The proof is very similar to the proof of those theorems, but making the analysis in the joins of belief basins, which makes it significantly more notationally cumbersome.

In summary, qualitatively the results from the main section of this paper still hold after allowing for more complex observational learning structures. The more agents someone observes, the easier it is to have (robust) asymptotic learning, because agents can leverage the joint distribution of actions of their predecessors.

If agents observe all their predecessors, then no information is lost across generations, and the conditions for asymptotic learning become much less stringent. All that is necessary is to rule out crazy types that only care about a state of the world that no one before cared about. An example of such a situation in which there is a different, independently distributed state of the world at each period of time, and each agent only cares about the state at his period of decision-making. This is ruled out, for instance, if the set of states of the world is finite.

Proposition 9. *If the signal structure is unbounded and agents can observe their entire string of predecessors, and if the signal structure is unbounded, then there will be asymptotic learning if , for all pairs $\theta, \theta' \in \Delta(\Theta)$, there is an infinite number of agents i for which θ and θ' belong to different elements of X_i.*

Proof. In Appendix 8. □

In other words, if there is an infinite number of agents whose action could possible differentiate between two states of the world at certainty, the distinction between these states will be learned asymptotically.

9 Conclusion

Heterogeneity of preferences can have severe implications for the process of social learning. If agents' behavior is different even if they were perfectly informed, asymptotic learning is impossible for non-trivial signal structures. The amount of heterogeneity of preferences for there to be asymptotic learning for all unbounded signal structures has to be limited by a coarsening process. In case this condition fails, social learning could completely collapse.

This failure happens because heterogeneity serves as an obstacle for agents' actions to reflect information accumulated by their predecessors. If agents can observe all their predecessor, that phenomenon is not present.

A myriad of applications can come from this framework. For example, future research could explore what is the optimal queuing of players for information to flow in heterogeneous settings. This can have large impacts in the design of policy, may it be spreading new technologies in rural Africa, or designing informationally-stable financial networks.

References

Acemoglu, D., Dahleh, M. A., and Lobel, I. (2011). Bayesian Learning in Social Networks. *Review of Economic Studies*, 78(4):1201–1236.

Akbarpour, M., Saberi, A., and Shameli, A. (2017). Information Aggregation in Overlapping Generations. *Working Paper*, pages 1–26.

Ali, S. N. (2018). On the Role of Responsiveness in Rational Herds. *Economic Letters*, 163:79–82.

Bala, V. and Goyal, S. (1995). A Theory of Learning with Heterogeneous Agents. *International Economic Review*, 36(2):303–323.

Banerjee, A. (1992). A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, 152(3).

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A Theory of Fads , Fashion , Custom , and Cultural Change as Informational Cascades. *Journal of Political Economy*, 100(5):992–1026.

Board, S. and Meyer-ter vehn, M. (2021). Learning Dynamics in Social Networks. *Econometrica*, 89(6):2601–2635.

Chen, J. Y. (2020). Sequential Learning under Informational Ambiguity. *Working Paper*, pages 1–46.

Goeree, J. K., Palfrey, T. R., and Rogers, B. W. (2006). Social learning with private and common values. *Economic Theory*, 28:245–264.

Golub, B. and Sadler, E. (2016). Learning in social networks. *The Oxford Handbook of the Economics of Networks*, pages 504–542.

Hirshleifer, D., Welch, I., Bikhchandani, S., and Tamuz, O. (2021). Information Cascades and Social Learning. *Working Paper*.

Kartik, N., Lee, S., and Rappoport, D. (2022). Observational Learning with Ordered States. *Working Paper*.

Krishnan, P. and Patnam, M. (2014). Neighbors and extension agents in ethiopia: Who matters more for technology adoption? *American Journal of Agricultural Economics*, 96(1):308–327.

Lawry, J. (2022). Heterogeneity and Robustness in Social Learning. *Working Paper*.

Lobel, I. and Sadler, E. (2016). Preferences, Homophily, and Social Learning. *Operations Research*, 64(3):564–584.

Meila, M. (2007). Comparing clusterings — an information based distance. *Journal of Multivariate Analysis*, 98:873–895.

Munshi, K. (2004). Social learning in a heterogeneous population : technology diffusion in the Indian Green Revolution. *Journal of Development Economics*, 73:185–213.

Oyetunde-usman, Z. (2022). Heterogenous Factors of Adoption of Agricultural Technologies in West and East Africa Countries : A Review. *Frontiers in Susteinable Food Systems*, 6(March):1–14.

Raz, I. T. (2020). Learning is Caring: Soil Heterogeneity, Social Learning and the Formation of Close-knit Communities. *Working Paper*, (December).

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

Smith, L. and Sørensen, P. (2000). Pathological Outcomes of Observational Learning. *Econometrica*, 68(2):371–398.

Tjernstrom, E. (2017). Learning from Others in Heterogeneous Environments. *Working Paper*, (April):1–46.

Williams, C. R. (2019). Echo Chambers : Social Learning under Unobserved Heterogeneity. *Working Paper*, pages 1–25.

Young, H. P. (2009). Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning. *American Economic Review*, 99(5):1899–1924.