Information Provision with Intertemporal Externalities*

Sanjari Kalantri†

4 November 2026

Click here for latest version

Abstract

I analyse a dynamic model of persuasion in which a long-lived Sender provides information about an evolving state to a sequence of short-lived Receivers who choose between two actions, high and low. In each period, the Sender has a weak ordinal preference for the high action over the low and the good state over the bad, whereas each Receiver wants the high action only in the good state. Importantly, each Receiver's action affects the state the next one faces, creating intertemporal externalities: the high action may either improve or worsen the future state. When the high action improves the future state, the Sender's ordinal preferences alone ensure that the optimal policy is the one that induces the high action with maximal probability each period. When the high action worsens the future state, the optimal policy depends on the cardinal features of the Sender's preferences: it may maximise the probability of the high action or of the low action, fully reveal the state (eliminating any uncertainty), or leave the Receiver uncertain of the state after every disclosure. Beyond *what* the Sender discloses, the optimal policy also determines *when* she discloses, and whether disclosure eventually stops.

*I am grateful to Margaret Meyer, Inés Moreno de Barreda, Daniel Quigley and Joel Watson for long discussions and generous guidance. I also thank Elizabeth Baldwin, César Barilla, Alex Bloedel and audiences at the Oxford Student Theory Workshop, HEC Paris PhD Economics Conference, ISI (Delhi) Conference on Economic Growth and Development, Nuffield Postdoctoral Fellows Seminar, Berlin Economics PhD Conference and speakers at the Nuffield Economic Theory Seminar for feedback.

†Nuffield College and Department of Economics, University of Oxford; sanjari.kalantri@economics.ox.ac.uk

1 Introduction

In many economic environments, today's actions affect not only today's outcomes but also the opportunities and payoffs available in the future. The quality of natural resources depends on current conservation decisions;[1] a lender's decision to finance a firm affects whether the firm survives and the terms on which later lenders can lend;[2] parents' investment in education affects their children's future earnings;[3] and current interpersonal contact affects the future spread of contagious diseases.[4] A common feature of these environments is that current actions generate *intertemporal externalities*: they impose benefits or costs on future decision-makers by changing the state of the world they will inherit.

Consider fisheries. Fish stocks are a common-property resource and are therefore vulnerable to over-exploitation, generating the classic "tragedy of the commons" (Gordon, 1954; Hardin, 1968). Because current extraction affects the future stock of fish, the current use of the resource creates an intertemporal trade-off: greater extraction raises current output but can reduce future availability. Responses to this problem include assigning property rights (Scott, 1955) or governing the resource through community institutions (Ostrom, 1990). In international waters, however, such solutions are hard to implement, and enforcement is comparatively limited. Where fishing decisions cannot be controlled directly, information about the stock becomes a key lever: regional fisheries advisory bodies provide the scientific assessments that inform these decisions. The International Council for the Exploration of the Sea (ICES), for example, collects data on fish stocks in the North Atlantic and adjacent seas and publishes its assessments. Fishing fleets observe this information and choose their level of fishing effort—how many ships to send out and how many nets to cast—to maximise their profit: high effort pays when the stock is abundant, but not when it is depleted. Thus, an expert's information disclosure can affect not only the decision taken today, but also the state of the resource that future decision-makers will face.

[1] This is the motivation underlying the notion of sustainable development; see World Commission on Environment and Development (1987).

[2] See, for example, Bizer and DeMarzo (1992), Dewatripont and Maskin (1995), and He and Xiong (2012).

[3] See Becker and Tomes (1979, 1986).

[4] See Kermack and McKendrick (1927); for economic analyses of policies that reduce interpersonal contact, see, for example, Geoffard and Philipson (1996) and Adda (2016).

This observation motivates the following question: *how should an expert provide information when the decision she induces today affects the state of the world tomorrow?* I study this question in a stylised infinite-horizon model of information provision. A long-lived Sender with commitment power interacts with a sequence of short-lived Receivers. At the beginning of each period, a payoff-relevant state is realised, which neither the Sender nor the Receiver observes. Instead, both share a common prior and observe the public history of signals and actions. In every period, the Sender chooses and commits to a Blackwell experiment. The resulting signal is publicly observed, and both players update their beliefs to a posterior, after which the Receiver takes a myopically optimal action. The current state and the Receiver's action then determine the transition probabilities governing the state in the next period. The Sender maximises her discounted payoff over the infinite horizon.

I focus on a binary state, which can be interpreted as a good or bad level of the resource, and a binary action, which can be interpreted as a high or low level of effort in exploiting it. The Sender has a weak ordinal preference for the good state over the bad and the high action over the low. This structure captures the trade-off faced by a resource-management body that values current production and employment but also cares about the long-run viability of the resource.[5] The Receiver, by contrast, wants to match his action to the current state: he takes the high action when the state is good and the low action when the state is bad.

The resulting problem has a simple interpretation. Within a period, the Sender would like to use information to maximise the probability of the high action, subject to a credibility constraint. Across periods, however, she would like to steer the state towards the good state. When there are *positive intertemporal externalities*, i.e. the high action improves the distribution of future states, these two objectives are aligned.[6] When there are *negative intertemporal externalities*, i.e. the high action worsens the distribution of future states, they pull in opposite directions.

[5] ICES, for example, states that it is "committed to better understanding marine ecosystems and securing the benefits that people derive from them." The body "generate[s] and share[s] the data, knowledge, and advice needed to meet current and emerging conservation, management, and sustainability goals" (International Council for the Exploration of the Sea, n.d.).

[6] For example, when diseases spread more easily in a dense wildlife population, hunting makes a healthy population more likely in the future, whether or not it is currently diseased. Since hunting pays when the population is healthy but not when it is diseased, hunters want the high action only in the good state, as in the model.

I show that this distinction sharply determines the optimal policy. The *greedy policy* applies the optimal static information structure in every period, thereby maximising the probability of the high action given the prevailing prior in each period. Under positive intertemporal externalities, the Sender's weak ordinal preferences for the high action and the good state are sufficient for the greedy policy to be optimal (Theorem 1). The reason is that positive externalities align the Sender's static and dynamic objectives, so there is no reason to sacrifice current persuasion for future state management.

Under negative intertemporal externalities, the Sender faces a trade-off between the current payoff from inducing the high action and the future benefit of improving the state distribution. Since the high action worsens the distribution of future states, inducing it more often comes at a dynamic cost. Therefore, the optimal policy depends on how the Sender's cardinal preferences weigh these competing objectives.

Theorem 2 characterises the optimal policy. I illustrate its logic using four extreme preferences. Suppose first that the Sender has *cardinal state-independent preferences*, that is, the high and low actions give her a constant payoff regardless of the state, so she cares only about the high action. In this case, the greedy policy is optimal, as in the static problem: a better future state matters to her only as a means to future high actions, and these are worth less than a high action today because they are discounted and the state keeps changing.

At the opposite extreme, suppose instead that the Sender has *cardinal action-independent preferences*, that is, the good and bad states give her a constant payoff regardless of the action, so she cares only about the state. She then prefers the low action in both states, because it makes the good state more likely tomorrow. In the bad state the Receiver agrees, but in the good state he does not, so she induces the low action as often as possible: the *judicious policy*, which maximises the probability of the low action every period given the prevailing prior, is optimal.

Consider now the case where the Sender values only one outcome: the good state with the high action. Then, she has no reason to mislead the Receiver in either state: inducing the high action in the bad state gives her nothing and makes the good state less likely tomorrow, and inducing the low action in the good state costs her more today than it gains tomorrow. *Full disclosure* is then

optimal.

Finally, suppose the Sender is equally satisfied with every outcome except one: the bad state with the low action, which gives her nothing. She then has a reason to induce both the high action in the bad state and the low action in the good state: the first raises her payoff today, and the second costs her nothing today and makes the good state more likely tomorrow. The optimal policy can then involve interior pooling. Unlike full disclosure, which reveals the state to the Receiver, interior pooling leaves him uncertain after either signal: the experiment generates two posterior beliefs, both strictly between zero and one. One posterior induces the high action and the other the low action, so the Receiver may take the high action in the bad state or the low action in the good state. It is this uncertainty that allows the Sender to strike a different balance from either of the two action-oriented policies. Compared to the greedy policy, the Sender gives up some chance of inducing the high action to induce the low action more often and improve the future state distribution. Compared to the judicious policy, she retains a chance of inducing the high action and securing its current payoff benefit. Interior pooling thus allows the Sender to pursue both objectives without focusing exclusively on either action or fully revealing the state.

Between these extremes, the optimal policy depends on two comparisons. In the bad state, the Sender weighs her current gain from the high action against the cost of worsening the state tomorrow. In the good state, she weighs her current loss from the low action against the benefit of improving the state. If the future cost of the high action in the bad state outweighs her gain today, and the future benefit of the low action in the good state falls short of her loss today, the externality removes every reason to mislead the Receiver, and full disclosure is optimal, unlike in the static problem. If only the high action in the bad state pays off, the greedy policy is optimal; if only the low action in the good state does, the judicious policy. If both do, the greedy or judicious policy remains optimal when one of these gains clearly outweighs the other, and otherwise the optimal policy involves interior pooling.

The analysis also yields comparative-statics insights into how patience and state transitions shape information provision (Proposition 3). They reveal that greater patience does not simply make the Sender more transparent. A more patient Sender becomes more willing to reveal the bad

state, where inducing the high action harms future prospects, but more willing to withhold information in the good state, where inducing the low action improves them. Holding her preferences fixed, changes in the transition probabilities have similarly state-specific effects.

Finally, I show that the optimal policy determines not only *what* information the Sender reveals, but also *when* she reveals it. This is important in applications such as fisheries management, where the fishing effort induced by one stock assessment changes the resource inherited by future fleets, thereby shaping when further information is needed. Under every policy other than full disclosure, the Sender communicates only when the current belief lies within a certain range. At each disclosure, her experiment generates one of two posterior beliefs: one induces the high action and the other the low action. The posterior and the action taken determine the next period's prior, after which beliefs evolve predictably if the Sender remains silent. Thus, after the first disclosure, beliefs follow one of two paths, depending on which posterior the experiment induces at each disclosure, and these paths determine when the Sender next communicates (Lemma 2).

Whether disclosure eventually stops depends on whether an action can be sustained without further information. Consider the greedy policy and suppose a signal is realised that induces the high action. If taking that action leaves the next Receiver willing to choose the high action without new information, disclosure is needed in only finitely many periods (Proposition 4). Otherwise, the Sender must disclose again immediately after inducing the high action, and disclosure recurs indefinitely (Proposition 5). The judicious policy exhibits the reverse pattern. These results show that the content of a message determines when the next one is needed.

The rest of the paper is organised as follows. Next, I review the related literature. Section 2 introduces a two-state, two-action, infinite-horizon model of information provision with intertemporal externalities. Section 3 characterises the static benchmark, derives the Sender's Bellman equation and the properties of optimal policies, and solves two dynamic benchmarks. Sections 4 and 5 characterise the optimal policy under positive and negative externalities, respectively. Section 6 describes the belief dynamics and disclosure patterns that the optimal policies generate. Section 7 concludes. All proofs are in the appendices. Appendices F and G generalise some of the results to mixed externalities and to finitely many states and actions, respectively.

1.1 Related Literature

This paper belongs to the literature on dynamic information provision, which builds on the insights in Aumann and Maschler (1995) and Kamenica and Gentzkow (2011). The closest papers to mine are Renault, Solan, and Vieille (2017), Ely (2017), and Ashkenazi-Golan et al. (2023).[7] As here, a single long-lived Sender communicates with a Receiver who behaves myopically. In all three, however, the state follows an *exogenous* Markov chain, so the Sender cares about how the state evolves but cannot influence it, and she has cardinal state-independent preferences: in every period, she only wants to maximise the probability that a particular action is taken. Renault, Solan, and Vieille (2017) allow a finite state space and a binary action space, and identify a sufficient condition under which the greedy policy is optimal in the dynamic game. Ashkenazi-Golan et al. (2023) allow a binary state space and a finite action space, and characterise the optimal policy as periods become short: it is greedy at beliefs below the stationary belief, but generally not above it. In the basic model of Ely (2017), the state and action spaces are binary, with one of the states being absorbing. In all three papers, if the state and action spaces are binary, then the greedy policy is optimal in the dynamic game.[8]

Keeping both spaces binary, I relax the two assumptions these papers share in turn: cardinal state-independent preferences and exogenous transitions (Section 3.3). With cardinal state-independent preferences, the greedy policy remains optimal whatever the effect of the Receiver's action on the state. With exogenous transitions, a weak ordinal preference for the high action and the good state suffices for the greedy policy to be optimal, and the same holds under positive externalities (Theorem 1). The greedy policy can therefore fail only when the Receiver's action affects the state and the Sender's cardinal preferences depend on it. Under negative externalities, it can: depending on how the Sender values current actions against the future state, the optimal policy may instead maximise the probability of the low action, fully reveal the state, or pool at interior beliefs (Theorem 2).

[7] Lehrer and Shaiderman (2026) also study persuasion about a state that follows an exogenous Markov chain, but ask a different question: for general finite state and action spaces and state-dependent payoffs, they characterise when a patient Sender attains the static value at the stationary belief.

[8] If the state space or the action space has more than two elements, the greedy policy need not be optimal in the dynamic game, even with exogenous transitions and cardinal state-independent Sender preferences; see Renault, Solan, and Vieille (2017) and Ashkenazi-Golan et al. (2023).

Other papers study dynamic information provision to long-lived Receivers, where information also acts as an incentive over time. Ely and Szydlowski (2020) show how a principal who knows the difficulty of a task can use information over time to keep an agent working on it. Ball (2023) analyses a model in which both players are long-lived and the state evolves exogenously. Sugaya and Wolitzky (2026) analyse a model with a single Sender and multiple Receivers, all long-lived. In the baseline model, the states are i.i.d., but the authors consider exogenous Markov transitions in an extension. In all three papers, the Sender commits at the outset to a policy for the whole game. In Ball (2023) and Sugaya and Wolitzky (2026), for example, the optimal policy is implemented by threatening to cut off all future information after any deviation. Here, because each Receiver is short-lived, such threats have no force, and the Sender's commitment power over time is equivalent to commitment within a period.[9]

A few papers also study a Sender who faces several short-lived Receivers in turn, but the Receivers are linked through different channels. In Kremer, Mansour, and Perry (2014), Receivers choose between two actions with unknown rewards, and each choice reveals its reward to the Sender. Thus, Receivers' actions improve future payoffs through learning. In the sequential procedure in Khantadze, Kremer, and Skrzypacz (2025), each Receiver cares only about his own state, but the states are correlated, so information given to one Receiver constrains the persuasion of the others. Here, Receivers are linked through the state: each Receiver's action changes the state that the next one faces.

Finally, from a modelling perspective, this paper can also be viewed as a stochastic game (Shapley, 1953; Solan & Vieille, 2015) in which one player's action affects the transition probabilities as in Renault (2012). There, however, the controlling player is also informed of the state, whereas here the Receiver controls the transitions and holds no private information.

[9]Without commitment, Renault, Solan, and Vieille (2013) characterise the limit equilibrium payoffs of a sender–receiver game in which the state follows a Markov chain observed only by the sender.

2 Model

In this section, I introduce a two-state, two-action, infinite-horizon model of information provision with intertemporal externalities. Time, indexed by $t \in \mathbb{N}$, is discrete.

Agents and payoffs. A long-lived Sender (she) interacts with a sequence of short-lived Receivers (he). Their stage game payoffs depend on the state of the world $\omega_t \in \Omega = \{G(\text{ood}), B(\text{ad})\}$ and the Receiver's action $a_t \in \mathcal{A} = \{h(\text{igh}), \ell(\text{ow})\}$ in that period, and are given as follows:

$u_S(\omega_t, a_t)$	$\omega_t = G$	$\omega_t = B$
$a_t = h$	1	x
$a_t = \ell$	y	0

$u_R(\omega_t, a_t)$	$\omega_t = G$	$\omega_t = B$
$a_t = h$	1	-1
$a_t = \ell$	0	0

where $x \in [0, 1]$ and $y \in [0, 1]$.

Clearly, within a period, both players weakly prefer state G to state B; the Sender's *ordinal* preference is strict if $x < 1$ and $y > 0$. The Sender also weakly prefers action h to action ℓ (strict *ordinal* preference if $x > 0$ and $y < 1$). The Sender has *cardinal state-independent* stage game preferences if and only if $x = 1$ and $y = 0$. Likewise, the Sender has *cardinal action-independent* stage game preferences if and only if $x = 0$ and $y = 1$. Meanwhile, the Receiver wants to match action to state: he strictly prefers taking action h in state G and action ℓ in state B.

Since the Receiver is short-lived, he chooses an action that maximises his stage game payoff. On the other hand, the Sender maximises her long-run expected discounted payoff:

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \mathbb{E}[u_S(\omega_t, a_t)],$$

where $\delta \in [0, 1)$ is the Sender's discount factor.

In the fisheries example from the Introduction, the Sender is the advisory body and the Receivers are the fishing fleets of successive seasons. The state is the stock, abundant (G) or depleted (B), and the action is the fleets' fishing effort: high effort pays when the stock is abundant and makes a loss when it is depleted. The parameter x measures how much the body values high effort at a depleted stock, and y how much it values low effort at an abundant one. At $(x, y) = (1, 0)$ it only cares about fishing activity, and at $(0, 1)$ it only cares about the stock.

Intertemporal externalities. The key departure from previous models of dynamic information provision is that the distribution of next period's state depends jointly on the current state and the Receiver's current action. Thus, the state evolves *endogenously* over time. Formally, transition probabilities are given by $\mathbb{P}(\omega_{t+1} \mid \omega_t, a_t)$, such that $\sum_{\omega_{t+1}\in\Omega} \mathbb{P}(\omega_{t+1} \mid \omega_t, a_t) = 1$. Since the state and action spaces are binary, the transition probabilities can be described by the following matrix:

$\mathbb{P}(\omega_{t+1} = G \mid \omega_t, a_t)$	$\omega_t = G$	$\omega_t = B$
$a_t = h$	p_G^h	p_B^h
$a_t = \ell$	p_G^ℓ	p_B^ℓ

(2.1)

I use the term *intertemporal externality* to describe the effect of a current action on future states and hence on future payoffs. Since Receivers are short-lived, these effects are external to the Receiver choosing the current action. In particular, the Receiver's period-t action may impose a benefit or a cost on future periods by changing the likelihood that the state is good in period $t+1$:

Definition 1 (Intertemporal externalities)**.** The high action imposes a **positive intertemporal externality** if $p_\omega^h > p_\omega^\ell \ \forall \omega \in \Omega$, and a **negative intertemporal externality** if $p_\omega^h < p_\omega^\ell \ \forall \omega \in \Omega$. The state evolution process is **exogenous** if $p_\omega^h = p_\omega^\ell \ \forall \omega \in \Omega$.

Therefore, under positive (negative) intertemporal externalities, conditional on either state, choosing the high action makes the good state more (less) likely in the following period than choosing the low action. In what follows, for the sake of brevity, I omit the term "intertemporal".

In the fisheries example, the externality is negative: high effort this season makes an abundant stock less likely next season, whether the stock is currently abundant or depleted. Since the stock is common property, this cost falls on the fleets of later seasons rather than on those fishing today.

Timing. At the beginning of a period, conditional on the public history of past signals and actions, both players share a common prior $p_t \coloneqq \mathbb{P}(\omega_t = G) \in [0, 1]$. Every period, the Sender publicly chooses an *information structure* (that is, a Blackwell experiment; Blackwell, 1951), whose signal realisation is publicly observed. A *policy* specifies the experiment that the Sender chooses in each period as a function of the public history.

Formally, an information structure is a pair (Σ, σ_t), where Σ is a set of signals, and $\{\sigma_t(\cdot \mid \omega)\}_{\omega \in \Omega}$ is a set of probability distributions over signals, conditional on the true state ω. The Receiver observes the chosen information structure and both players observe the signal realisation $\varsigma_t \in \Sigma$, update their belief about the state using Bayes' rule to $q_t = \mathbb{P}(\omega_t = G \mid \varsigma_t)$, and the Receiver chooses an action that maximises his stage game payoffs.[10] Based on the action chosen by the Receiver, the state evolves for the next period (see Figure 1). In the fisheries example, a period is a fishing season: the advisory body publishes a stock assessment, whose design it fixes before the data come in; the fleets then choose their effort; and the stock evolves for the next season.

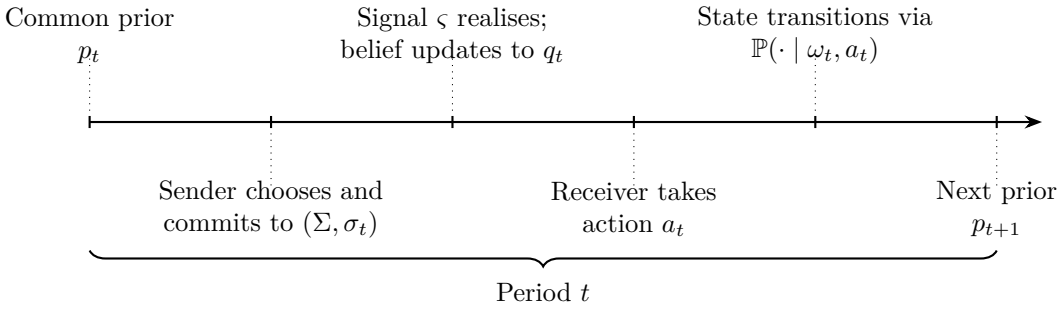


Figure 1: Timeline of events within a single period

Information. Both players know the stage game payoffs u_S, u_R and the transition probabilities $\{p^a_\omega\}$. Realised stage payoffs are not observed by either player (equivalently, they accrue only at the end of the game), so all learning about the state is through the Sender's signals. In every period t, both players observe the entire public history of past signals and past actions. Neither player ever directly observes the realised state ω_t, so the Sender has no private information in this setup.[11] She updates her belief in the same manner as the Receiver: after observing the signal realisation and after the Receiver takes an action. Thus, in each period t, conditional on the public history and

[10] I assume that the Receiver breaks ties in the Sender's favour taking into account her continuation payoff, so that the Sender's interim expected payoffs are upper semi-continuous. This is needed to ensure existence of the Sender's optimum.

[11] As mentioned in Renault, Solan, and Vieille (2017), an alternative interpretation would be that the Sender observes the sequence of states and chooses and commits to a policy *ex-ante*, i.e. prior to the beginning of the game. This is because the Receiver is short-lived, so the Sender's problem reduces to a single-agent stochastic dynamic optimisation problem in which she is constrained by the Receiver's best-response function. This has a recursive formulation: see Section 3.2.2 and equation (3.4) below. In this case, a policy for the Sender would map histories (past states and actions and current state) into distributions over signals.

the current signal realisation, both players share the same posterior belief q_t and the same prior for the next period p_{t+1}.

In the fisheries example, the stock cannot be observed directly: the advisory body learns about it only through its own assessments, which are public, so it knows no more than the fleets. The fleets' past effort is also public, but information from past catches reaches them only through the assessments.

3 Preliminaries

In this section, I first characterise the Sender's static optimum. I then derive the Sender's Bellman equation and the properties of optimal policies, and solve two dynamic benchmarks.

3.1 Static Benchmark: The Greedy Policy

In any period, the problem facing the Receiver is simple. Given his period-t posterior belief q_t, the Receiver gets an expected payoff of $q_t \cdot 1 + (1 - q_t) \cdot (-1)$ if he chooses action h, and a payoff of 0 if he chooses action ℓ. Thus, the Receiver's best response is to choose $a_t = h$ if $q_t \geq 1/2$ (at $q_t = 1/2$, he breaks his indifference in the Sender's favour, which in the static game means taking the high action) and $a_t = \ell$ otherwise. Given this best-response function, the Sender's interim expected payoffs are:

$$\mathbb{E}_{q_t}(u_S(q_t)) = \begin{cases} q_t \cdot 1 + (1 - q_t) \cdot x, & \text{if } q_t \geq \frac{1}{2} \\ q_t \cdot y, & \text{if } q_t < \frac{1}{2} \end{cases} \tag{3.1}$$

where $\mathbb{E}_{q_t}(u_S(q_t)) := \mathbb{E}_{q_t}(u_S(\omega_t, a(q_t)))$. In Figure 2, the Sender's interim expected payoff is plotted by the solid black line.

It is well-known from the literature on repeated games and Bayesian persuasion (Aumann & Maschler, 1995; Kamenica & Gentzkow, 2011) that within a period, the Sender's problem amounts to choosing a distribution over posteriors induced by the information structure to maximise her interim expected payoff. Let $\mu \in \Delta(\Delta(\Omega))$ denote such a distribution. Due to Bayesian updating, the feasible set of distributions over posteriors from which the Sender can choose comprises those

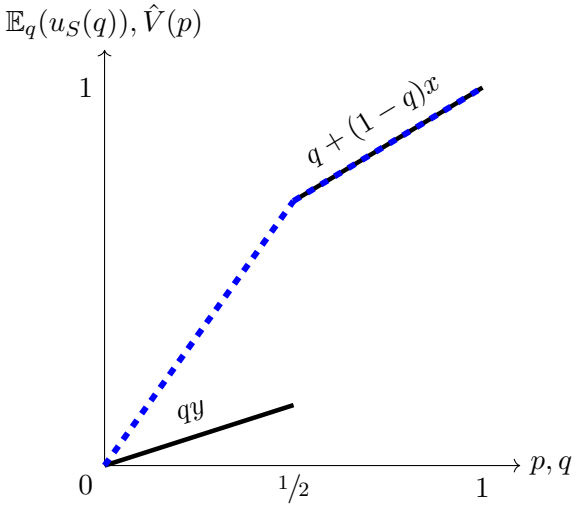


Figure 2: Sender's interim (black) and optimal (blue) expected payoff under the greedy policy

satisfying Bayes plausibility: $\mathbb{E}_\mu(q_t) = p_t$.[12] The Sender's optimal static value is the concave envelope[13] of her interim expected payoff (3.1), evaluated at p_t:

$$\hat{V}(p_t) = \min\{p_t(1+x), x + p_t(1-x)\},$$

which is plotted by the dashed blue line in Figure 2.

Thus, the optimal information structure for the Sender is to not disclose anything if $p_t \geq 1/2$; if $p_t < 1/2$, she induces a posterior of $1/2$ with probability $2p_t$ and 0 with complementary probability. Under this information structure, the Receiver never takes the low action in the good state: when he takes the low action, his posterior belief is 0, i.e. he is certain that the state is bad. The high action is taken with probability $\mu(p_t) = 1$ if $p_t \geq 1/2$ and with probability $\mu(1/2) = 2p_t$ if $p_t < 1/2$. Indeed, subject to Bayes plausibility, the Sender designs an information structure that maximally leads the Receiver to believe that the state is good, thereby maximising the probability that the high action is taken within a period. This is because myopically, she has an ordinal preference for action h over action ℓ. Consistent with the literature, I will refer to choosing the information

[12] A key result in Aumann and Maschler (1995) is that for a given distribution over posteriors such that they equal the prior in expectation, the Sender can always design an experiment that correlates the signals with the state in such a way that it induces the given distribution over posteriors.

[13] The concave envelope, or concavification, of a function f is the smallest concave function greater than or equal to f.

structure described above every period as the *greedy policy*.

3.2 The Dynamic Problem

3.2.1 Beliefs

From the previous subsection, we see that by choosing a distribution over posterior beliefs, the Sender indirectly chooses the probability with which an action is taken. This, in turn, influences the distribution of the next period's prior, due to the state transitions. This implies that when choosing a policy, the Sender chooses a (stochastic) sequence of players' beliefs. Hence, to solve the dynamic problem, we first need to track the effect of inducing a posterior in a period on the following period's prior.

Recall that p_t and q_t denote the common prior and posterior belief of the state being G in period t, respectively. After the Receiver chooses an action, the state evolves based on the transition probabilities described in (2.1). Specifically, the players' prior in period $t+1$ is given by:

$$p_{t+1}(q_t, a_t) = \begin{cases} q_t p_G^h + (1-q_t)p_B^h, & \text{if } a_t = h \\ q_t p_G^\ell + (1-q_t)p_B^\ell, & \text{if } a_t = \ell \end{cases} \tag{3.2}$$

Thus, the players' prior in a period is a conditional distribution over states given the entire public history, i.e. signals and actions taken in all previous periods.

Using the Receiver's best-response function, the above mapping from posteriors to priors can be re-written as:

$$p_{t+1}(q_t) = \begin{cases} f^h(q_t) \coloneqq p_B^h + q_t(p_G^h - p_B^h), & \text{if } q_t > 1/2 \\ f^\ell(q_t) \coloneqq p_B^\ell + q_t(p_G^\ell - p_B^\ell), & \text{if } q_t < 1/2 \end{cases} \tag{3.3}$$

and, if $q_t = 1/2$, the Receiver chooses action a_t under the Sender-favourable tie-breaking convention, and then $p_{t+1} = f^{a_t}(1/2)$.

3.2.2 Bellman Equation

We can now use the above to write the Bellman equation. Let $V : [0,1] \to [0,1]$ denote the Sender's maximum payoff from the dynamic optimisation problem. This function is characterised by the

solution to the following functional equation:

$$V(p_t) = \sup_{\mu \in \Delta(\Delta(\Omega))} \mathbb{E}_\mu\{(1-\delta)\mathbb{E}_{q_t}(u_S(q_t)) + \delta V(p_{t+1}(q_t))\} \quad \text{subject to } \mathbb{E}_\mu(q_t) = p_t \tag{3.4}$$

where $p_{t+1}(q_t)$ is as defined in equation (3.3). In Appendix A.2, I show (following Renault, Solan, & Vieille, 2017) that the value function $V(\cdot)$ is concave over its domain. Since the Sender's current payoff and $f^a(q)$ are affine in q for a fixed action a, concavity of $V(\cdot)$ implies that any collection of posteriors inducing the same action can be merged into a single posterior, their probability-weighted average, without lowering the Sender's payoff.[14] Hence, it is without loss to focus on policies that have at most one posterior inducing h and at most one inducing ℓ.

Let q^h and q^ℓ denote the posteriors that induce actions h and ℓ, respectively. Bayes plausibility implies that these need to be chosen such that $p = \alpha q^h + (1-\alpha)q^\ell$, where $\alpha := \mu(q^h) \in [0,1]$ is the probability with which the posterior q^h is induced. Then, the Bellman equation (3.4) can be re-written as (suppressing the time index for the current prior):[15]

$$V(p) = \max_{\alpha, q^h}\{(1-\delta)[\alpha(q^h + (1-q^h)x) + (1-\alpha)q^\ell y] + \delta\alpha V(f^h(q^h)) + \delta(1-\alpha)V(f^\ell(q^\ell))\} \quad \text{subject to } p = \alpha q^h + (1-\alpha)q^\ell, \alpha \in [0,1], q^h \in [1/2, 1], q^\ell \in [0, 1/2] \tag{3.5}$$

3.2.3 Properties of Optimal Policies

In this subsection, I characterise the properties of optimal policies, which greatly simplifies the problem.

The Bellman equation (3.5) shows that all histories leading to the same prior p give the Sender the same continuation payoff regardless of the time period in which this prior is reached. Since the maximum in equation (3.5) is attained at every prior,[16] it is without loss to restrict attention to

[14] That is, the Sender sends a single signal whenever the original information structure would send any of the signals leading to these posteriors. The Receiver still takes the same action at the merged posterior, since each action is optimal on an interval of beliefs. The Sender's expected current payoff is unchanged, and tomorrow's prior after the merged signal is the average of the original ones, so by concavity of V her continuation payoff cannot fall (see Corollary 3 in Appendix A.2).

[15] Given (α, q^h), q^ℓ is pinned down by Bayes plausibility and the constraint $q^\ell \in [0, 1/2]$ reflects the feasible (α, q^h) pairs. When $\alpha = 1$, q^ℓ is induced with probability 0 and its value is irrelevant.

[16] This follows from the continuity of V (Claim 1) and the tie-breaking convention of footnote 10, which together make the Sender's interim payoff upper semi-continuous; see the discussion after equation (A.5) in Appendix A.2.

stationary policies, which choose (α, q^h) as a function of the current prior only.

In this regard, consider any stationary policy of the Sender, which specifies the $(\alpha(p), q^h(p))$ pair that the Sender should choose at each prior p ($q^\ell(p)$ is determined by Bayes plausibility). Let $\phi(\cdot)$ denote the expected payoff to the Sender from following such a policy. The following lemma, which generalises Lemma 8 of Renault, Solan, and Vieille (2017), says that $(\alpha(p), q^h(p))$ is optimal if and only if $\phi(\cdot)$ is concave over $[0,1]$ and there is no profitable one-shot deviation to disclosing no information:

Lemma 1 (Properties of optimal policies)**.** A stationary policy $(\alpha(p), q^h(p))$ of the Sender is optimal if and only if $\phi(\cdot)$ is concave over $[0,1]$ and the following is satisfied:

$$\phi(p) \geq \begin{cases} (1-\delta)(p+(1-p)x) + \delta\phi(f^h(p)), & \text{if } p \geq 1/2 \\ (1-\delta)py + \delta\phi(f^\ell(p)), & \text{if } p \leq 1/2 \end{cases} \tag{3.6}$$

To summarise, we have shown that it is without loss to focus on stationary policies that induce at most one posterior for every action. Further, we have the above result, which is very useful: it can be used to verify the optimality of any candidate stationary policy.[17]

3.3 Two Dynamic Benchmarks

To facilitate a comparison to the existing literature, I analyse two special cases. Ely (2017), Renault, Solan, and Vieille (2017), and Ashkenazi-Golan et al. (2023) have all shown that in the binary-state, binary-action case, (i) if states evolve exogenously, and (ii) the Sender has cardinal state-independent preferences, then the greedy policy is optimal. I look at each of these in turn.

Proposition 1 (Exogenous transitions)**.** If states evolve exogenously, then the greedy policy is optimal for all values of $(x, y) \in [0,1]^2$.

In other words, if transitions are exogenous, then the greedy policy is optimal under weak ordinal preferences of the Sender. The reason for the above is as follows. Since actions do not affect future states, the expected prior tomorrow is the same whatever the Sender discloses today. Therefore, today's communication matters for the future only through the dispersion it creates in

[17] Lemma 1 generalises beyond the binary state, binary action case; see Lemma 20 in Appendix G.

tomorrow's prior, which, since the continuation value is concave, can only lower it. The greedy policy induces the high action as often as possible, and in the binary case this static gain is always at least as large as any loss in continuation value.

Proposition 2 (Cardinal state-independent preferences)**.** If $(x, y) = (1, 0)$, i.e. the Sender has cardinal state-independent preferences, then for all values of the transition probabilities, the greedy policy is optimal.

Intuitively, when the Sender only cares about the high action, a higher belief is just a means of getting that action. Holding the belief back for the future yields less than one-for-one, because it is discounted and diluted by state transitions. Thus, spending the belief now always wins.

Thus, Propositions 1 and 2 together generalise previous results: in the binary-state, binary-action case, if (i) states evolve exogenously, or (ii) the Sender has cardinal state-independent preferences, then the greedy policy is optimal.

4 Positive Externalities

In the dynamic game, the Sender has twin objectives: *within a period*, she wishes to maximise the probability of the high action, and *across periods*, she wants to maximise the probability of the good state. When there are positive externalities, her static and dynamic incentives align:

Theorem 1 (Positive externalities)**.** Under positive externalities, the greedy policy is optimal for all values of $(x, y) \in [0, 1]^2$.

The proof, in Appendices B and E, uses Lemma 1 and proceeds in two parts. First, I show that the payoff induced by the greedy policy is concave in p if and only if $x \geq \bar{x} \coloneqq \frac{\delta(p_B^\ell - p_B^h)}{1+\delta p_B^\ell - \delta p_G^h}$. Since $p_B^h > p_B^\ell$, $\bar{x} \leq 0$, and so the payoff from the greedy policy is concave for all $(x, y) \in [0, 1]^2$.

Second, I verify that the Sender does not gain from a one-shot deviation to no disclosure. The following observation provides a useful simplification. Since the Sender does not disclose any information when $p \geq 1/2$, this condition only needs to be verified when $p \in [0, 1/2)$. If the Sender one-shot deviates to no disclosure over this interval, the Receiver takes the low action, giving the

Sender an expected payoff of py from the current stage. This means that when $y = 1$, the Sender is maximally incentivised to deviate to no disclosure. Indeed, given x, the higher the value of y, the more the Sender prefers the good state. Thus, if the greedy policy is optimal when the Sender is most incentivised to deviate, then it is optimal for lower values of y as well. Moreover, the Sender's long-run payoff under the greedy policy, and hence the one-shot deviation condition, is affine in x. Extending to all real x, it therefore suffices to verify the condition at the extreme points $x = \bar{x}$ and $x = 1$ as $[0, 1] \subseteq [\bar{x}, 1]$. When $x = \bar{x}$, this condition is satisfied for all $y \leq \bar{y} \coloneqq \frac{1+\delta p_B^\ell - \delta p_G^\ell}{1+\delta p_B^\ell - \delta p_G^h}$; since $p_G^h > p_G^\ell$, $\bar{y} \geq 1$. The rest of the proof proceeds to show that when $(x, y) = (1, 1)$, under positive externalities, the Sender does not gain from a one-shot deviation to no disclosure.

Thus, the Sender's weak ordinal preference for the high action and the good state is sufficient for the greedy policy to be optimal under positive externalities.

5 Negative Externalities

In general, the Sender faces the following trade-off: the short-term gain from disclosing information about the current state, versus the benefit from influencing the evolution of the state, via the Receiver's action. Under negative externalities, her incentives today and tomorrow pull in opposite directions: the high action increases her static payoff, but it worsens the distribution of future states, and hence, lowers payoffs.

Recall that the Receiver wants to match his action to the state: he takes the high action when he believes the good state is more likely, and the low action otherwise. I refer to a mismatch between action and state, that is, the high action in the bad state or the low action in the good state, as a *mistake*. The policies that can be optimal differ in which of these mistakes they induce. In the next subsection, I introduce them through a few examples.

5.1 Policies

The four corners of the set of preference parameters $(x, y) \in [0, 1]^2$ illustrate the policies that can be optimal. At $(x, y) = (1, 0)$, the Sender has cardinal state-independent preferences: she only cares about the high action, and the greedy policy is optimal for all values of the transition probabilities

(Proposition 2). The greedy policy never induces the low action in the good state: its only mistake is the high action in the bad state. The other three corners lead to different policies.

Example 1. Suppose the Sender has cardinal action-independent preferences, that is, $(x, y) = (0, 1)$. This means that she is indifferent between the two actions in the stage game, and any disclosure is driven purely by dynamic incentives. In both states she prefers the low action, because it makes the good state more likely tomorrow. In the bad state, this is also what the Receiver wants; in the good state, it is not. She therefore induces the low action as often as possible.

I refer to the policy that does so as the *judicious policy*. Under the judicious policy, the Sender discloses nothing when $p_t \leq 1/2$, and when $p_t > 1/2$ she splits beliefs between $1/2$ and 1. In Figure 3, the black line is the Sender's interim payoff: her expected payoff from inducing the posterior q today and following the judicious policy from tomorrow on. The blue line is its concave envelope, her long-run payoff from the judicious policy as a function of the prior p. Note that the high action is never taken in the bad state: the high action follows a posterior belief of 1, when the state is good. This maximises the probability of the low action, which is taken with probability $\mu(p_t) = 1$ if $p_t \leq 1/2$ and with probability $\mu(1/2) = 2(1 - p_t)$ if $p_t > 1/2$. Equivalently, the judicious policy minimises the probability of the high action, and its only mistake is the low action in the good state. In this sense, the judicious policy completely reverses the greedy policy.

Example 2. Suppose $(x, y) = (0, 0)$, so that the Sender obtains a positive payoff only when the state is good and the Receiver takes the high action. Inducing the high action in the bad state gives her nothing today and makes the good state less likely tomorrow, so she has no reason to do it. Inducing the low action in the good state is a genuine trade-off: it sacrifices her payoff of 1 today, but raises the probability that tomorrow's state is good from p_G^h to p_G^ℓ. This future benefit is worth less than the payoff she gives up today (Section 5.3.1), so she has no reason to induce either mistake, and she reveals the state.

Full disclosure reveals the state every period: it is the only policy under which the Receiver never makes a mistake.

Example 3. Consider the following parameter specification: $\delta = 0.9$, $(x, y) = (1, 1)$, $p_G^h = 0.5$,

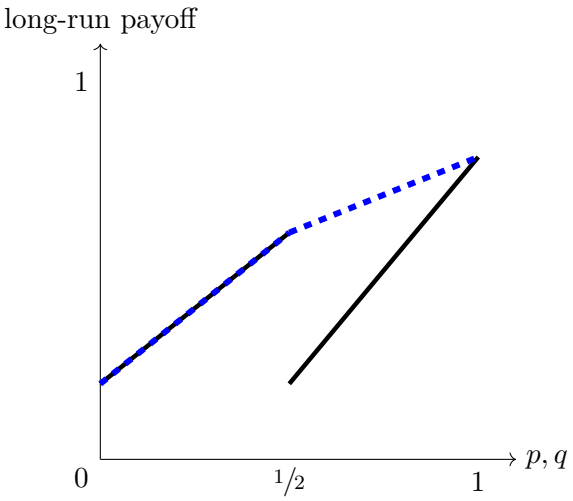


Figure 3: Sender's interim (black) and long-run (blue) expected payoff under the judicious policy

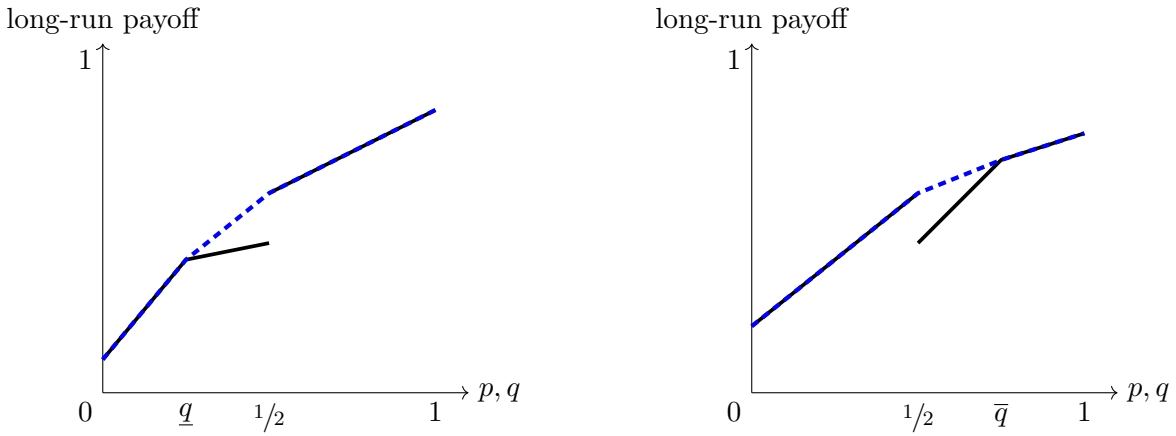
$p_B^h = 0.1$, $p_G^\ell = 0.9$, and $p_B^\ell = 0.4$. In the bad state, the high action gives the Sender 1 today instead of 0. In the good state, both actions give her 1 today, and the low action makes the good state more likely tomorrow. So she has a reason to induce both mistakes: the first raises her payoff today, and the second raises the probability of the good state in the future. At the prior $p = 0.3$, the optimal information structure induces the two posteriors $q^h = 0.5$ and $q^\ell = 0.2$, with $\alpha = 1/3$. At $q^h = 0.5$, the Receiver is indifferent and, under the Sender-favourable tie-breaking convention, takes the high action; at $q^\ell = 0.2$, he takes the low action. The Sender's long-run payoffs at $p = 0.3$ under this policy and the three policies discussed above are:

Policy	Long-run payoff at $p = 0.3$
Optimal policy	0.6667
Greedy policy	0.6610
Judicious policy	0.5792
Full disclosure	0.4286

The optimal policy in Example 3 is neither the greedy policy, nor the judicious policy, nor full disclosure: both of its posteriors are interior, so the Receiver is never certain of the state. Section 5.3.2 explains why.

I refer to the policy above as the *truncated greedy policy*, here with threshold $\underline{q} = 0.2$. In general, under the truncated greedy policy with threshold $\underline{q} \in (0, 1/2)$, the Sender discloses information only

at priors between $\underline{q}$ and $1/2$. There, as in the greedy policy, the Sender still raises the Receiver's belief to $1/2$ with some probability to obtain the high action today. However, unlike the greedy policy, she lowers the belief only to $\underline{q}$ instead of revealing the bad state, and at all other priors she discloses nothing (Figure 4). The *truncated judicious policy*, with threshold $\overline{q} \in (1/2, 1)$, is defined symmetrically: the Sender discloses information only at priors between $1/2$ and $\overline{q}$, where she splits the belief between $1/2$ and $\overline{q}$. Whenever the Sender discloses information under the greedy policy, the judicious policy or full disclosure, the Receiver is certain of at least one state with positive probability. The two truncated policies never make the Receiver certain of the state, and so induce both mistakes. Finally, I refer to the policy in which the Sender never discloses information as the *no-disclosure policy*.



(a) Truncated greedy policy

(b) Truncated judicious policy

Figure 4: Sender's interim (black) and long-run (blue) expected payoff under the truncated greedy and truncated judicious policies

In Appendix A.2, I show that it is without loss to restrict attention to policies in which the set of priors at which the Sender discloses information is either empty, the interval $(0, 1)$, or an interval with $1/2$ as an endpoint (see Lemmas 5 and 6). Each policy in this class is one of the six above. I use this to characterise the optimal policy in the next subsection.

5.2 Characterisation

Recall the thresholds from Section 4:

$$\bar{x} = \frac{\delta(p_B^\ell - p_B^h)}{1 + \delta p_B^\ell - \delta p_G^h}, \qquad \bar{y} = \frac{1 + \delta p_B^\ell - \delta p_G^\ell}{1 + \delta p_B^\ell - \delta p_G^h}$$

The following theorem characterises the optimal policy for every $(x, y) \in [0, 1]^2$:

Theorem 2 (Negative externalities)**.** Under negative externalities:

(i) Full disclosure is optimal if and only if $(x, y) \leq (\bar{x}, \bar{y})$,

(ii) The greedy policy is optimal if and only if $x \geq \bar{x}$ and $y \leq \min\{1, L_g(x)\}$,

(iii) The judicious policy is optimal if and only if $y \geq \bar{y}$ and $x \leq \min\{1, L_j(y)\}$,

(iv) Otherwise, there exists a no-disclosure line, with the optimal policy being truncated greedy below it and truncated judicious above it.

The functions $L_g(\cdot)$, $L_j(\cdot)$ and the no-disclosure line are affine, strictly increasing, and pass through $(\bar{x}, \bar{y})$. Their exact functional form depends on the values of the parameters: transition probabilities and discount factor.

Figure 5 illustrates the theorem.[18] Full disclosure is optimal in the rectangle below and to the left of the corner $(\bar{x}, \bar{y})$. The greedy policy is optimal to the right of $\bar{x}$ and below L_g, and the judicious policy above $\bar{y}$ and to the left of L_j. Above L_g and to the right of L_j, a truncated policy is optimal: truncated greedy below the no-disclosure line and truncated judicious above it. On the line, disclosing no information is optimal.

Thus, under negative externalities, various policies are optimal depending on the cardinal features of the Sender's preferences, captured through the value of (x, y). When a truncated policy is optimal, the Sender induces both mistakes, and the Receiver is never certain of the state. This interior pooling is the main qualitative difference from the existing literature: with exogenous transitions, a binary state space and a binary action space, the greedy policy is always optimal.

[18] Figure 5 shows one of four possible configurations: the no-disclosure line may coincide with L_g, with L_j, with neither (as shown), or with both, in which case the truncated policies are never optimal.

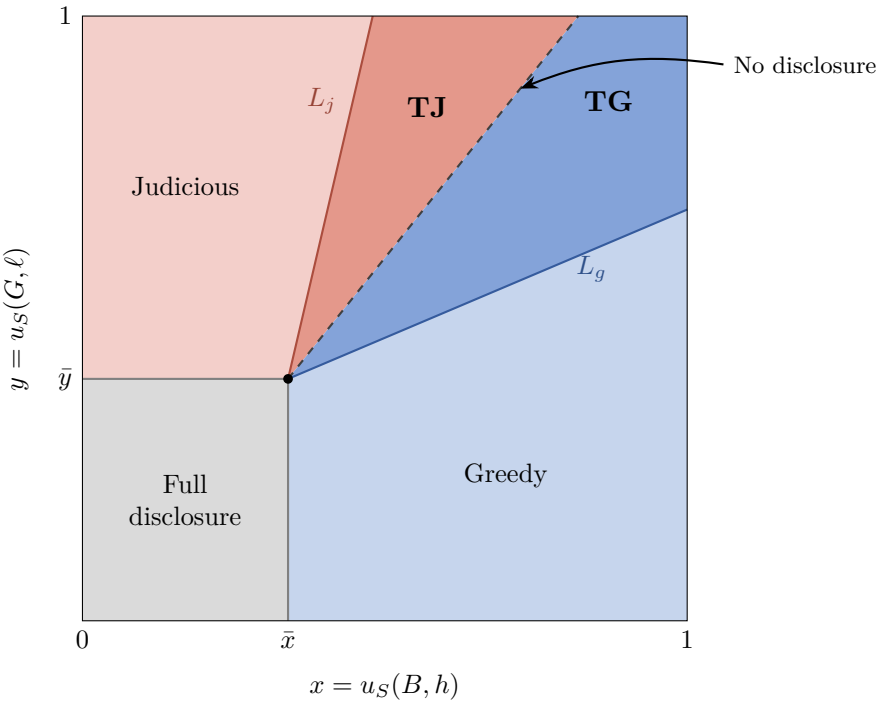


Figure 5: Optimal policies under negative externalities;
TG: truncated greedy; TJ: truncated judicious

5.3 Intuition

This subsection explains Theorem 2: first, why the regions in Figure 5 are arranged around the corner $(\bar{x}, \bar{y})$, and second, why a truncated policy can be optimal.

5.3.1 Where do Interests Conflict?

As in static persuasion, the Sender gains from withholding information only if her interests *conflict* with the Receiver's, that is, if in some state she prefers the action he would not choose there. She does not value mistakes as such, but a conflict in a state is exactly what makes a mistake in that state worth inducing. In the stage game, the only conflict is in the bad state, where she weakly prefers the high action; in the good state, both prefer it. However, with intertemporal externalities, today's action also changes tomorrow's state, and this changes the Sender's effective preferences. In the bad state, the high action gives her x today instead of 0, but makes the good state less likely tomorrow. In the good state, the low action costs her $1 - y$ today, but makes the good state more

likely tomorrow.

To weigh these effects against each other, suppose that from tomorrow on the Sender fully discloses the state. Then, compared with matching the action to the state, inducing the high action in the bad state today changes her payoff by $(1-\delta)(x-\bar{x})$, and inducing the low action in the good state changes it by $(1-\delta)(y-\bar{y})$.[19] Since $1-\bar{y} = \delta(p_G^\ell - p_G^h)/(1+\delta p_B^\ell - \delta p_G^h)$, both $\bar{x}$ and $1-\bar{y}$ take the same form: the discounted change in tomorrow's probability of the good state, multiplied by $1/(1+\delta p_B^\ell - \delta p_G^h)$. This factor is the value of the good state: under full disclosure, starting a period in the good rather than the bad state is worth as much to the Sender as a one-off payoff of $1/(1+\delta p_B^\ell - \delta p_G^h)$ today.[20] So $\bar{x}$ is the future cost of the high action in the bad state, and $1-\bar{y}$ is the future benefit of the low action in the good state. The Sender's interests therefore conflict with the Receiver's in the bad state if and only if $x > \bar{x}$, and in the good state if and only if $y > \bar{y}$.

Hence the corner $(\bar{x}, \bar{y})$ organises Figure 5: the Sender's interests conflict with the Receiver's in the bad state to its right, and in the good state above it. When interests are aligned in both states, she fully discloses the state. In Example 2, $(x, y) = (0, 0)$: the high action in the bad state gives her nothing today, and the future benefit of the low action in the good state, $1-\bar{y}$, is less than the payoff of 1 that she sacrifices today, since $\bar{y} > 0$, that is, $\delta(p_G^\ell - p_B^\ell) < 1$. In particular, the judicious policy sacrifices more today than it gains in the future. Because the full disclosure payoff is affine in the prior, no other information structure in the current period does better either, and by Lemma 1, full disclosure is optimal. The same argument applies whenever $x \leq \bar{x}$ and $y \leq \bar{y}$ (see Appendix D). Under negative externalities and $\delta > 0$, both thresholds are positive, so this region always contains a neighbourhood of $(0, 0)$.

When interests conflict only in the bad state, the Sender induces the high action there as often as possible, which is the greedy policy, as at $(x, y) = (1, 0)$. When they conflict only in the good state, she induces the low action there as often as possible, which is the judicious policy, as in

[19] Under full disclosure, the Sender's payoff from prior p is $\theta(p) = \frac{(1-\delta)p + \delta p_B^\ell}{1+\delta p_B^\ell - \delta p_G^h}$ (see Appendix D). In the bad state, the high action instead of the low action changes the payoff by $(1-\delta)x - \delta[\theta(p_B^\ell) - \theta(p_B^h)] = (1-\delta)(x-\bar{x})$; in the good state, the low action instead of the high action changes it by $(1-\delta)(y-1) + \delta[\theta(p_G^\ell) - \theta(p_G^h)] = (1-\delta)(y-\bar{y})$.

[20] This is because $\theta(1) - \theta(0) = \frac{1-\delta}{1+\delta p_B^\ell - \delta p_G^h}$.

Example 1, where $x = 0 \leq \bar{x}$ and $y = 1 \geq \bar{y}$. These are the one-sided persuasion problems familiar from the static literature.

The greedy and judicious policies remain optimal when interests also conflict in the other state, provided that this second conflict is weak. The two conflicts compete, because a policy that maximises the probability of one action also minimises that of the other action. Under the greedy policy, the Sender induces the high action whenever she can. As y rises above $\bar{y}$, the conflict in the good state grows, and above $L_g(x)$ she prefers to give up some high actions in the bad state in order to induce some low actions in the good state. The higher the value of x, the stronger the conflict in the bad state, and the higher y must be before she makes this trade; hence L_g is increasing. The same logic applies to the judicious policy and L_j.[21]

The same thresholds also explain Theorem 1. Under positive externalities, $\bar{x} \leq 0$ and $\bar{y} \geq 1$: the externality reinforces the static conflict in the bad state and never creates one in the good state. The Sender behaves like a static Sender biased towards the high action, and the greedy policy is optimal. Under negative externalities, $\bar{x} \geq 0$ and $\bar{y} \leq 1$, and the externality works the other way: it can remove the conflict in the bad state and create one in the good state. In this sense, the externality determines whose side the Sender is on in each state.[22]

The thresholds $\bar{x}$ and $\bar{y}$ assume full disclosure from tomorrow on. They therefore pin down when full disclosure is optimal, but not the whole solution: when interests conflict in both states, the Sender must also take into account where today's disclosure leaves tomorrow's belief. I turn to this second force next.

5.3.2 Conflict in Both States: The Truncated Policies

Example 3 (revisited)**.** In Example 3, $\bar{x} \approx 0.297$ and $\bar{y} \approx 0.604$, so $x > \bar{x}$ and $y > \bar{y}$, and the conflict in neither state is weak enough to ignore. At a prior between $\underline{q}$ and $1/2$, such as $p = 0.3$, the Sender splits the Receiver's belief between $1/2$ and $\underline{q} = 0.2$. To see why, recall that, with full disclosure from tomorrow onward, a mistake in the bad state is worth $(1 - \delta)(x - \bar{x})$ to the Sender

[21] See Appendices B, C and E.4 for details.

[22] In Appendix F, I apply the same reasoning to mixed externalities, where the externalities have different signs in the two states.

and a mistake in the good state is worth $(1-\delta)(y-\bar{y})$, relative to matching the action to the state. Since the common factor $1-\delta$ does not affect any comparison, I refer to $x-\bar{x}$ and $y-\bar{y}$ as the values of the two mistakes; in the example, they are approximately 0.70 and 0.40. At a posterior $q \geq 1/2$, the Receiver takes the high action, which results in a mistake with probability $1-q$; at a posterior $q < 1/2$, he takes the low action, which results in a mistake with probability q.

The upper posterior. Under negative externalities, the high action makes the good state less likely tomorrow in both states, so it may seem that the Sender should induce it less often, by using an upper posterior above $1/2$. Suppose the Sender uses a higher upper posterior. This raises the probability of the low action, but the low action is taken at $\underline{q} < 1/2$, where the state is bad with probability $1-\underline{q} > 1/2$. Thus, most of the increase falls in the bad state. In the bad state, the Sender prefers the high action, since its gain today, x, outweighs its cost tomorrow, $\bar{x}$. For each unit increase in the probability of the low action in the good state, the probability of the high action in the bad state falls by $(1-\underline{q})/\underline{q}$ units, four in the example.[23] Since each bad-state mistake is worth $x-\bar{x} \approx 0.70$ and each good-state mistake $y-\bar{y} \approx 0.40$, she loses 4×0.70 to gain 0.40. Thus, since $1/2$ is the lowest posterior at which the Receiver takes the high action, the Sender keeps the upper posterior at exactly $1/2$.

The lower posterior. With the upper posterior at $1/2$, the Receiver's chance of erring always equals his belief that the state is good: at $1/2$ he takes the high action, which is wrong half the time, and at the lower posterior he takes the low action, which is wrong with probability equal to the lower posterior. Since his beliefs average to the prior, he errs with probability p, whatever the lower posterior. Therefore, in the current period, the lower posterior determines only which mistake he makes. Under the greedy policy, where the lower posterior is 0, every mistake is the high action in the bad state. If the Sender discloses nothing, every mistake is the low action in the good state. Raising the lower posterior from 0 moves between these extremes, converting mistakes in the bad state into mistakes in the good state one-for-one. At $p = 0.3$, the greedy policy induces

[23] For example, at $p = 0.3$, raising the upper posterior from $1/2$ to 0.6 raises the probability of the low action from $2/3$ to $3/4$. The probability of the low action in the good state therefore rises by $1/60$, and that of the high action in the bad state falls by $1/15$.

the high action in the bad state with probability 0.3. On the other hand, the truncated greedy policy induces the high action in the bad state with probability $1/6$, and the low action in the good state with probability $2/15$.[24] Therefore, compared to the greedy policy, the truncated greedy policy gives up mistakes worth 0.70 each for mistakes worth only 0.40 each. Why would the Sender make this trade?

Tomorrow's belief. If each mistake had a fixed value, she would not: the one-for-one conversion would lead the Sender to an extreme, the greedy policy here. However, the values of the mistakes assumed full disclosure from tomorrow on, under which the value of tomorrow's belief is linear. Under the optimal policy it is not: its slope changes at $1/2$, where the Receiver's action switches.[25] In the example, raising the lower posterior raises tomorrow's prior after the low action. While this prior is below $1/2$, raising it is valuable, since the closer it is to $1/2$, the more often the Sender can induce the high action tomorrow. Once it reaches $1/2$, the Receiver takes the high action without any disclosure, while being as uncertain as possible. Raising it further gains less, since he takes the high action anyway, and in the example this no longer outweighs the cost today. The Sender therefore raises the lower posterior until tomorrow's prior is exactly $1/2$: $f^\ell(0.2) = 0.4 + 0.2(0.9 - 0.4) = 1/2$. This is a general feature of the truncated policies: starting from the prior $f^\ell(\underline{q})$, the belief either reaches exactly $1/2$, without disclosure along the way, or never returns to a prior between $\underline{q}$ and $1/2$.[26]

The truncated judicious policy makes the symmetric trade, starting from the judicious policy. Which of the two is optimal depends on which conflict is stronger: the truncated greedy policy below the no-disclosure line, and the truncated judicious policy above it. On the line the two are balanced, and disclosing no information at all is optimal.

When are the truncated policies optimal? The set of preferences at which none of the greedy, judicious and full disclosure policies is optimal is $\{(x, y) > (\bar{x}, \bar{y}) : y > L_g(x), x > L_j(y)\}$, and this

[24] By Bayes plausibility, the greedy policy induces the posterior $1/2$ with probability 0.6, and the state is then bad half the time. The truncated greedy policy induces $1/2$ with probability $1/3$ and 0.2 with probability $2/3$, giving $1/3 \cdot 1/2 = 1/6$ and $2/3 \cdot 0.2 = 2/15$. In both cases, the probabilities of the two mistakes sum to 0.3.

[25] This correction does not change the upper posterior: there, each unit of good-state mistakes costs four units of bad-state mistakes, and the margin (4×0.70 against 0.40) is wide enough to survive it.

[26] See Lemma 19 in Appendix E.4. Section 6 describes the resulting belief dynamics.

set may be empty.[27] As Example 3 shows, the truncated greedy policy relies on the Sender's ability to change the next period's *default action*—the action the Receiver takes when he receives no new information—through the posterior at which she induces the low action today.[28] In the example, the low action at any posterior below 0.2 leaves tomorrow's prior below $1/2$, so the next Receiver takes the low action by default; at 0.2 or above, it leaves the prior at or above $1/2$, so he takes the high action. If tomorrow's default action is instead determined by today's action alone, the Sender cannot change it through the posterior.[29] Tomorrow's value is then affine in today's posterior for each action, as under full disclosure, so each mistake has a fixed value to the Sender, and the one-for-one conversion between them leads her to one of the extremes.

5.3.3 Proof Outline

The proof of Theorem 2, covered in Appendices B–E, applies Lemma 1 to each candidate policy. The payoff from the greedy policy is concave if and only if $x \geq \bar{x}$, and since the one-shot deviation condition is affine in (x, y), it binds along a line through $(\bar{x}, \bar{y})$, which is L_g; the judicious policy and L_j are symmetric. In the remaining region, I show that the no-disclosure line exists, and characterise the thresholds $\underline{q}$ and $\overline{q}$.

5.4 Why is the Characterisation Asymmetric?

The judicious policy is the mirror image of the greedy policy, yet in Figure 5 the greedy region is larger, and L_g and L_j have different slopes. There are three reasons. First, the stage game favours the high action. Its static gain, x in the bad state and $1 - y$ in the good state, can be as large as 1 in each state, whereas its dynamic costs, $\bar{x}$ and $1 - \bar{y}$, are discounted and sum to less than 1:

$$\bar{x} + (1 - \bar{y}) = \frac{\delta[(p_B^\ell - p_B^h) + (p_G^\ell - p_G^h)]}{1 + \delta p_B^\ell - \delta p_G^h} < 1$$

[27] Lemmas 17 and 18 in Appendix E.4 give the necessary and sufficient conditions for non-emptiness.

[28] Recall from equation (3.2) that the period-$t+1$ prior p_{t+1} is a function of the period-t posterior q_t and period-t action a_t.

[29] This happens when: (a) the default action is always the high action, i.e. $f^h([1/2, 1])$ and $f^\ell([0, 1/2])$ both lie in $[1/2, 1]$; (b) the default action is always the low action, i.e. $f^h([1/2, 1])$ and $f^\ell([0, 1/2])$ both lie in $[0, 1/2]$; and (c) the default alternates, so that each action is followed by the other, i.e. $f^h([1/2, 1])$ lies in $[0, 1/2]$ and $f^\ell([0, 1/2])$ lies in $[1/2, 1]$. The remaining case, in which each action is its own default, is not possible under negative externalities as $f^\ell(1/2) - f^h(1/2) = (p_G^\ell - p_G^h + p_B^\ell - p_B^h)/2 > 0$.

since this inequality is equivalent to $\delta(p_G^\ell - p_B^h) < 1$. Thus, $\bar{y} > \bar{x}$: the corner $(\bar{x}, \bar{y})$ lies above the 45-degree line, and the rectangle below and to its right, where the greedy policy is always optimal, is larger than the rectangle above and to its left, where the judicious policy is always optimal. Second, the externalities may differ in strength across states: $\bar{x}$ and $1 - \bar{y}$ are proportional to the externalities in the bad and the good state, respectively, with the same factor.

Third, even when the externalities are equally strong, the two actions may affect the state differently. The slopes of L_g and L_j depend not only on how much each action changes tomorrow's probability of the good state, but also on the levels of the transition probabilities, and the greedy policy relies on the high action while the judicious policy relies on the low action. So L_g and L_j need not be mirror images. The two actions affect the state symmetrically if the high action keeps the good state as often as the low action keeps the bad state, and turns the bad state good as often as the low action turns the good state bad, that is, if $p_G^h + p_B^\ell = 1$ and $p_B^h + p_G^\ell = 1$. With such symmetric transitions, the externalities are also equally strong in the two states, so $\bar{x} = 1 - \bar{y}$, and L_g and L_j are mirror images of each other about the 45-degree line through $(\bar{x}, \bar{y})$. Thus, the second and third sources both disappear.[30]

The first source is the hardest to remove. With symmetric transitions, $\bar{y} - \bar{x}$ converges to p_B^h / p_B^ℓ as $\delta \to 1$, which is positive unless $p_B^h = 0$; so even a perfectly patient Sender typically faces a corner above the 45-degree line. The corner lies on it only in the limit in which, in addition, each mistake locks in the current state: the low action keeps the good state for sure and the high action keeps the bad state for sure ($p_G^\ell = 1$ and $p_B^h = 0$). In this limit, $(\bar{x}, \bar{y})$ converges to $(1/2, 1/2)$, the centre of the unit square, and none of the three sources of asymmetry remains.

5.5 Comparative Statics

How do the parameters affect the optimal policy? Recall the two conflicts: the Sender's interests conflict with the Receiver's in the bad state if and only if $x > \bar{x}$, and in the good state if and only if $y > \bar{y}$. The next result shows how the Sender's patience and the transition probabilities move

[30] Relabelling the states, the actions and the belief ($G \leftrightarrow B$, $h \leftrightarrow \ell$, $p \leftrightarrow 1 - p$) leaves symmetric transitions unchanged and turns the greedy policy into the judicious policy and each mistake into the other. Since the slopes of L_g and L_j depend only on how often each mistake occurs under these policies and their one-shot deviations (Appendices B and C), L_j is the reflection of L_g.

these two thresholds.

Proposition 3 (Comparative statics)**.** Under negative externalities, with $\delta \in (0,1)$:

(i) *Patience:* $\bar{x}$ is strictly increasing and $\bar{y}$ is strictly decreasing in δ.

(ii) *Transition probabilities after a mistake:* $\bar{x}$ is strictly decreasing in p_B^h and does not depend on p_G^ℓ; $\bar{y}$ is strictly decreasing in p_G^ℓ and does not depend on p_B^h.

(iii) *Transition probabilities after a matching action:* $\bar{x}$ and $\bar{y}$ are both strictly increasing in p_G^h and in p_B^ℓ.

By part (i), patience weakens the conflict in the bad state and strengthens the one in the good state. In other words, a more patient Sender needs a larger static gain before she prefers the high action in the bad state, and accepts a larger static loss to obtain the low action in the good state. So patience does not simply make the Sender more transparent: she reveals the bad state more readily, but is more willing to obscure the good one. In Figure 5, the corner $(\bar{x}, \bar{y})$ moves right and down, and the full disclosure region becomes wider and shorter.

Part (ii) concerns the transitions that follow a mistake, p_B^h after the high action in the bad state and p_G^ℓ after the low action in the good state. Each affects only the threshold for its own mistake. A lower p_B^h makes the high action more harmful in the bad state, and so weakens the conflict there. A higher p_G^ℓ makes the low action more beneficial in the good state, and so strengthens the conflict there.

Part (iii) concerns p_G^h and p_B^ℓ, the transition probabilities after an action that matches the state. Each has two effects. It changes the value of starting a period in the good rather than the bad state, $1/(1+\delta p_B^\ell - \delta p_G^h)$, which scales both $\bar{x}$ and $1-\bar{y}$, and it changes the externality in one state. A higher p_G^h raises this value, so $\bar{x}$ rises; it also makes the low action less beneficial in the good state, which outweighs the higher value (since $1-\bar{y}<1$), so $\bar{y}$ rises too. Symmetrically, a higher p_B^ℓ lowers the value, so $\bar{y}$ rises, and makes the high action more harmful in the bad state, which outweighs the lower value (since $\bar{x}<1$), so $\bar{x}$ rises too. Either way, the corner moves up and to the right, and both conflicts weaken.

The transition probabilities also affect the slopes of L_g and L_j, which are given in Appendices B and C.

6 Belief Dynamics and Disclosure Patterns

The previous two sections characterise which policy is optimal. This section shows that the optimal policy also determines *when* the Sender discloses. While the optimal policies are stationary, i.e. they depend on calendar time only through the current belief, the belief moves over time, partly because of the actions the Sender induces (equation (3.2)). Thus, except under full disclosure, the Sender may stay silent for some periods, during which Receivers act without new information. Each such silent stretch may end with a new disclosure, or last forever.

In this section, I assume that every transition probability lies in $(0, 1)$: from either state, and after either action, both states can occur in the following period. This implies that every belief reached after the first period lies in $(0, 1)$.

6.1 From Policies to Belief Paths

Each policy $\pi \in \{$full disclosure (fd), greedy (g), judicious (j), truncated greedy (tg), truncated judicious $(tj)\}$ in Section 5 can be described by a communication interval $I_\pi = (q_\pi^\ell, q_\pi^h)$, with $q_\pi^\ell \leq 1/2 \leq q_\pi^h$:

- At any prior $p \in I_\pi$, the Sender splits the belief between the endpoints of the interval. She induces the posterior q_π^h, at which the Receiver takes action h, with probability $\alpha_\pi(p) = (p - q_\pi^\ell)/(q_\pi^h - q_\pi^\ell)$, and the posterior q_π^ℓ, at which he takes action ℓ, otherwise.

- At any prior $p \notin I_\pi$, the Sender discloses nothing. The Receiver takes action h if $p > 1/2$, action ℓ if $p < 1/2$, and breaks ties in the Sender's favour at $p = 1/2$.

Table 1 gives the communication interval of each policy; under the no-disclosure policy, $I_\pi = \emptyset$.

Also recall that, given the current action, the function mapping the current posterior to the next period's prior is given by f^a. Thus, outside I_π when no information is disclosed, the following

Policy	I_π	q_π^ℓ	q_π^h
Full disclosure	$(0,1)$	0	1
Greedy	$(0,1/2)$	0	$1/2$
Judicious	$(1/2,1)$	$1/2$	1
Truncated greedy	$(\underline{q},1/2)$	$\underline{q}$	$1/2$
Truncated judicious	$(1/2,\overline{q})$	$1/2$	$\overline{q}$

Table 1: Policies as a communication interval with two posteriors

period's prior is given by $p_{t+1} = F_\pi(p_t) \coloneqq f^{a_\pi(p_t)}(p_t)$, where $a_\pi(p)$ is the action the Receiver takes at the prior p when nothing is disclosed. Starting from prior p_t, if no information is disclosed for $n \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ periods,[31] the prior in period $t+n$ is given by $\underbrace{F_\pi(F_\pi \cdots (F_\pi}_{n \text{ times}}(p_t))) \coloneqq F_\pi^{(n)}(p_t)$, where $F_\pi^{(0)}(p) = p$.[32] In the communication interval I_π, the next prior is one of two *landing points*: $z_\pi^h \coloneqq f^h(q_\pi^h)$ or $z_\pi^\ell \coloneqq f^\ell(q_\pi^\ell)$.

For $a \in \mathcal{A}$, let the length of the *quiet spell* that follows a disclosure inducing a be denoted by $n_\pi^a \coloneqq \min\{n \geq 0 : F_\pi^{(n)}(z_\pi^a) \in I_\pi\} \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$. I refer to a disclosure period together with the quiet spell that follows it as a *disclosure episode*. If the Sender induces a, a disclosure episode lasts $1 + n_\pi^a$ periods. Let $r_\pi^a \coloneqq F_\pi^{(n_\pi^a)}(z_\pi^a), n_\pi^a < \infty$ denote the prior at which disclosure resumes. Then:

Lemma 2 (Belief paths)**.** Fix a policy π and an initial prior $p_1 \in (0,1)$ and let $\tau \coloneqq \min\{t : p_t \in I_\pi\}$.

(i) Before τ, the Sender discloses nothing and the belief drifts deterministically: $p_{t+1} = F_\pi(p_t)$ for $t < \tau$.

(ii) From τ on, a disclosure that induces action a moves the belief to z_π^a. The Sender then discloses nothing for n_π^a periods, while the belief drifts deterministically. If $n_\pi^a < \infty$, then the belief reaches r_π^a, where the Sender discloses again.

(iii) If $n_\pi^h, n_\pi^\ell < \infty$, then $(p_t)_{t>\tau}$ is a Markov chain on finitely many beliefs. Its only random transitions are at r_π^h and r_π^ℓ: from r_π^a, it moves to z_π^h with probability $\alpha_\pi(r_\pi^a)$ and to z_π^ℓ otherwise.

[31] When $n = 0$, the interpretation is that information *is* disclosed in period t. When $n = \infty$, then information is *never* disclosed from period t onward.

[32] To be clear, the drift in beliefs during quiet spells can move across regions where the high and low actions are optimal under no information.

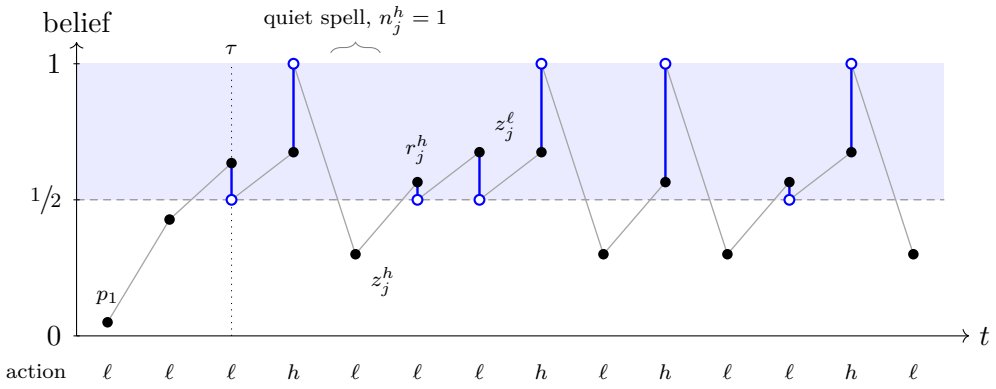


Figure 6: Belief path under the judicious policy: $p_G^h = 0.3$, $p_B^h = 0.1$, $p_G^\ell = 0.95$, $p_B^\ell = 0.4$. Black dots are priors and blue circles are the posteriors induced by disclosure. The shaded band is the communication interval $(1/2, 1)$, and the row below shows the action taken.

Lemma 2 says that beliefs follow a simple path. Before the first disclosure, the belief drifts deterministically. Each disclosure induces one of the two posteriors, so the next prior is one of the two landing points, z_π^h or z_π^ℓ. From there, the belief again drifts without disclosure, for n_π^a periods, until it re-enters the communication interval at r_π^a and the Sender discloses again. In Figure 6, the initial prior $p_1 = 0.05$ lies outside the communication interval $(1/2, 1)$, so the Sender discloses nothing until the belief first enters it, at $\tau = 3$. After the first disclosure, the belief only visits points on two paths, and the only randomness is which posterior each disclosure induces. The public history is therefore a sequence of disclosure episodes, whose lengths depend only on the actions induced.

6.2 When Does Disclosure Stop?

Each of the greedy, judicious and truncated policies uses $1/2$, where the Receiver is indifferent, as one of its two posteriors. I refer to the action he takes there as the *threshold action*. This is the high action under the greedy and truncated greedy policies, and the low action under the other two. Let p_*^a denote the fixed point of f^a. Then:

Proposition 4 (Termination)**.**

(i) *Greedy policy.* Suppose $p_G^h + p_B^h \geq 1$. Then from any prior, the Sender almost surely discloses in only finitely many periods. Afterwards the Receiver takes action h in every period and the

belief converges to $p_*^h \geq 1/2$.

(ii) *Judicious policy.* Suppose $p_G^\ell + p_B^\ell \leq 1$. Then from any prior, the Sender almost surely discloses in only finitely many periods. Afterwards the Receiver takes action ℓ in every period and the belief converges to $p_*^\ell \leq 1/2$.

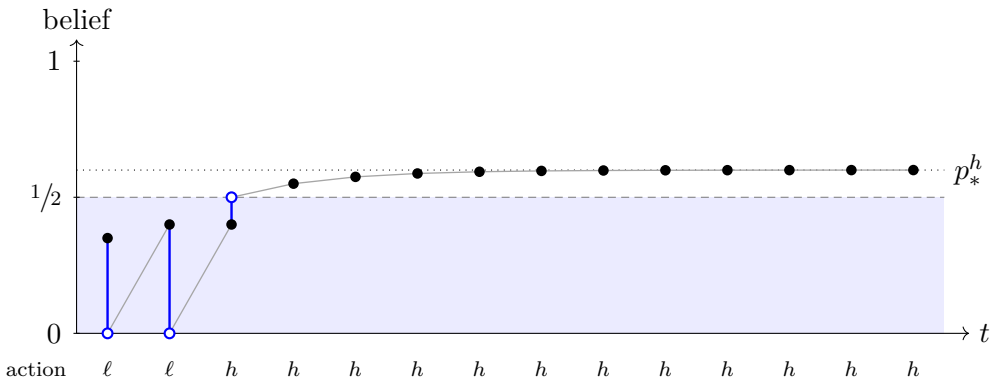


Figure 7: Disclosure stops under the greedy policy: $p_G^h = 0.8$, $p_B^h = 0.3$, $p_G^\ell = 0.9$, $p_B^\ell = 0.4$, so $p_G^h + p_B^h \geq 1$ and the belief converges towards p_*^h. Black dots are priors and blue circles are the posteriors induced by disclosure. The shaded band is the communication interval $(0, 1/2)$, and the row below shows the action taken.

Figure 7 illustrates part (i) of Proposition 4. The condition in part (i) is $f^h(1/2) \geq 1/2$: one period of the high action taken at the threshold leaves the next, uninformed Receiver willing to take the high action again.[33] In this sense the high action is *self-sustaining*: once the Sender sends a message that only just persuades the Receiver, she never needs to speak again, and greedy disclosure is a transitional phenomenon. In the fisheries application, this is the case when the stock can sustain high effort: if fleets fished at high effort every season, in the long run the stock would be abundant at least half the time. Then, once an assessment just persuades fleets to fish at high effort, later fleets keep doing so without new assessments. The condition in part (ii) says the same of the low action. On the other hand:

Proposition 5 (Recurring disclosure)**.**

[33] Propositions 4 and 5 are about the greedy and judicious policies themselves: given a policy, beliefs evolve by the transition probabilities alone, so the results do not depend on (x, y) or δ. Under negative externalities, however, each policy is optimal for some preferences (Theorem 2), so every case in the two propositions describes the optimal dynamics for some (x, y).

(i) *Greedy policy.* Suppose $p_G^h + p_B^h < 1$. Then from any prior, the Sender discloses in infinitely many periods. A disclosure that induces action h is followed by another disclosure in the next period. A disclosure that induces action ℓ is followed by $n_g^\ell \in \mathbb{Z}_{\geq 0}$ periods of action h without disclosure.

(ii) *Judicious policy.* Suppose $p_G^\ell + p_B^\ell > 1$. Then from any prior, the Sender discloses in infinitely many periods. A disclosure that induces action ℓ is followed by another disclosure in the next period. A disclosure that induces action h is followed by $n_j^h \in \mathbb{Z}_{\geq 0}$ periods of action ℓ without disclosure.

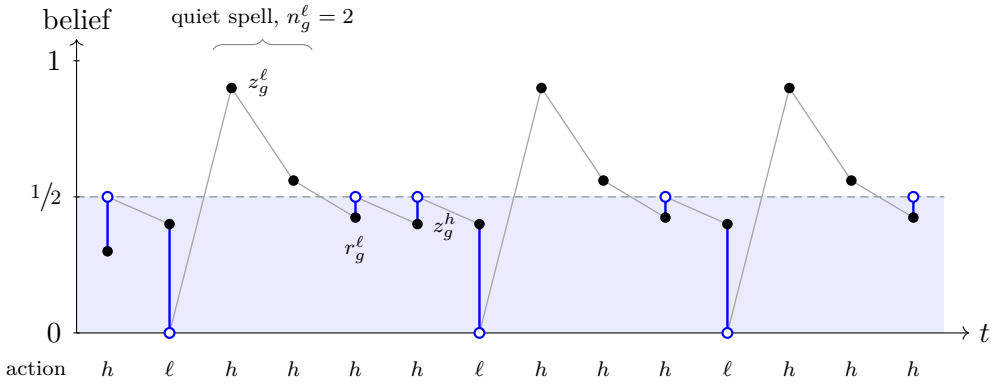


Figure 8: Disclosure recurs under the greedy policy: $p_G^h = 0.6$, $p_B^h = 0.2$, $p_G^\ell = 0.95$, $p_B^\ell = 0.9$, so $p_G^h + p_B^h < 1$. Black dots are priors and blue circles are the posteriors induced by disclosure. The shaded band is the communication interval $(0, 1/2)$, and the row below shows the action taken.

Figure 8 illustrates part (i) of Proposition 5. Disclosure recurs if the threshold action undermines itself. Under the greedy policy, $p_G^h + p_B^h < 1$ means $f^h(1/2) < 1/2$: a Receiver who is persuaded just enough to take the high action depletes the state, so that the next Receiver, left uninformed, would take the low action. The Sender therefore cannot rest on a message that induces h and must speak again in the following period. A message that induces the low action can work differently. The low action replenishes the state: if $p_B^\ell \geq 1/2$, the next Receivers start optimistic and take the high action without being told anything. The Sender stays silent while this optimism lasts, for n_g^ℓ periods, and speaks again once the drift under the high action has brought the belief back below $1/2$. (If $p_B^\ell < 1/2$, the Sender speaks again immediately after either message.) In this sense good

news is spent at once, whereas bad news can be an investment that pays out in periods of the high action obtained for free. In the fisheries application, under greedy disclosure an assessment that leads fleets to fish at high effort is followed by another assessment the next season, whereas an assessment that leads them to restrain their effort is followed by possibly several seasons of high effort with no new assessment. The judicious policy has the mirror-image rhythm.

Propositions 4 and 5 imply the following:

Corollary 1 (At most one self-sustaining action)**.** Under negative externalities, $p_G^\ell + p_B^\ell > p_G^h + p_B^h$. Hence at most one of the greedy and judicious policies has disclosure that eventually stops. Under full disclosure, disclosure never stops.

Under negative externalities, the low action makes the good state more likely than the high action does. So if the high action at $1/2$ leaves the next Receiver willing to take the high action, the low action at $1/2$ leaves him unwilling to take the low action, and vice versa. At most one of the two actions sustains itself.

The truncated policies have an additional feature. Under the truncated greedy policy, a disclosure that induces the posterior $\underline{q}$ leads to the prior $f^\ell(\underline{q})$. By Lemma 19 in Appendix E.4, the belief path that starts from this prior, absent further disclosure, either reaches exactly $1/2$ without first entering I_{tg}, or never returns to I_{tg}. In Example 3, it reaches $1/2$ at once: $f^\ell(0.2) = 1/2$. The truncated judicious policy is symmetric, starting from $f^h(\overline{q})$. Combined with Lemma 2, this gives the following:

Corollary 2 (Truncated policies)**.** Under a truncated greedy or truncated judicious policy with $n_\pi^h, n_\pi^\ell < \infty$, $r_\pi^h = r_\pi^\ell := r_\pi$. Every disclosure episode after the first therefore starts from the same prior r_π, and the actions they induce are i.i.d., with action h induced with probability $\alpha_\pi(r_\pi)$.

Figure 9 illustrates Corollary 2. The corollary requires both quiet spells to be finite, so that the belief re-enters the communication interval after every disclosure. After a disclosure that induces the truncated posterior ($\underline{q}$ or $\overline{q}$), the belief then reaches exactly $1/2$. Since $1/2$ is an endpoint of the communication interval, the Sender does not disclose there, and the next Receiver takes the threshold action without new information. From then on, the belief follows the same path as after

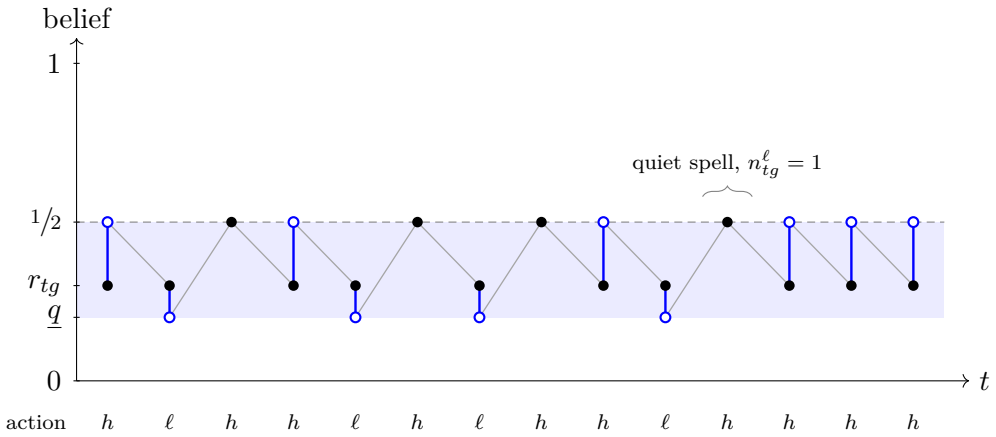


Figure 9: Belief path under the truncated greedy policy in Example 3. Black dots are priors and blue circles are the posteriors induced by disclosure. The shaded band is the communication interval $(\underline{q}, 1/2) = (0.2, 1/2)$, and the row below shows the action taken. Every disclosure episode starts from $r_{tg} = 0.3$.

a disclosure that induces $1/2$, so both quiet spells end at r_π. Therefore, after the first disclosure, every disclosure episode starts from the same prior, whatever the previous one induced. Under the greedy and judicious policies, by contrast, the prior at which disclosure resumes generally depends on the last action induced, so actions are correlated across episodes.

Example 3 (revisited)**.** Under the truncated greedy policy, $I_{tg} = (0.2, 1/2)$. A disclosure that induces the high action leads to the prior $f^h(1/2) = 0.3$, which lies in I_{tg}, so the Sender discloses again in the next period. A disclosure that induces the low action leads to the prior $f^\ell(0.2) = 1/2$. The next Receiver then takes the high action without new information, and the prior returns to $f^h(1/2) = 0.3$. So every disclosure episode starts from $r_{tg} = 0.3$ and induces the high action with probability $\alpha_{tg}(0.3) = 1/3$. A disclosure that induces the high action is followed by another in the next period, and one that induces the low action is followed by one period of no disclosure.

When are both quiet spells finite? Under the truncated greedy policy, the answer depends only on n^ℓ_{tg}. If n^ℓ_{tg} is finite, the belief passes through $1/2$ during the quiet spell after the low action, and then follows the same path as after a disclosure that induces $1/2$. So n^h_{tg} is finite whenever n^ℓ_{tg} is. There are therefore two cases. If $n^\ell_{tg} < \infty$, disclosure recurs forever once it starts, as in Corollary 2. If $n^\ell_{tg} = \infty$, disclosure stops at the first disclosure that induces the low action. Every

episode induces the low action with positive probability, so the Sender almost surely discloses in only finitely many periods.

The second case arises, for example, when the threshold action sustains itself, i.e. $p_G^h + p_B^h \geq 1$. After a high action at $1/2$, the belief then never falls below $1/2$ again, so neither quiet spell ends, and the Sender discloses at most once. The greedy policy, by contrast, can disclose several times before it stops, because the belief after a disclosure that induces the low action, p_B^ℓ, can return to the communication interval. The truncated posterior instead sends the belief to $1/2$, where the high action sustains itself.

The truncated judicious policy is symmetric, with n_{tj}^h in place of n_{tg}^ℓ and $p_G^\ell + p_B^\ell \leq 1$ in place of $p_G^h + p_B^h \geq 1$.

7 Conclusion

This paper studies an expert, the Sender, who informs a sequence of Receivers whose actions shape the state their successors inherit. These intertemporal externalities change what information is for: beyond steering today's action, the expert uses it to steer the state, and through it, tomorrow's actions. In effect, the externality determines, state by state, whether the expert's interests are aligned with the Receivers' or opposed to them.

When the externality is positive, the static and dynamic motives align, and the static benchmark survives intact: the greedy policy is optimal whatever the Sender's cardinal preferences. When it is negative, it can remove the conflict of interest in the bad state and create one in the good state. The Sender then fully discloses the state if neither conflict is present, uses the greedy or the judicious policy if one dominates, and pools at interior beliefs if both are strong, never making the Receiver certain of the state. This interior pooling is the main departure from persuasion about an exogenous state, where, with two states and two actions, the greedy policy is always optimal.

Three implications follow. First, an expert who cares about the future may withhold information because of that concern, not despite it: under negative externalities, she may hold back good news to restrain current activity. Second, patience does not simply make her more transparent: a more patient expert reveals bad states more readily but is more willing to obscure good ones. Third, the

timing of disclosure is part of the optimal policy, since the content of one message determines when the next one comes. For a fisheries body, the frequency and conclusiveness of stock assessments are therefore outcomes of the optimal policy, and a natural object for empirical work.

The analysis suggests several directions for future research. One is to allow the Sender and the Receivers to hold different prior beliefs about the state. Another is to let the Sender inform several Receivers within a period, so that she must coordinate or miscoordinate their actions while also steering the state across periods. A third is to characterise optimal policies with finitely many states and actions; Appendix G takes a first step.

References

Adda, J. (2016). Economic activity and the spread of viral diseases: Evidence from high frequency data. *Quarterly Journal of Economics*, *131*(2), 891–941. https://doi.org/10.1093/qje/qjw005

Ashkenazi-Golan, G., Hernández, P., Neeman, Z., & Solan, E. (2023). Markovian persuasion with two states. *Games and Economic Behavior*, *142*, 292–314. https://doi.org/10.1016/j.geb.2023.08.001

Aumann, R. J., & Maschler, M. B. (1995). *Repeated games with incomplete information*. MIT Press.

Ball, I. (2023). Dynamic information provision: Rewarding the past and guiding the future. *Econometrica*, *91*(4), 1363–1391. https://doi.org/10.3982/ECTA17345

Becker, G. S., & Tomes, N. (1979). An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy*, *87*(6), 1153–1189. https://doi.org/10.1086/260831

Becker, G. S., & Tomes, N. (1986). Human capital and the rise and fall of families. *Journal of Labor Economics*, *4*(3, Part 2), S1–S39. https://doi.org/10.1086/298118

Bizer, D. S., & DeMarzo, P. M. (1992). Sequential banking. *Journal of Political Economy*, *100*(1), 41–61. https://doi.org/10.1086/261806

Blackwell, D. (1951). Comparison of experiments. In J. Neyman (Ed.), *Proceedings of the second Berkeley symposium on mathematical statistics and probability* (pp. 93–102). University of California Press.

Dewatripont, M., & Maskin, E. (1995). Credit and efficiency in centralized and decentralized economies. *Review of Economic Studies*, *62*(4), 541–555. https://doi.org/10.2307/2298076

Ely, J. C. (2017). Beeps. *American Economic Review*, *107*(1), 31–53. https://doi.org/10.1257/aer.20150218

Ely, J. C., & Szydlowski, M. (2020). Moving the goalposts. *Journal of Political Economy*, *128*(2), 468–506. https://doi.org/10.1086/704387

Geoffard, P.-Y., & Philipson, T. (1996). Rational epidemics and their public control. *International Economic Review*, *37*(3), 603–624. https://doi.org/10.2307/2527443

Gordon, H. S. (1954). The economic theory of a common property resource: The fishery. *Journal of Political Economy*, *62*(2), 124–142. https://doi.org/10.1086/257497

Hardin, G. (1968). The tragedy of the commons. *Science*, *162*(3859), 1243–1248. https://doi.org/10.1126/science.162.3859.1243

He, Z., & Xiong, W. (2012). Dynamic debt runs. *Review of Financial Studies*, *25*(6), 1799–1843. https://doi.org/10.1093/rfs/hhs004

International Council for the Exploration of the Sea. (n.d.). Our strategy. Retrieved September 4, 2026, from https://www.ices.dk/about-ICES/how-we-work/Pages/Our-strategy.aspx

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, *101*(6), 2590–2615. https://doi.org/10.1257/aer.101.6.2590

Kermack, W. O., & McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London A*, *115*(772), 700–721. https://doi.org/10.1098/rspa.1927.0118

Khantadze, D., Kremer, I., & Skrzypacz, A. (2025). Persuasion with multiple actions. *Journal of Political Economy*, *133*(5), 1497–1526. https://doi.org/10.1086/734125

Kremer, I., Mansour, Y., & Perry, M. (2014). Implementing the "Wisdom of the Crowd". *Journal of Political Economy*, *122*(5), 988–1012. https://doi.org/10.1086/676597

Lehrer, E., & Shaiderman, D. (2026). Markovian persuasion. *Theoretical Economics*, *21*(1), 71–98. https://doi.org/10.3982/TE5372

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.

Renault, J. (2012). The value of repeated games with an informed controller. *Mathematics of Operations Research*, *37*(1), 154–179. https://doi.org/10.1287/moor.1110.0518

Renault, J., Solan, E., & Vieille, N. (2013). Dynamic sender-receiver games. *Journal of Economic Theory*, *148*(2), 502–534. https://doi.org/10.1016/j.jet.2012.07.006

Renault, J., Solan, E., & Vieille, N. (2017). Optimal dynamic information provision. *Games and Economic Behavior*, *104*, 329–349. https://doi.org/10.1016/j.geb.2017.04.010

Scott, A. (1955). The fishery: The objectives of sole ownership. *Journal of Political Economy*, *63*(2), 116–124. https://doi.org/10.1086/257653

Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences*, *39*(10), 1095–1100. https://doi.org/10.1073/pnas.39.10.1095

Solan, E., & Vieille, N. (2015). Stochastic games. *Proceedings of the National Academy of Sciences*, *112*(45), 13743–13746. https://doi.org/10.1073/pnas.1513508112

Stokey, N. L., Lucas, R. E., & Prescott, E. C. (1989). *Recursive methods in economic dynamics*. Harvard University Press.

Sugaya, T., & Wolitzky, A. (2026). Collusion with optimal information disclosure. *Quarterly Journal of Economics*, *141*(3), 2555–2595. https://doi.org/10.1093/qje/qjag020

World Commission on Environment and Development. (1987). *Our common future: Report of the world commission on environment and development*. Oxford University Press. https://sustainabledevelopment.un.org/content/documents/5987our-common-future.pdf

Appendix A Foundations

A.1 Notation

I first introduce some notation used throughout the appendices.

Define $s_\omega \coloneqq p^h_\omega - p^\ell_\omega$, $\omega \in \Omega$. Thus, there are positive (negative) intertemporal externalities if $\forall \omega \in \Omega$, $s_\omega > 0$ ($s_\omega < 0$). State transitions are exogenous if $s_\omega = 0$ $\forall \omega \in \Omega$.

A.1.1 Beliefs

Let $s^a \coloneqq p^a_G - p^a_B$, $a \in \mathcal{A}$ denote the slope of f^a. Since f^a is affine and maps $[0,1]$ into itself, its set of fixed points is non-empty; let p^a_* denote a fixed point of f^a, selected as follows.

$\underline{|\mathbf{s}^a| < \mathbf{1}.}$ If $|s^a| < 1$, then f^a is a contraction with a unique fixed point, that is,

$$f^a(p^a_*) = p^a_* \iff p^a_* = \frac{p^a_B}{1 - s^a} \tag{A.1}$$

For each $a \in \mathcal{A}$, let $f^{a(n)}(p) \coloneqq \underbrace{f^a(f^a \cdots (f^a}_{n \text{ times}}(p)))$, $n \in \mathbb{N}$, and define $f^{a(0)}(p) = p$ and $f^{a(\infty)}(p) = \lim_{n\to\infty} f^{a(n)}(p)$. Starting from prior p, if no information is disclosed for n periods and the Receiver repeatedly takes the same action a, the prior in period-$n+1$ is given by $f^{a(n)}(p)$. Using equation (A.1) and the fact that $f^{a(n)}(p) = f^a(f^{a(n-1)}(p))$, iterating the affine map f^a gives:

$$f^{a(n)}(p) = p^a_*\{1 - (s^a)^n\} + (s^a)^n p \tag{A.2}$$

so that $f^{a(\infty)}(p) = p^a_*$.

$\underline{|\mathbf{s}^a| = \mathbf{1}.}$ There are two cases to consider:

(i) If $s^a = 1$, then $p^a_G = 1$ and $p^a_B = 0$. Thus, f^a is the identity and does not have a unique fixed point. I set $p^a_* = 1/2$ in this case.

(ii) If $s^a = -1$, then $p^a_G = 0$ and $p^a_B = 1$. Thus, $f^{a(n)}(p)$ alternates between p and $1-p$. Equation (A.1) still gives the unique fixed point $p^a_* = 1/2$.

In both cases, equation (A.2) still holds. Additionally, for all values of $|s^a| \in [0,1]$, we have that:

$$p^a_* \gtreqless 1/2 \iff f^a(1/2) \gtreqless 1/2 \tag{A.3}$$

A.1.2 No Information Discounted Expected Payoff

Let $k^h(p) \coloneqq \mathbb{E}_p(u_S(\omega, h)) = p + (1-p)x = x + p(1-x)$ and $k^\ell(p) \coloneqq \mathbb{E}_p(u_S(\omega, \ell)) = py$ be the Sender's expected payoff from the high and low actions in the stage game, respectively, and let κ^a denote the slope of $k^a(\cdot)$. Starting from prior p, when the Sender discloses no information for $n \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ periods and Receiver repeatedly takes the same action a in those periods, the Sender's *no information discounted expected payoff* is denoted by $U^a(p, n)$. Define $U^a(p, 0) = 0$. Then:

$$U^a(p,n) = \begin{cases} 0, & \text{if } n = 0 \\ \sum_{i=0}^{n-1} \delta^i k^a(f^{a(i)}(p)), & \text{if } n > 0 \end{cases}$$

Using (A.2), the above can be re-written as follows:

$$U^a(p,n) = \begin{cases} 0, & \text{if } n = 0 \\ k^a(p^a_*)\frac{1-\delta^n}{1-\delta} + (p - p^a_*)\kappa^a \frac{1-(\delta s^a)^n}{1-\delta s^a}, & \text{if } n \in \mathbb{N} \\ \frac{k^a(p^a_*)}{1-\delta} + \frac{(p-p^a_*)\kappa^a}{1-\delta s^a}, & \text{if } n = \infty \end{cases} \tag{A.4}$$

If no information is disclosed and the same action a is taken repeatedly for $n \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ periods, let $m^a_n \coloneqq (1-\delta)\frac{\partial U^a(p,n)}{\partial p}$ denote the *marginal effect of a higher current prior on the Sender's payoff over those n periods*:

$$m^a_n = \begin{cases} 0, & \text{if } n = 0 \\ (1-\delta)\kappa^a \sum_{i=0}^{n-1} (\delta s^a)^i, & \text{if } n \in \mathbb{N} \\ \frac{(1-\delta)\kappa^a}{1-\delta s^a}, & \text{if } n = \infty \end{cases}$$

Finally, I adopt the convention that $\delta^\infty h \coloneqq \lim_{n\to\infty} \delta^n h = 0$ for any bounded h. Since $f^{a(n)}(p) \in [0,1]$ for every n, this sets $\delta^n f^{a(n)}(p) = 0$ whenever $n = \infty$.[34]

[34] When $s^a = -1$, the limit defining $f^{a(\infty)}(p)$ does not exist unless $p = 1/2$. This is immaterial below, because $f^{a(\infty)}$ enters only through products of the form $\delta^n f^{a(n)}(\cdot)$, which is set to zero at $n = \infty$.

A.2 Properties of the Value Function $V(\cdot)$

A.2.1 Uniqueness

Lemma 3. V is the unique bounded solution of equation (3.4).

Proof. For bounded W on $[0,1]$, let TW be the right-hand side of equation (3.4) with W in place of V. Then $V = TV$, and $|W - W'| \leq c$ implies $|TW - TW'| \leq \delta c$. So T is a contraction: V is the unique bounded solution of Bellman equation (3.4), and $T^n W \to V$ uniformly for every bounded W. ■

A.2.2 Concavity

Lemma 4. V is concave over its domain.

Proof. The proof of the result is in Renault, Solan, and Vieille (2017), reproduced here. To show concavity, we must show that $V(p) \geq \lambda V(p') + (1-\lambda)V(p'')$, where $p = \lambda p' + (1-\lambda)p''$ for $\lambda \in [0,1]$. At belief p, the Sender can guarantee an expected payoff of at least $\lambda V(p') + (1-\lambda)V(p'')$ by inducing beliefs of p' and p'' with probabilities λ and $1-\lambda$ respectively, and then following the optimal policy at the resulting belief. The result then follows since by definition, $V(p)$ denotes the maximum payoff to the Sender at the prior p. ■

Corollary 3. It is without loss to restrict attention to distributions $\mu \in \Delta(\Delta(\Omega))$ whose support has at most 2 elements.

Proof. Since $V(\cdot)$ is concave, if two or more posteriors induce the same action from the Receiver, the Sender weakly gains by inducing a single (combined) posterior that induces the said action. Since the Receiver's action space is binary, it is without loss for the Sender to induce at most two posteriors at any prior p. ■

A.2.3 Continuity

For $q \in [0,1]$, let $w^a(q) \coloneqq (1-\delta)k^a(q) + \delta V(f^a(q))$ denote the Sender's payoff from inducing the posterior q together with action $a \in \{h, \ell\}$ and then continuing optimally. Since V is concave by

Lemma 4 above, and k^a and f^a are affine, w^h and w^ℓ are concave on $[0,1]$, and hence continuous on $(0,1)$. Define:

$$w(q) \coloneqq \begin{cases} w^\ell(q), & \text{if } q < 1/2 \\ w^h(q), & \text{if } q > 1/2 \\ \max\{w^\ell(1/2), w^h(1/2)\}, & \text{if } q = 1/2 \end{cases}$$

Claim 1. *V is 2-Lipschitz. Hence V is continuous on $[0,1]$, and its one-sided derivatives on $(0,1)$ lie in $[-2,2]$.*

Proof. By Lemma 3, it suffices to show: if W is a function on $[0,1]$ that takes values in $[0,1]$ and is 2-Lipschitz, then TW has the same two properties. Fix such a W. Let w_W be w with W in place of V, so that $TW(p) = \sup_{\mathbb{E}_\mu(q)=p} \mathbb{E}_\mu[w_W(q)]$. Then w_W and TW take values in $[0,1]$. Since $y, |s^\ell| \leq 1$, for $q \geq 1/2$, $w_W(q) - w_W(0) \leq 1 \leq 2q$ and for $q < 1/2$, $w_W(q) - w_W(0) \leq (1-\delta)q + 2\delta q \leq 2q$. So $\mathbb{E}_\mu[w_W] \leq w_W(0) + 2p$ whenever $\mathbb{E}_\mu(q) = p$, and $TW(0) = w_W(0)$; hence $TW(p) - TW(0) \leq 2p$. Symmetrically, $TW(p) - TW(1) \leq 2(1-p)$. TW is concave (as in Lemma 4), so for $p < p'$,

$$-2 \leq \frac{TW(1) - TW(p)}{1-p} \leq \frac{TW(p') - TW(p)}{p' - p} \leq \frac{TW(p') - TW(0)}{p'} \leq 2. \qquad \blacksquare$$

A.2.4 Communication Interval

By Bellman equation (3.5),

$$\begin{gathered} V(p) = \max_{\alpha, q^h} \{\alpha w^h(q^h) + (1-\alpha) w^\ell(q^\ell)\} \\ \text{subject to } p = \alpha q^h + (1-\alpha) q^\ell, q^h \in [1/2, 1], q^\ell \in [0, 1/2] \end{gathered} \tag{A.5}$$

so that $V \geq w$, and $V(p) = w(p)$ if and only if disclosing no information is optimal at p. The maximum in (A.5) is attained because w is upper semi-continuous on $[0,1]$.[35]

Lemma 5. *Let $I \coloneqq \{p \in [0,1] : V(p) > w(p)\}$ be the set of priors at which the Sender provides information. Either $I = \emptyset$, or $I = (\underline{p}, \overline{p})$ for some $\underline{p} \leq 1/2 \leq \overline{p}$ with $\underline{p} < \overline{p}$, and V is affine on $[\underline{p}, \overline{p}]$. Moreover, if $\overline{p} = 1/2$ then $V(1/2) = w^h(1/2) > w^\ell(1/2)$, and if $\underline{p} = 1/2$ then $V(1/2) = w^\ell(1/2) > w^h(1/2)$.*

[35] Each w^a is continuous, since V is continuous (Claim 1). As $w(1/2) = \max\{w^\ell(1/2), w^h(1/2)\}$, w is upper semi-continuous.

Proof. At $p \in \{0,1\}$ no information can be disclosed, so $0,1 \notin I$. On $(0,1)$, V is continuous and w is upper semi-continuous, so I is open and is a disjoint union of open intervals. Let (c,d) be one of them, and take $p \in (c,d)$. Let (α, q^h) attain the maximum in (A.5). By concavity of V,

$$V(p) = \alpha w^h(q^h) + (1-\alpha)w^\ell(q^\ell) \leq \alpha V(q^h) + (1-\alpha)V(q^\ell) \leq V(p),$$

so both inequalities hold with equality. Hence the posteriors used satisfy $V = w$, i.e. they lie outside I, so $q^\ell \leq c < p < d \leq q^h$, and V is affine on $[q^\ell, q^h] \supseteq [c,d]$. Also $V(c) = w(c)$ and $V(d) = w(d)$, since $c, d \notin I$.

Suppose $d < 1/2$. Then $V(c) = w^\ell(c)$, $V(d) = w^\ell(d)$, and concavity of w^ℓ gives $w^\ell \geq V$ on $[c,d]$, contradicting $(c,d) \subseteq I$. The same contradiction arises if $d = 1/2$ and $V(1/2) = w^\ell(1/2)$. So if $d \leq 1/2$, then $d = 1/2$ and $V(1/2) = w^h(1/2) > w^\ell(1/2)$. Symmetrically, if $c \geq 1/2$, then $c = 1/2$ and $V(1/2) = w^\ell(1/2) > w^h(1/2)$. Every other interval satisfies $c < 1/2 < d$. Since the intervals are disjoint, and the two conditions on $V(1/2)$ are mutually exclusive, I contains at most one of them. ■

Thus, an optimal policy is to disclose nothing outside I, and to split every prior in I between $\underline{p}$ and $\overline{p}$; by (A.5), this policy attains V at every prior. In particular, if $\underline{p} = 0$ and $\overline{p} = 1/2$ it is the greedy policy, if $\underline{p} = 1/2$ and $\overline{p} = 1$ it is the judicious policy, and if $\underline{p} = 0$ and $\overline{p} = 1$ it is full disclosure. If $\underline{p} \in (0, 1/2)$ and $\overline{p} = 1/2$ it is the truncated greedy policy with $\underline{q} = \underline{p}$, and if $\underline{p} = 1/2$ and $\overline{p} \in (1/2, 1)$ it is the truncated judicious policy with $\overline{q} = \overline{p}$. So in Table 1, $\underline{p} = q^\ell_\pi$ and $\overline{p} = q^h_\pi$.

Lemma 6. If $1/2 \in I$, then full disclosure is optimal.

Proof. Suppose $1/2 \in I$. By Lemma 5, $I = (\underline{p}, \overline{p})$ with $\underline{p} < 1/2 < \overline{p}$, V is affine on $[\underline{p}, \overline{p}]$, and $V = w$ on $[0,1] \setminus I$. Since V, w^h and w^ℓ are concave (Lemma 4), they have non-increasing left and right derivatives on $(0,1)$, with $V'_- \geq V'_+$. For a closed interval $[c,d] \subset (0,1)$, possibly a single point, I refer to $V'_-(c) - V'_+(d) \geq 0$ as the *slope drop* of V across $[c,d]$. I refer to $[c,d]$ as *short* if $d - c < \overline{p} - \underline{p}$.

I first show that, for every short $[c,d] \subset (0,1)$, there exists a short $[c',d'] \subseteq f^h([0,1]) \cup f^\ell([0,1])$ such that

$$V'_-(c) - V'_+(d) \leq \delta \max_{a \in \mathcal{A}} |s^a| \big[V'_-(c') - V'_+(d')\big] \tag{A.6}$$

A short interval cannot meet both $[0, \underline{p}]$ and $[\overline{p}, 1]$. If it meets neither, it lies in I, where V is affine, so its slope drop is zero and equation (A.6) holds trivially. Suppose it meets $[0, \underline{p}]$; the other case is symmetric, with h in place of ℓ. Since V'_+ is constant on $[\underline{p}, \overline{p})$, replacing d by $\min\{d, \underline{p}\}$ does not change the slope drop, so we may assume $[c, d] \subseteq (0, \underline{p}]$. On $[0, \underline{p}]$, $V = w = w^\ell$, because $\underline{p} < 1/2$. So $V'_-(c) = (w^\ell)'_-(c)$, and $V'_+(d) \geq (w^\ell)'_+(d)$. The latter holds with equality if $d < \underline{p}$, and as an inequality if $d = \underline{p}$, because $V \geq w^\ell$ on $[\underline{p}, 1/2)$ with equality at $\underline{p}$. Hence

$$V'_-(c) - V'_+(d) \leq (w^\ell)'_-(c) - (w^\ell)'_+(d) = \delta |s^\ell| [V'_-(c') - V'_+(d')], \qquad [c', d'] \coloneqq f^\ell([c, d])$$

The equality holds because $w^\ell(p) = (1-\delta)k^\ell(p) + \delta V(f^\ell(p))$ with k^ℓ and f^ℓ affine, so $(w^\ell)'_\pm(p) = (1-\delta)\kappa^\ell + \delta s^\ell V'_\pm(f^\ell(p))$. (If $s^\ell < 0$, the two sides of V' are swapped, because f^ℓ reverses direction.) Finally, $d' - c' = |s^\ell|(d - c)$, so $[c', d']$ is short.

Let ξ be the supremum of the slope drop of V over short intervals in $(0, 1)$; by Claim 1, $\xi \leq 4$.[36] If $s^a \neq 0$, f^a maps $(0, 1)$ into $(0, 1)$, so equation (A.6) gives $\xi \leq \delta\xi$; if $s^a = 0$, w^a is affine and the slope drop in equation (A.6) is zero. Hence, $\xi = 0$.

Now split any $[c, d] \subset (0, 1)$ into finitely many consecutive short intervals. Since $V'_- \geq V'_+$ at each point where two consecutive intervals meet, the slope drop across $[c, d]$ is at most the sum of the slope drops across the pieces, which is zero. So $V'_-(c) = V'_+(d)$ for all $c \leq d$ in $(0, 1)$, and V is affine on $(0, 1)$.

It remains to show that $I = (0, 1)$. Suppose $\underline{p} > 0$. Since V is affine on $(0, 1)$ and continuous on $[0, 1]$, it is affine on $[0, 1]$, so w^ℓ is affine on $[0, 1]$. It equals V on $(0, \underline{p}]$, so it equals V on $(0, 1)$. This contradicts $V > w = w^\ell$ on $(\underline{p}, 1/2) \subseteq I$. Hence $\underline{p} = 0$, and symmetrically $\overline{p} = 1$. So $I = (0, 1)$, and the optimal policy splits every prior between the posteriors 0 and 1: full disclosure is optimal. ■

The contrapositive of the previous lemma gives the following:

Corollary 4. If full disclosure is not optimal, then $1/2 \notin I$, i.e. $V(1/2) = \max\{w^h(1/2), w^\ell(1/2)\}$.

To summarise, we have shown that: (i) the value function $V(\cdot)$ is unique, concave and continuous, (ii) it is without loss to induce at most one posterior for each action, (iii) the set of priors

[36] For general $(x, y) \in \mathbb{R}^2$, the same argument, with stage payoffs in $[\underline{u}, \overline{u}]$ where $\underline{u} \coloneqq \min\{0, x, y\}$ and $\overline{u} \coloneqq \max\{1, x, y\}$, shows that V is Lipschitz with constant $\max\{|1 - x|, |y|, 2(\overline{u} - \underline{u})\}$. Lemma 6 only uses that this constant is finite.

at which the Sender communicates information is an interval, where at every prior she splits the belief between the interval's endpoints, and she discloses nothing outside it, and (iv) if $1/2$ lies in the interior of this interval, then full disclosure is optimal.

A.3 Proof of Lemma 1

Consider a stationary policy $(\alpha(p), q^h(p))$ of the Sender and let $\phi(\cdot)$ denote the Sender's payoff from using this policy.

($\implies$): If $(\alpha(p), q^h(p))$ is optimal then $\phi(p) = V(p)$ and by Lemma 4, $\phi(\cdot)$ must be concave. Additionally, $(\alpha(p), q^h(p))$ must give a higher payoff than all deviations, including the one-shot deviation to providing no information in the current period, and then returning to the prescribed policy from the next period onward. This is precisely the inequality in (3.6).

($\impliedby$): By definition, the payoff from any policy is at most the maximum, so that:

$$\phi(p) \leq \max_{\tilde{\alpha}, \tilde{q}^h} \{(1-\delta)[\tilde{\alpha}(\tilde{q}^h + (1-\tilde{q}^h)x) + (1-\tilde{\alpha})\tilde{q}^\ell y] + \delta\tilde{\alpha}\phi(f^h(\tilde{q}^h)) + \delta(1-\tilde{\alpha})\phi(f^\ell(\tilde{q}^\ell))\}$$

Let $\phi(\cdot)$ be concave in p and (3.6) hold. Consider any arbitrary Bayes plausible experiment $(\tilde{\alpha}(p), \tilde{q}^h(p))$. Then:

$$\begin{aligned}
&\tilde{\alpha}\left[(1-\delta)(\tilde{q}^h + (1-\tilde{q}^h)x) + \delta\phi(f^h(\tilde{q}^h))\right] + (1-\tilde{\alpha})\left[(1-\delta)\tilde{q}^\ell y + \delta\phi(f^\ell(\tilde{q}^\ell))\right] \\
&\quad \leq \tilde{\alpha}\phi(\tilde{q}^h) + (1-\tilde{\alpha})\phi(\tilde{q}^\ell) && \text{(by (3.6))} \\
&\quad \leq \phi\Big(\tilde{\alpha}\tilde{q}^h + (1-\tilde{\alpha})\tilde{q}^\ell\Big) && \text{(concavity)} \\
&\quad = \phi(p) && \text{(Bayes plausibility)}
\end{aligned}$$

Thus, $\phi(p)$ solves the Bellman equation (3.4). Since equation (3.4) has a unique bounded solution by Lemma 3, $\phi = V$ and $(\alpha(p), q^h(p))$ is optimal. ■

Appendix B The Greedy Policy

Recall that the greedy policy maximises the probability with which the high action is taken by the Receiver each period, i.e.:

$$(\alpha_g(p), q_g^h(p)) = \begin{cases} (1, p), & p \geq 1/2 \\ (2p, 1/2), & p < 1/2 \end{cases}$$

Let $\gamma : [0,1] \to [0,1]$ denote the Sender's long-run payoff from using the greedy policy:

$$\gamma(p) = \begin{cases} (1-\delta)k^h(p) + \delta\gamma(f^h(p)), & p \geq 1/2 \\ 2p\gamma(1/2) + (1-2p)\gamma(0), & p < 1/2 \end{cases} \tag{B.1}$$

When $p = 0$, the Receiver is certain that the state is bad, so he chooses the low action. This implies that the next period's prior is $f^\ell(0) = p_B^\ell$, so that $\gamma(0) = \delta\gamma(p_B^\ell)$. Observe that, for $p < 1/2$, γ is affine, with a slope of $\gamma' := \gamma'_-(1/2) = 2[\gamma(1/2) - \gamma(0)]$.

For given values of transition probabilities and preference parameters x and y, to show that the greedy policy is optimal, we need to verify that both the conditions of Lemma 1 are satisfied. That is, we must show that γ is concave in p and that the Sender does not strictly gain from a one-shot deviation to no disclosure. I next analyse these in turn.

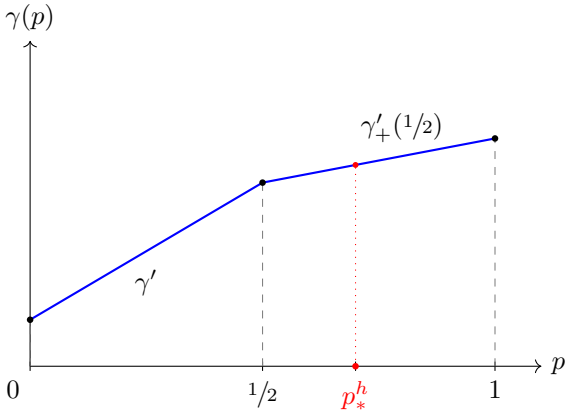
B.1 Concavity

First, I analyse the payoff from the greedy policy. I start by showing the following:

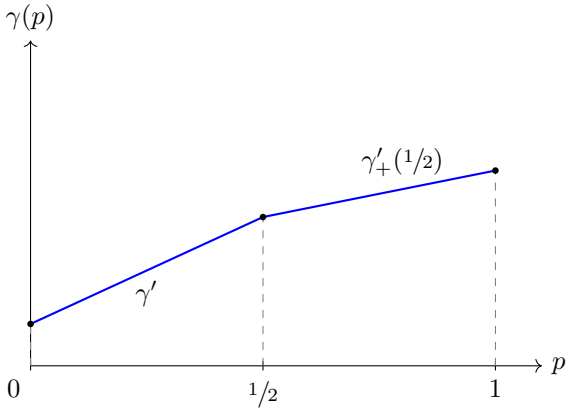
Lemma 7. The payoff from the greedy policy is concave if and only if $\gamma' \geq m_\infty^h$.

Proof. The proof is divided into 4 cases (see Figure 10), depending on the range of values of $f^h([1/2, 1])$.[37] In each of these, I argue that γ is piecewise affine. This means that for γ to be concave, the left-hand derivative must exceed the right-hand derivative at the kinks, i.e. $\gamma'_-(p) \geq \gamma'_+(p)$ at kink p. It turns out that this is equivalent to showing that $\gamma' \geq m_\infty^h$.

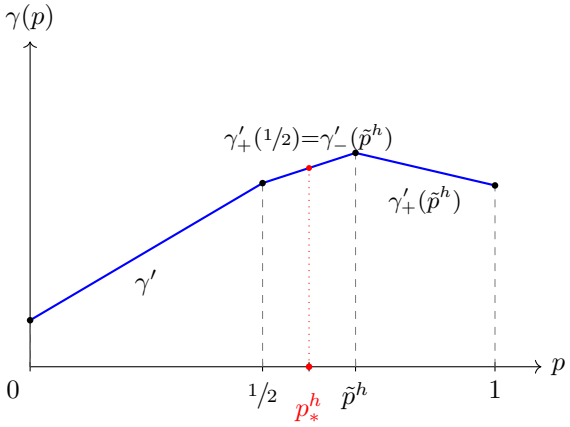
[37] Given the period-t posterior q_t, the prior in period $t+1$ lies either in the interval $f^h([1/2, 1])$ or in the interval $f^\ell([0, 1/2])$. The linearity of f^h and f^ℓ, and the fact that p_G^h, p_B^h, p_G^ℓ, p_B^ℓ are probabilities implies that $f^h([1/2, 1])$ and $f^\ell([0, 1/2])$ are interval subsets of $[0, 1]$.



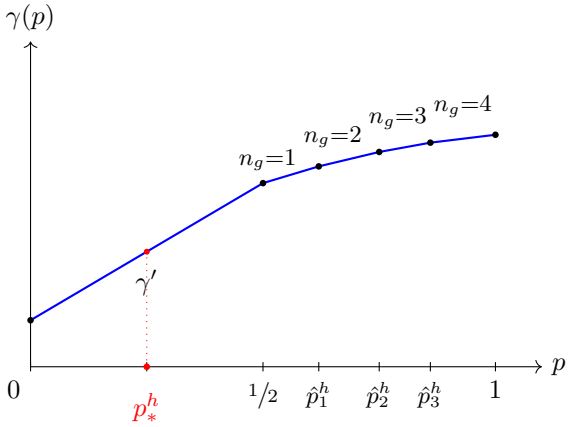
(a) Absorbing case: $f^h([1/2, 1]) \subseteq [1/2, 1]$



(b) Reverting case: $f^h([1/2, 1]) \subseteq [0, 1/2)$



(c) Partially absorbing case: $f^h(1/2) \geq 1/2 > f^h(1)$



(d) Cascading case: $f^h(1/2) < 1/2 \leq f^h(1)$

Figure 10: The Sender's payoff $\gamma(\cdot)$ under the greedy policy in each of the four cases of Section B.1; p^h_* is the fixed point of $f^h(\cdot)$, drawn for $|s^h| < 1$

Case 1: $\underline{f^h([1/2,1]) \subseteq [1/2,1].}$ The interval $[1/2,1]$ is absorbing: because no information is disclosed, once the belief is greater than $1/2$, it stays at or above $1/2$ forever (see panel (a) of Figure 10). For $p \geq 1/2$, since the static payoff $k^h(\cdot)$ and the future belief $f^h(\cdot)$ are affine in p, $\gamma(\cdot)$ is affine on $[1/2,1]$. Differentiating the first line of equation (B.1) with respect to p:

$$\gamma'_+(1/2) = (1-\delta)\kappa^h + \delta s^h \gamma'_+(1/2) \iff \gamma'_+(1/2) = m^h_\infty \tag{B.2}$$

It follows that γ is concave if and only if $\gamma' \geq m^h_\infty = \gamma'_+(1/2)$.

Case 2: $\underline{f^h([1/2,1]) \subseteq [0,1/2).}$ When $p \geq 1/2$, the Sender discloses no information and $f^h(p) < 1/2$, so $\gamma'_+(1/2) = m^h_1 + \delta s^h \gamma'$. Thus, γ is concave if and only if $\gamma' \geq \gamma'_+(1/2)$, or equivalently, $\gamma' \geq m^h_\infty$ (see panel (b) of Figure 10).

Case 3: $\underline{f^h(1/2) \geq 1/2 > f^h(1).}$ Since f^h is affine, $f^h(1/2) > f^h(1)$ implies that f^h is decreasing, i.e. $s^h < 0$. And since $f^h([1/2,1])$ is an interval containing $1/2$, by the intermediate value theorem, there exists a unique $\tilde{p}^h \in [1/2,1]$ such that $f^h(\tilde{p}^h) = 1/2$. This means that for $p \in [1/2, \tilde{p}^h]$, $f^h(p) \in [1/2, f^h(1/2)]$. Since $f^h(1/2) \leq \tilde{p}^h$, $[1/2, \tilde{p}^h]$ maps into itself. If $\tilde{p}^h > 1/2$, then $\gamma'_+(1/2)$ is given by equation (B.2) and $\gamma'_+(1/2) = m^h_\infty = \gamma'_-(\tilde{p}^h)$ (see panel (c) of Figure 10). If $\tilde{p}^h = 1/2$, then $\gamma'_+(1/2) = \gamma'_+(\tilde{p}^h)$, which is described next.

By the definition of $\tilde{p}^h$, for $p > \tilde{p}^h$, $f^h(p) < 1/2$. Hence $\gamma'_+(\tilde{p}^h)$ is obtained by applying the first-period recursion to the affine branch below $1/2$, giving $\gamma'_+(\tilde{p}^h) = m^h_1 + \delta s^h \gamma'$. It follows that γ is concave if and only if $\gamma' \geq \gamma'_+(1/2) = \gamma'_-(\tilde{p}^h) \geq \gamma'_+(\tilde{p}^h)$. These inequalities reduce to $\gamma' \geq m^h_\infty$, using the fact that $s^h < 0$.

Case 4: $\underline{f^h(1/2) < 1/2 \leq f^h(1).}$ Since f^h is affine, $f^h(1/2) < f^h(1)$ implies that f^h is increasing, i.e. $s^h > 0$. As in case 3, by the intermediate value theorem, there exists a unique $\hat{p}^h_1 \in [1/2,1]$ such that $f^h(\hat{p}^h_1) = 1/2$. Thus, for $p \in [1/2, \hat{p}^h_1)$, $f^h(p) < 1/2$.

Suppose $p \geq \hat{p}^h_1$ so that $f^h(p) \geq 1/2$. Again, the Sender discloses no information, so the new belief is $f^h(f^h(p))$, which could be less than or greater than $1/2$. Define $\hat{p}^h_2 \in [\hat{p}^h_1, 1]$ to be the point such that $f^h(\hat{p}^h_2) = \hat{p}^h_1$, so that for $p \geq \hat{p}^h_2$, $f^h(f^h(p)) \geq 1/2$. Continuing with the same logic, we can

define a sequence of thresholds $1/2 =: \hat{p}_0^h < \hat{p}_1^h < \hat{p}_2^h < \cdots < \hat{p}_{\eta-1}^h < 1$, where $f^h(\hat{p}_n^h) = \hat{p}_{n-1}^h$, and set $\hat{p}_\eta^h := 1$, where from an initial belief of $p = 1$, $\eta \in \mathbb{N}$ is the number of periods in which no information is disclosed before the belief falls below $1/2$. Since f^h is a contraction,[38] this sequence is finite and since $f^h(1/2) < 1/2$, beliefs always slide towards p_*^h, which is less than $1/2$ by equation (A.3). For example, in panel (d) of Figure 10, $\eta = 4$.

It follows that $[1/2, 1] = [\hat{p}_0^h, \hat{p}_\eta^h]$ is the union of disjoint intervals: $[\hat{p}_0^h, \hat{p}_1^h) \cup [\hat{p}_1^h, \hat{p}_2^h) \cup \cdots \cup [\hat{p}_{\eta-1}^h, \hat{p}_\eta^h]$. For $p \in [\hat{p}_{n-1}^h, \hat{p}_n^h)$, $n = \min\{i : f^{h(i)}(p) < 1/2\}$, where $n \in \{1, 2, \ldots, \eta\}$. Hence, for $n \in \{0, 1, \ldots, \eta - 1\}$, $\gamma'_+(\hat{p}_n^h) = m_{n+1}^h + (\delta s^h)^{n+1}\gamma' = \gamma'_-(\hat{p}_{n+1}^h)$.

Thus, $\gamma(\cdot)$ is concave if and only if $\gamma' \geq \gamma'_+(\hat{p}_0^h) = \gamma'_-(\hat{p}_1^h) \geq \gamma'_+(\hat{p}_1^h) = \gamma'_-(\hat{p}_2^h) \geq \cdots \geq \gamma'_-(\hat{p}_\eta^h)$. These inequalities reduce to $\gamma' \geq m_\infty^h$, using the fact that $s^h > 0$. This completes the proof. ■

The previous lemma expresses the concavity of the greedy policy as a condition on its slope. This can be further reduced to get a simple condition on the preference parameter x. The following claims provide some useful steps:

Claim 2. Consider any prior $p \in [0, 1]$. Let $n_g(p) \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ be the number of periods in which no information is disclosed before the belief falls below $1/2$, i.e. $n_g(p) = \min\{j : f^{h(j)}(p) < 1/2\}$. Using equation (A.4), for $p \in (0, 1]$, the long-run payoff from the greedy policy (B.1) can be re-written as:

$$\gamma(p) = (1 - \delta)U^h(p, n_g(p)) + \delta^{n_g(p)} \left\{ 2f^{h(n_g(p))}(p)\gamma(1/2) + (1 - 2f^{h(n_g(p))}(p))\gamma(0) \right\} \tag{B.3}$$

where $\gamma(0) = \delta\gamma(p_B^\ell)$.

Proof. If $n_g(p) = 0$, equation (B.3) reduces to the same expression for $\gamma(\cdot)$ when $p < 1/2$ as defined in equation (B.1). If $n_g(p) \in \mathbb{N}$, then by assumption, $f^{h(n_g(p)-1)}(p) \geq 1/2 > f^{h(n_g(p))}(p)$. Further, the prior in period $n_g(p) + 1$ is $f^{h(n_g(p))}(p)$, so that after $n_g(p)$ periods, information is disclosed by splitting beliefs between $1/2$ and 0. If $n_g(p) = \infty$, then information is never disclosed, so the second and third terms above equal 0 in the limit. ■

In the notation of Section 6.1, $n_g(f^h(1/2)) = n_g^h$ and $n_g(p_B^\ell) = n_g^\ell$.[39] Let $n_g(1/2) := n_g^{1/2}$. Then:

[38] This is because $s^h = 1$ forces $p_G^h + p_B^h = 1$ and hence $f^h(1/2) = 1/2$, contradicting $f^h(1/2) < 1/2$. Thus, $s^h < 1$ and f^h is a contraction.

[39] Recall from Section 6.1 that n_π^a measures the number of periods of the quiet spell that follows a disclosure that

Claim 3. The value of γ' is a function of n_g^ℓ and $n_g^{1/2}$, where $n_g^\ell \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ and $n_g^{1/2} \in \{1, \infty\}$:

$$\gamma' = 2(1-\delta)\frac{(1-\delta^{n_g^\ell+1})U^h(1/2, n_g^{1/2}) - \delta(1-\delta^{n_g^{1/2}})U^h(p_B^\ell, n_g^\ell)}{\beta_g(n_g^\ell, n_g^{1/2})} \tag{B.4}$$

where $\beta_g(n_g^\ell, n_g^{1/2}) \coloneqq (1-\delta^{n_g^\ell+1})(1-2\delta^{n_g^{1/2}} f^{h(n_g^{1/2})}(1/2)) + 2\delta^{n_g^\ell+1}(1-\delta^{n_g^{1/2}}) f^{h(n_g^\ell)}(p_B^\ell) > 0$.

Proof. Using equation (B.3), the value of γ' above is derived as the solution of a system of two equations in two unknowns, $\gamma(1/2)$ and $\gamma(0)$. From the proof of Lemma 7, observe that $n_g^{1/2} = \infty$ in Cases 1 and 3, while $n_g^{1/2} = 1$ in Cases 2 and 4. The value of n_g^ℓ, on the other hand, depends on the "branch" of γ on which p_B^ℓ lies. Finally, the denominator $\beta_g(n_g^\ell, n_g^{1/2})$ is the sum of a strictly positive and a non-negative term. ■

Lemma 8. The payoff from the greedy policy is concave if and only if $x \geq \bar{x}$.

Proof. By Lemma 7, it is sufficient to show that $\gamma' \geq m_\infty^h \iff x \geq \bar{x}$. For each $(n_g^\ell, n_g^{1/2})$ pair, $\gamma' - m_\infty^h = C(n_g^\ell, n_g^{1/2}) \cdot (x - \bar{x})$, where $C(n_g^\ell, n_g^{1/2}) = \frac{2(1-\delta)(1-\delta^{n_g^{1/2}})(1+\delta p_B^\ell - \delta p_G^h)}{\beta_g(n_g^\ell, n_g^{1/2})(1-\delta s^h)} > 0$. Therefore, the payoff from the greedy policy is concave if and only if $x \geq \bar{x}$. ■

From the above, it follows that when $x = \bar{x}$, $\gamma' = m_\infty^h$.

B.2 One-Shot Deviation Condition

I now turn to the one-shot deviation condition. Since the Sender already discloses no information when $p \geq 1/2$, we only need to verify the condition for $p < 1/2$. Let $D(p; x, y) \coloneqq \gamma(p) - \{(1-\delta)k^\ell(p) + \delta\gamma(f^\ell(p))\}$ $\forall p \in [0, 1/2)$.[40,41] The following lemmas provide some useful simplifications:

Lemma 9. For a given x, if the greedy policy is optimal for some $\hat{y}$, then it is optimal for all $y \leq \hat{y}$.

Proof. Fix a value of $x \in [0,1]$ and suppose that the greedy policy is optimal when $y = \hat{y}$. From equation (B.3), it follows that the value of y does not impact γ. Further, by assumption and by Lemma 1, $D(p; x, \hat{y}) \geq 0$ $\forall p \in [0, 1/2)$. The result then follows from the fact that $D(p; x, y)$ is decreasing in y. ■

induces action a. Recall further that under the greedy policy, the prior following the high and low actions is $f^h(1/2)$ and p_B^ℓ, respectively.

[40] Note that D depends on x and y via functions γ and k^ℓ, respectively.

[41] By continuity of γ, this also gives the condition at $p = 1/2$.

Lemma 10. For a given (x, y), the greedy policy is optimal if and only if γ is concave in p and $D'_+(0; x, y) \geq 0$.

Proof. Fix $(x, y) \in [0, 1]^2$. If the greedy policy is optimal, by Lemma 1, γ is concave in p and $D(p; x, y) \geq 0 \ \forall p \in [0, 1/2)$. Since $D(0; x, y) = \gamma(0) - \delta\gamma(f^\ell(0)) = 0$, $p = 0$ minimises $D(\cdot; x, y)$ over this interval. Hence, $D'_+(0; x, y) \geq 0$.

Conversely, suppose γ is concave in p and $D'_+(0; x, y) \geq 0$. This implies that $\gamma(f^\ell(\cdot))$ is concave as it is the composition of a concave and an affine function. Further, γ is affine on $[0, 1/2)$, so $D(\cdot; x, y)$ is convex on this interval. Since $D(0; x, y) = 0$ and $D'_+(0; x, y) \geq 0$, it follows that $D(p; x, y) \geq pD'_+(0; x, y) \geq 0$ for every $p \in [0, 1/2)$, so that by Lemma 1, the greedy policy is optimal. ■

Claim 4. It suffices to evaluate $D'_+(0; x, y)$ at the extreme values $x = \bar{x}$ and $x = 1$.

Proof. For fixed transition probabilities, y and p, by equation (B.3), $\gamma(p)$ is affine in x, so that $D'_+(0; x, y)$ is also affine in x. By Lemma 8, γ is concave if and only if $x \geq \bar{x}$, so the relevant extreme values are $\bar{x}$ and 1. ■

As with γ', the value of $D'_+(0; x, y)$ depends on the branch of γ to which p_B^ℓ belongs.[42] For $n_g^\ell \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$, $D'_+(0; x, y)$ is given by:

$$D'_+(0; x, y) = \gamma' - (1 - \delta)y - \delta s^\ell \left\{ m^h_{n_g^\ell} + (\delta s^h)^{n_g^\ell} \gamma' \right\}$$

so that:

$$D'_+(0; \bar{x}, y) = (1 - \delta)(\bar{y} - y) \geq 0 \iff y \leq \bar{y} \tag{B.5}$$

and

$$D'_+(0; 1, y) = \gamma'(1 - \delta s^\ell (\delta s^h)^{n_g^\ell}) - (1 - \delta)y. \tag{B.6}$$

where γ' is as defined in equation (B.4). Since $D'_+(0; \cdot, \bar{y})$ is affine in x and vanishes at $x = \bar{x}$ by

[42] If p_B^ℓ is a kink of γ, read n_g^ℓ in the display below and in equations (B.6)–(B.8) as $n_g(f^\ell(\varepsilon))$ for small $\varepsilon > 0$; equation (B.4) holds with either value, since γ is continuous.

equation (B.5), we have that:

$$D'_+(0;1,\bar{y}) = (1-\bar{x}) \cdot \frac{\partial D'_+(0;x,\bar{y})}{\partial x} \tag{B.7}$$

B.3 Optimality

Summarising the previous two subsections, by Lemma 8, the payoff from the greedy policy is concave if and only if $x \geq \bar{x}$. By Lemma 9, Lemma 10 and Claim 4, if $D'_+(0;\bar{x},\hat{y}) \geq 0$ and $D'_+(0;1,\hat{y}) \geq 0$, then the greedy policy is optimal for all $(x,y) \in [0,1]^2$ with $x \geq \bar{x}$ and $y \leq \hat{y}$.

Proposition 6. Suppose $s_G \geq 0$. Then the greedy policy is optimal if and only if $x \geq \bar{x}$.

Proof. If $s_G \geq 0$, then $\bar{y} \geq 1$. By Lemma 8, the greedy policy is not optimal if $x < \bar{x}$. By the discussion above with $\hat{y} = 1$, it is optimal for all $x \geq \bar{x}$ and $y \in [0,1]$ if $D'_+(0;\bar{x},1) \geq 0$ and $D'_+(0;1,1) \geq 0$. By equation (B.5), $D'_+(0;\bar{x},1) = (1-\delta)(\bar{y}-1) \geq 0$. It only remains to show that $D'_+(0;1,1) \geq 0$.

Suppose first that $n_g^{1/2} = \infty$. From the proof of Lemma 7, this implies $n_g^\ell \in \{0,1,\infty\}$. Using equations (B.4) and (B.6) to evaluate $D'_+(0;1,1)$, we get:

$$D'_+(0;1,1) = \frac{1-\delta}{\beta_g(n_g^\ell,\infty)}\left\{1 + \delta^{n_g^\ell+1} - 2\delta^{n_g^\ell+1} f^{h(n_g^\ell)}(p_G^\ell)\right\}$$

and the term in the curly brackets is strictly positive $\forall n_g^\ell \in \{0,1,\infty\}$.

Suppose instead that $n_g^{1/2} = 1$, i.e. $f^h(1/2) < 1/2$. From the proof of Lemma 7, this implies that $n_g^\ell \in \mathbb{Z}_{\geq 0}$:

(a) If $n_g^\ell = 0$, then $D'_+(0;1,1) = \frac{(1-\delta)^2}{\beta_g(0,1)}\left\{1 - \delta p_G^\ell + \delta p_B^h + \delta s_G\right\} > 0$ since $s_G \geq 0$ and $1 - \delta p_G^\ell > 0$.

(b) If $n_g^\ell = 1$, then $D'_+(0;1,1) = \frac{(1-\delta)^2}{\beta_g(1,1)}\left\{(1-\delta)(1+2\delta f^h(1/2)) + 2\delta^2[p_G^h(1-p_G^\ell) + p_B^h p_G^\ell]\right\}$, so that the term in the curly brackets is strictly positive.

(c) If $n_g^\ell \geq 2$, then $s^h > 0$,[43] and:

$$\begin{aligned} D'_+(0;1,1) &= \frac{(1-\delta)^2}{\beta_g(n_g^\ell,1)} \left\{ 1 + \delta^{n_g^\ell+1}(1 - 2f^{h(n_g^\ell)}(p_G^\ell)) - \delta(1-2f^h(1/2)) \sum_{i=0}^{n_g^\ell} \delta^i \right\} \\ &\geq \frac{(1-\delta)^2}{\beta_g(n_g^\ell,1)} \left\{ 1 - \delta^{n_g^\ell+1} v^h_{n_g^\ell+1} - \sum_{i=0}^{n_g^\ell} \delta^{i+1}(s^h v_i^h - v_{i+1}^h) \right\} \\ &= \frac{(1-\delta)^2}{\beta_g(n_g^\ell,1)} \left\{ (1-\delta s^h) \sum_{i=0}^{n_g^\ell} \delta^i v_i^h \right\} \end{aligned}$$

where $v_i^h \coloneqq 2f^{h(i)}(1) - 1 \in [-1,1]$. The above uses the facts that (i) since $s^h > 0$ and $s_G \geq 0$, $f^{h(n_g^\ell)}(p_G^\ell) \leq f^{h(n_g^\ell)}(p_G^h) = f^{h(n_g^\ell+1)}(1)$ and (ii) $1 - 2f^h(1/2) = s^h v_i^h - v_{i+1}^h$.

Claim 5. The term $\sum_{i=0}^{n_g^\ell} \delta^i v_i^h$ is bounded below by $1 - \delta^{n_g^\ell}$.

Proof. The claim is trivially true when $n_g^\ell \leq 1$. If $n_g^\ell \geq 2$, then $s^h > 0$, and recall that $n_g^\ell = \min\{i : f^{h(i)}(p_B^\ell) < 1/2\}$. Thus, for $i \in \{0,1,\ldots,n_g^\ell - 1\}$, $f^{h(i)}(p_B^\ell) \geq 1/2$ and $f^{h(n_g^\ell)}(p_B^\ell) < 1/2$. It follows that for $i \in \{0,1,\ldots,n_g^\ell - 1\}$, $f^{h(i)}(1) \geq f^{h(i)}(p_B^\ell) \geq 1/2$. Thus $v_i^h \geq 0$ for $i \in \{0,1,\ldots,n_g^\ell - 1\}$ with $v_0^h = 1$, so the only term in $\sum_{i=0}^{n_g^\ell} \delta^i v_i^h$ that could be negative is $\delta^{n_g^\ell} v^h_{n_g^\ell}$, which is bounded below by $-\delta^{n_g^\ell}$. Therefore, $\sum_{i=0}^{n_g^\ell} \delta^i v_i^h \geq 1 + \delta^{n_g^\ell} v^h_{n_g^\ell} \geq 1 - \delta^{n_g^\ell}$. Note that the proof did not use the fact that $s_G \geq 0$. ■

Thus, for $n_g^\ell \geq 2$, $\sum_{i=0}^{n_g^\ell} \delta^i v_i^h$ is bounded below by $1 - \delta^{n_g^\ell} > 0$, so $D'_+(0;1,1) > 0$ and both conditions of Lemma 1 are satisfied. ■

Proposition 6 provides a characterisation of the optimal policy when $s_G \geq 0$. What happens when $s_G < 0$?

Observe that $D'_+(0;x,y)$ is affine in both x and y, with all second-order partial derivatives equalling 0. Therefore, for every $x \geq \bar{x}$, the value of y satisfying $D'_+(0;x,y) = 0$ gives the maximum value of y for which the greedy policy is optimal. Moreover, $D'_+(0;x,y) = 0$ defines y as an affine function of x. Given x, let $L_g(x)$ denote the value of y for which $D'_+(0;x,y) = 0$. By equation (B.5),

[43] If $s^h \leq 0$, then for all $p \geq 1/2$, $f^h(p) \leq f^h(1/2) < 1/2$, so $n_g^\ell \leq 1$.

$D'_+(0;\bar{x},\bar{y}) = (1-\delta)(\bar{y}-\bar{y}) = 0$, so $L_g(\bar{x}) = \bar{y}$. Therefore, $L_g(x) = \bar{y} + m_g(x-\bar{x})$, where m_g is the slope of $L_g(\cdot)$.

Define Y to be the value of y such that $D'_+(0;1,Y) = 0$. Thus, $m_g = (Y-\bar{y})/(1-\bar{x})$. By equations (B.6) and (B.7), $(1-\delta)m_g = \frac{\partial D'_+(0;x,\bar{y})}{\partial x}$. Define:

$$\Xi_{n_g^\ell} \coloneqq \begin{cases} \sum_{i=0}^{n_g^\ell} \delta^i v_i^h, & \text{if } n_g^{1/2} = 1 \\ \dfrac{1-\delta^{n_g^\ell+1} v^h_{n_g^\ell+1}}{1-\delta s^h}, & \text{if } n_g^{1/2} = \infty \end{cases}$$

Using equation (B.4), it is then straightforward to show that:[44]

$$m_g = \frac{(1-\delta^{n_g^{1/2}})}{\beta_g(n_g^\ell, n_g^{1/2})} \left\{ (1-\delta s^\ell)\,\Xi_{n_g^\ell} + 2\delta(-s_G) \sum_{i=0}^{n_g^\ell - 1} (\delta s^h)^i \right\}. \tag{B.8}$$

The lemma below follows from the above discussion and the fact that $y \leq 1$:

Lemma 11. The greedy policy is optimal if and only if $x \geq \bar{x}$ and $y \leq \min\{1, L_g(x)\}$.

Appendix C The Judicious Policy

Recall that the judicious policy maximises the probability with which the low action is taken by the Receiver each period. Formally:

$$(\alpha_j(p), q_j^h(p)) = \begin{cases} (2p-1, 1), & p > 1/2 \\ (0, p), & p \leq 1/2 \end{cases}$$

Observe that the judicious policy completely reverses the greedy policy: the Sender reveals no information when $p \leq 1/2$, and splits beliefs between $1/2$ and 1 when $p > 1/2$. Let $\Lambda(\cdot)$ denote the Sender's long-run payoff from using the judicious policy:

$$\Lambda(p) = \begin{cases} 2(1-p)\Lambda(1/2) + (2p-1)\Lambda(1), & p > 1/2 \\ (1-\delta)k^\ell(p) + \delta\Lambda(f^\ell(p)), & p \leq 1/2 \end{cases} \tag{C.1}$$

When $p = 1$, the Receiver is certain that the state is good, so he chooses the high action. This

[44] This decomposition follows from the fact that $f^{h(i)}(p_G^\ell) = f^{h(i)}(p_G^h) + (-s_G)(s^h)^i = f^{h(i+1)}(1) + (-s_G)(s^h)^i$.

implies that the next period's prior is $f^h(1) = p_G^h$, so that $\Lambda(1) = (1-\delta) + \delta\Lambda(p_G^h)$. Moreover, for $p > 1/2$, $\Lambda(\cdot)$ is affine, with a slope of $\Lambda' \coloneqq \Lambda'_+(1/2) = 2[\Lambda(1) - \Lambda(1/2)]$.

Since the judicious policy is a complete reversal of the greedy policy, the arguments for concavity of $\Lambda(\cdot)$ and the one-shot deviation condition are symmetric to those for the greedy policy. Therefore, the judicious policy equivalents of Lemmas 7–10 are stated in the following subsections without proof.

C.1 Concavity

The first lemma is a condition on the slope Λ':

Lemma 12. The payoff from the judicious policy is concave if and only if $\Lambda' \leq m_\infty^\ell$.

Briefly, the proof rests on showing that $\Lambda(\cdot)$ is piecewise affine in each of four cases, depending on the range of values of $f^\ell([0, 1/2])$.[45] These are illustrated in Figure 11. It is then sufficient to show that Λ is concave at the kinks, i.e. $\Lambda'_-(p) \geq \Lambda'_+(p)$ at kink p; this reduces to the condition above.

Let $n_j(p) \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ be the number of periods in which no information is disclosed before the belief *rises* above $1/2$, i.e. $n_j(p) = \min\{i : f^{\ell(i)}(p) > 1/2\}$. Using equation (A.4), for $p \in [0, 1)$, the long-run payoff from the judicious policy (C.1) can be re-written as:

$$\Lambda(p) = (1-\delta)U^\ell(p, n_j(p)) + \delta^{n_j(p)}\left\{2(1 - f^{\ell(n_j(p))}(p))\Lambda(1/2) + (2f^{\ell(n_j(p))}(p) - 1)\Lambda(1)\right\} \tag{C.2}$$

where $\Lambda(1) = (1-\delta) + \delta\Lambda(p_G^h)$. In the notation of Section 6.1, $n_j(f^\ell(1/2)) = n_j^\ell$ and $n_j(p_G^h) = n_j^h$.[46] Let $n_j^{1/2} \coloneqq n_j(1/2)$. Using equation (C.2), the value of Λ' is a function of $n_j^{1/2}$ and n_j^h, where $n_j^{1/2} \in \{1, \infty\}$ and $n_j^h \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$:

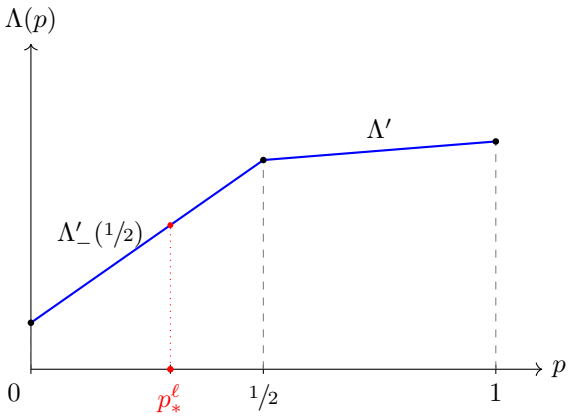
$$\Lambda' = 2(1-\delta)\frac{(1-\delta^{n_j^{1/2}})[1 + \delta U^\ell(p_G^h, n_j^h)] - (1-\delta^{n_j^h+1})U^\ell(1/2, n_j^{1/2})}{\beta_j(n_j^{1/2}, n_j^h)} \tag{C.3}$$

where $\beta_j(n_j^{1/2}, n_j^h) \coloneqq (1-\delta^{n_j^h+1})[1 - 2\delta^{n_j^{1/2}}(1 - f^{\ell(n_j^{1/2})}(1/2))] + 2\delta^{n_j^h+1}(1-\delta^{n_j^{1/2}})[1 - f^{\ell(n_j^h)}(p_G^h)] > 0$.

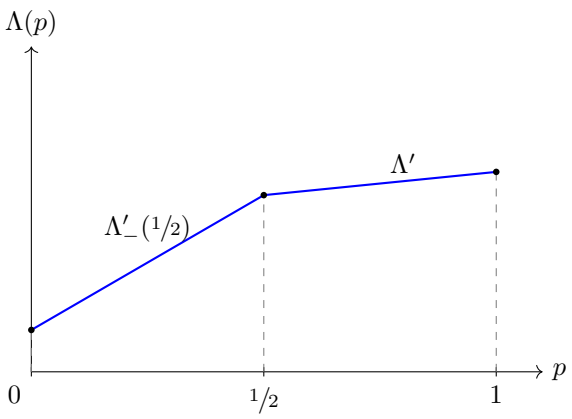
From the above, Lemma 12 reduces to a simple condition on the preference parameter y:

[45] Under the judicious policy, the Receiver takes the low action only if $p \leq 1/2$, so the relevant range of beliefs following the low action is $f^\ell([0, 1/2])$.

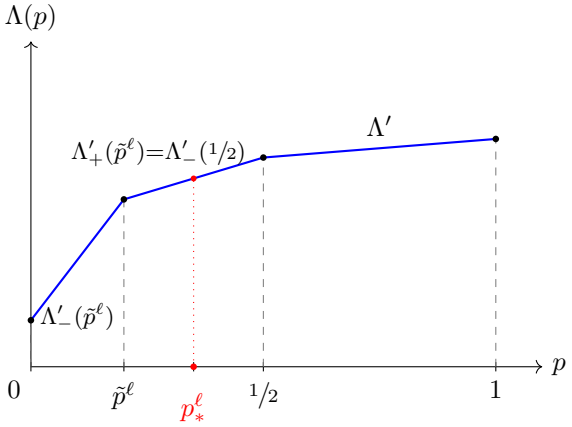
[46] Since the prior following a high action is p_G^h, n_j^h measures the length of the quiet spell following the *high* action.



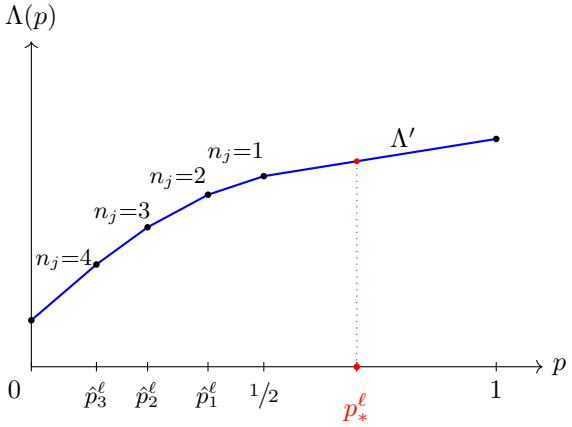
(a) Absorbing case: $f^\ell([0,1/2]) \subseteq [0,1/2]$



(b) Reverting case: $f^\ell([0,1/2]) \subseteq (1/2,1]$



(c) Partially absorbing case: $f^\ell(1/2) \leq 1/2 < f^\ell(0)$



(d) Cascading case: $f^\ell(1/2) > 1/2 \geq f^\ell(0)$

Figure 11: The Sender's payoff $\Lambda(\cdot)$ under the judicious policy of Section C.1; p^ℓ_* is the fixed point of $f^\ell(\cdot)$, drawn for $|s^\ell| < 1$

Lemma 13. The payoff from the judicious policy is concave if and only if $y \geq \bar{y}$.

As with the proof of Lemma 8, it turns out that $m^{\ell}_{\infty} - \Lambda' = \chi(n_j^{1/2}, n_j^h) \cdot (y - \bar{y})$ for each $(n_j^{1/2}, n_j^h)$ pair, where $\chi(n_j^{1/2}, n_j^h) = \frac{2(1-\delta)(1-\delta^{n_j^{1/2}})(1+\delta p_B^{\ell} - \delta p_G^h)}{\beta_j(n_j^{1/2}, n_j^h)(1-\delta s^{\ell})} > 0$. In particular, when $y = \bar{y}$, then $\Lambda' = m^{\ell}_{\infty}$.

C.2 One-Shot Deviation Condition

To show that the judicious policy is optimal, we now need to verify that the Sender does not strictly gain from a one-shot deviation to no disclosure, i.e. we must show that $\mathcal{D}(p; x, y) \coloneqq \Lambda(p) - \{(1-\delta)k^h(p) + \delta\Lambda(f^h(p))\} \geq 0 \ \forall p \in (1/2, 1]$.[47] The judicious policy equivalents of Lemmas 9 and 10 are as follows:

Lemma 14. For a given y, if the judicious policy is optimal for some $x = \hat{x}$, then it is optimal for all $x \in [0, \hat{x}]$.

Lemma 15. For a given (x, y), the judicious policy is optimal if and only if $\Lambda(\cdot)$ is concave in p and $\mathcal{D}'_{-}(1; x, y) \leq 0$.

Again, since $\mathcal{D}'_{-}(1; x, y)$ is affine in y, if $\bar{y} < 1$, it suffices to evaluate $\mathcal{D}'_{-}(1; x, y)$ at the extreme values $y = \bar{y}$ and $y = 1$. As with Λ', the value of $\mathcal{D}'_{-}(1; x, y)$ depends on the branch of Λ to which p_G^h belongs.[48] For $n_j^h \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$, $\mathcal{D}'_{-}(1; x, y)$ is given by:

$$\mathcal{D}'_{-}(1; x, y) = \Lambda' - (1-\delta)(1-x) - \delta s^h \left\{ m^{\ell}_{n_j^h} + (\delta s^{\ell})^{n_j^h} \Lambda' \right\}$$

where Λ' is as defined in equation (C.3). Thus,

$$\mathcal{D}'_{-}(1; x, \bar{y}) = (1-\delta)(x - \bar{x}) \leq 0 \iff x \leq \bar{x} \tag{C.4}$$

Finally, since $\mathcal{D}'_{-}(1; \bar{x}, \cdot)$ is affine in y and vanishes at $y = \bar{y}$ by equation (C.4), we have that:

$$\mathcal{D}'_{-}(1; \bar{x}, 1) = (1 - \bar{y}) \cdot \frac{\partial \mathcal{D}'_{-}(1; \bar{x}, y)}{\partial y} \tag{C.5}$$

[47] Continuity of Λ also gives the condition at $p = 1/2$.

[48] If p_G^h is a kink of Λ, read n_j^h in the display below and in equation (C.6) as $n_j(f^h(1-\varepsilon))$ for small $\varepsilon > 0$; equation (C.3) holds with either value, since Λ is continuous.

C.3 Optimality

By Lemma 13, the payoff from the judicious policy is concave if and only if $y \geq \bar{y}$. By Lemmas 14 and 15 and equations (C.4) and (C.5), an argument symmetric to that for the greedy policy shows that $\mathcal{D}'_{-}(1; x, y) = 0$ defines x as an affine function of y and given y, gives the maximum value of x for which the judicious policy is optimal. We can analogously define the line $L_j(y) = \bar{x} + m_j(y - \bar{y})$ where $(1-\delta)m_j = -\frac{\partial \mathcal{D}'_{-}(1;\bar{x},y)}{\partial y}$.

Let $v_i^\ell := 1 - 2f^{\ell(i)}(0) \in [-1, 1]$ denote the judicious policy counterpart of v_i^h, and define:

$$\Theta_{n_j^h} := \begin{cases} \sum_{i=0}^{n_j^h} \delta^i v_i^\ell, & \text{if } n_j^{1/2} = 1 \\ \dfrac{1 - \delta^{n_j^h+1} v_{n_j^h+1}^\ell}{1 - \delta s^\ell}, & \text{if } n_j^{1/2} = \infty \end{cases}$$

Using equation (C.3), it is then straightforward to show that:[49]

$$m_j = \frac{(1 - \delta^{n_j^{1/2}})}{\beta_j(n_j^{1/2}, n_j^h)} \left\{ (1 - \delta s^h)\, \Theta_{n_j^h} + 2\delta(-s_B) \sum_{i=0}^{n_j^h - 1} (\delta s^\ell)^i \right\}. \tag{C.6}$$

Thus, since $x \leq 1$, we have that:

Lemma 16. The judicious policy is optimal if and only if $y \geq \bar{y}$ and $x \leq \min\{1, L_j(y)\}$.

Appendix D Full Disclosure

Let $\theta(\cdot)$ denote the payoff to the Sender from fully revealing the state to the Receiver, i.e. $(\alpha_{fd}(p), q_{fd}^h(p)) = (p, 1)$, $\forall p \in [0, 1]$. For $p \in [0, 1]$, the Sender's long-run expected payoff is given by:

$$\begin{aligned} \theta(p) &= (1-\delta)p + \delta p \cdot \theta(f^h(1)) + \delta(1-p) \cdot \theta(f^\ell(0)) \\ &= p \cdot \theta(1) + (1-p) \cdot \theta(0) \end{aligned}$$

where $\theta(0) = \delta\theta(p_B^\ell)$ and $\theta(1) = (1-\delta) + \delta\theta(p_G^h)$. Thus, $\theta(\cdot)$ is affine over $[0, 1]$. Since the above is true for all $p \in [0, 1]$, it holds at $f^h(1) = p_G^h$ and $f^\ell(0) = p_B^\ell$ as well. It is then easy to solve for the

[49] This decomposition follows from the fact that: (i) $f^{\ell(i)}(p_B^h) = f^{\ell(i+1)}(0) - (-s_B)(s^\ell)^i$, and (ii) $2f^\ell(1/2) - 1 = s^\ell v_i^\ell - v_{i+1}^\ell$.

Sender's long-run payoff when she fully discloses the state every period:

$$\theta(p) = \frac{(1-\delta)p + \delta p_B^\ell}{1 + \delta p_B^\ell - \delta p_G^h}$$

Proposition 7. Full disclosure is optimal if and only if $x \leq \bar{x}$ and $y \leq \bar{y}$.

Proof. Since $\theta(\cdot)$ is affine and hence concave, using Lemma 1, it only remains to show that the Sender gets a weakly higher payoff than a one-shot deviation to providing no information. Thus, if $p \geq 1/2$, we require that:

$$\theta(p) \geq (1-\delta)(p + (1-p)x) + \delta\theta(f^h(p)) \iff x \leq \bar{x}$$

and if $p \leq 1/2$, we require that:

$$\theta(p) \geq (1-\delta)py + \delta\theta(f^\ell(p)) \iff y \leq \bar{y}$$ ■

Appendix E Omitted Proofs

E.1 Proof of Proposition 1

If states evolve exogenously, $s_G = 0$ and $s_B = 0$. Hence, $\bar{x} = 0$ and $\bar{y} = 1$. The result follows from Proposition 6. ■

E.2 Proof of Proposition 2

For all possible values of the transition probabilities, it is always the case that $\bar{x} \leq 1$ and $\bar{y} \geq 0$. Thus, when $(x, y) = (1, 0)$, by Lemma 8, the payoff induced by the greedy policy is concave.

Further, when $x = 1$, the numerator of γ' in equation (B.4) is $2(1-\delta)(1-\delta^{n_g^{1/2}}) > 0$. Thus, by Lemma 10 and equation (B.6), $D'_+(0; 1, 0) = \gamma'(1 - \delta s^\ell(\delta s^h)^{n_g^\ell}) > 0$, proving the result. ■

E.3 Proof of Theorem 1

Under positive externalities, $s_G > 0$ and $s_B > 0$. Hence, $\bar{x} \leq 0$ and $\bar{y} \geq 1$. The result then follows from Proposition 6. ■

E.4 Proof of Theorem 2

Under negative externalities, $s_G < 0$ and $s_B < 0$. Hence, $\bar{x} \geq 0$ and $\bar{y} \leq 1$.

Part (i): restates Proposition 7.

Part (ii): The statement of the result restates Lemma 11. It only remains to show that $m_g > 0$, so that L_g is strictly increasing.

In equation (B.8), we only need to show that the term in the curly brackets is positive. Under negative externalities, $-s_G > 0$, so the second term is non-negative. For the first term, since $|\delta s^a| < 1$, it suffices to show that $\Xi_{n_g^\ell}$ is positive. If $n_g^{1/2} = \infty$, this is immediate: $v^h_{n_g^\ell+1} \leq 1$, so that $\Xi_{n_g^\ell} \geq (1-\delta^{n_g^\ell+1})/(1-\delta s^h) > 0$.

If $n_g^{1/2} = 1$, i.e. $f^h(1/2) < 1/2$, then it suffices to show that $\sum_{i=0}^{n_g^\ell} \delta^i v_i^h$ is positive. The sum equals 1 when $n_g^\ell = 0$. It follows from Claim 5 that the sum is bounded below by $1-\delta^{n_g^\ell} > 0$ when $n_g^\ell \geq 1$. Thus, $m_g > 0$.

Part (iii): The statement of the result restates Lemma 16. It only remains to show that $m_j > 0$.

Again, in equation (C.6), we only need to show that the term in the curly brackets is positive. Under negative externalities, $-s_B > 0$, so the second term is non-negative. For the first term, since $|\delta s^a| < 1$, it only remains to show that $\Theta_{n_j^h}$ is positive. If $n_j^{1/2} = \infty$, this is immediate: $v^\ell_{n_j^h+1} \leq 1$, so that $\Theta_{n_j^h} \geq (1-\delta^{n_j^h+1})/(1-\delta s^\ell) > 0$. If $n_j^{1/2} = 1$, i.e. $f^\ell(1/2) > 1/2$, this implies $n_j^h \in \mathbb{Z}_{\geq 0}$ and we have the following:

Claim 6. The term $\sum_{i=0}^{n_j^h} \delta^i v_i^\ell$ is bounded below by $1-\delta^{n_j^h}$.

Proof. The claim is trivially true if $n_j^h \leq 1$. If $n_j^h \geq 2$, $s^\ell > 0$,[50] and recall that $n_j^h = \min\{i : f^{\ell(i)}(p_G^h) > 1/2\}$. Thus, for $i \in \{0, 1, \ldots, n_j^h - 1\}$, $f^{\ell(i)}(0) \leq f^{\ell(i)}(p_G^h) \leq 1/2$ and $f^{\ell(n_j^h)}(p_G^h) > 1/2$. An argument symmetric to Claim 5 proves the result. ■

Since the sum equals 1 when $n_j^h = 0$ and is at least $1-\delta^{n_j^h} > 0$ otherwise, we have shown that $m_j > 0$.

[50] If $s^\ell \leq 0$, then for all $p \leq 1/2$, $f^\ell(p) \geq f^\ell(1/2) > 1/2$, so $n_j^h \leq 1$.

Part (iv): Parts (i)–(iii) give necessary and sufficient conditions for the optimality of the greedy, judicious and full disclosure policies, and $m_g, m_j > 0$. Let $\mathcal{R} \coloneqq \{(x,y) \in [0,1]^2 : x > \bar{x}, y > \bar{y}, y > L_g(x), x > L_j(y)\}$ be the set of (x,y) for which none of these three policies are optimal.

Non-emptiness of $\mathcal{R}$. If $\mathcal{R}$ is empty, then parts (i)–(iii) cover the whole unit square, and there is nothing left to prove. The following result characterises when $\mathcal{R}$ is non-empty.

Lemma 17. $m_g m_j \leq 1$, and $\mathcal{R}$ is non-empty if and only if $m_g m_j < 1$.

Proof. By parts (ii) and (iii), $m_g > 0$ and $m_j > 0$. So, for $x > \bar{x}$, the conditions $y > L_g(x)$ and $x > L_j(y)$ together read $m_g(x - \bar{x}) < y - \bar{y} < (x - \bar{x})/m_j$ and they imply $y > \bar{y}$. If $(x,y) \in \mathcal{R}$, dividing by $x - \bar{x} > 0$ gives $m_g < 1/m_j$. Conversely, if $m_g m_j < 1$, pick $k \in (m_g, 1/m_j)$. Then $(\bar{x}+\varepsilon, \bar{y}+k\varepsilon)$ satisfies the inequalities above for every $\varepsilon > 0$, and it lies in $[0,1]^2$ for ε small, because $\bar{x}, \bar{y} \in (0,1)$ when $\delta > 0$; if $\delta = 0$, then $m_g = m_j = 1$.

It remains to show that $m_g m_j \leq 1$. Suppose instead that $m_g m_j > 1$, and pick $k \in (1/m_j, m_g)$. For $\varepsilon > 0$ small, $(x,y) \coloneqq (\bar{x} + \varepsilon, \bar{y} + k\varepsilon)$ lies in $[0,1]^2$ and satisfies $x > \bar{x}$, $y > \bar{y}$, $y < L_g(x)$ and $x < L_j(y)$. By parts (ii) and (iii), both the greedy and the judicious policies are optimal at (x,y). Since the judicious policy discloses nothing at priors below $1/2$, $V = w^\ell$ on $[0, 1/2)$. Since the greedy policy is optimal, $V = \gamma$. Hence $D(\cdot; x, y) = 0$ on $[0, 1/2)$, so $D'_+(0; x, y) = 0$, i.e. $y = L_g(x)$. This contradicts $y < L_g(x)$. ■

Equivalently, the following result characterises when $\mathcal{R}$ is non-empty in terms of the transition probabilities:

Lemma 18. Suppose $\delta > 0$. Then $\mathcal{R}$ is empty if and only if neither $f^\ell([0, 1/2])$ nor $f^h([1/2, 1])$ contains beliefs on both sides of $1/2$.

Proof. I say that f^ℓ *crosses* $1/2$ if $f^\ell([0, 1/2])$ contains beliefs both above and below $1/2$, and that f^h crosses $1/2$ if $f^h([1/2, 1])$ does. A policy whose split attains the maximum in equation (3.5) at every prior is optimal, since V is then the fixed point of the policy's own Bellman operator, a contraction.

($\implies$): If $\mathcal{R}$ is empty, by Lemma 17, $L_g = L_j$. Since $\bar{x}, \bar{y} \in (0,1)$ when $\delta > 0$, this line contains points $(x,y) \in [0,1]^2$ with $x > \bar{x}$, at which both the greedy and the judicious policies are optimal

by parts (ii) and (iii) above. There, V is affine on each half, $V = w^\ell$ on $[0, 1/2]$ and $V = w^h$ on $[1/2, 1]$. If f^ℓ crosses $1/2$, then $V \circ f^\ell = (w^\ell - (1-\delta)k^\ell)/\delta$ is affine on $[0, 1/2]$, although f^ℓ maps $[0, 1/2]$ onto beliefs on both sides of $1/2$; so V has no kink at $1/2$. The same holds if f^h crosses $1/2$. Then V is affine on $[0, 1]$, and full disclosure attains the maximum at every prior, so it is optimal although $x > \bar{x}$. This contradicts part (i) of Theorem 2, proved above.

($\Longleftarrow$): Let W be affine on $[0, 1/2]$ and on $[1/2, 1]$. Since $f^\ell([0, 1/2])$ and $f^h([1/2, 1])$ each lie in one half, w^ℓ_W is affine on $[0, 1/2]$ and w^h_W on $[1/2, 1]$ (with w_W as in the proof of Claim 1). So w_W lies below the piecewise-linear interpolation of its values at 0, $1/2$ and 1, and TW is the concave hull of these three points, which is again affine on each half. As $T^n 0 \to V$ uniformly (Lemma 3), V is affine on each half, and so are w^ℓ on $[0, 1/2]$ and w^h on $[1/2, 1]$. Recall that $V \geq w$, with equality at 0 and 1. If $V(1/2) = w^h(1/2)$, then $V - w^h$ is affine and non-negative on $[1/2, 1]$ and zero at both ends, so $V = w^h$ there; since V is affine on $[0, 1/2]$, the greedy policy attains the maximum at every prior. If $V(1/2) = w^\ell(1/2)$, the judicious policy does, symmetrically. Otherwise $V(1/2) > w(1/2)$, so V is affine around $1/2$ (Lemma 5), hence on $[0, 1]$, and full disclosure attains the maximum at every prior. So full disclosure, the greedy or the judicious policy is optimal at every (x, y), and $\mathcal{R}$ is empty. ■

Assume from now on that $m_g m_j < 1$. Let $\Gamma_g := \{(x, L_g(x)) : x > \bar{x}\} \cap [0, 1]^2$ and $\Gamma_j := \{(L_j(y), y) : y > \bar{y}\} \cap [0, 1]^2$ denote the two edges of $\mathcal{R}$ that start at $(\bar{x}, \bar{y})$. Both are segments of positive length; the greedy policy is optimal on Γ_g and the judicious policy is optimal on Γ_j.

Recall from Appendix A.2 that $w^a(q) = (1 - \delta)k^a(q) + \delta V(f^a(q))$ denotes the Sender's payoff from inducing the posterior q together with action $a \in \{h, \ell\}$ and then continuing optimally. By Corollary 4, for all $(x, y) \in \mathcal{R}$, $V(1/2) = \max\{w^h(1/2), w^\ell(1/2)\}$. Let $\mathcal{V}(x, y) := w^h(1/2) - w^\ell(1/2)$ be the Sender's gain from inducing action h rather than action ℓ at $q = 1/2$. Fix $(x, y) \in \mathcal{R}$.

The no-disclosure line. If $\mathcal{V}(x, y) = 0$, then $w^h(1/2) = w^\ell(1/2)$ and by Lemma 5 and Corollary 4, any interval in I would have $1/2$ as an endpoint with $w^h(1/2) \neq w^\ell(1/2)$. So $I = \emptyset$, i.e. $V = w$, and disclosing no information at any prior is optimal. Let $N(\cdot)$ denote the Sender's payoff from never disclosing information, with the Receiver taking h at the belief $1/2$. Along the no-disclosure path

from any prior the sequence of actions is fixed, so $N(p)$ is affine in (x, y) for each p. Let

$$\mathcal{N}(x, y) \coloneqq \left[(1-\delta)k^h(1/2) + \delta N\big(f^h(1/2)\big)\right] - \left[(1-\delta)k^\ell(1/2) + \delta N\big(f^\ell(1/2)\big)\right],$$

which is affine in (x, y). I refer to $\Gamma_N \coloneqq \{(x, y) : \mathcal{N}(x, y) = 0\}$ as the *no-disclosure line*. Two observations link it to the optimal policy.

(a) *If disclosing no information is optimal at (x, y), then $\mathcal{N}(x, y) = \mathcal{V}(x, y) = 0$.* Indeed, optimality of no disclosure implies $V = w$. Since V, w^h and w^ℓ are continuous at $1/2$, $w^\ell(1/2) = \lim_{p\uparrow 1/2} V(p) = V(1/2) = \lim_{p\downarrow 1/2} V(p) = w^h(1/2)$, so $\mathcal{V}(x, y) = 0$. Hence $V = w$ satisfies the recursion that defines N (with action h at $1/2$), so $V = N$, and therefore $\mathcal{N}(x, y) = \mathcal{V}(x, y) = 0$.

(b) *If the greedy policy is optimal at (x, y), then $\mathcal{V}(x, y) \geq 0$, and if $\mathcal{V}(x, y) = 0$ then disclosing no information is also optimal.* Indeed, by equation (B.1), $V(1/2) = \gamma(1/2) = w^h(1/2)$, while $V(1/2) \geq w^\ell(1/2)$ by equation (A.5). If, in addition, $\mathcal{V}(x, y) = 0$, then $V(1/2) = w(1/2)$ and $w^h(1/2) = w^\ell(1/2)$, so by Lemma 5, $I = \emptyset$. Symmetrically, using equation (C.1), if the judicious policy is optimal at (x, y) then $\mathcal{V}(x, y) \leq 0$, and if $\mathcal{V}(x, y) = 0$ then disclosing no information is also optimal.

In particular, at $(\bar{x}, \bar{y})$, both the greedy and the judicious policies are optimal, so $V = w$ on $[1/2, 1]$ (greedy) and on $[0, 1/2)$ (judicious). Disclosing no information is therefore optimal at $(\bar{x}, \bar{y})$, and by (a), the no-disclosure line passes through $(\bar{x}, \bar{y})$.

Truncated policies. If $\mathcal{V}(x, y) > 0$, then $V(1/2) \geq w^h(1/2) > w^\ell(1/2)$. By continuity of V and w^ℓ at $1/2$, $V(p) > w^\ell(p) = w(p)$ for all $p < 1/2$ close enough to $1/2$, so $I \neq \emptyset$. By Lemma 5 and Corollary 4, $I = (\underline{p}, 1/2)$ with $\underline{p} < 1/2$. If $\underline{p} = 0$, the greedy policy would be optimal, which is impossible on $\mathcal{R}$. So $\underline{p} \in (0, 1/2)$, and the truncated greedy policy with $q = \underline{p}$ is optimal. If $\mathcal{V}(x, y) < 0$, the symmetric argument shows that the truncated judicious policy with some $q \in (1/2, 1)$ is optimal.

Now, for each p, $V(p)$ is continuous in (x, y): changing (x, y) by at most ε in each coordinate changes the payoff of every policy by at most ε. Hence $\mathcal{V}$ is continuous on $[0, 1]^2$ and $\mathcal{V} \neq 0$ on $\mathcal{R}$

off the no-disclosure line. There are two cases.

Case A: the no-disclosure line Γ_N intersects $\mathcal{R}$. Since it passes through $(\bar{x}, \bar{y})$, its slope lies strictly between m_g and $1/m_j$, and it divides $\mathcal{R}$ into two non-empty convex (hence connected) sets: $\mathcal{R}_+$, between Γ_g and Γ_N, and $\mathcal{R}_-$, between Γ_N and Γ_j. $\mathcal{V}$ has a constant sign on each.

Suppose $\mathcal{V} < 0$ on $\mathcal{R}_+$. Since every point of Γ_g close to $(\bar{x}, \bar{y})$ is a limit of points in $\mathcal{R}_+$, continuity gives $\mathcal{V} \leq 0$ there, and by point (b) above $\mathcal{V} \geq 0$, so $\mathcal{V} = 0$ and disclosing no information is optimal at these points. This implies that by point (a) above, $\mathcal{N} = 0$ on L_g starting at $(\bar{x}, \bar{y})$, so the no-disclosure line is L_g, which does not intersect $\mathcal{R}$, a contradiction. Hence $\mathcal{V} > 0$ on $\mathcal{R}_+$, where a truncated greedy policy is optimal. Symmetrically, $\mathcal{V} < 0$ on $\mathcal{R}_-$, where a truncated judicious policy is optimal. Finally, at points of $\mathcal{R}$ on the line, $\mathcal{V} = 0$ by continuity, so disclosing no information is optimal by the previous discussion on the no-disclosure line.

Case B: the no-disclosure line does not intersect $\mathcal{R}$. Then $\mathcal{V} \neq 0$ on the connected set $\mathcal{R}$, so it has a constant sign. If $\mathcal{V} > 0$ on $\mathcal{R}$, the argument of Case A applied to Γ_j shows that disclosing no information is optimal on Γ_j and that the no-disclosure line coincides with L_j. Then all of $\mathcal{R}$ lies below the no-disclosure line, and a truncated greedy policy is optimal throughout. If $\mathcal{V} < 0$ on $\mathcal{R}$, the no-disclosure line coincides with L_g, all of $\mathcal{R}$ lies above it, and a truncated judicious policy is optimal throughout.

In either case, the no-disclosure line passes through $(\bar{x}, \bar{y})$ and has slope in $[m_g, 1/m_j]$, so it is strictly increasing. The optimal policy is truncated greedy below it and truncated judicious above it, and disclosing no information is optimal on it.

Characterising the truncated threshold. It only remains to characterise $\underline{q}$ and $\overline{q}$. Recall from Section 6.1 that, under the truncated greedy policy, $I_{tg} = (\underline{q}, 1/2)$, $F_{tg} = f^\ell$ on $[0, \underline{q}]$ and $F_{tg} = f^h$ on $[1/2, 1]$. Under the truncated judicious policy, $I_{tj} = (1/2, \overline{q})$, $F_{tj} = f^\ell$ on $[0, 1/2]$ and $F_{tj} = f^h$ on $[\overline{q}, 1]$.

Lemma 19. Suppose $I = (\underline{q}, 1/2)$ for some $\underline{q} \in (0, 1/2)$, so that the truncated greedy policy with threshold $\underline{q}$ is optimal. Then exactly one of the following holds:

(i) $F_{tg}^{(n)}(z_{tg}^\ell) = 1/2$ for some $n \geq 0$, and $F_{tg}^{(i)}(z_{tg}^\ell) \notin [\underline{q}, 1/2]$ for $i < n$;

(ii) the path $F_{tg}^{(n)}(z_{tg}^{\ell})$ never meets $[\underline{q}, 1/2]$.

Symmetrically, suppose $I = (1/2, \overline{q})$ for some $\overline{q} \in (1/2, 1)$ so that the truncated judicious policy with threshold $\overline{q}$ is optimal. Then either $F_{tj}^{(n)}(z_{tj}^{h}) = 1/2$ for some $n \geq 0$ and $F_{tj}^{(i)}(z_{tj}^{h}) \notin [1/2, \overline{q}]$ for $i < n$, or the path from z_{tj}^{h} never meets $[1/2, \overline{q}]$.

Proof. Consider the truncated greedy policy; the judicious case is symmetric. By Lemma 5, $I = (\underline{q}, 1/2)$, V is affine on $[\underline{q}, 1/2]$ with slope ρ say, and $V > w = w^{\ell}$ on I. Further, $V = w^{\ell}$ on $[0, \underline{q}]$, and $V = w^{h}$ on $[1/2, 1]$, because $V(1/2) = w^{h}(1/2)$. Recall that $z_{tg}^{\ell} = f^{\ell}(\underline{q})$.

We first show that V is not affine on any neighbourhood of z_{tg}^{ℓ}. If it were, then $w^{\ell}(p) = (1-\delta)k^{\ell}(p) + \delta V(f^{\ell}(p))$ would be affine on an open interval around $\underline{q}$. Its slope there equals $V'_{-}(\underline{q}) \geq V'_{+}(\underline{q}) = \rho$, because $w^{\ell} = V$ to the left of $\underline{q}$, and is at most ρ, because $w^{\ell} \leq V$ to the right of $\underline{q}$ with equality at $\underline{q}$. Hence $w^{\ell} = V$ just to the right of $\underline{q}$, contradicting $V > w = w^{\ell}$ on I.

Now suppose that neither (i) nor (ii) holds. Then the path from z_{tg}^{ℓ} first meets $[\underline{q}, 1/2]$ at some time n, at a point of $[\underline{q}, 1/2)$. Before time n it lies in $(0, \underline{q})$, where $V = w^{\ell}$, or in $(1/2, 1)$, where $V = w^{h}$. As in the proof of Lemma 6, for a closed interval in either set, the slope drop of V across it equals $\delta|s^{a}|$ times the slope drop of V across its image under $f^{a} = F_{tg}$. For $\zeta > 0$, let $\psi(\zeta)$ be the slope drop of V across $[z_{tg}^{\ell} - \zeta, z_{tg}^{\ell} + \zeta]$. Applying this n times, with ζ small, gives

$$\psi(\zeta) \leq \left(\delta \max_{a \in \mathcal{A}} |s^{a}|\right)^{n} [V'_{-}(c) - V'_{+}(d)],$$

where $[c, d]$ is an interval of length at most 2ζ centred at the point where the path meets $[\underline{q}, 1/2)$. If that point lies in I, then $[c, d] \subseteq I$ for ζ small, and the bracket is zero. If it is $\underline{q}$, then $[c, d] = [\underline{q} - \zeta', \underline{q} + \zeta']$ with $\zeta' \leq \zeta$. Since $\underline{q} + \zeta' \in I$, $V = w^{\ell}$ on $[0, \underline{q}]$, and $V'_{+}(\underline{q}) \geq (w^{\ell})'_{+}(\underline{q})$ as in the proof of Lemma 6,

$$\begin{aligned} V'_{-}(c) - V'_{+}(d) &\leq (w^{\ell})'_{-}(\underline{q} - \zeta') - (w^{\ell})'_{+}(\underline{q}) \\ &= \delta|s^{\ell}|\,[\text{slope drop of } V \text{ across } f^{\ell}([\underline{q} - \zeta', \underline{q}])] \\ &\leq \delta \max_{a \in \mathcal{A}} |s^{a}| \psi(\zeta), \end{aligned}$$

because $f^{\ell}([\underline{q} - \zeta', \underline{q}])$ contains z_{tg}^{ℓ} and has length at most ζ. In either case, $\psi(\zeta) \leq \delta \max_{a \in \mathcal{A}} |s^{a}| \psi(\zeta)$, so $\psi(\zeta) = 0$ and V is affine on $[z_{tg}^{\ell} - \zeta, z_{tg}^{\ell} + \zeta]$, a contradiction. ∎

Lemma 19 characterises the thresholds. Suppose (i) holds, and let $a_1, \ldots, a_n$ be the actions taken along the path from z^ℓ_{tg} to $1/2$: $a_i = \ell$ if $F^{(i-1)}_{tg}(z^\ell_{tg}) < \underline{q}$ and $a_i = h$ if $F^{(i-1)}_{tg}(z^\ell_{tg}) \geq 1/2$. Then $\underline{q}$ solves the affine equation

$$f^{a_n}(\cdots f^{a_1}(f^\ell(\underline{q}))) = \frac{1}{2} \tag{E.1}$$

So $\underline{q}$ belongs to a countable set of candidate thresholds that depends only on the transition probabilities, and not on δ, x or y; the preference parameters only select among the candidates. The simplest candidate, with $n = 0$, is $\underline{q} = (1/2 - p^\ell_B)/s^\ell$, which lies in $(0, 1/2)$ if and only if $1/2$ lies strictly between p^ℓ_B and $f^\ell(1/2)$. In Example 3, this gives $\underline{q} = (0.5 - 0.4)/0.5 = 0.2$. The same holds for $\overline{q}$, with the path from $z^h_{tj} = f^h(\overline{q})$ and h in place of ℓ in equation (E.1).

In case (ii) of Lemma 19, $n^\ell_{tg} = \infty$: after a disclosure that induces ℓ, the Sender never discloses again.

This completes the proof. ■

E.5 Proof of Proposition 3

Part (i): Differentiating,

$$\frac{\partial \bar{x}}{\partial \delta} = \frac{-s_B}{(1 + \delta p^\ell_B - \delta p^h_G)^2} > 0, \qquad \frac{\partial \bar{y}}{\partial \delta} = \frac{s_G}{(1 + \delta p^\ell_B - \delta p^h_G)^2} < 0.$$

Part (ii): The probability p^h_B only appears in the numerator of $\bar{x}$, so $\partial \bar{x}/\partial p^h_B = -\delta/(1 + \delta p^\ell_B - \delta p^h_G) < 0$ and does not affect $\bar{y}$. Symmetrically, p^ℓ_G only appears in the numerator of $\bar{y}$, so $\partial \bar{y}/\partial p^\ell_G = -\delta/(1 + \delta p^\ell_B - \delta p^h_G) < 0$ and does not affect $\bar{x}$.

Part (iii): Differentiating,

$$\frac{\partial \bar{x}}{\partial p^h_G} = \frac{-\delta^2 s_B}{(1 + \delta p^\ell_B - \delta p^h_G)^2}, \qquad \frac{\partial \bar{y}}{\partial p^h_G} = \frac{\delta[1 - \delta s^\ell]}{(1 + \delta p^\ell_B - \delta p^h_G)^2},$$

$$\frac{\partial \bar{x}}{\partial p^\ell_B} = \frac{\delta(1 - \delta s^h)}{(1 + \delta p^\ell_B - \delta p^h_G)^2}, \qquad \frac{\partial \bar{y}}{\partial p^\ell_B} = \frac{-\delta^2 s_G}{(1 + \delta p^\ell_B - \delta p^h_G)^2},$$

and by assumption, $\delta > 0$, $\delta s^a < 1$, $s_G < 0$ and $s_B < 0$, so all four expressions above are strictly positive. ■

E.6 Proof of Lemma 2

Recall that $\tau = \min\{t : p_t \in I_\pi\}$ is the first period in which the Sender provides information.

Part (i): Before τ the prior lies outside I_π, where the Sender discloses nothing and the belief moves by F_π.

Part (ii): In any period t with $p_t \in I_\pi$, the posterior is q_π^h with probability $\alpha_\pi(p_t)$ and q_π^ℓ otherwise, so the next prior is z_π^a if the disclosure induces action a. Starting from z_π^a, the belief moves by F_π through $F_\pi^{(n)}(z_\pi^a)$, $n < n_\pi^a$, all outside I_π, so the Sender discloses nothing for n_π^a periods. If $n_\pi^a < \infty$, the belief then reaches $r_\pi^a \in I_\pi$, and the same argument applies. Induction on the number of disclosures proves the result.

Part (iii): Let $P_\pi^a \coloneqq \{F_\pi^{(n)}(z_\pi^a) : 0 \le n \le n_\pi^a\}$ be the set of beliefs on the path from z_π^a to r_π^a, and $P_\pi \coloneqq P_\pi^h \cup P_\pi^\ell$. By part (ii), $p_t \in P_\pi$ for every $t > \tau$. If $n_\pi^h, n_\pi^\ell < \infty$, each P_π^a has at most $n_\pi^a + 1$ elements, so P_π is finite. Transitions from beliefs in P_π other than r_π^h and r_π^ℓ are deterministic, by F_π; from r_π^a, the belief moves to z_π^h with probability $\alpha_\pi(r_\pi^a)$ and to z_π^ℓ otherwise, by the first step of part (ii). ■

E.7 Proof of Proposition 4

By equation (A.3), $p_G^h + p_B^h \ge 1 \iff f^h(1/2) = (p_G^h + p_B^h)/2 \ge 1/2 \iff p_*^h \ge 1/2$. By equation (A.2), $f^{h(n)}(1/2) = p_*^h - (s^h)^n(p_*^h - 1/2) \ge p_*^h - (p_*^h - 1/2) = 1/2$. So under the greedy policy, once the posterior $1/2$ is realised, the belief never re-enters $I_g = (0, 1/2)$ and it converges to p_*^h; that is $n_g^{1/2} = \infty$.

Recall that $\tau = \min\{t : p_t \in I_\pi\}$ is the first period in which the Sender provides information. If $\tau = \infty$, the belief never falls below $1/2$, the Receiver always takes the high action, and $p_t \to p_*^h$ by equation (A.2). Suppose instead $\tau < \infty$. Each disclosure induces the posterior $1/2$ with probability $2p$, where p is the current prior, after which disclosure stops. Otherwise it induces the posterior 0, after which the belief is p_B^ℓ. If $n_g^\ell = \infty$, disclosure stops. If $n_g^\ell < \infty$, the next disclosure occurs at prior $r_g^\ell = f^{h(n_g^\ell)}(p_B^\ell) \in I_g$. This point is the same for every episode, so after the first disclosure, each subsequent one ends disclosure for good with the same probability $2r_g^\ell > 0$. Thus, the number

of disclosures is finite almost surely.

A mirror-image argument applies for the judicious policy. ■

E.8 Proof of Proposition 5

It suffices to prove Part (i); Part (ii) follows from a mirror-image argument. By equation (A.3), $p_G^h + p_B^h < 1 \iff p_*^h < 1/2$. By equation (A.2), since $f^{h(i)}(p) \to p_*^h$ for every p, $f^{h(i)}(p) < 1/2$ for some finite i, i.e. $n_g(p) < \infty$. In particular, $n_g^\ell < \infty$, and from any prior $p_1 \in (0,1)$ (an initial prior of 0 or 1 moves into $(0,1)$ after one period), the first disclosure occurs in period $\tau = 1 + n_g(p_1) < \infty$. Consider any period t in which the Sender discloses:

(i) If the disclosure induces action h, the posterior is $1/2$. The following period's prior is $f^h(1/2) < 1/2$ by assumption, so the Sender discloses again in period $t+1$.

(ii) If the disclosure induces action ℓ, the posterior is 0 and $p_{t+1} = p_B^\ell$. By definition of n_g^ℓ, $f^{h(i)}(p_B^\ell) \geq 1/2$ for $i \in \{0, 1, \ldots, n_g^\ell - 1\}$ and $f^{h(n_g^\ell)}(p_B^\ell) < 1/2$. So in periods $t+1, \ldots, t+n_g^\ell$ the Sender discloses nothing and the Receiver takes action h. The Sender discloses again in period $t + n_g^\ell + 1$.

In either case the next disclosure occurs after finitely many periods, and so the Sender discloses in infinitely many periods, on every path. ■

E.9 Proof of Corollary 1

Under negative externalities, $(p_G^\ell + p_B^\ell) - (p_G^h + p_B^h) = -(s_G + s_B) > 0$. So $p_G^h + p_B^h \geq 1$ implies $p_G^\ell + p_B^\ell > 1$, and the claim follows from Propositions 4 and 5. Under full disclosure, $I_{fd} = (0,1)$ and every prior after period 1 is in $\{p_B^\ell, p_G^h\} \subset (0,1)$. ■

E.10 Proof of Corollary 2

Consider the truncated greedy policy; the judicious case is symmetric. As in the proof of Theorem 2(iv), $I = I_{tg} = (\underline{q}, 1/2)$. Since $n_{tg}^\ell < \infty$, the path from z_{tg}^ℓ meets I, so case (ii) of Lemma 19 fails and (i) holds. So $F_{tg}^{(n)}(z_{tg}^\ell) = 1/2 \notin I_{tg}$ and $F_{tg}^{(n+1)}(z_{tg}^\ell) = f^h(1/2) = z_{tg}^h$. Hence $n_{tg}^\ell = n + 1 + n_{tg}^h$

and $r^{\ell}_{tg} = r^{h}_{tg}$. Every disclosure episode after the first therefore starts at $r_{tg} \coloneqq r^{h}_{tg}$, where action h is induced with probability $\alpha_{tg}(r_{tg})$, independently of the past. ■

Appendix F Mixed Externalities

By Definition 1, under positive (negative) externalities, the high action always improves (worsens) the future state compared to the low action, regardless of the current state. What happens when there are positive externalities in one state and negative externalities in the other, i.e. the action that improves the future state depends on the current state? There are two cases to consider: one where $s_G > 0$ and $s_B < 0$, and the other with the inequalities reversed.

Proposition 8. Suppose there are positive externalities in the good state and negative externalities in the bad state, i.e. $s_G > 0$ and $s_B < 0$. Then:

(i) Full disclosure is optimal if and only if $x \leq \bar{x}$,

(ii) The greedy policy is optimal if and only if $x \geq \bar{x}$.

Proof. Since $s_G > 0 > s_B$, we have that $\bar{x} \geq 0$ and $\bar{y} \geq 1$. Part (i) follows from Proposition 7. Part (ii) follows from Proposition 6. ■

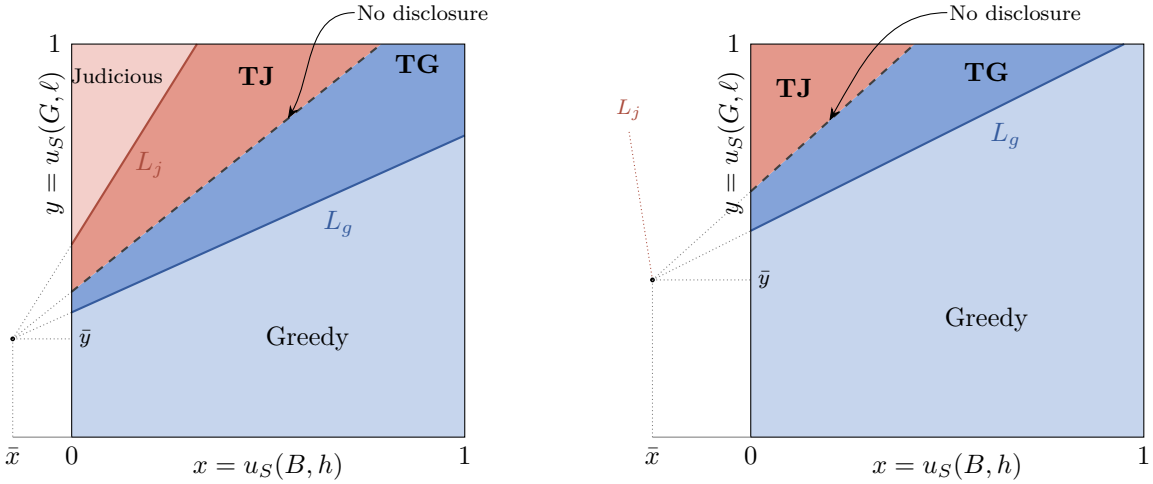
The intuition for Proposition 8 is as follows. When $s_G > 0$ and $s_B < 0$, the action that matches the state is also the one that is better for the future: the high action keeps the good state persistent, and the low action makes the bad state more likely to turn good. In the good state the high action therefore dominates: it gives the Sender her highest stage payoff of $1 \geq y$ and also leads to a better future state. Therefore, if $\delta > 0$, the Sender never wants to induce the low action in the good state, which is why $\bar{y} > 1$, so that the judicious policy is never optimal. All that remains is how often to induce the high action in the bad state. Each such mistake gives the Sender x today but makes the bad state more persistent. Because this trade-off is the same in every period, the solution is bang-bang. If $x \leq \bar{x}$, the current gain does not cover the discounted future loss, and the Sender fully discloses so that the Receiver never makes this mistake. If $x \geq \bar{x}$, she induces the high action as often as possible, which is the greedy policy.

Proposition 9. Suppose there are negative externalities in the good state and positive externalities in the bad state, i.e. $s_G < 0$ and $s_B > 0$. Then:

(i) The greedy policy is optimal if and only if $y \leq \min\{1, L_g(x)\}$,

(ii) The judicious policy is optimal if and only if $y \geq \bar{y}$ and $x \leq \min\{1, L_j(y)\}$,

(iii) Otherwise, there exists a no-disclosure line, with the optimal policy being truncated greedy below it and truncated judicious above it.



(a) With a judicious region

(b) Without a judicious region

Figure 12: Optimal policies when $s_G < 0 < s_B$;
TG: truncated greedy, TJ: truncated judicious

Proof. Since $s_G < 0$ and $s_B > 0$, we have that $\bar{x} \leq 0$ and $\bar{y} \leq 1$, with strict inequalities if $\delta > 0$. By Proposition 7, full disclosure is not optimal for $x > 0$ (and for any x if $\delta > 0$). Part (i) restates Lemma 11; in particular, since $s_G < 0$, $m_g > 0$ as in Theorem 2 (ii). Part (ii) restates Lemma 16. However, since $s_B > 0$ now, it is not necessary that $m_j > 0$. If $m_j \leq 0$, then for any $(x, y) \in [0, 1]^2$, the judicious policy is never optimal (see Figure 12).

For part (iii), we apply the same argument as part (iv) of the proof of Theorem 2 on $\mathbb{R}^2$, as each result it uses holds for all $(x, y) \in \mathbb{R}^2$ (for Claim 1, see footnote 36). On $\mathbb{R}^2$ the edge Γ_j, along which the judicious policy is optimal, always exists, even when it lies outside the unit square (see

Figure 12b). If $m_j \leq 0$, two steps of that argument change. First, $x > L_j(y)$ holds automatically when $x > \bar{x}$ and $y > \bar{y}$, so $\mathcal{R} = \{(x, y) : x > \bar{x}, y > \bar{y}, y > L_g(x)\}$ is non-empty and Lemma 17 is not needed. Second, in Case A the no-disclosure line passes through $(\bar{x}, \bar{y})$ with direction strictly between those of Γ_g and Γ_j, rather than with slope between m_g and $1/m_j$. It therefore need not be increasing, which part (iii) does not claim. Restricting to $[0, 1]^2$ gives the result. ■

When $s_G < 0$ and $s_B > 0$, the opposite is true: in each state, the action that does not match the state is the one that improves the future. Taking the low action in the good state, or the high action in the bad state, raises the probability that the next period's state is good. So Receiver mistakes now pay off twice: once through the current payoff (x or y), and again through a better future state. This is why when $\delta > 0$, full disclosure, which eliminates all mistakes, is never optimal here: $\bar{x} < 0$. For the same reason, the payoff from the greedy policy is always concave. The question is therefore not whether the Sender should induce mistakes but in which state. The greedy policy induces them in the bad state and the judicious policy in the good state. The greedy policy is optimal when the Sender's preference for the high action is strong relative to her preference for the good state, that is, when x is high and y is low. The judicious policy is optimal in the reverse case; unlike under negative externalities, however, the judicious region may be empty. The judicious policy induces the future-improving mistake in the good state, but it does so by taking the low action in the bad state, which is the worse action for the future there. In between, the Sender induces mistakes in both states by keeping the belief close to $1/2$ through the truncated policies.

Appendix G General Model

In this appendix, I adopt the general model of Ely (2017), allowing for intertemporal externalities. The Sender's and Receiver's period-t payoffs depend on a state of the world $\omega_t \in \Omega$ and the Receiver's action $a_t \in \mathcal{A}$ in that period, and are given by $u_S(\omega_t, a_t)$ and $u_R(\omega_t, a_t)$, respectively. The sets Ω and $\mathcal{A}$ are finite, with $|\Omega| \geq 2$ and $|\mathcal{A}| \geq 2$. Since the Receiver is short-lived, he plays myopic best responses that maximise his stage game payoff, while the Sender maximises her long-run discounted expected payoff.

As earlier, the state transitions depend on both the period-t action and the period-t state. Formally, the state transition probabilities are given by $\mathbb{P}(\omega' \mid \omega, a)$, such that $\sum_{\omega' \in \Omega} \mathbb{P}(\omega' \mid \omega, a) = 1$. This includes the following cases:

(i) Absorbing states: If, for a given state ω, $\mathbb{P}(\omega \mid \omega, a) = 1 \ \forall a \in \mathcal{A} \iff \mathbb{P}(\omega' \mid \omega, a) = 0$ $\forall \omega' \neq \omega, \forall a \in \mathcal{A}$, then state ω is absorbing: once the game reaches state ω, it stays there forever. This is the case analysed in the basic model in Ely (2017).

(ii) Action-independent evolution: If, for every distinct pair $\omega, \omega' \in \Omega$, $\mathbb{P}(\omega' \mid \omega, a) = \mathbb{P}(\omega' \mid \omega, a')$, $\forall a, a' \in \mathcal{A}$, then the state transition process is independent of action, and only depends on the state. This is the case explored in Renault, Solan, and Vieille (2017),[51] the general model of Ely (2017), Ashkenazi-Golan et al. (2023),[52] and Lehrer and Shaiderman (2026).

(iii) State-independent evolution: If, for every $a \in \mathcal{A}$, $\mathbb{P}(\omega \mid \omega', a) = \mathbb{P}(\omega \mid \omega'', a)$, $\forall \omega$, ω', $\omega'' \in \Omega$, then the state transition process is independent of state, and only depends on the action.

The players' common prior distribution in period t is represented by a vector $\mathbf{p_t} \in \Delta(\Omega)$. Let $p_t(\omega)$ denote the period-t prior probability of being in state ω, so that $\sum_{\omega \in \Omega} p_t(\omega) = 1$. The Sender chooses an experiment (Σ, σ_t), defined in the same manner as earlier. Using the same arguments as in the basic model, it is straightforward to show that it is without loss of generality to assume that $|\Sigma| \leq |\mathcal{A}|$, with each signal realisation inducing a different action. The players update to their posterior belief $q_t(\omega) := \mathbb{P}(\omega_t = \omega \mid \varsigma_t)$ after observing the signal realisation $\varsigma_t \in \Sigma$ and the chosen experiment, using Bayes' rule.

After the Receiver chooses an action, the state evolves based on the transition probabilities described above. Given the posterior $\mathbf{q_t}$ in period t and the action $a \in \mathcal{A}$ taken in period t, the prior in period $t+1$ is given by:

$$p_{t+1}(\omega') = \sum_{\omega \in \Omega} q_t(\omega) \cdot \mathbb{P}(\omega' \mid \omega, a) \tag{G.1}$$

Thus, the players' prior in a period is a conditional distribution over states given the signal realisations and actions in all periods preceding that period. Updating for every state $\omega \in \Omega$ using

[51] Renault, Solan, and Vieille (2017) analyse a special class of games in which the action space is binary.
[52] Ashkenazi-Golan et al. (2023) analyse a special class of games in which the state space is binary.

equation (G.1), we get the whole vector of period-$t+1$ beliefs. Given $\mathbf{q_t}$ and action $a \in \mathcal{A}$, let $f^a(\mathbf{q_t})$ denote period $t+1$'s prior.

Let the Receiver's best-response function be $a^*(\mathbf{q})$ and let $A^*(\mathbf{q})$ denote his set of best responses. I assume that the Receiver's best response is unique for almost every posterior belief q_t. When $A^*(\mathbf{q})$ has more than one element, he takes the action in $A^*(\mathbf{q})$ that the Sender prefers, taking her continuation payoff into account. Thus, the Sender's interim expected payoffs in the stage game and the next period's prior from inducing a posterior $\mathbf{q_t}$ are given by $\mathbb{E}_{\mathbf{q_t}}(u_S(\omega_t, a))$ and $f^a(\mathbf{q_t})$, respectively, where $a \in A^*(\mathbf{q_t})$.

Let $V(\cdot)$ denote the Sender's maximum payoff from the dynamic optimisation problem. This function is characterised by the solution to the following Bellman equation:

$$V(\mathbf{p_t}) = \sup_{\mu \in \Delta(\Delta(\Omega))} \mathbb{E}_\mu \left\{ \max_{a \in A^*(\mathbf{q_t})} \left[(1-\delta)\mathbb{E}_{\mathbf{q_t}}(u_S(\omega_t, a)) + \delta V(f^a(\mathbf{q_t})) \right] \right\} \quad \text{subject to } \mathbb{E}_\mu(\mathbf{q_t}) = \mathbf{p_t} \tag{G.2}$$

In the functional equation above, the expression in the curly brackets is the Sender's interim expected payoff in the dynamic game. Concavification of these interim payoffs gives us the Sender's optimal expected payoff, so that:

Theorem 3. The functional equation (G.2) characterising the Sender's optimal expected payoff has a unique fixed point and is given by:

$$V(\mathbf{p}) = \text{cav} \left\{ \max_{a \in A^*(\mathbf{q})} \left[(1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega, a)) + \delta V(f^a(\mathbf{q})) \right] \right\}.$$

Proof. The proof generalises that of Ely (2017). In order to prove the theorem, it is sufficient to show that the operator $TV(\mathbf{q}) = \text{cav}\left\{\max_{a \in A^*(\mathbf{q})}[(1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega, a)) + \delta V(f^a(\mathbf{q}))]\right\}$ is a contraction mapping in the space of bounded real-valued functions on $\Delta(\Omega)$ with the sup norm. Blackwell's sufficient conditions[53] are satisfied as follows:

(i) Consider two functions $V(\cdot)$ and $W(\cdot)$. If $V(\mathbf{q}) \geq W(\mathbf{q})$ $\forall \mathbf{q} \in \Delta(\Omega)$, then for every $\mathbf{q} \in \Delta(\Omega)$ and every $a \in A^*(\mathbf{q})$, $(1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega, a)) + \delta V(f^a(\mathbf{q})) \geq (1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega, a)) + \delta W(f^a(\mathbf{q}))$. Taking the maximum over $a \in A^*(\mathbf{q})$ on both sides preserves the inequality. Since the

[53] Theorem 3.3 of Stokey, Lucas, and Prescott (1989).

concavification of a function is the smallest concave function greater than or equal to the said function, this implies that

$$\begin{aligned}&\operatorname{cav}\left\{\max_{a\in A^*(\mathbf{q})}[(1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega,a))+\delta V(f^a(\mathbf{q}))]\right\}\\ \geq &\operatorname{cav}\left\{\max_{a\in A^*(\mathbf{q})}[(1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega,a))+\delta W(f^a(\mathbf{q}))]\right\}, \forall \mathbf{q}\in\Delta(\Omega).\end{aligned}$$

Equivalently, $TV \geq TW$.

(ii) Consider a constant $c \geq 0$. Then, $(V+c)(\mathbf{q}) = V(\mathbf{q}) + c$ which implies that:

$$\begin{aligned}T(V+c)(\mathbf{q}) = &\operatorname{cav}\left\{\max_{a\in A^*(\mathbf{q})}[(1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega,a))+\delta(V+c)(f^a(\mathbf{q}))]\right\}\\ = &\operatorname{cav}\left\{\max_{a\in A^*(\mathbf{q})}[(1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega,a))+\delta V(f^a(\mathbf{q}))]+\delta c\right\}\\ = &\operatorname{cav}\left\{\max_{a\in A^*(\mathbf{q})}[(1-\delta)\mathbb{E}_{\mathbf{q}}(u_S(\omega,a))+\delta V(f^a(\mathbf{q}))]\right\}+\delta c\\ = &\, TV+\delta c,\end{aligned}$$

since adding the constant δc to every term adds δc to the maximum. Thus, the operator T is a contraction (with Lipschitz constant $\delta < 1$, by assumption). By the contraction mapping theorem, the iterates of T converge to its unique fixed point V. ■

Ely (2017) proved a special case of the above theorem with $f(\mathbf{p})$ replacing $f^a(\mathbf{p})$. From the above, it follows that the Sender's value function is concave in the prior.

I use the above to generalise Lemma 1. The Sender's policy $(\mu(\mathbf{p}), \{\mathbf{q}^a(\mathbf{p})\}_{a\in\mathcal{A}})$ specifies the posteriors $\{\mathbf{q}^a\}_{a\in\mathcal{A}}$ and the distribution μ over them that the Sender would choose at each prior $\mathbf{p}$, such that $\mathbb{E}_\mu(\mathbf{q}^a) = \mathbf{p}$. Let $\phi(\mathbf{p})$ be the payoff to the Sender from following such a policy.

Lemma 20. A policy $(\mu(\mathbf{p}), \{\mathbf{q}^a(\mathbf{p})\}_{a\in\mathcal{A}})$ of the Sender is optimal if and only if $\phi(\cdot)$ is concave over $\Delta(\Omega)$ and the following is satisfied:

$$\phi(\mathbf{p}) \geq (1-\delta)\mathbb{E}_{\mathbf{p}}(u_S(\omega,a)) + \delta\phi(f^a(\mathbf{p})), \forall \mathbf{p}\in\Delta(\Omega), a\in A^*(\mathbf{p}) \tag{G.3}$$

The proof follows the same steps as the proof of Lemma 1 and is therefore omitted.