UNIVERSITAT DE BARCELONA
SCHOOL OF ECONOMICS

Working Papers

Col·lecció d'Economia E22/434

PERSUADING CROWDS

Caio Lorecchio

ISSN 1136-8365

UNIVERSITAT DE BARCELONA
SCHOOL OF ECONOMICS

Working Papers

UB Economics *Working Paper No. 434*

Title: Persuading crowds

Abstract: A sequence of short-lived agents must choose which action to take under a fixed, but unknown, state of the world. Prior to the realization of the state, the long-lived principal designs and commits to a dynamic information policy to persuade agents toward his most preferred action. The principal's persuasion power is potentially limited by the existence of conditionally independent and identically distributed private signals for the agents as well as their ability to observe the history of past actions. I characterize the problem for the principal in terms of a dynamic belief manipulation mechanism and analyze its implications for social learning. For a class of private information structure - the log-concave class, I derive conditions under which the principal should encourage some social learning and when he should induce herd behavior from the start (single disclosure). I also show that social learning is less valuable to a more patient principal: as his discount factor converges to one, the value of any optimal policy converges to the value of the single disclosure policy.

JEL Codes: D82, D83

Keywords: Observational learning, Bayesian persuasion, dynamic information design

Authors:

Caio Lorecchio

Universitat de Barcelona and BEAT

Email: paeslemelorecchio@ub.edu

Date: October 2022

Acknowledgements: I am thankful for conversations with Daniel Monte, Thomas Wiseman, Maxwell Stinchcombe, V. Bashkar and Vladimir Asriyan. I also thank participants at the 2021 North-American Winter Meeting of the Econometric Society, the 42nd Meeting of the Brazilian Econometric Society, the 6th World Congress of the Game Theory Society, the University of Barcelona Microeconomics Workshop, the 2022 European Economic Association Congress - Econometric Society European Meetings and seminar participants at Sao Paulo School of Economics and University of Texas at Austin. All remaining errors are my own.

Persuading Crowds

Caio Lorecchio*

October 12, 2022

Abstract

A sequence of short-lived agents must choose which action to take under a fixed, but unknown, state of the world. Prior to the realization of the state, the long-lived principal designs and commits to a dynamic information policy to persuade agents toward his most preferred action. The principal's persuasion power is potentially limited by the existence of conditionally independent and identically distributed private signals for the agents as well as their ability to observe the history of past actions. I characterize the problem for the principal in terms of a dynamic belief manipulation mechanism and analyze its implications for social learning. For a class of private information structure - the log-concave class, I derive conditions under which the principal should encourage some social learning and when he should induce herd behavior from the start (single disclosure). I also show that social learning is less valuable to a more patient principal: as his discount factor converges to one, the value of any optimal policy converges to the value of the single disclosure policy.

JEL classification: D82, D83.

Keywords: Observational learning, Bayesian persuasion, Dynamic information design.

*E-mail: paeslemelorecchio@ub.edu. University of Barcelona School of Economics and Barcelona Economic Analysis Team. I am thankful for conversations with Daniel Monte, Thomas Wiseman, Maxwell Stinchcombe, V. Bashkar and Vladimir Asriyan. I also thank participants at the 2021 North-American Winter Meeting of the Econometric Society, the 42nd Meeting of the Brazilian Econometric Society, the 6th World Congress of the Game Theory Society, the University of Barcelona Microeconomics Workshop, the 2022 European Economic Association Congress - Econometric Society European Meetings and seminar participants at Sao Paulo School of Economics and University of Texas at Austin. All remaining errors are my own.

1 Introduction

When, however, it is proposed to imbue in the mind of a crowd with ideas and beliefs [...] leaders have recourse to different expedients. The principal of them are three in number and clearly defined - affirmation, repetition and contagion. Their action is somewhat slow, but its effects, once produced, are very lasting.

— Gustave le Bon, The Crowd: A Study of the Popular Mind

People follow the wisdom of crowds. Consumers are more inclined to buy popular brands because they believe that popularity is an indicator of better quality. A New York Times best-selling book is more likely to remain on the list - and obtain good reviews from readers. A small number of people deciding to withdraw their money might be sufficient to trigger a huge bank run. Because some people do believe that a crowd knows best, manipulating herd behavior is the goal of some other people - marketers, digital influencers, and financial advisors, to name a few.

This study examines crowd manipulation through dynamic information disclosure. In other words, how "to imbue in the mind of a crowd with ideas and beliefs". The setting and the results shed light on interesting questions, such as: To what extent should an information designer care about the future crowd effects of his current public releases? Should he publicly leak critical information to induce (or avoid) herd behavior from the outset or should he withhold decisive releases for a later time?

I consider a standard model of observational learning with a binary action space, and I add an information designer with selfish interests. Specifically, an infinite sequence of myopic agents wish to match actions with an unknown state of the world. They rely both on public observation of past actions and current private information coming from independently and identically distributed signals to guide them. As long as it is believed that past agents had chosen according to their private information, the action history helps current agents to infer the state before deciding which action to take.

They also rely on the public observation of designer's past and current messages[1]. This designer is informed about the state, but not about the private information of the agents. I assume that he is patient and only cares about the discounted number of agents taking his preferred action. He chooses a public information policy consisting of a message space and an information rule - a map from states and public histories to a distribution over messages.

Because the designer is more informed than agents, his messages might influence agents' beliefs and, consequently, agents' actions. However, this influence is sometimes limited: some agent can obtain more informative private data than the one given by designer's communication; sufficiently informative to drive her choice in the opposite direction of the designer's intention. Since past messages are publicly observed, future agents will know that someone has got a good reason to not follow the designer's advice, making persuasion harder than before. Thus, the designer must choose between allowing agents to follow their own private information (thereby allowing future agents to learn from past actions) or shutting down the observational learning process through by sufficiently revealing public disclosure.

The features of agents' private information structure determine when it is optimal for the principal to persuade society into a herd from the start - the single disclosure case - and when it is optimal to encourage some social learning dynamics, that is, letting agents choose according to their own private information and public observation of past decisions. For a well-known class of private belief distributions generated by private signals - the log-concave class, I characterize when social learning is valuable to a selfish principal.

[1] The fact that agents observe the realization of past messages is not crucial to my results, as long as the principal can commit to a sequence of experiments and agents know such experiments. I show that public communication does not loose generality in section 5.

Specifically, when agents cannot perfectly learn from private information (that is, private beliefs are bounded), I show that single disclosure is optimal if and only if private information unfavorable to the principal's most preferred action is sufficiently frequent (Theorem 1). With unbounded private beliefs, I show that this possibility can never be too significant, so single disclosure is never optimal.

I also prove that social learning is less appealing to a more patient principal, regardless of the structure of private information agents might have. In the limiting case, that is, as the designer's discount factor goes to one, the optimal policy has the same value as a policy that discloses in the first period sufficiently revealing information to induce herd behavior (Theorem 2). This means that whenever the designer does not heavily discount current payoffs from persuasion, avoiding agents from learning through observation of past actions might be his best interest.

This paper brings together two research topics from the dynamic games literature. The first one deals with dynamic information disclosure. I consider a persuasion problem similar to the ones in Ely (2017) and Renault, Solan, and Vieille (2017), but I have a fixed state of the world and I allow agents to obtain private information. This generates an evolving public belief process, even if the state does not change over time. As in those papers, I show in section 4 that it is possible to reformulate the designer's problem as a Markov decision problem in which (i) the state space is the space of agents' public beliefs; (ii) transition functions are governing the public belief process; (iii) the action space is the set of information rules; (iv) the constraint set over the action space is the set of distributions of posteriors that are mean-preserving spreads of any given prior. By reformulating the problem, I show that the dynamic concavification algorithm is used to solve it.

Unlike those studies, there are multiple laws of motion governing the belief process - one for each agent's action. Together with the private information assumption, this happens because the designer in my model cannot censor information; that is, he cannot avoid current agents from observing past actions. Moreover, the designer's messages influence the probability of having each law of motion governing the transition from the current to the next period's public belief, because it influences the probability of taking each action. In this sense, my model deals with a stochastic dynamic concavification algorithm.

Because agents are privately informed, my model also joins the literature on private persuasion as well. Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) and Inostroza and Pavan (2017) are seminal references, although both deal with a static persuasion problem. In the first reference, both agents and the designer have utilities that are linear functions of the private information, because the designer has a payoff that is a weighted combination of his preferred action and the utility of agents. Additionally, the state of the world is the realization of a continuous distribution. I consider a simpler environment: one with a binary state space and the designer's payoff depending only on actions. However, as in that paper, I also seek to characterize the designer's optimal policy in terms of the distribution of the private information, and my characterization also comes from insights from Quah and Strulovici (2012) about the aggregation of single-crossing functions.

Even though Inostroza and Pavan (2017) studied persuasion applied to global games of regime change, their results are related to mine, mostly the finding about the optimality of the information policy coordinating market participants in the same course of action. I show that when the interaction is dynamic, this single disclosure policy is sometimes optimal, but not always. However, as the designer becomes infinitely patient, the once-and-for-all coordination policy becomes more appealing.

Au (2015) studies dynamic information disclosure with a privately informed receiver. His environment is different from mine because the receiver has her private information being realized once and for all. Thus, she cannot learn from observations of past actions. Moreover, she is patient and takes an irreversible action that might depend on the designer's communication strategy. Therefore, the agent's problem is an optimal stopping one, in the sense that she must choose when to end the sender-receiver's dynamic interaction. Nevertheless, such paper provides conditions under which the designer discloses no further information beyond the first period, that is, the designer chooses the single disclosure policy. Among other things, it proves that if full disclosure is not optimal in the one-shot interaction, the optimal mechanism sequentially discloses informative messages.

In my model, as long as agents do not have private access to perfectly informative signals, full disclosure is never optimal. Furthermore, even if it is optimal to disclose information in the static environment in a way that beliefs are outside cascade sets (public belief sets under which agents choose no matter their private signals), for a very patient designer, a single disclosure policy (one placing beliefs inside cascade sets) is always optimal.

Observational learning is the second research topic from the dynamic games literature that my paper mainly deals with. In section 2, I present an illustrative example using the simple symmetric binary private signal case from Bikhchandani, Hirshleifer, and Welch (1998). In section 3, I briefly present the standard observational learning model and discuss major findings from this literature that I seek to study under the additional assumption of an information designer. The references for the model and the findings come from Banerjee (1992), Bikhchandani et al. (1998), Smith and Sørensen (2000), Cao, Han, and Hirshleifer (2011) and Herrera and Hörner (2012). Rosenberg and Vieille (2019) also provides an excellent summary of results in the literature.

An important takeaway from those studies is the following. Agents eventually settle down on a correct herd with probability one, and they fully learn the true state if and only if private signals are boundedly informative. In my model, conditional on the state that favors the designer's preferred action, with bounded private signals, a correct herd starts with probability one, but the belief process does not converge to one of the extremes. Conditional on the other state, there is always a possibility of a wrong herd and with some probability the belief process converges to one of the extremes. Therefore, even with bounded private information, learning is partially correct (correct with certainty at least conditioned in one state) and partially complete (agents fully learn with positive probability, at least conditioned in one state). This implies that society benefits from obtaining information from the principal relative to the social learning model without an information designer.

Smith, Sørensen, and Tian (2021) discusses optimal persuasion mechanisms in the same observational learning environment, but with a benevolent social planner. Specifically, an information designer has the power to choose a map from private signals to action recommendations, to maximize the discounted sum of receivers' payoffs. That paper shows that (i) cascade sets strictly shrink in the discount factor and collapse to extreme points in the perfect patience limit; (ii) for any discount factor, the social planner always encourages agents to rely more on private signals, so that past actions are always more informative.

With a non-benevolent designer, without the possibility of eliciting agents' private information, under a binary action space, I prove that the cascade set toward his least preferred action is always a singleton, while the cascade set of his most preferred action does not change. This partial reduction in cascade sets does not depend on the discount factor. I also prove that, even though the designer does not care about letting information flow through agents, sometimes - but not always - it is optimal for him to encourage agents to rely on private signals.

This is not the first study to investigate observational learning with a non-benevolent planner. Sgroi (2002) considers an uninformed planner with the power to censor information in the market. His problem is then choosing the number of agents to decide using only their private signal after letting others have access to a history of past actions. I adopt a different approach. My planner is informed about the state, but cam commit ex-ante to an experiment. He cannot censor the observation of past actions. His choice is then to fine-tune his disclosure policy to be clear or vague in his communication. Nikiforov (2015) considers an informed manipulator with the power to costly influence only one agent along the sequence, in a symmetric binary private signal environment. I allow the manipulator to persuade as many agents as he wants, taking into consideration a general private information structure.

The remainder of this paper is organized as follows. Section 2 introduces a illustrative example to walk through the main results. Section 3 presents the benchmark model without an information designer. Section 4 describes the social learning problem as an information design problem. It then discusses belief dynamics; when social learning is valuable to the principal; and the role of patience in designing the policy. Section 5 considers private disclosure of information and section 6 concludes the paper. All proofs of lemmas and claims are in appendix A and all calculations for examples are in appendix B.

2 Illustrative example

Let me introduce a illustrative example as a way of walking through belief dynamics and the main results of this paper. Imagine that a financial advisor wishes to persuade his clients to buy a certain asset of an unknown return. These clients only care about their current gains from investing and they arrive sequentially at the advisor's office. If the asset yields a high return, clients obtain a payoff of 1 from investing and incur an opportunity cost of -1 from not doing so. If the asset yields a low return, payoffs are reversed. Every current client is partially informed: she observes a private signal about the asset's quality and the history of decisions. If the asset yields a high (low) return, she observes the signal $\bar{s}$ ($\underline{s}$) with probability $\sigma \in (.5, 1)$. The prior belief about the asset yielding a high return is .5.

Although clients do not observe the history of private signals, they can infer it from the history of actions. To understand this, consider the decision of the first client. Starting with a flat prior about the asset's quality, she will update her Bayesian belief to σ if she observes signal $\bar{s}$ and $1-\sigma$ if she observes signal $\underline{s}$. Given her payoffs, she will invest if and only if the posterior is at least half [2] [3]. This means that she will invest if and only if she receives signal $\bar{s}$.

The second client, after observing the first client's decision, will know what private signal she received. Thus, the second client's *public belief* (inference from past action) will be either σ from observing investment or $(1-\sigma)$ otherwise. Suppose it is σ. If this second client observes $\bar{s}$, her total belief (inference from past action and current private signal) will be $\frac{\sigma^2}{(1-\sigma)^2+\sigma^2}$, which exceeds her belief threshold .5; if she observes $\underline{s}$, her total belief goes back to .5, which also implies that she will invest. It follows that she will invest, regardless of her private signal. The third client will not be able to determine what private signal the second client received and will have an interim belief of σ, exactly like the second client. In other words, if the first client invests, all future clients will, independent of the realization of private signals.

Suppose that the second client has public belief $1-\sigma$. On the one hand, if she observes signal $\underline{s}$ as well, her posterior belief will be $\frac{(1-\sigma)^2}{(1-\sigma)^2+\sigma^2}$, which is below her belief threshold .5. Thus, she will choose not to invest. The third client will observe two consecutive decisions of no investment and will choose not to invest as well, regardless of her private signal. On the other hand, if the second client observes $\bar{s}$, her posterior belief will return to .5, implying that she will invest. The third client will be able to infer that the second client received a good signal, which offsets the bad signal $\underline{s}$ from the first client, and the analysis continues as if this third client does not have any additional information (i.e., as if she was the first client).

In the dynamics I have described so far, I have not said anything yet about the advisor's role, so think of it for a moment as a non-intervention benchmark. Assume that he receives 1 every time a client invests, and 0 otherwise. The first client will invest if and only if she receives the private signal $\bar{s}$, which occurs with probability σ if the asset yields a high return and $1-\sigma$ otherwise. Since the advisor is also ignorant about the asset's quality, the expected payoff from the first client coincides with the expected investment probability, which is .5.

If the first client invests, the second and every subsequent clients will invest for sure, so the advisor will receive 1 forever. If the first client does not invest, the second client will if and only if she receives signal $\bar{s}$, which continues to occur at a probability of σ if the asset has a high return, and a probability of $1-\sigma$ if the asset has a low return. However, because the public belief for the second client is $1-\sigma$, the expected investment probability for the second client (and advisor's expected payoff) is $2\sigma(1-\sigma)$, which is lower than .5. If the first client does not invest, the second client will dictate whether the advisor is dammed to a zero payoff forever. If the second client invests, the third client will have an expected investment probability equal to the first client, but if the second client does not invest, the third and every subsequent client will not invest, as no private signal can generate a belief in favor of investment.

[2] There is an underlying assumption that if the posterior is exactly .5, she will choose to invest. Breaking indifference towards the advisor's most preferred action in the Bayesian persuasion literature is common.

[3] All computations for the illustrative example are in appendix B.

To simplify the exposition, I present a visualization of the public belief dynamics and advisor's expected payoffs over time, as shown in figure 1(a). I also present the probability of clients choosing each action ignoring their private signals for each period in figure 1(b) - a phenomenon called *informational cascade*, assuming $\sigma = .8$. In figure $1(a)$, the blue line represents the change in public beliefs when investment is observed, and the red line represents the change when a non-investment decision is observed. The black dots represent the possible realizations of public beliefs, and the numbers above these points represent the associated expected investment probabilities at every belief. In figure $1(b)$, the x marks are the probabilities of clients taking investment decisions regardless of their private signals in each period, and the triangle marks are the probabilities of clients taking non-investment decisions with certainty. The blue line is the long-run probability of public beliefs hitting σ, the threshold over which an information cascade towards investment occurs. Similarly, the red line is the limiting probability of public beliefs hitting a value below $1 - \sigma$, the threshold below which an information cascade towards non-investment occurs. Note that as time goes by the probability of an information cascade towards investment or non-investment equals 1.

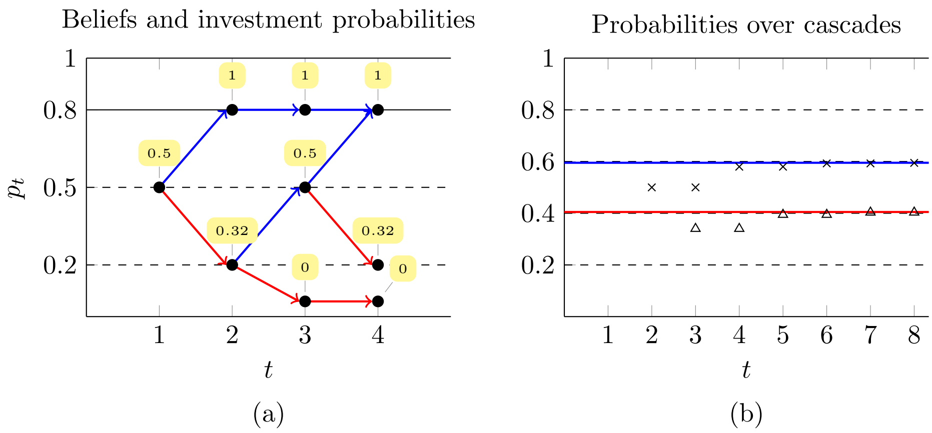

Figure 1: Dynamics without intervention, for $\sigma = .8$. Blue lines represent outcomes related to investment decisions, and red lines represent outcomes related to non-investment decisions. Figure 1(a) shows the possible realizations of public beliefs over time with numbers representing the designer's expected payoff at every belief. Figure 1(b) shows the probability of a cascade starting in each action (x marks correspond to an investment cascade; triangle marks to a non-investment cascade). Colored lines represent the long-run probability of each cascade.

The advisor can use his expertise to investigate whether the asset will yield a high or low return, but is legally obliged to report the outcome of this investigation. Specifically, this advisor will design *ex-ante* a contract specifying a set of messages and a probability distribution over messages conditional on what he knows at every period, that is, past messages and actions, as well as the true quality of the asset. I will refer to this contract as a public information policy and assume that every client knows the chosen policy. At the beginning of every period, the advisor sends some advice and a new inference about the asset's return is made.

Can the advisor perform better than the no-intervention benchmark? Unsurprisingly, he can. Consider the following policy. The messages are either Aaa (an investment with the lowest risk) or Caa (a junk with highest risk). The first client will observe Aaa for sure if the return is high and with probability $\frac{1-\sigma}{\sigma}$ if the return is low. In this way, after observing Aaa (Caa), the first client will have an *induced belief* of σ (0) and will invest (not invest) no matter private signals. The second client will have public beliefs of either σ (as she saw that the first client invested and the message was Aaa) or 0 (as she saw that the second client did not invest and the message was Caa). Therefore, after the first client, there is no room for intervention: if the first client invested, it suffices to send uninformative messages forever; if the first client did not invest, no message could refrain the second client from choosing not to invest. For this reason, I will refer to this policy as the *single disclosure* policy. With it, advisor uses his informational power to persuade society into cascades from the outset and no client learns from past actions.

The expected investment probability of such a rule at prior .5 equals the unconditional probability of message Aaa, which is $\frac{1}{2\sigma}$. This is higher than the value obtained without intervention. This is also the limiting probability of having an investment cascade, which is higher than that without any intervention. Figure $2(a)$ and $2(b)$ represent the public belief dynamics and the probabilities over cascades under this information policy. The solid black dots in figure $2(a)$ represent the possible public beliefs (p_t). The white dots represent are the induced beliefs (ρ_t).

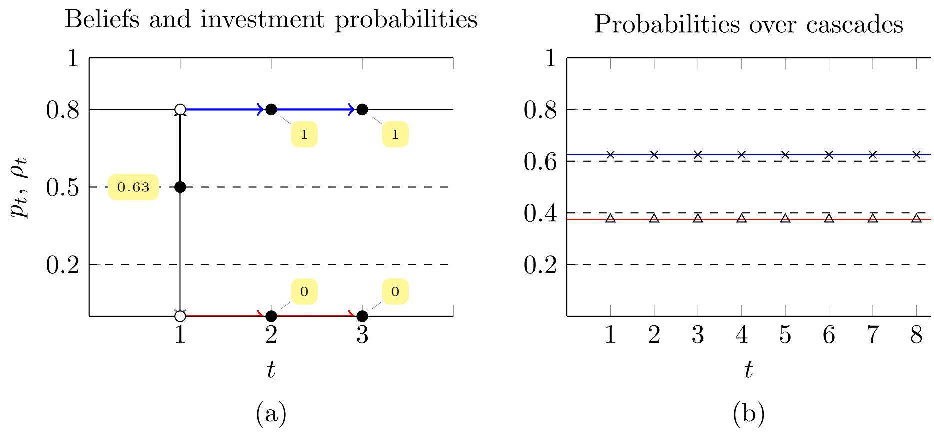

Figure 2: Dynamics with single disclosure. Blue lines represent outcomes related to investment, and red lines represent outcomes related to non-investment. Figure 2(a) shows the possible realizations of public (black dots) and induced (white dots) beliefs over time, with numbers representing the designer's expected payoff at every public belief. Figure 2(b) shows the probability of a cascade starting in each action (x marks correspond to an investment cascade; triangle marks to a non-investment cascade). Colored lines represent the long-run probability of each cascade.

Another policy is worth discussing. The message space now contains an intermediate message Baa. This message represents an investment with medium risk. Consider the set consisting of the null history at $t = 1$ and all history of actions that are repetitions of the pair "not invest/invest". After observing every history in this set, the public posterior is exactly the prior. At every such history, the advisor sends Aaa with probability σ if the asset has a high return and with probability $1 - \sigma$ if the asset has a low return. Therefore, if the current public history of actions leads to a public belief of .5, the advisor sends both Aaa and Baa with the same unconditional probability, although message Aaa is more likely if the return is high and Baa is more likely if the return is low. Note that the message Aaa induces belief σ and message Baa induces belief $1 - \sigma$.

After observing Aaa, the first client will invest no matter private signals. After observing Baa, the first client will invest if and only if she receives a private signal $\bar{s}$. If such signal occurs, the second client will start the period with public belief exactly like the prior, and the algorithm discussed in the previous paragraph applies. However, if the first client receives a private signal $\underline{s}$, the second client will hold unfavorable beliefs to investment, unless the advisor does something. In that case, the advisor communicates Baa for sure if the asset yields a high return and with probability $\frac{1-\sigma}{\sigma}$ otherwise. The alternative to Baa is Caa. This ensures that if the public belief is $\frac{(1-\sigma)^2}{(1-\sigma)^2+\sigma^2}$, clients will have induced beliefs of $1 - \sigma$ under message Baa and 0 under message Caa. Medium-grade Baa works as an advice for clients to follow their private signals; Aaa and Baa work as recommendations to choose irrespective of private information. With this alternative rule, the advisor allows some clients to learn from their predecessors. The probability of investment at the prior is $(0.5)[1 + 2\sigma(1 - \sigma)]$, which again is higher than in the case without intervention. There is now a positive probability of investment, even when two non-investment decisions are observed. This probability is equal to the unconditional probability of sending message Baa under public belief $\frac{(1-\sigma)^2}{(1-\sigma)^2+\sigma^2}$ times the investment probability when public belief is $1 - \sigma$. This was not possible in the case without intervention.

Figure 3(a) and 3(b) represent the belief dynamics and the probabilities over cascades under this information rule, respectively. Note that the probability of having cascades equals one as time goes by, but the belief convergence is not immediate.

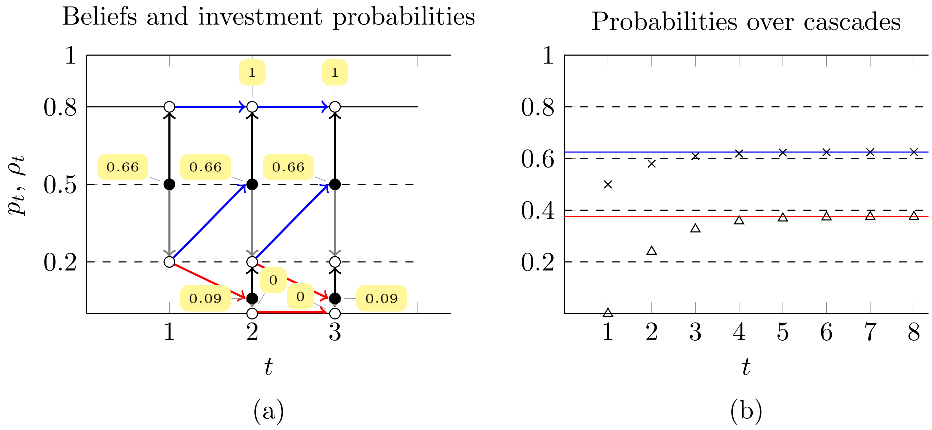

Figure 3: Dynamics with the alternative policy. Blue lines represent outcomes related to investment, and red lines represent outcomes related to non-investment. Figure 3(a) shows the possible realizations of public (black dots) and induced (white dots) beliefs over time, with numbers representing the designer's expected payoff at every public belief. Figure 3(b) shows the probability of a cascade starting in each action (x marks correspond to an investment cascade; triangle marks correspond to a non-investment cascade). Colored lines represent the long-run probability of each cascade.

Which policy is better for the advisor? Inspecting investment probabilities, one can see that if there is only one client, the second rule yields a higher value as long as private signals are sufficiently informative, that is, as long as $\sigma \geq \frac{1}{\sqrt{2}}$. This is the case for $\sigma = .8$: the first client's investment probability with the single disclosure rule is .63 versus .66 with the alternative rule. In a repeated interaction with a very patient advisor, both policies would look the same to him, because both lead to the same long-run probability of having a cascade toward investment. For non-extreme discount factors, the analysis is not straightforward. For instance, if the advisor is impatient, it might be that he prefers to increase the probability of investment in every period and sacrifice the speed of belief convergence towards cascades. But every time he discloses additional information, he also gives away part of his informational advantage to future clients, as messages are public. Additionally, private signals can make future clients less easily to be persuaded. Moreover, both policies considered here are stationary in public beliefs and do not depend on advisor's discount factor. Is there a more complex policy that improves upon these two?

Perhaps surprisingly, I will show that single disclosure will be optimal in this example if and only if $\sigma \leq \frac{1}{\sqrt{2}}$, regardless of the discount $\delta \in (0, 1)$. Although this threshold is specific to this example, I will show that for a broader class of private information structures, there is a simple test to verify optimality of single disclosure. This test depends on how informative private signals can be. As in the example, I will prove that if private signals are very revealing, then some social learning is always valuable to the advisor.

For $\sigma > \frac{1}{\sqrt{2}}$, the illustrative example is sufficiently manageable to characterize the related optimal policy. It is the alternative policy presented here, indeed. This is the policy that minimizes the amount of information given at every public belief, subject to the expected investment probability being maximal. Note that this leads to the same long-run cascade probabilities. As such, no matter the parameter σ, social learning is less appealing the more patient the advisor is. I will prove that this observation holds for every private information structure. Finally, note that both rules benefit society relative to the non-intervention case. This happens because the selfish advisor always discloses information to move clients away from the public belief set $(0, 1 - \sigma]$. Without him, belief dynamics would be forever trapped in there. This observation also generalizes.

3 A model of crowds

This section discusses how the wisdom of a crowd evolves without any intervention. Thus, it serves as a no-policy benchmark. I will present a standard model of observational learning and add a long-lived principal ("he") who derives instantaneous payoffs from actions taken by a sequence of identical short-lived agents $t \in \mathbb{N}$ (each one referred to as "she"). Then, I will emphasize some relevant results from the observational learning literature for the optimal information design that I seek to characterize in the next section.

At the beginning of the interaction between the principal and the agents, Nature draws a state: either H or L. No player observes this realization of Nature, but everyone shares a flat common prior belief that it is H: $p_1 = 1/2$. Every agent t must choose an action $a \in \{h, \ell\}$ to obtain either $u(a, H)$ or $u(a, L)$ as instantaneous payoffs. It is assumed that $u(h, H) = u(\ell, L) = 1$ and $u(\ell, H) = u(h, L) = 0$. This means that agents want to match actions with the unknown state. It also means that any agent with some belief $r \in [0, 1]$ about the state H will find action h optimal if and only if $r \geq 1/2$. Every time an agent chooses h, the principal receives 1 regardless of the state; otherwise, he receives nothing[4].

Whenever possible, agents compute beliefs using two sources of information. The first one comes from the observation of a private signal[5]. Conditional on the state, the signals are independently and identically distributed. Combining signals with the common prior, agents compute private beliefs $\{\tilde{q}_t\}_{t\in\mathbb{N}}$ about the state being L [6]. Because private signals are conditionally i.i.d., the private belief process will have the same feature. Let G represent the unconditional distribution function for private beliefs. I assume that G is absolutely continuous with density g. Note that $G \coloneqq (1/2)[G^H + G^L]$ where G^H and G^L denote the distribution functions over private beliefs conditional on the states. Thus, assuming absolute continuity of G is equivalent to assuming absolute continuity of G^H and G^L. It also ensures that no observation of private beliefs perfectly reveals the state and that both distributions share a common support.

The second source of information comes from the public observation of action histories. Since past private signals are non-observable, but past actions might be taken conditional on specific realizations of such signals, the action history might help inference about the state. A strategy for each agent t is a map from the set of private signals and the set of public action histories up to $t-1$ to a choice over $\{h, \ell\}$. A strategy profile for the agents is a collection of each map. A strategy profile, the private information structure, and the prior belief generate a probability distribution over the set of outcomes of the game.

Agents' rationality is common knowledge, so they can compute probabilities for every possible history of actions. Let $\tilde{p}_t$ represent the conditional probability of the state being H, given the observation of some action history up to $t-1$. Likewise, let $\{\tilde{p}_t\}_{t\in\mathbb{N}}$ be a stochastic process of the *public beliefs*. If agent t obtains a realization q_t of a private belief and a realization p_t of the public belief, she will have a Bayesian total belief r_t and choose action h if and only if

$$r_t = \frac{(1-q_t)p_t}{(1-q_t)p_t + q_t(1-p_t)} \geq 1/2 \quad \Leftrightarrow \quad p_t \geq q_t. \tag{1}$$

[4]I also assume that at belief $r = 1/2$, agents choose h, that is, principal's preferred choice, but this will be innocuous, because I will consider only continuous distributions over private beliefs, such that points of indifference will have zero measure. I will discuss private belief distributions later.

[5]Each signal s_t takes value on space S and its domain is the sample set of a probability space capturing all exogenous uncertainty in the interaction. Appendix A provides a more detailed description of this space.

[6]Throughout the text, I will identify a random variable by a tilde superscript and a value it can assume by its symbol. Additionally, the subscript t will represent a random variable with index t in a stochastic process, and the symbol indexed by t a realization of such variable. Thus, each $\tilde{q}_t$ is a random variable taking values in $[0, 1]$ and q_t is a realization of $\tilde{q}_t$.

Let $\underline{q}$ be the infimum value of $q \in [0,1]$ such that $G(q) > 0$ and $\bar{q}$ be the supremum value of $q \in [0,1]$ such that $G(q) < 1$. I will impose $\underline{q} < 1/2 < \bar{q}$ to avoid uninteresting situations in which public beliefs converge from the start. When $[\underline{q}, \bar{q}] = [0,1]$, I will say that private beliefs are bounded, and when $[\underline{q}, \bar{q}] \subset [0,1]$, I will say they are unbounded[7]. The agent's strategy is now a function of the private and the public beliefs. As q_t is not observed, principal conditionally and unconditionally expect that action h is taken at t according to the probabilities below.

$$\alpha^H(p_t) := G^H(p_t) \text{ and } \alpha^L(p_t) := G^L(p_t) \tag{2}$$
$$\alpha(p_t) := p_t G^H(p_t) + (1-p_t) G^L(p_t). \tag{3}$$

One can show[8] that $\alpha^L(p) \leq \alpha(p) \leq \alpha^H(p)$ with strict inequalities for every $p \in (\underline{q}, \bar{q})$. Moreover, $\alpha(p) = 0$ for $p \leq \underline{q}$, $\alpha(p) = 1$ for $p \geq \bar{q}$ and $\alpha(p)$ strictly increases in p for $p \in (\underline{q}, \bar{q})$. Intuitively, if Ms. t is very convinced that state is H by looking at past actions, she needs a very high private belief about the state being L to make her choose action ℓ.

Because agent $t+1$ is a rational Bayesian player, after observing some history (a, a^{t-1}), she can infer that t had public belief p_t and can compute the probability of her choosing action a under any state. This is exactly the probability that agent t had a private belief that led her to choose a under p_t in any state, that is, probabilities described by equations 2 and 3. Thus, agent 's $t+1$ inference from public history will lead to a public belief update:

$$p_{t+1} = \varphi_a(p_t) := \begin{cases} \left[\frac{\alpha^H(p_t)}{\alpha(p_t)}\right] p_t & \text{if } a_t = h, \\ \\ \left[\frac{1-\alpha^H(p_t)}{1-\alpha(p_t)}\right] p_t & \text{if } a_t = \ell. \end{cases} \tag{4}$$

One can show[9] for every $p \in (\underline{q}, \bar{q})$, $\varphi_\ell(p) < \min\{p, 1/2\}$ and $\varphi_h(p) > \max\{p, 1/2\}$. Therefore, for public beliefs in such set, (i) agents will choose actions according to private beliefs; therefore, past actions convey valuable information; (ii) observing action ℓ is always perceived as "bad news" about state being H (thus reducing the public belief) and observing action h is always "good news" (thus increasing it).

However, depending on how convinced an agent is about the state being H, her private inference might not change her choice of action at all. That is, she chooses according to her public information, regardless of the private information she receives. The next agent will infer that observing her action conveys no additional information about the state and will find optimal as well to choose the same action regardless of possible private beliefs. This process will go on infinitely, and there will be no more learning from the observation of past actions. To better describe this phenomenon, first consider $C_\ell := [0, \underline{q}]$ and $C_h := [\bar{q}, 1]$. Whenever $p_t \in C_a$, agent t chooses action a without considering private signals, so the belief dynamics stop in the next period and no further belief updating occurs. Second, note that the public belief process $\{\tilde{p}_t\}_{t \in \mathbb{N}}$ is a martingale. Indeed, consider any public history $a^t = (a, a^{t-1})$ that leads to a public belief p_t after the observation of a^{t-1}. From equation 4,

$$\mathbb{E}[\tilde{p}_{t+1} | a^t] = \alpha(p_t) \left[\frac{\alpha^H(p_t)}{\alpha(p_t)}\right] p_t + (1 - \alpha(p_t)) \left[\frac{1-\alpha^H(p_t)}{1-\alpha(p_t)}\right] p_t = p_t.$$

[7] I will focus on these two symmetric cases. Note that $[\underline{q}, \bar{q}]$ is the support of the distribution G. Indeed, because G is absolutely continuous, it is continuous. As such, its support is an interval.

[8] See claim 1 in Appendix A or Lemma 1 in Smith and Sørensen (1996).

[9] See claim 2 in Appendix A or Lemma 7 in Smith and Sørensen (1996).

Being a martingale, a well-known theorem ensures that it converges to a random variable $\tilde{p}_\infty$ almost surely. Moreover, it is possible to show that every realization of this random variable must belong to $C_\ell \cup C_h$. Intuitively, the stationary public belief process must reach an absorbing set, one for which no further update takes place; otherwise, public beliefs keep changing infinitely often, contradicting almost sure convergence[10].

What does this convergence imply for the non-interventionist principal? If he discounts future payoffs according to the discount factor $\delta \in (0, 1)$, his welfare is the expectation of the discounted number of agents taking action h. Let $\mathbb{P}_{np}$ be the probability measure over action histories without any intervention from the principal. The "np" abbreviation stands for "no policy." This non-interventionist principal obtains:

$$V_\delta^{np} = \sum_{t \in \mathbb{N}} (1 - \delta)\delta^{t-1} \mathbb{P}_{np}[a_t = h].$$

Let $\{\lambda_t^{np}\}_{t \in \mathbb{N}}$ be a sequence of probability measures over the belief space representing, at each t, the probability of the public belief process belonging to some subset of the Borel σ-algebra of $[0, 1]$. Note that at each t, the probability of agent t taking action h is the expected value of α with respect to λ_t. Because the public belief process converges almost surely to $\tilde{p}_\infty$, the sequence of probability measures must converge weakly to the limiting measure λ_∞^{np}. Because $\alpha(\cdot)$ is a continuous function, the sequence of the expected values of α with respect to each λ_t must converge to the expected value of α with respect to λ_∞^{np}. However, λ_∞ must place positive probability only on events that intersect the cascade sets and principal receives a positive payoff only on points that belong to C_h. Thus, the limiting expected value of α under λ_∞^{np} must be $\lambda_\infty^{np}(C_h)$.

One can show[11] that, as the principal becomes very patient, his long-run value of the no-policy interaction must approach the stationary probability of having the public belief process trapped in C_h:

$$\lim_{\delta \to 1} V_\delta^{np} = \lim_{t \to \infty} \mathbb{E}_{\lambda_t^{np}}[\tilde{\alpha}] = \lambda_\infty^{np}(C_h).$$

Because the state of the world is fixed throughout the dynamics, I can split the unconditional measure λ_∞^{np} into the conditional measures $\lambda_\infty^{H,np}$ and $\lambda_\infty^{L,np}$, such that $\lambda_\infty^{np} = (1/2)(\lambda_\infty^{H,np} + \lambda_\infty^{L,np})$. Say that learning is *correct* if $\lambda_\infty^{H,np}(C_h) = \lambda_\infty^{L,np}(C_\ell) = 1$; that is, agents eventually settle down on the correct actions. In addition, say that learning is *complete* if $\lambda_\infty^{H,np}(\{1\}) = \lambda_\infty^{L,np}(\{0\}) = 1$; that is, agents learn the true state. If learning is complete, it is correct, but the converse is not necessarily true.

When private beliefs are unbounded, learning is complete - thus, correct. In this case, the principal's only hope of getting some positive payoff infinitely often is the state of world being H; otherwise, he gets nothing as the dynamic interaction proceeds. Thus, for a very patient principal, the value of a no-policy interaction is $1/2$. When private beliefs are bounded, learning is both incomplete and incorrect. It is incomplete because there are no perfectly informative private beliefs that drive the belief process to either zero or one, by assumption. It is incorrect because there is always a positive probability of society settling down on incorrect actions, conditional on the true state. Thus, even if the state is L, the principal can obtain a positive payoff infinitely often.

Let me summarize all the primitives of the model and the relevant lessons from the observational learning literature. The principal takes as given the (common) private information structure of each agent. As discussed, it is sufficient to describe this information structure in terms of the unconditional distribution of private beliefs G and the associated support $[\underline{q}, \bar{q}]$. Agents would like to act according to the states of the world; principal only cares about one of the actions.

[10] See claim 3 in Appendix A or Theorem 1 in Smith and Sørensen (1996).
[11] See claim 4 in the Appendix A or Lemma 1 in Cao et al. (2011).

Each agent observes past actions and current signals; the total inference about the state comes from a combination of a private and a public beliefs. The public belief process follows a martingale and converges to a random variable whose support belongs to cascade sets. A very patient principal that does not intervene in the public belief process expects to earn a positive payoff with same probability of the limiting probability of the process reaching cascade set C_h.

4 Persuading crowds

This section assumes away the inability of the principal to intervene in the dynamic interaction and addresses the question of optimal public information provision from his point of view. Now he can commit to any information policy before the realization of the state of the world. In practical terms, think of the principal as being responsible for designing an experiment to learn the value of the state and is legally required to report its outcomes, even though the structure of the report and the timing of the releases are principal's decision.

Each agent still has access to private beliefs and past history of actions as sources of information, but a strategic communicator now provides a third source. The commitment assumption gives him more persuasion power because agents can infer something about the state of the world without worrying whether the principal deviates from his communication strategy. However, the public (past messages are not erased from public histories) and transparent (the principal cannot censor past records of actions as well, so he conditions messages on the same public history that agents observe) communication restrictions might increase or reduce this power.

An information policy π consists of a message space M and two information rules μ^H, μ^L. These rules specify conditional probabilities on the set of probability measures over the message space, given public histories, that is, $\mu^H : \cup_{t\in\mathbb{N}} X^t \to \Delta(M)$ and $X^t := (A \times M)^{t-1}$ for each t (the set of public histories at $t = 1$ is the null history). The rule μ^L is similarly defined. This policy is chosen before the realization of the state and agents know unambiguously know how to interpret what the principal is communicating in each period. Figure 4 below describes the timing of the events.

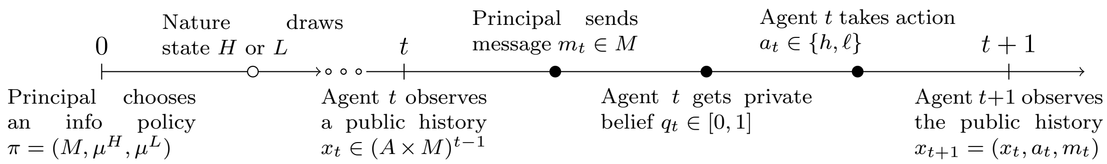

Figure 4: Timing of events.

Along with agents' strategies and the prior belief, the policy generates a probability measure $\mathbb{P}_\pi$ over the set of outcomes. Thus, principal's value from the information policy π is

$$V_\delta^\pi := \sum_{t\in\mathbb{N}} (1-\delta)\delta^{t-1}\mathbb{P}_\pi[a_t = h].$$

Upon observing some history x_t, Ms. t will have a public belief p_t about the state being H. However, because the principal's message at t might provide some information about the state, agent t will also have a Bayesian *induced belief* ρ_t. The expected value of induced beliefs $\tilde{\rho}_t$ conditional on the information obtained in t must equal the public belief process obtained from that information. Formally,

$$\mathbb{E}_\pi[\tilde{\rho}_t|x_t] = \mathbb{E}_\pi[\mathbb{E}_\pi[\mathbb{1}_H|m_t, x_t]|x_t] = \mathbb{E}_\pi[\mathbb{1}_H|x_t] = p_t,$$

where $\mathbb{E}_\pi$ is the expectation operator over outcomes with respect to $\mathbb{P}_\pi$. Ms. t then combines this induced belief with some realization of the private belief to choose which action to take. Thus, t chooses h upon observing ρ_t and q if and only if $\rho_t \geq q_t$. The (conditional and unconditional) probabilities of taking action h are computed according to equations 2 and 3, but using ρ_t instead of p_t. Agent $t+1$ can compute these probabilities, so she starts the period with an interim belief p_{t+1} according to equation 4, but uses ρ_t instead of p_t.

The results so far shows that any information policy generates stochastic processes $\{\tilde{\rho}_t\}_{t\in\mathbb{N}}$ and $\{\tilde{p}_t\}_{t\in\mathbb{N}}$. They are they connected in the following sense. First, for every realization p_t, the conditional law of the induced beliefs $\tilde{\rho}_t$ equals p_t in expectation. This follows from agents updating induced beliefs after the principal's message according to Bayes rule. Second, for every realization ρ_t, there exists some action a taken with positive probability such that $p_{t+1} = \varphi_a(\rho_t)$. This happens from agents updating induced beliefs after the observation of the history of actions (but not private beliefs). Lemma 1 below[12] shows that the converse also holds.

Lemma 1. *Consider any stochastic processes $\{\tilde{\rho}_t\}_{t\in\mathbb{N}}$ and $\{\tilde{p}_t\}_{t\in\mathbb{N}}$, with initial prior belief p_1 given, such that (i) for every realization of a public belief p_t, the law of the induced belief $\tilde{\rho}_t$ conditional on p_t equals p_t in expectation; (ii) for every realization of an induced belief ρ_t, there exists some action a taken with positive probability such that the next's period public belief is $p_{t+1} = \varphi_a(\rho_t)$. These processes can be generated by an information policy for which the message space is the belief space $[0,1]$, and the information rules depend only on the current public belief.*

The principal's problem is now greatly simplified. He chooses stochastic processes $\{\tilde{\rho}_t\}_{t\in\mathbb{N}}$ and $\{\tilde{p}_t\}_{t\in\mathbb{N}}$, satisfying the requirements of Lemma 1. If Ms. t enters the period with belief p_t, the principal tells her that her induced belief should be some $\rho_t \in \text{supp}(\tau)$, where τ is a probability measure over induced beliefs whose expected value equals p_t. Let $\mathcal{S}(p_t)$ be the set of all such probability measures. The public belief in the next period will be either $\varphi_\ell(\rho_t)$ with probability $1-\alpha(\rho_t)$ or $\varphi_h(\rho_t)$ with probability $\alpha(\rho_t)$. Therefore, if V_δ^{op} is the principal's value function from an optimal policy, then conditional on each p_t, the continuation value can be written as

$$V_\delta^{op}(p_t) \coloneqq \sup_{\tau\in\mathcal{S}(p_t)} \mathbb{E}_\tau\left[(1-\delta)\alpha(\tilde{\rho}_t) + \delta\Big((1-\alpha(\tilde{\rho}_t))V_\delta^{op}(\varphi_\ell(\tilde{\rho}_t)) + \alpha(\rho_t)V_\delta^{op}(\varphi_h(\tilde{\rho}_t))\Big)\right]. \quad (5)$$

One can show that the right-hand side of the above equation is a contraction. As such, a unique value function exists as a fixed point. Moreover, this function is continuous. This, in turn, implies that there exists a probability measure $\tau \in \mathcal{S}(p_t)$ that generates $V_\delta^{op}(p_t)$. Therefore, the supremum is the maximum and there exists an optimal stationary policy. Finally, one can show that the optimal value function must be concave in public beliefs and that an optimal policy needs to generate at most two induced beliefs with positive probability, for any given realization of a public belief[13].

The above equation shows the trade-off principal faces. On the one hand, he can avoid agents following private beliefs by inducing posteriors on cascade sets. This leads to the maximum value of the expected continuation value - the term multiplied by δ, because V_δ^{op} is concave. However, unless the maximum current payoff $\mathbb{E}_\tau[\tilde{\alpha}]$ is achieved by splitting beliefs over C_ℓ and C_h, maximizing tomorrow's value of information implies that the maximum value of information today is not obtained. On the other hand, the principal can minimize the information he shares today to maximize his current payoff. However, if this implies letting Ms. t follow private beliefs to some extent, he might be receiving a lower future value of information than what he could get by inducing future agents into cascades.

[12]The proof of this lemma is an almost exact reproduction of the proof of the obfuscation principle in Ely (2017).
[13]All those claims are proved in Appendix A.

The optimal policy for the illustrative example

To fix ideas, let us reexamine the illustrative example[14]. The private signal space is $S = \{\underline{s}, \bar{s}\}$, and the probability distributions are $f^H(\bar{s}) = f^L(\underline{s}) = \sigma$, for $\sigma \in (1/2, 1)$. Therefore, the belief space is $\{1-\sigma, \sigma\}$ with unconditional probability $g(1-\sigma) = g(\sigma) = 1/2$. The cascade sets are $C_\ell = [0, 1-\sigma)$ and $C_h = [\sigma, 1]$. The conditional and unconditional probabilities of action h (investment) given p are

$$\alpha^H(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ \sigma & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases} \qquad \alpha^L(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ (1-\sigma) & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases}$$

$$\alpha(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ p\sigma + (1-p)(1-\sigma) & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases}$$

The system moves to another public belief according to the transition functions

$$\varphi_h(p) := \begin{cases} p & \text{if } p \in C_\ell, \\ \frac{\sigma p}{p\sigma + (1-p)(1-\sigma)} & \text{if } p \notin C_\ell \cup C_h. \end{cases} \qquad \varphi_\ell(p) := \begin{cases} p & \text{if } p \in C_h, \\ \frac{(1-\sigma)p}{p(1-\sigma) + (1-p)\sigma} & \text{if } p \notin C_\ell \cup C_h. \end{cases}$$

The figures on the right and on the left below present the transition functions and the principal's expected payoff for $\sigma = .8$, respectively. Observe that as long as $p < 1/2$ ($p \geq 1/2$), a single observation of action ℓ (h) brings the posterior to the cascade set C_ℓ (C_h). So for $p_1 = 1/2$, if the first agent chooses investment because she receives a good signal, all subsequent agents will do the same, as the public belief for the second agent will be at the boundary of cascade h. However, if the first agent chooses not to invest due to an observation of a bad signal, the public belief for the second agent will be such that she still gets to follow her private signal, even if she is at the threshold of cascade ℓ.

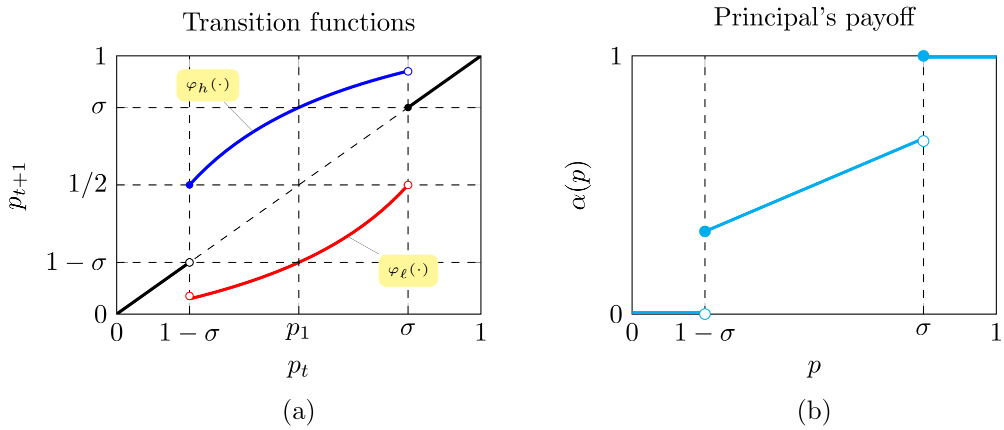

Figure 5: Analysis of the transition functions and principal's expected payoff of the illustrative example. Figure 5(a) represents the law of motion over public beliefs (the blue line corresponds to the one for the investment decision and the red line corresponds to the one for the non-investment decision). Figure 5(b) represents the expected investment probability for each public belief. The figures assume $\sigma = .8$.

[14] Even though I present the theory assuming a continuous private information structure, so far, there is no reason not to use it to analyze an example with a discrete information structure.

The value of the Bayes plausible distribution over beliefs that maximizes the investment probability α at p is called the concave closure of α at p (Aumann, Maschler, and Stearns, 1995). I refer to this value as cav$[\alpha]$. Direct computation shows that

$$\text{cav}[\alpha](p) = \begin{cases} \frac{p}{\sigma} & \text{if } p \notin C_h, \\ 1 & \text{if } p \in C_h. \end{cases} \qquad \text{for } \frac{1}{2} < \sigma \leq \frac{1}{\sqrt{2}};$$

$$\text{cav}[\alpha](p) = \begin{cases} 2\sigma p & \text{if } p \in C_\ell, \\ \left[\frac{(1-\sigma)^2+\sigma^2}{2\sigma-1}\right] p + \left[\frac{2\sigma^2-1}{2\sigma-1}\right](1-\sigma) & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases} \qquad \text{for } \frac{1}{\sqrt{2}} < \sigma < 1.$$

The figures below present the concave closures cav$[\alpha]$ of α for $\sigma = .6$ and $\sigma = .8$, respectively.

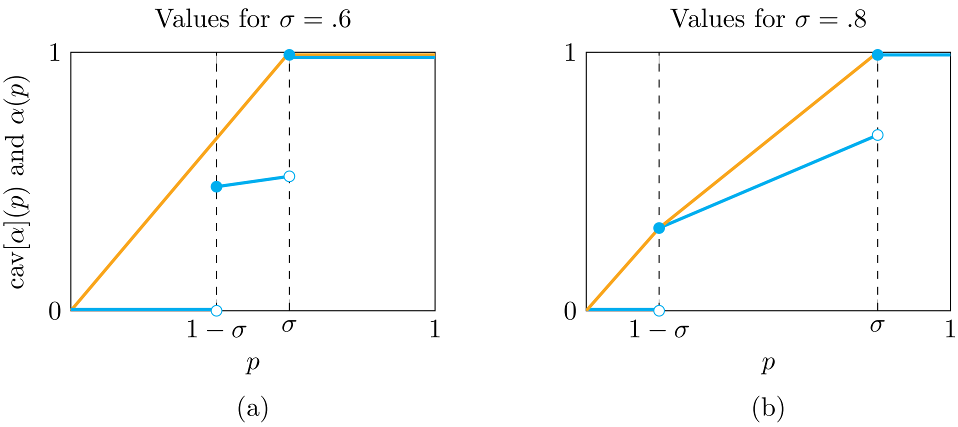

Figure 6: Values of selected functions for different values of σs in the illustrative example. The yellow function is the value of the one-shot concavification and the blue one is the advisor's expected investment probability.

In the repeated interaction, if $1/2 < \sigma \leq 1/\sqrt{2}$, it is straightforward to see from equation 5 that no trade-off arises between maximizing current payoffs and maximizing belief convergence toward cascade set C_h. Indeed, the one-shot optimal splitting of p_1 refrains all future agents from learning from past actions, so a single informative disclosure suffices to reach the value of an optimal policy. This value is $V_\delta^{sd}(p_1) = \text{cav}[\alpha](p_1) = 1/(2\sigma)$.

If $1/\sqrt{2} < \sigma < 1$, the single disclosure strategy does not maximize the advisor's current payoff. Applying the algorithm in equation 5, it is possible to check graphically that inducing belief convergence from the outset cannot be optimal when the precision of the private belief is sufficiently high. Indeed, consider the candidate value function $V^{sd}(p)$, where

$$V^{sd}(p) = \begin{cases} \frac{p}{\sigma} & \text{if } p \notin C_h, \\ 1 & \text{if } p \in C_h. \end{cases}$$

The candidate value function generates other two other value functions: $V^{sd}(\varphi_\ell(p))$ and $V^{sd}(\varphi_h(p))$. These values and the expected continuation value - $\bar{V}^{sd}(p) := \alpha(p)V^{sd}(\varphi_h(p)) + (1-\alpha(p))V^{sd}(\varphi_\ell(p))$ - are

$$V^{sd}(\varphi_\ell(p)) = \begin{cases} \frac{p}{\sigma} & \text{if } p \in C_\ell, \\ \frac{\varphi_\ell(p)}{\sigma} & \text{if } p \in [1-\sigma, \sigma); \end{cases} \qquad V^{sd}(\varphi_h(p)) = \begin{cases} \frac{\varphi_h(p)}{\sigma} & \text{if } p \in [1-\sigma, 1/2), \\ 1 & \text{if } p \in [1/2, 1]; \end{cases}$$

$$\bar{V}^{sd}(p) = \begin{cases} \frac{p}{\sigma} & \text{if } p \in [0, 1/2), \\ p\left[\frac{1-2\sigma(1-\sigma)}{\sigma}\right] + (1-\sigma) & \text{if } p \in [1/2, \sigma), \\ 1 & \text{if } p \in C_h. \end{cases}$$

I represent the compositions below in figure 7(a) below, for $\sigma = .8$. In figure 7(b), I also represent the convex combination between the investment probability α and the candidate function V^{sd} using $(1-\delta) = .5$ and $\delta = .5$ as weights respectively - call it Z^{sd}_δ. If V^{sd} is the value of an optimal policy, this candidate must be the fixed point of equation 5; that is, it must be $\text{cav}[Z^{sd}_\delta](p) = V^{sd}(p)$ for every p. The concavification of Z^{sd}_δ is the dashed line in figure 7(c). From this figure, it can be seen that $\text{cav}[Z^{sd}_\delta]$ and V^{sd} differ outside C_h.

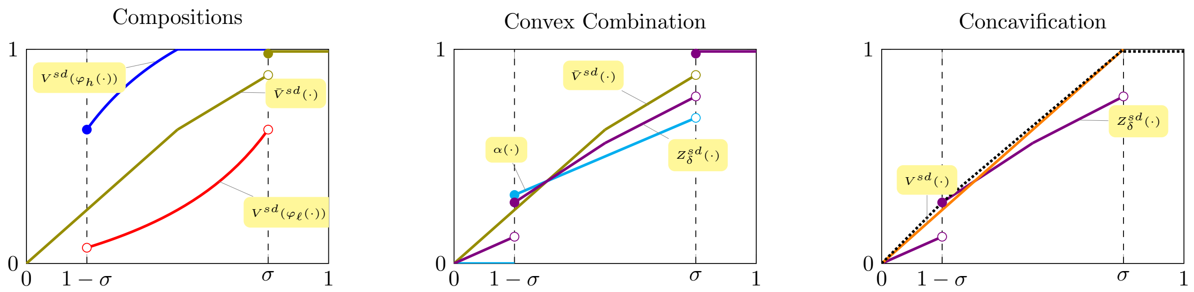

(a) (b) (c)

Figure 7: Testing the optimality of a single disclosure policy for the illustrative example. The values of the parameters are $\sigma = .8$ and $\delta = .5$. Figure 7(a) on the left shows the compositions of V^{sd} with the laws of motion φ_ℓ (red line) and φ_h (blue line). It also represents the convex combination of $V^{sd}(\varphi_\ell(\cdot))$ and $V^{sd}(\varphi_h(\cdot))$ using weights $1-\alpha(\cdot)$ and $\alpha(\cdot)$, respectively (the olive line). The figure 7(b) in the middle represents $Z^{sd}_\delta(\cdot)$ - the convex combination of $\alpha(\cdot)$ and $\bar{V}^{sd}(\cdot)$ using $1-\delta$ and δ as weights, respectively (the violet line). Figure 7(c) represents the concave closure of the composition Z^{sd}_δ (the dashed line) and contrasts with the candidate value function V^{sd} (the orange line). It can be observed that $\text{cav}[Z^{sd}_\delta](p) \neq V^{sd}(p)$ for $p \notin C_h$.

What is the optimal value function for the illustrative example? As Proposition 1 shows - whose proof is in Appendix B, it is the value function arising from minimizing the information disclosed to maximize the expected current payoff at every realization of the public belief process. In other words, for every p_t, the principal induces posteriors according to the probability $\tau \in \mathcal{S}(p_t)$ that maximizes $\mathbb{E}_{\tau(p_t)}[\alpha]$. At $p_1 = 1/2$, this leads to the values of an optimal policy below. For uninformative private information, the greedy and the single disclosure policies coincide; for informative private information, it is optimal to induce some investors to follow their private signals.

Proposition 1. *In the illustrative example, the value of an optimal policy for $\sigma > \frac{1}{\sqrt{2}}$ is*

$$V_\delta(p) = \begin{cases} p\left(\frac{\sigma(2-\delta)}{1-\delta+\delta\sigma^2}\right) & \textit{if } p \in C_\ell, \\ p\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) + (1-\sigma)\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) & \textit{if } p \notin C_\ell \cup C_h, \\ 1 & \textit{if } p \in C_h. \end{cases}$$

This value function is achieved through a greedy policy, that is, a policy that induces posterior beliefs to generate cav$[\alpha](p)$ at every public belief p. This means that whenever $p < 1-\sigma$, the principal induces posteriors 0 and $1-\sigma$, and whenever $p \in (1-\sigma, \sigma)$, the principal induces posteriors $1-\sigma$ and σ. For beliefs $p \geq \sigma$, principal does not disclose any additional information.

4.1 Belief dynamics

In the illustrative example, for any value of the private signal's precision, it is always optimal not to change disclose any additional information for $p \geq \sigma$, the lower bound of the cascade set on the principal's most preferred action. For any private information structure, whenever $p_t \in C_h$, it is optimal to the principal not to release any additional information. Indeed, without any disclosure, agent t will take action h, regardless of private beliefs. Therefore, Ms. $t+1$ will not learn anything new from the observation of the agent's t action and will choose action h as well. Releasing additional information in this case can only potentially harm the principal. To see this, consider any $p \in C_h$ and any $\tau \in \mathcal{S}(p)$. Let $\mathbb{1}_p$ be the probability measure that assigns probability one to $\rho = p$. Because V_δ^{op} is concave and α is at most 1,

$$1 - \delta + \delta V_\delta^{op}(p) = \mathbb{E}_{\mathbb{1}_p}\left[Z_\delta^{op}\right] \geq \mathbb{E}_\tau\left[Z_\delta^{op}\right].$$

This means that $V_\delta^{op}(p) = 1$ for every $p \in C_h$. Note that there are no conflicting effects, that is, the principal maximizes his current expected payoff without sacrificing the continuation value or drifting society away from action h.

In addition, in the illustrative example, for any value of σ, the principal always splits beliefs $p < 1-\sigma$, into 0 and other belief outside C_ℓ. This also generalizes to any private information structure. To drive the public belief process away from this cascade set with some probability, the principal must induce a higher belief $\rho^+ > p$ that makes action h at least considerable - and a lower belief $\rho^- < p$ that does not change the decision to choose ℓ no matter the private belief. Proposition 2 shows that, to recommend agent t to choose action ℓ irrespective of private beliefs, the principal must partially avoid any release of future information, by setting $\rho^- = 0$.

Proposition 2. *Suppose that private beliefs are bounded. For any positive public belief $p > 0$ belonging to the principal's cascade set on the least preferred action C_ℓ, it is optimal to induce beliefs $\rho^- = 0$ and ρ^+ outside C_ℓ.*

Proof. If private beliefs are unbounded, $C_\ell = \{0\}$ and there is nothing else to release. Suppose then C_ℓ is a proper interval and pick any $p > 0 \in C_\ell$. Assume by way of contradiction that any policy leading to the optimal value function splits p into at most two points ρ^- and ρ^+, both within C_ℓ and such that $0 < \rho^- \leq p \leq \rho^+ \leq \underline{q}$. This leads to $V_\delta^{op}(p) = 0$. Indeed, because V_δ^{op} is concave, any optimal strategy yields

$$V_\delta^{op}(p) \leq \delta V_\delta^{op}(p).$$

This means that $V_\delta^{op}(p) = \delta V_\delta^{op}(p) = 0$, for $\delta < 1$. Now take $\rho^- > \varepsilon > 0$ small enough and define $\varepsilon' := \min\{\rho^- - \varepsilon, \underline{q} - \rho^+ + \varepsilon\}$. Note that $\varepsilon' > 0$. Consider two points: $\hat{\rho}^- = \rho^- - \varepsilon'$ and $\hat{\rho}^+ := \rho^+ + \varepsilon'$. Note that $0 < \hat{\rho}^- < \rho^-$ and $\hat{\rho}^+ > \underline{q}$. As such, $\hat{\rho}^- < \hat{\rho}^+$. Consider the probability distribution $\hat{\tau} := (\hat{\beta}, 1-\hat{\beta})$, where

$$\hat{\beta} = \frac{p - \hat{\rho}^-}{\hat{\rho}^+ - \hat{\rho}^-}.$$

The distribution $\hat{\tau}$ belongs to $\mathcal{S}(p)$. However, this contradicts the split placing both posteriors in C_ℓ being optimal, because

$$\begin{aligned}\mathbb{E}_{\hat{\tau}}\left[Z_\delta^{op}\right] &= \hat{\beta}[(1-\delta)\alpha(\hat{\rho}^+) + \delta\bar{V}_\delta^{op}(\hat{\rho}^+)] + (1-\hat{\beta})[\delta\bar{V}_\delta^{op}(\hat{\rho}^-)], \\ &> \delta\mathbb{E}_{\hat{\tau}}\left[\bar{V}_\delta^{op}\right], \\ &\geq 0.\end{aligned}$$

□

Note that there might be conflicting effects when $p \in C_\ell$. The principal wants to drift society away from action ℓ and to do so he must disclose additional information. He could provide sufficient information to make all future agents take action h no matter the private beliefs, by inducing $\rho^+ \in C_h$. However, depending on the distribution of private beliefs, this could lead to a lower *ex-ante* probability of agent t choosing h, because the principal can only induce a higher ρ^+ by recommending h less often when the state is L. Alternatively, he could minimize the information released to maximize the *ex ante* probability of agent t choosing h, by inducing $\rho^+ \notin C_\ell \cup C_h$ that maximizes $(1/\rho^+)\alpha(\rho^+)$. However, because agent t will choose according to her private beliefs, agent $t+1$ might learn something beyond what was disclosed to agent t and this might reduce the principal's expected continuation value. I will investigate this trade-off in deeper later sections.

The discussion thus far leads to the following corollary. Define $C_\ell^{op} := \{p : \alpha(\rho) = 0 \ \forall\rho \in \text{supp}(\tau^{op}(p))]$ and $C_h^{op} := \{p : \alpha(\rho) = 1; \forall\rho \in \text{supp}(\tau^{op}(p))\}$, where each $\tau^{op}(p) \in \mathcal{S}(p)$ is a probability measure that leads to the optimal continuation value at p. Note that $C_a^{op} \subseteq C_a$ for every $a \in \{h, \ell\}$. That is, the principal can only shrink the cascade sets. Under any optimal policy, C_ℓ^{op} is always minimal and C_h^{op} is maximal. Intuitively, a non-degenerate cascade set C_ℓ has no value to the principal: he can always persuade society out of it if $p \in C_\ell$, provided that persuasion is at least possible - that is, if $p > 0$.

Corollary 1. *Under any optimal policy, the principal always induces the minimal cascade set on the least preferred action ℓ and the maximal on his most preferred action h: $C_\ell^{op} = \{0\}$ and $C_h^{op} = C_h$.*

Although manipulated, the public belief process $\{\tilde{p}_t\}_{t\in\mathbb{N}}$ continues to be a martingale. To see this, consider any public history x_t - recall that $x_t = \{a_\tau, \rho_\tau\}_{\tau=1}^{t-1}$ - that leads to a public belief of p_t. Consider any information policy π with the associated $\tau \in \mathcal{S}(p_t)$. From equation 4,

$$\mathbb{E}_\pi[\tilde{p}_{t+1}|x_t] = \mathbb{E}_\tau\left[\alpha(\tilde{\rho}_t)\left(\frac{\alpha^H(\tilde{\rho}_t)}{\alpha(\tilde{\rho}_t)}\right)\tilde{\rho}_t + (1-\alpha(\tilde{\rho}_t))\left(\frac{1-\alpha^H(\tilde{\rho}_t)}{1-\alpha(\tilde{\rho}_t)}\right)\tilde{\rho}_t\right] = \mathbb{E}_\beta[\tilde{\rho}_t] = p_t.$$

Being a martingale, the process almost surely converges to $\tilde{p}_\infty$. Similar to the no-policy case, every realization of this random variable must belong to the absorbing sets. However, these sets are now $C_\ell^{op} = \{0\}$ and $C_h^{op} = C_h$, as I show in the next proposition.

Proposition 3. *Under any optimal policy, the public belief process almost surely converges to the induced cascade sets $C_\ell^{op} = \{0\}$ and $C_h^{op} = C_h$.*

Proof. First, note that the induced belief process $\{\tilde{\rho}_t\}_{t\in\mathbb{N}}$ is a martingale as well. Indeed, fix an optimal policy π. Because $\mathbb{E}_\pi[\tilde{p}_{t+1}|x_t, \rho_t] = \rho_t$ and $\mathbb{E}_\pi[\tilde{\rho}_{t+1}|x_{t+1}] = p_{t+1}$, the law of total expectation implies that

$$\mathbb{E}_\pi[\tilde{\rho}_{t+1}|x_{t+1}] = p_{t+1} \Rightarrow \mathbb{E}_\pi[\tilde{\rho}_{t+1}|x_t, \rho_t] = \mathbb{E}_\pi[\mathbb{E}_\pi[\tilde{\rho}_{t+1}|\tilde{x}_{t+1}]|x_t, \rho_t] = \mathbb{E}_\pi[\tilde{p}_{t+1}|x_t, \rho_t] = \rho_t.$$

Because the process is a martingale, it converges almost surely to a random variable $\tilde{\rho}_\infty$. Considering this, assume by way of contradiction that there exists some p_∞ in the support of $\tilde{p}_\infty$ that does not belong to $\{0\} \cup C_h$. Let $\tau \in \mathcal{S}(p_\infty)$ be the associated optimal Bayes plausibility measure for p_∞. There must exist some ρ in the support of τ such that $\underline{q} < \rho < \bar{q}$. So consider an open interval I around ρ such that $I \subset (\underline{q}, \bar{q})$. It is possible to find some $\varepsilon > 0$ with the following property. For every $\rho' \in I$, either (i) $\alpha(\rho') > \varepsilon$ and $|\varphi_h(\rho') - \rho'| > \varepsilon$ or (ii) $\alpha(\rho') < 1 - \varepsilon$ and $|\varphi_\ell(\rho') - \rho'| > \varepsilon$. This follows from α being continuous as well as from $0 < \alpha^L(\rho') < \alpha^H(\rho') < 1$. Claim 3 from Appendix A implies that I does not contain any induced beliefs in the support of $\tilde{\rho}_\infty$. This is turn proves the existence of an open set I' containing ρ with measure zero with respect to the law of $\tilde{\rho}_\infty$. However, because p_∞ belongs to the support of $\tilde{p}_\infty$, this is only possible if I' also has measure with respect to τ, contradicting $\rho \in \text{supp}(\tau)$.

□

As an implication of Proposition 3, learning will be *partially* complete and correct under bounded private beliefs[15]. To see this, suppose the true state is H. Because the public belief process converges, the stationary public belief must place positive probability on points in C_h and/or in $\{0\}$. However, because the belief process is a martingale, agents cannot be dead wrong about the state, that is, they cannot hold belief 0 in equilibrium. Therefore, all beliefs must belong to the correct cascade set. However, this process cannot jump to the extreme belief 1. Thus, when the state is H, learning is correct, but not complete. Suppose now that the true state is L. Learning can be incorrect with positive probability if private beliefs are boundedly informative. However, learning can also be complete with positive probability, because the cascade set on action ℓ is degenerate. Corollary 2 summarizes these observations.

Corollary 2. *Assume that the private beliefs are bounded. Under any optimal policy, learning is always correct but incomplete if the true state is H. Conversely, learning can be incorrect, but it can also be complete, if the true state is L.*

Comparing the learning outcomes with the no-policy case, one sees that the selfish principal actually makes society better off. This occurs because one set in which no additional information is generated (the set C_ℓ) shrinks to a singleton. Thus, the principal eliminates one set of informational inefficiencies. When the true state is H, only a correct, good herd can arise. When the state is L, unlike the belief dynamics without intervention, there is a probability of complete learning even with bounded private beliefs.

4.2 Valuable social learning

Going back to the illustrative example, with a binary and symmetric private information structure, the principal optimally allows agents to learn from past actions if and only if private signals are very revealing. Does this observation generalize to a broader class of private information structures? This section addresses such inquiry. For log-concave private belief densities, I will show that single disclosure is optimal if and only if the right tail of such density is quite fat. One interpretation of this result is that social learning is valuable to the advisor if and only if private information unfavorable to investment is rare or contrarian behavior on high public beliefs is unlikely. For unbounded private beliefs, single disclosure will never be optimal.

[15] Recall that, with unbounded private information, learning is always complete - thus correct.

Before proceeding, let me return to the trade-off between maximizing the value of information today and the value of information tomorrow. Recall that the optimal value function V_δ^{op} must be concave in public beliefs. Let τ^{op} be an associated optimal probability measure over posteriors, for any public belief. Because Bayes plausibility is required, the value of the dynamic interaction is bounded above by the value of the static interaction:

$$\begin{aligned} V_\delta^{op}(p) &= (1-\delta)\mathbb{E}_{\tau^{op}(p)}[\alpha] + \delta\mathbb{E}_{\tau^{op}(p)}\left[\bar{V}_\delta^{op}\right], \\ &\leq (1-\delta)\text{cav}[\alpha](p) + \delta V_\delta^{op}(p), \\ \therefore V_\delta^{op}(p) &\leq \text{cav}[\alpha](p). \end{aligned}$$

For $p \in C_h$, this upper bound is achieved. This is just another way of seeing that social learning does not impose conflicting effects when beliefs are in the cascade set C_h. Now recall that, for every $p > 0$ and $p \notin C_h$, the single disclosure splitting induces beliefs $\rho^- = 0$ and $\rho^+ = \bar{q}$ with probabilities $1 - p/\bar{q}$ and $p/\bar{q}$, respectively. Because the optimal value function is a fixed point of the contraction algorithm in equation 5, it is also true that the value of information outside C_h is bounded below by the value of shutting down learning. Formally,

$$V_\delta^{op}(p) \geq \frac{p}{\bar{q}}\left[(1-\delta)1 + \delta V_\delta^{op}(\bar{q})\right] + \left(1 - \frac{p}{\bar{q}}\right)\left[(1-\delta)0 + \delta V_\delta^{op}(0)\right] = V^{sd}(p).$$

Since $V^{sd} \leq V_\delta^{op} \leq \text{cav}[\alpha]$, whenever $V^{sd} = \text{cav}[\alpha]$ it is the case that $V^{sd} = V_\delta^{op}$. In other words, if the maximization of the static value of information implies inducing posteriors in the extreme points of cascade sets, then shutting down learning from the outset is feasible and desirable from the principal's viewpoint. In fact, it is easier to check whether the single disclosure strategy is optimal: just comparing α and V^{sd}. The next proposition proves that this effectively characterizes when social learning has no value to the principal.

Proposition 4. *Single disclosure is optimal if and only $\alpha(p) \leq V^{sd}(p)$ for every $p \in (\underline{q}, \bar{q})$.*

Proof. Suppose first that $\alpha \leq V^{sd}$. Because V^{sd} is an affine function majorizing α, it must be that

$$\text{cav}[\alpha](p) := \inf\{f(p) \;\text{ s.t. }\; f \in \mathbb{R}^{[0,1]} \;\text{ affine and }\; f \geq \alpha\} \leq V^{sd}(p).$$

This implies that single disclosure is optimal for every public belief, because $V^{sd} \leq V^{op} \leq \text{cav}[\alpha]$. Suppose now that there exists some $p \notin C_h$ such that $\alpha(p) > V^{sd}(p)$. Because $V_\delta^{op}(p) \geq Z_\delta^{op}(p)$ - the value of not disclosing anything at p and resorting to the optimal policy next period, it follows that

$$\begin{aligned} V_\delta^{op}(p) &\geq (1-\delta)\alpha(p) + \delta\left[\alpha(p)V_\delta^{op}(\varphi_h(p)) + (1-\alpha(p))V_\delta^{op}(\varphi_\ell(p))\right], \\ &> (1-\delta)V^{sd}(p) + \delta\left[\alpha(p)V^{sd}(\varphi_h(p)) + (1-\alpha(p))V^{sd}(\varphi_\ell(p))\right], \\ &\geq \alpha(p)V^{sd}(\varphi_h(p)) + (1-\alpha(p))V^{sd}(\varphi_\ell(p)). \end{aligned}$$

The second inequality follows from $\alpha(p) > V^{sd}(p)$ and $V_\delta^{op} \geq V^{sd}$. The third inequality follows from V^{sd} being concave and $\mathbb{E}[\tilde{p}'|p] = p$. There are two cases to consider. In the first case, $\varphi_h(p) \notin C_h$. Then, $V^{sd}(\varphi_a(p)) = \varphi_a(p)/\bar{q}$ for $a \in \{h, \ell\}$. This implies that $V_\delta^{op}(p) > V^{sd}(p)$. In the second case, $\varphi_h(p) \in C_h$. Then,

$$\alpha(p)V^{sd}(\varphi_h(p)) + (1-\alpha(p))V^{sd}(\varphi_\ell(p)) = \alpha(p)1 + (1-\alpha(p))V^{sd}(\varphi_\ell(p)) > V^{sd}(p).$$

This again implies that $V_\delta^{op}(p) > V^{sd}(p)$; that is, it is not optimal to shut down social learning at p. □

At this point, some further assumptions are necessary to derive new insights. From now on, I restrict the analysis to a rich class of probability densities: the log-concave class. I will also impose a technical condition - differentiability - to simplify the exposition. Assuming the unconditional g to be log-concave over $(\underline{q}, \bar{q})$ means that $\ln g$ is a concave function over $(\underline{q}, \bar{q})$. Equivalently, this means that the ratio g'/g is non-increasing in its domain. Many distributions commonly used in economics have log-concave densities: uniform, normal and exponential, to name a few. An (1998) and Bagnoli and Bergstrom (2005) are excellent surveys of nice properties of log-concave densities[16]. Log-concavity here will be useful for generating regularity in the expected probability α.

Assumption 1. *The private belief density g is log-concave and differentiable on $(\underline{q}, \bar{q})$.*

The differentiability of g implies that α is twice differentiable over $(\underline{q}, \bar{q})$. By doing so and simplifying the result, the following expression is obtained:

$$-\alpha''(p) = 4p(1-p)g(p)\left[\frac{3}{2}\left(\frac{2p-1}{p(1-p)}\right) - \frac{g'(p)}{g(p)}\right]. \tag{6}$$

The term multiplying $3/2$ has a single-crossing property, that is, it crosses the horizontal axis only once and from below[17]. If $-\alpha''$ inherits the same property on $(\underline{q}, \bar{q})$, then the *ex ante* expected probability α will be convex up to a point and concave after it. Because the concave closure of α in C_ℓ is linear, the static problem then will be to find the maximum inclination ι such that ιp touches $\alpha(p)$ at some point ρ^+. In other words, the static persuasion problem breaks down to maximize $\alpha(\rho)/\rho$. Note that the point ρ^+ must be at least higher than the inflection point. Moreover, the single disclosure policy will be optimal if and only if $\rho^+ = \bar{q}$.

Quah and Strulovici (2012) proved that a linear combination of two single-crossing functions has the single-crossing property if and only if they satisfy what they called signed-ratio monotonicity. Briefly, if a form of monotonicity of the ratio of those functions holds even when the signs of the functions are different[18]. Hence, for $-\alpha''$ to have the single-crossing property, $-g'(q)/g(q)$ would have to have the single-crossing property as well. Moreover, $-(\ln q(1-q))'$ and $(\ln g(q))'$ must satisfy the signed-ratio monotonicity. This will be the case for g log-concave, as Lemma 2 demonstrates.

Lemma 2. *If the private belief density g is log-concave, then α is convex-concave on $(\underline{q}, \bar{q})$.*

Profiting from Lemma 2, Theorem 1 characterizes the optimality of the single disclosure policy in terms of the private information structure solely, provided that the private belief density is log-concave. Specifically, social learning is not valuable to the principal if and only if there is a high mass concentration of belief at the right tail of the density.

[16]Log-concavity of the private density does not imply neither is implied by log-concavity of private signals. I will have nothing to say about the general conditions for which distributions over private signals generate unconditional log-concave densities, but Roesler (2014) offers some insights about this. Recall, however, that the boundedness of private signals does translates it into the boundedness of private beliefs. Moreover, discrete signal distributions cannot be log-concave, as they are not atomless.

[17]A function f satisfies this if $f(s') \geq 0 \Rightarrow f(s'') \geq 0$ whenever $s'' > s'$ and $f(s') > 0 \Rightarrow f(s') > 0$ whenever $s'' > s'$.

[18]As Quah and Strulovici (2012) defines it, two functions f and $\hat{f}$ satisfy the signed-ratio monotonicity if (i) at any $r' : \hat{f}(r') < 0$ and $f(r') > 0$, $(-\hat{f}(r')/f(r')) \geq (-\hat{f}(r'')/f(r''))$ whenever $r'' > r'$; (ii) at any $r' : f(r') < 0$ and $\hat{f}(r') > 0$, $(-f(r')/\hat{f}(r')) \geq (-f(r'')/\hat{f}(r''))$ whenever $r'' > r'$.

Here is the intuition for this result. Suppose that the private information structure is boundedly revealing. Recall that private beliefs close to $\bar{q}$ mean higher beliefs about the state being L. If higher private beliefs are likely, this acts against the principal's interest. If he allows agents to follow private beliefs, even if he induces a high posterior belief, a private realization of q near $\bar{q}$ could drive down the public belief process. Thus, outside C_h, the *ex-ante* expected probability α is higher under the single disclosure policy than under any other policy. Note that single disclosure achieves the highest value cav$[\alpha]$ in this case.

However, if higher private beliefs about the state being L are not likely, then the principal can expect that agents will follow a recommendation to choose action h with a high probability. Behavior contrary to the principal's recommendation is possible, but relatively unexpected. Thus, outside C_h, there is a strategy that leads to a higher expected probability α than the single disclosure one.

Theorem 1. *Assume that private beliefs are bounded and that the density of private beliefs is log-concave in $(\underline{q}, \bar{q})$. Single disclosure is optimal if and only if the right tail of the private belief density is sufficiently fat. Formally, for any $\delta \in (0, 1)$,*

$$V_\delta^{op}(p) = V^{sd}(p) \Leftrightarrow \lim_{q \uparrow \bar{q}} g(q) \geq \frac{1}{4(1-\bar{q})\bar{q}^2} \quad \forall p < \bar{q}.$$

Proof. Suppose that single disclosure is optimal, that is, $\alpha \leq p/\bar{q}$ for every $p < \bar{q}$. Then

$$\frac{1}{\bar{q}} \leq \frac{1-\alpha(p)}{\bar{q}-p}.$$

Taking the left limit of the right side of the inequality at $\bar{q}$ - the limit exists because α is concave near $\bar{q}$ - leads to $\alpha'(\bar{q}_-) \geq 1/\bar{q}$. In Appendix A, I show that this left limit equals $4\bar{q}(1-\bar{q})g(\bar{q}_-)$. Rearranging the inequality, it follows that

$$g(\bar{q}_-) \geq \frac{1}{4(1-\bar{q})\bar{q}^{\,2}}.$$

Now assume that the above inequality is reversed. I need to show that this leads to single disclosure not being optimal. From the computation of $\alpha'(\cdot)$, one can show that $\alpha'(\bar{q}_-) < 1/\bar{q}$. Because α is concave on an interval near $\bar{q}$, there exists some p close enough to $\bar{q}$ (but below $\bar{q}$) such that $\alpha'(\bar{q}_-) \leq \alpha'(p) < 1/\bar{q}$. Also, it must be that

$$\alpha(p) + \alpha'(p)(p'-p) \geq \alpha(p') \;\; \forall p'.$$

In particular, for $p' = \bar{q}$, $\frac{1-\alpha(p)}{\bar{q}-p} \leq \alpha'(p)$. Therefore,

$$\frac{1-\alpha(p)}{\bar{q}-p} < \frac{1}{\bar{q}} \;\text{ or }\; \alpha(p) > \frac{p}{\bar{q}}.$$

□

One final remark regarding the log-concavity assumption is that log-concave densities have exponential tails (An, 1997; Cule and Samworth, 2010). This means that the right tail goes to zero fast as q goes to 1 and the threshold inequality for single disclosure being optimal does not hold. Therefore, for the log-concave class, single disclosure will never be optimal when private beliefs are unbounded, as corollary 3 evidences.

Corollary 3. *Assume that private beliefs are unbounded and that the density of private beliefs is log-concave in* $(0, 1)$. *Single disclosure is never optimal: there is always some public belief above which* $V^{op} > V^{sd}$.

Collecting results, social learning has some value to the principal whenever the expected investment probability at some public belief near his preferred cascade set is higher than the expected investment probability from single disclosure. If this is the case, the principal can come up with a better split at this public belief to maximize α and resort to single disclosure at a later time. With log-concave private belief density, this condition can be characterized in terms of the right tail of g only. Social learning is valuable if and only if private beliefs unfavorable to action h are rare. With unbounded private beliefs, this is always the case because, although private information can be fully revealing, the probability of a contrarian agent in public beliefs near 1 is small.

An example with log-concave private belief density

Discrete private belief distributions cannot be log-concave, so the illustrative example fails to capture the results in this section. Let me introduce then another example to fix ideas[19]. Let $\underline{q} := (1/2)(1 - \sigma)$ and $\bar{q} := (1/2)(1 + \sigma)$ where $\sigma \in [0, 1]$. As in the first example, the parameter σ controls to which extent the private beliefs can be unbounded. The unconditional density is uniform over $[\underline{q}, \bar{q}]$. Under this uniform density, I compute the the expected probability of taking action h and the transition functions in Appendix B. Here, I provide a visual representation of these functions as well as the single disclosure policy in figure 8, for different values of σ. Specifically, the first line shows the functions for $\sigma = .4$ and the second line shows the functions for $\sigma = 8$.

In this example, the single disclosure strategy is optimal whenever $\sigma \leq \sigma^* \approx 0.54$. In this case, depicted in the figures of the first line, the value of a one-shot concavification and the single disclosure policy coincide, so at $p_1 = 1/2$, $V^{op}_{\delta} = 1/(2\sigma)$. Whenever $\sigma > \sigma^*$, one can see that $\alpha(p) > p/\bar{q}$ for any $p < \bar{q}$ above a threshold p^*, represented in the graphs. Therefore, at p_1, it is safe to say that $V^{op}_{\delta} > V^{sd}$, that is, some social learning is valuable to the principal.

[19]This example comes from Herrera and Hörner (2012).

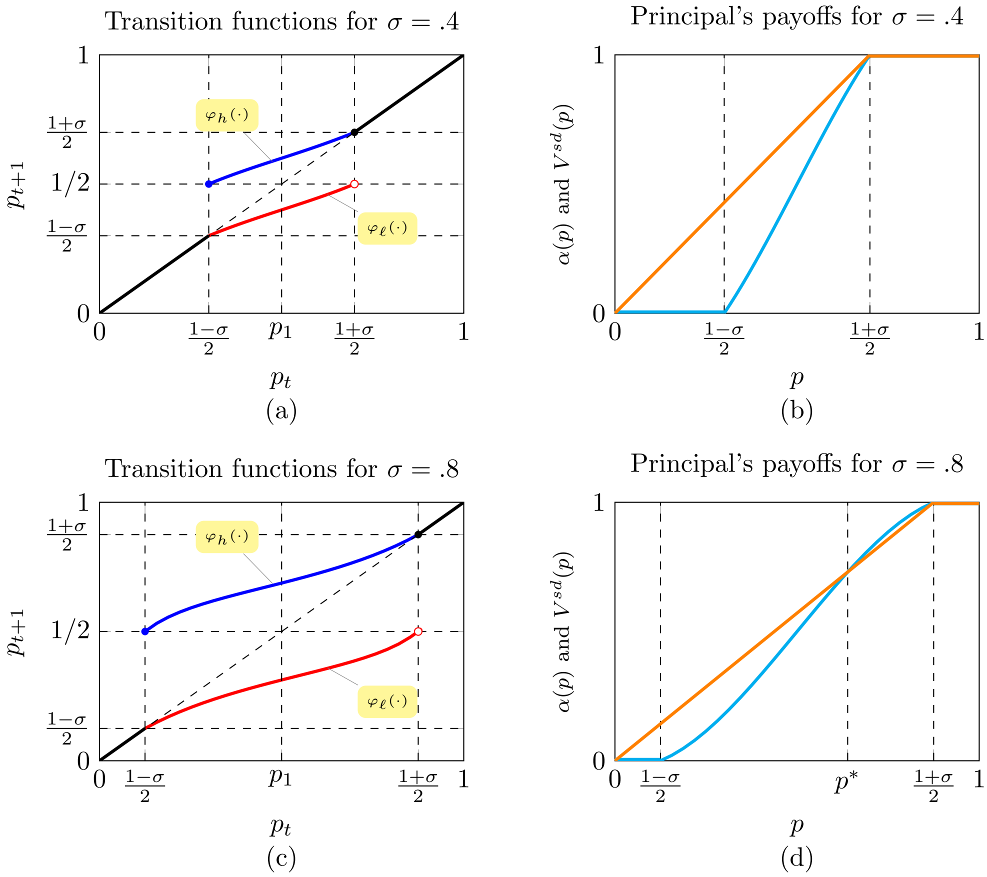

Figure 8: Expected probability vs. single disclosure value function for different values of σ. Figures 8(a) and 8(b) represent relevant functions with $\sigma = .4$, and figures 8(c) and 8(d) show the same functions with $\sigma = .8$. The blue lines in figures 8(b) and 8(d) show the investment probabilities, and the orange lines show the value of a single disclosure policy.

4.3 The role of patience

In the case that single disclosure is not the optimal policy for the second example, then what policy is? The greedy one? It turns out that an explicit computation of the value function for continuous private signals is a daunting task. This is most often true whenever α is concave or convex outside C_H or when there is only one law of motion that is exogenous to agents' action. However,, with multiple laws and expected investment probability being convex-concave, it will not necessarily be the case. Nevertheless, I will have a few things to say about the long-run value of information.

As the principal becomes infinitely patient, the optimal value function converges pointwise to the single disclosure value function. This does not depend on the private information structure - and it can be seen from the policy derived in the illustrative example. The result follows mainly from the stationarity of the optimal policy. Intuitively, the more patient he is, the more he cares about the stationary probability of herds. As he always has informational power to induce herd behavior, the short-run value of social learning is less important to him. Thus, for high values of δ, the simplest strategy - not caring about social learning - might be reasonably close the highest possible payoff the principal can receive in the long run.

Let me present this result formally. Any optimal policy π - given the initial prior - induces a sequence of probability measures $\{\hat{\lambda}_t^\pi\}_{t\in\mathbb{N}}$ over the induced belief space. Therefore, I can write the principal's value from an optimal policy as a function of the induced belief process:

$$V_\delta^{op} = \sum_{t\in\mathbb{N}}(1-\delta)\delta^{t-1}\mathbb{E}_{\hat{\lambda}_t^\pi}[\alpha].$$

Moreover, because the public belief process converges to the new cascade sets (from Proposition 3), so does the induced one. This means that informative communication must eventually settle down. Indeed, if $p = 0$ with positive probability in the long-run, the principal cannot split beliefs further; if $p \in C_h$, there is no reason to split beliefs. Therefore, in the limiting case (that is, as the discount factor goes to 1), I can interchangeably talk about either the laws of induced beliefs $\{\hat{\lambda}_t^\pi\}_{t\in\mathbb{N}}$ or public beliefs $\{\lambda_t^\pi\}_{t\in\mathbb{N}}$. Lemma 3 then follows.

Lemma 3. *Let π be an optimal policy. The associated value function must converge to the stationary value of the public belief process hitting C_h under π. Precisely,*

$$\lim_{\delta\to 1} V_\delta^{op} = \lim_{t\to\infty}\mathbb{E}_{\hat{\lambda}_t^\pi}[\alpha] = \lambda_\infty^\pi(C_h).$$

Define now the belief $\rho_+^\pi = \mathbb{E}_{\lambda_\infty^\pi}[p|p \in C_h]$. Because the public belief process is a martingale and $C_\ell^\pi = \{0\}$, it must be that $\rho_+^\pi\lambda_\infty^\pi(C_h) = p_1$. Clearly, the split of p_1 in ρ_+^π with probability $\lambda_\infty^\pi(C_h)$ and 0 with probability $1-\lambda_\infty^\pi(C_h)$ is Bayes plausible at $t = 1$. Moreover, it places posteriors at cascade sets from the outset, undermining any necessity of disclosing additional information at $t = 2$. Note that this strategy yields the same long-run value $\lim_{\delta\to 1} V_\delta^{op}$. Moreover, it has to give principal a lower value than the single disclosure strategy, because the latter is the best strategy among those that disclose informative messages only at the beginning. Thus $\lim_{\delta\to 1} V_\delta^{op} \le V^{sd}$. Reverse inequality must also be true. Indeed, by definition, $V_\delta^{op} \ge V^{sd}$ for every discount factor $\delta < 1$; in particular, it must hold for δ close to one. In summary, I have proved the following theorem.

Theorem 2. *The value of an optimal policy approaches the value of the single disclosure one, as principal becomes increasingly patient:*

$$\lim_{\delta\to 1} V_\delta^{op} = V^{sd}.$$

5 Private communication

Suppose the principal still can publicly commit to an information policy, but now can restrain current agents from observing past realizations of messages. In this sense, communication is private. Because agents' strategy will only depend on the observation of the action history - not message histories and current messages, it is without loss to consider information rules that are maps from action histories to a distribution over messages. As in the previous sections, along with agents' strategies and the prior belief, the policy generates a probability measure over the set of public outcomes (the set of action histories).

For a given information policy and a given strategy for the agents, upon the observation of a history a^t, Ms. $t+1$ will have an interim belief p_{t+1} about the state being H. However, because she does not observe what message Ms. t received, she needs to average out all possible realizations of induced beliefs ρ_t that led t to take the observed action a_t, given that t observed history a^{t-1}. Therefore - and by law of total expectation - it is conditionally and unconditionally expected that t takes action h with the probabilities below, respectively.

$$\hat{\alpha}^{\theta}(p_t, \tau_t) := \mathbb{E}_{\tau_t^H(p_t)}[\alpha(\rho_t)] \text{ for } \theta \in \{H, L\}; \tag{7}$$

$$\hat{\alpha}(p_t, \tau_t) := p_t \hat{\alpha}^H(p_t, \tau_t) + (1 - p_t)\hat{\alpha}^L(p_t, \tau_t). \tag{8}$$

After observing action $a_t = a$, agent $t + 1$ updates her public belief according to

$$\tilde{p}_{t+1} = \hat{\varphi}_a(p_t, \tau_t) = \begin{cases} \left[\frac{\hat{\alpha}^H(p_t,\tau_t)}{\hat{\alpha}(p_t,\tau_t)}\right] p_t & \text{if } a = h, \\ \\ \left[\frac{1-\hat{\alpha}^H(p_t,\tau_t)}{1-\hat{\alpha}(p_t,\tau_t)}\right] p_t & \text{if } a = \ell. \end{cases} \tag{9}$$

The simplifications discussed in previous subsections still hold here. Specifically, it is without loss to focus on direct (the message space is the belief space and principal tells agents exactly what their beliefs should be) and Markov (information rule only depends on public history through the realization of interim beliefs) information policies. Because these simplifications hold, I can reformulate the principal's problem in terms of Markov decision problem, same way as before. However, the realization of a current public belief does not have a direct effect on the next period's belief. The principal must consider the average effect a given current distribution over messages will have on the next agent's inference about the state of the world. Thus, the value of an optimal policy must satisfy

$$V_\delta^{op}(p) = \underset{\tau \in \mathcal{S}(p)}{max}\left[(1-\delta)\hat{\alpha}(p,\tau) + \delta\Big(\hat{\alpha}(p,\tau)V_\delta^{op}(\hat{\varphi}_h(p,\tau)) + (1-\hat{\alpha}(p,\tau))V_\delta^{op}(\hat{\varphi}_\ell(p,\tau))\Big)\right] \;\; \forall \;\; p \in [0,1]. \tag{10}$$

The equation above still is an operator and satisfies Blackwell's sufficient conditions for a contraction. However, the fact that the value function is still concave is not straightforward, as this is not a dynamic concavification operator the same way as in previous sections. Nevertheless, Lemma 4 below proves that concavity is preserved in a private persuasion mechanism.

Lemma 4. *With private communication, the function V_δ^{op} is concave in* $(0, 1)$.

Lemma 4 implies that one of the crucial features of the results in the previous sections is preserved under a private communication mechanism. Namely, the concavity of the value function. Recall that concavity ensures that the principal's expectation of future continuation values is weakly lower than his continuation value under the expected value of future beliefs. The second crucial feature - principal's best prediction of the next public belief given a current p is exactly p - also holds. Under private communication, the public belief process evolves according to a new transition kernel that still equals p on average, for every p. As such, the public belief process continues to converge almost surely to the same (induced) cascade sets as before. All the results from the previous section are valid.

6 Conclusion

People rely on the wisdom of the crowds to make decisions. Because they do, using information disclosure to induce or avoid herd behavior is the goal of many professionals and institutions. This study investigates the optimal ways to persuade crowds. Specifically, I consider an observational learning model and add a non-benevolent information designer who can commit to an information disclosure strategy, but cannot censor public information in society nor observe each agent's private information. The designer's problem then is basically when to be strategically vague - thus letting agents follow their own signals to some extent - and when to be strategically clear - thus triggering informational cascades.

This study has two main results. First, the features of agents' private information structure determine when it is optimal to persuade a single agent - single disclosure case - and when it is optimal to allow some social learning dynamics. For a well-known class of private belief distributions - the log-concave class, I give a characterization in terms of one of the tails of the unconditional private belief density (theorem 1). Some social learning is optimal if and only if private information unfavorable to the principal's most preferred action is sufficiently rare. With unbounded private beliefs, this possibility can never be too significant, so single disclosure is never optimal.

Second, social learning is less valuable to a more patient principal. In the limiting case - that is, as the designer's discount factor goes to one - the optimal policy has the same value as the single disclosure policy (Theorem 2). This means that whenever designer does not heavily discount current payoffs from persuasion, avoiding agents from learning through actions might be in his best interest.

An auxiliary result is worth mentioning. For bounded private beliefs, under any optimal policy, conditional on the state being high, there can be no herds toward the worst action for agents. Without an information intermediary, there is also a chance of society getting trapped in the bad herd. Conditional on state being low, complete learning occurs with positive probability. Again, this could not happen without intervention. Thus, the information policy from the selfish designer benefits society, as it eliminates informationally inefficient outcomes.

As an extension, I also prove that allowing the principal to censor past messages to current agents does not provide him with any additional benefit. This happens because agents know the information rule in every period, even though they might not be sure about the realization of past messages. As such, the public belief process is still a martingale and the principal's value function is still concave in those beliefs.

References

An, M. Y. (1997). Log-concave probability distributions: Theory and statistical testing. *Working Paper* (95-03).

An, M. Y. (1998). Logconcavity versus logconvexity: a complete characterization. *Journal of economic theory 80*(2), 350–369.

Au, P. H. (2015). Dynamic information disclosure. *The RAND Journal of Economics 46*(4), 791–823.

Aumann, R. J., M. Maschler, and R. E. Stearns (1995). *Repeated games with incomplete information.* MIT press.

Bagnoli, M. and T. Bergstrom (2005). Log-concave probability and its applications. *Economic theory 26*(2), 445–469.

Banerjee, A. V. (1992). A simple model of herd behavior. *The quarterly journal of economics 107*(3), 797–817.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1998). Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of economic perspectives 12*(3), 151–170.

Cao, H. H., B. Han, and D. Hirshleifer (2011). Taking the road less traveled by: Does conversation eradicate pernicious cascades? *Journal of Economic Theory 146*(4), 1418–1436.

Cule, M. and R. Samworth (2010). Theoretical properties of the log-concave maximum likelihood estimator of a multidimensional density. *Electronic Journal of Statistics 4*, 254–270.

Ely, J. C. (2017). Beeps. *American Economic Review 107*(1), 31–53.

Herrera, H. and J. Hörner (2012). A necesssary and sufficient condition for information cascades. *Working paper*.

Inostroza, N. and A. Pavan (2017). Persuasion in global games with application to stress testing. *Working paper*.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017). Persuasion of a privately informed receiver. *Econometrica 85*(6), 1949–1964.

Nikiforov, D. (2015). *On the belief manipulation and observational learning.* Ph. D. thesis, Citeseer.

Quah, J. K.-H. and B. Strulovici (2012). Aggregating the single crossing property. *Econometrica 80*(5), 2333–2348.

Renault, J., E. Solan, and N. Vieille (2017). Optimal dynamic information provision. *Games and Economic Behavior 104*, 329–349.

Rockafellar, R. T. (1970). *Convex analysis.* Princeton university press.

Roesler, A.-K. (2014). Mechanism design with endogenous information. *Working paper*.

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

Sgroi, D. (2002). Optimizing information in the herd: Guinea pigs, profits, and welfare. *Games and Economic Behavior 39*(1), 137–166.

Smith, L. and P. Sørensen (1996). Pathological outcomes of observational learning. *Working paper*.

Smith, L. and P. Sørensen (2000). Pathological outcomes of observational learning. *Econometrica 68*(2), 371–398.

Smith, L., P. Sørensen, and J. Tian (2021). Informational herding, optimal experimentation, and contrarianism. *the Review of Economic Studies*.

Stokey, N. L. (1989). *Recursive methods in economic dynamics.* Harvard University Press.

Williams, D. (1991). *Probability with martingales.* Cambridge university press.

Appendix A Technical details and omitted proofs

A model of crowds

This subsection reproduces key results about private beliefs and the public belief process in a standard observational learning model, for the sake of completeness. All claims are adaptations from results that have already appeared in the literature. Claims 1, 2 and 3 are taken from Smith and Sørensen (1996) and Rosenberg and Vieille (2019). Claim 4 is taken from Cao et al. (2011).

I have said that private beliefs come from the observation of a private signal, but I have remained silent about what those signals might be. Let me give now a detailed description of the private inference process and let me explain why it is sufficient to impose assumptions directly on the unconditional distribution of private beliefs. First, let me summarize all possible outcomes from this repeated interaction by the sample space $\Omega := \Theta \times (A \times S)^{\mathbb{N}}$. The S is a space of private signals. Agent t's set of public histories is A^{t-1}; the first agent's public history is the null set. A strategy profile for the agents and the common prior belief over the states generates a probability measure $\mathbb{P}$ over $\mathcal{F}$, the σ-algebra generated by Ω. The sample space admits the partition $\Omega^H := \{H\} \times (A \times S)^{\mathbb{N}}$ and $\Omega^L := \{L\} \times (A \times S)^{\mathbb{N}}$ with $\mathbb{P}(\Omega^H) = \mathbb{P}(\Omega^L) = 1/2$. I also refer to $\mathbb{P}^\theta$ as the conditional probability measure over $(\Omega, \mathcal{F})$ given θ.

Every agent observes the realization of a measurable function $\tilde{s}_t : (\Omega, \mathcal{F}) \to (S, \mathcal{S})$. The conditional law of $\tilde{s}_t$ is thus $\mathbb{F}^\theta_t = \mathbb{P}^\theta \circ \tilde{s}_t^{-1}$ for $\theta \in \{H, L\}$. The assumption of signals being conditionally i.i.d. means that $\mathbb{F}^\theta_t = \mathbb{F}^\theta$ for every $t \in \mathbb{N}$. To ensure that no private signal perfectly reveals the state, I impose $\mathbb{F}^H$ and $\mathbb{F}^L$ to be mutually absolutely continuous. This means that every subset of $\mathcal{S}$ has measure zero under $\mathbb{F}^H$ if and only if it has measure zero under $\mathbb{F}^L$. The Radon-Nikodym theorem ensures the existence of a non-negative measurable function ζ such that $\mathbb{F}^H = \zeta \mathbb{F}^L$. This function is almost surely unique, positive and finite. For every agent t, consider now the measurable function $\tilde{q}_t : (S, \mathcal{S}) \to ((0, 1], \mathcal{B})$ such that

$$\tilde{q}_t(s) := \frac{1}{1 + \zeta(s)},$$

where $\mathcal{B}$ is the Borel σ-algebra of the unit interval. Note that $\tilde{q}_t$ is the conditional probability of $\theta = L$ given the (σ-algebra generated by the) private signals. That is why I refer to $\tilde{q}_t$ as the private belief variable. Because $\{\tilde{s}_t\}_{t\in\mathbb{N}}$ is conditionally i.i.d., so it will be $\{\tilde{q}_t\}_{t\in\mathbb{N}}$. The associated conditional measures are $\mathbb{G}^\theta = \mathbb{F}^\theta \circ \tilde{q}^{-1}$. Because $\mathbb{F}^H$ and $\mathbb{F}^L$ are mutually absolutely continuous, $\mathbb{G}^H$ and $\mathbb{G}^L$ will also have this property. Therefore, there exists a non-negative measurable function η such that $\mathbb{G}^L = \eta \mathbb{G}^H$. Observe that

$$\mathbb{G}^L(B) = \int_{\tilde{q}^{-1}(B)} d\mathbb{F}^L = \int_{\tilde{q}^{-1}(B)} \frac{1}{\zeta} d\mathbb{F}^H = \int_B \left[\frac{q}{1-q}\right] d\mathbb{G}^H.$$

This means that the density η equals $q/(1-q)$ almost surely. In particular, it is true for the conditional cumulative distribution functions G^H and G^L. This is called the *no introspection condition* in Smith and Sørensen (1996). Now, consider only the assumptions that private signals are conditionally i.i.d and that the unconditional distribution over private beliefs G is absolutely continuous with density g. Then G^H and G^L will be mutually absolutely continuous with each other and will have densities g^H and g^L respectively. If it holds that $g^L/g^H = q/(1-q)$ almost surely, then the whole private information structure is determined by g. This is so because I can set $g^H(q) := 2(1-q)g(q)$; $g^L(q) := 2qg(q)$ and define conditionally i.i.d. distributions of private signals that generates G: just set $F^\theta = G^\theta$ for every $\theta \in \{H, L\}$.

Let me show that $g^L/g^H = q/(1-q)$ indeed. Set $\eta = g^L/g^H$. If an agent could see directly a private belief q, her inference about state L would be

$$\tilde{q}(q) = \frac{g^L(q)}{g^L(q) + g^H(q)} = \frac{\eta(q)}{1+\eta(q)}.$$

But $\tilde{q}(s) = \mathbb{E}[\mathbb{1}_{\tilde{\theta}=L}|s]$. It follows that $\tilde{q}(q) = \mathbb{E}[\mathbb{1}_{\tilde{\theta}=L}|q] = \mathbb{E}[\mathbb{E}[\mathbb{1}_{\tilde{\theta}=L}|s]|q] = \mathbb{E}[\tilde{q} = q] = q$. Thus,

$$\eta(q) = \frac{q}{1-q}.$$

Claim 1. *The difference $\alpha^L(p) - \alpha^H(p)$ is non-decreasing in $p \in [1/2, 1]$ and strictly increasing in $p \in (1/2, \bar{q})$. Likewise, it is non-increasing in $p \leq [0, 1/2]$ and strictly decreasing in $p \in (\underline{q}, 1/2)$. Moreover, $\alpha^H(p) > \alpha^L(p)$ for every $p \in (\underline{q}, \bar{q})$.*

Proof. Recall that it is possible to rewrite conditional densities in terms of g only: $g^H(q) = 2(1-q)g(q)$ and $g^L(q) = 2qg(q)$. Integrating by parts, I can rewrite $\alpha^H(p_t)$ and $\alpha^L(p_t)$ as

$$\alpha^H(p_t) = 2\left[(1-p_t)G(p_t) + \int_{\underline{q}}^{p_t} G(q)dq\right],$$
$$\alpha^L(p_t) = 2\left[p_t G(p_t) - \int_{\underline{q}}^{p_t} G(q)dq\right],$$

The difference between α^L and α^H is

$$\alpha^L(p_t) - \alpha^H(p_t) = 2\left[(2p_t - 1)G(p_t) - 2\int_{\underline{q}}^{p_t} G(q)dq\right].$$

Suppose $p_t \geq 1/2$. Take any $p'_t > p_t$. Because $G(q) \leq G(p'_t)$ for every $q \in (p_t, p'_t]$, it follows that

$$\begin{aligned}(\alpha^L(p'_t) - \alpha^H(p'_t)) - (\alpha^L(p_t) - \alpha^H(p_t)) &= 2\left[2(p'_t - p_t)G(p'_t) + (G(p'_t) - G(p_t))(2p_t - 1) - 2\int_{p_t}^{p'_t} G(q)dq\right], \\ &\geq [G(p'_t) - G(p_t)](2p_t - 1), \\ &\geq 0.\end{aligned}$$

This means that the difference $\alpha^L - \alpha^H$ is non-decreasing for beliefs above $1/2$. The difference is strict if $\bar{q} > p_t > 1/2$, because G is strictly increasing in its support (recall that G is continuous). Suppose now that $p_t \leq 1/2$. Take any $p'_t < p_t$. It follows that

$$\begin{aligned}(\alpha^L(p_t) - \alpha^H(p_t)) - (\alpha^L(p'_t) - \alpha^H(p'_t)) &= 2\left[2(p_t - p'_t)G(p_t) + (G(p'_t) - G(p_t))(1 - 2p_t) - 2\int_{p'_t}^{p_t} G(q)dq\right], \\ &\leq [G(p'_t) - G(p_t)](1 - 2p_t), \\ &\leq 0.\end{aligned}$$

This means that the difference $\alpha^L - \alpha^H$ is non-increasing for beliefs above 1/2. Again, the difference is strict if $\underline{q} < p_t < 1/2$. Let me now show that α^L stochastically dominates α^H. From the difference $\alpha^L - \alpha^H$, it is possible to see that this is certainly true for $p_t \leq 1/2$. Suppose now $p_t \geq 1/2$. Because $\bar{q} - 1/2 = \int_{\underline{q}}^{\bar{q}} G(q)dq = \int_{\underline{q}}^{p} G(q)dq + \int_{p}^{\bar{q}} G(q)dq$, another way of writing the difference is

$$\alpha^L(p_t) - \alpha^H(p_t) = 2\left[(2p_t - 1)G(p_t) + 1 - 2\bar{q} + 2\int_{p}^{\bar{q}} G(q)dq\right].$$

Because $G(q) \leq G(\bar{q}) = 1$, it follows that

$$\begin{aligned}\alpha^L(p) - \alpha^H(p) &\leq 2\left[(2p-1)G(p) + 1 - 2p\right],\\ &= 2(1-2p)(1-G(p)),\\ &\leq 0.\end{aligned}$$

Note that whenever $p \leq \underline{q}$, $\alpha^H(p) = \alpha^L(p) = 0$; whenever $p \geq \bar{q}$, $\alpha^H(p) = \alpha^L(p) = 1$. This proves that $\alpha^L(p) < \alpha^H(p)$ for every $p \in (\underline{q}, \bar{q})$. □

Claim 2. *For every $p \in (\underline{q}, \bar{q})$, the laws of motion $\varphi_h(p)$ and $\varphi_\ell(p)$ for public beliefs given the past observation of actions satisfy $\varphi_h(p) > \max\{1/2, p\}$ and $\varphi_\ell(p) < \min\{1/2, p\}$.*

Proof. That $\varphi_h(p) > p$ and $\varphi_\ell(p) < p$ for every $p \in (\underline{q}, \bar{q})$ follows from claim 1. Let me show that $\varphi_h(p) > 1/2 > \varphi_\ell(p)$ as well. Since $(g^L/g^H) = q/(1-q)$ and $q/(1-q)$ is a strictly increasing function, it follows that

$$\alpha^L(p) = 2\int_{\underline{q}}^{p} g^L(q)dq = \int_{\underline{q}}^{p}\left(\frac{q}{1-q}\right)g^H dq < \left(\frac{p}{1-p}\right)\int_{\underline{q}}^{p} g^H dq = \left(\frac{p}{1-p}\right)\alpha^H(p).$$

This implies that $\varphi_h(p) > 1/2$, because

$$\varphi_h(p) = \frac{\alpha^H(p)p}{\alpha^H(p)p + \alpha^L(p)(1-p)} = \frac{1}{1 + \frac{\alpha^L(p)}{\alpha^H(p)}\frac{1-p}{p}} > \frac{1}{2}.$$

Similarly,

$$1 - \alpha^L(p) = 2\int_{p}^{1} g^L(q)dq = \int_{p}^{1}\left(\frac{q}{1-q}\right)g^H dq > \left(\frac{p}{1-p}\right)\int_{p}^{1} g^H dq = \left(\frac{p}{1-p}\right)(1-\alpha^H(p)).$$

Implying that $\varphi_\ell(p) < 1/2$, because

$$\varphi_\ell(p) = \frac{(1-\alpha^H(p))p}{(1-\alpha^H(p))p + (1-\alpha^L(p))(1-p)} = \frac{1}{1 + \frac{1-\alpha^L(p)}{1-\alpha^H(p)}\frac{1-p}{p}} < \frac{1}{2}.$$

□

To proceed, let me formally argue that the public belief process is a martingale. Recall that each public belief is a random variable $\tilde{p}_t : \Omega \to [0,1]$ measurable with respect to $\mathcal{F}$. Let $\mathcal{A}_t$ be the sigma-algebra generated by A^{t-1}, for every t (at $t = 1$, agent 1 does not observe any history). Each $\tilde{p}_t$ is measurable with respect to $\mathcal{A}_t$ and $(\mathcal{A}_t)_{t\in\mathbb{N}}$ is an increasing family of sub-σ-algebras of $\mathcal{F}$. Thus, $(\tilde{p}_t)_{t\in\mathbb{N}}$ is adapted. Moreover, because it is a version of the conditional probability of Ω^H given events in $\mathcal{A}_t$ and $\mathcal{A}_t \subset \mathcal{A}_{t+1}$, it follows that $\mathbb{E}[\tilde{p}_{t+1}|\mathcal{A}_t] = \mathbb{E}[\mathbb{E}[\mathbb{1}_{\theta=H}|\mathcal{A}_{t+1}]\mathcal{A}_t] = \mathbb{E}[\mathbb{1}_{\theta=H}|\mathcal{A}_t] = \tilde{p}_t$ almost surely. The belief process is a martingale indeed.

Being a martingale, it must converge almost surely to a random variable $\tilde{p}_\infty$ (see for instance Williams, 1991, section 11.5). The proof here - an almost exact reproduction of theorem B.1 and B.2 in Smith and Sørensen (1996) - relies on $\alpha(\cdot)$ and $\varphi_a(\cdot)$ being continuous functions outside cascade sets, but this is not crucial, as it can be see in the proofs of the theorems in the referred paper.

Claim 3. *The limiting public belief $\tilde{p}_\infty$ has all points of its support belonging to cascade sets.*

Proof. First, I need to prove the following. If an open interval $I \subset [0,1]$ has the property that there exists a number $\varepsilon > 0$ such that, $\forall\ p \in I$, either (i) $\alpha(p) > \varepsilon$ and $|\varphi_h(p) - p| > \varepsilon$ or (i) $\alpha(p) < 1 - \varepsilon$ and $|\varphi_\ell(p) - p| > \varepsilon$, then I cannot contain any point in the support of $\tilde{p}_\infty$. Indeed, assume by way of contradiction that this is not the case. Consider any point $p^*_\infty \in I \cap \text{supp}(\tilde{p}_\infty)$ and define the set $I' := (p^*_\infty - \varepsilon/2, p^*_\infty + \varepsilon/2) \cap I$. For any p in I', either (i) $\alpha(p) > \varepsilon$ and $\varphi_h(p) \notin I'$ or (ii) $\alpha(p) < 1 - \varepsilon$ and $\varphi_\ell(p) \notin I'$. On the one hand, $p^*_\infty \in \text{supp}(\tilde{p}_\infty)$, so it must be that there is a positive probability that the event $\{\tilde{p}_t \in I'\}$ occurs for infinitely many t. On the other hand, conditional on the event $\{\tilde{p}_t \in I'\}$, the event $\{\tilde{p}_{t+1} \notin I'\}$ has probability at least ε. Thus, $\sum_{t\in\mathbb{N}} \mathbb{P}[\tilde{p}_{t+1} \notin I' | \tilde{p}_t \in I'] = \infty$. The (conditional) second Borel-Cantelli lemma (see for instance Williams, 1991, section 12.15) implies then that $\{\tilde{p}_{t+1} \notin I'\}$ happens infinitely often, conditional on $\{\tilde{p}_t \in I'\}$ infinitely often. But then probability of the event $\{\tilde{p}_t \in I'\}$ happening for infinitely many t is zero, a contradiction.

With the above claim, I can continue with the proof that the support of $\tilde{p}_\infty$ contains only points in $C_\ell \cup C_h$. Assume by way of contradiction that there exists some point in the support of $\tilde{p}_\infty$ - say, p^*_∞ - such that $p^*_\infty \notin C_\ell \cup C_h$. Then there exists $\varepsilon > 0$ s.t. either $\alpha(p^*_\infty) > \varepsilon$ and $|\varphi_h(p^*_\infty) - p^*_\infty| > \varepsilon$ or $\alpha(p^*_\infty) < 1 - \varepsilon$ and $|\varphi_\ell(p^*_\infty) - p^*_\infty| > \varepsilon$. Without loss, suppose the first case holds. Because $\alpha(\cdot)$ is continuous at p^*_∞, there exists an open neighborhood around p^*_∞ - call it I - such that $\alpha(p) > \varepsilon$ and $|\varphi_h(p) - p| > \varepsilon$ for every $p \in I$. But then I cannot contain any point in the support of $\tilde{p}_\infty$, a contradiction. □

To prove the last claim in this subsection, let $\{\lambda^{np}_t\}_{t\in\mathbb{N}}$ be a sequence of probability measures, each t giving the probability of the public belief process belonging to any event at t. Because $\{\tilde{p}_t\}_{t\in\mathbb{N}}$ converges almost surely to $\tilde{p}_\infty$, it must be the case that $\mathbb{E}_{\lambda^{np}_t}[f]$ converges almost surely to $\mathbb{E}_{\lambda^{np}_\infty}[f]$, for every bounded, continuous function f. Because $\text{supp}(\tilde{p}_\infty) \subseteq C_\ell \cup C_h$, it must also be the case that $\text{supp}(\lambda^{np}_\infty) \subseteq C_\ell \cup C_h$. The function α is continuous, so $\lim_{t\to\infty} \mathbb{E}_{\lambda^{np}_t}[\alpha] = \lambda^{np}_\infty(C_h)$. It remains to show that $\lim_{\delta\to 1} V^{np}_\delta = \lim_{t\to\infty} \mathbb{E}_{\lambda^{np}_t}[\alpha]$.

Claim 4. $\lim_{\delta\to 1} V^{np}_\delta = \lim_{t\to\infty} \mathbb{E}_{\lambda^{np}_t}[\alpha]$.

Proof. Let $V^* = \lim_{t\to\infty} \mathbb{E}_{\lambda^{np}_t}[\alpha]$. Because $\mathbb{E}_{\lambda^{np}_t}[\alpha] \to V^*$, for all $(\varepsilon/2) > 0$, there exists some $N \in \mathbb{N}$ such that for $t \geq N$, $|\mathbb{E}_{\lambda^{np}_t}[\alpha] - V^*| \leq \varepsilon/2$. This leads to

$$|V^{np}_\delta - V^*| = \left| \sum_{t\in\mathbb{N}} (1-\delta)\delta^{t-1}(\mathbb{E}_{\lambda^{np}_t}[\alpha] - V^*) \right|,$$
$$\leq \sum_{t<N} (1-\delta)\delta^{t-1} |\mathbb{E}_{\lambda^{np}_t}[\alpha] - V^*| + \delta^{N-1}\varepsilon/2.$$

Now consider $\bar{V} := \sum_{t<N} |\mathbb{E}_{\lambda^{np}_t}[\alpha] - V^*|$ and $\bar{\delta} := 1 - \varepsilon/(2\bar{V})$. Then, for any $1 > \delta' > \bar{\delta}$,

$$\sum_{t<N}(1-\delta')\delta'^{t-1}|\mathbb{E}_{\lambda_t^{np}}[\alpha]-V^*|+\delta'^{N-1}\varepsilon/2\leq\varepsilon/2+\varepsilon/2=\varepsilon.$$

As the choice of ε was arbitrary, this means that $\lim_{\delta\to 1}V_\delta^{np}=\lim_{t\to\infty}\mathbb{E}_{\lambda_t^{np}}[\alpha]$. □

Persuading crowds

In this subsection, the set of all possible outcomes of the infinite interaction is $\Omega=\Theta\times(A\times S\times M)^{\mathbb{N}}$. To simplify notation, I set $X:=A\times M$. A strategy profile for the agents, the common prior belief over the states, the prior information structure and the information policy generate a probability measure over $\mathcal{F}$, the σ-algebra generated by Ω.

Lemma 1. *Consider any stochastic processes $\{\tilde{\rho}_t\}_{t\in\mathbb{N}}$ and $\{\tilde{p}_t\}_{t\in\mathbb{N}}$ - with initial prior belief p_1 given - such that (i) for every realization of a public belief p_t, the law of the induced belief $\tilde{\rho}_t$ conditional on p_t equals p_t in expectation; (ii) for every realization of an induced belief ρ_t, there exists some action a taken with positive probability such that next period's public belief is $p_{t+1}=\varphi_a(\rho_t)$. Therese processes can be generated by an information policy for which the message space is the belief space $[0,1]$ and the information rules depend only on the current public belief.*

Proof. Consider any stochastic processes $\{\tilde{\rho}_t\}_{t\in\mathbb{N}}$ and $\{\tilde{p}_t\}_{t\in\mathbb{N}}$ s.t. the expected value of the conditional law of $\tilde{\rho}_t$ given a realization p_t equals p_t, for each t. Call this conditional law $\tau(\cdot;p_t)$. Let the message space be $M=[0,1]$ and let the associated σ-algebra be the Borel σ-algebra $\mathcal{B}$ of M. Consider the $\mathcal{M}$-measurable mappings:

$$\kappa^H(m,p_t):=\begin{cases}\frac{m}{p_t} & \text{if } p_t\in(0,1),\\ 1 & \text{if } p_t\in\{0,1\};\end{cases}\qquad \kappa^L(m,p_t):=\begin{cases}\frac{1-m}{1-p_t} & \text{if } p_t\in(0,1),\\ 1 & \text{if } p_t\in\{0,1\}.\end{cases}$$

Consider as well the following set functions on $\mathcal{M}$,

$$\mu^H(B;p_t):=\int_{m\in B}\kappa^H(m,p_t)\tau(dm;p_t),\qquad \mu^L(B;p_t):=\int_{m\in B}\kappa^L(m,p_t)\tau(dm;p_t).$$

I claim that they are probability measures, given p_t. Suppose that $p_t\in(0,1)$ (otherwise this is trivially true). That they are non-negative is immediate. Moreover, $\mu^\theta(M;p_t)=1$. Finally, they are σ-additive. Indeed, for any sequence $(B_n)_{n\in\mathbb{N}}$ of pairwise disjoint subsets of $\mathcal{M}$ with $B=\cup_{n\in\mathbb{N}}B_n$,

$$\begin{aligned}\mu^\theta(B;p_t)&=\int_{m\in\cup_{n\in\mathbb{N}}B_n}\kappa^\theta(m,p_t)\tau(dm;p_t),\\ &=\int_M\left[\sum_{n\in\mathbb{N}}\kappa^\theta(m,p_t)\cdot\mathbb{1}_{B_n}(m)\right]\tau(dm;p_t),\\ &=\sum_{n\in\mathbb{N}}\left[\int_M\kappa^\theta(m,p_t)\cdot\mathbb{1}_{B_n}(m)\tau(dm;p_t)\right],\\ &=\sum_{n\in\mathbb{N}}\mu^\theta(B_n;p_t).\end{aligned}$$

The value $\mathbb{1}_{B_n}(m)$ above represents an indicator function, equal to one whenever $m \in B_n$ and zero otherwise. Observe that, from the point of agent t that does not know θ but observes the information policy and the realization $p_t \in (0,1)$ (again, if $p_t \in \{0,1\}$ the proof is trivial), the probability of any $B \in \mathcal{M}$ is given by

$$\begin{aligned}\mu(B;p_t) &= p_t\mu^H(B;p_t) + (1-p_t)\mu^L(B;p_t),\\ &= p_t\int_B \frac{m}{p_t}\tau(dm;p_t) + (1-p_t)\int_B \frac{1-m}{1-p_t}\tau(dm;p_t),\\ &= \tau(B;p_t).\end{aligned}$$

The information policy generates the same conditional probability measure over induced posteriors. Let me now show that under this policy, the realization of posterior beliefs coincide with the posterior beliefs $\rho_t \in \text{supp}(\tau(p_t))$. To do so, suppose first that $p_t \in (0,1)$. If the principal sends ρ_t to the agent, her posterior belief is

$$\tilde{\rho}(\rho_t;p_t) := \frac{p_t\kappa^H(\rho_t;p_t)}{p_t\kappa^H(\rho_t;p_t) + (1-p_t)\kappa^L(\rho_t;p_t)} = \rho_t.$$

If $p_t = 0$, then $\mathbb{E}[\tilde{\rho}_t|p_t] = 0$ implies that the only possible $\tilde{\rho}_t$ is 0. Then trivially $\tilde{\rho}(B;p_t)$ induces 0 for any message B that has positive probability. Similar analysis holds for $p_t = 1$. Note as well that under this information policy the expected value of induced beliefs conditional on the realization of belief p_t equals p_t:

$$\mathbb{E}[\tilde{\rho}_t|p_t] = \int \rho_t\mu(d\rho_t;p_t) = \int \rho_t\tau(d\rho_t;p_t) = p_t.$$

□

In what follows, it will be convenient to review some results about the concave closure of a bounded function $f : X \to Y$, with $X \subseteq \mathbb{R}$ convex and $Y \subseteq \mathbb{R}$. This is given by

$$\text{cav}[f](x) = \sup\{y : (x,y) \in \text{co}(\text{hyp}(f))\},$$

where co(hyp(f)) is the convex hull of the hypograph of f. The concave closure of a bounded function is concave. Indeed, let hyp(cav[f]) be the hypograph of cav[f]. Take any $(x,t),(x',t')$ in it. There exists probability weights τ and τ', both over X, such that $\mathbb{E}_\tau[\tilde{x}] = x$ and $\mathbb{E}_{\tau'}[\tilde{x}] = x'$ as well as $\mathbb{E}_\tau[f(\tilde{x})] = t$ and $\mathbb{E}_{\tau'}[f(\tilde{x})] = t' \leq f(x')$. Consider now an arbitrary $\lambda \in [0,1]$. Define $x'' := \lambda x + (1-\lambda)x'$ as well as $\tau'' := \lambda\tau + (1-\lambda)\tau'$. There exists probability weights such that $\mathbb{E}_{\tau''}[\tilde{x}] = x''$, $\mathbb{E}_{\tau''}[f(\tilde{x})] = \lambda t + (1-\lambda)t' := t'' \leq \text{cav}[f](x'')$. That implies hyp(cav[$f$]) is convex and cav[$f$] is concave.

It will be convenient as well to recast the problem in terms of a Markov chain over the belief space. Define a transition probability $P : [0,1] \times \mathcal{B} \to [0,1]$ such that for every $p \in [0,1]$ and every $B \in \mathcal{B}$,

$$P(p,B) = \mathbb{1}\{\varphi_h(p) \in B\}\alpha(p) + \mathbb{1}\{\varphi_\ell(p) \in B\}(1-\alpha(p)).$$

Note that, for every p, the expected value of $P(p)$ is exactly p:

$$\int p' P(p, dp') = \alpha(p)\varphi_h(p) + (1 - \alpha(p))\varphi_\ell(p) = p$$

Associated with it, there is a transformation mapping the space of bounded functions f on the belief space to the same space, defined as below. This is the expected value of a function f given that the current belief is p.

$$\int f(p')P(p, dp') = \alpha(p)f(\varphi_h(p)) + (1 - \alpha(p))f(\varphi_\ell(p)).$$

Associated with this operator, there is an adjoint operator P^* mapping the space of probability measures ν over the belief space to this same space, defined as below. This is the probability of next belief belonging to B if the current belief is drawn according to ν.

$$(P^*\nu)(B) := \int P(p, B)\nu(dp).$$

One can show[20] that P and P^* are connected trough the following relation:

$$\int (Pf)(p)\nu(dp) = \int f(p')(P^*\nu)(dp').$$

Using the above notation, I can define another transformation T from the space of bounded functions V to itself. This transformation is the concave closure of the function $(1 - \alpha)(p) + \delta(PV)(p)$. The transformation is given below. From it, a series of claims follow.

$$(TV)(p) := \sup_{\tau \in \mathcal{S}(p)} \mathbb{E}_\tau \left[(1 - \delta)\alpha(\tilde{\rho}) + \delta(PV)(\tilde{\rho})\right] = \sup_{\tau \in \mathcal{S}(p)} \left\{(1 - \delta)\mathbb{E}_\beta[\alpha(\tilde{\rho})] + \delta \int V(p')(P^*_\tau)(dp')\right\},$$

Claim 5. *For every p and every bounded, continuous function V, there exists a solution $\tau \in \mathcal{S}(p)$ to the problem:*

$$\sup_{\tau \in \mathcal{S}(p)} \mathbb{E}_\tau \left[(1 - \delta)\alpha(\tilde{\rho}) + \delta(PV)(\tilde{\rho})\right]$$

Proof. The assumption of an absolutely continuous unconditional distribution of private beliefs imply that both α and PV will be bounded and continuous, if V is bounded and continuous. In particular, the expression in brackets will be upper semi-continuous, so its hypograph is convex. Therefore, any element on the convex hull of the hypograph of will be attainable, and I can interchange the sup by the max.

□

[20] See for instance Stokey (1989), theorem 8.3.

Claim 6. *For every every bounded function V, the transformation function TV is concave in beliefs.*

Proof. TV is the concave closure of a bounded function. From previous discussion, the concave closure of a bounded function is concave. $\square$

Claim 7. *The transformation T is a contraction.*

Proof. First note that $(TV)(p)$ is equivalent to $\text{cav}[(1-\delta)\alpha + \delta(PV)](p)$. From Blackwell sufficient conditions, to show that the operator is a contraction, it suffices to show that it satisfies continuity[21] and discounting[22]. Continuity follows from $(PV') \geq PV''$ for every $V' \geq V''$ and cav being itself a operator that satisfies continuity. Discounting follows from $(Pf+d)(p) = (Pf)(p)+d$ and $\text{cav}[f+d](p) = \text{cav}[f]+d$. Therefore, $(Tf+d)(p) = (Tf)(p) + \delta d$. $\square$

Claim 8. *The optimal value function $V_\delta^{op}(p)$ is continuous in beliefs.*

Proof. The transformation T maps the space of bounded functions to itself. Because it is a contraction, it suffices to observe that for every continuous function, the image of the operator will be continuous as well. $\square$

Claim 9. *For every p, any optimal policy with associated optimal probability measure over posteriors at p places positive probability on at most two induced beliefs ρ^-, ρ^+ s.t. $\rho^- \leq p \leq \rho^+$.*

Proof. This is a straightforward application of Carathéodory's theorem on any point of the convex hull of graph of $Z_\delta^{op} : [0,1] \to \mathbb{R}_+$ with $Z_\delta^{op}(p) := (1-\delta)\alpha(p) + \delta(PV_\delta^{op})(p)$. See for instance Rockafellar, 1970, corollary, 17.1.5. $\square$

Valuable social learning

Lemma 2. *If the private belief density g is log-concave, then α is convex-concave on $(\underline{q}, \bar{q})$.*

Proof. The function $c_1(q) := -(\ln q(1-q))'$ satisfies the single-crossing property. Likewise, if the density g is log-concave, then $c_2(q) = -(\ln g(q))'$ satisfies it as well: the log-concavity implies that c_2 monotonically increases in $(\underline{q}, \bar{q})$. Following Quah and Strulovici (2012), say that two functions f and $\hat{f}$ satisfy signed-ratio monotonicity if (i) at any $r' : \hat{f}(r') < 0$ and $f(r') > 0$, $(-\hat{f}(r')/f(r')) \geq (-\hat{f}(r'')/f(r''))$ whenever $r'' > r'$; (ii) at any $r' : f(r') < 0$ and $\hat{f}(r') > 0$, $(-f(r')/\hat{f}(r')) \geq (-f(r'')/\hat{f}(r''))$ whenever $r'' > r'$. Let me show that c_1 and c_2 satisfy the signed-ratio monotonicity.

Pick any $q' : c_2(q') < 0$ and $c_1(q') > 0$. As remarked, c_2 is monotonically increasing because g is log-concave, so $-c_2(q') \geq -c_2(q'')$ whenever $q'' > q'$. Likewise, because the function c_1 is increasing, $1/c_1(q') \geq 1/c_1(q'')$ whenever $q'' > q'$. Therefore, $(-c_2(q')/c_1(q')) \geq (-c_2(q'')/c_1(q''))$ whenever $q'' > q'$, as required. Now pick any $q' : c_1(q) < 0$ and $c_2(q') > 0$. Because c_1 is increasing, $-c_1(q') \geq c_1(q'')$ whenever $q'' > q'$. Similarly, because c_2 is decreasing, $(1/c_2(q')) \geq (1/c_2(q'')$ whenever $q'' > q'$. Therefore, $(-c_1(q')/c_2(q')) \geq (-c_1(q'')/c_2(q''))$ whenever $q'' > q'$, as required.

Because those functions satisfy the signed-ratio monotonicity, I can apply proposition 1 from Quah and Strulovici (2012) to conclude that $-\alpha''$ satisfies the single-crossing property as well. That means there exists a value $m \in (\underline{q}, \bar{q})$ such that $\alpha(p)$ is convex for $p < m$ and concave for $p > m$. $\square$

[21] That is, for any V', V'' in the space of bounded functions and s.t. $V' \leq V''$, $(TV') \leq (TV'')$.

[22] That is, there exists a discount factor $\gamma \in (0,1)$ such that $(TV+d)(p) \leq (TV)(p) + \gamma d$ for every $d \geq 0$.

The role of patience

Lemma 3. *Let π be an optimal policy. The value of the optimal value function must converge to the stationary value of the public belief process hitting C_h under π. Precisely,*

$$\lim_{\delta\to 1} V_\delta^{op} = \lim_{t\to\infty} \mathbb{E}_{\hat{\lambda}_t^\pi}[\alpha] = \lambda_\infty^\pi(C_h).$$

Proof. Because informative communication eventually stops, $\hat{\lambda}_t$ converges to λ_t as t goes to infinity. Claim 3 then implies

$$\lim_{\delta\to 1} V_\delta^{op} = \lim_{t\to\infty} \mathbb{E}_{\lambda_t^\pi}[\alpha].$$

Because the public belief process converges almost surely to the new cascade sets, the above limiting expected probability must equal

$$\lim_{t\to\infty} \mathbb{E}_{\lambda_t^\pi}[\alpha] = \lambda_\infty^\pi(C_h).$$

□

Private communication

Lemma 4. *With private communication, the function V_δ^{op} is concave in $(0,1)$.*

Proof. Because we have a contraction algorithm, it suffices to show that equation 10 is concave for any function V concave. To do so, pick any belief $p \in (0,1)$, any two interior beliefs $p' < p''$ and any value $\xi \in (0,1)$ such that $p = \xi p'' + (1-\xi)p'$. Consider $\tau_\xi := \xi\tau'' + (1-\xi)\tau'$ where τ'' (τ') is the Bayes plausible distribution solving equation 10 at p'' (p') for V. Moreover, consider $\tau_\xi^H(B) = \int_B (\rho/p)\tau_\xi(d\rho)$ if state is H as well as $\tau_\xi^L = \int_B[(1-\rho)/(1-p)]\tau_\xi(d\rho)$ if state is L, for any $B \subseteq [0,1]$. This splitting satisfies Bayes plausibility and $\tau_\xi = p\tau_\xi^H + (1-p)\tau_\xi^L$. Observe that under τ_ξ, the laws of motion as in equation 9 satisfy

$$\begin{aligned}
\hat{\varphi}_h(p,\tau_\xi) &= \frac{\hat{\alpha}^H(p,\tau_\xi)p}{\hat{\alpha}(p,\tau_\xi)}, \\
&= \xi\left[\frac{\hat{\alpha}^H(p'',\tau'')}{\hat{\alpha}(p,\tau_\xi)}\right]p'' + (1-\xi)\left[\frac{\hat{\alpha}^H(p',\tau')}{\hat{\alpha}(p,\tau_\xi)}\right]p', \\
&= \xi\left[\frac{\hat{\alpha}(p'',\tau'')}{\hat{\alpha}(p,\tau_\xi)}\right]\hat{\varphi}_h(p'',\tau'') + (1-\xi)\left[\frac{\hat{\alpha}(p',\tau')}{\hat{\alpha}(p,\tau_\xi)}\right]\hat{\varphi}_h(p',\tau');
\end{aligned}$$

$$\begin{aligned}
\hat{\varphi}_\ell(p,\tau_\xi) &:= \left[\frac{1-\hat{\alpha}^H(p,\tau_\xi)}{1-\hat{\alpha}(p,\tau_\xi)}\right]p, \\
&= \xi\left[\frac{1-\hat{\alpha}^H(p'',\tau'')}{1-\hat{\alpha}(p,\tau_\xi)}\right]p'' + (1-\xi)\left[\frac{1-\hat{\alpha}^H(p',\tau')}{1-\hat{\alpha}(p,\tau_\xi)}\right]p', \\
&= \xi\left[\frac{1-\hat{\alpha}(p'',\tau'')}{1-\hat{\alpha}(p,\tau_\xi)}\right]\hat{\varphi}_\ell(p'',\tau'') + (1-\xi)\left[\frac{1-\hat{\alpha}(p',\tau')}{1-\hat{\alpha}(p,\tau_\xi)}\right]\hat{\varphi}_\ell(p',\tau').
\end{aligned}$$

Because V is concave, it follows that

$$
\begin{aligned}
\hat{\alpha}(p,\tau_\xi)V(\hat{\varphi}_h(p,\tau_\xi)) &\geq \xi\alpha(p'',\tau'')V(\hat{\varphi}_h(p'',\tau'')) + (1-\xi)\hat{\alpha}(p',\tau')V(\hat{\varphi}_h(p',\tau')),\\
(1-\hat{\alpha}(p,\tau_\xi))V(\hat{\varphi}_\ell(p,\tau_\xi)) &\geq \xi(1-\hat{\alpha}(p'',\tau''))V(\hat{\varphi}_\ell(p'',\tau'')) + (1-\xi)(1-\hat{\alpha}(p',\tau'))V(\hat{\varphi}_\ell(p',\tau')).
\end{aligned}
$$

Combining the above results with the fact that $\hat{\alpha}(p,\tau_\xi) = \hat{\alpha}(p'',\tau'')\xi + \hat{\alpha}(p',\tau')(1-\xi)$, we get

$$
\begin{aligned}
&\max_{\tau\in\mathcal{S}(p)}\left[(1-\delta)\hat{\alpha}(p,\tau) + \delta\Big(\hat{\alpha}(p,\tau)V(\hat{\varphi}_h(p,\tau)) + (1-\hat{\alpha}(p,\tau))V(\hat{\varphi}_\ell(p,\tau))\Big)\right]\\
&\quad\geq (1-\delta)\hat{\alpha}(p,\tau_\xi) + \delta\Big(\hat{\alpha}(p,\tau_\xi)V(\hat{\varphi}_h(p,\tau_\xi)) + (1-\hat{\alpha}(p,\tau_\xi))V(\hat{\varphi}_\ell(p,\tau_\xi))\Big),\\
&\quad\geq \xi\left[(1-\delta)\hat{\alpha}(p'',\tau'') + \delta\Big(\hat{\alpha}(p'',\tau'')V(\varphi_h(p'')) + (1-\alpha(p''))V(\varphi_\ell(p''))\Big)\right]+\\
&\quad+(1-\xi)\left[(1-\delta)\alpha(p') + \delta\Big(\alpha(p')V(\varphi_h(p')) + (1-\alpha(p'))V(\varphi_\ell(p'))\Big)\right],\\
&\quad= \xi\max_{\tau\in\mathcal{S}(p'')}\left[(1-\delta)\hat{\alpha}(p'',\tau) + \delta\Big(\hat{\alpha}(p'',\tau)V(\hat{\varphi}_h(p'',\tau)) + (1-\hat{\alpha}(p'',\tau))V(\hat{\varphi}_\ell(p'',\tau))\Big)\right]+\\
&\quad+(1-\xi)\max_{\tau\in\mathcal{S}(p')}\left[(1-\delta)\hat{\alpha}(p',\tau) + \delta\Big(\hat{\alpha}(p',\tau)V(\hat{\varphi}_h(p',\tau)) + (1-\hat{\alpha}(p',\tau))V(\hat{\varphi}_\ell(p',\tau))\Big)\right].
\end{aligned}
$$

□

Appendix B Calculations for the examples

Illustrative example

Recall that the private signal space is $S = \{\underline{s}, \bar{s}\}$ and the probability distributions are $f^H(\bar{s}) = f^L(\underline{s}) = \sigma$, for $\sigma \in (1/2, 1)$. Therefore, the belief space is $\{1-\sigma, \sigma\}$ with unconditional prob. $g(1-\sigma) = g(\sigma) = 1/2$. The cascade sets are $C_\ell = [0, 1-\sigma)$ and $C_h = [\sigma, 1]$. The conditional and unconditional probabilities of action h (investment) given p are

$$\alpha^H(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ \sigma & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases} \qquad \alpha^L(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ (1-\sigma) & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases}$$

$$\alpha(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ p\sigma + (1-p)(1-\sigma) & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases}$$

The system moves to another public belief according to the transition functions

$$\varphi_h(p) := \begin{cases} p & \text{if } p \in C_\ell, \\ \frac{\sigma p}{p\sigma + (1-p)(1-\sigma)} & \text{if } p \notin C_\ell \cup C_h. \end{cases} \qquad \varphi_\ell(p) := \begin{cases} p & \text{if } p \in C_h, \\ \frac{(1-\sigma)p}{p(1-\sigma) + (1-p)\sigma} & \text{if } p \notin C_\ell \cup C_h. \end{cases}$$

Let me compute the probability measures $(\lambda_t^{np})_{t\in\mathbb{N}}$ over public beliefs in each period in this example. Recall that $P(p, B)$ refers to the transition kernel from p to a public belief within B. At $t = 1$, $\lambda_1^{np}(1/2) = 1$. At $t = 2$, there are two possible public beliefs $1-\sigma$ and σ. Their probabilities are

$$\lambda_2^{np}(1-\sigma) = P(1/2, 1-\sigma) = 1 - \alpha(1/2) = 1/2,$$
$$\lambda_2^{np}(\sigma) = P(1/2, \sigma) = \alpha(1/2) = 1/2.$$

At $t = 3$, there are three possible public beliefs: $\varphi_\ell(1-\sigma)$, $1/2$ and σ, because $\varphi_h(1-\sigma) = 1/2$. The probabilities over beliefs are

$$\lambda_3^{np}(\varphi_\ell(1-\sigma)) = P(1-\sigma, \varphi_\ell(1-\sigma))\lambda_2^{np}(1-\sigma) = (1/2)(1-\alpha(1-\sigma)) = (1/2)[(1-\sigma)^2 + \sigma^2],$$
$$\lambda_3^{np}(1/2) = P(1-\sigma, 1/2)\lambda_2^{np}(1-\sigma) = (1/2)\alpha(1-\sigma) = \sigma(1-\sigma),$$
$$\lambda_3^{np}(\sigma) = P(\sigma, \sigma)\lambda_2^{np}(\sigma) = 1/2.$$

At $t = 4$, there are three possible beliefs : $\varphi_\ell(1-\sigma)$, $1-\sigma$ and σ with probabilities

$$\lambda_4^{np}(\varphi_\ell(1-\sigma)) = \lambda_3^{np}(\varphi_\ell(1-\sigma)) = (1/2)[(1-\sigma)^2 + \sigma^2],$$
$$\lambda_4^{np}(1-\sigma) = P(1/2, 1-\sigma)\lambda_3^{np}(1/2) = (1/2)\sigma(1-\sigma),$$
$$\lambda_4^{np}(\sigma) = P(1/2, \sigma)\lambda_3^{np}(1/2) + \lambda_3^{np}(\sigma) = 1/2[1 + \sigma(1-\sigma)].$$

At $t=5$, there are three possible beliefs: $\varphi_\ell(1-\sigma)$, $1/2$ and σ with probabilities

$$\begin{aligned}
&\lambda_5^{np}(\varphi_\ell(1-\sigma)) = P(1-\sigma,\varphi_\ell(1-\sigma))\lambda_4^{np}(1-\sigma)+\lambda_4^{np}(\varphi_\ell(1-\sigma)) = (1/2)[(1-\sigma)^2+\sigma^2](1+\sigma(1-\sigma)),\\
&\lambda_5^{np}(1/2) = P(1-\sigma,1/2)\lambda_4^{np}(1-\sigma) = \sigma^2(1-\sigma)^2,\\
&\lambda_5^{np}(\sigma) = \lambda_4^{np}(\sigma) = 1/2[1+\sigma(1-\sigma)].
\end{aligned}$$

By now a pattern is clear. For $t>2$ even, there are three possible public beliefs: $\varphi_\ell(1-\sigma)$, $1-\sigma$ and σ with probabilities

$$\begin{aligned}
&\lambda_t^{np}(\varphi_\ell(1-\sigma)) = \lambda_{t-1}^{np}(\varphi_\ell(1-\sigma)) = (1/2)[(1-\sigma)^2+\sigma^2]\sum_{\tau=0}^{\frac{t-2}{2}-1}\sigma^\tau(1-\tau)^\tau = \frac{1}{2}\left[\frac{(1-\sigma)^2+\sigma^2}{(1-\sigma)^2+\sigma}\right](1-\sigma^{\frac{t-2}{2}}(1-\sigma)^{\frac{t-2}{2}}),\\
&\lambda_t^{np}(1-\sigma) = P(1/2,1-\sigma)\lambda_{t-1}^{np}(1/2) = (1/2)\sigma^{\frac{t-2}{2}}(1-\sigma)^{\frac{t-2}{2}},\\
&\lambda_t^{np}(\sigma) = P(1/2,\sigma)\lambda_{t-1}^{np}(1/2)+\lambda_{t-1}^{np}(\sigma) = (1/2)\sum_{\tau=0}^{\frac{t-2}{2}}\sigma^\tau(1-\sigma)^\tau = \frac{1}{2}\left[\frac{1-\sigma^{\frac{t}{2}}(1-\sigma)^{\frac{t}{2}}}{(1-\sigma)^2+\sigma}\right].
\end{aligned}$$

For $t>1$ odd, there are three possible beliefs: $\varphi_\ell(1-\sigma)$, $1/2$ and σ with probabilities

$$\begin{aligned}
&\lambda_t^{np}(\varphi_\ell(1-\sigma)) = K(1-\sigma,\varphi_\ell(1-\sigma))\lambda_{t-1}^{np}(1-\sigma)+\lambda_{t-1}^{np}(\varphi_\ell(1-\sigma)) = \frac{1}{2}\left[\frac{(1-\sigma)^2+\sigma^2}{(1-\sigma)^2+\sigma}\right](1-\sigma^{\frac{t-1}{2}}(1-\sigma)^{\frac{t-1}{2}}),\\
&\lambda_t^{np}(1/2) = K(1-\sigma,1/2)\lambda_{t-1}^{np}(1-\sigma) = \sigma^{\frac{t-1}{2}}(1-\sigma)^{\frac{t-1}{2}},\\
&\lambda_t^{np}(\sigma) = \lambda_{t-1}^{np}(\sigma) = (1/2)\sum_{\tau=0}^{\frac{t-1}{2}-1}\sigma^\tau(1-\sigma)^\tau = \frac{1}{2}\left[\frac{1-\sigma^{\frac{t-1}{2}}(1-\sigma)^{\frac{t-1}{2}}}{(1-\sigma)^2+\sigma}\right].
\end{aligned}$$

The probabilities $\lambda_t^{np}(\sigma)$ and $\lambda_t^{np}(\varphi_\ell(1-\sigma))$ for each period t are represented in figure 1(b) for $\sigma=.8$, together with the limiting probability measures (red and blue lines). Figure 1(a) represents the possible interim beliefs in each period together with the values $\alpha(p_t)$ for each p_t. Let me compute principal's average discounted payoff without any information policy. Let $\lambda_\delta^{np}(p') = \sum_{t\in\mathbb{N}}(1-\delta)\delta^{t-1}\lambda_t(p')$, for $p'\in\{\varphi_\ell(1-\sigma),1/2,\sigma\}$. The value V_δ^{np} satisfies

$$V_\delta^{np} = \alpha(\varphi_\ell(1-\delta))\lambda_\delta^{np}(\varphi_\ell(1-\sigma))+\alpha(1/2)\lambda_\delta^{np}(1/2)+\alpha(\sigma)\lambda_\delta^{np}(\sigma).$$

Note that $\lim_{\delta\to1}V_\delta^{np} = \lim_{t\to\infty}\mathbb{E}_{\lambda_t^{np}}[\alpha] = \lambda_\infty^{np}(C_h) = (1/2)/[(1-\sigma)^2+\sigma]$. Let me now compute the value of greedy policy V_δ^{gp}. Suppose first that $1/2<\sigma\le1/\sqrt{2}$. Then whenever $p<\sigma$, principal splits posteriors between 0 and σ with probabilities $1-(p/\sigma)$ and p/σ respectively; otherwise, he does not disclose any additional information. Suppose now $1>\sigma>1/\sqrt{2}$. Whenever $p\in[0,1-\sigma)$, principal splits posterior between 0 and $1-\sigma$ and places weight $p/(1-\sigma)$ on $1-\sigma$. Whenever $p\in[1-\sigma,\sigma)$, principal splits posterior between $1-\sigma$ and σ and places weight $(p-(1-\sigma))/(2\sigma-1)$ on σ. Therefore, the concave closure of α is

$$\text{cav}[\alpha](p) = \begin{cases} \frac{p}{\sigma} & \text{if } p \notin C_h, \\ 1 & \text{if } p \in C_h. \end{cases} \qquad \text{for } \frac{1}{2} < \sigma \leq \frac{1}{\sqrt{2}};$$

$$\text{cav}[\alpha](p) = \begin{cases} 2\sigma p & \text{if } p \in C_\ell, \\ \left[\frac{(1-\sigma)^2+\sigma^2}{2\sigma-1}\right]p + \left[\frac{2\sigma^2-1}{2\sigma-1}\right](1-\sigma) & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases} \qquad \text{for } \frac{1}{\sqrt{2}} < \sigma < 1.$$

If $1/2 < \sigma \leq 1/\sqrt{2}$, the greedy policy dictates that the principal should induce beliefs on the extreme of the cascade sets for every initial belief $p_1 \notin C_\ell \cup C_h$ and he should not say anything for $p_1 \in C_h$. Thus, the value of a greedy policy and the value of a one-shot concavification coincide for every initial prior: $V_\delta^{gp}(p) = \text{cav}[\alpha](p)$. As this is actually the upper bound of every optimal policy, the greedy strategy reaches the optimal value.

If $1/\sqrt{2} < \sigma < 1$, it is not immediate to observe that the greedy policy is optimal, for every initial prior belief p_1. I prove this is the case in the next proposition. Letting $p_1 = 1/2$ leads to the value function given in proposition 1 in the the persuading crowds section.

Proposition 1. *In the illustrative example, the value of an optimal policy for $\sigma > \frac{1}{\sqrt{2}}$ is*

$$V_\delta(p) = \begin{cases} p\left(\frac{\sigma(2-\delta)}{1-\delta+\delta\sigma^2}\right) & \text{if } p \in C_\ell, \\ p\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) + (1-\sigma)\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases}$$

This value function is achieved through a greedy policy, that is, a policy that induces posteriors beliefs to generate $cav[\alpha](p)$ at every public belief p. This means that whenever $p < 1-\sigma$, principal induces posteriors 0 and $1-\sigma$ and whenever $p \in (1-\sigma, \sigma)$, principal induces posteriors $1-\sigma$ and σ. For beliefs $p \geq \sigma$, principal does not disclose any additional information.

Proof. First note that this value function is concave. Second, I need to show that the the greedy strategy actually leads to V_δ or $\mathbb{E}_{\tau^{gp}(p)}[Z_\delta] = V_\delta(p)$ for every p, where Z_δ is defined below.

$$Z_\delta(p) = (1-\delta)\alpha(p) + \delta\Big[\alpha(p)V_\delta(\varphi_h(p)) + (1-\alpha(p))V_\delta(\varphi_\ell(p))\Big].$$

To do so, let me define two compositions of the value function:

$$V_\delta(\varphi_\ell(p)) = \begin{cases} p\left(\frac{\sigma(2-\delta)}{1-\delta+\delta\sigma^2}\right) & \text{if } p \in C_\ell, \\ \varphi_\ell(p)\left(\frac{\sigma(2-\delta)}{1-\delta+\delta\sigma^2}\right) & \text{if } p \in [1-\sigma, 1/2), \\ \varphi_\ell(p)\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) + (1-\sigma)\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) & \text{if } p \in [1/2, \sigma), \end{cases}$$

$$V_\delta(\varphi_h(p)) = \begin{cases} \varphi_h(p)\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) + (1-\sigma)\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) & \text{if } p \in [1-\sigma, 1/2), \\ 1 & \text{if } p \in [1/2, 1], \end{cases}$$

and the expected continuation value:

$$\bar{Z}_\delta(p) = \begin{cases} p\left(\frac{\sigma(2-\delta)}{1-\delta+\delta\sigma^2}\right) & \text{if } p \in C_\ell, \\ p\left[\left(\frac{\sigma(1-\sigma)(2-\delta)}{1-\delta+\delta\sigma^2}\right) + \sigma\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right)\right] + \alpha(p)(1-\sigma)\left[\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right] & \text{if } p \in [1-\sigma, 1/2), \\ p(1-\sigma)\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) + (1-\alpha(p))(1-\sigma)\left[\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right] + \alpha(p) & \text{if } p \in [1/2, \sigma), \\ 1 & \text{if } p \in C_h. \end{cases}$$

Let me rearrange this expression to evidence the terms multiplying p:

$$\bar{Z}_\delta(p) = \begin{cases} p\left(\frac{\sigma(2-\delta)}{1-\delta+\delta\sigma^2}\right) & \text{if } p \in C_\ell, \\ p\left[\left(\frac{\sigma(1-\sigma)(2-\delta)}{1-\delta+\delta\sigma^2}\right) + \sigma\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) + (1-\sigma)\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{1-\delta+\delta\sigma^2}\right)\right] + & \\ +\frac{(1-\sigma)^2[\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)]}{(2\sigma-1)(1-\delta+\delta\sigma^2)} & \text{if } p \in [1-\sigma, 1/2), \\ p\left[(1-\sigma)\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) - (1-\sigma)\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{1-\delta+\delta\sigma^2}\right) + 2\sigma - 1\right] + & \\ +\sigma(1-\sigma)\left[\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right] + 1 - \sigma & \text{if } p \in [1/2, \sigma), \\ 1 & \text{if } p \in C_h. \end{cases}$$

Finally, $Z_\delta(p)$ is given by

$$Z_\delta(p) = \begin{cases} p\left(\frac{\sigma\delta(2-\delta)}{1-\delta+\delta\sigma^2}\right) & \text{if } p \in C_\ell, \\ p\left[\left(\frac{\delta\sigma(1-\sigma)(2-\delta)}{1-\delta+\delta\sigma^2}\right) + \delta\sigma\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) + (1-\sigma)\delta\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{1-\delta+\delta\sigma^2}\right) + (2\sigma-1)(1-\delta)\right] & \\ +(1-\delta)(1-\sigma) + \frac{\delta(1-\sigma)^2[\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)]}{(2\sigma-1)(1-\delta+\delta\sigma^2)} & \text{if } p \in [1-\sigma, 1/2), \\ p\left[\delta(1-\sigma)\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) - \delta(1-\sigma)\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{1-\delta+\delta\sigma^2}\right) + (2\sigma-1)\right] + & \\ (1-\sigma) + \delta\sigma(1-\sigma)\left[\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right] & \text{if } p \in [1/2, \sigma), \\ 1 & \text{if } p \in C_h. \end{cases}$$

Consider first $p \in C_\ell$. The greedy splitting implies inducing beliefs 0 and $1-\sigma$ with probabilities $1 - p/(1-\sigma)$ and $p/(1-\sigma)$, respectively. Because $Z_\delta(0) = 0$, this leads to $(pZ_\delta(1-\sigma))/(1-\sigma) = V_\delta(p)$ and consequently $\mathbb{E}_{\tau^{gp}(p)}[Z_\delta(p)] = V_\delta(p)$. Indeed,

$$\begin{aligned} \frac{Z_\delta(1-\sigma)}{1-\sigma} &= \left(\frac{\delta\sigma(1-\sigma)(2-\delta)}{1-\delta+\delta\sigma^2}\right) + \sigma\delta\left(\frac{1-\delta+\delta\sigma^2-\sigma(1-\sigma)(2-\delta)}{(2\sigma-1)(1-\delta+\delta\sigma^2)}\right) + (1-\sigma)\delta\left(\frac{\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)}{(1-\delta+\delta\sigma^2)}\right) + \\ &+ (2\sigma-1)(1-\delta) + (1-\delta) + \frac{\delta(1-\sigma)[\sigma^2(2-\delta)-(1-\delta+\delta\sigma^2)]}{(2\sigma-1)(1-\delta+\delta\sigma^2)}, \\ &= (1-\delta)2\sigma - \delta(1-\sigma) + \delta + \frac{\delta\sigma(1-\sigma^2)(2-\delta)}{1-\delta+\delta\sigma^2}, \\ &= \frac{\sigma(2-\delta)}{1-\delta+\delta\sigma^2}. \end{aligned}$$

Similar analysis holds for $p \geq 1-\sigma$. This shows that the greedy strategy generates $V_\delta(p)$. It remains to show that the greedy policy leads to the concave closure of Z_δ or $\text{cav}[Z_\delta](p) = V_\delta((p)$. Again, suppose first that $p \in C_\ell$. Principal could either set $\rho^+ = 1-\sigma$, $\rho^+ = 1/2$ or $\rho^+ = \sigma$ (those are the possible kinks of the optimal value function). He would choose ρ^+ to maximize $Z_\delta(\rho^+)/\rho^+$. Each ones leads to

$$\frac{Z_\delta(1-\sigma)}{1-\sigma} = \frac{\sigma(2-\delta)}{1-\delta+\delta\sigma^2},$$
$$\frac{Z_\delta(1/2)}{1/2} = (1-\sigma)\left(\frac{(2\sigma-1-\delta\sigma)(1-\delta+\delta\sigma^2)+\delta\sigma^3(2-\delta)}{2(2\sigma-1)(1-\delta+\delta\sigma^2)}\right)+$$
$$+\left(\frac{2(1-\delta+\delta\sigma^2)[2\sigma(1-\sigma)\delta+(2\sigma-1)^2]-2(2-\delta)\delta\sigma(1-\sigma)(1-2\sigma+2\sigma^2)}{2(2\sigma-1)(1-\delta+\delta\sigma^2)}\right),$$
$$\frac{Z_\delta(\sigma)}{\sigma} = \frac{1}{\sigma}.$$

With some algebra, it follows that

$$\frac{Z_\delta(1-\sigma)}{1-\sigma} - \frac{Z_\delta(\sigma)}{\sigma} \geq 0 \Leftrightarrow 2\sigma^2 \geq 1.$$
$$\frac{Z_\delta(\sigma)}{\sigma} - \frac{Z_\delta(1/2)}{1/2} \geq 0 \Rightarrow \sigma^3(1+2\delta) - \sigma^2(1+2\delta) + 3\sigma - 1 \geq 0.$$

The first inequality is true by assumption; the second is true for every δ because $\sigma \geq 1/2$. Therefore, whenever $p \in C_\ell$, it is optimal to split beliefs according to the greedy strategy. Now suppose that $p \in [1-\sigma, 1/2)$. Principal could set $\rho^- = 0, \rho^+ = 1/2$, $\rho^- = 0, \rho^+ = \sigma$, $\rho^- = 1-\sigma, \rho^+ = 1/2$ or $\rho^- = 1-\sigma, \rho^+ = \sigma$. The splitting between $\rho^- = 1-\sigma$ and $\rho^+ = \sigma$ is better than the splitting between $\rho^- = 0$ and $\rho^+ = \sigma$, because

$$(1-\sigma)\left[\frac{\sigma-p}{2\sigma-1}\right]\frac{Z_\delta(1-\sigma)}{1-\sigma} + \sigma\left[\frac{p-(1-\sigma)}{2\sigma-1}\right]\frac{Z_\delta(\sigma)}{\sigma} \geq \left((1-\sigma)\left[\frac{\sigma-p}{2\sigma-1}\right]+\sigma\left[\frac{p-(1-\sigma)}{2\sigma-1}\right]\right)\frac{Z_\delta(\sigma)}{\sigma},$$
$$= p\frac{Z_\delta(\sigma)}{\sigma}.$$

Because $Z_\delta(\sigma)/\sigma \geq Z_\delta(1/2)/(1/2)$, the splitting between $\rho^- = 1-\sigma$ and ρ^+ is also better than the splitting between $\rho^- = 0$ and $\rho^+ = 1/2$. Moreover, the splitting between $\rho^- = 1-\sigma$ and $\rho^+ = \sigma$ is better than the splitting between $\rho^- = 1-\sigma$ and $\rho^+ = 1/2$, because

$$\left[\frac{1-2p}{2\sigma-1}\right]Z_\delta(1-\sigma) + 2\left[\frac{p-(1-\sigma)}{2\sigma-1}\right]Z_\delta(1/2) \leq \frac{1}{2}\left[\frac{\sigma-p}{2\sigma-1}\right]Z_\delta(1-\sigma) + \sigma\left[\frac{p-(1-\sigma)}{2\sigma-1}\right]Z_\delta(\sigma),$$
$$\leq \sigma\mathbb{E}_{\tau^{gp}(p)}[Z_\delta],$$
$$\leq \mathbb{E}_{\tau^{gp}(p)}[Z_\delta].$$

Finally, suppose $p \in [1/2, \sigma)$. In this case, principal could set $\rho^- = 1-\sigma$, $\rho^+ = \sigma$; $\rho^- = 1/2$, $\rho^+ = \sigma$ or $\rho^- = 0$, $\rho^+ = \sigma$. I have already showed that the splitting between $\rho^- = 1-\sigma$ and $\rho^+ = \sigma$ is better than the splitting between $\rho^- = 0$ and $\rho^+ = \sigma$. It remains to show that is also better than the splitting between $\rho^- = 1/2$ and $\rho^+ = \sigma$. Indeed,

$$2\left[\frac{\sigma-p}{2\sigma-1}\right]Z_\delta(1/2) + \left[\frac{2p-1}{2\sigma-1}\right]Z_\delta(\sigma) \leq (1-\sigma)\left[\frac{\sigma-p}{2\sigma-1}\right]Z_\delta(1-\sigma) + \frac{1}{2}\left[\frac{p-(1-\sigma)}{2\sigma-1}\right]Z_\delta(\sigma),$$
$$\leq \frac{1}{2}\mathbb{E}_{\tau^{gp}(p)}[Z_\delta],$$
$$\leq \mathbb{E}_{\tau^{gp}(p)}[Z_\delta].$$

□

Finally, let me compute the stationary distribution of public beliefs under the greedy strategy and $p_1 = 1/2$. At $t = 1$, principal induces two posteriors σ and $1-\sigma$ with probabilities $\tau^{gp}(1-\sigma; 1/2) = 1/2$ and $\tau^{gp}(\sigma; 1/2) = 1/2$. Thus, $\tau^{gp}(\sigma; 1/2)$ is the probability of a cascade towards action 2 starts by $t = 1$.

At $t = 2$, there are two possible interim beliefs: $\varphi_\ell(1-\sigma)$, $1/2$ and σ with probabilities

$$\begin{aligned}\lambda_2^{gp}(\varphi_\ell(1-\sigma)) &= P(1-\sigma, \varphi_1(1-\sigma))\tau^{gp}(1-\sigma; 1/2) = (1/2)[(1-\sigma)^2 + \sigma^2],\\ \lambda_2^{gp}(1/2) &= P(1-\sigma, 1/2)\tau^{gp}(1-\sigma; 1/2) = \sigma(1-\sigma),\\ \lambda_2^{gp}(\sigma) &= P(\sigma,\sigma)\tau^{gp}(\sigma; 1/2) = 1/2.\end{aligned}$$

Principal induces possible beliefs 0, $1-\sigma$ and σ with probabilities:

$$\begin{aligned}\hat\lambda_2^{gp}(0) &= \tau^{gp}(0; \varphi_1(1-\sigma))\lambda_2^{gp}(\varphi_\ell(1-\sigma)) = (1/2)\sigma[2\sigma - 1],\\ \hat\lambda_2^{gp}(1-\sigma) &= \tau^{gp}(1-\sigma; \varphi_\ell(1-\sigma))\lambda_2^{gp}(\varphi_\ell(1-\sigma)) + \tau^{gp}(1-\sigma; 1/2)\lambda_2^{gp}(1/2) = (1/2)(1-\sigma^2),\\ \hat\lambda_2^{gp}(\sigma) &= \tau^{gp}(\sigma;\sigma)\lambda_2^{gp}(\sigma) + \tau^{gp}(\sigma; 1/2)\lambda_2^{gp}(1/2) = (1/2)[1 + \sigma(1-\sigma)].\end{aligned}$$

If agent 2 has interim belief σ, she will take action 2 no matter the private signal. All other agents will do the same. Thus, $\hat\lambda_2^{gp}(\sigma)$ is the probability of a cascade towards action 2 has started at $t = 2$. Same reasoning leads to $\lambda_2^{gp}(0)$ being the probability of a cascade towards action 1 starts by $t = 2$.

Are $t = 3$, there are three possible interim beliefs: 0, $\varphi_\ell(1-\sigma)$ and $1/2$, σ with probabilities

$$\begin{aligned}\lambda_3^{gp}(0) &= P(0,0)\hat\lambda_2^{gp}(0) = (1/2)\sigma(2\sigma - 1),\\ \lambda_3^{gp}(\varphi_\ell(1-\sigma)) &= P(1-\sigma, \varphi_1(1-\sigma))\hat\lambda_2^{gp}(1-\sigma) = (1/2)[(1-\sigma)^2 + \sigma^2](1-\sigma^2),\\ \lambda_3^{gp}(1/2) &= P(1-\sigma, 1/2)\hat\lambda_2^{gp}(1-\sigma) = \sigma(1-\sigma)(1-\sigma^2),\\ \lambda_3^{gp}(\sigma) &= \hat\lambda_2^{gp}(\sigma) = (1/2)[1 + \sigma(1-\sigma)].\end{aligned}$$

Principal induces beliefs 0, $1-\sigma$ and σ with probabilities

$$\begin{aligned}\hat\lambda_3^{gp}(0) &= \lambda_3^{gp}(0) + \tau^{gp}(0; \varphi_1(1-\sigma))\lambda_3^{gp}(\varphi_\ell(1-\sigma)) = (1/2)\sigma(2\sigma - 1)[1 + (1-\sigma^2)],\\ \hat\lambda_3^{gp}(1-\sigma) &= \tau^{gp}(1-\sigma; \varphi_1(1-\sigma))\lambda_3^{gp}(\varphi_1(1-\sigma)) + \tau^{gp}(1-\sigma; 1/2)\lambda_3^{gp}(1/2) = (1/2)(1-\sigma^2)^2,\\ \hat\lambda_3^{gp}(\sigma) &= \tau^{gp}(\sigma,\sigma)\lambda_3^{gp}(\sigma) + \tau^{gp}(\sigma; 1/2)\lambda_3^{gp}(1/2) = (1/2)[1 + \sigma(1-\sigma) + \sigma(1-\sigma)(1-\sigma^2)].\end{aligned}$$

If agent $t = 3$ has induced belief σ, she will take action 2 no matter the private signal and other all agents will do so as well. So $\hat\lambda_3^{gp}(\sigma)$ is the probability of a cascade towards action 2 has started by $t = 3$. Same reasoning holds for $\hat\lambda_3^{gp}(0)$ being the probability of a cascade towards action 1 has started by $t = 3$.

At $t = 4$, the possible interim beliefs are 0, $\varphi_\ell(1-\sigma)$, $1/2$, σ with probabilities

$$\begin{aligned}\lambda_4^{gp}(0) &= P(0,0)\hat\lambda_3^{gp}(0) = (1/2)\sigma(2\sigma - 1)[1 + (1-\sigma^2)],\\ \lambda_4^{gp}(\varphi_\ell(1-\sigma)) &= P(1-\sigma, \varphi_1(1-\sigma))\hat\lambda_3^{gp}(1-\sigma) = (1/2)[(1-\sigma)^2 + \sigma^2](1-\sigma^2)^2,\\ \lambda_4^{gp}(1/2) &= P(1-\sigma, 1/2)\hat\lambda_3^{gp}(1-\sigma) = \sigma(1-\sigma)(1-\sigma^2)^2,\\ \lambda_4^{gp}(\sigma) &= \hat\lambda_3^{gp}(\sigma) = (1/2)[1 + \sigma(1-\sigma) + \sigma(1-\sigma)(1-\sigma^2)].\end{aligned}$$

Principal then induces beliefs in 0, $1-\sigma$ and σ with probabilities

$$\begin{aligned}\hat{\lambda}_4^{gp}(0) &= \lambda_4^{gp}(0) + \tau^{gp}(0;\varphi_1(1-\sigma))\lambda_4^{gp}(\varphi_1(1-\sigma)) = (1/2)\sigma(2\sigma-1)[1+(1-\sigma^2)+(1-\sigma^2)^2],\\ \hat{\lambda}_4^{gp}(1-\sigma) &= \tau^{gp}(1-\sigma;\varphi_1(1-\sigma))\lambda_4^{gp}(\varphi_1(1-\sigma)) + \tau^{gp}(1-\sigma;1/2)\lambda_4^{gp}(1/2) = (1/2)(1-\sigma^2)^3,\\ \hat{\lambda}_4^{gp}(\sigma) &= \tau^{gp}(\sigma,\sigma)\lambda_4^{gp}(\sigma) + \tau^{gp}(\sigma;1/2)\lambda_4^{gp}(1/2) = (1/2)[1+\sigma(1-\sigma)[1+(1-\sigma^2)+(1-\sigma^2)^2].\end{aligned}$$

By now a pattern is clear. At $t \geq 1$, principal induces beliefs 0, $1-\sigma$ and σ with probabilities

$$\begin{aligned}\hat{\lambda}_t^{gp}(0) &= (1/2)\sigma(2\sigma-1)\sum_{\tau=0}^{t-2}(1-\sigma^2)^\tau = \frac{1}{2}\left[\frac{2\sigma-1}{\sigma}\right](1-(1-\sigma^2)^{t-1}),\\ \hat{\lambda}_t^{gp}(1-\sigma) &= (1/2)(1-\sigma^2)^{t-1},\\ \hat{\lambda}_t^{gp}(\sigma) &= (1/2)[1+\sigma(1-\sigma)\sum_{\tau=0}^{t-2}(1-\sigma^2)^\tau] = \frac{1}{2}\left[1+(1-\sigma)\left(\frac{1-(1-\sigma^2)^{t-1}}{\sigma}\right)\right].\end{aligned}$$

The probabilities $\hat{\lambda}_t^{gp}(0)$ and $\hat{\lambda}_t^{gp}(\sigma)$ represent the probabilities of a cascade towards action 1 and action 2 starting by t, respectively. The probabilities of interim beliefs at $t+1$ are

$$\begin{aligned}\lambda_{t+1}^{gp}(0) &= \frac{1}{2}\left[\frac{2\sigma-1}{\sigma}\right](1-(1-\sigma^2)^{t-1}),\\ \lambda_{t+1}^{gp}(\varphi_1(1-\sigma)) &= (1/2)[(1-\sigma)^2+\sigma^2](1-\sigma^2)^{t-1},\\ \lambda_{t+1}^{gp}(1/2) &= \sigma(1-\sigma)(1-\sigma^2)^{t-1},\\ \lambda_{t+1}^{gp}(\sigma) &= \frac{1}{2}\left[1+(1-\sigma)\left(\frac{1-(1-\sigma^2)^{t-1}}{\sigma}\right)\right].\end{aligned}$$

Note that the limiting probability of having a cascade towards action 2 is given by $\hat{\lambda}_\infty^{gp}(\sigma) = \lim_{t\to\infty}\hat{\lambda}_t^{gp}(\sigma) = 1/(2\sigma)$. Likewise, the limiting probability of having a cascade towards action 1 is given by $\hat{\lambda}_\infty^{gp}(0) = \lim_{t\to\infty}\hat{\lambda}_t^{gp}(0) = (1/2)[(2\sigma-1)/\sigma]$.

Example with uniform distribution

The private belief space is $[\underline{q},\bar{q}]$ where $\underline{q} := (1/2)(1-\sigma)$ and $\bar{q} := (1/2)(1+\sigma)$. The parameter σ thus governs how revealing private information can be, just as it was the case in the illustrative example. The unconditional density is $g(q) = 1/\sigma$ for $q \in [\underline{q},\bar{q}]$; the conditional densities are $g^h = 2(1-q)(1/\sigma)$ and $g^L = 2q(1/\sigma)$. The cascade sets are $C_\ell = [0, \frac{1}{2}(1-\sigma))$ and $C_h = [\frac{1}{2}(1+\sigma), 1]$. That leads to the following expected probabilities of action h:

$$\alpha^H(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ \frac{1}{\sigma}\left[2p - p^2 - (2\underline{q} - \underline{q}^2)\right] & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases} \qquad \alpha^L(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ \frac{1}{\sigma}\left[p^2 - \underline{q}^2\right] & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases}$$

$$\alpha(p) = \begin{cases} 0 & \text{if } p \in C_\ell, \\ \frac{p}{\sigma}[2p - p^2 - (2\underline{q} - \underline{q}^2)] + \frac{1-p}{\sigma}[p^2 - \underline{q}^2] & \text{if } p \notin C_\ell \cup C_h, \\ 1 & \text{if } p \in C_h. \end{cases}$$

The system moves to another public belief according to the transition functions

$$\varphi_h(p) := \begin{cases} p & \text{if } p \in C_h, \\ \frac{p(2p-p^2-2\underline{q}+\underline{q}^2)}{p(2p-p^2-2\underline{q}+\underline{q}^2)+(1-p)(p^2-\underline{q}^2)} & \text{if } p \notin C_\ell \cup C_h. \end{cases}$$

$$\varphi_\ell(p) := \begin{cases} p & \text{if } p \in C_\ell, \\ \frac{p\left[1-\frac{1}{\sigma}(2p-p^2-2\underline{q}+\underline{q}^2)\right]}{p\left[1-\frac{1}{\sigma}(2p-p^2-2\underline{q}+\underline{q}^2)\right]+(1-p)\left[1-\frac{1}{\sigma}(p^2-\underline{q}^2)\right]} & \text{if } p \notin C_\ell \cup C_h. \end{cases}$$

From theorem 1, single disclosure is optimal if and only if $4(1-\bar{q})\bar{q}^2 g(\bar{q}) \geq 1$. When the distribution is uniform, this comes down to

$$\bar{q} - \underline{q} \leq 4(1-\bar{q})\bar{q}^2.$$

Because $\bar{q} - \underline{q} = \sigma$ and $\bar{q} = (1/2)(1+\sigma)$ in this example, single disclosure will be optimal iff

$$\sigma \leq (1-\sigma)(1+\sigma)^2 \Leftrightarrow \sigma \leq \sigma^* \approx 0.54.$$

The cut-off p^* above which $\alpha > V^{sd}$ is given by

$$p^* = \frac{1}{2} - \frac{\sigma}{4} + \frac{\sqrt{\frac{16}{\sigma+1} + \sigma^2 - 8}}{4}.$$