FUNDAÇÃO GETULIO VARGAS

ESCOLA DE ECONOMIA DE SÃO PAULO

CAIO PAES LEME LORECCHIO

ESSAYS ON INFORMATION DESIGN

SÃO PAULO

2021

CAIO PAES LEME LORECCHIO

ESSAYS ON INFORMATION DESIGN

Tese apresentada à Escola de Economia de São Paulo da Fundação Getulio Vargas como requisito para obtenção do título de Doutor em Economia de Empresas.

Campo de Conhecimento:
Microeconomia - Teoria dos Jogos

Orientador: Prof. Dr. Daniel Monte

SÃO PAULO

2021

Lorecchio, Caio Paes Leme.
Essays on information design / Caio Paes Leme Lorecchio. - 2021.
119 f.

Orientador: Daniel Monte.
Tese (doutorado CDEE) – Fundação Getulio Vargas, Escola de Economia de São Paulo.

1. Visualização da informação. 2. Comunicação na tecnologia. 3. Teoria dos jogos. 4. Persuasão (Retorica). 5. Aprendizagem social. I. Monte, Daniel. II. Tese (doutorado) – Escola de Economia de São Paulo. III. Fundação Getulio Vargas. IV. Título.

CDU 330.115

Ficha Catalográfica elaborada por: Isabele Oliveira dos Santos Garcia CRB SP-010191/O
Biblioteca Karl A. Boedecker da Fundação Getulio Vargas - SP

CAIO PAES LEME LORECCHIO

ESSAYS ON INFORMATION DESIGN

Tese apresentada à Escola de Economia de São Paulo da Fundação Getulio Vargas como requisito para obtenção do título de Doutor em Economia de Empresas.

Campo de Conhecimento:
Microeconomia - Teoria dos Jogos

Data de aprovação:
11/06/2021

Banca Examinadora:

Prof. Dr. Daniel Monte (Orientador)
FGV-EESP

Prof. Dr. Braz Ministerio de Camargo
FGV-EESP

Prof. Dr. Marcos Yamada Nakaguma
FGV-EEESP

Prof. Dr. Eduardo Faingold
INSPER

Prof. Dr. Thomas Wiseman
UT Austin

ACKNOWLEDGMENTS

I am grateful to my dissertation committee: Braz Camargo, Daniel Monte, Eduardo Faingold, Marcos Nakaguma and Thomas Wiseman. I am also thankful for conversations with V. Bashkar and Maxwell Stinchcombe. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

ABSTRACT

This dissertation is a study in simple information design. In the first chapter, I examine dynamic information design under constrained communication rules, motivated by record keeping regulations and targeted transparency policies. Constrained rules require that messages cannot be conditioned on private information. A long-run and uninformed designer wishes to persuade short-lived agents to invest in a project of fixed, but unknown, quality. The restriction on the possible communication rules hinders the designer without benefiting the agents. It unambiguously decreases welfare if the designer has altruistic motives. The optimal policy depends on the conditional payoff distribution and if failures are very informative, constrained rules are less restrictive: simple policies approximate the designer's first-best payoff.

In the second chapter, I consider dynamic information design with bad reputational concerns. Customers increasingly rely on rating systems when hiring experts. If a rating system is ill designed, the reputational gain from taking a certain action might be excessively high, inducing the experts to over-choose it. In such a case, markets could even cease to exist altogether. I show how to design simple, optimal rating systems for both customers and experts. Such rating systems overcome market failures and improve upon both the full memory case and the case with no memory at all. Customers benefit from a higher number of ratings, while binary systems are sufficient to achieve the expert's highest value.

In the third chapter, I study information design in an observational learning environment. 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.

Keywords: dynamic information design, rating systems, social learning.

LIST OF FIGURES

CONTENTS

1 INFORMATION DESIGN UNDER CONSTRAINED COMMUNICATION RULES

1.1 Introduction

In dynamic relationships, information design is crucial in determining the success of an information intermediary in motivating actions[1]. Online platforms employ a variety of forms to communicate with their customers, for instance. Credit bureaus provide information about previous behavior of debtors, and are fundamental for the well functioning of credit markets. Importantly, though, in many instances, not all information generated is used. First, the length of memory that can be made publicly available might be limited. Second, some or all information used by a sender might be required to be made public. Indeed, there is a growing concern with targeted transparency: the requirement to disclose specific information in a standardized format.[2] There are many cases of targeted transparency restrictions, such as the Dodd-Frank Act, which, among other things, mandate credit agencies to publicly disclose their methodologies and data relied on to define grades.[3] These imposed restrictions raise the following questions: How is communication affected if the information used is restricted, either by memory or by transparency? Do these restrictions on communication benefit or hinder the agents?

In our model, a long-run designer faces an infinite sequence of short-lived agents, who must each decide whether or not to invest in a project. The project is of a fixed quality, which we assume to be either good or bad. This underlying quality is unknown to the designer and the agents. In our benchmark model, the designer's objective function is state-independent: she gets a positive payoff whenever an agent invests and zero otherwise, regardless of the true underlying quality. If an agent invests, the project randomly generates a payoff to that agent. A good quality project yields positive expected payoffs, while a bad quality project yields negative expected payoffs. If the agent does not invest, she gets a zero payoff. Any given agent will invest if and only if her belief that the project is of bad quality is sufficiently low. The realized payoff is public to the current agent and designer, but not to future agents. Instead, each agent observes only the current message sent by the designer before deciding on her action.

The designer can credibly commit to a constrained information policy: a map from the current message and current payoffs to a distribution over future messages. This class of communication rules includes the case in which every agent observes the entire history of messages. Importantly, though, it excludes rules that condition messages on the project's type.[4] Designer's future recommendation is a map from the same information that is available to the current agent. In this sense, we may say that these rules are *transparent*.

[1] This is a joint work with Daniel Monte.

[2] See, for example, Weil, Graham and Fung (2013).

[3] There are many other examples, such as the No Child Left Behind Act, which requires certain schools to publish their report cards.

[4] In other words, the designer is choosing the set of messages (finite or countably infinite set) and a Markovian transition rule.

Our definition of constrained rules allows for an alternative interpretation as well. The size of the chosen message space is a measure of complexity. In particular, finite message spaces imply that communication rules have finite complexity. Thus, our restriction on the set of communication rules is intended to *either* represent transparent communication rules or (and) capture potential restrictions on the complexity of communication systems.[5] We can compute the value of the informational advantage that the optimal unrestricted communication protocol has over communicating through these restricted rules. We also provide conditions under which simple rules approximate the designer's first best payoff.

Three key ingredients of our environment are: (1) the designer is uninformed about the state of the world and shares a common prior with the agents, so she must learn in order to be able to persuade; (2) learning must be elicited, that is, realized payoffs signal the true quality, but will only be generated whenever the agents invest, implying that messages serve a dual purpose: they generate both payoffs and signals. Finally, (3) we assume that agents do not know their place in the queue, that is, they do not know the calendar time and, instead, they form beliefs about the actual history.[6]

One way of thinking about our constrained rule is as if it is a rating system, as in the commonly used star-rating systems. In this case, we must interpret each message as a different rating. Indeed, for ease of visualization, from now on we will refer to our communication rules as rating systems, and each message as a rating. In the special case where the designer chooses to use only two messages, the output of our system is simply a direct recommendation.

We obtain the optimal communication system and we show that it mirrors the optimal static information design, as we discuss below, but with some fundamental differences. We also show the implications of these communication restrictions on the designer's welfare, the agents' welfare and how the optimal policy depends on the conditional payoff distributions.

To understand the problem more intuitively, let us consider the same underlying stage game within a Bayesian persuasion framework.[7] Think of the static problem in which the designer wishes to maximize the agent's probability of investing in the project. Is she could credibly commit to a map from the true state of the world to a set of messages, she would choose to "split the posteriors" of the agent. If we return to our designer with finite messages (interpreted hereafter as finite ratings), ideally she would like to have two extreme ratings to reproduce the Bayesian persuasion benchmark. However, our designer cannot commit to a map from states of the world to recommendations, and it is only through the induced actions of the agents that she can (partially) learn the true state of the world.

We show that the belief spread induced by our designer's communication rule cannot be as large as in the static Bayesian persuasion literature. The intuition is as follows.[8] Once the communication rule (hereafter, rating system) is designed, the set of messages (hereafter, ratings) can be partitioned into two subsets: ratings in which the induced belief is such that agents invest and ratings in which the agents do not invest.[9]

---

[5]Restrictions on complexity, such as finite memory, may come about due to legal restrictions. For example, in criminal courts and credit score agencies, many countries impose legal constraints on how past information can be used.

[6]When agents do not know the calendar time, there is an issue on how to compute beliefs. We postpone this discussion to section 1.2. Note also that if agents knew the calendar time, there would be a cold-start problem if the prior is above a given threshold: the first agent would not invest, so there would be no belief updating and no other agents would invest either.

[7]See, for example, Kamenica and Gentzkow (2011) or Rayo and Segal (2010).

[8]For now, let us focus on irreducible systems, but in the main part of the paper we prove the result for the general case.

[9]Indifference can be accommodated easily, so for now, let us ignore indifference.

At least one rating from one subset must interact with at least one rating from the other subset, a result that comes directly from irreducibility. The beliefs of these two ratings in consideration can be far apart only insofar as the actions played in each of them can give enough information about the underlying state of the world. However, in one of these ratings agents are not investing, so no information is being generated. This allows us to compute a maximum bound on the belief spread as a function of the parameters of the model.

Most importantly, we show that, for this benchmark specification, only two ratings are needed for the optimal rating system, which we interpret as a direct recommendation: either recommend to invest or to not invest. Note that although this result looks similar to the revelation principle in the Bayesian persuasion literature, it is actually not a consequence in our environment. Ratings serve the dual purpose of learning and inducing payoff relevant actions, but for this specific benchmark environment, whenever the agents do not invest, nothing is learned. Moreover, in order to induce investment it suffices to have a rating with a belief on the agent's indifference condition, and learning further is not optimal for the designer.

In order to fully understand the role that ratings are playing, we also solve for the optimal system in a world in which learning occurs regardless of the induced action, that is, signals are generated even when there is no investment. In this case, we do not have a revelation principle of this sort; instead, more ratings imply a higher payoff. Such result also holds when the designer's objective function is to maximize social welfare, since in this case, learning is important even in the set of ratings that induce investment.[10]

Our model also enables us to pin down an important relation between the signal structure and Bayesian persuasion. Specifically, the relevant statistic for persuading is the strength of bad news, that is, the informativeness of the signal with highest likelihood between the bad and the good state of the world. As this informativeness increases, the designer's payoff becomes closer to the Bayesian persuasion payoff. In contrast, good news is not relevant for persuasion.

The intuition here is that persuasion requires agents to be very convinced of the bad state of the world, but only partially convinced by the good state. Thus, the designer would like to have an inflated belief for the bad state and an indifference belief for the good state. Because signals are not generated in the bad rating, the only way to have a high belief in this no-signal rating is if it is possible to have an informative transition from the good rating to the bad one. This is only possible insofar as a signal might help you, and since there is no learning when there is no investment, the only useful signal is the one coming from the investment rating. In contrast, in the no experimentation case in which the designer is capable of learning even when there is no investment, good news is also relevant, and the more informative the signals are (regardless if good or bad news), the closer to the Bayesian persuasion payoff the designer will be.

[10] An interesting implication of optimality and Bayes-plausibility, is that even when the designer uses the many ratings to improve her payoff, the optimal rule implies staying most often at the extreme ratings. This result is consistent with the results in the bounded memory literature, such as Hellman and Cover (1970) and Wilson (2014).

Our paper combines Bayesian persuasion with complexity considerations. Thus, perhaps the most related papers to ours are the ones that combine dynamic games with limited record keeping. Ekmekci (2011) studies a rating system in a moral hazard game, and the system structure is very close to ours - yet with some fundamental differences. For instance, in his model, there is a permanent flow of informative signals about the long-run player's type. In credence goods markets, investment is needed in order to have informative signals, so there might not exist such a permanent flow of informative signals.

Our paper is related to the literature of learning under limited memory. Hellman and Cover (1970) and Wilson (2014) consider a bounded memory agent that has to design the optimal memory for a two-hypothesis testing problem. Monte (2013, 2014) studies a model in which a bounded memory player is playing a reputation game against an informed player. We contribute to this literature since we study a similar underlying problem, but under the perspective of a designer that must persuade agents to invest at the same time that it must learn. In our model there is a subset of the memory system in which there is no information revelation and the designer must choose how long to spend in those ratings.

Kovbasyuk and Spagnolo (2018), Liu (2011) and Liu and Skrzypacz (2014) study information transmission in dynamic markets with limited record keeping as well. Types are constantly changing over time in Kovbasyuk and Spagnolo (2018)'s environment, thus it is natural to look at stationary rules, like they do. In our case, types are fixed throughout time, but the stationarity comes from the constraint on the set of possible communication rules. In Liu (2011), past information is costly to observe and in Liu and Skrzypacz (2014), the focus is on how reputation evolves in a dynamic market with limited record keeping.

We also connect with the literature on dynamic information design and Bayesian persuasion. Aumann, Maschler and Stearns (1995) study this question in repeated games with incomplete information. Recent papers on the literature include Brocas and Carrillo (2007), Rayo and Segal (2010), Kamenica and Gentzkow (2011), Kamenica and Gentzkow (2014), Bergemann and Morris (2016), Ely (2017), Best and Quigley (2017), Mathevet, Perego and Taneva (2020), Matyskova (2018), Lipnowski, Ravid and Shishkin (2018), Taneva (2019), Li and Norman (2020) and Doval and Ely (2020).[11] We tackle a related problem, but our designer has much less information than the canonical information designer: instead of observing the actual type and sending a recommendation given the type, our designer learns about the true type only by the signals received over time, which are constrained by the recommendations that she can give. In particular, a negative recommendation implies less learning, since in those periods the customers do not hire and the designer learns nothing.[12]

By studying optimal rating structures from the principal's point of view, our work also relates to the literature of reputation and optimal information design (Hörner and Lambert, 2016; Smolin, 2017; and Bhaskar and Thomas, 2017). Also related is Halac, Kartik and Liu (2017), which studies the design of optimal disclosure policies in dynamic contests. Our model is somewhat similar to Kremer, Mansour and Perry (2014) and Che and Horner (2017), in which a recommender system must learn from the product feedback. The difference between our paper and theirs is the restriction on the communication systems that we impose. Moreover, we consider a different signal structure and different objective function for the designer.

[11] See also Bergemann and Morris (2019) and Kamenica (2019) for two recent surveys on information design.

[12] There is a literature in computer science and algorithmic game theory studying information design and complexity. Dughmi (2017) provides a recent survey on the topic.

Our environment is also similar to Glazer, Kremer and Perry (2015), but in their paper, agents must pay a cost to observe the signal. Vong (2021) studies optimal ratings in a repeated game in which there is adverse selection and moral hazard. Other recent related papers are Lillethun (2017) and Sperisen (2018). Laiho, Murto and Salmi (2021) study a rather analogous problem, in which a decision maker chooses in every period whether or not to increase the consumption of a variable of an unknown, fixed type. The more she experiments, the more she learns about the type.

Our rating system might resemble Best and Quigley (2017)'s review aggregator, but there are fundamental differences. In their paper, the long-run designer cannot commit to a rule and the state of the world is changing over time. Thus, their review aggregator is meant to provide the designer a way to approach the Bayesian persuasion payoff in a world with no commitment and changing types (states of the world). In contrast, our rating system is the information policy to which the designer credibly commits to and is intended to solve a joint learning problem between the designer and the sequence of agents.

1.2 Model

An infinite sequence of myopic agents enter the market, one at a time. Each agent chooses whether or not to invest in a given project. The project is one of two types: $\Omega = \{B, G\}$. The type is fixed throughout the game and unknown to the agents. Each type randomly generates a payoff to an agent every period in which an agent invests. If the agent does not invest, her payoff is zero. Suppose that there are $M \geq 2$ possible payoff realizations: $X = \{x_1, x_2, ..., x_M\}$. If the state of the world is B, payoffs are drawn according to $\Pr(x_m|B) = \gamma_m^B$, $m \in \{1, 2, ..., M\}$ whereas if the state of the world is G, payoffs are drawn according to $\Pr(x_m|G) = \gamma_m^G$.[13] Therefore, the conditional expectation of investing if the state is B is $\sum_{m=1}^M \gamma_m^B x_m$, and similarly if the state is G. We assume that the conditional expected payoff of investing is such that

$$\sum_{m=1}^{M} \gamma_m^G x_m > 0 > \sum_{m=1}^{M} \gamma_m^B x_m.$$

Every agent's initial prior probability that the type of the project is bad is denoted by ρ, with $\rho \in (0,1)$. In a one-shot interaction, an agent decides to invest if and only if

$$\rho \sum_{m=1}^{M} \gamma_m^B x_m + (1-\rho) \sum_{m=1}^{M} \gamma_m^G x_m \geq 0.$$

This inequality defines an indifference threshold, ρ^*, which is given by

$$\rho^* = \frac{\sum_{m=1}^M \gamma_m^G x_m}{\sum_{m=1}^M \gamma_m^G x_m - \sum_{m=1}^M \gamma_m^B x_m}.$$

[13] Our results extend naturally to more general distributions F^B and F^G, such as the case in which F^B and F^G are absolutely continuous with respect to the Lebesgue measure, with corresponding continuous densities f^B and f^G uniformly bounded away from 0 and ∞. The discrete case that we use in the paper has the advantage of allowing us to construct the exact optimal rating system, instead of the ε-optimal system that may be required if the payoff distribution is continuous.

Whenever $\rho > \rho^*$, the agent in a static problem does not invest. In our benchmark model, we make the simplifying assumption that agents do not observe calendar time. Formally, we assume that they attribute an uniform (improper) prior distribution to all possible positions in the sequence. We will consider alternatives ways of modeling position uncertainty later on. For now, it suffices to say that their only information is the message (rating) send by the designer. Note that if agents could observe the entire sequence of previous actions and payoff realizations, the market would break down, because the first agent would not invest and since she would not do so, no information would be generated and subsequent agents would also choose not to invest.

If $\rho \leq \rho^*$, it is optimal for the first agent to invest. If subsequent agents have access to the full history of the game, they will update their beliefs on the quality of the project as new information arrives. As long as the posterior is smaller than ρ^*, agents keep investing. As soon as the posterior exceeds ρ^*, agents cease to invest and the market collapses. A rating system might overcome this problem if agents have limited information about past outcomes. A straightforward implication of our model is that if $\rho \leq \rho^*$, then a rating system that conceals all information will induce investment every period. For this reason, we focus on the more interesting case, that is, when $\rho > \rho^*$.

We will consider the information design problem faced by a designer restricted to a system with finitely many messages and a stationary transparent rule. Thus, his information technology is constrained in two ways. First, he cannot disclose every possible history of past actions, payoffs and messages, because a finite message space is smaller than the set of all public histories of this game. Second and most important, he cannot convey information conditioned on what is not observable to agents. It is in this sense that we say that information must be transparent. Therefore, messages must depend only on the message space and the payoff space. We will denote this restricted system as a rating system.

Definition 1.1. *A rating system is a tuple $\phi = (\mathscr{I}, \varphi_0, \varphi)$, where $\mathscr{I} = \{1, 2, ..., I\}$ is a finite set of ratings; $\varphi_0 \in \Delta(\mathscr{I})$ is an initial probability distribution ratings; and $\varphi : \mathscr{I} \times [X \cup \{\emptyset\}] \to \Delta(\mathscr{I})$ is a transition rule, where $\emptyset$ represents the event in which there was no investment and X is the set of returns to the investment, which are public to the designer and the current agent.*

Because agents do not observe calendar time, but only the current rating,[14] they have the same strategy - a map from the current rating to a probability of investment. We refer to their strategies as $\alpha : \mathscr{I} \to [0, 1]$.

Agents compute the probability distribution over the histories as if the game had been going on for a long time, and their beliefs are computed using steady-state probabilities (or time-average convergence). To give a more detailed description of how beliefs are computed, first note that any given rating system ϕ together with a given strategy α defines Markov matrices $T^G = (\tau^G_{i,j})$ and $T^B = (\tau^B_{i,j})$ for the good and bad state of the world, respectively. The transition $\tau^\theta_{i,j}$ from rating i to rating j given θ is

$$\tau^\theta_{i,j} = \alpha_i \sum_{m=1}^{M} \gamma^\theta_m \varphi^m_{i,j} + (1 - \alpha_i)\varphi_{i,j},$$

[14] Strictly speaking, they also observe the resulting payoff of their investment, but since they are short-lived, this is irrelevant for what follows.

where $\varphi_{i,j}^m$ represents the transition from i to j upon the observation of payoff x_m and $\varphi_{i,j}$ stands for the transition when there is no investment. A well-known result in Markov processes ensures that the sequences $\{(1/t)\sum_{n=0}^{t-1}(T^B)^n\}_{t=1}^{\infty}$ and $\{(1/t)\sum_{n=0}^{t-1}(T^G)^n\}_{t=1}^{\infty}$ converge to stochastic matrices, each row of them being an invariant distribution. Thus, for every initial distribution φ_0 and for each θ, there will be an unique invariant distribution $f^\theta = (f_i)_{i\in\mathscr{I}}$ over the ratings.[15]

The invariant distributions will be used to calculate both the designer's expected payoff and the agents' beliefs at any given rating. The designer seeks to maximize the *ex-ante* probability of investment, so his expected payoff given a rating system ϕ and a strategy α, is given by

$$\Pi = \rho\sum_{i=1}^{I} f_i^B\alpha_i + (1-\rho)\sum_{i=1}^{I} f_i^G\alpha_i.$$

For every rating i that is reachable, that is, $f_i^\theta > 0$ for some θ, the updated belief is given by

$$\rho_i = \frac{\rho f_i^B}{\rho f_i^B + (1-\rho)f_i^G}.$$

For any given rating system we define the equilibrium as a strategy α and beliefs $(\rho_i)_{i\in\mathscr{I}}$ such that every investor is taking her optimal action given her beliefs. We state the concept below.

Definition 1.2. *Given a rating system ϕ, an equilibrium is a strategy α as well as beliefs $(\rho_i)_{i\in\mathscr{I}}$ such that, for every $i \in \mathscr{I}$:*

1. $\alpha_i > 0 \Leftrightarrow \rho_i\sum_{m=1}^{M}\gamma_m^B x_m + (1-\rho_i)\sum_{m=1}^{M}\gamma_m^G x_m \geq 0;$
2. ρ_i *is consistent whenever reachable (that is, derived from Bayes' rule as well as the stationary distributions generated by ρ, α, and ϕ).*

For now, we focus on irreducible rating systems, i.e., systems in which all ratings are visited for at least one of the project's type. This means that all ratings will have induced beliefs derived from Bayes's rule as well as that the invariant distributions will not depend on the initial distribution.[16] We will show later on that this focus on irreducible systems is without loss of generality.

In order to better understand the mechanics of our model and how to compute beliefs, let us consider the following example.

Example 1.1. *Consider a binary payoff environment in which there is a good realization, denoted by h, and a bad realization, denoted by l. Then, assume that $\Pr(h|G) = \gamma_h > \frac{1}{2}$, and $\Pr(h|B) = \gamma_l < \frac{1}{2}$. Construct a rating system with two ratings and a deterministic transition rule. That is, $\varphi_{21}^h = 0, \varphi_{21}^l = 1$, and $\varphi_{12} = 1$. Below we depict this system for the purpose of illustration.*

[15] For the proof see, for example, Stokey and Lucas (1989), Theorem 11.2.
[16] See Stokey and Lucas (1989), Theorem 11.4.

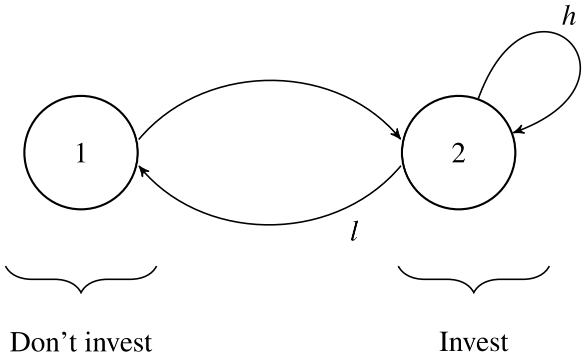

Figure 1.1 – A binary, deterministic and irreducible system in which agents only invest in rating 2.

For such a system, the induced Markov matrices are given by

$$T^G = \begin{pmatrix} 0 & 1 \\ 1-\gamma_h & \gamma_h \end{pmatrix} \text{ and } T^B = \begin{pmatrix} 0 & 1 \\ 1-\gamma_l & \gamma_l \end{pmatrix}.$$

Due to irreducibility, we compute stationary distribution by solving $f^\theta \cdot T^\theta = f^\theta$*, leading to*

$$f_1^G = \frac{1-\gamma_h}{2-\gamma_h}, \qquad f_1^B = \frac{1-\gamma_l}{2-\gamma_l}.$$

Finally, beliefs are given by

$$\rho_1 = \frac{\rho\,(1-\gamma_l)\,(2-\gamma_h)}{\rho\,(1-\gamma_l)\,(2-\gamma_h)+(1-\rho)\,(1-\gamma_h)\,(2-\gamma_l)};$$
$$\rho_2 = \frac{\rho\,(2-\gamma_h)}{\rho\,(2-\gamma_h)+(1-\rho)\,(2-\gamma_l)}.$$

This rating system is irreducible and satisfies the obedience constraints (i.e., it is incentive compatible to invest in rating 2 and to not invest in rating 1) if $\rho_2 \leq \rho^*$ *and* $\rho_1 > \rho^*$*. We note that this is just an example of rating system, but it is generically not the optimal one.*

1.3 Optimal rating system

In this section, we derive a bound on the belief spread induced by designer's constrained information rule. This in turn leads to an upper bound on his optimal value of information, which is proved in theorem 1.1. We contrast these results with what the designer could get if he could condition messages on hidden information; specifically, on the project's type. We also discuss how the informativeness of the payoffs plays a role in designer's value of information.

Without loss of generality, we can label the ratings such that $\rho_1 \geq \rho_2 \geq ... \geq \rho_I$. This implies that

$$\frac{f_1^B}{f_1^G} \geq \frac{f_2^B}{f_2^G} \geq ... \geq \frac{f_I^B}{f_I^G}.$$

For what follows, it will also be convenient to define

$$\lambda = \frac{\rho^*}{1-\rho^*}\frac{1-\rho}{\rho},$$

where note that $\lambda < 1$ if $\rho > \rho^*$. If i is a rating in which there is investment, then it must be that $\rho_i \leq \rho^*$. Note as well that $\rho_i \leq \rho^*$ if and only if $\frac{f_i^B}{f_i^G} \leq \lambda$. This represents the prior bias against investment. Let $\mathcal{N}$ represent the set of non-investment ratings. We can write incentive compatibility inequalities as:

$$\frac{f_i^B}{f_i^G} > \lambda \;\; \forall\, i \in \mathcal{N}; \tag{1.1}$$

$$\frac{f_i^B}{f_i^G} \leq \lambda \;\; \forall\, i \in \mathcal{N}^{\complement}. \tag{1.2}$$

We will now work out a third restriction for equilibria with at least one investment rating. We will show that there is a maximum spread in the induced beliefs. The intuition for the fact that beliefs cannot be too far apart is that if the designer tries to construct a rating that induces a belief that is too high (meaning that it is likely that the project is bad), it will affect the beliefs in the other rating, violating the obedience constraint (incentives for investing) in the other rating. The next proposition formalizes this intuition, but before we state it, we need some additional notation. Let us define

$$\eta = \max_m \frac{\gamma_m^B}{\gamma_m^G}.$$

This represents the strength of bad news, or the informativeness of the payoff with the highest likelihood between the bad and the good state of the world. For the remainder of this section, let us assume that $\eta < \infty$, so that all payoffs that occur with positive probability under state B also occur with positive probability under state G. Note that $\eta > 1$. The case with $\eta = \infty$ will be discussed separately (see example 1.2). It will also be useful to define

$$\nu = \min_m \frac{\gamma_m^B}{\gamma_m^G}.$$

As we will see, perhaps surprisingly, our results do not depend on whether ν is zero or positive. We will postpone this discussion. Before we prove our proposition, let us state the following well-known result for Markov processes:

Lemma 1.1. *If the space $\mathscr{I}$ is partitioned into two sets S and S', then in the steady state the probability of transition from S to S' must equal the probability of transition from S' to S. Precisely stated, for each θ,*

$$\sum_{i\in S}\sum_{j\in S'} f_i^\theta \tau_{i,j}^\theta = \sum_{i\in S'}\sum_{j\in S} f_j^\theta \tau_{i,j}^\theta .$$

Proposition 1.1. *In any rating system with investment, all non-investment ratings obey*

$$\frac{f_i^B}{f_i^G} \leq \lambda\eta .$$

Proof. Let $\mathscr{N}_1$ be the set of non-investment ratings i such that $\frac{f_i^B}{f_i^G} > \lambda\eta$ and $\mathscr{N}_2$ the set of non-investment ratings such that $\frac{f_i^B}{f_i^G} \leq \lambda\eta$. We want to show that $\mathscr{N}_1 = \emptyset$. Assume by way of contradiction that $\mathscr{N}_1$ is non-empty. With a slight abuse of notation, let $\tau_{i,S}^\theta$ be the transition chance from rating i to a set of ratings S conditional on $\theta \in \{B,G\}$. For every non-investment rating, we eliminate the superscript θ because the transition chance will be uninformative. From lemma 1.1, the steady state probabilities obey the following:

$$\frac{\sum_{i\in\mathscr{N}_1} f_i^B \left(\tau_{i,\mathscr{N}_2} + \tau_{i,\mathscr{N}^\complement}\right)}{\sum_{i\in\mathscr{N}_1} f_i^G \left(\tau_{i,\mathscr{N}_2} + \tau_{i,\mathscr{N}^\complement}\right)} = \frac{\sum_{i\in\mathscr{N}_2} f_i^B \tau_{i,\mathscr{N}_1} + \sum_{i\in\mathscr{N}^\complement} f_i^B \tau_{i,\mathscr{N}_1}^B}{\sum_{i\in\mathscr{N}_2} f_i^G \tau_{i,\mathscr{N}_1} + \sum_{i\in\mathscr{N}^\complement} f_i^G \tau_{i,\mathscr{N}_1}^G} .$$

The LHS of the above equation exceeds $\lambda\eta$ by assumption since it is a weighted average of the ratios $\frac{f_i^B}{f_i^G}$, $i \in \mathscr{N}_1$ - each weight is $\frac{f_i^G(\tau_{1,\mathscr{N}_2}+\tau_{i,\mathscr{N}^\complement})}{\sum_{i\in\mathscr{N}_1} f_i^G(\tau_{1,\mathscr{N}_2}+\tau_{i,\mathscr{N}^\complement})}$. If the RHS is at most $\lambda\eta$, a contradiction follows. But the RHS is a weighted average of ratios $\frac{f_i^B}{f_i^G}$, $i \in \mathscr{N}_1$, $i \in \mathscr{N}_2$ and ratios $\frac{f_i^B \tau_{i,\mathscr{N}_1}^B}{f_i^G \tau_{i,\mathscr{N}_1}^G}$, $i \in \mathscr{N}^\complement$. The former ratios are at most $\lambda\eta$ by construction and the latter ratios are at most $\lambda\eta$ by the incentive compatibility constraint (investment requires $\frac{f_i^B}{f_i^G} \leq \lambda$ for every $i \in \mathscr{N}^\complement$) and information constraint (investment payoff likelihoods satisfy $\frac{\tau_{i,\mathscr{N}_1}^B}{f_{i,\mathscr{N}_1}^G} \leq \lambda$ for every $i \in \mathscr{N}^\complement$). □

Proposition 1.1 implies that in all non-investment ratings, agents will hold induced beliefs about a bad project no higher than $\bar{\rho}$, defined by

$$\bar{\rho} = \frac{\eta\rho^*}{1-\rho^*+\eta\rho^*} .$$

We also derive an upper bound on information designer's payoff and prove that this upper bound can be achieved, provided that the rating system induces all investment beliefs to be the lowest possible (that is, equal to ρ^*) and all non-investment beliefs to be highest possible (equal to $\bar{\rho}$). This is done in theorem 2.1 below.

Theorem 1.1. *For any rating system, there exists an upper bound on the information designer's payoff:*

$$\Pi \leq \frac{(1-\rho)}{(1-\rho^*)}\frac{\lambda\eta - 1}{\lambda(\eta-1)}.$$

This upper bound is achieved iff beliefs in all non-investment ratings equal $\bar{\rho}$ and beliefs in all investment ratings equal ρ^*.

Proof. Incentive compatibility of investing in $\mathscr{N}^{\complement}$ implies

$$\frac{\sum_{i\in\mathscr{N}^{\complement}} f_i^B}{\sum_{i\in\mathscr{N}^{\complement}} f_i^G} = \frac{1-\sum_{i\in\mathscr{N}} f_i^B}{1-\sum_{i\in\mathscr{N}} f_i^G} \leq \lambda.$$

Therefore,

$$\sum_{i\in\mathscr{N}} f_i^B \geq 1-\lambda+\lambda\sum_{i\in\mathscr{N}} f_i^G. \tag{1.3}$$

By proposition 1.1,

$$\frac{\sum_{i\in\mathscr{N}} f_i^B}{\sum_{i\in\mathscr{N}} f_i^G} \leq \lambda\eta.$$

Therefore,

$$\sum_{i\in\mathscr{N}} f_i^G \geq \frac{1}{\lambda\eta}\sum_{i\in\mathscr{N}} f_i^B. \tag{1.4}$$

Substituting equation 1.4 into equation 1.3 yields

$$\sum_{i\in\mathscr{N}} f_i^B \geq \eta\left[\frac{1-\lambda}{\eta-1}\right]. \tag{1.5}$$

Substituting equation 1.5 back into equation 1.4 yields

$$\sum_{i\in\mathscr{N}} f_i^G \geq \frac{1}{\lambda}\left[\frac{1-\lambda}{\eta-1}\right]. \tag{1.6}$$

By equation 1.5 and equation 1.6, the information designer's payoff - the steady state investment probability - has an upper bound, as it can be seen below.

$$\rho\left(1-\sum_{i\in\mathcal{N}}f_i^B\right)+(1-\rho)\left(1-\sum_{i\in\mathcal{N}}f_i^G\right)\leq\rho\left(1-\eta\left[\frac{1-\lambda}{\eta-1}\right]\right)+(1-\rho)\left(1-\frac{1}{\lambda}\left[\frac{1-\lambda}{\eta-1}\right]\right).$$

Using $\lambda\equiv\frac{\rho^*}{1-\rho^*}\frac{1-\rho}{\rho}$, this simplifies to the desired upper bound. Such upper bound is achieved if and only if equations 1.3 and 1.4 both hold with equality. Finally, 1.3 holds with equality if and only if every investment rating i has $\frac{f_i^B}{f_i^G}=\lambda$, which holds iff ρ_i is the exactly cutoff belief for investment, i.e. $\rho_i=\rho^*$. Likewise, 1.4 holds with equality if and only if every non-investment rating i has $\frac{f_i^B}{f_i^G}=\lambda\eta$, which holds if and only if ρ_i is exactly the upper bound belief for non-investment, i.e. $\rho_i=\bar{\rho}$.

□

As an implication, we can reproduce the optimal rating system with I ratings by a two-rating system that induces beliefs $\rho_1=\bar{\rho}$ and $\rho_2=\rho^*$. Thus, a simple recommendation of investing or not is sufficient to implement the optimal rule. But before constructing such system, let us discuss three main differences from the results from our model with results from the unconstrained information design literature, specifically the Bayesian persuasion framework.

First, if the designer could commit to a distribution of messages conditional on the project's type, it would also be sufficient for him to induce only two beliefs in equilibrium: one equal to the cutoff for investment and the other as strongly as possible in favor of not investing. As agents only rely on designer's message, this would be the optimal way to maximize the overall probability of investment. The designer could induce these beliefs by (i) recommending investment for sure when the type is G and (ii) randomizing between an investment and a non-investment recommendation when the type is B.

However, the reason for the optimality of a binary information in our model differs from the one in the Bayesian persuasion framework. More specifically, our result is *not* a revelation principle. It is instead a consequence of the fact that (i) in non-investment ratings no hard data is produced and posterior beliefs are induced through such data; (ii) in investment ratings there is no need to provide additional incentives through more information. Indeed, we will show in the next section that, for the case in which observable data is produced even without investment, the designer's payoff is increasing in the number of ratings.

Second, an unconstrained designer has power to induce the strongest belief in favor of not investing, that is, the extreme belief 1. Abusing notation, let f_i^θ represent the probability that this unconstrained designer sends message $i\in\{1,2\}$ conditional on type θ and $i=2$ be the recommendation to invest (the investment rating). Then he would induce

$$f_2^B=\lambda; \qquad\qquad f_2^G=1.$$

Leading to $\frac{f_2^B}{f_2^G}=\lambda$ and $\frac{f_1^B}{f_1^G}=\infty$. From proposition 1.1, our constrained designer cannot obtain the ratio of rating 1. As a result, the designer's investment payoff Π^{BP} from Bayesian persuasion is higher than designer's constrained payoff Π from theorem 1.1:

$$\Pi^{BP} = \rho\lambda + (1-\rho) = \frac{1-\rho}{1-\rho^*} > \frac{1-\rho}{1-\rho^*}\left[\frac{\lambda\eta - 1}{\lambda(\eta-1)}\right] = \Pi.$$

Nevertheless, as $\eta \to \infty$, Π converges to Π^{BP}. This means that the relevant statistic for the designer in the likelihood η. Intuitively, a perfectly informative signal (payoff) for the bad state does make it possible to fully convince the agents that the project is bad. Note that, because the upper bound on proposition 2.1 does not depend on ν, a very informative signal (payoff) towards the good state of the world adds nothing to the designer's payoff.

Third, in the Bayesian persuasion framework, for every initial prior $\rho \in (\rho^*, 1)$, the information designer can induce posteriors leading to the optimal value of information at ρ. This is true because the optimal value of information is the concave closure of designer's payoff without any communication.[17] This is not the case in the constrained framework. To better understand this, we need the following result.

Lemma 1.2. *In any rating system with investment, at least one rating has a belief (weakly) lower than the original prior, and at least one rating must have a belief (weakly) higher than the prior. Formally* $\rho_i \leq \rho$ *and* $\rho_j \geq \rho$ *for some* $i, j \in \mathscr{I}$.

Proof. Suppose not. That is, suppose that $\rho_i > \rho$, $\forall i$. We can write: $\frac{\rho f_i^B}{\rho f_i^B + (1-\rho) f_i^G} > \rho$, which implies that $f_i^B > f_i^G$,$\forall i$. If we sum for all i, we have that $\sum_i f_i^B > \sum_i f_i^G$. However, $\sum_i f_i^B = \sum_i f_i^G = 1$, so we have a contradiction. □

A corollary of this result is that whenever $\rho > \bar{\rho}$ there is no equilibrium in a rating system in which there is investing in at least one rating and no investing in at least other rating. From now on, we will focus on this more interesting case, that is,

Assumption 1.1. *The prior belief about the bad type is such that* $\rho \leq \bar{\rho}$.

The figures below summarize the comparison between our constrained framework and the persuasion framework. In the left figure, the dashed blue line represents additional payoff designer gets from Bayesian persuasion, relative to no intervention. In the right figure, the dashed red line represents the additional payoff designer gets from an optimal rating system. As $\eta \to \infty$, the prior belief region for which there is an equilibrium with a rating system expands and the payoff converges to the Bayesian persuasion payoff.

We now turn to the construction of the optimal rating system. We do this by defining transition rules so that f_1^B and f_1^G are as low as possible. Applying equation 1.3 for only one non-investment rating, we have

$$f_1^B \geq 1 - \lambda + \lambda f_1^G \tag{1.7}$$

[17] That is, the point-wise infimum of affine functions that are weakly higher than the no-communication payoff.

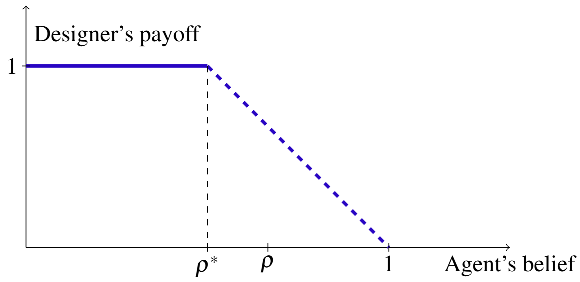

Figure 1.2 – Bayesian persuasion payoff

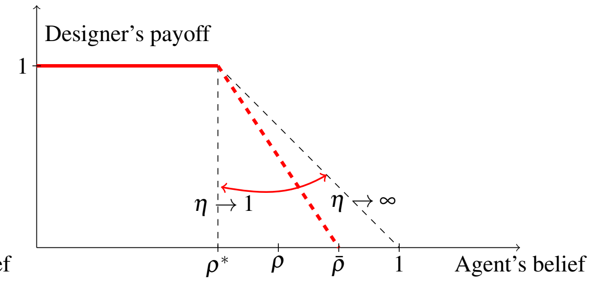

Figure 1.3 – Rating system payoff

Because $\lambda < 1$, if the steady state probabilities satisfy equation 1.7, then the non-investment compatibility constraint 1.1 is satisfied and we can ignore it. From the investment compatibility constraint 1.2 and the constraint from proposition 1.1, we derive

$$f_1^B\left(1-f_1^G\right) \leq f_1^G\eta\left(1-f_1^B\right) \Rightarrow f_1^B \leq \frac{\eta f_1^G}{1+f_1^G(\eta-1)}. \tag{1.8}$$

The pair of stationary distributions $\left(f_1^B, f_1^G\right)$ that maximizes the probability of investing is obtained when equations 1.7 and 1.8 are binding. Thus, the distributions must satisfy

$$1-\lambda+\lambda f_1^G = \frac{\eta f_1^G}{1+f_1^G(\eta-1)}.$$

Indeed, we can see from the figure below that there are two points that satisfy the above equation, namely, when $f_1^G = f_1^B = 1$ (although this does not interest us), and another point at which $f_1^B < 1$ and $f_1^G < 1$.

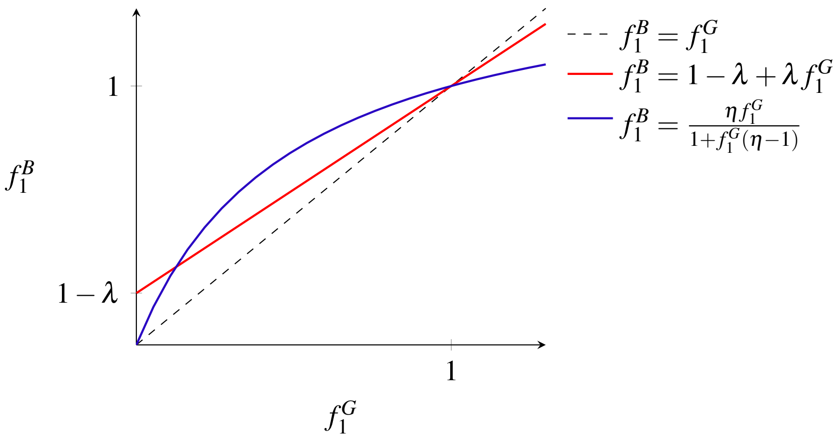

Figure 1.4 – Optimal stationary probabilities. The blue represents equation 2.8 and red line represents equation 2.7. The optimal values of f_1^B and f_1^G is the interior point in which the colored lines intersect.

In the next proposition, we construct rating systems that achieve this upper bound. The particular choice of transition rules will depend on the parameters, but it is possible to construct the optimal system for any given set of parameters.

Proposition 1.2. *For any set of parameters, there exists a binary rating that achieve the payoff bound from theorem 1.1.*

Proof. Let m be such that $\frac{\gamma_m^B}{\gamma_m^G} = \eta$. From rating 1 there is a random exit probability $\varphi_{12} = \tau$ and from rating 2 the transition rule is: (i) $\varphi_{21}^m = \kappa$ and (ii) $\varphi_{21}^n = 0 \; \forall n \neq m$. Then, these transitions induce the following Markov transition matrices:

$$T^B = \begin{pmatrix} 1-\tau & \tau \\ \gamma_m^B \kappa & (1-\gamma_m^B) + \gamma_m^B (1-\kappa) \end{pmatrix} \text{ and } T^G = \begin{pmatrix} 1-\tau & \tau \\ \gamma_m^G \kappa & (1-\gamma_m^G) + \gamma_m^G (1-\kappa) \end{pmatrix}.$$

Because we are focusing on irreducible systems, we can compute the stationary distributions by solving $f^\theta \cdot T^\theta = f^\theta$ for each $\theta \in \{B, G\}$. This leads to

$$f_2^B = \frac{\tau}{\tau + \gamma_m^B \kappa}; \qquad f_2^G = \frac{\tau}{\tau + \gamma_m^G \kappa}.$$

We will choose τ and k appropriately so that $\frac{f_2^B}{f_2^G} = \lambda$. This gives us $\frac{\tau + \gamma_j^G k}{\tau + \gamma_j^B k} = \lambda$, which in turn implies that:

$$\frac{\tau}{\kappa} = \frac{\lambda \gamma_m^B - \gamma_m^G}{1-\lambda}.$$

Recall from assumption 1 that we are only interested in the case in which $\rho \leq \bar{\rho}$, which implies that $\lambda \eta > 1$, and in turn implies that $\lambda \gamma_m^B - \gamma_m^G > 0$. We can conclude that both denominator and numerator are between 0 and 1. Thus, we can set $\tau = \lambda \gamma_m^B - \gamma_m^G$ and $\kappa = 1 - \lambda$. This leads to the desired ratios $\frac{f_2^B}{f_2^G} = \lambda$ and $\frac{f_1^B}{f_1^G} = \eta \lambda$. □

The optimal rating system derived in proposition 1.2 is illustrated in the figure below.

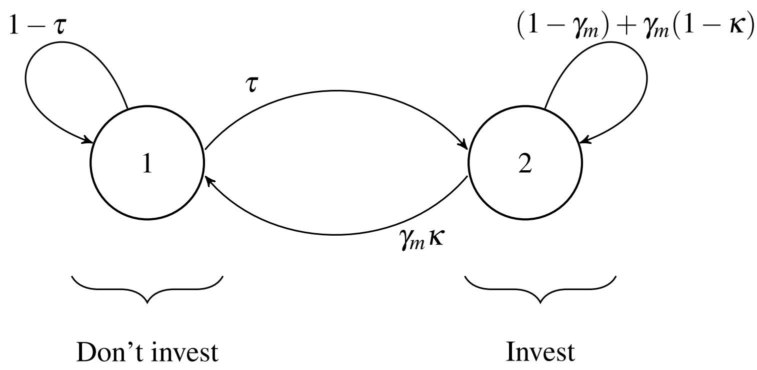

Figure 1.5 – Optimal rating system. Only the payoff corresponding to the highest likelihood of project being bad matters. Parameters τ and κ are carefully chosen to generate the desired posterior beliefs.

We still have to deal with $\eta = \infty$. This is easily done, as it can be seen in example 1.2. In this case, the rating system achieves the Bayesian persuasion payoff.

Example 1.2. *Consider the case in which $\eta = \infty$, that is, there is a payoff that happens with positive probability under B, but not under G. Since in this case a bad state of the world can be fully learned, we might call this special case as the bad news case.*

A simple construction will suffice to achieve the optimal Bayesian persuasion payoff. Let x_m be a payoff for which $\gamma_m^B > 0$ and $\gamma_m^G = 0$. Consider a binary rating system with transitions given by $\varphi_{21}^m = \kappa$, $\varphi_{21}^n = 0$ for all $n \neq m$, and $\tau_{12} = \tau$, where, recall, that this last transition must be independent of the payoff realization, since on-equilibrium path, there is no investment in rating 1. With this rating system and this signal structure, the stationary distributions are given by $f_1^B = \frac{\gamma_m^B \kappa}{\tau + \gamma_m^B \kappa}$ and $f_1^G = 0$. It follows immediately that $\rho_1 = 1$. To get the Bayesian persuasion payoff, we need to choose τ and k such that $\rho_2 = \rho^*$. It must be that

$$\frac{\kappa}{\tau} = \frac{\rho - \rho^*}{\gamma_m^B \rho^* (1 - \rho)}.$$

Since both numerator and denominator are positive and strictly less than one, it suffices to set $\kappa = \rho - \rho^*$ and $\tau = \gamma_m^B \rho^* (1 - \rho)$.

Before concluding this section, we prove that restricting attention to irreducible systems was without loss of generality. First, note that a rating system that has a unique recurrent class (i.e., one with a unique absorbing set of ratings) must be such that eventually the system reaches that rating and remains there thereafter, regardless of the underlying state of the world. This implies that this smaller set of ratings - the set of recurrent ratings - works in a fashion similar to an irreducible system, and we have shown that for such systems it is sufficient to consider two ratings.

Our next task is to show that the multiple recurrent classes cannot all be solely composed of investment ratings. Suppose, by contradiction, that all recurrent classes are composed exclusively of ratings in which the recommendation is to invest. The system will eventually reach some of these ratings, and the belief of at least one of these ratings must be at least as high as the original prior, which we are assuming to be a non-investment prior. Thus, at least one recurrent class has at least one non-investment rating, which is a contradiction. This result is similar to the one in lemma 1.2.

In recurrent classes with both investment and non-investment ratings, we repeat the analysis of our general rating system with irreducible ratings. Thus, in each recurrent class, it must be that beliefs are bounded as shown in proposition 1. Thus, the designer cannot improve upon a two-rating machine. Finally, it remains to show that any system with a recurrent class composed exclusively of non-investment ratings does not improve on our general irreducible system. This is shown in proposition 2.3, whose proof is in the appendix.

Proposition 1.3. *A reducible system with a recurrent class of non-investment ratings cannot improve upon the optimal irreducible system.*

1.4 No experimentation

In this section we look at the case where hard data is independent of actions. That is, if the agent does not invest, she gets a payoff of zero, but there is still an observable realization from the set X following the same distribution as specified in section 2.2. We are interested in this environment since it helps us understand the forces driving our main results in the previous section.

We show there exists an upper bound on beliefs than can be achieved in equilibrium, but this bound is increasing in both the number of ratings and in the strength of good news. Then we derive an upper bound on designer's payoff and construct a rating system that gets arbitrarily close to it. The construction will work for any prior $\rho \in (\rho^*, 1)$ - not just $\rho \leq \bar{\rho}$ from assumption 1 - provided that the number of ratings is big enough. Finally, we show this rating also approximates the Bayesian persuasion payoff.

As before, we label ratings in a decreasing belief order, that is, $\rho_1 \geq \rho_2 \geq \ldots \geq \rho_I$. This implies that the likelihood ratios of the steady state probabilities are decreasing as well. Hellman and Cover (1970) provided the following two lemmas[18], which will be useful in this section. The first derives a bound on the ratio of the transition rules. The second derives an upper bound on the likelihood ratios from rating i to rating $i+1$.

Lemma 1.3. *The ratio of the transition from rating i to rating j under state of the world B to the same transition under G satisfies*

$$\eta \geq \frac{\tau_{ij}^B}{\tau_{ij}^G} \geq \nu.$$

Lemma 1.4. *For every rating $i \leq I-1$,*

$$\frac{f_i^B}{f_i^G} \leq \frac{\eta}{\nu} \frac{f_{i+1}^B}{f_{i+1}^G}.$$

Proposition 1.4. *In any rating system with investment, it must be that*

$$\frac{f_1^B}{f_1^G} \leq \left(\frac{\eta}{\nu}\right)^{I-1} \lambda.$$

Proof. Because there must be investment in at least one rating, it is true that $\frac{f_I^B}{f_I^G} \leq \lambda$. From lemma 1.4, it is also true that $\frac{f_{I-1}^B}{f_{I-1}^G} \leq \left(\frac{\eta}{\nu}\right)\lambda$. Proceeding recursively, we find that $\frac{f_1^B}{f_1^G} \leq \left(\frac{\eta}{\nu}\right)^{I-1}\lambda$, as desired. □

Proposition 1.2 implies an upper bound on the highest belief induced in equilibrium, but now this bound is increasing in the number of ratings I and decreasing in the value of ν - the strength of good news. Indeed, this bound is given by

[18] Lemmas 3 an 4 in our paper are lemmas 1 and 2 in Hellman and Cover (1970), respectively.

$$\hat{\rho} = \frac{\eta^{I-1}\rho^*}{\eta^{I-1}\rho^* + \nu^{I-1}(1-\rho^*)}.$$

Note that $\hat{\rho} \to 1$ when $\eta \to 1$, a result that is similar to the one we had in the previous section. However, we also have $\hat{\rho} \to 1$ when $\nu \to 0$ or when $I \to \infty$. The parameter ν affects the belief bound in a non-experimentation environment because induced beliefs in non-investment ratings need not be all the same anymore. Thus, a high belief in non-investment rating is possible even if transitions from good ratings to bad ratings are not very informative, provided that I is big enough.

The number of ratings affects the belief bound because every non-investment rating $i \leq I-1$ can keep track of the history of signals (payoffs) up to $I-1$. Therefore, the higher the number of ratings, the easier it is to convince the agents that only the bad type of project can visit infinitely often the first ratings, by assigning downgrades to more informative signals (payoffs) about the state being bad and upgrades to more informative signals (payoffs) about state being good.

Theorem 2.2 below is the equivalent result of theorem 1.1 for this section.

Theorem 1.2. *For any rating system with some rating i^* such that there is no investment in any rating $i \leq i^*-1$ and investment in any rating $i \geq i^*$, there exists an upper bound on the information designer's payoff:*

$$\Pi \leq \frac{1-\rho}{1-\rho^*}\left[\frac{\lambda - \left(\frac{\nu}{\eta}\right)^{i^*-1}}{\lambda\left(1-\left(\frac{\nu}{\eta}\right)^{i^*-1}\right)}\right].$$

Proof. As before, $\mathcal{N}$ refers to the set of non-investment ratings and $\mathcal{N}^{\complement}$ to the set of investment ratings. Let $i^* = \min\{i : i \in \mathcal{N}^{\complement}\}$. Repeating the same steps of the proof of theorem 1, but now considering the bound on every $\frac{f_i^B}{f_i^G}$ from proposition 2.4; in particular, $\sum_{i\in\mathcal{N}} f_i^G \geq \frac{1}{\lambda}\left(\frac{\nu}{\eta}\right)^{i^*-1}\sum_{i\in\mathcal{N}} f_i^B$, the information designer's payoff has the following upper bound:

$$\Pi \leq \rho\left(1-\left[\frac{1-\lambda}{1-\left(\frac{\nu}{\eta}\right)^{i^*-1}}\right]\right) + (1-\rho)\left(1-\frac{(\frac{\nu}{\eta})^{i^*-1}}{\lambda}\left[\frac{1-\lambda}{1-\left(\frac{\nu}{\eta}\right)^{i^*-1}}\right]\right).$$

Using $\lambda \equiv \frac{\rho^*}{1-\rho^*}\frac{1-\rho}{\rho}$, this simplifies to the desired upper bound. □

To achieve this upper bound, it must be that every investment rating i has $\frac{f_i^B}{f_i^G} = \lambda$, which holds if and only ρ_i is the exactly cutoff belief for investment, i.e. $\rho_i = \rho^*$. Therefore, it is sufficient to have only one investment rating with an indifference belief, that is, $\mathcal{N}^{\complement} = \{I\}$ and $\rho_I = \rho^*$. However, it is not true anymore that all non-investment ratings must have the same induced beliefs. To see this, note that, from

lemma 2.4, for every $i \leq I-1$ without investment, steady state likelihood ratios satisfy $\frac{f_i^B}{f_i^G} \leq \left(\frac{\eta}{\nu}\right)^{I-i}\lambda$. If, say, $\frac{f_2^B}{f_2^G} = \frac{f_1^B}{f_1^G}$, then such constraint would be violated, because

$$\frac{f_2^B}{f_2^G} = \left(\frac{\eta}{\nu}\right)^{I-1}\lambda > \left(\frac{\eta}{\nu}\right)^{I-2}\lambda,$$

as $\eta > \nu$. Nevertheless, we can construct a rating system that gets arbitrarily close to the upper bound derived in theorem 1.2 by making intermediary steady state probabilities arbitrarily close to zero - that is, $f_I^B + f_I^G = f_1^B + f_1^G \approx 1$ - and extreme beliefs close to the optimal bounds - that is, $\rho_I \approx \rho^*$ and $\rho_1 \approx \hat{\rho}$.

The construction of such system is done in the next proposition and the proof of it is in the appendix, but we give an intuition here.

We first choose $\bar{I}$ high enough so that $\rho \leq \hat{\rho}$. Note that we can do this for every prior $\rho \in (\rho^*, 1)$ - not just $\rho \leq \bar{\rho}$, so we can have an equilibrium even without assumption 1. Then, we partition the payoff space in two subsets such that the likelihood of aggregated payoffs in one set is higher than the likelihood of aggregated payoffs in the other set. These aggregated payoffs will work as a proxy for two relevant statistics for the transition rules. The constructed rating system for given $\bar{I}$ is such that: (i) at intermediate ratings, every signal in one payoff subset leads to a downgrade and every signal in another payoff subset leads to an upgrade; (ii) at extreme ratings 1 and $\bar{I}$, upgrades and downgrades are governed by parameters τ and κ, respectively. Those parameters are chosen appropriately so that even if $\tau \approx 0$ and $\kappa \approx 0$, the belief in $\bar{I}$ is kept at ρ^* and the belief in 1 is well defined. For a number $I \geq \bar{I}$ large enough, the designer's payoff from such system will be close enough to the bound derived in theorem 2.2.

Because the upper bound on designer's payoff in that theorem is increasing in the number of ratings and converges to the Bayesian persuasion payoff as I grows large, a corollary of proposition 1.5 below is that the constructed system also approximates the Bayesian persuasion payoff $\frac{1-\rho^*}{1-\rho}$. In that case, extreme induced beliefs will be such that $\rho_I \to \rho^*$ and $\rho_1 \to 1$ as $I \to \infty$.

Proposition 1.5. *For any set of parameters and for every $\varepsilon > 0$, there exists a number $\bar{I}$ and a rating system with $I \geq \bar{I}$ ratings that is ε-close to the payoff bound from theorem 1.2, that is,*

$$\Pi \geq \frac{1-\rho}{1-\rho^*}\left[\frac{\lambda - \left(\frac{\nu}{\eta}\right)^{I-1}}{\lambda\left(1-\left(\frac{\nu}{\eta}\right)^{I-1}\right)}\right] - \varepsilon.$$

In such system, intermediary ratings are almost never visited, that is, $f_1^\theta + f_I^\theta \approx 1$ for $\theta \in \{B, G\}$.

1.5 Altruistic designer

We argued before that the reason for the optimality of two ratings is a consequence of the fact that (i) no hard data is produced in non-investment ratings and (ii) there is no need to produce more information in investment ratings. In previous section, we showed that the number of ratings matters indeed when (i) is relaxed. In this section, we show this will also be the case when (ii) is relaxed. We do this by assuming the designer is now a benevolent social planner, so he cares about agents learning the state of the world.

Let us return to the case in which experimentation is needed to generate information. In it, the agents' *ex-ante* payoff under designer's optimal rating system can be written as follows:

$$\begin{aligned}
\Pi^A &= \rho f_2^B \sum_{m=1}^{M} \gamma_m^B x_m + (1-\rho) f_2^G \sum_{m=1}^{M} \gamma_m^G x_m, \\
&= f_2 \left[\rho_2 \sum_{m=1}^{M} \gamma_m^B x_m + (1-\rho_2) \sum_{m=1}^{M} \gamma_m^G x_m \right].
\end{aligned}$$

We get the second equality by dividing and multiplying Π^A by $f_2 = \rho f_2^B + (1-\rho) f_2^G$. We know from theorem 1.1 that at rating 2, the agent is indifferent between investing and not investing. Therefore, $\Pi^A = 0$. Consider now the case in which the designer is altruistic. That is, his payoff is equivalent to the *ex-ante* payoff of the agents, which we write as

$$\Pi = \Pi^A = \rho \sum_{i \in \mathcal{N}^\complement} f_i^B \sum_{m=1}^{M} \gamma_m^B x_m + (1-\rho) \sum_{i \in \mathcal{N}^\complement} f_i^G \sum_{m=1}^{M} \gamma_m^G x_m.$$

Ideally, the altruistic designer would like to have extreme beliefs - zero for investment ratings and one for non-investment ratings. However, from proposition 1.1 we know that the highest belief is bounded above $\bar{\rho}$ and from theorem 1.1 we know that the altruistic designer's expected payoff will also have a bound. This bound is given in theorem 2.3 below.

Theorem 1.3. *For any $\rho \leq \bar{\rho}$, there exists an upper bound on the altruistic designer's payoff:*

$$\Pi \leq \sum_{m=1}^{M} \gamma_m^G x_m \left[\frac{\bar{\rho} - \rho}{\bar{\rho}} \right].$$

Proof. This follows from

$$\begin{aligned}
\Pi &= \rho \sum_{i \in \mathcal{N}^\complement} f_i^B \sum_{m=1}^{M} \gamma_m^B x_m + (1-\rho) \sum_{i \in \mathcal{N}^\complement} f_i^G \sum_{m=1}^{M} \gamma_m^G x_m, \\
&= \sum_{i \in \mathcal{N}^\complement} f_i \left[\rho_i \sum_{m=1}^{M} \gamma_m^B x_m + (1-\rho_i) \sum_{m=1}^{M} \gamma_m^G x_m \right], \\
&\leq \sum_{m=1}^{M} \gamma_m^G x_m \left[\sum_{i \in \mathcal{N}^\complement} f_i \right].
\end{aligned}$$

The second line follows from the definition of f_i and ρ_i. The third line follows from $\sum_{m=1}^{M} \gamma_m^G x_m$ being the upper bound for the expression in brackets. Such value is maximal if and only if every investment rating i has ρ_i. However, note that every irreducible rating system satisfies[19]

$$\sum_{i \in \mathscr{I}} f_i \rho_i = \rho. \tag{1.9}$$

Using equation 1.9 but considering the split into non-investment ratings $i \in \mathscr{N}$ - for which $\rho_i \leq \bar{\rho}$ from proposition 2.1 - and investment ratings $\mathscr{N}^{\complement}$ - for which we want to set $\rho_i = 0$, we get

$$\rho = \sum_{i \in \mathscr{N}^{\complement}} f_i \rho_i \leq \bar{\rho} \left[\sum_{i \in \mathscr{N}} f_i \right] \Rightarrow \sum_{i \in \mathscr{N}} f_i \geq \frac{\rho}{\bar{\rho}}.$$

Therefore,

$$\sum_{i \in \mathscr{N}^{\complement}} f_i \leq \frac{\bar{\rho} - \rho}{\bar{\rho}}.$$

Substituting the above inequality in the inequality derived for the payoff Π at the beginning of the proof, we have the desired result. $\square$

This is best seen in the figure below. The solid blue line represents the expected payoff an altruistic designer would get without any intervention, as a function of beliefs. The dashed blue represents what he would get if he could condition messages directly on the states of the world, that is, what he would get under a Bayesian persuasion framework. The dashed red line represents the payoff upper bound derived in previous theorem, under our constrained framework. Examining the figure, we can conclude that the altruistic designer and, consequently, the agents, are worse-off under transparency requirements.

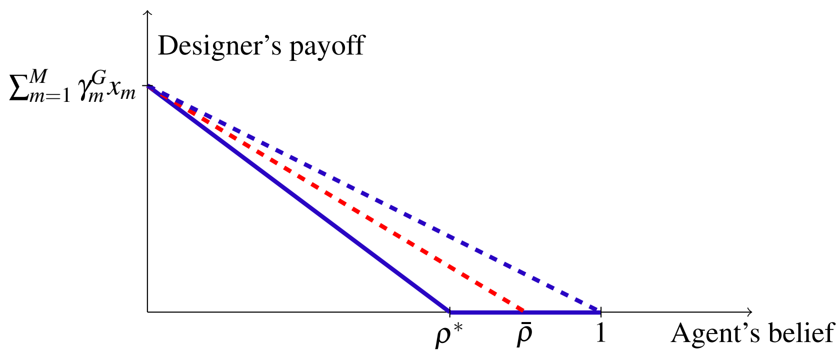

Figure 1.6 – Altruistic Designer

[19]This is analogous to the Bayes plausibility result in the Bayesian persuasion literature (Kamenica and Gentzkow, 2011).

We can construct a rating system that approximates the payoff upper bound given by theorem 1.3. The construction is very similar to the optimal rating system under no experimentation; therefore, intermediary ratings will be required to bring extreme beliefs ρ_1 and ρ_I close to $\bar{\rho}$ and 0, respectively, but the system will stay most often on ratings 1 and I. The main difference now is that in rating 1, the exit rate must be independent of payoffs, since there is no investment in this state. The higher the number of ratings, closer this system will be to the upper bound.[20]

1.6 Conclusion

In many of our regular transactions, we rely on information intermediaries of some form. However, communication is often impaired by legal restrictions that impose a limited record keeping, or that rules must be transparent. Here, we consider a designer restricted to use communication rules that satisfy the Markov property of being dependent only on the current publicly available history.

We construct the optimal communication system under such restrictions, which we refer to as rating system, and we obtain one negative result and a set of positive results. Our negative result is that beliefs cannot be too far apart, limiting the scope of Bayesian persuasion. Our designer shares a common prior with the agents, so that the rating system must also be used for learning the project's correct type. When agents do not invest, nothing is learned. This need for learning generates a bound on how different a bad belief must be from a good belief, and this bound is what generates the maximum prior for which there is an informative equilibrium.

In our set of positive results, we show that (i) direct recommendation (viewed as a two-rating system) is the optimal design when experimentation is needed, but in the case of no experimentation, each rating matters; and (ii) for specific signal structures, simple rules can approximate Bayesian persuasion. Our results show that these communication restrictions decrease the designer's payoff, but do not benefit the agents. Additionally, from our model we now know that bad-news type of signals are very useful for persuasion, but good-news signals are not.

1.7 Appendix

Proposition 1.3 (Proof)**.** *A reducible system with a recurrent class of non-investment ratings cannot improve upon the optimal irreducible system.*

Proof. Let us proceed in three main steps. First, let $\mathscr{R}_1$ be a recurrent class of $j-$ratings, all of which induce no investment. Then, it must be the case that $\rho_1 = \rho_2 = ... = \rho_j$. Since no investments are made in these ratings, the transitions among them are independent of the payoffs (because there are no payoffs). And given that this subset $\mathscr{R}_1$ is irreducible, the stationary distribution is independent of the initial distribution.

[20] We omit the construction of such system here, but it is available upon request.

The second step is to argue that the prior on these ratings must be at least as high as the highest posterior of the ratings out of $\mathscr{R}_1$ that transition into $\mathscr{R}_1$. Ratings that transition into $\mathscr{R}_1$ are either investment ratings or non-investment ratings. If they are non-investment ratings, their prior is at least the posterior of an investment rating that led to this rating; since there are no payoffs in those non-investment ratings, the priors in those ratings are the same as the posteriors.

We also know that in any investment rating, the highest prior possible is ρ^* and thus, the highest possible posterior out of this rating is

$$\frac{\rho \sum_{m=1}^{M} \gamma_m^B \varphi_{21}^m}{\rho \sum_{m=1}^{M} \gamma_m^B \varphi_{21}^m + (1-\rho) \sum_{m=1}^{M} \gamma_m^G \varphi_{21}^m} \leq \frac{\eta \rho^*}{1+\rho^*(\eta-1)} = \bar{\rho}.$$

Given that the prior in $\mathscr{R}_1$ must be at least as high as the highest posterior of ratings transitioning into it. Thus, the maximum belief that a rating system can induce in equilibrium is $\rho_1 \leq \bar{\rho}$. This concludes the proof. □

Proposition 1.5. *For any set of parameters and for every $\varepsilon > 0$, there exists a number $\bar{I}$ and a rating system with $I \geq \bar{I}$ ratings that is ε-close to the payoff bound from theorem 3.2, that is,*

$$\Pi \geq \frac{1-\rho}{1-\rho^*} \left[\frac{\lambda - \left(\frac{\nu}{\eta}\right)^{I-1}}{\lambda \left(1 - \left(\frac{\nu}{\eta}\right)^{I-1}\right)} \right] - \varepsilon.$$

In such system, intermediary ratings are almost never visited, that is, $f_1^\theta + f_I^\theta \approx 1$ for $\theta \in \{B, G\}$.

Proof. First, note that we can always find some $\bar{I}$ such that

$$\rho \leq \frac{\eta^{\bar{I}-1} \rho^*}{\nu^{\bar{I}-1}(1-\rho^*) + \eta^{\bar{I}-1} \rho^*}.$$

Therefore, for every prior value ρ, we can construct a finite, irreducible rating system. Now partition the payoff space into X_ℓ and X_h and the index space into $\mathscr{M}_h = \{m : x_m \in X_h\}$ and $\mathscr{M}_\ell = \{m : x_\ell \in X_\ell\}$ such that

$$\frac{\sum_{m \in \mathscr{M}_\ell} \gamma_m^B}{\sum_{m \in \mathscr{M}_\ell} \gamma_m^G} > \frac{\sum_{m \in \mathscr{M}_h} \gamma_m^B}{\sum_{m \in \mathscr{M}_h} \gamma_m^G}.$$

In words, we divide the payoff space into two aggregation of payoffs, such that one aggregation leads to higher aggregated likelihood than the other. Such aggregations will work as an aggregate payoff (signal). Denote by φ_{ij}^n the chosen transition rule from rating i to rating j conditional on the observation of any $x \in X_n$, with $n \in \{\ell, h\}$. Define as well $\delta^\theta = \sum_{m \in \mathscr{M}_h} \gamma_m^\theta$, for each θ.

Consider the following (irreducible) system with $\bar{I}$ ratings: (1) $\varphi^h_{12}=\tau$, $\varphi^h_{11}=1-\tau$, $\varphi^\ell_{11}=1$; (2) $\varphi^\ell_{\bar{I}\bar{I}-1}=\kappa$, $\varphi^\ell_{\bar{I}\bar{I}}=1-\kappa$, $\varphi^h_{\bar{I}\bar{I}}=1$; (3) $\varphi^h_{ii+1}=\varphi^\ell_{ii-1}=1$ for every $i\notin\{1,\bar{I}\}$. With these transition rules, we obtain the following steady state probabilities, for each θ:

$$f_2^\theta=\tau\left(\frac{\delta^\theta}{1-\delta^\theta}\right)f_1^\theta,\qquad f_i^\theta=\left(\frac{\delta^\theta}{1-\delta^\theta}\right)f_{i-1}^\theta\ \forall\, i\notin\{1,\bar{I}\},\qquad f_{\bar{I}}^\theta=\frac{1}{\kappa}\left(\frac{\delta^\theta}{1-\delta^\theta}\right)f_{\bar{I}-1}^\theta.$$

In particular, we have $f_{\bar{I}}^\theta=\frac{\tau}{\kappa}\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^{\bar{I}-1}f_1^\theta$. This generates a system of equations in which

$$f_1^\theta+\tau\left(\frac{\delta^\theta}{1-\delta^\theta}\right)f_1^\theta+\tau\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^2f_1^\theta+...+\tau\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^{\bar{I}-2}f_1^\theta+\frac{\tau}{\kappa}\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^{\bar{I}-1}f_1^\theta=1.$$

Solving for f_1^θ leads to

$$f_1^\theta=\left[1+\tau\left(\frac{\delta^\theta}{1-\delta^\theta}\right)+\tau\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^2+...+\tau\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^{\bar{I}-2}+\frac{\tau}{\kappa}\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^{\bar{I}-1}\right]^{-1}. \tag{A.2.1}$$

We want to set $\frac{f_{\bar{I}}^B}{f_{\bar{I}}^G}=\lambda$. To achieve this, we find that $\frac{f_1^B}{f_1^G}$ must be such that

$$\frac{f_{\bar{I}}^B}{f_{\bar{I}}^G}=\left(\frac{\delta^B}{\delta^G}\frac{1-\delta^G}{1-\delta^B}\right)^{\bar{I}-1}\frac{f_1^B}{f_1^G}=\lambda\qquad\Rightarrow\frac{f_1^B}{f_1^G}=\lambda\left(\frac{\delta^G}{\delta^B}\frac{1-\delta^B}{1-\delta^G}\right)^{\bar{I}-1}.$$

Note that $\frac{f_1^B}{f_1^G}\leq\left(\frac{\eta}{\nu}\right)^{\bar{I}-1}\lambda$ from the definition of δ^θ. Moreover, $\frac{f_2^B}{f_2^G}=\lambda\left(\frac{\delta^G}{\delta^B}\frac{1-\delta^B}{1-\delta^G}\right)^{\bar{I}-2}\leq\left(\frac{\eta}{\nu}\right)^{\bar{I}-2}\lambda$ as well as $\frac{f_i^B}{f_i^G}\leq\left(\frac{\eta}{\nu}\right)^{\bar{I}-i}\lambda$ for every $i\notin\{1,\bar{I}\}$, because $f_i^\theta=\left(\frac{\delta^\theta}{1-\delta^\theta}\right)f_{i-1}^\theta$. This means that all non-investment steady state probabilities satisfy the constrained obtained from lemma 1.4. Also note that substituting equation A.2.1 into the expression for f_2^θ, we have

$$f_2^\theta=\tau\left(\frac{\delta^\theta}{1-\delta^\theta}\right)f_1^\theta=\frac{\tau}{\frac{1-\delta^\theta}{\delta^\theta}+\tau+\left(\frac{\delta^\theta}{1-\delta^\theta}\right)\tau+...+\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^{\bar{I}-3}\tau+\left(\frac{\delta^\theta}{1-\delta^\theta}\right)^{\bar{I}-2}\frac{\tau}{\kappa}}.$$

Therefore, $\lim_{\tau\to 0}f_2^\theta=0$ as well as $\lim_{\tau\to 0}f_i^\theta=0$, for $i=3,4,...,N-1$. Furthermore, if $\tau\to 0$ an $\kappa\to 0$, but $\frac{\tau}{\kappa}>0$, then

$$\lim_{\tau\to 0,\kappa\to 0}\frac{f_1^B}{f_1^G}=\frac{1+\left(\frac{\delta^G}{1-\delta^G}\right)^{\bar{I}-1}\frac{\tau}{\kappa}}{1+\left(\frac{\delta^B}{1-\delta^B}\right)^{\bar{I}-1}\frac{\tau}{\kappa}}.$$

Given that we want $\frac{f_{\bar{I}}^B}{f_{\bar{I}}^G} = \lambda$, let us find the appropriate ratio $\frac{\tau}{\kappa}$ such that $\lim_{\tau\to 0,\kappa\to 0} \frac{f_1^B}{f_1^G} = \lambda\left(\frac{\delta^G}{\delta^B}\frac{1-\delta^B}{1-\delta^G}\right)^{\bar{I}-1}$. With some algebra, we get

$$\frac{\tau}{\kappa} = \frac{\lambda}{1-\lambda}\left(\frac{1-\delta^B}{\delta^B}\right)^{\bar{I}-1} - \frac{1}{1-\lambda}\left(\frac{1-\delta^G}{\delta^G}\right)^{\bar{I}-1}. \tag{A.2.2}$$

This value must be positive, so we need to guarantee that $\lambda > \left(\frac{\delta^B}{\delta^G}\frac{1-\delta^G}{1-\delta^B}\right)^{\bar{I}-1}$. But $\left(\frac{\delta^B}{\delta^G}\frac{1-\delta^G}{1-\delta^B}\right) < 1$ by construction, so as long as we set some $I \geq \bar{I}$ sufficiently large, this inequality must be satisfied. Therefore, we have the result. Take, for instance, the sequence$\kappa_t = \frac{1}{t}$ for each $t \in \mathbb{N}$ and set

$$\tau_t = \frac{1}{t}\left[\frac{\lambda}{1-\lambda}\left(\frac{1-\delta^B}{\delta^B}\right)^{\bar{I}-1} - \frac{1}{1-\lambda}\left(\frac{1-\delta^G}{\delta^G}\right)^{\bar{I}-1}.\right].$$

As $t\to\infty$, we have that $\kappa_t \to 0, \tau_t \to 0$ and the ratio $\frac{\tau_t}{\kappa_t}$ equals the ratio A.2.2 for every t. Thus,

$$\begin{aligned}
\lim_{t\to\infty}\Pi &= \rho f_I^B + (1-\rho) f_I^G, \\
&= f_I^G(\rho\lambda + 1 - \rho), \\
&= \frac{1-\rho}{1-\rho^*}\left[\frac{\tau}{\kappa}\left(\frac{\delta^G}{1-\delta^G}\right)^{I-1} f_1^G\right], \\
&= \frac{1-\rho}{1-\rho^*}\left[\frac{\frac{\tau}{\kappa}\left(\frac{\delta^G}{1-\delta^G}\right)^{I-1}}{1+\frac{\tau}{\kappa}\left(\frac{\delta^G}{1-\delta^G}\right)^{I-1}}\right].
\end{aligned}$$

The second equality follows from $f_I^B = \lambda f_I^G$; the third from the definition of λ and $f_I^G = \frac{\tau}{\kappa}\left(\frac{\delta^G}{1-\delta^G}\right)^{I-1} f_1^G$; the fourth from equation A.2.1. Using the value of $\frac{\tau}{\kappa}$ as in A.2.2,

$$\frac{\tau}{\kappa}\left(\frac{\delta^G}{1-\delta^G}\right) = \frac{\lambda}{1-\lambda}\left(\frac{\delta^G}{\delta^B}\frac{1-\delta^B}{1-\delta^G}\right)^{I-1} - \frac{1}{1-\lambda}.$$

Substituting this back to the limiting value of Π,

$$\lim_{t\to\infty}\Pi = \frac{1-\rho}{1-\rho^*}\left[\frac{\lambda\left(\frac{\delta^G}{\delta^B}\frac{1-\delta^B}{1-\delta^G}\right)^{I-1} - 1}{\left(\frac{\delta^G}{\delta^B}\frac{1-\delta^B}{1-\delta^G}\right)^{I-1} - 1}\right]\frac{1}{\lambda} = \frac{1-\rho}{1-\rho^*}\left[\frac{\lambda - \left(\frac{\delta^B}{\delta^G}\frac{1-\delta^G}{1-\delta^B}\right)^{I-1}}{\lambda\left(1-\left(\frac{\delta^B}{\delta^G}\frac{1-\delta^B}{1-\delta^B}\right)^{I-1}\right)}\right].$$

Now, $\left(\frac{\delta^B}{\delta^G}\frac{1-\delta^G}{1-\delta^B}\right) \geq \frac{\nu}{\eta}$, but $\lim_{I\to\infty}\left(\frac{\delta^B}{\delta^G}\frac{1-\delta^G}{1-\delta^B}\right) = \lim_{I\to\infty}\left(\frac{\nu}{\eta}\right)^{I-1} = 0$. Thus, for every $\varepsilon > 0$, we can find some $I \geq \bar{I}$ that still keeps equation A.2.2 positive and gets $\left(\frac{\delta^B}{\delta^G}\frac{1-\delta^G}{1-\delta^B}\right)^{I-1}$ ε-close to $\left(\frac{\nu}{\eta}\right)^{I-1}$. $\square$

2 BAD REPUTATION WITH RATING SYSTEMS

2.1 Introduction

People rely on online reputation systems for a variety of daily activities.[1] Rating systems help consumers choose their goods, but increasingly also help fulfill clients' request for expert services. Prior to visiting a doctor's office, patients can check physicians' ratings and reviews at RateMDs. Students can select courses according to the teachers' evaluations from other students at Rate My Professors or Rate My Teachers. Americans can hire, rate, and review home service providers on Angi, and motorists can find the best-rated mechanics all around the world on Yelp.

A great deal is known about designing rating systems that rate products, but much less is known about systems that rate economic agents. The substantial difference between the two classes of rating systems is that, in the latter the expert being rated can strategically influence her rating, so the design must account for such strategic effects as well. For concreteness, consider the case of an expert who faces an online feedback record. Some clients might imperfectly ascertain their true needs and evaluate the expert's service negatively, even in the case in which the diagnosis and the proposed treatment are correct. This is often the case when physicians propose a high-cost intervention for a patient who feels that a lower cost treatment would lead to same result, or when teachers are badly evaluated for being rigorous.[2] This imperfect assessment by clients can generate a perverse effect, because the expert might have incentives to provide wrong solutions in exchange for a better review. For example, a mechanic could propose a low-cost tune-up for a driver in order to avoid being reviewed as an expensive professional, when the most adequate solution would be a costly engine replacement. Additionally, a customer might look at past feedbacks as deceptive, because the platform could censor information or benefit premium service providers.[3] For those reasons, the benefits of such mediated interactions could be limited.

In this paper, we study the design of recommender systems in a reputation game. We consider a stylized interaction between a long-run expert and a sequence of short-run customers. In each period, a customer has the option of hiring the expert's services. If hired, the expert observes a problem and proposes a treatment. A severe problem requires a high cost treatment and a mild one is solvable through a low cost treatment. The customer, however, cannot determine how severe the problem is. Moreover, each customer has a common prior belief that an expert is a "bad" commitment type that always provides the expensive treatment. Our benchmark setting is the bad reputation model by Ely and Välimäki (2003), but with an intermediary committed to a special class of information policy. We chose this model as a benchmark in order to isolate how the design of a recommender system interacts with its strategic implications in a reputation game. In this game, reputational concerns have a striking effect on the outcome.

[1] This is a joint work with Daniel Monte.

[2] Some companies even specialized in helping doctors conceal patients' reviews and devise "anti-ratings" contracts, in which patients are asked to sign away their right to provide additional information of a medical service online (Goldman, 2010). In France, a legal decision banned Rate My Professors and related sites from naming national teachers (Tech Dirt, 2008).

[3] Recently, Angi - former Angie's List - got caught up in a scandal concerning the aggregation of home service providers that advertise on the site in a "top-rate pros" category, even though such professionals were not highly rated among users (CBS News, 2019).

In Ely and Välimäki's model, the reputational concern of the strategic expert becomes so strong that the market collapses and customers only rarely hire the expert (we reproduce this result here as theorem 2.1). Due to this bad reputation effect, information censoring might improve efficiency. In fact, the outcomes of the game with full censoring - essentially an infinite repetition of the same static game - Pareto dominates the outcomes of the game with no censoring at all (full information), for all players.

The class of information policies we consider consists of (i) a finite message space,[4] (ii) an initial distribution over messages, and (iii) a transition rule, mapping the current message and any possible observable outcome (high treatment, low treatment or no treatment at all) to a probability distribution over the message space. Technically, information policies in our model are finite automata. Empirically, one can think of them as rating systems, commonly observed in many online reputation systems.

We focus on this class of information policies to capture complexity issues that online platforms face when dealing with huge amounts of data. It is also motivated by the fact that this class of information policies might better represent the way in which customers draw their inferences from such relationships. Under full memory, customers from different periods must be aware of all possible treatment histories and remember the precise realization up to their period of choice. A more reasonable assumption is that customers rely on ratings as summary statistics and aggregate histories into common equivalent classes[5].

Our recommender systems capture a plausible account of the way platform users consume available data. Star ratings matter a great deal to online buyers. Recent surveys show that the star rating is the most or the second most important factor in online reviews that consumers pay attention to when judging a business or a product.[6] Customers in our model rely exclusively on the design of the rating system and the observation of the current rating to infer the type of expert they are hiring.

Due to the fact that we are restricting the class of policies a platform can design, there will be a constraint on what the intermediary can achieve with a rating system. We characterize the upper bounds on the expert's and customers' equilibrium payoffs and construct rating systems that reach such upper bounds.

We show that, to maximize customers' value from the interactions (or if the online platform only cares about the customers), the higher the number of ratings is, the closer the customers' equilibrium payoff will be to the upper bound (theorem 2.2). Interestingly, however, in equilibrium, the system is almost always not in intermediary ratings. In other words, a mass concentration of data in extreme ratings arises endogenously from our characterization of an optimal rating system. Such "rating inflation"is consistent with what is empirically observed on popular rating systems.[7]

[4]Actually, the message space can be finite or countably infinite. Although we will present present results for finite systems only, countable systems could be accommodated. In particular, the message space could be the entire space of infinite sequences of public outcomes, reproducing the full information environment from Ely and Välimäki (2003).

[5]Compte and Postlewaite (2015) motivate restricting players to play similarly across all histories in the same group of equivalence classes as a more realistic description of cooperative behavior in the long-run. We share the same view, although we differ from their paper because we construct optimal systems for each player.

[6]See for instance surveys on Bright Local (2018, 2019).

[7]For example, eBay sellers have a median score of perfect rating (Nosko and Tadelis, 2015); Airbnb has an overwhelming percentage of properties with 4.5 stars or more (Zervas, Proserpio and Byers, 2021). Other well known rating systems, such as Amazon, Yelp, Uber, Lyft, and online labor marketplaces also exhibit this pattern. We stress that there can be other explanations for this phenomenon, such as reporting biases; see, for example, Dellarocas and Wood (2008) and Filippas, Horton and Golden (2018).

The optimal rating system for customers requires carefully designed transition rules to trigger the correct incentives for the expert to tell the truth. Specifically, the expert's discount factor matters for the construction. The intuition for this is that, ideally, the intermediary would like to have a rating (say, rating 1) in which customers obey a non-hiring recommendation and other ratings in which hiring is optimal. The intermediary does not observe the expert's type so it uses hiring ratings to learn more about him.

Without loss of generality, let us assume that the posterior belief about the expert being bad is higher at lower ratings. The intermediary would like to have the lowest posterior belief arbitrarily close to zero, and the system staying most often at extreme ratings, at which the posterior belief spread is maximal.

This is achieved if ratings are upgraded with higher probability under a low cost treatment than a high cost one. However, as the expert's patience rate increases, the more he cares about being caught up in the extreme rating in which customer hires. He then has incentives to game the system by providing cheap treatment to a severe problem.

Since with no memory there is no conflict of interest (strategic expert and customers' utilities coincide), it suffices to make transitions between ratings less frequent as discount factor rises. Using a suitable choice of transition rules, we have well-defined extreme posterior beliefs that are incentive compatible for both players, even for $\delta \approx 1$. Therefore the optimal rating system for customers also involves the expert being hired often and telling the truth in equilibrium.

To maximize the expert's value from the interaction, we only need two ratings (theorem 2.3). The optimal system is a significant improvement in comparison to both the no-censoring and the full-censoring environments. In this optimal system, there is no bad reputation effect and the expert is often hired in the long run.

Both customers' and experts' equilibrium payoffs are constrained mainly because our class of information policies constrains the posterior belief spread. This happens because the Markov restriction generates a strong connection between hiring and non-hiring ratings. At hiring ratings, the strategic expert must play a truth-telling strategy with at least some positive probability; otherwise, customers would not find optimal to hire. This means that he must generate the same signal of a bad expert. So the strategic expert visits non-hiring ratings sometimes in equilibrium. At non-hiring ratings, there is no additional evidence to separate types. Thus, the beliefs in such ratings account for the possibility of expert being strategic, leading to an upper bound on the highest posterior belief about the expert the being bad type.

Our paper shows that rating systems generate enough data to support a trustful interaction between experts and customers, relative to no information (that is, the repetition of the one-shot interaction) and full information (Ely and Välimäki's game). Thus, our paper has some policy implications. For instance, well designed systems can alleviate the disarrangement between a patient's perception of a problem and a physician's diagnosis of it, a problem we discussed at the beginning of this introduction.

We consider information censoring through finite memory as a device against bad reputational concerns. Ely and Välimäki (2003) construct a score mechanism that solves the bad reputation effect, but it depends on relaxing the assumption that customers are short-run. Being long-run, a customer has more incentives to hire the expert even if he sometimes provides the wrong treatment: by doing so today, the customer collects additional information about the expert's type to support better hiring decisions later. Mailath and Samuelson (2006) show that bad reputation is also avoidable when some customers are captive, that is, they choose to hire the expert regardless of his reputation. Our customers are short-lived and strategic.

Conceptually, the class of information policies we consider resembles finite automata designed to generate some learning about an unknown state of the world. Thus we connect with the literature of learning with limited memory. We use techniques from Hellman and Cover (1970) and Wilson (2014), both of which study optimal finite memory allocation when information is incomplete. As in Monte (2013), we apply such techniques in a reputation game. However, the design of the memory system comes from an intermediary in our model, whereas the cited papers discuss memory design from the perspective of the uninformed player. This means that all players in our model take the transition rules over memory states as given, but the actions in each memory state must be incentive compatible.

Lorecchio and Monte (2021) study optimal rating systems to maximize stationary probability of buyers buying a product of unknown quality. The environment is similar to ours, with the difference that reputational concerns in our model generate an additional incentive compatibility constraint. In other words, in Lorecchio and Monte (2021) the rating system is rating a product and the signals depend on the type of the product, whereas here the rating system rates an expert, and the expert's behavior is itself dependent on the particular rating system.

Ekmekci (2011) also relates to our paper. In it, a rating system generates the right incentives for a long-run player to always take short-run players' most preferred action, in order to sustain a "good" reputation. Thus, information censoring through finite memory is efficient in restoring good reputational concerns in repeated games with one-sided incomplete information.[8] Besides discussing mechanisms to avoid bad reputational concerns, our system has some fundamental differences from Ekmekci's system. Most importantly, in his model, there is a permanent flow of informative signals about long-run players' type. In our model, whenever customers do not hire, there is no signal about his type. Technically, we deal with an interactive automaton.

Other papers discuss information censoring as a way to overcome the bad reputation effect.[9] Sperisen (2018) studies bounded recall, that is, when customers are only able to remember histories of a certain length. He also considers bad reputation under fading recall, that is, customers' memory decays over time. Bad reputation is avoidable provided that memory decays quickly enough.

[8]Fudenberg and Levine (1989; 1992) have shown that if there is a small common knowledge probability that the long-run player is committed to the short-run players' most preferred action, it can be sequentially rational for this long-run player to always play such preferred action. Cripps, Mailath and Samuelson (2004) however proved that reputational concerns are insufficient to curb incentives for the long-run player to deviate when his actions are imperfectly monitored.

[9]More broadly, Vong (2021) studies rating systems in a reputation game similar to Ekmekci (2011), but with a commitment type of long-lived player that never exerts effort. In it, the intermediary knows the type of the long-lived player prior to the construction of the system. Instead, our designer only observes the outcomes of the interactions with short-run players, which are signals about expert's type.

Lillethun (2017) studies dynamic information disclosure in repeated games with reputational concerns and consider the bad reputation game as an application. In a finite time horizon, bad reputation is avoidable if early customers are kept uninformed about expert's choices and late customers get full information from early periods.

Our class of systems is time independent and we focus on stationary strategies for players, since we assume that customers are unaware of calendar time. In some sense, we view this as a design of a simple information rule. Even though we show that it does generates a bound on the maximum payoff players can get in equilibrium, we highlight that simpler rules have some advantages over more complicated ones. First, they simplify the inference process for the uninformed players. Dynamic information policies can depend on public histories in a quite complex fashion, and requires customers to keep track of all possible outcomes of past interactions. Similar to Compte and Postlewaite (2015), we view rating systems as a more plausible description of how customers reason about the expert's type upon observing a given message. [10]

We organize the rest of this paper as follows. In section 3.2, we present Ely and Välimäki's setting and comment on their bad reputation result. We then present our definition of a rating system and construct a simple example to fix ideas. In Section 3.3, we start the analysis with binary, irreducible rating systems. Section 3.4 extends the results for finite rating systems, keeping irreducibility. Section 3.5 proves that focusing on irreducible systems is actually without loss of generality. The main results of our paper are presented in Section 3.6.

2.2 Model

The setting is a repeated game between one long-run player (the expert) who interacts with a sequence of identical time t myopic players (the customers). We reproduce the underlying model of Ely and Välimäki (2003), but we restrict strategies to be functions of ratings and not on finer details of the public history.

At the beginning of every stage game, Nature draws problem $\theta \in \{H,L\}$, each happening with probability $1/2$. Without observing θ, a customer decides first whether to hire the expert (In) or not (Out). If she does not hire, both players get an outside option payoff which is normalized to zero. If she does hire, the expert then perfectly observes θ and decides the level of service to provide. An appropriate treatment for H would be t_H and an appropriate treatment for L would be t_L. The customer observes the treatment and the payoffs are realized: the right treatment generates a payoff of u and the wrong one generates a payoff of $-w$. The stage game is summarized in figure 3.1.

[10] In contract theory and monetary theory, policy design can also be complicated. Thus, some authors have studied how simpler rules can approximate optimal outcomes from more complex ones (Levin and Williams, 2003; Rogerson, 2003; Garrett, 2004; Carroll, 2015).

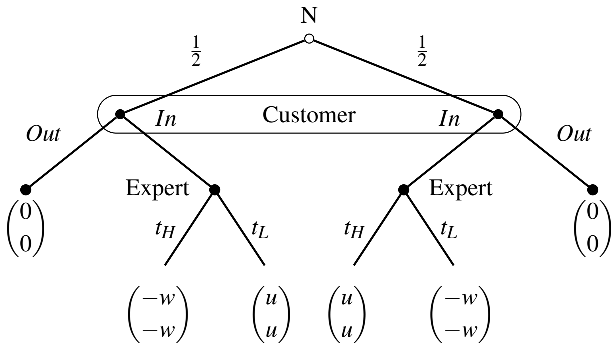

Figure 2.1 – The stage game. Nature draws a problem H or L with equal probability. A customer chooses whether to hire the expert or not. If hired, expert observes the problem and proposes treatment t_H or t_L. A right treatment generates u for both players; a wrong one generates $-w$.

There is incomplete information about the expert's type. Specifically, with common knowledge probability $\rho \in (0,1)$, the expert is a type committed to a pure strategy: he always provides the high-cost treatment t_H. We denote this behavioral "bad" type by B. With probability (1-ρ), he is a strategic type,[11] denoted by S, and his payoffs at the end of each period are identical to customers' payoffs.

Assume $w > u > 0$. It is straightforward to see that the expert strictly prefers to provide the correct treatment when hired in the stage game in absence of a repeated setting. Let $v(\rho)$ denote a customer's expected payoff from hiring in the stage game as a function of the initial prior. This is given by:

$$v(\rho) = \rho\left(\frac{u-w}{2}\right) + (1-\rho)u. \tag{A.2.1}$$

The non-hiring payoff is zero to both players, so a customer will not hire if $v(\rho) < 0$. A simple manipulation of the above equation shows that this is equivalent to ρ being higher than a belief threshold $\rho^* \in (0,1)$, defined below.

$$\rho^* := \frac{2u}{u+w}. \tag{A.2.2}$$

2.2.1 Bad reputation

Let us now discuss the bad reputation result in Ely and Välimäki (2003). Assume first that $\rho > \rho^*$. Then customers never hire and the expert never gets the chance to reveal himself. Thus, his average discounted payoff in equilibrium is zero. Assume now that $\rho \leq \rho^*$. They show that if customers get to see all past treatments chosen by the expert (but not the associated problems) and choose according to this public history, the expert will rarely be hired in any (Nash) equilibrium if he is sufficiently patient. We state below the theorem.

Theorem 2.1 (Ely and Välimäki, 2003)**.** *Fix any prior $\rho \in (0,1)$ and any discount $\delta \in (0,1)$. Let $\bar{V}_\delta(\rho)$ be the supremum of discounted average equilibrium payoffs for the strategic type of expert. Then,* $\lim_{\delta\to 1} \bar{V}_\delta(\rho) = 0$.

[11]Unless stated otherwise, when we refer to the expert, we will be referring to this strategic type.

Moreover, if incentives to play the wrong treatment to avoid this zero payoff equilibrium are strong enough, he will never be hired. History-depend strategies and reputation effects harm both the expert and the customers in this game.

Here is the intuition for this result. At every period t, hiring is possible only if the expert is playing correct actions with positive probability for every problem, including treatment t_H for problem H. This implies that the posterior belief ρ_{t+1} of any customer about expert being bad is increasing in the prior ρ_t. Thus, a long sequence of problems H inevitably leads to the posterior surpassing the critical level ρ^*, even if the expert frequently plays t_L in θ_H with some probability. If the expert discounts future payoffs at a rate $\delta \in (0,1)$, this means that he can expect to earn the zero payoff on average, if he is sufficiently patient.

Furthermore, if we assume that the expert is hired at any history on equilibrium path at which he is revealed to be good (that is, whenever he plays t_L; Ely and Valimaki refer to this assumption as renegotiation proofness), then hiring does not take place at all if the expert is sufficiently patient: the temptations to play t_L are so high that he will play it before reaching the critical belief ρ^*. But if a customer knows that the incentives to lie are strong enough, she will not hire at the critical history in the first place. The expert knows about this customer's unwillingness to hire and lies whenever close to a belief which is slightly below the critical level. A backward induction argument leads to the non-hiring result.

2.2.2 Rating systems

We relax Ely and Välimäki's full memory assumption and instead assume that all customers have access to only one source of information - a rating system. Specifically, a rating system $\mathscr{M} := (M, \varphi, \varphi_0)$ consists of (i) a finite set of ratings M; (ii) a transition function φ determining a probability distribution over ratings at period t as a function of the rating at period $t-1$, the choice customer made at that period and the outcome of her interaction with the expert (t_H or t_L in case she hired him; *Out* otherwise) and (iii) an initial probability distribution φ_0 over ratings.

Customers can only condition their behavior on ratings, not on finer details of the public history. Their strategies do not even depend on calendar time. They take the system as given and their strategies are a map from the current rating to a probability of hiring the expert:

$$\alpha : M \to [0,1].$$

There are many underlying motivations for assuming this fairly simple constraint on information technology. For example, we can think of an information designer restricted to disclose specific information in standardized format and conditioned only on data that customers can verify. We can also think of a rating as a summary statistic for the repeated interaction and a rating system as an information policy that simplifies the implementation and the inference process (Dellarocas, 2010). Finally, we can think of ratings as customers' moods at which a specific decision is always optimal and a rating system as a plausible description of how customers and the expert cooperate to overcome the bad reputation effect (Compte and Postlewaite, 2015).

The expert also observes the realized rating at every period. Since this is the only information customers have, we will focus on equilibria in which the expert's choice of treatment will vary only across ratings and problems, not across time. Thus, expert's strategies are maps from current ratings and current problems to a probability of providing correct treatment:

$$\beta_H : M \to [0,1],$$
$$\beta_L : M \to [0,1].$$

We assume that customers compute the probability distribution over the public histories as if the game had been going on for a long time. Thus, they compute beliefs using steady state probabilities (or time-average convergence). To understand how these probabilities arise, note that any given rating system $\mathscr{M}$, strategy profile $(\alpha, \beta_H, \beta_L)$ and prior ρ define Markov matrices T^S and T^B for the strategic and the bad expert, respectively. To simplify notation, let $\varphi^y_{mm'}$ represent the probability of transitioning from rating m to m' upon the observation of $y = H$ (corresponding to t_H), $y = L$ (corresponding to T_L) or $y = Out$. Then the entries of T^S and T^B are:

$$\tau^S_{mm'} = \alpha_m[\gamma_m \varphi^H_{mm'} + (1-\gamma_m)\varphi^L_{mm'}] + (1-\alpha_m)\varphi^{Out}_{mm'}, \tag{A.2.3}$$
$$\tau^B_{mm'} = \alpha_m \varphi^H_{mm'} + (1-\alpha_m)\varphi^{Out}_{mm'}, \tag{A.2.4}$$

where $\gamma_m := 1/2 \times \beta_H(m) + 1/2 \times (1-\beta_L(m))$ represents the probability of observing the t_H treatment at m and $1-\gamma_m$ the probability of observing treatment t_L, if expert is strategic. From a well-known result in Markov processes[12], there will be unique invariant distributions $f^B := (f^B_m)_{m\in M}$ and $f^S := (f^S_m)_{m\in M}$. Customers use such distributions to compute updated beliefs at every rating reached with positive probability, as defined below.

$$\rho_m := Pr[B|m] = \frac{\rho f^B_m}{\rho f^B_m + (1-\rho) f^S_m} \qquad \forall m \text{ s.t. } f_m := \rho f^B_m + (1-\rho) f^S_m > 0. \tag{A.2.5}$$

Define $\beta_m := 1/2 \times \beta^H(m) + 1/2 \times \beta^L(m)$ as the expert's probability of telling the truth at rating m. Customers' expected payoff from hiring is

$$v(\rho_m) = \rho_m \left(\frac{u-w}{2}\right) + (1-\rho_m)[\beta_m u - (1-\beta_m)w]. \tag{A.2.6}$$

Customers' strategy must be incentive compatible, that is, whenever they are willing to hire, the expected payoff from doing so must be non-negative. Thus, customers find optimal to hire for every m reached with positive probability in the long-run such that $v(\rho_m) \geq 0$.

[12] See for instance Stokey and Lucas (1989), theorem 11.2.

For any discount factor δ, let $V_\delta^\theta(m)$ denote expert's continuation value from a rating system and a strategy profile, at rating m, after being hired and after observing problem $\theta \in \{H, L\}$. These continuation values are given by

$$V_\delta^H(m) := (1-\delta)[\beta_H(m)u - (1-\beta_H(m))w] + \delta \sum_{m' \in M} [\beta_H(m)\varphi_{mm'}^H + (1-\beta_H(m))\varphi_{mm'}^L] V_\delta(m'),$$
$$V_\delta^L(m) := (1-\delta)[\beta_L(m)u - (1-\beta_L(m))w] + \delta \sum_{m' \in M} [\beta_L(m)\varphi_{mm'}^L + (1-\beta_L(m))\varphi_{mm'}^H] V_\delta(m'). \tag{A.2.7}$$

Expert's strategy must be optimal, that is, whenever hired and upon the observation of θ, β_θ must maximize the continuation value defined above.

For any given system $\mathscr{M}$ and prior ρ, we define an equilibrium as a strategy profile $(\alpha, \beta_H, \beta_L)$ and a posterior belief distribution $(\rho_m)_{m\in M}$ such that (i) customers are taking optimal actions given her posterior beliefs; (ii) the expert is taking optimal actions whenever hired and (iii) posterior beliefs are consistent with Bayes rationality and the invariant distributions f^B and f^S. We state the concept below.

Definition 2.1. *An equilibrium is a strategy profile $(\alpha, \beta_H, \beta_L)$ as well as a posterior belief distribution $(\rho_m)_{m\in M}$ such that, for every $m \in M$ with $f_m^S > 0$,*

1. *ρ_m is consistent, that is, derived from Bayes' rule and the stationary distributions f^B and f^S, as in equation A.2.5;*
2. *$\alpha_m > 0 \Rightarrow \nu(\rho_m) \geq 0$, where $\nu(\rho_m)$ is defined in equation A.2.6;*
3. *for every $\theta \in \{H, L\}$, $\beta_\theta(m)$ maximizes $V_\delta^\theta(m)$ whenever $\alpha_m > 0$, as in equation A.2.7.*

The strategic expert computes the probability distribution over histories as if the game has been going for a long time, but discount periods using δ. There will also be a unique invariant distribution $(g^S) = (g_m^S)_{m\in M}$ to a slightly modified Markov matrix that depends on the discount factor[13]. We will be mostly interested in results for a very patient expert. One can show that, whenever the transition matrix T^S is irreducible, g^S converges to f^S as $\delta \to 1$. The value from the interaction that interest us is

$$\lim_{\delta \to 1} V_\delta = \sum_{m\in M} f_m^S \alpha_m [\beta_m u - (1-\beta_m)w]. \tag{A.2.8}$$

We are also interested in the value from the interaction for the customers. This is given by the expected value of the payoffs at each rating with respect to the probability of reaching such ratings. Specifically, customers' value from the interaction is

$$v = \sum_{m\in M} f_m \alpha_m \nu(\rho_m). \tag{A.2.9}$$

[13] See Wilson (2014), lemma 3.

We will look for rating systems that maximize expert's or customers' value from the interaction. An optimal system for the expert is one for which there is an equilibrium and this equilibrium leads to the maximum value of equation A.2.8. Likewise, an optimal system for customers is one for which there is an equilibrium and this equilibrium leads to the maximum value of equation A.2.9. We state this concept below.

Definition 2.2. *An optimal system for the expert is a rating system that induces a strategy profile* $(\alpha, \beta_H, \beta_L)$ *leading to the highest value of equation A.2.8 subject to* $(\alpha, \beta_H, \beta_L)$ *being an equilibrium. An optimal system for the customers is a rating system that induces a strategy profile* $(\alpha, \beta_H, \beta_L)$ *leading to the highest value of equation A.2.9 subject to* $(\alpha, \beta_H, \beta_L)$ *being an equilibrium.*

Evidently, there are some trivial systems that lead to highest value of information for one of the players. For instance, if $\rho \leq \rho^*$, then a rating system with just one rating or a binary system with uninformative ratings leads to the highest value, since customers always hire in the infinite repetition of the one-shot game. But if $\rho > \rho^*$, then both the no memory and the full memory settings cannot generate hiring at all. We will show that rating systems improve expert's payoff upon these extreme information settings because it creates incentives for customers to hire the expert, even if the prior is above ρ^*. For customers, even if $\rho \leq \rho^*$, a rating system might lead to a higher equilibrium payoff relative to no memory and full memory. This is true because customers benefit from a better separation of types, even if hiring is optimal in the one-shot interaction.

To understand the mechanics of our model, let us consider the following example.

Example 2.1. *The rating set is binary and there is a deterministic transition rule. Customers hire in rating* 2*, but not in rating* 1*. Starting at rating* 2*, every time the expert provides treatment* t_H*, the rating moves to* 1*; otherwise, the system remains in* 2*. Once in* 1*, the system moves to* 2*. Our equilibrium candidate is (i)* $\alpha_1 = 0$ *and* $\alpha_2 = 1$*; (ii)* $\beta^H(2) = \beta^L(2) = 1$*. One way of interpreting this is as if rating 1 is a "bad" market perception and rating 2 is a "good" one. This system and the strategies that are part of an equilibrium candidate are represented in figure 2 below.*

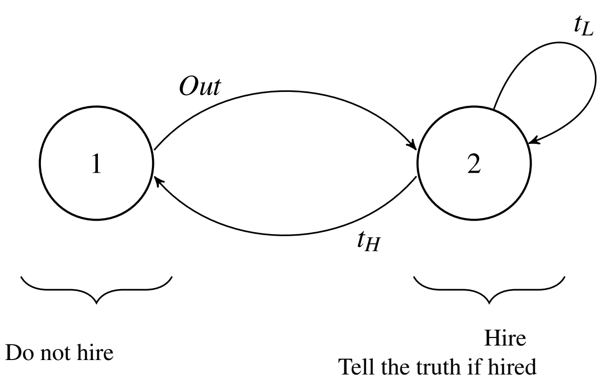

Figure 2.2 – A binary, deterministic and irreducible system in which expert tells the truth if hired and customers hire only in rating 2.

For such a system, recalling that the entries of the transition matrices T^S and T^B are given by equation 3.3 and 3.4 respectively, we get

$$T^B = \begin{array}{c} \\ 1 \\ 2 \end{array}\!\!\begin{array}{c} \begin{array}{cc} 1 & 2 \end{array} \\ \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \end{array} \qquad\qquad T^S = \begin{array}{c} \\ 1 \\ 2 \end{array}\!\!\begin{array}{c} \begin{array}{cc} 1 & 2 \end{array} \\ \begin{pmatrix} 0 & 1 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \end{array}$$

Because both matrices do not have absorbing subsets, we can compute the invariant distributions using the identities $f^B \cdot T^B = f^B$ and $f^S \cdot T^S = f^S$. That leads to the distributions $f^B = (1/2, 1/2)$ and $f^S = (1/3, 2/3)$. The equilibrium beliefs will be

$$\rho_1 = \frac{3\rho}{3\rho + 2(1-\rho)} \qquad\qquad \rho_2 = \frac{3\rho}{3\rho + 4(1-\rho)}.$$

It remains to check whether the strategy profile and the beliefs form an equilibrium. First, note that $\rho_2 \leq \rho^*$ and $\rho_1 > \rho^*$ if and only if $\frac{2\rho^*}{3-\rho^*} < \rho \leq \frac{4\rho^*}{3+\rho^*}$. For this range of prior belief values, $\alpha_1 = 0$ and $\alpha_2 = 1$ is optimal for customers. To see that truth-telling is optimal in this example, note that continuation values in rating 2 are

$$\begin{aligned} V_\delta^H(2) &= (1-\delta)[\beta_H(2)u - (1-\beta_H(2))w] + \delta[(1-\beta_H(2))V_\delta(2) + \beta_H(2)V_\delta(1)], \\ V_\delta^L(2) &= (1-\delta)[\beta_L(2)u - (1-\beta_L(2))w] + \delta[\beta_L(2)V_\delta(2) + (1-\beta_L(2))V_\delta(1)]. \end{aligned}$$

It is optimal for the expert to play t_L for sure upon the observation of problem L in rating 2. This maximizes the current expected payoff $\beta_L(2)u - (1-\beta_L(2))w$ and the likelihood of remaining in rating 2. To have $\beta_H(2) = 1$, it must be that

$$(1-\delta)u + \delta V_\delta(1) \geq (1-\delta)(-w) + \delta V_\delta(2). \tag{A.2.10}$$

We need to derive the values of $V_\delta(1)$ and $V_\delta(2)$. They are given by the following system of equations.

$$V_\delta(1) = \delta V_\delta(2), \qquad\qquad V_\delta(2) = (1-\delta)u + (\delta/2)[V_\delta(1) + V_\delta(2)].$$

One can show that $V_\delta(1)$ and $V_\delta(2)$ are

$$V_\delta(1) = \frac{2u\delta}{2+\delta} \qquad\qquad V_\delta(2) = \frac{2u}{2+\delta}.$$

Therefore, equation 3.10 is satisfied for every $\delta \in (0,1)$, because

$$(1-\delta)u+\delta V_\delta(1) \geq (1-\delta)(-w)+\delta V_\delta(2) \Leftrightarrow 2+\delta(1-\rho^*) \geq 0.$$

For this specific system, for any prior $\frac{2\rho^*}{3-\rho^*} < \rho \leq \frac{4\rho^*}{3+\rho^*}$*, the strategy profile and beliefs form an equilibrium, because (i) customers optimally hire in rating* 2 *but not in rating* 1*; (ii) truth-telling is optimal in rating* 2 *and (iii) beliefs are consistent. Note that, in this equilibrium - at least for a range of prior values, reputational concerns do not tempt the expert to play the wrong treatment to reveal himself and expert's equilibrium payoff is bounded away from zero, even as he becomes increasingly patient. We stress that this is just an example of a rating system, but it is generically not the optimal one neither for the customers nor the expert.*

2.3 Binary systems

Let us now focus on the special case of binary and irreducible rating systems. We will show that there exists a bound on the maximum belief spread that a rating system can generate in equilibrium. This is turn leads to upper bounds on the maximum value from interaction the customers and the expert can get. Nevertheless, binary systems eliminate the bad reputation effect: we prove that expert finds it optimal to tell the truth whenever hired. We will construct rating systems that achieve the obtained upper bounds on the equilibrium payoffs.

For our results, we will follow a similar logic as in Lorecchio and Monte (2021). The essential difference between that paper and this is that here a strategic agent (the expert) is being rated, so the induced stationary probability depends on the rating system but also on the strategy of the expert, which, in turn, is a best response to the strategy of the customers and the rating system. In that paper, the rating system induced a type-dependent stationary probability, and the signals were generated whenever the product was bought. In particular, propositions 2.1, 2.2 and 2.7 in this paper were proved in Lorecchio and Monte (2021). We adapt those proofs to our (strategic) environment and include them in this paper.

Without loss of generality, let us denote by 1 the rating for which there is no hiring and by 2 the rating for which hiring is optimal. Before proceeding, let us define an useful statistic:

$$\lambda := \frac{\rho^*}{(1-\rho^*)}\frac{(1-\rho)}{\rho}.$$

If $\rho \leq \rho^*$, we have $\lambda \geq 1$; if $\rho > \rho^*$, then $\lambda < 1$. This represents the prior bias in favor of hiring the expert. We can write customers' incentive compatibility constraints as

$$\frac{f_1^B}{f_1^S} > \beta_1\lambda - (1-\beta_1)\left(\frac{2-\rho^*}{1-\rho^*}\right)\left(\frac{1-\rho}{\rho}\right) := \lambda_1, \tag{A.2.11}$$

$$\frac{f_2^B}{f_2^S} \leq \beta_2\lambda - (1-\beta_2)\left(\frac{2-\rho^*}{1-\rho^*}\right)\left(\frac{1-\rho}{\rho}\right) := \lambda_2. \tag{A.2.12}$$

There will be a third restriction, coming from expert's incentive compatibility constraint. Here is the intuition. To have hiring in rating 2, the expert must be willing to provide the correct treatment with a minimum probability. This in turn implies that there is an upper bound on the probability of observing treatment t_H when expert is strategic. But then he sometimes must visit rating 1, since we would like to have some positive transition from 2 to 1 that induces relative frequency (f_1^B/f_1^S) to be high enough to justify not hiring there.

Proposition 2.1. *In a binary system with hiring only in rating 2, the expert must play t_H with positive probability in equilibrium and stationary distributions must satisfy*

$$\frac{f_1^B}{f_1^S} \leq \frac{1}{\gamma_2}\frac{f_2^B}{f_2^S}.$$

Proof. Note that the stationary distributions must be such that

$$f_1^S = f_1^S\tau_{11}^S + f_2^S\tau_{21}^S, \qquad f_2^B = f_1^B\tau_{12}^B + f_2^B\tau_{22}^B.$$

Manipulating these identities and noting that $\tau_{12}^B = \tau_{12}^S$ because there is no hiring in 1,

$$\frac{f_1^B}{f_1^S} = \frac{f_2^B}{f_2^S}\left[\frac{\varphi_{21}^H}{\gamma_2\varphi_{21}^H + (1-\gamma_2)\varphi_{21}^L}\right].$$

Inspecting equation A.2.6, we see that hiring in rating 2 implies that expert must be playing a truth-telling in it with minimum probability given by

$$\beta_2 \geq \frac{1}{(u+w)}\left[w + \frac{\rho_2}{1-\rho_2}\left(\frac{w-u}{2}\right)\right] \Rightarrow \beta_2 \geq \frac{w}{u+w}.$$

Because $\beta_2 = 1/2 \times \beta_H(2) + 1/2 \times \beta_L(2)$, there will a lower bound on the probability of playing t_H when problem is H:

$$\frac{1}{2}\beta_H(2) \geq \frac{w}{u+w} - \frac{1}{2}\beta_L(2) \Rightarrow \beta_H(2) \geq \frac{w-u}{w+u}.$$

From the definition of the belief threshold ρ^* in equation 2, the lower bound above equals $1-\rho^*$. Thus, $\gamma_2 = 1/2 \times \beta_H(2) + 1/2 \times (1-\beta_L(2)) \geq (1/2)[1-\rho^*] > 0$. Therefore,

$$\frac{f_1^B}{f_1^S} = \frac{f_2^B}{f_2^S}\left[\frac{\varphi_{21}^H}{\gamma_2\varphi_{21}^H + (1-\gamma_2)\varphi_{21}^L}\right] \leq \frac{1}{\gamma_2}\frac{f_2^B}{f_2^S}.$$

□

This implies that the highest posterior belief is bounded above. Indeed, by combining inequality from proposition 2.1 with compatibility constraint A.2.12, we have

$$\frac{\rho_1}{1-\rho_1} = \frac{\rho f_1^B}{(1-\rho)f_1^S} \leq \frac{\rho\lambda_2}{(1-\rho)\gamma_2} \Rightarrow \rho_1 \leq \frac{\rho\lambda_2}{\rho\lambda_2+(1-\rho)\gamma_2}.$$

Because $\gamma_2 > 0$, it must that this upper bound is lower than 1. It is also the case that it is non-negative, because $\lambda_2 \geq 0$. Indeed, because hiring requires truth-telling in 2 with probability at least $\frac{w}{u+w}$,

$$\begin{aligned}
\lambda_2 &= \beta_2\lambda - (1-\beta_2)\left(\frac{2-\rho^*}{1-\rho^*}\right)\left(\frac{1-\rho}{\rho}\right),\\
&\geq \frac{1}{2}\left(\frac{2w}{u+w}\right)\lambda - \frac{1}{2}\left(\frac{2u}{u+w}\right)\left(\frac{2-\rho^*}{1-\rho^*}\right)\left(\frac{1-\rho}{\rho}\right),\\
&= \frac{1}{2}(2-\rho^*)\lambda - \frac{1}{2}(2-\rho^*)\lambda,\\
&= 0.
\end{aligned}$$

Without loss, we will restrict to systems in which there is no reason not to play t_L when the problem is L. The ratio $\frac{\lambda_2}{\gamma_2}$ is then maximized when expert provides the correct treatment in state t_H as well. To see this is true, note that since $\beta_2 = 1/2 \times \beta_H(2) + 1/2$ and $\gamma_2 = 1/2 \times \beta_H(2)$, we can rewrite $\frac{\lambda_2}{\gamma_2}$ as

$$\begin{aligned}
\frac{\lambda_2}{\gamma_2} &= \lambda\left[\frac{\beta_2}{\gamma_2} - \frac{1-\beta_2}{\gamma_2}\left(\frac{2-\rho^*}{\rho^*}\right)\right],\\
&= \lambda\left[\left(1+\frac{1}{\beta_H(2)}\right) - \left(\frac{1}{\beta_H(2)}-1\right)\left(\frac{2-\rho^*}{\rho^*}\right)\right],\\
&= 2\lambda\left[\frac{1}{\rho^*} - \frac{1}{\beta_H(2)}\left(\frac{1-\rho^*}{\rho^*}\right)\right].
\end{aligned}$$

The highest value is achieved by setting $\beta_H(2) = 1$. This leads to the highest value of $\frac{\lambda_2}{\gamma_2}$ being 2λ. Therefore, $\frac{f_1^B}{f_1^S} \leq 2\lambda$ and $\rho_1 \leq \frac{2\rho^*}{1+\rho^*}$. Proposition 2.2 below summarizes this result.

Proposition 2.2. *The highest posterior belief about expert being bad is bounded above:*

$$\rho_1 \leq \frac{2\rho^*}{1+\rho^*} := \bar{\rho}. \tag{A.2.13}$$

Note that the bound does not depend on the strategy chosen by the expert. But is it optimal for the expert to play a truth-telling strategy in rating 2? We show this is the case in next proposition. Therefore, there is no bad reputation effect in the environment we consider in this section.

Proposition 2.3. *Playing truthfully in the hiring rating is always optimal for the expert.*

Proof. Assume by way of contradiction that telling the truth is not optimal in the hiring rating, upon the observation of problem H. The continuation values at 2 and 1 will be

$$V_\delta(2) = (1-\delta)\left(\frac{u-w}{2}\right) + \delta[\varphi_{21}^{L} V_\delta(1) + \varphi_{22}^{L} V_\delta(2)],$$
$$V_\delta(1) = \delta[\varphi_{11}^{Out} V_\delta(1) + \varphi_{12}^{Out} V_\delta(2)].$$

Solving this system of equations for $V_\delta(2)$, we get

$$V_\delta(2) = \left[\frac{1-\delta+\delta\varphi_{12}^{Out}}{1-\delta+\delta\varphi_{12}^{Out}+\delta\varphi_{21}^{L}}\right]\left(\frac{u-w}{2}\right) < 0.$$

This contradicts not telling the truth being optimal. □

To have an equilibrium, it must be that the prior ρ is lower than the upper bound on the highest belief. If this is not true, then both posterior beliefs ρ_2 and ρ_1 will be higher than ρ and will imply that $f_m^B > f_m^S$ for $m \in \{1,2\}$. But this leads to a contradiction as both distributions must sum 1 over ratings. For the rest of the section, we assume the following.

Assumption 2.1. *The prior is lower than the upper bound on the highest posterior: $\rho \leq \bar{\rho}$.*

It is evident from example 2.1 and proposition 2.3 that there exists a system that sustains a good interaction between customers and the expert, at least for a set of values of the prior belief. But what system maximizes customers' equilibrium payoff? And what system maximizes the expert's equilibrium payoff? We tackle those issues in the next subsections.

2.3.1 Optimal system for customers

From customers' viewpoint, at the hiring rating, it would be optimal if they could know for sure that the expert is strategic. For this to happen, we would need to have $f_2^B = 0$ and $f_2^S > 0$. But from proposition 2.1, that would imply $(f_1^B/f_1^S) \leq 0$ or $f_1^B = 0$, a contradiction. The best thing customer can hope for is having (f_1^B/f_1^S) to be the highest possible and having (f_2^B/f_2^S) to be lowest possible. Recall that

$$\frac{f_2^B}{f_2^S} \geq \gamma_2 \frac{f_1^B}{f_1^S}.$$

From the proof of proposition 2.1, we know that there is a lower bound on the probability of expert playing treatment t_H whenever hired: $\frac{1}{2}(1-\rho^*)$. Therefore, $\frac{f_2^B}{f_2^S} \geq \frac{1}{2}(1-\rho^*)2\lambda$. Rearranging, we get a lower bound for ρ_2.

Proposition 2.4. *The belief about the expert being bad is bounded below:*

$$\rho_2 \geq \frac{\rho^*}{1+\rho^*}.$$

Again, to have an equilibrium, it must be that the prior ρ is higher than the lower bound on the lowest belief. Indeed, if this is not the case, then both posterior beliefs ρ_2 and ρ_1 will be lower than ρ. This will imply that $f_m^B < f_m^S$ for $m \in \{1,2\}$. But this leads to a contradiction as both distributions must sum 1 over ratings. For the rest of the subsection, we impose another assumption.

Assumption 2.2. *The prior is higher than the lower bound on the lowest posterior:* $\rho \geq \frac{\rho^*}{1+\rho^*}$.

In order to find the optimal system for the customers, we would like to set $(f_1^B/f_1^S) = 2\lambda$ and $(f_2^B/f_2^S) = (1/2)(1-\rho^*)2\lambda$. Thus, we have to solve the system of equations

$$f_2^B = 1 - 2\lambda + 2\lambda f_2^S, \qquad f_2^B = (1-\rho^*)\lambda f_2^S.$$

The solution of this system is

$$f_2^B = \frac{\rho^*(1-\rho)+\rho^*-\rho}{\rho(1+\rho^*)}, \qquad f_2^S = \frac{\rho^*-\rho+\rho^*(1-\rho)}{(1-\rho)\rho^*(1+\rho^*)}.$$

Because $\frac{\rho^*}{1+\rho^*} \leq \rho \leq \bar{\rho}$ from assumption 2.1 and assumption 2.2 respectively, the steady-state probabilities are well defined. Recall that v from equation A.2.9 denotes customers' equilibrium payoff or the value from the interaction mediated by a rating system. Let $v(\rho)$ be this value for a given prior ρ. We find this value in proposition 2.5 below.

Proposition 2.5. *Customers' maximum equilibrium payoff is given by*

$$v(\rho) = \begin{cases} \rho\left(\frac{u-w}{2}\right) + (1-\rho)u & \text{if } \rho < \frac{\rho^*}{1+\rho^*}, \\ \left(\frac{\rho^*(1-\rho)+\rho^*-\rho}{\rho^*}\right)\left[\left(\frac{\rho^*}{1+\rho^*}\right)\frac{u-w}{2} + u\left(\frac{1}{1+\rho^*}\right)\right] & \text{if } \frac{\rho^*}{1+\rho^*} \leq \rho \leq \bar{\rho}, \\ 0 & \text{if } \rho > \bar{\rho}. \end{cases}$$

Proof. For $\frac{\rho^*}{1+\rho^*} < \rho \leq \bar{\rho}$, expert gets hired in rating 2, but not in rating 1. Then customers' equilibrium payoff will be

$$\begin{aligned} f_2 v(p_2) &= (\rho f_2^B + (1-\rho) f_2^S)[\rho_2\left(\frac{u-w}{2}\right) + (1-\rho_2)u], \\ &= \left(\frac{\rho^*(1-\rho)+\rho^*-\rho}{\rho^*}\right)\left[\left(\frac{\rho^*}{1+\rho^*}\right)\frac{u-w}{2} + u\left(\frac{1}{1+\rho^*}\right)\right]. \end{aligned}$$

□

We represent customers' equilibrium payoff in figure 3.3 below. The red line is the payoff in the one-shot interaction and the dashed line their equilibrium payoff under this optimal binary system.

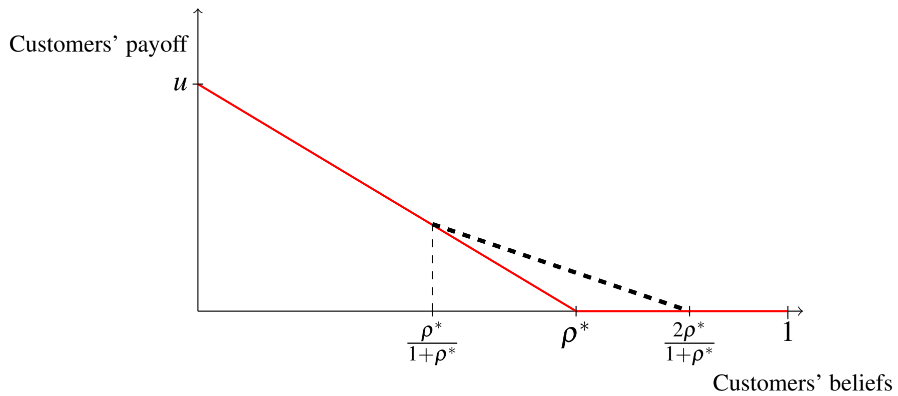

Figure 2.3 – customers' payoffs for the one-shot interaction and the equilibrium payoff of a binary, irreducible system. Red lines corresponds to the one-shot payoffs; the dashed thick line corresponds to what they get with the optimal system.

We can construct a binary, irreducible system that leads to the upper bound for customers' payoff. However, the particular construction will depend on parameters. We will have to deal with two cases. Whenever hiring is optimal in the one-shot interaction (that is, $\rho \leq \rho^*$), the system allows the strategic expert to stay more often in rating 2 if he provides treatment t_H, provided that he transitions from rating 1 to rating 2 with some probability. Whenever hiring is not optimal in the one-shot interaction (that is, $\rho > \rho^*$), the system always sends the expert to the non-hiring rating upon the observation of treatment t_H and the transition out of rating 1 occurs with a lower probability.

Remark 2.1. *For any given value of $\frac{\rho^*}{1+\rho^*} \leq \rho \leq \bar{\rho}$, we can construct an irreducible, binary system that (i) leads to an equilibrium in which customers hire in rating 2 but not in 2 and expert provides the correct treatment in rating 2; (ii) gives the highest equilibrium value from the interaction for the customers. We consider a system in which there is a random exit probability τ from rating 1 and a random exit probability κ from rating 2 upon the observation of t_H. Whenever t_L is played in rating 2, the system stays in 2*[14]*. The transition matrices will be*

$$T^B = \begin{array}{c c} & \begin{array}{cc} 1 & 2 \end{array} \\ \begin{array}{c} 1 \\ 2 \end{array} & \begin{pmatrix} 1-\tau & \tau \\ \kappa & 1-\kappa \end{pmatrix} \end{array} \qquad T^S = \begin{array}{c c} & \begin{array}{cc} 1 & 2 \end{array} \\ \begin{array}{c} 1 \\ 2 \end{array} & \begin{pmatrix} 1-\tau & \tau \\ \frac{\kappa}{2} & 1-\frac{\kappa}{2} \end{pmatrix} \end{array}.$$

This leads to the following values for the invariant distributions

$$f_2^B = \frac{\tau}{\kappa+\tau}, \qquad f_2^S = \frac{2\tau}{\kappa+2\tau}.$$

We need to find values of κ and τ such that $f_2^B/f_2^S = (1-\rho^*)\lambda$. This is satisfied when

$$\frac{\kappa}{\tau} = \frac{1-(1-\rho^*)\lambda}{(1-\rho^*)\lambda - 1/2}.$$

[14]The initial distribution φ_0 will not matter here because we will construct an irreducible system.

From assumption 2.1, we know that $\lambda \geq 1/2$; since $\rho^* > 0$, both observations lead to the denominator of the equation above being positive. From assumption 2.2, we have that $(1-\rho^*)\lambda \leq 1$, so the numerator is non-negative. There are some cases to consider. First, suppose $\lambda \geq 1$, that is, $\rho \leq \rho^*$. Then the both the numerator and the denominator will be lower than one, so it suffices to set $\kappa = 1-(1-\rho^*)\lambda$ and $\tau = (1-\rho^*)\lambda - 1/2$. Second, suppose $\lambda < 1$. Then we could set $\kappa = 1$ and

$$\tau = \frac{(1-\rho^*)\lambda - 1/2}{1-(1-\rho^*)\lambda}.$$

We depict the optimal system for the two cases in the figure below.

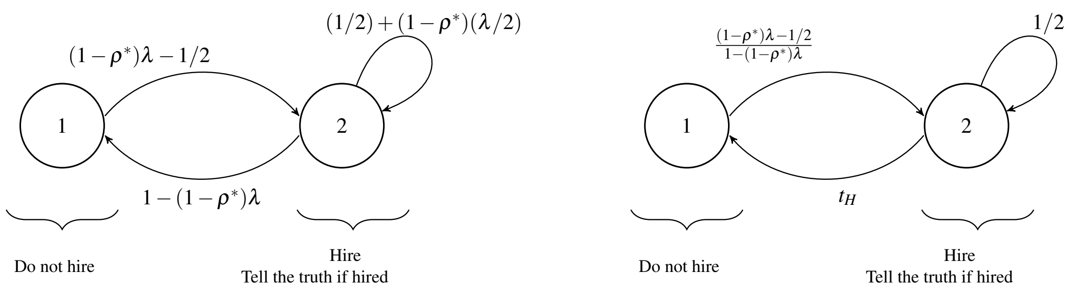

(a) Case 1: $\lambda \geq 1$

(b) Case 2: $\lambda < 1$

Figure 2.4 – Optimal binary, irreducible system for the customers. If hiring is optimal in the one-shot interaction (case 1), the expert stays in 2 with positive probability after the observation of treatment t_H. If hiring is not optimal in the one-shot game (case 2), the expert always transitions out of rating 2 after the observation of t_H.

2.3.2 Optimal system for the expert

If $\lambda \geq 1$, customer hires in the one-shot interaction ($\rho \leq \rho^*$). In this case, any system that does not induce any additional information is optimal for the principal. Take for instance $\varphi^H_{mm'} = \varphi^L_{mm'} = 1$ for $m \neq m'$. This yields $f^B = f^S = (1/2, 1/2)$ and $\rho_1 = \rho_2 = \rho$. Because customer gets hired in every rating, there is no reason not to tell the truth, so his equilibrium payoff is u, the highest possible indeed.

The more interesting case is when $\lambda < 1$. The expert would like to spend as much time as possible in rating 2 and the lowest possible in 1. Therefore, for the irreducible system, we would like to set $(f^B_2/f^S_2) = \lambda$ and $(f^B_1/f^S_1) = 2\lambda$. Thus, we have to solve the system of equations

$$f^B_1 = 1-\lambda+\lambda f^S_1, \qquad f^B_1 = 2\lambda f^S_1.$$

The solution of this system is

$$f^B_1 = 2(1-\lambda), \qquad f^S_1 = \frac{1-\lambda}{\lambda}.$$

This leads to the main result in this subsection. For a range of prior belief values - larger than $[0,\rho^*]$ - expert's maximum equilibrium payoff is bounded away from zero, even if he is sufficiently patient and if he is playing the truth-telling strategy.

Proposition 2.6. *For any prior $\rho \le \bar{\rho}$, expert's maximum equilibrium payoff is bounded away from zero, even as $\delta \to 1$. In this equilibrium, he always plays the truth-telling strategy in the hiring rating. The maximum payoff is given by*

$$\lim_{\delta \to 1} V_\delta(\rho) = \begin{cases} u & \text{if } \rho \le \rho^*, \\ \left[2 - \frac{\rho}{1-\rho}\frac{1-\rho^*}{\rho^*}\right] u & \text{if } \rho^* < \rho \le \bar{\rho}, \\ 0 & \text{if } \rho > \bar{\rho}. \end{cases}$$

Proof. Expert gets hired in rating 2 but not in rating 1. We also know that for $\rho > \bar{\rho}$, there will be no equilibrium in a binary, irreducible system, as no customer will hire. From proposition 2.3, we know that he will provide the correct treatment in 2, so his payoff will be u. As $\delta \to 1$, as discussed before, for an irreducible matrix T^S, expert's average discounted long-run frequency in every rating coincides with the distribution f^S. So for a patient expert, the equilibrium payoff will be $(f_2^S)u$. But

$$f_2^S = 1 - f_1^S = 2 - \frac{1}{\lambda} = 2 - \frac{\rho}{1-\rho}\frac{1-\rho^*}{\rho^*}.$$

□

We represent expert's equilibrium payoff for $\delta \approx 1$ in figure 3.5 below. The blue line gives the payoff in the one-shot interaction and the dashed thick line gives the average discounted payoff under the optimal binary system.

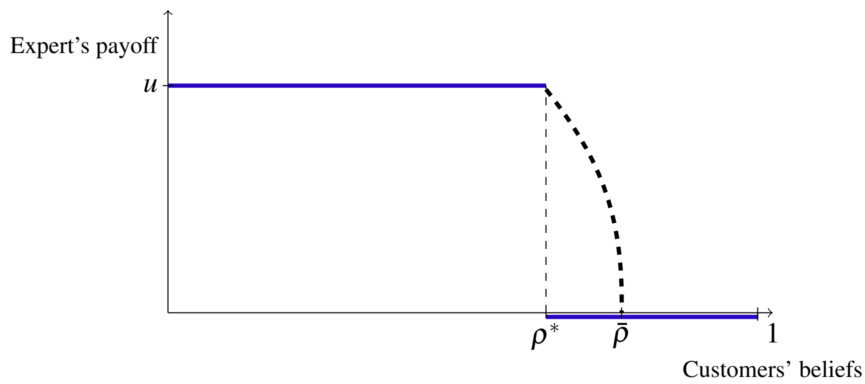

Figure 2.5 – Expert's payoffs for the one-shot interaction and the long-run interaction mediated through a binary, irreducible system. Blue lines corresponds to the one-shot payoffs; the dashed line corresponds to what he gets with the optimal system as $\delta \to 1$.

Can we construct a binary, irreducible system that actually leads to the upper bound for expert's payoff we derived in proposition 2.6? Next remark shows that indeed we can.

Remark 2.2. *For any given value of $\rho^* < \rho \le \bar{\rho}$, we can construct a binary and irreducible system that leads to the equilibrium and gives the highest value from the interaction for the expert. Consider the following system. There is a random exit probability τ from rating 1 and a random exit probability κ from rating 2 upon the observation of t_H. Whenever t_L is played in rating 2, the system stays in 2.*

If the strategy profile of not hiring in 1 and telling the truth in 2 is optimal, then the transition matrices will be

$$T^B = \begin{array}{c c} & \begin{array}{cc} 1 & 2 \end{array} \\ \begin{array}{c} 1 \\ 2 \end{array} & \begin{pmatrix} 1-\tau & \tau \\ \kappa & 1-\kappa \end{pmatrix} \end{array} \qquad T^S = \begin{array}{c c} & \begin{array}{cc} 1 & 2 \end{array} \\ \begin{array}{c} 1 \\ 2 \end{array} & \begin{pmatrix} 1-\tau & \tau \\ \frac{\kappa}{2} & 1-\frac{\kappa}{2} \end{pmatrix} \end{array}$$

This leads to the following values for the invariant distributions

$$f_2^B = \frac{\tau}{\kappa+\tau}, \qquad f_2^S = \frac{2\tau}{\kappa+2\tau}.$$

We need to find values of κ and τ such that $f_2^B/f_2^S = \lambda$. This is satisfied when

$$\frac{\tau}{\kappa} = \frac{\lambda-1/2}{1-\lambda}.$$

From assumption 2.1, we know that $\lambda \geq 1/2$; from the values of λ for which we are considering here, $\lambda < 1$. so both the numerator and the denominator are between 0 and 1. Therefore, it suffices to set $\tau = \lambda - 1/2$ and $\kappa = 1-\lambda$. We depict this optimal system in the figure below.

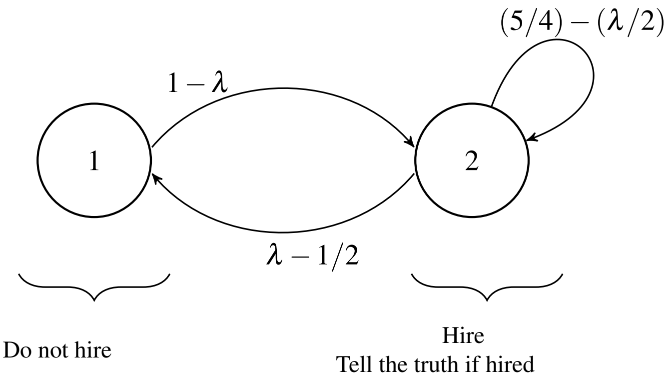

Figure 2.6 – Optimal binary, irreducible system for the expert. With probability $1-\lambda$, the system moves from the non-hiring rating to the hiring one and with probability $\lambda-1/2$, the expert is sent to the non-hiring rating after providing treatment t_H. The observation of treatment t_L never moves the system from 2 to 1. The strategy profile in which customers hire only in rating 2 and expert tells the truth whenever hired is an equilibrium.

2.4 Finite systems

We have characterized optimal binary, irreducible systems for the customers and the expert. Maintaining the irreducibility assumption, is it possible to improve players' equilibrium payoffs with more ratings? We investigate this now. We are going to show that more ratings are beneficial to customers: we characterize their maximum equilibrium payoff and show that it is increasing in the number of ratings. However, we will show that more ratings play no role on the maximization of expert's payoff (theorem 2.3).

We can construct a system that gets arbitrarily close to this maximum equilibrium payoff, but expert's discount factor matters when there is more than two ratings. To achieve the upper bound, expert must be playing truthfully whenever hired; but as he gets more patient, the system must be adjusted to give him the right incentives. For every discount $\delta \in [0,1)$, we can construct such system (theorem 2.2).

Abusing notation, we let $|M| = M$ and again without loss of generality we label ratings such that $\rho_1 \geq ... \geq \rho_M$. We now present a more general result regarding the upper bound in the posterior belief that a rating system can induce. The proposition was proved in Lorecchio and Monte (2021), but we include it here for completeness, since minor adaptations were required to adapt the proof to our bad reputation environment.

Proposition 2.7. *The highest posterior belief about expert being bad is bounded above: $\rho_1 \leq \bar{\rho}$.*

Below is the sketch of the proof. A formal proof is in the appendix. Note that the above mentioned labeling of the states is equivalent to assuming that

$$\frac{f_1^B}{f_1^S} \geq ... \geq \frac{f_M^B}{f_M^S}.$$

We will also use a well-known result - which we state it below - for Markov processes about transitioning in and out of a partition of the state space.

Lemma 2.1. *If the space M is partitioned into two sets C and C', then in the steady-state the probability of transition from C to C' must equal the probability of transition from C' to C:*

$$\sum_{m\in C}\sum_{n\in C'} f_m \tau_{mn} = \sum_{m\in C'}\sum_{n\in C} f_m \tau_{mn}.$$

Let us partition the rating space into a set $\{1,...,n-1\}$ for which hiring will not be optimal and another set $\{n,...,M\}$ for which hiring will take place. Lemma 1 implies that

$$\frac{\sum_{m=1}^{n-1}\sum_{m'=n}^{M} f_m^B \tau_{mm'}^B}{\sum_{m=1}^{n-1}\sum_{m'=n}^{M} f_m^S \tau_{mm'}^S} = \frac{\sum_{m=n}^{M}\sum_{m'=1}^{n-1} f_m^B \tau_{mm'}^B}{\sum_{m=n}^{M}\sum_{m'=1}^{n-1} f_m^S \tau_{mm'}^S}. \tag{A.2.14}$$

Recall that for every hiring rating m, we have $f_m^B \leq \lambda_m f_m^S$, where λ_m is defined as in equation A.2.12. Moreover, from equations A.2.3 and A.2.4, we know that $\tau_{mm'}^B \leq (1/\gamma_m)\tau_{mm'}^S$ for every $m' \neq m$. Finally, if m is a hiring rating, then we can be sure that $\gamma_m > 0$. Collecting results, we find that there is an upper bound on the right-hand side of the above equation, given by

$$\frac{\sum_{m=n}^{M}\sum_{m'=1}^{n-1} f_m^B \tau_{mm'}^B}{\sum_{m=n}^{M}\sum_{m'=1}^{n-1} f_m^S \tau_{mm'}^S} \leq \frac{\sum_{m=n}^{M}\left(\frac{\lambda_m}{\gamma_m}\right)\sum_{m'=1}^{n-1} f_m^S \tau_{mm'}^s}{\sum_{m=n}^{M}\sum_{m'=1}^{n-1} f_m^S \tau_{mm'}^S} \leq 2\lambda.$$

The last inequality follows from the highest value of the ratio (λ_m/γ_m) being 2λ, from previous discussions. From our assumption at the beginning of this section, $(f_m^B/f_m^S) \geq (f_{n-1}^B/f_{n-1}^S)$ for every $m \in \{1,...,n-1\}$. Because $\tau_{mm'}^B = \tau_{mm'}^S$ at no-hiring ratings, we find that there is a lower bound on the left-hand side of equation 3.14, given by

$$\frac{\sum_{m=1}^{n-1}\sum_{m'=n}^{M} f_m^B \tau_{mm'}^B}{\sum_{m=1}^{n-1}\sum_{m'=n}^{M} f_m^S \tau_{mm'}^S} \geq \frac{f_{n-1}^B}{f_{n-1}^S}\left[\frac{\sum_{m=n}^{M}\sum_{m'=1}^{n-1} f_m^S \tau_{mm'}^S}{\sum_{m=n}^{M}\sum_{m'=1}^{n-1} f_m^S \tau_{mm'}^S}\right] = \frac{f_{n-1}^B}{f_{n-1}^S}.$$

Combining results, we conclude that there is an upper bound on the relative frequency of the first rating for which hiring is not optimal.

$$\frac{f_{n-1}^B}{f_{n-1}^S} \leq 2\lambda. \tag{A.2.15}$$

Now partition the rating set into $\{1,...,n-1\}$ and $\{n,...,M\}$. Following the same steps of the argument above and using the bound in equation 15, we can show that the relative frequency (f_{n-2}^B/f_{n-2}^S) will also be lower than 2λ. Proceeding recursively, we find that the highest possible relative frequency (f_1^B/f_1^S) is bounded above by 2λ.

2.4.1 Optimal system for customers

Ideally, customers would like to hold extreme posterior beliefs about the expert being bad: they would benefit from knowing for sure that they are hiring strategic experts and not hiring bad ones. However, from proposition 2.7, we know that the highest belief about expert being bad we can possible achieve is $\bar{\rho}$. This creates an upper bound on customers' maximum equilibrium payoff. We characterize such bound in next proposition.

Proposition 2.8. *In a finite system, for every prior $\rho \leq \bar{\rho}$, customers' maximum equilibrium payoff is bounded above:*

$$v(\rho) \leq u\left[\frac{\bar{\rho}-\rho}{\bar{\rho}}\right].$$

Proof. As before, suppose non-hiring ratings belong to the set $\{1,...,n-1\}$ and hiring ones to $\{n,...,M\}$, where $n \geq 2$. Whenever customers do not hire, they get a zero payoff. Whenever they hire, they get payoff $v(\rho_m)$, as given by equation A.2.9. Therefore, for a system with $M \geq 2$ ratings,

$$
\begin{aligned}
v &= \sum_{m=n}^{M} f_m \nu(\rho_m), \\
&= \sum_{m=n}^{M} f_m \left[\rho_m \left(\frac{u-w}{2}\right) + (1-\rho_m)\,(\beta_m u - (1-\beta_m)w)\right], \\
&\leq u \sum_{m=n}^{M} f_m.
\end{aligned} \tag{A.2.16}
$$

The inequality follows from u being the highest value for $\nu(\rho_m)$ - the expression in brackets - for any ρ_m. Ideally, this is achieved if every hiring rating has $\rho_m = 0$ and $\beta_m = 1$, that is, customers know they are dealing with a strategic expert and expert tells the truth whenever hired. Note that the irreducible equilibrium must satisfy $\sum_{m \in M} f_m \rho_m = \rho$. If we partition the rating space into non-hiring ratings and hiring ones, every posterior belief in the non-hiring ratings must be weakly lower than $\bar{\rho}$; every posterior belief in the hiring ratings must ideally equal 0. This implies

$$
\rho = \sum_{i=1}^{n-1} f_m \rho_m \leq \bar{\rho}\left[1 - \sum_{m=n}^{M} f_m\right] \Rightarrow \sum_{m=n}^{M} f_m \leq \frac{\bar{\rho}-\rho}{\bar{\rho}}. \tag{A.2.17}
$$

Substituting equation 3.16 into equation 3.17 leads to the desired bound.

□

The upper bound is best seen in figure 3.7 below. The red line represents the payoff under the one-shot interaction and the dashed line the bound derived in previous proposition. Comparing it with figure 2.3, we see that a binary rating constraints the lowest posterior belief that could arise in equilibrium. We will argue in the remaining of this subsection that we can come up with a system in which the lowest possible posterior belief decreases with the number of ratings.

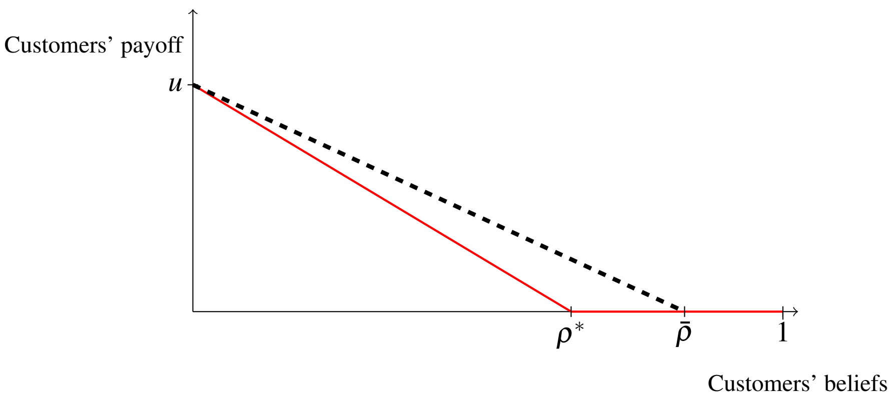

Figure 2.7 – Customers' equilibrium payoffs for the one-shot interaction and the upper bound equilibrium payoff for finite, irreducible systems. Red lines corresponds to the one-shot payoffs; the dashed line corresponds to the highest value they can get from the interaction with an expert.

For every $\rho \in (0, \bar{\rho}]$, we can construct a system that gets arbitrarily close to the payoff bound and provides right incentives for the expert to tell the truth whenever hired, for every discount factor $\delta < 1$.

What is more, this system improves upon a binary system and extreme memory systems (the one-shot and the full-memory interactions), as customers are hiring with more information in the long run, without triggering incentives for the expert to provide the wrong treatment. However, expert's discount factor will mater for constructing the system. The construction is given in the proof of next theorem, which is in the appendix, but we provide a sketch here.

Consider a system with M ratings (M will be chosen endogenously later on) in which the expert gets hired in every rating except the first one. At rating 1, there is an exit probability τ to rating 2. At rating M, there exists a downgrade probability κ to $M-1$ whenever the high-cost treatment is observed. Under a low-cost treatment at M, the system stays there. Every intermediary rating is downgraded (upgraded) with same probability under a high cost (low cost) treatment. Those probabilities are chosen so that (i) the expert is downgraded with the same probability as he is upgraded, if he tells the truth; (ii) the probability of an upgrade (downgrade) is higher (lower) than the probability of a downgrade (upgrade) if t_L (t_H) is observed.

We can set τ, κ and the intermediary transitions to have the highest posterior $\rho_1 = \bar{\rho}$ and the lowest posterior $\rho_M = \underline{\rho}(M)$ even if $\tau \approx 0$ and $\kappa \approx 0$. The lower bound $\rho(M)$ is defined as

$$\underline{\rho}(M) := \frac{\rho^* \left(\frac{1-\rho^*}{2}\right)^{M-2}}{1+\rho^* \left(\frac{1-\rho^*}{2}\right)^{M-2}}.$$

Note that, for $M=2$, we have $\underline{\rho}(2) = \frac{\rho^*}{1+\rho^*}$, the lower bound on posterior belief we have obtained when discussing optimal binary systems for customers. Note as well that, for any $\rho \in (0,\bar{\rho}]$, we can find a value of M high enough to have $\rho \geq \underline{\rho}(M)$. Recall that there was no equilibrium for prior beliefs below $\underline{\rho}(2)$. If that was the case, both posterior beliefs would be higher than the prior, leading to a contradiction. The same argument holds for M ratings, as proven in lemma 2.2 in the appendix.

To give the expert incentives to tell the truth for every discount factor, we increase the time he spends in each rating as δ increases. As long as the intermediary transitions are chosen carefully, the posterior beliefs in extreme ratings are well defined. We know that truth-telling is optimal in the one-shot interaction, so by making intermediary ratings relatively more persistent (but always less than extreme ratings), expert finds optimal to provide right treatment whenever hired. After adjusting all parameters to satisfy incentive compatibility of both players, by increasing the number of ratings M, we can get customers' equilibrium payoff to be arbitrarily close to the upper bound from proposition 2.8.

Theorem 2.2. *For any prior belief $\rho \in (0,\bar{\rho}]$, any discount factor $\delta \in (0,1)$ and every $\varepsilon > 0$, there exists a number $\bar{M}$ and a finite rating system with $M \geq \bar{M}$ ratings that is ε-close to the customers' payoff bound:*

$$v(\rho) \geq u\left(\frac{\bar{\rho}-\rho}{\bar{\rho}}\right) - \varepsilon.$$

In it, customers hire in every rating except 1 and expert tells the truth whenever hired.

2.4.2 Optimal system for the expert

The bound on the highest posterior in proposition 2.7 is exactly the same one for a binary system. So more ratings cannot benefit the expert: simple hiring recommendations are sufficient to implement the optimal rule for the expert. This leads to one of our main results, which is a summary of previous propositions.

Theorem 2.3. *The optimal system for the expert requires only two ratings and eliminates the bad reputation effect: expert plays a truth-telling strategy whenever hired, customers hire in one rating. For $\rho \leq \bar{\rho}$, expert's equilibrium payoff is bounded away from zero, even if $\delta \to 1$.*

The main intuition here comes from the interaction of two underlying reasons. First, having only two non-trivial ratings means that in one of them there is hiring, but not in the other. Second, the expert plays truthfully in both states. Thus, the expert's long-run payoff depends on her frequency of being hired, but is independent of the actual induced belief. In simple terms, if she is hired, she does the right thing and gets a payoff accordingly.

2.5 Reducible systems

We have focused on irreducible systems, but can we come up with reducible systems that benefit the expert or the customers? In this section, we argue that there is no gain from such systems, so the irreducibility assumption was without loss of generality indeed.

We begin by discussing reducible systems for the expert. Note that if there is only one recurrent class, the highest possible posterior belief about expert being bad must be bounded above by $\bar{\rho}$, as this class must have the same properties we derive in previous sections. So suppose there are multiple recurrent classes. Again, if $\rho \leq \rho^*$, any uninformative system would be optimal for the expert, so we focus on $\rho > \rho^*$.

It cannot be that recurrent classes are composed of hiring ratings only. If that was the case, all posterior beliefs would be lower than the prior, and lemma 2.2 in the appendix shows that this is not possible. Thus, there exists at least one rating in which not hiring is optimal for customers in each recurrent class. If all recurrent classes are composed of hiring and non-hiring ratings, then they all share the same bound on the belief spread we derived in previous sections. Therefore, the reducible system cannot increase expert's equilibrium payoff above the payoff he gets from the binary, irreducible optimal system.

It remains to argue that there is no improvement with a recurrent class of only non-hiring ratings. In it, all posterior beliefs must be equal. This happens because the transitions from one rating to another does not depend on any additional information (expert does not provide any treatment). Because this is a recurrent class, the stationary distribution of it does not depend on the initial distribution of the system.

Ratings that transition into this class are either hiring or non-hiring ones. If it is a non-hiring rating, then the transition to the recurrent class brings no information; and the posterior in this non-hiring rating must be equal to the posteriors in the recurrent class. If it is a hiring one, the highest possible posterior in it is ρ^* (there is no need to induce a higher posterior).

Abusing notation, let $\tau^i_{m,\mathscr{R}_1}$ represent the transition from this hiring rating m to any rating in the non-hiring recurrent class $\mathscr{R}_1$ if expert is of type $i \in \{B,G\}$. The likelihood probability of expert being bad (relative to being strategic) conditional on the system being at $\mathscr{R}_1$ and the rating that led to $\mathscr{R}_1$ was m is

$$\frac{\rho_m}{1-\rho_m}\frac{\tau^B_{m,\mathscr{R}_1}}{\tau^G_{m,\mathscr{R}_1}} \leq \frac{2\rho^*}{1-\rho^*} = \frac{\bar{\rho}}{1-\bar{\rho}}.$$

This means that the posterior belief in $\mathscr{R}_1$ must be bounded above by the same bound $\bar{\rho}$ we derived with irreducible systems. Thus, it proves that a recurrent class of non-hiring ratings only cannot benefit the expert and concludes the argument about irreducibility being without loss of generality for the expert.

We now discuss reducible systems to maximize customers' payoff. Similarly, we only need to investigate multiple recurrent classes. We start with $\rho > \rho^*$. Same argument in previous paragraphs leads to the conclusion that there exists at least one rating in which not hiring is optimal for customers in each recurrent class. Analogous argument also proves that multiple recurrent classes with hiring and non-hiring ratings in it cannot improve upon the optimal irreducible, finite system.

Let $\mathscr{R}_1$ represent a recurrent class composed of non-hiring ratings only. If there exists a non-hiring rating outside $\mathscr{R}_1$ that transitions to this set, the posterior beliefs in $\mathscr{R}_1$ cannot be higher than the posterior belief in m, which is at most $\bar{\rho}$. If there exists a hiring rating m' outside $\mathscr{R}_1$ that transitions to this set, the posterior belief must at most ρ^*. Therefore, we can conclude that any posterior belief in $\mathscr{R}_1$ must be at most $\bar{\rho}$, implying that reducible systems cannot improve customers' payoff.

Finally, suppose that $\rho \leq \rho^*$. It cannot be that all recurrent classes are composed of non-hiring ratings exclusively. We only need to investigate reducible systems in which there exists a recurrent class $\mathscr{R}_2$ composed of hiring ratings. We will argue that, if there is an improvement with $\mathscr{R}_2$ composed of $|M_2| < \infty$ hiring ratings, then we could reproduce such improvement by increasing the number of rating in our irreducible, finite system.

Let $v_2(\rho)$ be customers' equilibrium payoff in recurrent class $\mathscr{R}_2$ and $v(\rho)$ be customers' equilibrium payoff from an irreducible system. Suppose $v_2(\rho) > v(\rho)$. Any posterior belief in $\mathscr{R}_2$ must be weakly lower than ρ^* and higher than 0. Indeed, to have a zero belief, it must be that the bad type of expert never visits one rating m in this class. But this is only possible if strategic expert never plays the high-cost treatment in ratings that transition to m - which contradicts hiring as an optimal response - or if strategic expert never visits this rating - which contradicts $\mathscr{R}_2$ being a recurrent class.

Because the highest possible posterior in $\mathscr{R}_2$ is weakly lower than ρ^* and the lowest possible belief is higher than zero, we can adjust the parameters and the number of ratings in the optimal irreducible system to be ε-higher than $v_2(\rho)$, since the system can be arbitrarily close to the maximum belief spread 0 and $\bar{\rho}$. This concludes that a reducible system cannot increase customers' equilibrium payoff relative to the optimal irreducible system.

2.6 Conclusion

People rely on online reputation systems to hire experts, but such systems can generate wrong incentives for experts and mistrust from customers. We consider the work by Ely and Välimäki (2003) regarding the stylized interaction between those players, but with an information intermediary. Without any informational restriction, the expert cannot avoid being rarely hired in equilibrium since either the belief about him being bad is eventually high enough that customers no longer hire, or the temptation to reveal himself by lying to a customer is so strong that market collapses from the start.

With rating systems, we have shown that there is a bound on the equilibrium payoffs - which we characterize it in terms of posterior belief spread (proposition 2.2). Nevertheless, bad reputation is avoidable, that is, the strategic expert tells the truth whenever she is hired and gets hired with positive probability, even in the long run.

From the customers' perspective, more ratings generate higher equilibrium payoffs. We then construct finite and irreducible systems to arbitrarily approach customers' equilibrium payoff upper bound (theorem 2.2). The system endogenously generates inflation: the interaction takes place most times at extreme ratings, even though intermediate ratings are important to further separate types. This is consistent with empirical observation.

From the expert's perspective, a simple binary and irreducible rating system is sufficient for optimality (theorem 2.3). This follows from non-hiring ratings being uninformative - no treatment is observed - and the strategic expert only caring about being hired in equilibrium.

2.7 Appendix

Proposition 2.7. *The highest posterior belief about expert being bad is bounded above: $\rho_1 \leq \bar{\rho}$.*

Proof. Partition the rating space into $\{1,...,n-2\}$ and $\{n-1,...,M\}$. From lemma 3.1,

$$\frac{\sum_{m=1}^{n-2}\sum_{m'=n-1}^{M} f_m^B \tau_{mm'}^B}{\sum_{m=1}^{n-2}\sum_{m'=n-1}^{M} f_m^S \tau_{mm'}^S} = \frac{\sum_{m=n-1}^{M}\sum_{m'=1}^{n-2} f_m^B \tau_{mm'}^B}{\sum_{m=n-1}^{M}\sum_{m'=1}^{n-2} f_m^S \tau_{mm'}^S}. \tag{A.3.1}$$

For every hiring rating $m \in \{n,..,M\}$, we have $f_m^B \leq \lambda_m f_m^S$ and $\tau_{mm'}^B \leq \frac{1}{\gamma_m}\tau_{mm'}^S$. Since the highest value of the ratio $\frac{\lambda_m}{\gamma_m}$ is 2λ, it follows that $f_m^B \tau_{mm'}^B \leq 2\lambda f_m^S \tau_{mm'}^S$ for every $m \in \{n,...,M\}$ and $m' \in \{1,..,n-2\}$. From equation A.2.15 and the fact that $\tau_{n-1m'}^B = \tau_{n-1m'}^S$, we have $f_{n-1}^B \tau_{n-1m'}^B \leq 2\lambda f_{n-1}^S \tau_{n-1m'}^S$ for every $m' \in \{1,...,n-2\}$. Therefore, the right-hand side of the above equation is bounded above:

$$\frac{\sum_{m=n-k}^{M}\sum_{m'=1}^{n-k-1} f_m^B \tau_{mm'}^B}{\sum_{m=n-k}^{M}\sum_{m'=1}^{n-k-1} f_m^S \tau_{mm'}^S} \leq 2\lambda.$$

By construction, $f_m^B \geq \left(\frac{f_{n-2}^B}{f_{n-2}^S}\right) f_m^S$ for every $m \in \{1,..,n-2\}$. Because there is no hiring in these ratings, $\tau_{mm'}^B = \tau_{mm'}^S$ for every $m' \in \{n-1,...,M\}$. Therefore, the left hand side of equation A.3.1 is bounded below:

$$\frac{\sum_{m=1}^{n-2}\sum_{m'=n-1}^{M} f_m^B \tau_{mm'}^B}{\sum_{m=1}^{n-2}\sum_{m'=n-1}^{M} f_m^S \tau_{mm'}^S} \geq \frac{f_{n-2}^B}{f_{n-2}^S}.$$

Combining results, we see that

$$\frac{f_{n-2}^B}{f_{n-2}^S} \leq 2\lambda.$$

Proceeding by induction, one can show that $\frac{f_{n-k}^B}{f_{n-k}^S} \leq 2\lambda$ for every $k \in \{1,...,n-1\}$. This leads to

$$\frac{\rho_1}{1-\rho_1} = \left(\frac{\rho}{1-\rho}\right)\frac{f_1^B}{f_1^S} \leq 2\left(\frac{\rho}{1-\rho}\right)\frac{\rho^*}{1-\rho^*}\frac{1-\rho}{\rho} \Rightarrow \rho_1 \leq \frac{2\rho^*}{1+\rho^*} := \bar{\rho}.$$

□

Lemma 2.2. *At least one rating has a belief lower than the prior (there exists some m such that $\rho_m \leq \rho$) and at least one rating has a belief higher than the prior (there exists some m' such that $\rho_{m'} \geq \rho$)*[15].

Proof. Assume by way of contradiction that $\rho_m > \rho$ for every $m \in M$. This implies that $f_m^B > f_m^S$. Summing over the ratings, that leads to $1 = \sum_{m\in M} f_m^B > \sum_{m\in M} f_m^S = 1$, a contradiction. Similar argument holds when assuming by way of contradiction that $\rho_m < \rho$ for every $m \in M$. □

Theorem 2.3. *For any prior belief $\rho \in (0,\bar{\rho}]$, any discount factor $\delta \in (0,1)$ and every $\varepsilon > 0$, there exists a number $\bar{M}$ and a finite rating system with $M \geq \bar{M}$ ratings that is ε-close to the customers' payoff bound:*

$$v(\rho) \geq u\left(\frac{\bar{\rho}-\rho}{\bar{\rho}}\right) - \varepsilon.$$

In it, customers hire in every rating except 1 and expert tells the truth whenever hired.

Proof. We proceed through the following steps: (I) for M_0 high enough, we construct the system with $M \geq M_0$ that approaches the belief spread - $\rho_1 = \bar{\rho}$ and $\underline{\rho}(M_0)$; (II) we adjust parameters so that truth-telling is optimal whenever hired for every $\delta < 1$; (III) with the adjusted parameters, we show that the induced payoff is ε-close to customers' maximum equilibrium payoff.

I.i - constructing the system that approaches the maximum belief spread. Choose M_0 big enough so that $\rho \geq \underline{\rho}(M_0)$, where $\underline{\rho}(M_0)$ is given by

$$\underline{\rho}(M_0) := \frac{\rho^*\left(\frac{1-\rho^*}{2}\right)^{M_0-2}}{1+\rho^*\left(\frac{1-\rho^*}{2}\right)^{M_0-2}}. \quad \text{(A.3.2)}$$

[15] This is an adaptation from a result in Monte (2013).

As a corollary of lemma 2.2, as long as $\rho \in [\underline{\rho}(M_0), \bar{\rho}]$, an equilibrium with hiring in every rating except 1 is at least possible. We will use the following quantities throughout the exposition:

$$\varphi_+^H = \varphi_-^L := \left(\frac{1-\rho^*}{2}\right)\psi, \qquad \varphi_-^H = \varphi_+^L := \psi, \tag{A.3.3}$$

where $\psi \in (0,1)$ will be chosen appropriately throughout the proof.

We would like to implement a strategy in which customers hire in every rating except 1 and expert plays truthfully at every hiring rating. To do so, we will work with the following transition rules. At rating 1, there is a random exit probability τ to rating 2. At any $m \notin \{1, M_0\}$, (a) whenever treatment t_H is observed, the rating is upgraded with probability φ_+^H and downgraded with probability φ_-^H; (b) whenever t_L is observed, the rating is upgraded with probability φ_+^L and downgraded with probability φ_-^L. At rating M_0, treatment t_H leads to a downgrade with probability 2κ and treatment t_L does not change the rating. Similar to ψ, the parameters τ and κ are chosen appropriately through the proof.

The transition probabilities satisfy the following properties for every rating $m \notin \{1, M_0\}$: (a) $\frac{\tau^B_{mm+1}}{\tau^B_{m+1m}} = \frac{\varphi_+^H}{\varphi_-^H} = \frac{1-\rho^*}{2}$; (b) $\tau^S_{mm+1} = \frac{1}{2}\varphi_+^H + \frac{1}{2}\varphi_+^L = \frac{1}{2}\varphi_-^H + \frac{1}{2}\varphi_-^L = \tau^S_{m+1m}$ and (c) $\varphi_+^L - \varphi_+^H = \varphi_-^H - \varphi_-^L$. This third property will be useful when adjusting parameters so that strategic expert has the right incentives to tell the truth whenever hired. We will refer to the equal probabilities of strategic expert moving to adjacent rating as τ^S. We have the following identities.

$$f_1^B\tau = \psi f_2^B, \qquad f_{m+1}^B = \left(\frac{1-\rho^*}{2}\right) f_m^B, \qquad f_{M_0}^B\kappa = \left(\frac{1-\rho^*}{2}\right) f_{M_0-1}^B;$$
$$f_1^S\tau = \tau^S f_2^S, \qquad f_{m+1}^S = f_m^S, \qquad f_{M_0}^S\kappa = \tau^S f_{M_0-1}^S.$$

That leads to the following relative frequencies.

$$\frac{f_2^B}{f_2^S} = \frac{1}{2}\frac{f_1^B}{f_1^S}, \qquad \frac{f_m^B}{f_m^S} = \left(\frac{1-\rho^*}{2}\right)^{m-2}\frac{f_2^B}{f_2^S}, \qquad \frac{f_{M_0}^B}{f_{M_0}^S} = \frac{1}{\tau^S}\left(\frac{1-\rho^*}{2}\right)\frac{f_{M_0-1}^B}{f_{M_0-1}^S}.$$

We can find the values of f_1^B and f_1^S through the following system of equations.

$$\begin{aligned} 1 =& f_1^B + \frac{\tau}{\psi} f_1^B + \left(\frac{1-\rho^*}{2}\right)\left[\frac{\tau}{\psi}\right] f_1^B + \dots + \left(\frac{1-\rho^*}{2}\right)^{M_0-3}\left[\frac{\tau}{\psi}\right] f_1^B + \frac{\tau}{\kappa}\left(\frac{1-\rho^*}{2}\right)^{M_0-1} f_1^B, \\ 1 =& f_1^S + \left[\frac{\tau}{\tau^S}\right] f_1^S + \left[\frac{\tau}{\tau^S}\right] f_1^S + \dots + \left[\frac{\tau}{\tau^S}\right] f_1^S + \left(\frac{\tau}{\kappa}\right) f_1^S. \end{aligned}$$

The solutions are

$$f_1^B = \frac{1}{1+\frac{\tau}{\psi}\left(\frac{1-\left(\frac{1-\rho^*}{2}\right)^{M_0-2}}{1-\left(\frac{1-\rho^*}{2}\right)}\right)+\frac{\tau}{\kappa}\left(\frac{1-\rho^*}{2}\right)^{M_0-1}}, \tag{A.3.4}$$

$$f_1^S = \frac{1}{1+\frac{\tau}{\tau^S}(M_0-2)+\frac{\tau}{\kappa}}. \tag{A.3.5}$$

We want to have $\tau \approx 0$, $\kappa \approx 0$ but we also want to keep $\frac{\tau}{\kappa} > 0$ as well as $\frac{\tau}{\psi} \approx 0$. That will lead to intermediary ratings been rarely visited by both types in equilibrium. Indeed, because f_2^B and f_2^S respectively are

$$f_2^B = \frac{\frac{\tau}{\psi}}{1+\frac{\tau}{\psi}\left(\frac{1-\left(\frac{1-\rho^*}{2}\right)^{M_0-2}}{1-\left(\frac{1-\rho^*}{2}\right)}\right)+\frac{\tau}{\kappa}\left(\frac{1-\rho^*}{2}\right)^{M_0-1}},$$

$$f_2^S = \frac{\frac{\tau}{\tau^S}}{1+\frac{\tau}{\tau^S}(M_0-2)+\frac{\tau}{\kappa}}.$$

As $\frac{\tau}{\psi} \approx 0$, we will have $\frac{\tau}{\tau^S} \approx 0$ and $f_2^B \approx f_2^S \approx 0$. From our derivation of the stationary probabilities, it will also be the case $f_m^B \approx f_m^S \approx 0$ for $m \in \{3,...,M_0-1\}$. However, for rating M_0,

$$f_{M_0}^B = \frac{\frac{\tau}{\kappa}\left(\frac{1-\rho^*}{2}\right)^{M_0-1}}{1+\frac{\tau}{\psi}\left(\frac{1-\left(\frac{1-\rho^*}{2}\right)^{M_0-2}}{1-\left(\frac{1-\rho^*}{2}\right)}\right)+\frac{\tau}{\kappa}\left(\frac{1-\rho^*}{2}\right)^{M_0-1}}, \qquad f_{M_0}^S = \frac{\frac{\tau}{\kappa}}{1+\frac{\tau}{\tau^S}(M_0-2)+\frac{\tau}{\kappa}}.$$

So even if $\frac{\tau}{\psi} \approx 0$, as long as $\frac{\tau}{\kappa} > 0$, those probabilities will be positive. We also want to set $\frac{f_1^B}{f_1^S} = 2\lambda$. Using equations A.3.4 and A.3.5, we can find the value of the ratio $\frac{\tau}{\kappa}$ that leads to ratio 2λ for $\frac{\tau}{\psi} \approx 0$. This value is

$$\frac{\tau}{\kappa} = \frac{2\lambda-1}{1-2\lambda\left(\frac{1-\rho^*}{2}\right)^{M_0-1}}. \tag{A.3.6}$$

Because $\rho \leq \bar{\rho}$ implies $\lambda \geq 1/2$, the numerator is positive. For $M_1 \geq M_0$ big enough, the denominator will be positive as well. Thus, this ratio is well defined. With this system, we achieve $\rho_1 = \bar{\rho}$ and $\rho_{M_1} = \underline{\rho}(M_1)$ as $\tau \approx 0$, $\kappa \approx 0$ and $\frac{\tau}{\psi} \approx 0$. To see that the belief ρ_{M_1} is equal to the lower bound as in equation A.3.2, note that the relative frequency $\frac{f_{M_1}^B}{f_{M_1}^S}$ is

$$
\begin{aligned}
\frac{f^B_{M_1}}{f^S_{M_1}} &= \left(\frac{1-\rho^*}{2}\right)^{M_1-1}\left[\frac{1+\frac{\tau}{\kappa}}{1+\frac{\tau}{\kappa}\left(\frac{1-\rho^*}{2}\right)^{M_1-1}}\right],\\
&= \left(\frac{1-\rho^*}{2}\right)^{M_1-1}\left[\frac{1-2\lambda\left(\frac{1-\rho^*}{2}\right)^{M_1-1}+2\lambda-1}{1-\left(\frac{1-\rho^*}{2}\right)^{M_1-1}}\right],\\
&= \left(\frac{1-\rho^*}{2}\right)^{M_1-1}2\lambda\left[\frac{1-\left(\frac{1-\rho^*}{2}\right)^{M_1-1}}{1-\left(\frac{1-\rho^*}{2}\right)^{M_1-1}}\right],\\
&= \left(\frac{1-\rho^*}{2}\right)^{M_1-1}2\lambda.
\end{aligned}
$$

Recalling that $\rho_{M_1} = \frac{\rho f^B_{M_1}}{(1-\rho)f^S_{M_1}}$ and $\lambda = \frac{(1-\rho)\rho^*}{\rho(1-\rho^*)}$, simple manipulation shows that

$$
\rho_{M_1} = \frac{\rho^*\left(\frac{1-\rho^*}{2}\right)^{M_1-2}}{1+\rho^*\left(\frac{1-\rho^*}{2}\right)^{M_1-2}} := \underline{\rho}(M_1).
$$

II - adjusting parameters so that truth-telling is optimal for every $\boldsymbol{\delta < 1}$. We have assumed as an equilibrium candidate a strategy profile in which customers hire in every rating except 1 and expert plays the truth-telling strategy whenever hired. It remains to adjust parameter values so that this profile is an equilibrium indeed.

First, note that at rating 1 not hiring is optimal indeed, because $\frac{f^B_1}{f^S_1} > \lambda$. In rating 2, if expert provides the right treatment, then hiring is optimal, because $\frac{f^B_2}{f^S_2} = \lambda$. For every other rating $m \geq 2$, we will have $\frac{f^B_2}{f^S_2} < \lambda$, so hiring is optimal in all of them. Therefore, customers' strategy is incentive compatible.

Before deriving parameter values for which truth-telling is an optimal strategy for the strategic expert, the following claim will be useful. We claim that expert's continuation value in every rating is bounded below by $-w$. To see this, let $\underline{m} \in \arg\min\{V_\delta(m) : 1 \leq m \leq M_1\}$. From equations A.2.3 and A.2.7, the continuation value at $\underline{m}$ is such that

$$
\begin{aligned}
V_\delta(\underline{m}) &= (1-\delta)\left[\alpha\left(\beta_m u-(1-\beta_m)w\right)\right]+\sum_{m=1}^{M_1}\tau^S_{m,m'}V_\delta(m'),\\
&\geq (1-\delta)(-w)+\delta V_\delta(\underline{m}).
\end{aligned}
$$

This implies that $V_\delta(\underline{m}) \geq -w$; so expert's continuation value in every rating is bounded below by $-w$. With that in mind, we start the analysis with M_1. We have

$$
V_\delta(M_1) = (1-\delta)u+\delta\left[(1-\kappa)V_\delta(M_1)+\kappa V_\delta(M_1-1)\right].
$$

Subtracting $V_\delta(M_1-1)$ on both sides and rearranging,

$$V_\delta(M_1)-V_\delta(M_1-1)=\left(\frac{1-\delta}{1-\delta+\delta\kappa}\right)[u-V_\delta(M_1-1)].$$

The right hand side of the above equation is at most $u+w$. If expert provides the correct treatment t_H at problem H, he gets

$$(1-\delta)u+\delta\left\{2\kappa[V_\delta(M_1-1)-V_\delta(M_1)]+V_\delta(M_1)\right\}.$$

and if he provides treatment t_L instead, he gets

$$(1-\delta)(-w)+\delta V_\delta(M_1).$$

Therefore, truth-telling is optimal if and only if $\frac{1-\delta}{\delta}(u+w)\geq 2\kappa[V_\delta(M_1)-V_\delta(M_1-1)]$. Because $[V_\delta(M_1)-V_\delta(M_1-1)]\leq u+w$, as long as $2\kappa\leq\frac{1-\delta}{\delta}$, telling the truth at rating M_1 will be optimal indeed. Let us check now whether telling the truth is optimal at rating M_1-1. We have

$$V_\delta(M_1-1)=(1-\delta)u+\delta\left\{\tau^S[V_\delta(M_1-2)-V_\delta(M_1-1)]+V_\delta(M_1-1)+\tau^S[V_\delta(M_1)-V_\delta(M_1-1)]\right\}.$$

Subtracting $V_\delta(M_1-2)$ on both sides and rearranging,

$$V_\delta(M_1-1)-V_\delta(M_1-2)=\left(\frac{1-\delta}{1-\delta+\delta\tau^S}\right)[u-V_\delta(M_1-2)]+\left(\frac{\delta\tau^S}{1-\delta+\delta\tau^S}\right)[V_\delta(M_1)-V_\delta(M_1-1)].$$

The right hand side of this equation is at most $u+w$ because $\tau^S>0$, $V_\delta(M_1-2)\geq -w$ and $V_\delta(M_1)-V_\delta(M_1-1)\leq u+w$. If expert provides the correct treatment t_H at problem H, he gets

$$(1-\delta)u+\delta\left\{\varphi_-^H[V_\delta(M_1-2)-V_\delta(M_1-1)]+V_\delta(M_1-1)+\varphi_+^H[V_\delta(M_1)-V_\delta(M_1-1)]\right\}.$$

And if he provides treatment t_L instead, he gets

$$(1-\delta)(-w)+\delta\left\{\varphi_-^L[V_\delta(M_1-2)-V_\delta(M_1-1)]+V_\delta(M_1-1)+\varphi_+^L[V_\delta(M_1)-V_\delta(M_1-1)]\right\}.$$

Therefore, truth-telling is optimal if and only if

$$\frac{1-\delta}{\delta}(u+w) \geq (\varphi_+^L - \varphi_+^H)\left[V_\delta(M_1) - V_\delta(M_1-1)\right] + (\varphi_-^H - \varphi_-^L)\left[V_\delta(M_1-1) - V_\delta(M_1-2)\right].$$

From our previous discussion, we know that $[V_\delta(M_1) - V_\delta(M_1-1)] \leq u+w$ as well as $[V_\delta(M_1-1) - V_\delta(M_1-2)] \leq u+w$. From the values of φ_-^L and φ_+^H, we know that $\varphi_+^L - \varphi_+^H = \varphi_-^H - \varphi_-^L = \psi\left(\frac{1+\rho^*}{2}\right)$. Thus, to satisfy the incentive constraint, it must be that

$$\psi \leq \frac{1}{1+\rho^*}\left(\frac{1-\delta}{\delta}\right). \tag{A.3.7}$$

Assume that this is the case. We now proceed by induction. Suppose we have proved that truth-telling is optimal for $M_1, M_1-1, ..., M_1-m$. We would like to show that truth-telling is optimal for $M_1-(m-1)$. We have

$$\begin{aligned} V_\delta(M_1-m+1) = (1-\delta)u+\delta\{\tau^S[V_\delta(M_1-m+2) - V_\delta(M_1-m+1)]+ \\ + V_\delta(M_1-m+1) + \tau^S(V_\delta(M_1-m) - V_\delta(M_1-m+1))\}. \end{aligned}$$

Subtracting $V_\delta(M_1-m+2)$ on both sides and rearranging,

$$\begin{aligned} V_\delta(M_1-m+1) - V_\delta(M_1-m+2) = \left(\frac{1-\delta}{1-\delta+\delta\tau^S}\right)[u - V_\delta(M_1-m+2)]+ \\ + \left(\frac{\delta\tau^S}{1-\delta+\delta\tau^S}\right)[V_\delta(M_1-m) - V_\delta(M_1-m+1)]. \end{aligned}$$

The right hand side of the above equation is at most $u+w$ because $\tau^S > 0$ as well as $V_\delta(M_1-m+2) \geq -w$ and $V_\delta(M_1-m) - V_\delta(M_1-m+1) \geq u+w$. From the same argument in previous paragraphs, truth-telling is optimal if and only if

$$\begin{aligned} \frac{1-\delta}{\delta}(u+w) \geq (\varphi_+^L - \varphi_+^H)[V_\delta(M_1-m+1) - V_\delta(M_1-m+2)]+ \\ + (\varphi_-^H - \varphi_-^L)[V_\delta(M_1-m) - V_\delta(M_1-m+1)]. \end{aligned}$$

We know $[V_\delta(M_1-m) - V_\delta(M_1-m+1)] \leq u+w$ and $[V_\delta(M_1-m+1) - V_\delta(M_1-m+2)] \leq u+w$. From the values of φ_-^L, φ_+^H and ψ, we know that $\varphi_+^L - \varphi_+^H = \varphi_-^H - \varphi_-^L = \psi\left(\frac{1+\rho^*}{2}\right) \leq \frac{1}{2}\left(\frac{1-\delta}{\delta}\right)$. Therefore, the above inequality is satisfied.

III - setting parameters so that the customers' equilibrium payoff is ε-close to the maximum payoff. Consider the following sequence $\{\kappa_t, \psi_t, \tau_t\}_{t\in\mathbb{N}}$ of parameters, where

$$\kappa_t = \frac{1}{2}\left(\frac{1}{t}\right)^2, \qquad \psi_t = \frac{1}{1+\rho^*}\left(\frac{1}{t}\right), \qquad \tau_t = \frac{1}{2}\left(\frac{1}{t}\right)^2\left[\frac{2\lambda-1}{1-2\lambda\left(\frac{1-\rho^*}{2}\right)^{M_1-1}}\right].$$

There exists some T high enough such that, for every $t \geq T$, we have $\kappa_t \leq \frac{1}{2}\left(\frac{1-\delta}{\delta}\right)$ and $\psi_t \leq \frac{1}{1+\rho^*}\left(\frac{1-\delta}{\delta}\right)$. From previous discussions, this means that truth-telling will be optimal for the expert whenever hired. We can also set t high enough so that $\tau_t^S = \psi_t\left[\frac{3-\rho^*}{4}\right] \in (0,1)$. This means that all probabilities are well defined and truth-telling is optimal, for t high enough. Moreover, observe that, for every $t \in \mathbb{N}$,

$$\frac{\tau_t}{\kappa_t} = \left[\frac{2\lambda - 1}{1 - 2\lambda\left(\frac{1-\rho^*}{2}\right)^{M_1 - 1}}\right], \qquad \frac{\tau_t}{\psi_t} = \frac{1+\rho^*}{2}\left(\frac{1}{t}\right)\left[\frac{2\lambda - 1}{1 - 2\lambda\left(\frac{1-\rho^*}{2}\right)^{M_1 - 1}}\right].$$

This means that $\frac{\tau_t}{\psi_t} \to \frac{\tau_t}{\tau_t^S} \to 0$ as $t \to 0$, but $\frac{\tau_t}{\kappa_t}$ is kept constant to have $\rho_1 = \bar{\rho}$ and $\rho_{M_1} = \underline{\rho}(M_1)$. With such choice of parameters, as $t \to \infty$, we have

$$\lim_{t\to\infty} v(\rho) = f_{M_1}\left[\rho_{M_1}\left(\frac{u-w}{2}\right) + (1-\rho_{M_1})u\right] = \left(\frac{\bar{\rho}-\rho}{\bar{\rho}-\underline{\rho}(M_1)}\right)\left[\underline{\rho}(M_1)\left(\frac{u-w}{2}\right) + (1-\underline{\rho}(M_1))u\right].$$

So consider any $\varepsilon > 0$. Because $v(\rho)$ converges as $t \to \infty$, we can find some number T such that, for every $t \geq T$,

$$v(\rho) \geq \left(\frac{\bar{\rho}-\rho}{\bar{\rho}-\underline{\rho}(M_1)}\right)\left[\underline{\rho}(M_1)\left(\frac{u-w}{2}\right) + (1-\underline{\rho}(M_1))u\right] - \varepsilon.$$

Because the first term on the right-hand side of the above equation converges to $u\left(\frac{\bar{\rho}-\rho}{\bar{\rho}}\right)$ as $M_1 \to \infty$, we can find some $\bar{M} \geq M_1$ high enough such that, for every $M \geq \bar{M}$,

$$v(\rho) \geq u\left(\frac{\bar{\rho}-\rho}{\bar{\rho}}\right) - \varepsilon.$$

□

3 PERSUADING CROWDS

3.1 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 prone 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 paper 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.

[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 3.5.

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 the 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.

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 3.1). 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 3.2). 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 3.4 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 et al. (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 disclosure of information 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 3.2, I present an illustrative example using the simple symmetric binary private signal case from Bikhchandani, Hirshleifer and Welch (1998). In section 3.3, 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, Hirshleifer and Welch (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) discuss as well optimal persuasion mechanisms in this 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 3.2 introduces an illustrative example to walk through the main results. Section 3.3 presents the benchmark model without an information designer. Section 3.4 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 3.5 considers private disclosure of information and section 3.6 concludes the paper. All proofs of lemmas and claims are in appendix 3.7 and all calculations for examples are in appendix 3.8.

3.2 Illustrative example

Let me introduce an 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* (that is, the 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.

[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 3.8.

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.

To simplify the exposition, I present a visualization of the public belief dynamics and advisor's expected payoffs over time, as shown in figure 4.1(a). I also present the probability of clients choosing each action ignoring their private signals for each period in figure 4.1(b) - a phenomenon which is known as *informational cascade*, assuming $\sigma = .8$. In figure $4.1(a)$, 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 $4.1(b)$, 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 the 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.

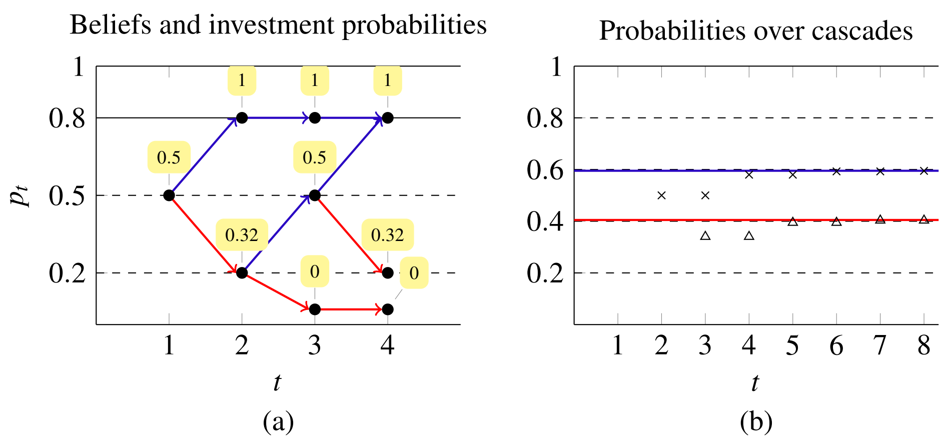

Figure 3.1 – Dynamics without intervention, for $\sigma = .8$. The blue lines represent outcomes related to investment decisions, and red lines represent outcomes related to non-investment decisions. Figure 4.1(a) shows the possible realizations of public beliefs over time with numbers representing the designer's expected payoff at every belief. Figure 4.1(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 this 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 4.2(a) and 4.2(b) represent the public belief dynamics and the probabilities over cascades under this information policy. The solid black dots in figure 4.2(a) represent the possible public beliefs (p_t). The white dots represent are the induced beliefs (ρ_t).

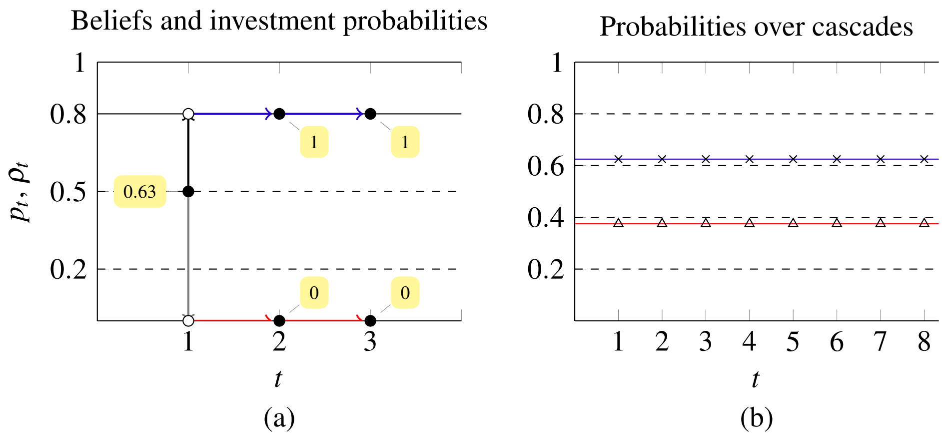

Figure 3.2 – Dynamics with single disclosure. Blue lines represent outcomes related to investment, and red lines represent outcomes related to non-investment. Figure 4.2(a) shows the possible realizations of public (black dots) and induced (white dots) beliefs over time, with numbers representing designer's expected payoff at every public belief. Figure 4.2(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). The 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 $4.3(a)$ and $4.3(b)$ 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.

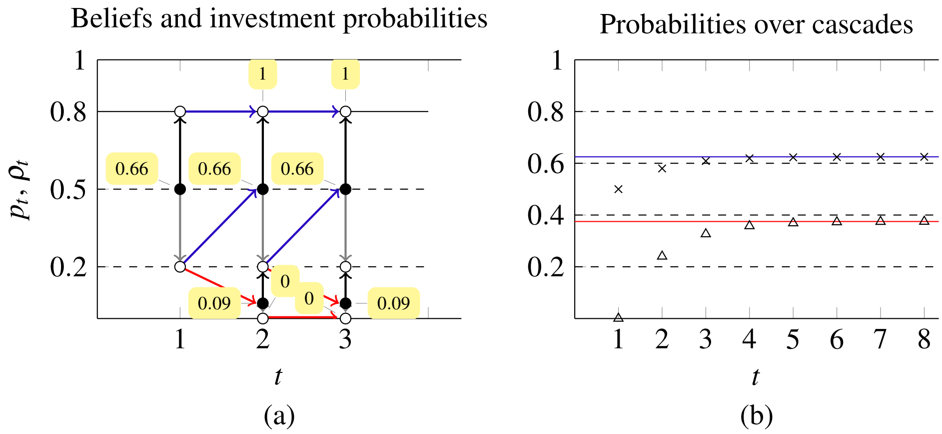

Figure 3.3 – Dynamics with the alternative policy. Blue lines represent outcomes related to investment, and red lines represent outcomes related to non-investment. Figure 4.3(a) shows the possible realizations of public (black dots) and induced (white dots) beliefs over time, with numbers representing the designer's expected payoff at every public belief. Figure 4.3(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 example is sufficiently manageable to characterize the related optimal policy. It is the alternative policy presented here, indeed. This policy minimizes the amount of information given at every public belief, subject to the expected investment probability being maximal (proposition 3.1). 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.

3.3 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, the 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 $G := (1/2)[G^H + G^L]$ 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 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.

[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 3.7 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 $\tilde{q}_t$.

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

$$r_t = \frac{(1-q_t)p_t}{(1-q_t)p_t + q_t(1-p_t)} \geq 1/2 \;\Leftrightarrow\; p_t \geq q_t. \tag{3.1}$$

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.

$$\alpha^H(p_t) := G^H(p_t) \text{ and } \alpha^L(p_t) := G^L(p_t) \tag{3.2}$$

$$\alpha(p_t) := p_t G^H(p_t) + (1-p_t)G^L(p_t). \tag{3.3}$$

One can show[8] that $\alpha^L(p) \leq \alpha(p) \leq \alpha^H(p)$ with strict inequalities for every belief $p \in (\underline{q},\bar{q})$. 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 4.2 and 4.3. Thus, agent 's $t+1$ inference from public history will lead to a public belief update:

$$p_{t+1} = \varphi_a(p_t) := \begin{cases} \left[\frac{\alpha^H(p_t)}{\alpha(p_t)}\right] p_t & \text{if } a_t = h, \\ \\ \left[\frac{1-\alpha^H(p_t)}{1-\alpha(p_t)}\right] p_t & \text{if } a_t = \ell. \end{cases} \tag{3.4}$$

[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 3.1 in appendix 3.7 or lemma 1 in Smith and Sørensen (1996).

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 4.4,

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

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}.

[9] See claim 3.2 in appendix 3.7 or lemma 7 in Smith and Sørensen (1996).
[10] See claim 3.3 in appendix 3.7 or theorem 1 in Smith and Sørensen (1996).

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, 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 the primitives of the model I consider as well as the relevant lessons from this 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.

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.

[11] See claim 3.4 in the appendix 3.7 or lemma 1 in Cao, Han and Hirshleifer (2011).

3.4 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 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 4.4 below describes the timing of the events.

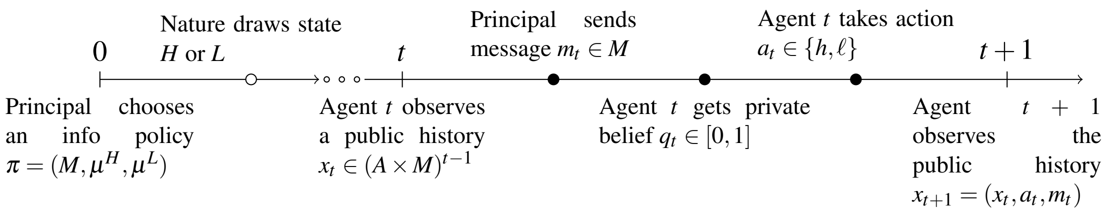

Figure 3.4 – Timing of events.

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 3.2 and 3.3, 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 3.4, 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 4.1 below[12] shows that the converse also holds.

Lemma 3.1. *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 value $\rho_t \in \operatorname{supp}(\tau)$, where τ is a probability measure over induced beliefs whose expected value equals p_t. Let $\mathscr{S}(p_t)$ be the set of all such probability measures. The public belief in the next period will be $\varphi_\ell(\rho_t)$ 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

$$V_\delta^{op}(p_t) := \sup_{\tau\in\mathscr{S}(p_t)} \mathbb{E}_\tau\left[(1-\delta)\alpha(\tilde{\rho}_t) + \delta\left((1-\alpha(\tilde{\rho}_t))V_\delta^{op}(\varphi_\ell(\tilde{\rho}_t)) + \alpha(\rho_t)V_\delta^{op}(\varphi_h(\tilde{\rho}_t))\right)\right]. \tag{3.5}$$

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 $\tau \in \mathscr{S}(p_t)$ 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].

[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 3.7.

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. But 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.

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.

[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.

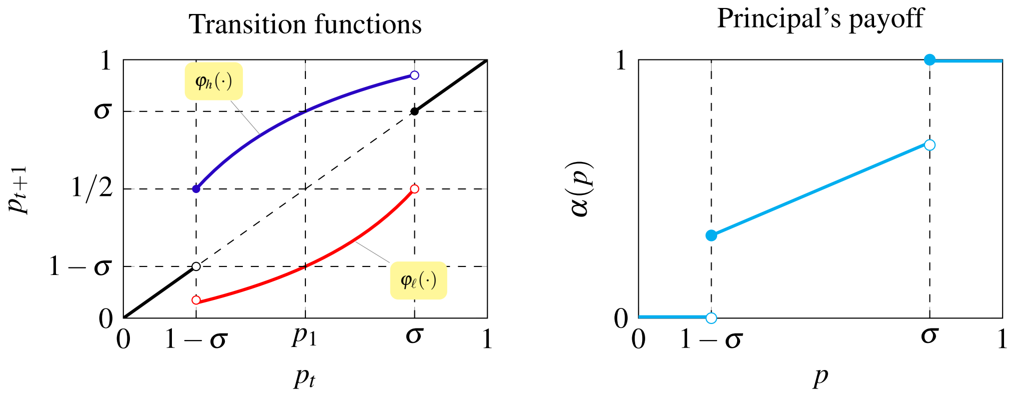

(a) (b)

Figure 3.5 – Analysis of the transition functions and the principal's expected payoff of the illustrative example. Figure 4.5(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 4.5(b) represents the expected investment probability for each public belief. The figures assume $\sigma = .8$.

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 $\mathrm{cav}[\alpha]$. Direct computation shows that $\mathrm{cav}[\alpha]$ is defined as below. The figures 4.6 (a) and 4.6 (b) present the concave closures $\mathrm{cav}[\alpha]$ of α for $\sigma = .6$ and $\sigma = .8$, respectively.

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

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

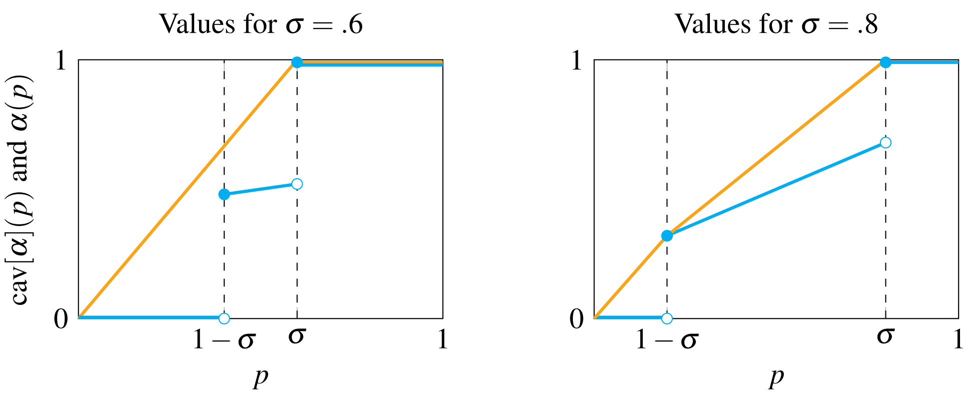

(a) (b)

Figure 3.6 – Values of selected functions for different values of σ 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 3.5 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 $1/\sqrt{2} < \sigma < 1$ instead, the single disclosure strategy does not maximize the advisor's current payoff. Applying the algorithm in equation 3.5, 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 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 given below. I represent such compositions in figure 4.7 (a), for $\sigma = .8$. In figure 4.7(b), I plot 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_δ^{sd}. 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 $\text{cav}[Z_\delta^{sd}](p) = V^{sd}(p)$ for every p. The concavification of Z_δ^{sd} is the dashed line in figure 4.7(c). From this figure, it can be seen that $\text{cav}[Z_\delta^{sd}]$ and V^{sd} differ outside C_h.

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

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

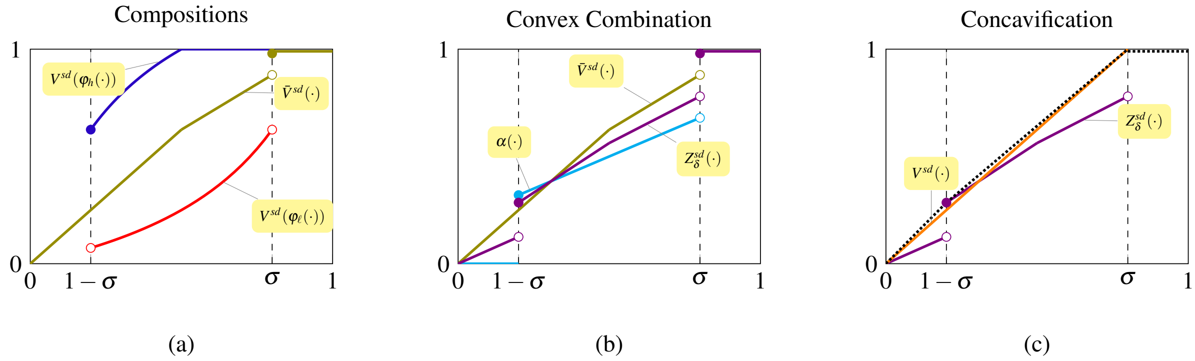

Figure 3.7 – Testing the optimality of a single disclosure policy for the illustrative example. The values of the parameters are $\sigma = .8$ and $\delta = .5$. Figure 4.7 (a) on the left shows the compositions of V^{sd} with the laws of motion φ_ℓ (red line) and φ_h (blue line). It also represents 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 4.7 (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 4.7 (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 3.8, 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 $\tau \in \mathscr{S}(p_t)$ 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 3.1. *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 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.

3.4.1 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 $\tau \in \mathscr{S}(p)$. 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 4.2 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 3.2. *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 $\mathscr{S}(p)$. 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 $\tau^{op}(p) \in \mathscr{S}(p)$ 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 3.1. *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 $\tau \in \mathscr{S}(p_t)$. From equation 3.4,

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

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 3.3. *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 $\tilde{\rho}_\infty$. Taking this into consideration, 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 $\tau \in \mathscr{S}(p_\infty)$ 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 3.3 from appendix 3.7 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 the above proposition, 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 4.2 summarizes these observations.

Corollary 3.2. *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.

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

3.4.2 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.

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:

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

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 3.5, 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 $V^{sd} \leq V_\delta^{op} \leq \mathrm{cav}[\alpha]$, whenever $V^{sd} = \mathrm{cav}[\alpha]$ 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 3.4. *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

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

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 $V^{op}_{\delta}(p) \geq Z^{op}_{\delta}(p)$ - the value of not disclosing anything at p and resorting to the optimal policy next period, it follows that

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

The second inequality follows from $\alpha(p) > V^{sd}(p)$ and $V^{op}_{\delta} \geq V^{sd}$. 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 $V^{op}_{\delta}(p) > V^{sd}(p)$. 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 implies that $V^{op}_{\delta}(p) > V^{sd}(p)$; 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 3.1. *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:

$$-\alpha''(p) = 4p(1-p)g(p)\left[\frac{3}{2}\left(\frac{2p-1}{p(1-p)}\right) - \frac{g'(p)}{g(p)}\right]. \tag{3.6}$$

[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.

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 4.2 demonstrates.

Lemma 3.2. *If the private belief density g is log-concave, then α is convex-concave on $(\underline{q},\bar{q})$.*

Profiting from lemma 4.2, theorem 4.1 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. More 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.

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 $\text{cav}[\alpha]$ 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 3.1. *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 $\delta \in (0,1)$,*

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

[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 $r'' > r'$.

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,

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

□

Let me stress one final remark regarding the log-concavity assumption. 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 4.3 evidences.

Corollary 3.3. *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 is valuable 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 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 3.8. Here, I provide a visual representation of these functions as well as the single disclosure policy in figure 4.8, 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 $p_1 = 1/2$, $V_\delta^{op} = 1/(2\sigma)$. 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 $V_\delta^{op} > V^{sd}$, that is, some social learning is valuable to the principal.

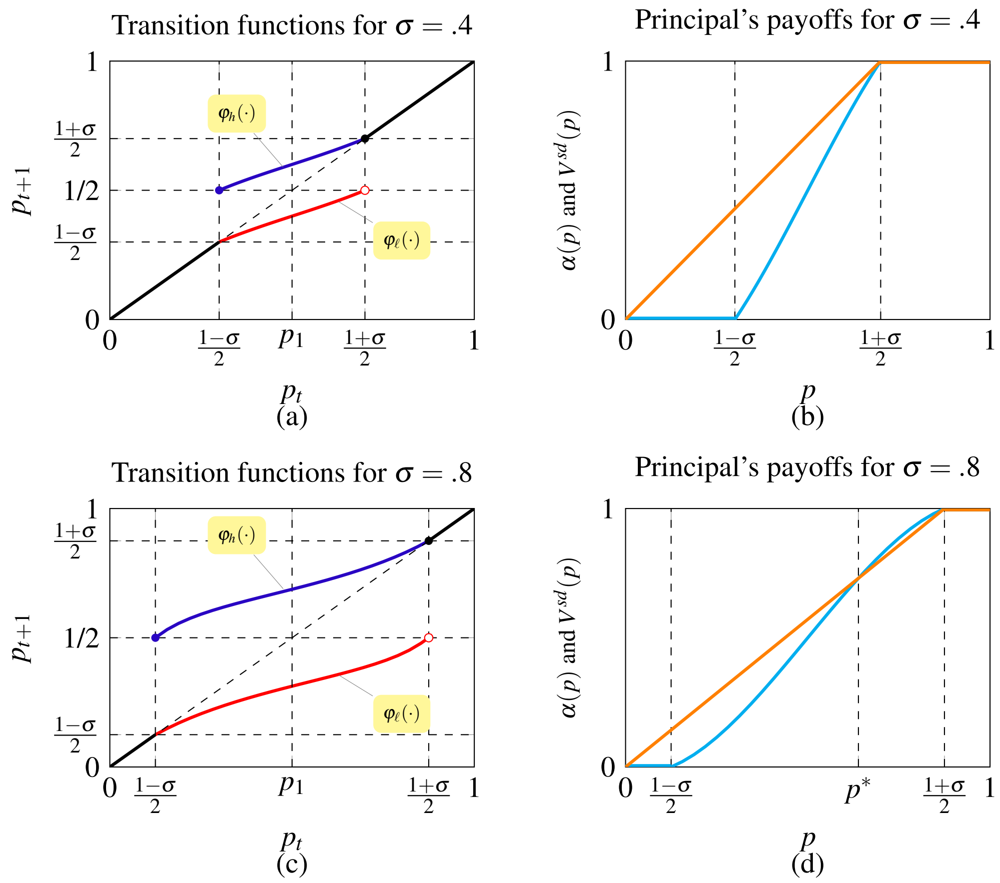

Figure 3.8 – Expected probability vs. single disclosure value function for different values of σ. Figures 4.8 (a) and 4.8 (b) represent relevant functions with $\sigma = .4$, and figures 4.8 (c) and 4.8 (d) show the same functions with $\sigma = .8$. The blue lines in figures 4.8 (b) and 4.8 (d) show the investment probabilities, and the orange lines show the value of a single disclosure policy.

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

3.4.3 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 3.3), 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 4.3 then follows.

Lemma 3.3. *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 $\lim_{\delta\to 1} V_\delta^{op} \leq V^{sd}$.

Reverse inequality must also be true. Indeed, by definition, $V_\delta^{op} \geq V^{sd}$ 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 3.2. *The value of an optimal policy approaches the value of the single disclosure one, as principal becomes increasingly patient:*

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

3.5 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.

$$\hat{\alpha}^\theta(p_t, \tau_t) := \mathbb{E}_{\tau_t^H(p_t)}[\alpha(\rho_t)] \text{ for } \theta \in \{H, L\}; \tag{3.7}$$

$$\hat{\alpha}(p_t, \tau_t) := p_t \hat{\alpha}^H(p_t, \tau_t) + (1 - p_t)\hat{\alpha}^L(p_t, \tau_t). \tag{3.8}$$

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

$$\tilde{p}_{t+1} = \hat{\varphi}_a(p_t, \tau_t) = \begin{cases} \left[\frac{\hat{\alpha}^H(p_t,\tau_t)}{\hat{\alpha}(p_t,\tau_t)}\right] p_t & \text{if } a = h, \\ \\ \left[\frac{1-\hat{\alpha}^H(p_t,\tau_t)}{1-\hat{\alpha}(p_t,\tau_t)}\right] p_t & \text{if } a = \ell. \end{cases} \tag{3.9}$$

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 for every $p \in [0,1]$

$$V_{\delta}^{op}(p) = \max_{\tau \in \mathscr{S}(p)} \left[(1-\delta)\hat{\alpha}(p,\tau) + \delta \left(\hat{\alpha}(p,\tau) V_{\delta}^{op}(\hat{\varphi}_h(p,\tau)) + (1-\hat{\alpha}(p,\tau)) V_{\delta}^{op}(\hat{\varphi}_{\ell}(p,\tau)) \right) \right]. \quad (3.10)$$

The equation above 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 4.4 below proves that concavity is preserved in a private persuasion mechanism.

Lemma 3.4. *With private communication, the function V_{δ}^{op} is concave in* $(0,1)$.

Lemma 4.4 implies that one of the essential 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.

3.6 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 paper 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 3.1). 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 3.2). 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.

3.7 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 3.1, 3.2 and 3.3 are taken from Smith and Sørensen (1996) and Rosenberg and Vieille (2019). Claim 3.4 is taken from Cao, Han and Hirshleifer (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 $\mathscr{F}$, 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 $(\Omega, \mathscr{F})$ given θ.

Every agent observes the realization of a measurable function $\tilde{s}_t : (\Omega, \mathscr{F}) \to (S, \mathscr{S})$. 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 $\mathscr{S}$ 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 $\tilde{q}_t : (S, \mathscr{S}) \to ((0,1], \mathscr{B})$ such that

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

where $\mathscr{B}$ 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 3.1. *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

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

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

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 3.2. *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 3.1. 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

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

$\square$

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 $\mathscr{F}$. Let $\mathscr{A}_t$ 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 $\mathscr{A}_t$ and $(\mathscr{A}_t)_{t\in\mathbb{N}}$ is an increasing family of sub-σ-algebras of $\mathscr{F}$. 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 $\mathscr{A}_t$ and $\mathscr{A}_t \subset \mathscr{A}_{t+1}$, it follows that $\mathbb{E}[\tilde{p}_{t+1}|\mathscr{A}_t] = \mathbb{E}[\mathbb{E}[\mathbb{1}_{\theta=H}|\mathscr{A}_{t+1}]\mathscr{A}_t] = \mathbb{E}[\mathbb{1}_{\theta=H}|\mathscr{A}_t] = \tilde{p}_t$ 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 3.3. *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 $p^*_\infty \in I \cap \mathrm{supp}(\tilde{p}_\infty)$ 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, $p^*_\infty \in \mathrm{supp}(\tilde{p}_\infty)$, 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. $\square$

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 $\mathrm{supp}(\tilde{p}_\infty) \subseteq C_\ell \cup C_h$, it must also be the case that $\mathrm{supp}(\lambda^{np}_\infty) \subseteq C_\ell \cup C_h$. 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 3.4. $\lim_{\delta\to 1} V^{np}_\delta = \lim_{t\to\infty} \mathbb{E}_{\lambda^{np}_t}[\alpha]$.

Demonstração. 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

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

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 $1 > \delta' > \bar{\delta}$,

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

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

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 $\mathscr{F}$, the σ-algebra generated by Ω.

Lemma 3.1. *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.*

Demonstração. 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 $\mathscr{B}$ of M. Consider the $\mathscr{M}$-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 $\mathscr{M}$,

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

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 $\mathscr{M}$ 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 $B \in \mathscr{M}$ 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 $\rho_t \in \operatorname{supp}(\tau(p_t))$. 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:

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

$\square$

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

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

where $\operatorname{co}(\operatorname{hyp}(f))$ is the convex hull of the hypograph of f. The concave closure of a bounded function is concave. Indeed, let $\operatorname{hyp}(\operatorname{cav}[f])$ be the hypograph of $\operatorname{cav}[f]$. 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 $\mathbb{E}_{\tau''}[\tilde{x}] = x''$, $\mathbb{E}_{\tau''}[f(\tilde{x})] = \lambda t + (1-\lambda)t' := t'' \leq \operatorname{cav}[f](x'')$. That implies $\operatorname{hyp}(\operatorname{cav}[f])$ is convex and $\operatorname{cav}[f]$ 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 $P: [0,1] \times \mathscr{B} \to [0,1]$ such that for every $p \in [0,1]$ and every $B \in \mathscr{B}$,

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

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.

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

Claim 3.5. *For every p and every bounded, continuous function V, there exists a solution $\tau \in \mathscr{S}(p)$ to the problem:*

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

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 and Lucas, 1989, theorem 8.3.

Claim 3.6. *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. □

Claim 3.7. *The transformation T is a contraction.*

Proof. First note that $(TV)(p)$ is equivalent to $\mathrm{cav}[(1-\delta)\alpha+\delta(PV)](p)$. 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 $\mathrm{cav}[f+d](p) = \mathrm{cav}[f]+d$. Therefore, $(Tf+d)(p) = (Tf)(p)+\delta d$. □

Claim 3.8. *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. □

Claim 3.9. *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. □

Valuable social learning

Lemma 3.2. *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.

[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 $d \geq 0$.

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 $p > m$. □

The role of patience

Lemma 3.3. *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,*

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

Demonstração. Because informative communication eventually stops, $\hat{\lambda}_t$ converges to λ_t as t goes to infinity. Claim 3.3 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 3.4. *With private communication, the function V_δ^{op} is concave in $(0,1)$.*

Demonstração. Because we have a contraction algorithm, it suffices to show that equation 3.10 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 3.10 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 3.9 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

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

$\square$

3.8 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

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

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

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

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

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

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

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

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

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

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

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

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 $\lim_{\delta\to 1}V_\delta^{np} = \lim_{t\to\infty}\mathbb{E}_{\lambda_t^{np}}[\alpha] = \lambda_\infty^{np}(C_h) = (1/2)/[(1-\sigma)^2+\sigma]$. Let me now compute the value of greedy policy V_δ^{gp}. Suppose first that $1/2<\sigma\le 1/\sqrt{2}$. 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 3.1. *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 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.

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

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

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

With some algebra, it follows that

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

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

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

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

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

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,

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

□

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

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

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

$$\hat{\lambda}_2^{gp}(0) = \tau^{gp}(0; \varphi_1(1-\sigma))\lambda_2^{gp}(\varphi_\ell(1-\sigma)) = (1/2)\sigma[2\sigma-1],$$
$$\hat{\lambda}_2^{gp}(1-\sigma) = \tau^{gp}(1-\sigma; \varphi_\ell(1-\sigma))\lambda_2^{gp}(\varphi_\ell(1-\sigma)) + \tau^{gp}(1-\sigma; 1/2)\lambda_2^{gp}(1/2) = (1/2)(1-\sigma^2),$$
$$\hat{\lambda}_2^{gp}(\sigma) = \tau^{gp}(\sigma; \sigma)\lambda_2^{gp}(\sigma) + \tau^{gp}(\sigma; 1/2)\lambda_2^{gp}(1/2) = (1/2)[1 + \sigma(1-\sigma)].$$

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 3.1, 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

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

BIBLIOGRAPHY

AN, M. Y. Log-concave probability distributions: Theory and statistical testing. *Working paper*, 1997.

AN, M. Y. Logconcavity versus logconvexity: a complete characterization. *Journal of economic theory*, New York: Academic Press., v. 80, n. 2, p. 350–369, 1998.

AU, P. H. Dynamic information disclosure. *The RAND Journal of Economics*, Wiley Online Library, v. 46, n. 4, p. 791–823, 2015.

AUMANN, R. J.; MASCHLER, M.; STEARNS, R. E. *Repeated games with incomplete information*. Cambridge: MIT press, 1995.

BAGNOLI, M.; BERGSTROM, T. Log-concave probability and its applications. *Economic theory*, Springer, v. 26, n. 2, p. 445–469, 2005.

BANERJEE, A. V. A simple model of herd behavior. *The quarterly journal of economics*, MIT Press, v. 107, n. 3, p. 797–817, 1992.

BERGEMANN, D.; MORRIS, S. Information design, Bayesian persuasion, and Bayes correlated equilibrium. *American Economic Review*, v. 106, n. 5, p. 586–91, 2016.

BERGEMANN, D.; MORRIS, S. Information design: A unified perspective. *Journal of Economic Literature*, v. 57, n. 1, p. 44–95, March 2019.

BEST, J.; QUIGLEY, D. Persuasion for the long run. *Working paper*, 2017.

BHASKAR, V.; THOMAS, C. Community Enforcement of Trust with Bounded Memory. *The Review of Economic Studies*, v. 86, n. 3, p. 1010–1032, 10 2018.

BIKHCHANDANI, S.; HIRSHLEIFER, D.; WELCH, I. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of economic perspectives*, v. 12, n. 3, p. 151–170, 1998.

Bright Local. *Local Consumer Review Survey 2018*. 2018. `https://www.brightlocal.com/research/local-consumer-review-survey-2018/`. Accessed: 2021-03-22.

Bright Local. *Local Consumer Review Survey 2019*. 2019. `https://www.brightlocal.com/research/local-consumer-review-survey-2019/`. Accessed: 2021-03-22.

BROCAS, I.; CARRILLO, J. D. Influence through ignorance. *RAND Journal of Economics*, v. 38, n. 4, p. 931–947, 2007.

CAO, H. H.; HAN, B.; HIRSHLEIFER, D. Taking the road less traveled by: Does conversation eradicate pernicious cascades? *Journal of Economic Theory*, Elsevier, v. 146, n. 4, p. 1418–1436, 2011.

CARROLL, G. Robustness and linear contracts. *American Economic Review*, v. 105, n. 2, p. 536–63, 2015.

CBS News. *Angie's List favors its own advertisers*. 2019. `https://www.cbsnews.com/news/angies-list-reviews-favor-its-own-advertisers-consumer-group-warns/`. Accessed: 2021-03-22.

CHE, Y.-K.; HORNER, J. Recommender Systems as Mechanisms for Social Learning*. *The Quarterly Journal of Economics*, v. 133, n. 2, p. 871–925, 12 2017.

COMPTE, O.; POSTLEWAITE, A. Plausible cooperation. *Games and Economic Behavior*, v. 91, p. 45–59, 2015.

CRIPPS, M. W.; MAILATH, G. J.; SAMUELSON, L. Imperfect monitoring and impermanent reputations. *Econometrica*, Wiley Online Library, v. 72, n. 2, p. 407–432, 2004.

CULE, M.; SAMWORTH, R. Theoretical properties of the log-concave maximum likelihood estimator of a multidimensional density. *Electronic Journal of Statistics*, The Institute of Mathematical Statistics and the Bernoulli Society, v. 4, p. 254–270, 2010.

DELLAROCAS, C. Designing reputation systems for the social web. *Working paper*, p. 1–10, 2010.

DELLAROCAS, C.; WOOD, C. A. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. *Management Science*, v. 54, p. 460–476, 2008.

DOVAL, L.; ELY, J. Sequential information design. *Econometrica*, n. forthcoming, 2020.

DUGHMI, S. Algorithmic information structure design: A survey. *SIGecom Exch.*, v. 15, n. 2, p. 2–24, 2017.

EKMEKCI, M. Sustainable reputations with rating systems. *Journal of Economic Theory*, Elsevier, v. 146, n. 2, p. 479–503, 2011.

ELY, J. C. Beeps. *The American Economic Review*, American Economic Association, v. 107, n. 1, p. 31–53, 2017.

ELY, J. C.; VÄLIMÄKI, J. Bad reputation. *The Quarterly Journal of Economics*, v. 118, n. 3, p. 785–814, 2003.

FILIPPAS, A.; HORTON, J. J.; GOLDEN, J. Reputation inflation. *Proceedings of the 2018 ACM Conference on Economics and Computation*, p. 483–484, 2018.

FUDENBERG, D.; LEVINE, D. K. Reputation and equilibrium selection in games with a patient player. *Econometrica*, JSTOR, v. 57, n. 4, p. 759–778, 1989.

FUDENBERG, D.; LEVINE, D. K. Maintaining a reputation when strategies are imperfectly observed. *The Review of Economic Studies*, Oxford University Press, v. 59, n. 3, p. 561–579, 1992.

GARRETT, D. F. Robustness of simple menus of contracts in cost-based procurement. *Games and Economic Behavior*, v. 87, p. 631–641, 2014.

GLAZER, J.; KREMER, I.; PERRY, M. Crowd learning without herding: A mechanism design approach. *Working Paper*, 2015.

GOLDMAN, E. The regulation of reputational information. In: SZOKA, B.; MARCUS, A. (Ed.). *The Next Digital Age: Essays on the Future of the Internet*. Washington, D.C.: TechFreedom, 2010. cap. 4, p. 293–304.

HALAC, M.; KARTIK, N.; LIU, Q. Contests for experimentation. *Journal of Political Economy*, v. 125, n. 5, p. 1523–1569, 2017.

HELLMAN, M. E.; COVER, T. M. Learning with finite memory. *The Annals of Mathematical Statistics*, p. 765–782, 1970.

HERRERA, H.; HÖRNER, J. A necesssary and sufficient condition for information cascades. *Working paper*, 2012.

HÖRNER, J.; LAMBERT, N. S. Motivational ratings. *Working paper*, 2016.

INOSTROZA, N.; PAVAN, A. Persuasion in global games with application to stress testing. *Working paper*, 2017.

KAMENICA, E. Bayesian persuasion and information design. *Annual Review of Economics*, v. 11, n. 1, 2019.

KAMENICA, E.; GENTZKOW, M. Bayesian persuasion. *The American Economic Review*, American Economic Association, v. 101, n. 6, p. 2590–2615, 2011.

KAMENICA, E.; GENTZKOW, M. Costly persuasion. *American Economic Review*, v. 104, n. 5, p. 457–462, 2014.

KOLOTILIN, A. et al. Persuasion of a privately informed receiver. *Econometrica*, Wiley Online Library, v. 85, n. 6, p. 1949–1964, 2017.

KOVBASYUK, S.; SPAGNOLO, G. Memory and markets. *Working paper*, 2018.

KREMER, I.; MANSOUR, Y.; PERRY, M. Implementing the wisdom of the crowd. *Journal of Political Economy*, v. 122, n. 5, p. 988–1012, 2014.

LAIHO, T.; MURTO, P.; SALMI, J. Gradual learning from incremental actions. *Working paper*, 2021.

LEVIN, A. T.; WILLIAMS, J. C. Robust monetary policy with competing reference models. *Journal of Monetary Economics*, v. 50, p. 945–975, 2003.

LI, F.; NORMAN, P. Sequential persuasion. *Theoretical Economics*, n. forthcoming, 2020.

LILLETHUN, E. Optimal information design for reputation building. *Working paper*, 2017.

LIPNOWSKI, E.; RAVID, D.; SHISHKIN, D. Persuasion via weak institutions. *Working paper*, 2018.

LIU, Q. Information acquisition and reputation dynamics. *The Review of Economic Studies*, Oxford University Press, v. 78, n. 4, p. 1400–1425, 2011.

LIU, Q.; SKRZYPACZ, A. Limited records and reputation bubbles. *Journal of Economic Theory*, Elsevier, v. 151, p. 2–29, 2014.

LORECCHIO, C.; MONTE, D. Dynamic information design under constrained communication rules. *Working Paper*, Sao Paulo School of Economics, 2021.

MAILATH, G. J.; SAMUELSON, L. *Repeated games and reputations: long-run relationships*. New York, NY: Oxford university press, 2006.

MATHEVET, L.; PEREGO, J.; TANEVA, I. On information design in games. *Journal of Political Economy*, v. 128, n. 4, p. 1370–1404, April 2020.

MATYSKOVA, L. Bayesian persuasion with costly information acquisition. *Working Paper*, 2018.

MONTE, D. Bounded memory and permanent reputations. *Journal of Mathematical Economics*, v. 49, p. 345–354, 2013.

MONTE, D. Learning with bounded memory in games. *Games and Economic Behavior*, Elsevier, v. 87, p. 204–223, 2014.

NIKIFOROV, D. On the belief manipulation and observational learning. *Working paper*, 2015.

NOSKO, C.; TADELIS, S. The limits of reputation in platform markets: An empirical analysis and field experiment. *Working Paper*, 2015.

QUAH, J. K.-H.; STRULOVICI, B. Aggregating the single crossing property. *Econometrica*, Wiley Online Library, v. 80, n. 5, p. 2333–2348, 2012.

RAYO, L.; SEGAL, I. Optimal information disclosure. *Journal of Political Economy*, v. 118, n. 5, p. 949–987, 2010.

RENAULT, J.; SOLAN, E.; VIEILLE, N. Optimal dynamic information provision. *Games and Economic Behavior*, Elsevier, v. 104, p. 329–349, 2017.

ROCKAFELLAR, R. T. *Convex analysis*. Princeton, New Jersey: Princeton university press, 1970.

ROESLER, A.-K. Mechanism design with endogenous information. *Working paper*, 2014.

ROGERSON, W. P. Simple menus of contracts in cost-based procurement and regulation. *The American Economic Review*, v. 93, n. 3, p. 919–926, 2003.

ROSENBERG, D.; VIEILLE, N. On the efficiency of social learning. *Econometrica*, Wiley Online Library, v. 87, n. 6, p. 2141–2168, 2019.

SGROI, D. Optimizing information in the herd: Guinea pigs, profits, and welfare. *Games and Economic Behavior*, Elsevier, v. 39, n. 1, p. 137–166, 2002.

SMITH, L.; SØRENSEN, P. Pathological outcomes of observational learning. *Working paper*, Massachusetts Institute of Technology, 1996.

SMITH, L.; SØRENSEN, P. Pathological outcomes of observational learning. *Econometrica*, Wiley Online Library, v. 68, n. 2, p. 371–398, 2000.

SMITH, L.; SØRENSEN, P.; TIAN, J. Informational herding, optimal experimentation, and contrarianism. *the Review of Economic Studies*, 2021.

SMOLIN, A. Dynamic evaluation design. *Working paper*, 2017.

SPERISEN, B. Bad reputation under bounded and fading memory. *Economic Inquiry*, v. 56, n. 1, p. 138–157, 2018.

STOKEY, N.; LUCAS, R. *Recursive methods in economic dynamics*. Cambridge, Massachusetts: Harvard University Press, 1989.

TANEVA, I. Information design. *American Economic Journal: Microeconomics*, v. 11, n. 4, p. 151–185, November 2019.

Tech Dirt. *France Decides That Expressing An Opinion About Your Teachers Should Be Illegal*. 2008. https://www.techdirt.com/articles/20080304/005526425.shtml. Accessed: 2021-03-22.

VONG, A. Firm certification. *Working Paper*, 2021.

WEIL, D.; GRAHAM, M.; FUNG, A. Targeting transparency. *Science*, v. 340, n. 6139, p. 1410–1411, 2013.

WILLIAMS, D. *Probability with martingales*. Cambridge, United Kingdom: Cambridge university press, 1991.

WILSON, A. Bounded memory and biases in information processing. *Econometrica*, v. 82, n. 6, p. 2257–2294, 2014.

ZERVAS, G.; PROSERPIO, D.; BYERS, J. W. A first look at online reputation on airbnb, where every stay is above average. *Marketing Letters*, Springer, v. 32, n. 1, p. 1–16, 2021.