TLDNR: Inattentive Learning on the Internet

John W.E. Cremin*†

December 15, 2025

Latest Version

Abstract

Our ever greater access to information has not produced a perfectly informed society of political consensus. In this article, I study the role of rational inattention in explaining this, within a model of sequential social learning. In so doing, I illustrate how to tractably model a very general class of 'social cost functions': functions that give the cost of observing any given subset of predecessors. In such a model, where there are costs to learning both from these social signals and private information, I find that making access to both forms of information cheaper (either by making the cost of private signals lower, or making it easier to observe the actions of predecessors) can reduce the asymptotic probability with which agents correctly match the state. Finally, I use my model to study the impact of the internet on our media environment, showing how greater access to the opinions of others on social media (for example, those of influencers) can remove the incentives for news organisations to produce high quality news in equilibrium.

Keywords— Sequential Social Learning, Endogenous Social Networks, Network Theory, Information Economics

1 Introduction

"Orwell feared those who would deprive us of information. Huxley feared those who would give us so much that we would be reduced to passivity and egoism. Orwell feared that the truth would be concealed from us. Huxley feared the truth would be drowned in a sea of irrelevance." Postman (2005)

Access to information is cheaper than ever before, and our ability to share that information with each other has never been so great, yet the resulting glut of information does not seem to have produced a hyper-rational paradise of informed consensus. Rather, polling data shows that even on questions of basic fact voters are not able to reliably discern the true state of the world. As I note in

*Aix-Marseille School of Economics, Université d'Aix-Marseille, CNRS, john-walter-edward.cremin@univ-amu.fr.

†I would like to thank my defense committee of Evan Sadler, Navin Kartik, Jacopo Perego, Mark Dean and Andrea Prat for many helpful comments and suggestions.

Cremin (2023), the '*Polarization of Reality*' observed by Alesina et al. (2020) seems a fixed feature of modern politics, and a seemingly paradoxical one in view of the glut of available information in the modern world. However, whereas Cremin (2023) considers the role of partisan motivated reasoning in impeding social learning, here I show how even without partisanship, an informational glut can actively degrade learning: failures of learning can directly *result* from this excess of information brought about by the internet. Even without any partisan noise agents (as in Cremin (2025)), and without any partisan bias on behalf of agents, the *reduction* in the cost of obtaining information, and observing the actions of others, can itself damage learning and reduce the probability with which agents manage to correctly match the state.

That information is costly to obtain, both in terms of money and cognitive resources, is clearly relevant to social learning, and one of the most obvious consequences of the development of the internet is that the costs of doing so have fallen in recent decades. In this paper, I set out a model of sequential social learning in which agents are rationally inattentive, and face costs to observe both private signals, and also the actions of their predecessors. My model is a notationally minimal representation of a more general specification of information costs, and one of its contributions is to show how one can analyse learning in the context of such a general model by considering weighted observation networks that provide a partial representation of these costs. The insights of my analysis can be relayed by supposing that when observing predecessors, each agent must simply pay a price to observe a specific predecessor, and has only a fixed 'link budget' to spend observing them. To be precise, the budget constrains the *number* of links they may observe, and is distinct from the costs of observing them, which are reflected in the agents' utility function.

Armed with this model, I find necessary and sufficient conditions for complete learning to obtain, but more importantly discover comparative statics results that show that increasing our access to information can reduce the asymptotic accuracy of agents, as I mention above. More specifically, I find that we can reduce the asymptotic probability with which agents match the true state by giving them access to unambiguously more informative private signals. We can achieve this by reducing the cost of observing every subset of predecessors. This establishes that the intuitive claim that making it cheaper for agents to observe private and/or social information *must* increase their accuracy is not true in this setting. The basic intuition behind this fact is that if early agents are made cheaper to observe to a greater extent than later agents, or if they are given access to much cheaper and more informative private signals whilst other agents are given only small improvements, this can change the endogenous observation network to the detriment of information aggregation. In models of sequential social learning, asymptotic learning is dependent on long *information paths* - chains of agents in which each observes their predecessor in the chain - as each agent in such a path can improve on their predecessors. If we shift the cost parameters such that we move from an equilibrium in which agents are all choosing to observe recent agents and generating an observation network with expanding observations to one in which they all choose to observe a very early agent, this prevents information aggregation. Informational externalities driving breakdowns of learning in sequential models are as old as the literature itself: in the classic complete network setting (Bikhchandani et al., 1992; Banerjee, 1992), agents optimally following their social belief ensures subsequent agents learn nothing about their private signal. It would be better in such a setting for them to vary their action with their private signal, thus allowing information to continue to accumulate. In this paper, these externalities are even more pernicious, as in addition to ignoring externalities when choosing their action, they also ignore them in choosing what information to observe before doing so.

In addition to this, I also consider how we can model the media landscape within this set-up, and find that social media can directly cause a collapse in the readership of newspapers. When agents can too cheaply observe influencers (agents who are not able to conduct private investigations any more cheaply that their followers), they are easily tempted to do so instead of reading newspapers. This can produce a situation in which all agents need to conduct their own investigations, achieving less accuracy at greater cost.

Literature Review: The literature on sequential social learning is substantial, going back to classic articles by Bikhchandani et al. (1992) and Banerjee (1992). The contribution of Smith and Sorensen (2000) provides a general analysis of sequential social learning on a complete network,[1] nesting much of the earlier work, and was the first to discover the importance of the distinction between bounded and unbounded beliefs. A number of articles in the sequential literature have considered costly information before, though none have yet made the observation of neighbors and private information costly simultaneously. This limitation ensures that none allow the analyst to evaluate the impact of social media as I do. Burguet and Vives (2000), Mueller-Frank and Pai (2016), and Lomys (2020) each consider costly private signals, but the articles whose approaches mine resembles most are those of Ali (2018) and Song (2016). Whereas Lomys and Mueller-Frank and Pai both allow agents to conduct costly investigation of the pay-off of one risky action, which they may or may not choose over a certain-payoff safe action, Ali's model involves a binary state of the world, agents who wish to match this state, and the availability of costly private signals, or 'experiments', that reveal information about the state of the world. My notion of affordability follows on from Ali, though my private-signal setting is more general in some ways that require me to differentiate between affordability and a stronger notion (uniform affordability) that are equivalent in his setting. My Proposition 2 partially nests his Theorem 1. The major difference between my paper and his, however, it that I allow general network topologies (a feature my model shares with that of Lomys), and make it costly to observe predecessors. Conversely, Song (2016) allows free observation of a private signal but makes the *social* signal costly and strategic. In his paper, agents observe a private signal for free before doing anything else[2] They then can pay a cost c to obtain capacity $K(n)$, where once obtained this allows them to observe up to $K(n) \in \mathbb{N}$ agents. My model is similar but not the same, as agents have both a *link budget*, and pay a different price for each predecessor they might want to observe. Appendix C studies a more general model, and Section

[1] In which all agents observe all those who came before.

[2] He does also note that making them choose whether or not to observe a social signal first changes results. With unbounded signals he finds that there is not complete learning if social signals are not free, but that there can be with low enough social cost if private signal observation occurs afterwards.

4 further considers when private and social costs are nonseparable. Song finds that decreasing the social cost c can improve learning, just as I find here that either increasing social or private costs can improve learning by changing the properties of the endogenous observation network. As far as I am aware, no article yet studies a model in which both private and social signals are costly. Nor does any existing article observe (as I do here) the importance of whether these two costs are separable or not, or use this as a model of learning with rational inattention.

The standard references of Acemoglu et al. (2011), Lobel and Sadler (2015), Smith and Sorensen (2000) are ancestors of much of this costly observation literature, and the techniques of these papers are still of use here as can be seen in the proof techniques I use. I have noted the importance of Smith and Sorensen (2000) above, but Acemoglu et al. (2011) is important in generalising beyond the complete network, and allowing for arbitrary social network structures, with the caveat that agents' neighbourhoods are independent of each other. In this setting, they find that *Expanding Observations*- a minimal connectivity condition- is necessary and sufficient for asymptotic learning with unbounded beliefs.[3] A small literature following Acemoglu et al. (2011) has developed, for example containing Lobel and Sadler (2015, 2016) and Lomys (2020). The first of these removes the neighborhood independence assumption of Acemoglu et al. (2011), unlike the model I present here,[4] and the second studies learning in a setting where agents have different preferences over the two actions. This turns out to be sufficient to break learning in general networks, and some of the learning problems in my model are reminiscent of it.

In addition to commenting on the relationship of this paper to the literature in sequential social learning, it will also be worth commenting on that of rational inattention. Rational inattention has been much studied in political economy already. Matějka and Tabellini (2021) model rationally inattentive voters, and show that this gives more influence to voters with strong preferences, since they are correspondingly more likely to pay close attention to issues in which they are very invested. Yuksel (2022) and Hu and Li (2018) demonstrate that rationally attentive voters can cause political platforms to polarize, complementing the themes of this thesis, and even more pertinently, Maćkowiak et al. (2023) note that the internet can exacerbate by providing '*a finer granularity of information, allowing voters to focus even more on narrow topics of their particular interest.*' The literature literature linking rational inattention to polarization also includes articles that show how it can produce the polarization of beliefs directly (rather than simply political platforms), such as Nimark and Sundaresan (2019) and Novák et al. (2021). This paper does not show rational inattention producing polarization, but rather shows that in its presence an increase in the availability of information can paradoxically produce a collapse in agents' access to information.

[3] In addition to these articles studying general network topologies with the Acemoglu et al. (2011) framework, there are of course articles such as that by Çelen and Kariv (2004) that study specific non-complete network topologies such as the line network. I follow the Acemoglu et al. (2011) approach, as real world networks are inevitably going to contain all sorts of arbitrary patterns, making results on general network topologies of much more use in studying social learning.

[4] Note that an implication of this it that motivated reasoning can damage social learning even *without* type homophily and echo chambers.

An interesting bridge between the the literatures on social learning and rational inattention is Caplin et al. (2016); this article studies a model of social learning with rational inattention, in which a continuum of agents enter a market each period and choose between a set of options. Like Ali 2018, they assume that firms observe social information (the fraction of agents that chose each option in each previous period) without cost, and assume a complete network, though they do not use networks terminology to describe this.

Caplin, Leahy and Matějka model the cost of acquiring private information using mutual information (the expected reduction in the entropy of beliefs over the state), and this is standard in the rational inattention literature. It can also be seen in the famous Matějka and McKay (2015) article microfounding the multinomial choice model as a product of rational inattention; that of Caplin et al. (2019), which connects the theory of consideration sets and rational inattention based on Shannon mutual information; and Sim's seminal 2003 article (Sims, 2003) that first set out the rational inattention model using costs reflecting the mutual information between prior and posterior beliefs. In their literature review on rational inattention, Maćkowiak et al. (2023) list the use of a cost function based on mutual information as the third of three defining features of the benchmark model, though the first two assumptions - that information is available in a wide variety of forms, and that agents choose information optimally - are the main ones.

Despite this, mutual information costs are not essential, and in my particular setting I believe it is important to be able to analyze behavior without imposing a specific cost function. To see why, consider the cost for any agent of observing their immediate predecessor in a symmetric[5] environment. If we assume agents' cost function for social information is based on mutual information, and agent n matches the state with probability α_n, then the cost an agent will pay to observe the action of their immediate predecessor could vary wildly as a function of their index. If we intend this cost to reflect the cost of logging into Twitter or Facebook and observing the last comment, we might reasonably object to this property, instead wanting the cost of observing $n-1$ for n to be the same for all n in $\mathbb{N}$: mutual information based costs are thus arguably inappropriate here. On the other hand, why might the α_n be higher for higher values of n? In such sequential models as I study here, an important reason for this is that later agents come with higher probability at the end of long chains of agents, each observing some predecessor, and this ensures their action 'reflects' a large number of independent signals. Given this, mutual information costs that make higher α neighbours more costly to observe could be seen as a useful heuristic reflecting the greater cognitive effort required to think through such a long chain of observations, and what value of α_n it implies. This reflects a specific sequential social learning twist on the classic criticism of rational inattention that it is odd to suppose agents who face cognitive costs are nonetheless going through the elaborate calculations required to derive the optimal attention strategy: computing α_n can be an absolutely fiendish task in equilibrium, depending on the exact network structure and available signals. Given that there are reasonable arguments both in favor of and specifically against mutual information, it is important

[5]By this I mean the private signal structures available to all agents are symmetric across the two states of the world.

to work with general cost functions, and find results that do not depend on their precise form.

The rest of the article is structured as follows. In Section 2 I set out the model, along with several important concepts concerning the network topology, available private signals, and orders over these objects. Section 3 presents the necessary and sufficient conditions, among them that if the network of free-to-observe neighbors has expanding observations and private information is sufficiently available (*overturning information is uniformly affordable*) (Proposition 2). Proposition 3 then establishes that a basic condition on private signals is not necessary for learning, and Proposition 4 that without such free neighborhoods we cannot achieve learning even with free unbounded private signals. More importantly, however, are the comparative statics results in Propositions 5 and 6. Since the motivation of this paper is to evaluate the impact of the internet on political discourse, these results on the changes that can result from cheaper signals and observation are the clear results that have something to say on this. These establish that the impact of the internet is not obvious, but that it can counter-intuitively reduce asymptotic accuracy to make information and observations cheaper. In Section 4, I consider how nonseparable information costs could be used to model agents keeping up with the news via social media news feeds. I note that in this instance, social media could help learning, and in fact the improvement of off-platform private information (for example via the development of AI and ChatGPT) could threaten any equilibrium producing complete learning on this basis. Finally, I consider the impact on the media landscape in Section 5 and find that social media generates a public good problem in Proposition 7. Section 6 concludes.

2 Model

In this model, I adopt a canonical sequential social learning model with a binary state $\theta \in \Theta = \{0, 1\}$, in which an infinite sequence of agents $n \in \mathbb{N}$ arrive and must make an irreversible action choice $x_n \in \{0, 1\}$ in order to maximise their expected utility, with utility function:

$$u_n(x_n, \theta) = \begin{cases} 1 & \text{if } x_n = \theta, \\ 0 & \text{if } x_n \neq \theta, \end{cases}$$

Each agent n can buy a single experiment,[6] X_n, from a finite set of experiments $\mathcal{E}$ at a price given by the private cost function $C_n^P : \mathcal{E} \to \mathbb{R}_+$. Each agent's private cost function is drawn from a distribution over cost functions by nature at the beginning of their turn $C_n^P : \mathcal{E} \to \mathbb{R}_+ \sim \Delta_n^P(\mathbb{R}_+^{\mathcal{E}})$. I normalise the realizations of all experiments to the posteriors they induce, and assume that all available experiments are informative but never perfectly so (they are mutually absolutely continuous). I assume there is a uniform common prior for convenience; this assumption is without loss.

As I have noted in my introduction, the model I study here is a simple implementation of a much more general model, that sacrifices some of its richness for greater parsimony. The model of the main

[6] I shall refer to these as either *experiments* or *private signals* as convenient.

body of this text is rich enough to replicate all the results I have established in my more general model, in a vastly less complicated framework. I discuss this more general model in Appendix C.

In models of social learning in which the observational network is exogenous, there is no need to distinguish between different notions of the network topology. Here, however, this is not the case: each agent n chooses his own neighborhood $B(n)$. In this model there will be three important elements to the social costs: (1) all agents can form at most L links; (2) each agent i will have to pay cost c_{ij} to observe any predecessor j, and the vector of individual costs for each predecessor will be drawn from a distribution $\Delta_n^S(\mathbb{R}_+^{n-1})$; (3) the cost of observing a set $\mathcal{B}$ of predecessors is simply the sum of the individual observation costs for each agent in $\mathcal{B}$, thus the cost of observing any subset of predecessors is completely determined by this vector of realized costs C_i. For the sake of this present paper, I can assume that all agents face the same link budget L, though all results hold in the more general model, which is also more general than would even be represented by giving each agent their own individual L. L can take any value within $\mathbb{N}\cup\{\infty\}$, where setting $L=\infty$ and studying a degenerate distribution of network costs where each cost is equal to zero with probability one gives the benchmark complete network.

Definition 1 (Social Cost Structure)**.** *The social cost structure of a given game is the sequence of all agents' social cost vector distributions:* $\Delta^S = \{\Delta_n^S\}_{n\in\mathbb{N}}$ *and the universal link budget* L*.*

The 'network topology' is then an endogenous object, precisely I call it the Endogenous Observation Network Topology for Equilibrium σ, $\{\{\mathbb{Q}^n_{\sigma,\theta}\}_{\theta\in\Theta}\}_{n\in\mathbb{N}}$. This is the sequence of pairs (one for each state) of neighborhood distributions induced by the *social cost structure* and strategy profile in equilibrium σ. One network topology is *strictly more cheaply connected* than another, if each individual cost c_{ij} is lower with probability 1 for all i and all j, and the link budget of the two is the same.

The timing of each agent's turn is as follows: agent n arrives in the game, and may choose either to observe a private signal $X_n \in \mathcal{E}$, some subset of his predecessors $B(n) \in 2^{\{1,\dots,n-1\}}$, or nothing. In the case he chooses to observe nothing, agent n simply selects his irreversible action $x_n \in \{0,1\}$ and his turn ends. Otherwise, if he observed a private (social) signal, he updates his belief to form an interim belief. He can then choose to either observe a social (private) signal, or nothing at all. If the latter, again he chooses $x_n \in \{0,1\}$ and ends his turn; if the former he updates his interim belief to form a final posterior belief, and chooses his action to end his turn.

I denote the action-history of the game up to and including agent n as $h^n := \{x_1, ..., x_n\}$. The equilibrium concept I use here is Perfect Bayesian Equilibrium, and equilibrium existence is straightforward to establish as is standard in such models.

Proposition 1 (Equilibrium Existence)**.** *A Perfect Bayesian Equilibrium exists.*

Proof. See Appendix B. ☐

Finally, the major outcome of interest in this game is whether or not *Complete Learning* obtains:

Definition 2 (Complete Learning)**.** ***Complete learning*** *obtains in equilibrium σ if x_n converges to θ in probability (according to measure $\mathbb{P}_\sigma$), i.e. if $\lim_{n\to\infty} \alpha_n = 1$.*

2.1 Affordability

Following Ali (2018), it is necessary to introduce the concept of *affordability* to analyse this environment. A given experiment $X \in \mathcal{E}$ is *affordable* for agent n if for any $k \in \mathbb{R}_+$ there is strictly positive probability that n can afford experiment X on a budget of k, i.e. $\Delta_n^P(C_n^P(X) \leq k) > 0$. *Information* is said to be *affordable* for n if there exists some experiment $X \in \mathcal{E}$ that is affordable, and we say that *overturning information is affordable for n* if for every $\underline{b}, \overline{b} \in (0,1)$ with $\underline{b} < \overline{b}$ there exists both an affordable experiment with support extending above $\overline{b}$, and one with support extending below $\underline{b}$. If information is not affordable for n it is *unaffordable* for n, and similarly if overturning information is not affordable for n it is also *unaffordable for n*.

An important difference between this setting and that of Ali (2018) is that agents are not assumed to be homogeneous in their access to private information here. I thus introduce another concept, *uniform affordability*, that is equivalent to affordability in the setting of Ali (2018). Information is said to be *uniformly affordable* for a set, A, of agents if there is some common distribution over cost functions $\Delta^S(C^P(X))$ such that for any $n \in A$ we have that for any $k \in \mathbb{R}_+$, $\Delta_n^P(C_n^P(X) \leq k) \geq \Delta^P(C^P(X) \leq k) > 0$. Naturally, one can also speak of overturning information being uniformly affordable for a set of agents. In Ali (2018), if information is affordable, it is uniformly affordable for the set $\mathbb{N}$.

In addition to defining affordability, it will also be useful to define an order over experiment set-cost function distribution pairs. This shall be that one experiment set-cost function distribution pair shall be *Blackwell-preferred* to another for agent n if we can define a pair an injective functions, $\mathcal{F}^{\mathcal{E}}$, from the latter to the former such that every experiment-cost function $\{X, C_n^P(X)\}$, X is Blackwell-dominated by its corresponding experiment, $\mathcal{F}^{\mathcal{E}}(X)$, and costs less with probability 1.

Definition 3 (Blackwell-Preferred Experiment Set-Cost Function Distribution Pairs $\succeq_B$)**.** *One experiment set-cost function distribution pair $\{\mathcal{E}_1, \Delta_n^{P,1}(\cdot)\}$ is Blackwell-Preferred $\succeq_B$ to another $\{\mathcal{E}_2, \Delta_n^{P,2}(\cdot)\}$, $\{\mathcal{E}_1, \Delta_n^{P,1}(\cdot)\} \succeq_B \{\mathcal{E}_2, \Delta_n^{P,2}(\cdot)\}$, if there exists an injective function $\mathcal{F}^{\mathcal{E}}$ such that $\mathcal{F}^{\mathcal{E}}(X_2) \in \mathcal{E}_1$ for any $X_2 \in \mathcal{E}_2$, $\mathcal{F}^{\mathcal{E}}(X_2) \succeq_B X_2$, and $\Delta_n^{P,1}(C_n^P(\mathcal{F}^{\mathcal{E}}(X_2)) < \underline{C}_n^P(X_2)) = 1$ where $\underline{C}_n^P(X_2)$ is the lower bound of the support of the cost of experiment X_2 according to $\Delta_n^{P,2}$.*

3 Necessary and Sufficient Conditions for Learning

In what circumstances should we expect learning to occur, and to what extent? A basic necessary[7] condition for learning that is ubiquitous in this literature is *expanding observations*, which guarantees that early agents do not have too much influence asymptotically. We must adjust this condition to reflect that observation networks are endogenous in this context, which produces the following:

[7]It is also sometimes sufficient with unbounded private signals, for example in Acemoglu et al. (2011)

Condition 1 (Expanding Observations)**.** *In equilibrium σ, the network topology satisfies expanding observations if for all $\theta \in \{0,1\}$, and any $K \in \mathbb{N}$:*

$$\lim_{n\to\infty} \mathbb{Q}^n_{\sigma,\theta}\left(\max_{b\in B(n)} b < K\right) = 0$$

Where I adopt the convention that $max_{b\in B(n)} b = 0$ if $B(n) = \emptyset$.

Much as it is useful in this framework to speak of different varieties of network topology, it is also useful to define multiple versions of expanding observations. In particular, I will describe a given network topology as exhibiting c-expanding observations if it would satisfy expanding observations were every link with cost $c_{ij} < c$ to be present.

Condition 2 (c-Expanding Observations)**.** *Let $\{\mathbb{W}^c_n\}_{n\in\mathbb{N}}$ be the sequence of distributions over graphs that result from adding a link between every pair of agents i and j ($i > j$) connected by observation cost $c_{ij} < c$, but adding no others. The network topology satisfies c-Expanding observations if:*

$$\lim_{n\to\infty} \mathbb{W}^c_n\left(\max_{b\in B(n)} b < K\right) = 0$$

Firstly, it is natural to consider the benchmark cases in which only one of private and social signals have non-zero costs. If observing some neighbors is actually free, what conditions on private signals guarantee learning and vice versa? If agents all observe the complete network for free, we know from Ali (2018, Theorem 1, Part 2) that if overturning information is uniformly affordable for all agents,[8] then complete learning obtains, and if it is unaffordable for all agents we have incomplete learning.[9] The first aspect of this can be generalised in this setting, as I show in Proposition 2, though unaffordable overturning information does not guarantee that complete learning does not obtain.

Proposition 2. *If we have 0-expanding observations, and overturning information is uniformly affordable for $\mathbb{N}$, complete learning obtains.*

Proof. See Appendix B. □

The 0-expanding observations and uniformly affordable overturning information ensure that we can apply an improvement principle. Where in standard improvement principles agents can improve on any neighbor, here we instead consider them improving on one of their freely-observable neighbors specifically. The uniformly affordable overturning information allows us to mimic the free unbounded experiment agents are normally assumed to observe. That we need uniformly affordable overturning information is not obvious, since in Ali 2018 this concept is not defined, and it would be natural

[8]Since he does not have the notion of uniform affordability defined, his result is expressed in terms of affordability

[9]In Ali's paper, complete learning is defined as obtaining if the public belief *almost surely* converges to certainty on the true state, whereas here it is defined as convergence in probability to the correct action.

to try and extend Ali's result by instead simply insisting that each agent has access to affordable overturning information. This is not sufficient, however, and I illustrate this fact in the following remark.

> **Remark 3.1** (Importance of *Uniformly* affordable overturning information)**.** *If we replace 'overturning information is uniformly affordable for* $\mathbb{N}$*' in Proposition 2 with 'all agents in* $\mathbb{N}$ *have affordable information', the result no longer holds.*

Proof. See Appendix B. □

The problem here is that we can conceive of information cost structures such that though each individual agent n has affordable overturning information, the probability with which they can afford such overturning information converges to zero fast enough to prevent the accuracy of agents α_n ever converging to 1. In the proof of Remark 3.1 I consider one particular example in detail, where the probability with which private information is cheap enough to be worth observing goes to zero at an exponential rate as the index of agents climbs. I plot this alongside an example in which this probability converges linearly to zero in Figure 1. Figures 1a and 1b show the same curves, though the former plots them over a very large domain to exhibit the limit, and the latter shows only the first five agents to show that both curves are gradually increasing, though it may not seem so looking at 1a.

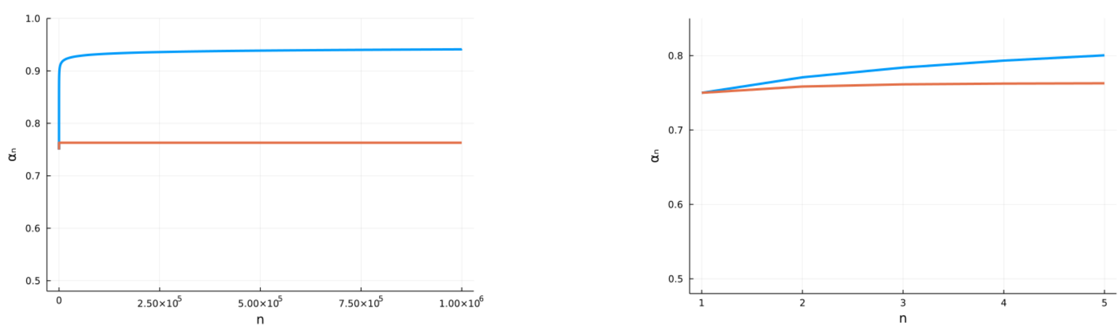

(a) Path over the first million agents

(b) Path over the first five agents

Figure 1: Figure 1: In the scenarios represented by both of these curves, agents all have access to affordable overturning information, but nonetheless they do not learn.

Though uniformly affordable overturning information can ensure complete learning, it is not necessary for it. We can instead achieve learning with sacrificial lambs for whom observing predecessors is prohibitively expensive.

Proposition 3 (Necessity of Affordable Information)**.** *Uniformly affordable overturning information for* $\mathbb{N}$ *is not necessary for complete learning to obtain, though affordable information is. Even if no agents have affordable overturning information, complete learning is still possible.*

Proof. See Appendix B. □

Having considered under what conditions on private signals we get learning with 0-expanding observations, we can also ask in which network topologies we get learning when agents are all assumed to have access to one free unbounded private signal. Clearly, by Proposition 2, 0-expanding observations is sufficient here since with this private signal structure all agents have uniformly affordable overturning information, but it turns out that the moment we dispense with 0-expanding observations and impose minimum observation costs with some probability, even free unbounded signals are not enough to achieve learning.

Proposition 4. *If there exists some $c > 0, \delta > 0$ such that all neighbors cost at least c to observe with probability at least δ, we do not have complete learning if $\mathcal{E} = \{X_1\}$ where X_1 is a free, unbounded signal. Hence, uniformly affordable overturning information is not sufficient for learning without 0-expanding observations. Furthermore, to ensure that there is not complete learning it is sufficient that for infinitely many agents, with at least probability δ, that they cannot observe any neighbors cheaper than $c > 0$. That the above condition holds for* ***all*** *agents is not necessary.*

Proof. See Appendix B. □

To see the intuition for this result, consider the simpler statement that if all agents must spend at least c to observe any neighbor, and have access to $\mathcal{E} = \{X_1\}$ where X_1 is a free, unbounded signal, then there is not complete learning. This follows from the fact that agents with access to only a single free, unbounded signal will necessarily observe it before observing any social signal, and that for some private signal realizations will choose not to observe the social signal. This ensures that with some probability every agent manages to match the state only with a probability bounded away from 1.

3.1 Comparative Statics

Beyond establishing under what conditions we should expect learning, we can now consider some comparative statics. Given our desire to understand what the effect of the internet will be on social learning when agents are rationally inattentive, we should clearly seek to know whether making networks either more connected or less connected, or increasing or decreasing agents' access to private information, will straightforwardly improve or damage the asymptotic accuracy of agents. Unfortunately, we answer this question in the negative: increasing or decreasing either the level of connectivity or access to private information can both increase and decrease the asymptotic accuracy of agents. The impact of the internet is thus not evident in such a model, and will require more assumptions on the exact nature of the network topology before and after the internet. Firstly, I consider the impact of making the network topology more or less connected in Proposition 5.

Proposition 5 (Network Topology Comparative Statics)**.** *Making the network either strictly less or strictly more cheaply connected can reduce the asymptotic probability with which agents match their action to the state.*

Proof. See Appendix B. □

In this proposition, I use two examples to demonstrate that increasing the connectivity of the network can either increase or decrease asymptotic accuracy (and thus decreasing it can do the same). Figure 2a illustrates a scenario in which increasing connectivity increase asymptotic accuracy, and Figure 2b the opposite.

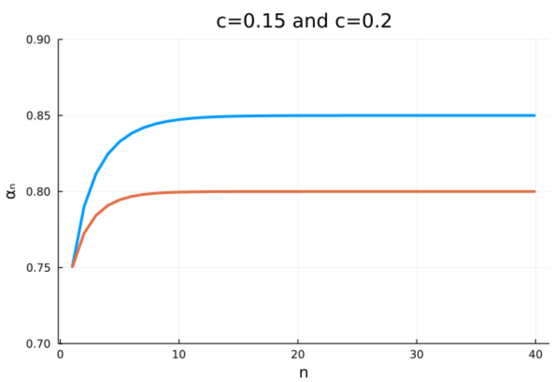

(a) The blue line here is with $c = 0.15$, and the orange $c = 0.2$.

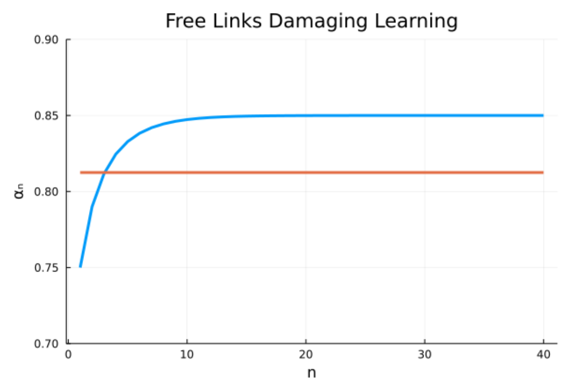

(b) The blue line here is with $c = 0.15$ before free links to agent 1 are added, and the orange after.

Figure 2: Greater connectivity can help or hinder.

In the first example, I take a setting in which agents have free access to a fixed, unbounded private signal. I also assume agents have a link budget of $L = 1$, and can pay price c to observe their immediate predecessor, but must pay a price greater than 1 to observe others. Given the specific private signal structure I assume, I show that agents will only choose to observe their predecessor if this cost is below 0.25, $c \leq 0.25$, and that below this point decreasing the cost of observing immediate predecessors smoothly increase asymptotic accuracy: $\alpha = 1 - c$. Above $c = 0.25$, there is no social learning at all, and all agents perform exactly as well as the first.

Whilst this example establishes that increasing connectivity can help asymptotic accuracy, that of Figure 2b shows the opposite. Taking the setting of our first example with $c = 0.15$ as a starting point, I then suppose that each agent can observe the first agent of the game for free (though the link budget prevents them observing both this agent and their immediate predecessor). Courtesy of this, agents are tempted away from observing their immediate predecessor, which imposes negative informational externalities on all following agents: where before this change every agent could choose to observe the last element of an information path containing $n-1$ agents, now they can either freely observe an information path of length 1, or pay c to observe one of length 2. Asymptotic accuracy drops from 0.85 to 0.8125.

Both of these examples are defined such that the accuracy of agents does in fact converge, but this is of course not necessarily going to be the case in an arbitrary network topology. In fact, by merging the two examples above one can easily create a game in which one subset of agents benefits as the network topology becomes more connected whilst the other becomes worse off. To see this, take the first agent (who faces the same information environment in all cases), and make him the first agent of a new game. Then, let each successive agent in the game of Figure 2a have the lowest even index remaining, and similarly let each agent in Figure 2b have the lowest odd index remaining. One can choose the social cost vectors such that agents in the first set never observe agents in the second and vice versa. Then, moving from the first scenario in each example to the second involves making the network topology strictly more cheaply connected, and producing a higher accuracy for all even-indexed agents, whilst ensuring a lower accuracy for all sufficiently late odd-indexed agents. Hence, it is clear that the consequences of making the network more connected will very much depend on the details of the environment.

Proposition 6. *Providing every agent with a strictly Blackwell-preferred experiment set-cost function pair can increase or decrease the asymptotic probability with which agents match the state.*

Proof. See Appendix 6 . □

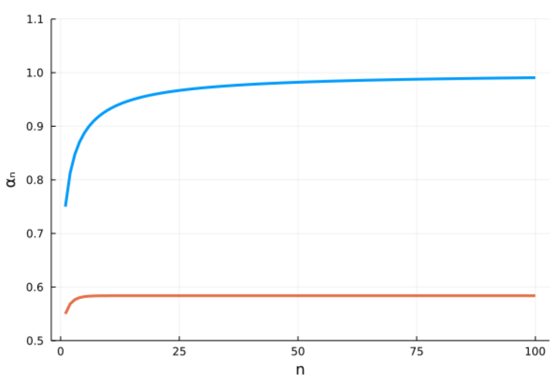

(a) In this example the blue curve displays a scenario in which agents all have a strictly strictly Blackwell-preferred experiment set-cost function pair to that of the orange curve. Here this improves accuracy, allowing complete learning in the higher-information case, but limited learning in the lower-information case.

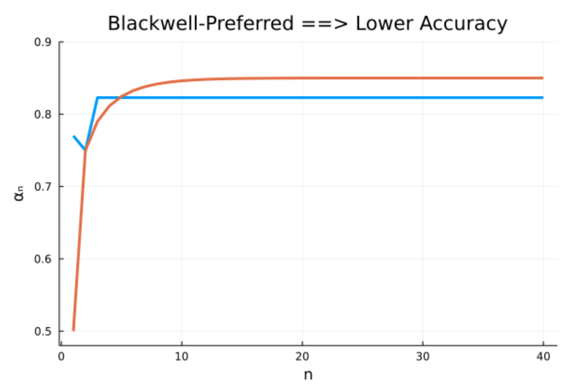

(b) The blue line here shows the path of accuracy in a scenario where every agent has a strictly Blackwell-preferred experiment set-cost function pair to that of the orange line, and identical network topologies. Learning is damaged here as agents end up all observing only the first agent.

Figure 3: More informative signals can help or hinder.

Finding an example in which increasing the access agents have to private information improves the asymptotic probability which they match the state is relatively straightforward. If we assume a network topology such that agents can only observe their immediate predecessors, this rules out any problem with agents switching to observing exclusively early neighbors. The specific example I use to establish the first part of this proposition is one in which agents can observe their immediate predecessor for free, but are unable to observe anyone else. In such a setting, it is easy to make asymptotic accuracy increase as we move to Blackwell-preferred signals, as if we give agents access to a bound signal in the first instance, but move to an unbounded Blackwell dominant signal structure in the second, it is clear that asymptotic accuracy will improve (from being strictly below 1 to 1 itself). In Figure 3a I show the accuracy paths that one can generate doing this.

Conversely, when the network topology is more flexible, the implications of improving agents' private signals are no longer unambiguously good. Figure 3b shows an example in which increasing the informativeness of all agents' signals leads to a lower level of asymptotic accuracy, precisely because it pushes the game away from an equilibrium which produces an observation network with expanding observations, to one which doesn't. Precisely, I choose a set-up here in which agents can either observe the first agent of the game free of charge, or their immediate predecessor at cost 0.15 (again with a link budget of $L = 1$). I provide agent 1 with a Bernoulli signal, and all other agents with unbounded signals $\{f_0(s), f_1(s)\} = \{2(1-s), 2s\}$; this yields the equilibrium represented by the orange line in Figure 3b. If we then increase the informativeness of agent 1's Bernoulli trial, whilst also giving all other agents a Blackwell-preferred signal structure (but one that is only marginally better) we can shift them from the immediate predecessor equilibrium to one in which all agents after $n = 2$ choose instead to observe agent 1. Whilst each agent is individually better off than if they had observed their immediate predecessor, these decisions together inflict a negative informational externality on the society since they prevent information ever accumulating. The endogenous observation structure is thus absolutely key to these comparative statics, and demonstrates the value of modelling these costs, and doing so in such general terms.

A broad conclusion we can draw from propositions 5 and 6 is that the impact of the internet on learning is not obvious with rationally inattentive agents. A general model does not allow us to make unambiguous statements about the consequences of our ever easier access to information, and more and more connected online discussion and social networks. In order to comment on the consequences of this, we will need to commit to a more specific model, and so in the next section I set out such a model of the information environment. Through this, I comment on the likely effects of the internet on the media landscape, and explain recent trends in opinion vs news journalism.

4 Social Media and Nonseparable Information Costs

The model of this paper assumes that searching for private and social information are separate activities, and hence that the attention costs of conducting them are separable. However, particularly in the world of social media, one could argue that they can be conducted simultaneously. There is evidence that, as is frequently observed in the media and our political discourse, that many people actually get their news from social media websites in the first instance. Pew Research Center polling from September 2025, for example, shows that at least 53% of Americans get at least some news from social media (PRC, 2025).

If this is the case, one could imagine that choosing to observe a private signal in reading one's newsfeed might make it very low cost to then observe what other agents have said on the topic of interest. Indeed, perhaps in some circumstances one can argue that searching for private information in this fashion makes observing the statements of others inevitable. In such a world, one could model the cost of observing a social signal *conditional* on having observed a particular private signal as being different to the cost of observing it alone.

Suppose we augment our model by supposing that in addition to private and social signals, agents can choose instead to observe a *joint* signal, for example $\{x_{n-1}, \{f_0(s), f_1(s)\}\}$, which involves observing their immediate predecessor and an unbounded private signal. Using Proposition 2, we can see a new path to complete learning. If this joint signal is sufficiently attractive (relative to the separate social and private signals that are available) that all agents will choose to observe it, then it produces both a network with expanding observations and provides unbounded private signals. Hence, it will produce learning.

What's more, we could see a threat to learning via any such channel in the improvement of separate social *or* private signals alone. If at one point in time agents are all consuming joint signals and this produces learning, but we then increase the quality of private signals, perhaps by inventing ChatGPT and increasing the effectiveness of search off social media, then in equilibrium agents may no longer consume the joint signal. Just as with Proposition 6, improving private signals alone can ensure that the equilibrium network no longer exhibits expanding observations. This would ensure the high-quality social signals available to high-index agents in the initial equilibrium will no longer exist, since agents along the relevant improvement path will simply have engaged in private search instead.

Many of the results of this paper, and indeed of Section 5 next, suggest that social media likely has a negative impact on social learning. This section raises a notable counterpoint to that. As I have discussed, with nonseparable costs of information one could even envisage a golden age of social-media-based social learning being broken by the advent of AI-enhanced private search.

5 The Internet, Influencers, and a Changing Media Landscape

Here we can think of a model in which the first M agents are newspapers. They conduct private investigations (a private signal of a certain cost), but are unable to observe each other since they publish at the same time. Then the following agents are people, who can observe newspapers for a price, but also each other. To represent the impact of the internet, I will consider the consequences of moving from a set-up in which it is extortionately expensive for individuals to conduct private investigations, and even observe each others' opinions on the true state, to one in which both private investigation and observing predecessors become cheaper.

I will suppose that the media companies are capable of of either conducting an investigation (a Bernoulli trial with success parameter q_I) or not reporting at all $\{I, \emptyset\}$ at costs $\{c_I, 0\}$ where $0 \leq q_I < c_I$ (the cost of the report is more than its value to a single agent, and multiple agents must observe it before it is worthwhile to produce). I will model media company j as making an expected revenue of $\sum_{n=M+1}^{\infty} \delta^{n-M} p_j \mathbb{I}(j \in B(n))$, where $\delta \in [0, 1]$ reflects the probability with which the game will end each period (the news cycle will move on, no more agents will search this hashtag). Suppose that agents can conduct an individual Bernoulli experiment $\{X_1\}$ with success parameter q_X, but that initially the cost of doing so is extortionate: $C_n^P(X_1) > 1$ for all $n > M$. In the language of my broader model, $\mathcal{E} = \{X_1, I\}$, $C_n(I) = c_I$ for $n \leq M$. Agents must pay a price p_m to observe media company $m \in \{1, ..., M\}$, in addition to any cognitive cost, so the total cost of doing so is strictly greater than $p_m + p_{m-1}$. For the sake of parsimony, let us suppose that the entire cost of consuming a newspaper report is simply its price, and that this cost is therefore additive, and moreover let us set the link budget to infinity. I suppose that the media companies cannot observe each other $c_{m,m-j} = \infty$ for $m \in \{1, ..., M\}, j \in \{1, ..., m-1\}$. I solve for the Perfect Bayesian Equilibrium of this game and discuss how it changes as we adjust parameters to reflect the advent and development of the internet.

The immediate impact of the internet was to make searching for information much cheaper and easier than it was before, so in this model its first consequences will be reductions in $C_n^P(X_I)$ and c_I. The advent of social media shall instead reduce the cost of observing predecessors.

Definition 4 (Stages of the Development of the Internet)**.** *I divide the history of the internet into the following stages:*

1. *Pre-Internet:* $C_n^P(X_1) > 1$ *for all* $n > M$, $c_I = c_I(0)$, *neighbor observation costs are above* 1 *for all agents.*

2. *Early Internet*[a]*:* $C_n^P(X_1) = C^P(X_1) < q_x - 0.5 < 1$[b] *for all* $n > M$, $c_I = c_I(1) < c_I(0)$, *neighbor observation costs are still above* 1 *for all agents.*

3. *Post Social Media: neighborhood observation costs decrease, so that each agent can cheaply observe either recent predecessors, or influencers (early agents observed by all or many successors), or both. For simplicity, set these costs to zero:* $c_{n,n-1}, ..., c_{n,n-j} = 0$ *or* $c_{n,M+1}, ..., c_{n,M+\#Influencers}$ *or both; all other observation costs remain the same.*

[a]Crucially this represents the internet *before* Social Media comes about.

[b]Without this condition the agents prefer to observe no information anyway, and this stage is identical to the Pre-Internet Stage.

Proposition 7. *If $M = 1$, the following are true:*

1. *In the Pre-Internet stage, the newspaper investigates and agents all read the report if $c_I \leq \frac{\delta}{1-\delta}(q - 0.5)$, and agents match the state with accuracy $q_I > 0.5$. Otherwise agents receive no information, and their asymptotic accuracy is $\alpha_n = 0.5$ for all n. If they report, the newspaper sets its price at $p_1 = q_I - 0.5$.*

2. *In the Early Internet stage: the newspaper investigates if $c_I \leq \frac{\delta}{1-\delta}((q_I - q_X) + C^P(X_1))$ and charges price $(q_I - q_X) + C^P(X_1)$. Agents have accuracy q_I if the report is produced (and strictly positive consumer surplus), and accuracy q_X otherwise (and positive consumer surplus).*

3. *In the Post Social Media stage: The media company conducts no investigation, agent $n = 2$ observes a private experiment, and all following agents copy his action after observing him either directly or indirectly. Their accuracy is q_I.*

Proof. The Pre-Internet Equilibrium: In the first instance, no consumers will conduct their own private investigations or observe any neighbors, since these are so expensive. Either they will consume no information at all, or they will read the newspaper if the cost of doing so is cheap enough, i.e. if:

$$\mathbb{E}U(B(n) = \{1\}) = 1 \times \alpha_{News}(n) - p_1 \geq 0.5$$

The news company will choose to conduct an investigation if their expected revenue is higher than c_I:

$$\sum_{n=2}^{\infty} \delta^{n-1} p_1 \mathbb{I}(1 \in B(n)) - c_I \geq 0$$
$$\sum_{n=2}^{\infty} \delta^{n-1} p_1 \mathbb{I}(q - p_1 \geq 0.5) - c_I \geq 0$$

Where I use that if the news company chooses to investigate and observes a Bernoulli signal, and chooses $x_n = \varsigma$ to communicate this information in its news report, the success probability of agents reading this report will be q. Clearly the news organisation will choose $p_1 = q - 0.5$ in any equilibrium where they choose to investigate, since this maximises their profit whilst ensuring that all agents consume the news report. Their profit in the event they investigate will then be:

$$\sum_{n=2}^{\infty} \delta^{n-1}(q - 0.5) - c_I = \left(\frac{1}{1-\delta} - 1\right)(q - 0.5) - c_I$$

The media company investigates if $c_I \leq \frac{\delta}{1-\delta}(q - 0.5)$, where the RHS is increasing in δ and q.

The Early Internet Equilibrium:

The advent of the early internet can give agents a better outside option than before. Agents will never consumer observe both their private experiment and the news report. If $q_X > q_I$, and the agent knows that upon observing the outcome of q_I (which will produce a posterior they can anticipate, and which is symmetric) they will want to consumer X and follow that, it is dominant strategy to skip the media investigation and observe X_1. If $q_X > q_I$ the reverse is true, and if $q_X = q_I$ observing a second piece of information will either put them in the position of indifference between the two actions (in which case they may as well have chosen the action recommended first and spared themselves the price of the second piece of information), or will simply confirm what they were going to do anyway.

Therefore, in equilibrium the media company will either price its report such that agents consume it instead of their signal, or choose not to produce one at all. The equilibrium therefore involves individuals choosing to read the report, and the media company choosing its price such that:

$$q_X - C^P(X_1) = q_I - p_1$$
$$p_1 = (q_I - q_X) + C^P(X_1)$$

The media company then makes profit $\frac{\delta}{1-\delta}((q_I - q_X) + C^P(X_1)) - c_I$, and thus produces the report if this is greater than 0. Agents again have accuracy q_I if the report is produced, and q_X otherwise.

The Post Social Media Equilibrium:

The Social Media Stage ensures that once agent has read the news report (if one is produced), all following agents will observe either them or their immediate predecessor to learn the action recommendation of the report. Hence the media company cannot sell to more than 1 agent, and since $c_I > q_I$, this is not enough to produce a positive profit. The media company does not report, and agent $n = 2$ conducts his own investigation and all following agents observe him (directly or indirectly, for the newsfeed and influencer networks respectively). If individuals are less effective than the media at investigating, all agents are less accurate.

□

There are a number of points that merit discussion here. Firstly, the conclusion of this result is that, in the monopoly case, the advent of the internet is good for consumers, since it prevents the monopoly extracting all of the surplus that their report generates. It can, however, reduce the overall accuracy of agents if their private investigation is less informative than the media report, and cheap enough to make the media company's business model unviable. The advent of social media, however, is clearly negative in this model. The ability to share the media report's action recommendation takes a non-rivalrous good, the information in that report, and renders it non-excludable as well: i.e. it creates a public good problem. Since agents can no longer be made to pay for the media report, it is no longer produced, and the consumer surplus and accuracy of all agents drop.

It is also worth noting that 'failing to report' in this model is not necessarily analogous to news companies going out of business, though it could be. Another interpretation, given that in reality agents also read newspapers for entertainment and the satisfaction of opinion journalism, is that the media company simply stops producing a good with any informational value. News divisions could be shuttered in favour of commentary that appeals to the political bent of readers. Noticing the relevance of ideology to the decision of which newspaper agents choose to read, this also suggests the monopoly setting is more relevant to the real world than it might seem: if the Wall Street Journal and the New York Times are the two newspapers available, and Republican and Democratic readers dislike to see non-congenial commentary, each newspaper could be thought of as playing the above monopoly game within their own political tribe.

The basic problem here is in some senses robust to the nature of the social network. Firstly, it does not matter if agents are connected to each other in some sort of line (reflecting a series of agents reading their newsfeeds as I suggest above), or something closer to a start network in which agent $n = 2$ is playing the role of an influencer. Were these agents not to form a connected graph, this would increase the revenue of the media company (if agent 300 is not connected to any of his predecessors, he will buy the report), but it would still be much lower than its pre-social media revenue. Similarly, the complete collapse of their readership depends here on the fact that I have assumed that social observation costs are exactly zero, and this is an extreme assumption. However, were I to relax this, and simply place a very low cost of these observations, this would still ensure that the newspaper would need to price below the cost of observing a single neighbor. For example, if an agent would pay one dime to avoid the trouble of logging into Twitter, this is as much as the newspaper can charge. δ and c_I would not need to take high values before the same market collapse occurs here as well. If we take higher values of M, and suppose that media firms choose their strategies simultaneously (since sequential moves on behalf of these companies seems a little artificial), what will remove the ability of the monopolist in the pre-internet stage to extract the entirety of the surplus a single report produces. This implies that the impact of introducing the early internet will not be so beneficial, but does nothing to prevent the public goods problem that comes with the social media. Finally, we can also observe here that whilst in this model individuals would be rational in observing an influencer who had observed a newspaper report, the same breakdown in investigation incentives will occur if agents simply misperceive influencers as having more information. Proposition 7 demonstrates that even without people particularly valuing the opinions of influencers, their presence can still cause this collapse.

6 Conclusion

That individuals have greater access to information than ever does not imply that they are consuming more of it. In this article, I have considered the implications of Rational Inattention for social learning on the internet. Even in a setting in which agents are otherwise completely rational, I have shown that reducing the barriers to searching for information and observing that shared by others does not necessarily produce equilibria in which agents match the state with higher probability.

Whilst it is possible that reducing these costs leads to agents achieving higher accuracy asymptotically, Propositions 5 and 6 establish that we can see the exact opposite of this: giving agents access to strictly more informative private signals, and making the network topology unambiguously more connected, can in fact reduce the amount of information agents observe. As influencers become easier to observe or more informative, the incentive to observe them instead of more recent agents can impede the formation of the large connected observation networks necessary for information aggregation to take place. This holds even in a model of rational agents with a correctly specified model of the world, and as I have noted, if in reality agents misperceive influencers as being more informed than they in fact are, such problems will be all the more severe.

Moving beyond a model of information aggregation in the standard social learning mould, I have also shown that the development of the internet and social media can create a market failure in the media industry. By transforming news reports in a sort of public good, the ability to share their information content on social media can cause the equilibrium readership of such companies to collapse, as I show in Proposition 7. A move away from investigative to opinion journalism can be interpreted as a symptom of this. With rationally inattentive agents, in short, Huxley's dystopia carries the day.

References

(2025): "Social Media and News Fact Sheet," Accessed: (2025-12-14).

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

Alesina, A., A. Miano, and S. Stantcheva (2020): "The Polarization of Reality," *AEA Papers and Proceedings*, 110, 324–28.

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

Banerjee, A. V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*.

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

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

Caplin, A., M. Dean, and J. Leahy (2019): "Rational Inattention, Optimal Consideration Sets, and Stochastic Choice," *The Review of Economic Studies*, 86, pp. 1061–1094.

Caplin, A., J. Leahy, and F. Matejka (2016): "Rational inattention and inference from market share data," *Unpublished working paper. NYU and CERGE-EI*.

Cremin, J. W. (2023): "Too Much Information & the Death of Consensus," .

——— (2025): "Bot Got Your Tongue: Social Learning with Timidity and Noise," .

Hu, L. and A. Li (2018): "The politics of attention," *arXiv preprint arXiv:1810.11449*.

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

——— (2016): "Preferences, Homophily and Social Learning," *Operations Research*.

Lomys, N. (2020): "Collective search in networks," *Available at SSRN 3197244*.

Maćkowiak, B., F. Matějka, and M. Wiederholt (2023): "Rational inattention: A review," *Journal of Economic Literature*, 61, 226–273.

Matějka, F. and G. Tabellini (2021): "Electoral competition with rationally inattentive voters," *Journal of the European Economic Association*, 19, 1899–1935.

Matějka, F. and A. McKay (2015): "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 105, 272–98.

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

Nimark, K. P. and S. Sundaresan (2019): "Inattention and belief polarization," *Journal of Economic Theory*, 180, 203–228.

Novák, V., A. Matveenko, and S. Ravaioli (2021): "The status quo and belief polarization of inattentive agents: Theory and experiment," .

Postman, N. (2005): *Amusing ourselves to death: Public discourse in the age of show business*, Penguin.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50, 665–690.

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

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

Yuksel, S. (2022): "Specialized learning and political polarization," *International Economic Review*, 63, 457–474.

Çelen, B. and S. Kariv (2004): "Observational learning under imperfect information," *Games and Economic Behavior*, 47, 72–86.

A Useful Lemmas

The first two lemmas here reproduce important results in Acemoglu et al. (2011), which I also use in Cremin (2023, 2025). In Acemoglu et al. 2011, each agent is endowed with the same private signal and an exogenous neighborhood; here, therefore, they do not directly apply. There are some examples I construct in this paper where they do apply almost directly, as I give agents only a single private signal and a choice of only a very limited variety of neighborhoods.

Lemma 1 (Acemoglu et al. Proposition 2)**.** *Agent n will choose $x_n = 1$ upon observing neighborhood $B(n)$ and private signal s_n if:*

$$\mathbb{P}\big(\theta = 1 | B(n)\big) + \mathbb{P}(\theta = 1 | s_n) > 1$$

Proof. See (Acemoglu et al., 2011, Proposition 2). □

This next lemma is (Acemoglu et al., 2011, Lemma 1), and gives several useful properties of the belief distributions.

Lemma 2 (Acemoglu et al. 2011 Lemma 1)**.** *The private belief distributions, $\mathbb{G}_0$ and $\mathbb{G}_1$, satisfy the following properties:*

(a) For all $r \in (0,1)$, $d\mathbb{G}_0(r)/d\mathbb{G}_1 = (1-r)/r$

(b) For all $0 < z < r < 1$, $\mathbb{G}_0(r) \geq ((1-r)/r)\mathbb{G}_1(r) + ((r-z)/2)\mathbb{G}_1(z)$

(c) For all $0 < r < w < 1$, $1 - \mathbb{G}_1(r) \geq (r/(1-r))(1-\mathbb{G}_0(r)) + ((w-r)/2)(1-\mathbb{G}_0(z))$

(d) The term $\mathbb{G}_0(r)/\mathbb{G}_1(r)$ is nonincreasing in r and is strictly larger than 1 for all $r \in (\underline{\beta}, \overline{\beta})$

Proof. See (Acemoglu et al., 2011, Lemma 1). □

This next lemma is a result I derive in Cremin (2023) for my model of motivated reasoning. Here it is useful in examples where I provide either a single free neighborhood, or at least give agents a very limited choice of neighborhoods:

Lemma 3 (Bayesian Social Belief Distribution Relationship)**.** *If the Bayesian social belief of agent n in state θ has PMF $h^n_\theta(\cdot)$, they obey the following relation:*

$$h^n_1(SB_n)(1 - SB_n) = h^n_0(SB_n)SB_n$$

Proof. This follows almost exactly the proof of (Acemoglu et al., 2011, Lemma A1 (a))- adjusted in necessary ways. By then definition of a Bayesian social belief, we have for any $sb_n \in (0,1)$:

$$\mathbb{P}(\theta = 1|SS_n) = \mathbb{P}(\theta = 1|SB_n)$$

Using Bayes' Rule, it follows that:

$$SB_n = \mathbb{P}_\sigma(\theta = 1|SB_n) = \frac{\mathbb{P}_\sigma(SB_n|\theta = 1)\mathbb{P}_\sigma(\theta = 1)}{\sum_{j=0}^{1}\mathbb{P}_\sigma(SB_n|\theta = j)\mathbb{P}_\sigma(\theta = j)}$$

(*Note this differs from the analogous expression in Acemoglu et al. (2011) since there are only a finite number of possible Bayesian social beliefs at any point.)

$$SB_n = \frac{\mathbb{P}_\sigma(SB_n|\theta = 1)}{\mathbb{P}_\sigma(SB_n|\theta = 0) + \mathbb{P}_\sigma(SB_n|\theta = 1)}$$
$$\mathbb{P}_\sigma(SB_n|\theta = 1) = [\mathbb{P}_\sigma(SB_n|\theta = 0) + \mathbb{P}_\sigma(SB_n|\theta = 1)]SB_n$$

Using the notation that h^n_θ is the probability mass function for the Bayesian social beliefs of agent n in state θ:

$$h^n_1(SB_n)(1 - SB_n) = h^n_0(SB_n)SB_n$$

□

B Omitted Proofs

Proof of Proposition 1. An agent makes at most three decisions in sequence: (1) to observe a private (social) signal, (2) to observe a social (private) signal, and (3) to choose $x_n \in \{0, 1\}$; he can simply choose x_n after observing either just one signal, or no signal at all. If he observes a private signal in the first sub-period, he has a finite set of social signals to observe (since he has a finite number of predecessors, and thus the power set is also finite), each of which implies an expected value for his decision problem, as well as the option to choose x_n straight away. The finite set of options guarantees the existence of an optimal action. If instead he chooses to observe a social signal first, he can again observe one of finitely many private signals (this is finite by assumption), each of which implies an expected value for his decision problem. In the first sub-period, he can choose amongst finitely many private or social signals, each of which has an expected value implied by the optimal second stage choice for each possible signal realization. Hence an optimal strategy always exists.[10]

There is no strategic interaction in social learning games, so recursively applying the above argument establishes equilibrium existence.

□

Proof of Proposition 2. To prove this, all that is necessary is to establish that standard improvement principles can apply here (that from Acemoglu et al. (2011) is the template I have in mind). There are two differences that must be addressed:

1. The assumptions on the network topology are in terms of the neighborhood of zero cost agents.

2. Agents do not have free access to a single free unbounded signal.

For the first, the adjustment to be made is straightforward: whereas standard improvement principle involve agents improving on the best performing agent in their neighborhood, here we must think of them improving on the best neighbor that they can observe for free. Where Acemoglu et al. prove an 'Information Monotonicity' Lemma in which agents must outperform the best of their neighborhood $\mathcal{B}$, here they must outperform the best of $\mathcal{B}^0$, where this includes only free agents:

$$\mathbb{P}_\sigma(x_n = \theta | B^0(n) = \mathcal{B}^0) \geq \max_{b \in \mathcal{B}} \mathbb{P}_\sigma(x_b = \theta)$$

[10]Which tie-breaking rule agents use when different strategies given the same expected utility is unimportant; we can assume they uniformly randomise between them.

For the second, for any level of accuracy $\beta \in [0.5, 1)$, we can apply the standard improvement principle to show that accuracy must grow beyond that β. To do this, we can simply note that there is at least some (unbounded in the 'overturning direction') private signal that provides over-turning information for any binary social signal of strength β that is uniformly affordable (by the assumption of uniform affordability). This may not be the same signal for $x_b = 0$ as $x_b = 1$. Let $p_S > 0$ be the minimum probability with which an agent observes the relevant over-turning private signal given the value of x_b (note that this is conditionally independent of the value of their chosen predecessors action, and that uniform affordability tells us that this is strictly positive). In this instance, they will do at least as well as an agent observing a signal that produces the value $\#N/A$ with probability $1 - p_S$, and realises the relevant overturning private signal with probability p_S: the available signal is necessarily unbounded.[11] Equipped with this translation, we can see that the standard improvement principle argument in Acemoglu et al. 2011 ensures that $\alpha_n \to_n \alpha > \beta$, and that since this is true for arbitrary $\beta \in [0, 1)$, this establishes that $\alpha_n \to_n 1$.

□

Proof of Remark 3.1. As I note in the main text, the problem here is that we can conceive of information cost structures such that though each individual agent n has affordable overturning information, the probability with which they can afford such overturning information converges to zero fast enough to prevent the accuracy of agents α_n ever converging to 1. To see this, consider the following scenario.

Let agents exist in a line network, where each agent can certainly observe their predecessor for free, but would have to pay more than 1 (their entire possible payoff) to observe anyone else. Suppose that the set of experiments available to agents is $\mathcal{E} = \{X_1\}$ where X_1 has density functions $\{f_0(\varsigma), f_1(\varsigma)\} = \{2(1-\varsigma), 2\varsigma\}$. Assume further that with probability $\frac{1}{e^n}$ agent n can observe X_1 for free, and with probability $1 - \frac{1}{e^n}$ it costs more than 1: $\Delta_P^n(C_P^n(X_1) = 0) = \frac{1}{e^n}$ and $\Delta_P^n(C_P^n(X_1) = 0) = \frac{1}{e^n}$. Note that in this set-up overturning information is affordable for each agent, since for any social belief their is some probability they observe the unbounded signal of X_1, and in each state their is some fixed strictly positive probability this signal overturns any belief.

In this setting, we can first note that if an agent does in fact get to observe X_1, their accuracy (conditional on this) will be related to α_{n-1} by the function $H(\alpha_{n-1}) = \alpha_{n-1}^2 - \alpha_n + 1$ (the working for this can be seen in *The Unbounded Signal Case* in the proof of Proposition 6). If they do not get to observe X_1, they will match the state with the same probability as their predecessor, as they will simply copy him. Hence:

$$\begin{aligned} \alpha_n &= (1 - \frac{1}{e^n})\alpha_{n-1} + \frac{1}{e^n} H(\alpha_{n-1}) \\ &= \alpha_{n-1} + \frac{1}{e^n}\Big(H(\alpha_{n-1}) - \alpha_{n-1}\Big) \end{aligned}$$

First of all, note that $H(\alpha) > \alpha$ for any $\alpha \in [0.5, 1]$ (any agent can never do worse than tossing a coin, so this is the domain of α), so agents are always managing to improve on their predecessors,

[11] The agent could face one affordable signal that can generate arbitrarily strong private beliefs in one direction, and a separate signal that generates arbitrarily strong private beliefs in the other, without having access to any individual signal that is unbounded.

and this improvement is always strict.

$$\alpha_n < \alpha_{n-1} + \frac{1}{e^n}\Big(0.8125 - 0.75\Big)$$

The gradient of H is strictly decreasing, so the improvement from α_2 to α_3 given by H is strictly below that of the improvement from α_1 to α_2, hence $H(\alpha_{n-1} - \alpha_{n-1}) < 0.8125 - 0.75 = 0.0625$ for any $n \in \{3, ...\}$.

$$\alpha_n < \alpha_{n-1} + \frac{1}{e^n}\Big(0.0625\Big)$$
$$\alpha_n < \alpha_1 + 0.0625 \sum_{i=2}^{n} \frac{1}{e^i}$$

Let us call the limiting value of α_n, α, then we have that:

$$\alpha = 0.75 + 0.0625 \frac{1}{e-1} = 0.786(3d.p.)$$

Thus we do not have complete learning here, even though every agent has affordable overturning information. The path of α is display in Figures 1a and 1b by the orange lines. In the blue line I show the path of alpha if we replace the probability $\frac{1}{e^n}$ with $\frac{1}{n+1}$, which also fails to produce complete learning.

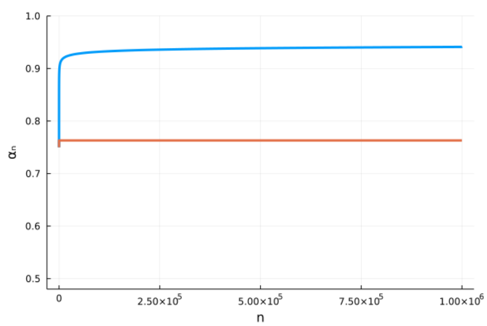

(a) Path over the first million agents

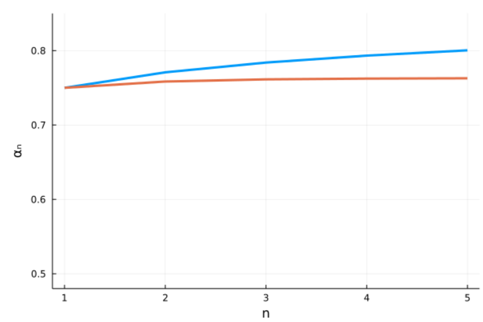

(b) Path over the first five agents

Figure 4: In the scenarios represented by both of these curves, agents all have access to affordable overturning information, but nonetheless they do not learn.

□

Proof of Proposition 3. The necessity of affordable information is trivial; if agents observe no private information they cannot learn at all. To see that affordable overturning information is not necessary, consider a network of agents in which each agent n can observe all of his predecessors at zero cost with probability p_n, but otherwise must pay some very large cost C to observe even a single predecessor[12]. Suppose that each agent can observe a bounded private experiment at zero cost, but no other information is available. If we further impose that $p_n \to 1$ as $n \to \infty$, but that it does so

[12] C must be high enough that the agents never choose to pay this cost to observe predecessors.

sufficiently slowly (e.g. $p_n = 1 - \frac{1}{n}$) that agents face cost C to observe predecessors infinitely often, then we have complete learning by the same martingale arguments that establish Theorem 3 in Cremin (2025). The argument behind Theorem 4 in Cremin (2023) also works to establish this. □

Proof of Proposition 4. Before moving on to the complete statement, let us consider instead settings in which agents must certainly spend at least c to observe any neighbor. This can then be easily extended.

To see the truth of this proposition, first note that agents will observe the private signal first, since it is free, and may spare them the cost of observing the social signal at all. Since this signal is unbounded, with non-zero probability it will generate a private belief in the set $[c, 1-c]^C$, where c is the minimum cost of observing a neighbor. Having formed this strong belief, the agent could simply choose their preferred action and gain an expected utility greater than $1 - c$. Given this, they can only perform worse by spending c (or more) to observe a predecessor, even if that predecessor were a perfectly revealing signal of the state of the world.

Denote as $\mathbb{P}_c(\theta)$ the probability that X_1 produces a signal within $[c, 1-c]^C$ in state θ, it follows that the ex-ante probability with which the agent matches the true state is bounded above as follows:

$$\begin{aligned}
\alpha_n \leq\ & \mathbb{P}(\theta = 0)\big\{\mathbb{P}_c(0)\mathbb{P}(\varsigma < 0.5 | \varsigma \in [c, 1-c]^C, \theta = 0) + \big(1 - \mathbb{P}_c(0)\big) \times 1\big\} \\
& + \mathbb{P}(\theta = 1)\big\{\mathbb{P}_c(1)\mathbb{P}(\varsigma \geq 0.5 | \varsigma \in [c, 1-c]^C, \theta = 1) + \big(1 - \mathbb{P}_c(1)\big) \times 1\big\} \\
\alpha_n \leq\ & \frac{1}{2}\big\{\mathbb{P}_c(0)\mathbb{P}(\varsigma < 0.5 | \varsigma \in [c, 1-c]^C, \theta = 0) + 1 - \mathbb{P}_c(0)\big\} \\
& + \frac{1}{2}\big\{\mathbb{P}_c(1)\mathbb{P}(\varsigma \geq 0.5 | \varsigma \in [c, 1-c]^C, \theta = 1) + 1 - \mathbb{P}_c(1)\big\}
\end{aligned}$$

Each bracketed term is strictly less than 1, so it follows that α_n is bounded away from 1 for arbitrary n.

Note that if instead we had that agents only need to pay a minimum neighbor-observation cost of c with some probability $\delta > 0$, we can make the same argument. The probability that is bounded away from one is now their success probability in the event that they receive a private signal outside of $[c, 1-c]^C$ *and* must pay this cost of c. The relevant probabilities are then $\mathbb{P}_c(0) \times \delta$ and $\mathbb{P}_c(1) \times \delta$.

Finally, to prevent complete learning, we only need to be able to make this argument for infinitely many agents, rather than all of them. This produces the final statement of the proposition.

□

Proof of Proposition 5. Each statement can be established with an example:

(1) Making a network ***more*** *connected can* ***increase*** *asymptotic accuracy*:

Let us suppose that the private information environment is as follows:

- $\mathcal{E} = \{X_1\}$
- $C_n^P(X_1) = 0$

- X_1 has probability density functions $\{f_0(s), f_1(s)\} = \{2(1-s), 2s\}$

In words, I suppose that all agents have free access to a single unbounded experiment with the signal structure given above, but no other private signal is available. As for the network topology, let us say that at first it involves agents being able to observe only their immediate predecessor at cost $c_1 > 0$, but in the second instance shall cost $c_2 < c_1$.

- At first $\forall n \; c_{n,n-1} = c_1 > 0$, but $c_{n,n-j} > 1 \; \forall j \in \{2, n-1\}$
- Then the network topology shall change such that $\forall n \; c_{n,n-1} = c_2 < c_1$, but otherwise the social cost functions are the same.

Observe that clearly the second topology is strictly more connected than the first, given definition 8. Also note that every agent will always choose to observe their private signal, since it is free, and will necessarily do so first, since for some sufficiently strong signal realizations it will spare them the expense of observing the social signal. Given the symmetry of the private signal structure, it is without loss to suppose that $\theta = 1$.

For which private signal realizations will our agent choose not to observe their predecessor? If they observe signal realization ς, by the assumed normalization their interim belief is ς. Suppose without loss that this is greater than 0.5 and the agent is choosing $x_n = 1$. Their social signal choice problem is to decided whether or not to ignore the social signal given that belief. If they choose to ignore the social signal, they gain expected utility:

$$\mathbb{E}U(B(n) = \{n-1\}) = \varsigma \times 1 + (1-\varsigma)0 = \varsigma$$

If they choose to observe it, they instead get:

$$\mathbb{E}U(B(n) = \{n-1\}) = \Big\{ \mathbb{P}(x_n = 0, \theta = 0 | \text{Observed } x_{n-1}, s = \varsigma) \times 1 \tag{B.1}$$

$$+ \mathbb{P}(x_n = 1, \theta = 1 | \text{Observed } x_{n-1}, s = \varsigma) \times 1 \Big\} + 0 - c_1 \tag{B.2}$$

Suppressing the 'Observed' for brevity, we can compute these probabilities:

$$\begin{aligned}\mathbb{P}(x_n = 0, \theta = 0 | x_{n-1}, s = \varsigma) &= \mathbb{P}(x_n = 0 | \theta = 0, x_{n-1}, s = \varsigma)\mathbb{P}(\theta = 0 | x_{n-1}, s = \varsigma) \\ &= \big[\mathbb{P}(sb_n(0) < 1 - \varsigma | \theta = 0)\mathbb{P}(x_{n-1} = 0 | \theta = 0) \\ &\quad + 0 \times \mathbb{P}(x_{n-1} = 1 | \theta = 0)\big]\mathbb{P}(\theta = 0 | s = \varsigma)\end{aligned}$$

The first probability on the right-hand side here can be expressed in this fashion using Lemma 1, where $sb_n(0)$ denotes the social belief of n after observing $x_{n-1} = 0$. If agent $n-1$ chose $x_{n-1} = 1$, n will clearly choose $x_n = 1$, hence the zero. In the second probability "Observed x_{n-1}" is a choice variable, and thus clearly carries no additional information for the agent. We can also see immediately that $\mathbb{P}(\theta = 0|s = \varsigma) = 1 - \varsigma$, use from above that $sb_n(0) = 1 - \alpha_{n-1}$, and note that the symmetry of the problem gives that $\mathbb{P}(x_{n-1} = 0|\theta = 0) = \alpha_{n-1}$.

$$\begin{aligned}
\mathbb{P}(x_n = 0, \theta = 0|x_{n-1}, s = \varsigma) &= \mathbb{I}\big(1 - \alpha_{n-1} < 1 - \varsigma\big)\alpha_{n-1}(1 - \varsigma) \\
&= \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)\alpha_{n-1}(1 - \varsigma) \\
\mathbb{P}(x_n = 1, \theta = 1|\text{Observed } x_{n-1}, s = \varsigma) &= \mathbb{P}(x_n = 1|\theta = 1, x_{n-1}, s = \varsigma)\mathbb{P}(\theta = 1|x_{n-1}, s = \varsigma) \\
&= \big[\mathbb{P}(sb_n(0) \geq 1 - \varsigma|\theta = 1)\mathbb{P}(x_{n-1} = 0|\theta = 1) \\
&+ 1 \times \mathbb{P}(x_{n-1} = 1|\theta = 1)\big]\mathbb{P}(\theta = 1|s = \varsigma) \\
&= \big[\mathbb{I}(1 - \alpha_{n-1} \geq 1 - \varsigma)(1 - \alpha_{n-1}) + 1 \times \alpha_{n-1}\big]\varsigma \\
&= \big[\mathbb{I}(1 - \alpha_{n-1} \geq 1 - \varsigma)(1 - \alpha_{n-1}) + \alpha_{n-1}\big]\varsigma
\end{aligned}$$

Feeding this back into Equation B.2, we get:

$$\begin{aligned}
\mathbb{E}U(B(n) = \{n-1\}) &= \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)\alpha_{n-1}(1 - \varsigma) + \big[\big(1 - \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)\big)(1 - \alpha_{n-1}) + \alpha_{n-1}\big]\varsigma - c_1 \\
&= \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)\alpha_{n-1}(1 - \varsigma) + \big[(1 - \alpha_{n-1}) - \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)(1 - \alpha_{n-1}) + \alpha_{n-1}\big]\varsigma - c_1 \\
&= \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)\alpha_{n-1}(1 - \varsigma) + \big[1 - \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)(1 - \alpha_{n-1})\big]\varsigma - c_1 \\
&= \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)\alpha_{n-1}(1 - \varsigma) - \big[\mathbb{I}\big(\alpha_{n-1} > \varsigma\big)(1 - \alpha_{n-1})\big]\varsigma + s - c_1 \\
&= \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)(\alpha_{n-1} - \alpha_{n-1}\varsigma) - \big[\mathbb{I}\big(\alpha_{n-1} > \varsigma\big)(\varsigma - \alpha_{n-1}\varsigma)\big] + \varsigma - c_1 \\
&= \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)(\alpha_{n-1} - \alpha_{n-1}\varsigma - \varsigma + \varsigma\alpha_{n-1}) + \varsigma - c_1 \\
&= \mathbb{I}\big(\alpha_{n-1} > \varsigma\big)\big\{\alpha_{n-1} - \varsigma\big\} + \varsigma - c_1
\end{aligned}$$

Hence, upon observing $\varsigma > 0.5$ the agent will choose to observe their immediate predecessor after observing ς if:

$$\mathbb{I}\big(\alpha_{n-1} > \varsigma\big)\big\{\alpha_{n-1} - \varsigma\big\} - c_1 \geq 0$$

or if both $\alpha_{n-1} > \varsigma$ and $\alpha_{n-1} \geq \varsigma + c_1$. Clearly the latter implies the former, and highlights the distortion introduced by adding a neighbor-observation cost of c_1, this introduces a wedge between α_{n-1} and ς such that the agent will not observe the signal whenever $\alpha_{n-1} > \varsigma$, but instead needs this stricter condition to be satisfied.

By symmetry, if they observe a private signal below 0.5, they will observe their predecessor if $(1 - \alpha_{n-1}) \leq \varsigma - c_1 \iff \varsigma \geq 1 - \alpha_{n-1} + c_1$.

Here I am assuming they break ties in favour of observing their predecessor.

There are three distinct paths to choosing correctly:

- The get such a strong ς in favour of 1 that they choose $x_n = 1$ and do not observe the social signal.

- They get a signal in favour of 1 that can be over-turned by the social signal, but the social signal is correct too.

- They get a signal against one, but weak enough that they look at the signal, and it is for the correct state.

For $n > 1$, assuming that $0.5 < \alpha_{n-1} - c_1$,[13]

$$
\begin{aligned}
\alpha_n &= \mathbb{P}(\varsigma \geq \alpha_{n-1} - c_1 | \theta = 1) + \int_{0.5}^{\alpha_{n-1}-c_1} \mathbb{P}(x_{n-1} = 1) g_1(s) ds \\
&+ \int_{1-\alpha_{n-1}+c_1}^{0.5} \mathbb{P}(x_{n-1} = 1) g_1(s) ds + 0 \\
\alpha_n &= \mathbb{P}(\varsigma \geq \alpha_{n-1} - c_1 | \theta = 1) + \int_{1-\alpha_{n-1}+c_1}^{\alpha_{n-1}-c_1} \mathbb{P}(x_{n-1} = 1) g_1(s) ds
\end{aligned}
$$

$$
\begin{aligned}
\alpha_n &= \mathbb{P}(\varsigma \geq \alpha_{n-1} - c_1 | \theta = 1) + \int_{1-\alpha_{n-1}+c_1}^{\alpha_{n-1}-c_1} \mathbb{P}(x_{n-1} = 1) g_1(s) ds \\
&= 1 - \mathbb{G}_1(\alpha_{n-1} - c_1) + \alpha_{n-1} \Big\{ \mathbb{G}_1(\alpha_{n-1} - c_1) - \mathbb{G}_1(1 - \alpha_{n-1} + c_1) \Big\}
\end{aligned}
$$

Given the private signal structure I have assumed here, $\mathbb{G}_1(s) = s^2$. Hence this boils down to:

$$
\begin{aligned}
\alpha_n &= 1 - (\alpha_{n-1} - c_1)^2 + \alpha_{n-1} \Big\{ (\alpha_{n-1} - c_1)^2 - (1 - \alpha_{n-1} + c_1)^2 \Big\} \\
&= 1 - (1 - \alpha_{n-1})(\alpha_{n-1} - c_1)^2 - \alpha_{n-1}(1 - \alpha_{n-1} + c_1)^2 \\
&= 1 - (1 - \alpha_{n-1})(\alpha_{n-1} - c_1)^2 - \alpha_{n-1}(1 - (\alpha_{n-1} - c_1))^2 \\
&= 1 - (1 - \alpha_{n-1})(\alpha_{n-1} - c_1)^2 - \alpha_{n-1}(1 - 2(\alpha_{n-1} - c_1) + (\alpha_{n-1} - c_1)^2) \\
&= 1 - (1 - \alpha_{n-1})(\alpha_{n-1} - c_1)^2 - \alpha_{n-1} + 2\alpha_{n-1}(\alpha_{n-1} - c_1) - \alpha_{n-1}(\alpha_{n-1} - c_1)^2 \\
&= 1 - (\alpha_{n-1} - c_1)^2 - \alpha_{n-1} + 2\alpha_{n-1}(\alpha_{n-1} - c_1) \\
&= 1 - \alpha_{n-1}^2 + 2\alpha_{n-1}c_1 - c_1^2 - \alpha_{n-1} + 2\alpha_{n-1}^2 - 2\alpha_{n-1}c_1 \\
&= 1 + \alpha_{n-1}^2 - \alpha_{n-1} - c_1^2
\end{aligned}
$$

Note that c_1 must be less than 0.25 for any learning to take place at all here. To find the asymptotic value of α_n, let us set $\alpha_n = \alpha_{n-1} = \alpha$ here:

$$
\begin{aligned}
0 &= \alpha^2 - 2\alpha + (1 - c_1^2) \\
\alpha &= 1 - c_1
\end{aligned}
$$

Hence in this example, accuracy will converge in the long run to exactly $1 - c_1$, giving us a neat example in which reducing the cost c_1 and making the network strictly more connected will cause a smooth increase in asymptotic accuracy (below $c_1 = 0.25$). This establishes the result. □

(2) Making a network ***more*** *connected can* ***reduce*** *asymptotic accuracy*:

For this part, let's consider the network topology in the previous example. Now, however, let

[13] Otherwise every agent has accuracy $0.5^2 - 0.5 + 1 =: \alpha_1$.

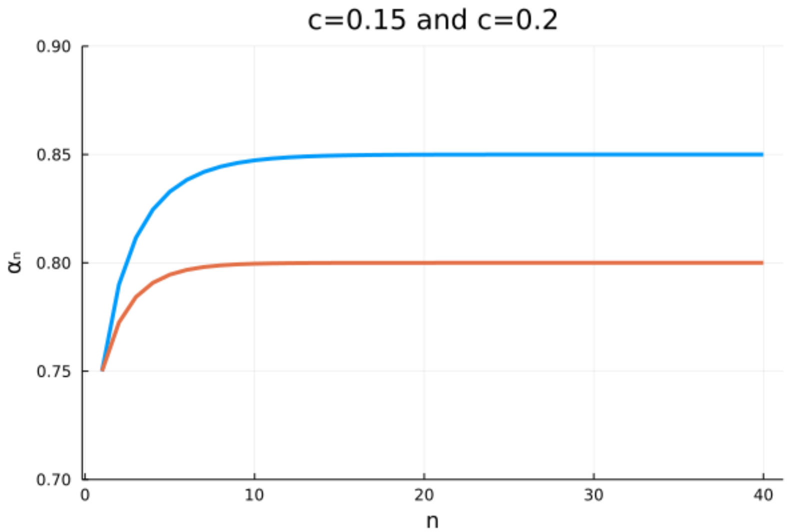

Figure 2a: The blue line here is with $c = 0.15$, and the orange $c = 0.2$.

us imagine introducing a free link between each agent after 2 and agent 1. This clearly makes the network strictly more connected.

Specifically, in the second network topology, let us assume that:

- For all $n > 2$, $c_{n,1} = 0$ and L=1.
- For all $n > 1$, $c_{n,n-1} = c_1 = 0.15$.

We can see the following:

- Agent 1 has an accuracy of $\alpha_1 = 0.75$
- Agent 2 has an accuracy of $\alpha_2 = 0.79$ using $\alpha_n = 1 + \alpha_{n-1}^2 - \alpha_{n-1} - c_1^2$.

If player 3 observes a private signal that generates a belief stronger than 0.79 (higher than 0.79 or lower than 0.21), they will not gain anything observing either predecessor (though they may as well observe agent 1 as this is free). If their private belief were weaker than 0.79 but stronger than 0.75, they would gain something observing agent 2, but nothing observing agent 1: agent 2 costs 0.15 to observe though. If player 3 observes player 2 with private signal realization ς, they get the correct state with probability α_2 rather than ς. Since in this range $alpha_2 - \varsigma < 0.15$, they will not observe agent 2. If $\varsigma \in [0.25, 0.75]$, they gain by observing either agent 1 or agent 2, and in the event of observing them will have success probability 0.75 and 0.79 respectively, this difference does not justify the cost of observing 2, so they will observe 1. In other words, in computing the accuracy of agent 3 we can act as if only agent 1 is available, and they will certainly observe agent 1. Their accuracy is then $\alpha_3 = 1 + 0.75^2 - 0.75 = 0.8125$.

For agent 4 we can deploy a similar line of reasoning, the only difference being that the accuracy gap between their two options (3 and 4) is slightly larger ($0.8125 - 0.75 = 0.0625$ rather than $0.79 - 0.75 = 0.04$). This does not invalidate any step in the previous line of reasoning though, so agent 4 observes 1 instead of 3 as well. Agent 5 in turn observes 1 instead of 4, and so on and so forth. The negative informational externalities of these decisions ensure that we achieve a lower asymptotic accuracy even with a strictly more connected network. □

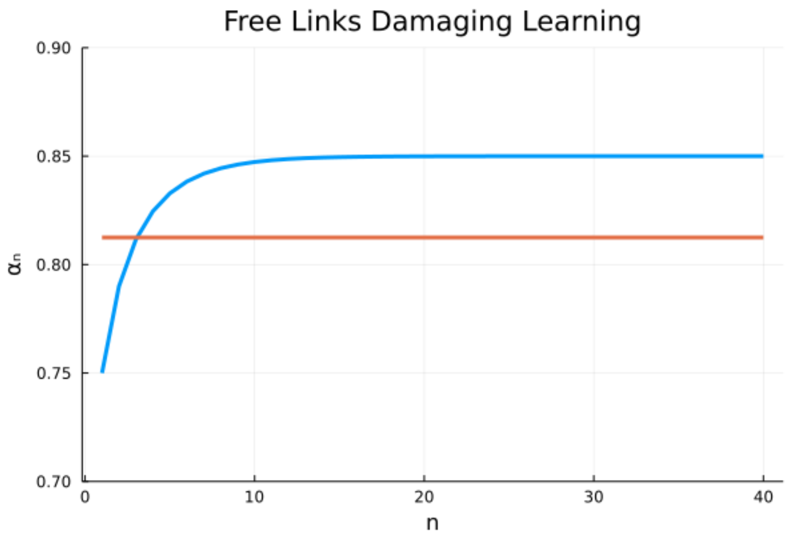

Figure 2b: The blue line here is with $c = 0.15$ before free links to agent 1 are added, and the orange after.

Proof of Proposition 6. (1) Making private signals ***more*** *informative can* ***increase*** *asymptotic accuracy*:

For this example, suppose that the social cost functions are the similar to those from the first example in the proof of Proposition B, except that observing one's immediate predecessor is now free. Our cost functions are thus such that $\forall n\ C_n^S(\{n-1\}) = 0$, but $C_n^S(n-j) > 1\ \forall j \in \{2, n-1\}$.

Suppose that in the first instance agents have access to a the following bounded private signal:

- $\mathcal{E} = \{X_1\}$
- $C_n^P(X_1) = 0$
- X_1 has probability density functions $\{f_0^*(s), f_1^*(s)\}$ with an atom on 0.5 with probability $1 - \mathbb{F}_1(0.8) + \mathbb{F}_1(0.2)$:

$$f_0^*(s) = \begin{cases} 2(1-s) & \text{if } s \in [0.2, 0.8] \setminus \{0.5\} \\ 1 - \mathbb{F}_1(0.8) + \mathbb{F}_1(0.2) & \text{if } s = 0.5 \end{cases}$$

and

$$f_1^*(s) = \begin{cases} 2s & \text{if } s \in [0.2, 0.8] \setminus \{0.5\} \\ 1 - \mathbb{F}_1(0.8) + \mathbb{F}_1(0.2) & \text{if } s = 0.5 \end{cases}$$

Note that this is simply the private signal distribution from before, except I have stripped its ability to produce beliefs above outside $[0.2, 0.8]$, and reassigned that probability to 0.5 (this probability has the same expression due to their symmetry).

Now suppose that we move from this world to one in which the agents instead observe the private signal distribution given by: $\{f_0(s), f_1(s)\} = \{2(1-s), 2s\}$. In the original case, agents in the line network will eventually ignore there private signal in favour of copying their predecessor (and choose the same action in perpetuity), and there is a non-zero probability that they choose the incorrect action forever.

The Bounded Signal Case:

First, let us consider the bounded signal case. If a predecessor has $\alpha \in [0.2, 0.8]^C$, the alpha of our agent is clearly the same α since they ignore their own private signal. If $\alpha \in [0.2, 0.8]$, it is possible that they will observe an overturning private signal. Due to the symmetry of this example, we can take the case that $\theta = 1$ without loss.

$$\begin{aligned}
\alpha_n &= \mathbb{P}(x_{n-1} = 1|\theta = 1) \times \mathbb{P}(s \geq 1 - \alpha_{n-1}|\theta = 1) + \mathbb{P}(x_{n-1} = 0|\theta = 1) \times \mathbb{P}(s \geq \alpha_{n-1}|\theta = 1) \\
&= \alpha_{n-1}\mathbb{P}(s \geq 1 - \alpha_{n-1}|\theta = 1) + (1 - \alpha_{n-1}) \times \mathbb{P}(s \geq \alpha_{n-1}|\theta = 1) \\
&= \alpha_{n-1}(1 - \mathbb{G}_1^*(1 - \alpha_{n-1})) + (1 - \alpha_{n-1})(1 - \mathbb{G}_1^*(\alpha_{n-1})) \\
&= \frac{132}{100}\alpha_{n-1} + (1 - \alpha_{n-1})^2 - \frac{36}{100}
\end{aligned}$$

$$\alpha_n = \alpha_{n-1}^2 - \frac{68}{100}\alpha_{n-1} + \frac{64}{100}$$

The Unbounded Signal Case: The reasoning here is exactly the same, except that we are using the distribution functions corresponding to the unbounded signal: $\{\mathbb{G}_0, \mathbb{G}_1\}$.

$$\begin{aligned}
\alpha_n &= \alpha_{n-1}(1 - \mathbb{G}_1(1 - \alpha_{n-1})) + (1 - \alpha_{n-1})(1 - \mathbb{G}_1(\alpha_{n-1})) \\
&= \alpha_{n-1}(1 - (1 - \alpha_{n-1})^2) + (1 - \alpha_{n-1})(1 - (\alpha_{n-1})^2) \\
&= \alpha_{n-1}(1 - (1 - \alpha_{n-1})^2) + (1 - \alpha_{n-1})^2(1 + \alpha_{n-1}) \\
&= \alpha_{n-1} - \alpha_{n-1}(1 - \alpha_{n-1})^2 + (1 - \alpha_{n-1})^2 + \alpha_{n-1}(1 - \alpha_{n-1})^2 \\
&= \alpha_{n-1} + (1 - \alpha_{n-1})^2 \\
&= \alpha_{n-1}^2 - \alpha_{n-1} + 1
\end{aligned}$$

In other words, we just have a normal line network of Bayesians here, and we achieve complete learning.

Comparison:

We can thus graph the path of α_n in these two settings, and see the result in Figure 3a. Setting $\alpha_n = \alpha_{n-1} = \alpha$ gives us that the limit of the bounded signal process is 0.584 (3d.p.), and we can see the graph converging to this. The unbounded process, as discussed, instead gives complete learning as is clearly displayed by the blue line.

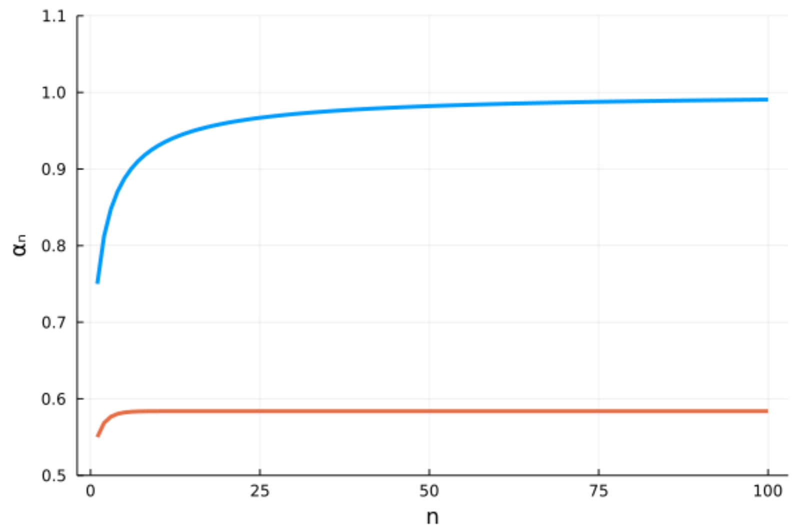

Figure 3a: The orange line shows the accuracy of agents receiving the less informative, bounded signal. We have complete learning in the unbounded case, and it is clear that reducing the informativeness of the signals hurts learning here.

*(2) Making private signals **more** informative can **decrease** asymptotic accuracy*:

Consider a network topology involving the following:

- For all $n > 2$, $c_{n,1} = 0$, and $L = 1$.
- For all $n > 2$, $c_{n,n-1} = 0.15$.
- For $n = 2$, $c_{n,1} > 1$.

In the two scenarios, let the private information structure be as follows:

- $\mathcal{E}_1 = \{X_1, X_2\}$ where X_1 is an unbounded private experiment with density functions $\{f_0(s), f_1(s)\} = \{2(1-s), 2s\}$, and X_2 is a symmetric Bernoulli trial with success parameter $p_1 = 0.5$. We have the costs functions $C_1^P(X_1) > 1$ and $C_1^P(X_2) = 0$ for agent one, and $C_n^P(X_1) = 0$ and $C_n^P(X_2) = 0$ for all $n > 1$.
- $\mathcal{E}_2 = \{X_1, X_2\}$ where X_1 is an unbounded private experiment with density functions $\{f_0^{\dagger,\epsilon}(s), f^{\dagger_1,\epsilon}(s)\}$, specified momentarily, and X_2 is a symmetric Bernoulli trial with success parameter $p_2 = 0.77$. The cost functions for these experiments are the same as in scenario 1.

$$f_\theta^{\dagger,\epsilon}(s) = \begin{cases} f_\theta(s) & \text{if } s \in [0.5-\epsilon, 0.5+\epsilon]^C \\ 0 & \text{if } s \in [0.5-\epsilon, 0.5+\epsilon] \\ (0.5+\epsilon) \times \Big\{ \mathbb{F}_\theta(0.5+\epsilon) - \mathbb{F}_\theta(0.5-\epsilon) \Big\} & s = \theta \\ (0.5-\epsilon) \times \Big\{ \mathbb{F}_\theta(0.5+\epsilon) - \mathbb{F}_\theta(0.5-\epsilon) \Big\} & s = 1-\theta \end{cases}$$

Note that it does not matter whether we use $\mathbb{F}_1$ or $\mathbb{F}_0$ here due to the symmetry of this signal distribution. I have constructed this signal structure so that $\{f_0^{\dagger,\epsilon}(s), f_1^{\dagger,\epsilon}(s)\}$ strictly Blackwell dominates $\{f_0^{\dagger,\epsilon'}(s), f_1^{\dagger,\epsilon'}(s)\}$ if $\epsilon > \epsilon'$. This implies that any such signal structure with $\epsilon > 0$ strictly Blackwell dominates $\{f_0(s), f_1(s)\}$, and with some abuse of terminology can be seen to do so by less and less as $\epsilon \to 0$. Hence, in the second scenario, all agents observe a private signal that strictly Blackwell dominates what they were able to observe in the first.

The idea here will be the same as that of part 2 of Proposition 5, except that we shall break the line-network improvement path by increasing the informativeness of agent 1's private signal, rather than by reducing the cost of observing him as before.

In the first scenario, this situation is the same as that of the blue line in Figure 2b, except that here agent $n+1$ corresponds to agent n, so the asymptotic accuracy of agents here is 0.85.

In the second scenario, the accuracy of agent 1 is $\alpha_1 = 0.77$, so agent 3 will of course observe 1 rather than 2 since it costs less and is more informative (if his signal is strong enough, both add no value but 1 is at least free). The improvement functions from earlier do not quite hold now, as the private signals of agents 2 onwards are different, but we can choose low enough ϵ that the functions are arbitrarily close to the true values. Agent 3's accuracy will be approximately (due to the $\epsilon > 0$) $0.77^2 - 0.77 + 1 = 0.8229$. This is not 0.23 greater than 0.77, so all subsequent agents shall observe agent 1 instead of their predecessor as well.

So with such a set-up, in scenario 1 they converge to 0.85 and we have complete learning. In scenario 2, every agent 3 onwards observes agent 1 and the asymptotic success rate is $0.8229 < 0.85$. Hence even though all we have done is strictly improved the Blackwell informativeness of everyone's signals, we have reduced asymptotic learning. This can be seen in the Figure below.

□

C A More General Model of Social Costs

Here I present the more general model of social costs of which my main model is a restriction. Every choice of parameters in the model of my main article can be re-expressed in the terminology of this one, and all results proved in this paper apply to this more general model as well.

In models of social learning in which the observational network is exogenous, there is no need to distinguish between different notions of the network topology. Here, however, this is not the case

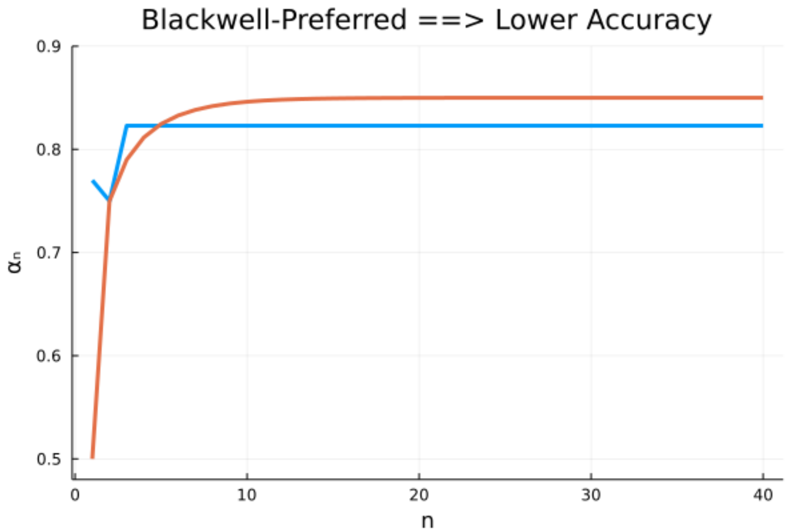

Figure 5: Figure 3b: Blackwell-Preferred Signals can Produce Lower Accuracy

(each agent n chooses his own neighborhood $B(n)$), and I discuss several useful network topology definitions in Section C.0.1. The exogenous object in this model to which I refer as the 'network topology' is simply the set of social cost function distributions for all agents, where nature draws for each agent a function that specifies how costly it is for that agent to observe any subset of their predecessors $C_n^S : 2^{\{1,\ldots,n-1\}} \to \mathbb{R}_+ \sim \Delta_n^S(\mathbb{R}_+^{2^{\{1,\ldots,n-1\}}})$. I assume that the cost of observing predecessors is at least weakly increasing in set inclusion, i.e. that if $A \subseteq B$, $C_n^S(A) \leq C_n^S(B)$, and that each agent's social cost function is independent of both the state of the world, θ, and the cost functions of all other agents.

Definition 5 ((Cost Function) Network Topology)**.** *The network topology of a given game is the sequence of all agents' social cost function distributions:* $\Delta^S = \{\Delta_n^S\}_{n\in\mathbb{N}}$.

I discuss several useful network topology definitions next in this appendix, in Section C.0.1. The exogenous object in this model to which I refer as the 'network topology' is simply the set of social cost function distributions for all agents, where nature draws for each agent a function that specifies how costly it is for that agent to observe any subset of their predecessors $C_n^S : 2^{\{1,\ldots,n-1\}} \to \mathbb{R}_+ \sim \Delta_n^S(\mathbb{R}_+^{2^{\{1,\ldots,n-1\}}})$. I assume that the cost of observing predecessors is at least weakly increasing in set inclusion, i.e. that if $A \subseteq B$, $C_n^S(A) \leq C_n^S(B)$, and that each agent's social cost function is independent of both the state of the world, θ, and the cost functions of all other agents.

Definition 6 (Cost Function Network Topology)**.** *The network topology of a given game is the sequence of all agents' social cost function distributions:* $\Delta^S = \{\Delta_n^S\}_{n\in\mathbb{N}}$.

C.0.1 Network Topologies

Given a specific cost function network topology, and a given equilibrium $\sigma \in \Sigma$, agents choose which predecessors they observe in order to maximise their expected utility. In the event that they choose to observe a private signal first, they may choose their neighbors as a function of their private signal realization. Courtesy of this, the endogenous probability with which an agent observes a given neighborhood in equilibrium depends on the state. Hence, for each agent this implies a pair of endogenous neighborhood distributions $\{\mathbb{Q}^n_{\sigma,\theta}\}_{\theta\in\Theta}$, and the sequence of these distributions make up the *Endogenous Observation Network Topology for Equilibrium* σ: $\{\{\mathbb{Q}^n_{\sigma,\theta}\}_{\theta\in\Theta}\}_{n\in\mathbb{N}}$.

In addition to this, it will be useful to consider the weighted network in which we include a link between each agent and all predecessors they can observe at finite cost, where the weight for each link is the cost of observing only that individual predecessor. For each agent, let us denote the distribution over weighted neighborhoods implied by this $\mathbb{W}_n$. The collection of all such distributions $\{\mathbb{W}_n\}_{n\in\mathbb{N}}$ shall make up the *Weighted Single-Neighbor Network Topology.* In addition to this it will be useful to consider the network containing only the cheapest link for each agent in each realization of the graph (or links if multiple neighbors have the same weight for some agent), $\{\mathbb{W}^*_n\}_{n\in\mathbb{N}}$. Finally, let us use a superscript c to denote the network topology produced by stripping out all links with cost strictly greater than c. I summarise all these terms in the following definition.

Definition 7. *Network Topolopies The following network topology concepts shall be useful:*

- *The* ***Network Topology*** *is as defined in Definition 6.*
- *The* ***Endogenous Observation Network Topology for Equilibrium*** σ*,* $\{\{\mathbb{Q}^n_{\sigma,\theta}\}_{\theta\in\Theta}\}_{n\in\mathbb{N}}$*, is the sequence of pairs (one for each state) of neighborhood distributions induced by the Network Topology and strategy profile in equilibrium* σ*.*
- *The* ***Weighted Single-Neighbor Network Topology****,* $\{\mathbb{W}_n\}_{n\in\mathbb{N}}$*, is the sequence of distributions over weighted graphs (one distribution* $\mathbb{W}_n$ *for each agent, with one realization* $W_n \in \mathbb{W}_n$*) in which the link between* n *and* $n-j$ *implied by a given realization of the cost functions has weight* $C^S_n(\{n-j\})$ $\forall j \in \{1, ..., n-1\}$*.*
- *The* ***Cheapest Neighbor Network Topology****,* $\{\mathbb{W}^*_n\}_{n\in\mathbb{N}}$*, is the same as the weighted single-neighbor network topology, except that for every agent* n *and neighborhood realization for that agent* n *we delete all links except those with the minimum weight* $\arg\min_{j\in\{1,...,n-1\}}\{C^S_n(\{n-j\})\}$*.*

Beyond defining different notions of network topology, it will also be useful to define in what sense I will speak of one network topology being more or less *connected* than another.

Definition 8 (Connectedness Order over Network Topologies $\succeq_C$)**.** *One network topology* $\Delta^S_{n,1}$ *is more connected than a second* $\Delta^S_{n,2}$*,* $\Delta^S_{n,1} \succeq_C \Delta^S_{n,2}$*, if for any* $n \in \mathbb{N}$*, and any set of predecessors* $A \subseteq 2^{\{1,...,n-1\}}$*, the cost of observing* A *in the first network topology is at least weakly less than the cost of doing so in the second with probability* 1*.* $\Delta^S_{n,1}$ *is strictly more connected than* $\Delta^S_{n,2}$ *if* $\Delta^S_{n,1} \succeq_C \Delta^S_{n,2}$ *but not* $\Delta^S_{n,2} \succeq_C \Delta^S_{n,1}$*.*

Note that for one network topology to be strictly more connected than another in this sense is a very strong requirement in a few ways, since we need every possible neighborhood for every single agent to be at least as cheap almost surely, with at least some non-zero probability than some neighborhood of some agent is strictly cheaper. On the other hand, we only need one neighborhood of one agent to be strictly cheaper with some (possibly arbitrarily small) probability $\epsilon > 0$.