

Dynamic Persuasion with Outside Information*

Jacopo Bizzotto,[†] Jesper Rüdiger[‡] and Adrien Vigier[§]

First draft: August 2015

This version: February 2020

Abstract

A principal seeks to persuade an agent to accept an offer of uncertain value before a deadline expires. The principal can generate information, but exerts no control over exogenous outside information. The combined effect of the deadline and outside information creates incentives for the principal to keep uncertainty high in the first periods so as to persuade the agent close to the deadline. We characterize the equilibrium, compare it to the single-player decision problem in which exogenous outside information is the agent's only source of information, and examine the welfare implications of our analysis.

JEL classification: C72, C73, D83

Keywords: Information Design, Bayesian Persuasion, Dynamic Games

*We thank Jeff Ely, Yingni Guo, Johannes Hörner, Anton Kolotilin, Meg Meyer, Eduardo Perez-Richet, Jerome Renault, Paolo Piacquadio, Christoph Schottmüller, Peter Norman Sørensen and various seminar audiences for helpful comments. Any remaining error is our own.

[†]Oslo Business School, Oslo Metropolitan University. Email: jacopobi@oslomet.no.

[‡]Department of Business Administration, Universidad Carlos III de Madrid, Email: jrudiger@emp.uc3m.es.

[§]Department of Economics, BI Norwegian Business School. Email: a.h.vigier@gmail.com.

1 Introduction

The analysis of persuasion has lately received a lot of attention following Kamenica and Gentzkow (2011). How is the persuasion problem affected by the possibility for the receiver to wait in order to learn from exogenous sources?

Consider a principal who would like an agent to take a particular action. For instance, the principal might be a lawyer aiming to make partner at a law firm. The firm's policy requires that she is either made partner within T years, or fired. Each year until this deadline is reached, the firm can choose between promoting her, firing her, and delaying its final decision in order to better evaluate her merits. Delaying the final decision gives the lawyer more time to build a track record in court. Information generated in court is beyond the lawyer's control to the extent that better lawyers are able to build, on average, a better record. Yet a lawyer is not without agency in this process: she can request from judges outcomes extremely favorable for her customers, and by doing so test her ability to perform, or else she can aim for low-risk, low-reward agreements with the prosecutors, and by doing so make it harder to evaluate her ability.

We develop in this paper a simple model to analyze problems of the kind described above. A principal ("she") and an agent ("he") interact over T periods. The agent has to choose between accepting and rejecting an offer that the principal would like him to accept. Rejecting the offer allows the agent to secure a known positive payoff. The agent's payoff from accepting the offer depends on the other hand on an unknown binary state of the world, such that the agent would like to accept in one state but to reject in the other. The agent can postpone making his final decision until period T , and waiting allows him to accumulate information. The central novelty of our model is to allow *inside* as well as *outside* information:

- at any point in time, the principal can generate information by designing the experiments of her choice (inside information);
- in addition, exogenous signals over which the principal exerts no control are observed over time (outside information).

Both players discount time, neither can commit to future decisions, and all information is symmetric.

Although in this kind of environment waiting is socially wasteful, we show that the agent sometimes waits T periods before making his final decision.¹ The basic tension is as fol-

¹We show that waiting is socially wasteful even if the agent's payoff from a type II error is negative.

lows. Each period t , information that the agent expects to obtain –from future signals *and* experiments– determines a cutoff belief above which the agent chooses to accept; we call it the agent’s period- t threshold of acceptance.² The closer the deadline the smaller the amount of useful information the agent expects to obtain in the future. Hence, the threshold of acceptance decreases over time. This in turn creates incentives for the principal to keep uncertainty high in the first periods with a view to persuading the agent closer to the deadline.³ We call this mechanism the *deadline effect*. Yet, in order to try persuading the agent close to the deadline the principal must let the agent observe exogenous signals. The caveat is that favorable signal realizations could lead the belief to “overshoot” the agent’s period-1 threshold of acceptance. The greater the overshooting the fewer mistakes the principal can induce the agent to make in periods ahead, thus incentivizing the principal to try persuading the agent in the first period. We call this mechanism the *overshooting effect*. Which one of the deadline and overshooting effects dominates the other pins down the principal’s choice of experiment in the first period and, via this choice, whether the agent sometimes waits before making his final decision.

We examine two types of signals: “perfect good news” and “perfect bad news”. Under perfect good news (respectively, bad news) the state in which accepting the offer is optimal (resp. suboptimal) for the agent is perfectly revealed with positive probability each period. With perfect good news, or if the deadline is sufficiently far into the future, the overshooting effect is the dominant force. The principal then generates a sufficient amount of information in the first period to induce the agent to make an immediate final decision. However, under perfect bad news, if signal accuracy is intermediate and the deadline is not too far into the future, the deadline effect then becomes the dominant force. In this case the agent might wait up to T periods before making his final decision.

Welfare hinges on the equilibrium strategy of the principal. In one regime, the principal generates information so as to induce the agent to act in the first period; we say in this case that the principal is aggressive. In the other regime, the principal generates less information, and in the first periods seeks to sustain uncertainty so as to persuade the agent closer to the deadline; we say in this case that the principal is conservative. Pareto efficiency obtains if and only if the principal is aggressive. Furthermore, as long as no regime switch occurs, the

²The “belief” refers to the probability attached to the state in which accepting the offer is optimal for the agent.

³The lower the threshold of acceptance, the more errors of type II the principal can induce the agent to make.

agent's equilibrium expected payoff as well as the quality of the final decision are monotonically increasing in the amount of exogenous outside information, be it in the form of more accurate outside information, or a deadline further away in time (allowing the agent to observe more exogenous signals). However, any regime switch from aggressive to conservative causes the agent's welfare and the quality of the final decision to drop, and vice versa.

Our analysis reveals a rich interplay between inside information and exogenous outside information, that contrasts sharply with settings in which exogenous outside information is the agent's only source of information (as in Wald (1947), for example). For instance, in our setting, extending the deadline can accelerate the agent's final decision. The reason is that pushing the deadline further away in time increases the amount of information generated by the principal.

The rest of the paper is organized as follows. The related literature is discussed below. The model is presented in Section 2. The core of the analysis is in Section 3. Several extensions of the model are examined in Section 4. Section 5 concludes.

Related Literature. We contribute to the literature on Bayesian persuasion by introducing outside information in the canonical framework of Kamenica and Gentzkow (2011), that is, by relaxing the assumption that the sender (or principal) fully controls the flow of information to the receiver (or agent). This approach connects our work to two active strands of research.

A first strand of research examines the case in which multiple senders compete to persuade the agent. This includes Gentzkow and Kamenica (2016), Li and Norman (2018) and Board and Lu (2018). The models and applications are different from ours: we study situations in which a single principal designs multiple experiments over time whereas these papers examine situations in which multiple principals design one experiment each.

The second strand of research focuses like we do on the dynamic persuasion of an agent, and begins with Au (2015) and Honryo (2018). The contemporaneous work of Orlov, Skrzypacz and Zryumov (2020) is the study most related to ours. In their model, an evolving state affects the principal's and the agent's payoff from exercising an option. This process is exogenous, and creates an incentive for the agent to wait. However, the payoffs also depend on a second state. The principal controls the flow of information concerning the second state, but the evolution of the first state is publicly observable. The environment is stationary; in particular, there is no deadline by which the agent must act. In our model, in the absence of a deadline, the agent's threshold of acceptance is the same in all periods. This means, in turn, that the principal generates information inducing the agent to make a final decision in the first period.

In equilibrium, the agent therefore never waits. In Orlov *et al.* (2020), by contrast, waiting can be socially optimal, since the principal is unable to generate information about one of the two states. The key tradeoffs in the two papers are thus different. While several other papers examine the dynamic persuasion of an agent, including Henry and Ottaviani (2019), Che, Kim and Mierendorff (2020), Ely and Szydlowski (2020), Smolin (2020), and Zhao, Renou and Tomala (2020), the focus in all of them is different from ours since in these models the principal fully controls the flow of information.

A few additional papers are related to specific aspects of our work. Our finding that the agent’s equilibrium expected payoff is a non-monotonic function of signal accuracy is linked in spirit to a related result in Kolotilin (2018). Gratton, Holden and Kolotilin (2017) examine the problem of a principal deciding when to start a public flow of information about her type and is one of very few papers which, like us, analyze the role of deadlines in contexts of persuasion. The decision problem of the agent naturally links our analysis to the literature on experimentation starting with Rothschild (1974), Bolton and Harris (1999), and Keller, Rady and Cripps (2005). However, whereas we study the interplay of inside and outside information, there is no inside information in that literature.

2 Model

A principal (“she”) and an agent (“he”) interact over $T \geq 2$ periods. We refer to the final period as the deadline of our game. The agent has to choose between accepting and rejecting an offer that the principal would like him to accept. By rejecting, the agent secures a (undiscounted) payoff $V_R > 0$; accepting yields him V_ω , where $\omega \in \{G, B\}$ represents an unknown state of the world. To make the model interesting, $V_G > V_R > V_B$. In order to learn about the realized state, the agent can postpone making his *final decision* (accept or reject) until $t = T$. Both outside and inside information is observed over time: the former is exogenous whereas the latter is strategically generated by the principal. All information being public, the players share common beliefs about the state. The (evolving) probability assigned to $\omega = G$ will be referred to as *the belief*. The payoff of the principal is 0 in case of rejection, and is normalized to 1 in case of acceptance. Both players discount time at rate $\delta \in (0, 1)$.⁴

⁴Section 4.2 examines the case in which players discount time differently.

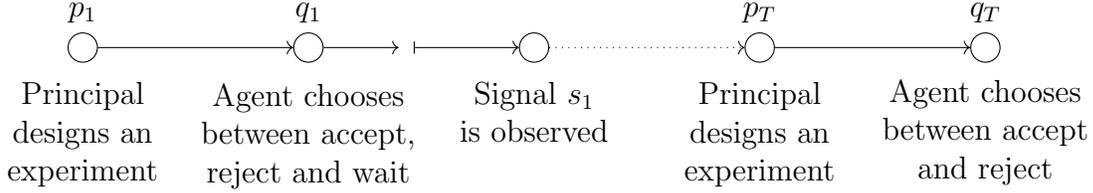


FIGURE 1

Timing. The state of the world is drawn by nature according to $\mathbb{P}(\omega = G) = p_1$. We suppose for expository purposes that the agent initially leans towards rejection, that is, $p_1 \in (0, \underline{b})$, where $\underline{b} := \frac{V_R - V_B}{V_G - V_B}$ denotes the belief at which the agent is indifferent between accepting and rejecting. The principal designs an experiment inducing the end-of-period-1 belief q_1 . The agent then chooses between *accept*, *reject* and *wait*. If the agent makes a final decision, payoffs are realized; if the agent waits, the exogenous signal s_1 is observed, inducing the beginning-of-period-2 belief p_2 . This sequence repeats until the agent makes a final decision, with the caveat that, at $t = T$, the agent has to make a final decision. Figure 1 summarizes the timeline; the broken arrow between the second and third node indicates that the game may terminate at the second node.

Inside Information. The principal's experiment in period t is a probability distribution $\tau_t \in \Delta([0, 1])$ governing the end-of-period- t belief q_t ; the only constraint imposed on each experiment is Bayes plausibility: $\mathbb{E}_{\tau_t}[q_t] = p_t$. The support M_t of τ_t therefore uniquely determines this experiment as long as $|M_t| \leq 2$. It will be convenient, whenever possible, to use M_t in order to represent τ_t .

Outside Information. The signal in period t , denoted s_t , is drawn from the conditional probability distribution $\pi(\cdot | \omega)$ over $\{g, b\}$. The signal-generating process is assumed i.i.d. across time periods. As is common in the literature on strategic experimentation,⁵ we focus for tractability on conclusive signal-generating processes. Under *perfect bad news*, $\pi(b | B) = \gamma$ and $\pi(g | G) = 1$. In this case $s_t = b$ informs players that $\omega = B$, whereas the belief drifts upwards as long as $s_t = g$. By contrast, under *perfect good news*, $\pi(b | B) = 1$ and $\pi(g | G) = \gamma$. The signal realization g then informs players that $\omega = G$, whereas the belief drifts downwards as long as $s_t = b$. The parameter $\gamma \in [0, 1]$ capturing the informativeness of the signal-generating process will be referred to as the signal accuracy.

⁵See for instance Keller and Rady (2015), Bonatti and Hörner (2017), or Frick and Ishii (2018).

Strategies and Equilibrium. A t -history consists of experiments, end-of-period beliefs and signal realizations for the first $t-1$ periods, that is $\{\tau_k, q_k, s_k\}_{k=1}^{t-1}$; an augmented t -history contains in addition the experiment τ_t and the belief q_t . A strategy for the principal maps each t -history to an experiment τ_t . A strategy for the agent maps each augmented t -history to a decision in $\{\text{accept, reject, wait}\}$ for $t < T$, and to a decision in $\{\text{accept, reject}\}$ for $t = T$. The equilibrium concept is Perfect Bayesian Equilibrium (PBE): the player at each decision node maximizes her/his expected payoff conditional on (a) the other player's strategy and (b) the belief obtained using Bayes' rule.

3 Analysis

Subsection 3.1 characterizes the equilibrium of our game. A general discussion of the main theorem is provided in Subsection 3.2, and a sketch of its proof is presented in Subsection 3.3. All omitted proofs of this section are in Appendices A, B and C.

3.1 Main Result

As usual in models of Bayesian persuasion, equilibrium multiplicity arises from the fact that, for a subset of beliefs, several experiments ultimately induce identical outcomes. We thus focus throughout the paper on PBE such that: (i) whenever the principal is indifferent between experiments ordered according to Blackwell's criterion, she chooses the least informative experiment; (ii) whenever indifferent, the agent makes the decision preferred by the principal. These refinements simplify the exposition, but are inessential for our results. The first deals with the kind of multiplicity mentioned above;⁶ the second rules out inconsequential multiplicity off the equilibrium path.⁷ Henceforth, PBE satisfying (i) and (ii) will be referred to as equilibria for short.

Proposition 1. *There exists a unique equilibrium.*

⁶For instance, imagine that in a given period the agent accepts for q_t in an interval $[x, y]$. Then, for $p_t \in (x, y)$, the principal is indifferent between designing the uninformative experiment or $M_t = \{x, y\}$. In this case, we assume that the principal chooses the uninformative experiment.

⁷For instance, imagine that, in a given period, at $q_t = x$ the agent is indifferent between rejecting and waiting, but that irrespective of whether the agent does one or the other, any period- t experiment with x in its support is strictly dominated for the principal by some other experiment. Then what the agent does at $q_t = x$ is inconsequential, as $q_t = x$ never occurs on the equilibrium path. In this case, we assume that the agent waits at $q_t = x$.

We henceforth refer to the threshold \underline{b} at which the agent is indifferent between accepting and rejecting as the *static threshold of acceptance*. Note that \underline{b} is independent of the signal-generating process. At $t=T$, the agent accepts if $q_T \geq \underline{b}$ and rejects otherwise. At $t < T$ however, information which the agent expects to obtain in periods ahead (from the experiments and from the signals) determines an interval of beliefs at which the agent chooses to wait.

Lemma 1. *Each period, cutoffs $0 < a_t \leq b_t < 1$ exist such that, in equilibrium the agent rejects if $q_t < a_t$, waits if $q_t \in [a_t, b_t)$, and accepts if $q_t \geq b_t$.*

We henceforth refer to the cutoff b_t as the agent's *period- t threshold of acceptance*. As information which the agent expects to obtain can only increase his incentive to wait, $b_t \geq \underline{b}$ regardless of the period t . One shows more generally that the agent's threshold of acceptance decreases with t .

Lemma 2. *The agent's period- t threshold of acceptance b_t decreases with t .*

We turn next to the principal. Each period the principal can either try to persuade the agent immediately or aim to keep uncertainty high (i.e. aim for $q_t \in [a_t, b_t)$) so as to try persuading the agent in a future period. The optimal choice of the principal is illustrated in Figure 2. In both panels the gray solid curve represents the principal's equilibrium continuation payoffs given the end-of-period- t belief q_t .⁸ The black dashed curve depicts the concavification of the former curve (Aumann, Maschler and Stearns, 1995), and represents the principal's equilibrium continuation payoffs given the beginning-of-period- t belief p_t . The case in which the principal optimally tries to persuade the agent in period t is depicted in Panel I. In this case,

$$M_t = \begin{cases} \{0, b_t\} & \text{if } p_t \in (0, b_t); \\ \{p_t\} & \text{otherwise,} \end{cases}$$

and we say that the principal is *aggressive* in period t . The case in which the principal optimally keeps uncertainty high so as to try persuading the agent in a future period is depicted in Panel II. In this case $a_t < b_t$ and

$$M_t = \begin{cases} \{0, a_t\} & \text{if } p_t \in (0, a_t); \\ \{a_t, b_t\} & \text{if } p_t \in (a_t, b_t); \\ \{p_t\} & \text{otherwise.} \end{cases}$$

⁸If $q_t \geq b_t$ (respectively $q_t < a_t$) the agent accepts (resp. rejects), yielding payoff 1 (resp. 0) to the principal; if $q_t \in [a_t, b_t)$ the agent waits, yielding a payoff in $(0, 1)$ determined by the continuation strategies.

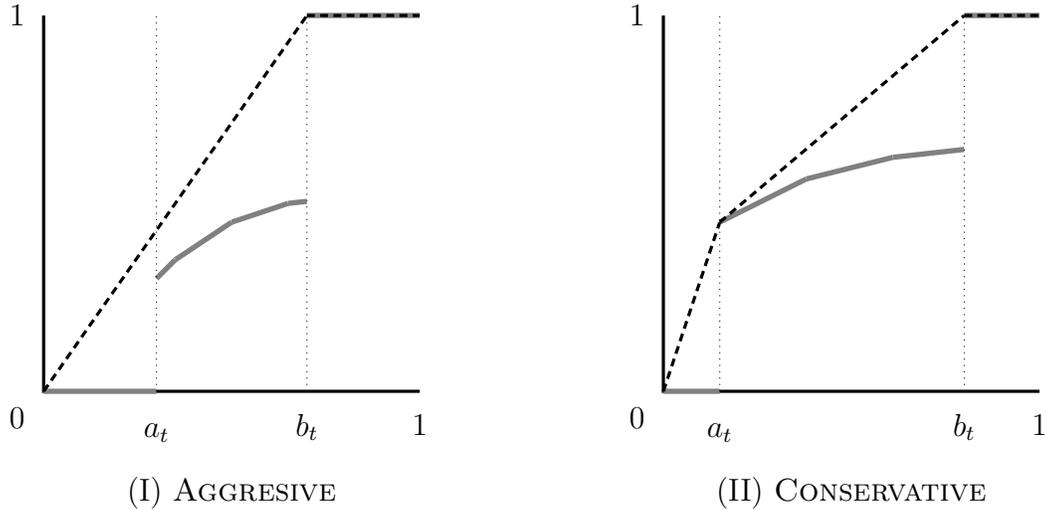


FIGURE 2: PRINCIPAL'S PERSUASION STRATEGY

We then say that the principal is *conservative* in period t . The experiments described in the previous paragraph are the only experiments ever designed by the principal in equilibrium.

Lemma 3. *Each period, in equilibrium, either the principal is aggressive, or the principal is conservative.*

The following theorem is the central result of our analysis.

Theorem 1. *In equilibrium, with perfect good news the principal is aggressive at $t=1$. However, with perfect bad news the principal is aggressive at $t=1$ if and only if T is strictly greater than some threshold $\tilde{T}(\gamma, \delta)$. The set of parameters γ and δ such that, for T sufficiently small, the principal is conservative at $t=1$ (i.e. such that $\tilde{T}(\gamma, \delta) \geq 2$) is given by*

- $\delta > \underline{\delta}$, where $\underline{\delta} \in (0, 1)$,
- and $\gamma \in (\underline{\gamma}(\delta), \bar{\gamma}(\delta))$, where $0 < \underline{\gamma}(\delta) < \bar{\gamma}(\delta) < 1$.

The perfect bad news case is illustrated in Figure 3.⁹ For example, each parameter pair (δ, γ) that belongs to the vertically dashed region of the figure is such that in equilibrium the principal is conservative at $t=1$ if either $T=2$ or $T=3$, whereas the principal is aggressive whenever $T \geq 4$. The rest of this section is organized as follows. We discuss below the key

⁹The figure is drawn for $V_G=2$, $V_R=1$ and $V_B=0$. The code is available on the authors' website.

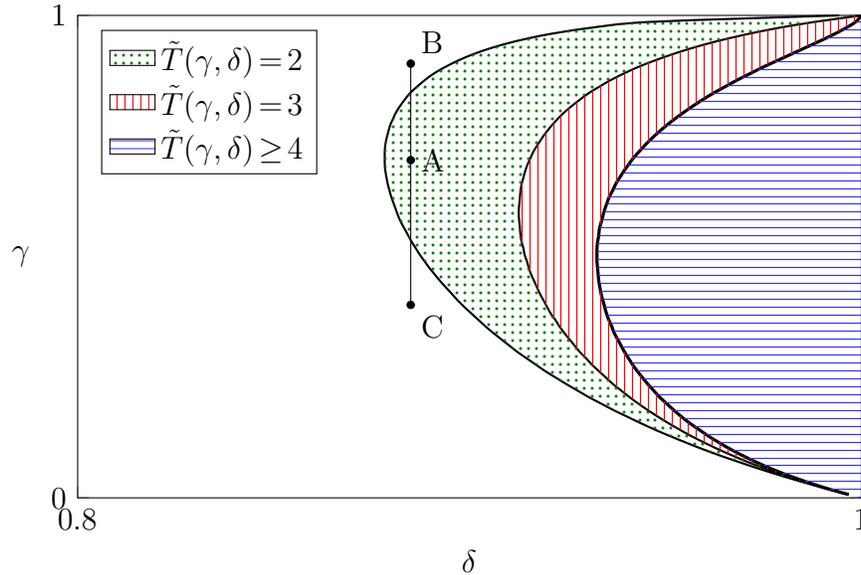


FIGURE 3: THE PERFECT BAD NEWS CASE

tension at the heart of our model and how this tension explains Theorem 1. In Subsection 3.2, we link Theorem 1 to the welfare properties of the equilibrium. We also examine the impact of information supplied by the principal’s experiments, by contrasting our model and results with the benchmark setting in which exogenous outside information is the agent’s only source of information. A sketch of the Proof of Theorem 1 is provided in the final subsection.

Information which the agent expects to obtain determines each period the agent’s threshold of acceptance b_t . The lower b_t , the more mistakes the principal can induce the agent to make. Thus, if $b_1 > b_T$, the principal is incentivized to maintain enough uncertainty in the first periods in order to try persuading the agent at $t = T$. We refer to this as the *deadline effect*. The caveat is the following: to persuade the agent in period T , the principal must let the agent observe $T - 1$ exogenous signals. However, if g signal realizations are sufficiently conclusive, letting the agent observe $T - 1$ exogenous signals may lead the belief to “overshoot” the agent’s period-1 threshold of acceptance, as illustrated in Figure 4.¹⁰ We refer to this as the *overshooting effect*. The greater the overshooting the fewer (future) mistakes the principal can induce the agent to make. For the principal, the overshooting effect thus creates countervailing incentives relative to the deadline effect.

¹⁰In Panel I, to prevent the overshooting, the principal must generate information inducing the agent to accept in the first period with positive probability. In Panel II, the principal has more freedom, but to avert overshooting the principal must generate information inducing the agent to accept with positive probability before $t = T$.

Which effect dominates the other pins down the principal's choice of experiment at $t = 1$.¹¹ Roughly, with perfect good news the overshooting effect dominates the deadline effect because $s_1 = g$ then induces $p_2 = 1$.¹² By contrast, with perfect bad news, the deadline effect can dominate the overshooting effect (provided γ and T are sufficiently small, so as to avoid the scenarios illustrated in, respectively, panel I and panel II of Figure 4). To see that neither γ nor δ can be too small for this mechanism to work, observe that $\gamma \approx 0$ and $\delta \approx 0$ both imply $b_1 \approx \underline{b} = b_T$, in which case the deadline effect becomes vanishingly small: for small γ , this is because the agent does not expect to obtain much information by waiting; for small δ , this is because the agent does not value future information much.

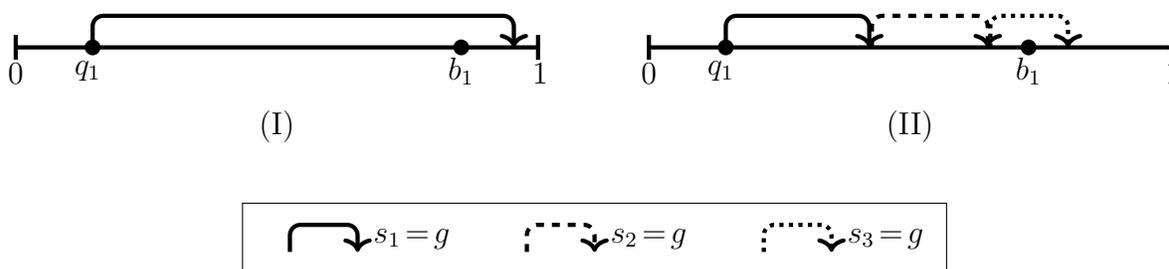


FIGURE 4: THE OVERSHOOTING EFFECT

3.2 Discussion

Pareto efficiency. When is the equilibrium Pareto efficient, and when is it not? In our model, Pareto efficiency obtains if and only if the agent (a) accepts with probability 1 conditional on state G and (b) makes his final decision at $t = 1$ with probability 1.¹³ By Lemmata 1 and 3, condition (a) is always satisfied in equilibrium.¹⁴ The key question then is whether in equilibrium the agent's final decision occurs at $t = 1$ with probability 1. Notice that if in equilibrium the principal is aggressive at $t = 1$ then $M_1 = \{0, b_1\}$, and so (b) holds in this case (by Lemma 1). On the other hand, if at $t = 1$ the principal is conservative in equilibrium then $a_1 < b_1$ and either $M_1 = \{0, a_1\}$ or $M_1 = \{a_1, b_1\}$. Either way, $q_1 = a_1$ with positive probability,

¹¹To be sure, the principal's time discounting provides her with an additional incentive to try persuading the agent immediately. However, the overshooting effect alone can provide sufficient incentives for the principal to be aggressive at $t = 1$.

¹²Moreover, we show in Appendix B, Proposition B.1, that the deadline effect is weaker in the perfect good news case than in the perfect bad news case, in the sense that the difference $b_1 - b_T$ is smaller in the former case than in the latter.

¹³See Proposition C.1 in Appendix C.

¹⁴Lemma 1 ensures that in equilibrium the agent rejects in period t if and only if $q_t \in [0, a_t)$. Lemma 3 ensures that in equilibrium $q_t \notin (0, a_t)$. Thus if the agent rejects in period t , it must be the case that $q_t = 0$.

and so in this case (b) does not hold. We conclude from Lemma 3 and the previous remarks that, in equilibrium, Pareto efficiency obtains if and only if the principal is aggressive at $t=1$. Theorem 1 thus pins down the conditions under which the equilibrium satisfies Pareto efficiency.

Comparison with the single-player setting. Our main theorem also offers interesting contrasts with the corresponding single-player setting in which exogenous signals are the agent's only source of information (as in Wald (1947)). First, in the single-player setting, if the agent waits given a certain amount of outside information then the agent also waits for all greater amounts of outside information (greater T , greater γ , or both). In our model on the other hand, pushing the deadline further away in time can increase the amount of information generated by the principal, and thereby cause the agent to make his final decision earlier on. By the same token, increasing γ may accelerate the agent's final decision with probability 1. For example, at point A in Figure 3, in equilibrium the agent sometimes waits if $T=2$, but never waits if $T=3$; similarly, if $T=2$, the agent sometimes waits at point A but never waits at point B, albeit $\gamma_B > \gamma_A$ (and $\delta_B = \delta_A$).

Second, whereas in the single-player setting increasing the amount of outside information always improves the expected quality of the agent's final decision and raises the agent's expected payoff, in our model increasing γ may increase the probability of type II errors made and lower the agent's expected payoff.¹⁵ The reason is as follows. By switching from aggressive to conservative, the principal causes delay in the agent's decision to accept. Since the principal discounts time, she must then be compensated by a higher probability of acceptance. But we earlier pointed out that in equilibrium the agent accepts with probability 1 conditional on state G . So the higher probability of acceptance must be coming from state B . In consequence, any change of parameters leading the principal to switch from aggressive to conservative can induce a higher probability of type II errors and a reduction of the agent's expected payoff. In Figure 3 for example, in equilibrium the expected quality of the agent's final decision and the agent's expected payoff are higher at point C than at point A , though $\gamma_A > \gamma_C$ (and $\delta_A = \delta_C$).

¹⁵See Propositions C.2 and C.3 in Appendix C.

3.3 Sketch of the Proof of Theorem 1

We present here the main steps of the proof of Theorem 1. Readers uninterested in the technical details may skip this subsection.

In equilibrium the principal ensures that the agent makes no type I error. However, the principal would like to maximize the number of type II errors. The first part of Theorem 1 is founded upon the observation that with perfect good news, in equilibrium, making the agent wait induces him to base his final decision on more information (in Blackwell's sense) than if the principal were aggressive and optimally triggered the agent's final decision at $t=1$. The principal therefore chooses to be aggressive in the first period.

Lemma 4. *With perfect good news, in equilibrium the principal is aggressive at $t=1$.*

Proof: Consider a period $t' < T$ such that, in equilibrium, the principal is aggressive in period $t'+1$. Notice that the latter requirement is satisfied if $t'=T-1$. Observe as well that given $q_t = b_t$, the belief $p_{t'+1}$ induced by $s_{t'} = b$ has to be strictly smaller than $b_{t'+1}$; if this were not the case then, by Lemma 1, at $q_{t'} = b_{t'}$ the agent would prefer accepting to waiting, contradicting the definition of $b_{t'}$. Straightforward algebra then establishes

$$b_{t'} < \frac{b_{t'+1}}{1 - \gamma(1 - b_{t'+1})}. \quad (1)$$

Next, we claim that for all $z \in [a_{t'}, b_{t'})$, in equilibrium, given $p_{t'} = z$ the principal is strictly better off designing the experiment $M_{t'} = \{0, b_{t'}\}$ than the uninformative experiment. This in turn will imply that, in equilibrium, the principal is aggressive in period t' and, by induction, also at $t=1$.

We now prove the previous claim. Let X denote the random variable representing the belief at which the agent makes his final decision given $p_{t'} = z$ and the equilibrium strategies in the continuation game, assuming that the principal designs the experiment $M_{t'} = \{0, b_{t'}\}$. Let Y denote the corresponding random variable assuming that the principal designs the uninformative experiment. One shows, using (1), that Y is a mean-preserving spread of X .¹⁶ Let ϕ denote the piecewise linear function with a kink at $\min\{b_{t'}, b_{t'+1}\}$ such that $\phi(0) = 0$ and $\phi(\min\{b_{t'}, b_{t'+1}\}) = \phi(1) = 1$. Given the equilibrium strategies in the continuation game, the principal's expected payoff from designing the experiment $M_{t'} = \{0, b_{t'}\}$ can be written as

¹⁶Since $\text{supp}(X) = \{0, b_{t'}\}$ and $\text{supp}(Y) = \{0, b_{t'+1}, 1\}$ we only need to show that $\mathbb{P}(X=0) < \mathbb{P}(Y=0)$. This inequality can be written $1 - \frac{z}{b_{t'}} < [(1-z) + z(1-\gamma)] \left[1 - \frac{1}{b_{t'+1}} \cdot \frac{z(1-\gamma)}{(1-z) + z(1-\gamma)} \right]$, which is equivalent to (1).

$\mathbb{E}[\phi(X)]$. On the other hand, as $\delta < 1$, her expected payoff from designing the uninformative experiment is bounded from above by $\mathbb{E}[\phi(Y)]$. Since ϕ is concave, $\mathbb{E}[\phi(Y)] \leq \mathbb{E}[\phi(X)]$. This concludes the proof of the claim which, in turn, by the arguments laid out in the second paragraph, concludes the proof of the theorem. ■

Lemma 4 establishes the first part of Theorem 1. In the rest of this subsection, the focus is on the perfect bad news case. We start by showing that information generated by the principal is such that, at $q_t = b_t$, any benefit accruing to the agent from waiting must come from information generated by the following period's exogenous signal.¹⁷ Therefore, the agent's standard of acceptance is the same at all $t < T$. This, in turn, implies (by Lemma 2) that either the agent's threshold of acceptance equals the static threshold of acceptance in all periods, or the agent exhibits two thresholds of acceptance: a high threshold of acceptance prior to the deadline, which then drops to the static threshold of acceptance at $t = T$.

Lemma 5. *With perfect bad news, either $a_t = b_t = \underline{b}$ in every period, or for all $t < T$: $b_t = b_{T-1} > b_T = \underline{b}$.*

Lemma 5 shows that the agent's threshold of acceptance is the same at all $t < T$. Then, suppose that in some period $t < T - 1$ the principal knows that she will try to persuade the agent in period $t + 1$ (i.e. she will be aggressive in period $t + 1$). Since next period's threshold of acceptance is the same as this period's threshold of acceptance, in equilibrium the principal has to try persuading the agent this period. We therefore obtain the following result.

Lemma 6. *Let $t < T - 1$. With perfect bad news, in equilibrium if the principal is aggressive in period $t + 1$, then the principal is also aggressive in period t .*

We infer from Lemma 6 that if in a game of given length, in equilibrium the principal is aggressive in period 1, then the same must be true in all longer games. Building on Lemma 5 enables us to show in addition that, for sufficiently long games, the principal has to be aggressive in period 1. We thus obtain the following result.

Lemma 7. *There exists $\tilde{T}(\gamma, \delta) < \infty$ such that, in equilibrium, the principal is aggressive in period 1 if and only if $T > \tilde{T}(\gamma, \delta)$.*

We next record the conditions under which, in equilibrium, the principal is aggressive in period $T - 1$.¹⁸

¹⁷The qualification "at $q_t = b_t$ " is essential here. At more pessimistic beliefs, the agent usually strictly benefits from information generated by the principal's experiments and signals two or more periods ahead.

¹⁸Note that, by Lemma 4, $\tilde{T}(\gamma, \delta) = 1$ in the perfect good news case.

Lemma 8. *There exist a cutoff $\underline{\delta} \in (0, 1)$ and, for $\delta > \underline{\delta}$, functions $0 < \underline{\gamma}(\delta) < \bar{\gamma}(\delta) < 1$ such that, with perfect bad news, in equilibrium the principal is conservative in period $T - 1$ if and only if $\delta > \underline{\delta}$ and $\gamma \in (\underline{\gamma}(\delta), \bar{\gamma}(\delta))$.*

Lemmata 7 and 8 together yield the second half of Theorem 1.

4 Extensions

4.1 Frequent Signals

Our framework is founded upon the assumption that exogenous signals are observed at discrete points in time. This assumption is not without loss of generality. In our setting, to observe any signal the agent must incur the cost of waiting a discrete amount of time. This, in turn, assures that the agent's threshold of acceptance is always strictly below 1 (no matter the signal accuracy). To take advantage of this wedge, for γ close to 1, the principal chooses to be aggressive, thereby inducing the agent to make his final decision in the first period.

In our discrete time setting, a natural question is to inquire about the impact of the frequency at which exogenous outside information is observed. In this subsection we recast our model by letting $\Delta_n = \frac{1}{2^{n-1}}$ capture the period length, and refer to $n \in \mathbb{N}^*$ as the signal frequency; T_n will denote the total number of periods until the deadline. The game length (in units of time) is thus $L := T_n \Delta_n$. The signal-generating process is such that $\pi(b|B) = 1 - e^{-\lambda \Delta_n}$, where $\lambda \geq 0$, and $\pi(g|G) = 1$ (we focus on the perfect bad news case; it is easy to show that with perfect good news the principal is aggressive at $t = 1$ regardless of the signal frequency). The per-period discount factor is $e^{-r \Delta_n}$, where $r > 0$.

Keeping n fixed, the baseline model (Section 2) is obtained by setting $\gamma_n = 1 - e^{-\lambda \Delta_n}$, and $\delta_n = e^{-r \Delta_n}$. Relabelling appropriately, the analysis in Section 3 shows that, irrespective of the signal frequency n , in equilibrium: the agent's final decision is made in period 1 with probability 1 if and only if one of the following conditions holds: the game is sufficiently long ($L_n > \tilde{L}_n$), signals are sufficiently inaccurate ($\lambda_n < \underline{\lambda}_n$), signals are sufficiently accurate ($\lambda_n > \underline{\lambda}_n$), players are sufficiently impatient ($r_n > \underline{r}_n$).

However, a question arises regarding the model's behavior in the limit as n tends to infinity, since both γ_n and δ_n then tend to 0. The first effect pushes the principal to be aggressive at $t = 1$ (Theorem 1), while the second effect pushes the principal to be conservative (Lemma

C.3).¹⁹ The question then is whether the dichotomy between aggressive and conservative regimes that our analysis uncovered continues to exist at very high frequency: namely, if $\lim_{n \rightarrow \infty} \tilde{L}_n = 0$ (respectively, $\lim_{n \rightarrow \infty} \tilde{L}_n = \infty$) then at very high frequency the principal is aggressive (respectively, conservative) at $t=1$ irrespective of the game length.

We show in the next proposition (proved in Online Appendix 1) that, provided the signals are sufficiently informative, the aforementioned dichotomy continues to exist at very high frequency. A sufficient condition is $\lambda > \varphi(r)$, where

$$\varphi(r) := \ln \left(\frac{e^{-r}(V_G - V_R)(V_R - V_B)}{e^{-r}[V_R(2V_B - V_R) - V_G V_B] - V_R(V_G - V_B)} \right).$$

This condition is equivalent to requiring sufficiently informative signals for the agent to prefer waiting at low signal frequency ($n=1$) when $p_1 = \underline{b}$.

Proposition 2. *With perfect good news, the agent's final decision is made in period 1 with probability 1. Suppose $\lambda > \varphi(r)$. Then, there exist N and $0 < \tilde{L} < L^\dagger \leq \infty$ such that, with perfect bad news, for all $n > N$:*

- (i) *the agent's final decision is made in period 1 with probability 1 if $L > L^\dagger$;*
- (ii) *the agent's final decision is made in period 1 with probability strictly less than 1 if $L < \tilde{L}$.*

4.2 Different Discount Factors

Here we let the discount rates of the two players differ. We denote the agent's discount rate $\delta_A \in (0, 1)$, and the principal's discount rate $\delta_P \in (0, 1)$. Our baseline model corresponds to $\delta_P = \delta_A$. The findings listed in Theorem 1 hold qualitatively unchanged with different discount factors, as recorded in the following proposition.

Proposition 3. *In equilibrium, with perfect good news the principal is aggressive at $t=1$. However, with perfect bad news the principal is aggressive at $t=1$ if and only if T is strictly greater than some threshold $\tilde{T}(\gamma, \delta_A, \delta_P) \geq 1$. In particular, the principal is conservative for some T (that is, $\tilde{T}(\gamma, \delta_A, \delta_P) \geq 2$) if and only if:*

- *signal accuracy is intermediate, that is, $\gamma \in (\underline{\gamma}(\delta_A, \delta_P), \bar{\gamma}(\delta_A, \delta_P))$,*
- *at least one player is sufficiently impatient, that is, $\delta_A < \underline{\delta}_A$, or $\delta_P < \underline{\delta}_P(\delta_A)$,*

¹⁹Keeping L fixed, increasing n also increases T_n of course.

where $\underline{\delta}_A \in (0, 1)$, $\underline{\delta}_P(\delta_A) \in (0, 1]$ and $0 < \underline{\gamma}(\delta_A, \delta_P) < \bar{\gamma}(\delta_A, \delta_P) < 1$ for $\delta_A > \underline{\delta}_A$ and $\delta_P > \underline{\delta}_P(\delta_A)$.

The proof is in Online Appendix 2. In our baseline model, whenever players are sufficiently impatient the principal is aggressive at $t = 1$. Proposition 3 shows that for this result to hold, it is enough that one of the players be sufficiently impatient. If the principal is sufficiently impatient then she is aggressive regardless of the period-1 threshold of acceptance b_1 . If instead the agent is sufficiently impatient, then he does not wait, regardless of the information generated by the principal's experiment at $t = 1$ ($a_1 = b_1$). This, in turn, results in the principal being aggressive at $t = 1$.

4.3 Costly Experiments

In this section we extend the baseline model by assuming that the principal incurs a cost $C > 0$ for each new experiment. With costly experiments, the principal's payoff (expressed in period-1 units) from acceptance in period t can be written as

$$\delta^{t-1} - \sum_{k=1}^t \delta^{k-1} C = \delta^{t-1} \left(1 + \frac{C\delta}{1-\delta} \right) - \frac{C}{1-\delta}.$$

Similarly, the principal's payoff from rejection in period t becomes

$$\delta^{t-1} \left(\frac{C\delta}{1-\delta} \right) - \frac{C}{1-\delta}.$$

The game with costly experiments may thus be viewed as a modified version of the baseline model in which the principal's (undiscounted) payoff is $U_R := \frac{C\delta}{1-\delta}$ in case of rejection and $U_A := 1 + U_R$ in case of acceptance. Intuitively, costly experiments add an extra incentive for the principal to generate information provoking the agent's final decision early on, since the principal now prefers early rejection over late rejection ($U_R > 0$). One shows that Theorem 1 holds unchanged, except perhaps for the exact values of the cutoffs in the statement of the theorem.²⁰

²⁰The analysis follows the same steps as in the baseline model. The reader is referred to an earlier version of our paper for the details of this analysis (Bizzotto, Rüdiger and Vigier (2018)).

5 Concluding Remarks

We develop a model of a principal seeking to persuade an agent to accept an offer before a deadline. Whether accepting the offer is optimal for the agent depends on an unknown state of the world. The agent can wait in order to accumulate information. That information might come from the principal (inside information), and/or exogenous signals over which the principal exerts no control (outside information). The combination of this outside information and the deadline by which the agent needs to act yields a non-stationary environment in which the agent's threshold of acceptance evolves over time, providing in turn incentives for the principal to keep uncertainty high in the first periods so as to persuade the agent close to the deadline.²¹ We characterize the conditions under which in equilibrium the agent makes his final decision in the first period, those in which the agent sometimes waits until the deadline, link these results to the welfare properties of the model, and contrast our analysis with the setting in which exogenous outside information is the agent's only source of information.

²¹Without a deadline, the principal always generates information inducing the agent to make his final decision in the first period. The same observation applies in the absence of outside information.

Appendix A: Proof of Proposition 1

In this appendix we prove equilibrium existence and uniqueness (Proposition 1). We start with a very general result that will be used repeatedly in this and the next appendices.

Proposition A.1. *Let $t < T$. If $\phi: [0, 1] \rightarrow \mathbb{R}$ is convex (respectively concave) then, irrespective of the signal-generating process, $\mathbb{E}_{s_t} [\phi(p_{t+1})|q_t]$ is convex (resp. concave) in q_t .*

Proof of Proposition A.1: Consider an arbitrary signal-generating process, with realizations in $S = \{s^i\}_{i=1}^K$. Let $\gamma_{Gi} := \mathbb{P}(s_t = s^i | \omega = G)$, $\gamma_{Bi} := \mathbb{P}(s_t = s^i | \omega = B)$, and $p^i(q) := \frac{q\gamma_{Gi}}{q\gamma_{Gi} + (1-q)\gamma_{Bi}}$. To shorten notation, in what follows we use p^i to refer to $p^i(q_t)$.

We have, with this notation,

$$\mathbb{E}_{s_t} [\phi(p_{t+1})|q_t] = \sum_{i=1}^K \mathbb{P}(s_t = s^i | q_t) \phi(p^i).$$

Thus

$$\frac{d\mathbb{E}_{s_t} [\phi(p_{t+1})|q_t]}{dq_t} = \sum_{i=1}^K \left(\frac{d\mathbb{P}(s_t = s^i | q_t)}{dq_t} \phi(p^i) + \mathbb{P}(s_t = s^i | q_t) \frac{d\phi(p^i)}{dp^i} \frac{dp^i}{dq_t} \right),$$

while

$$\begin{aligned} \frac{d^2\mathbb{E}_{s_t} [\phi(p_{t+1})|q_t]}{dq_t^2} &= \sum_{i=1}^K \left(\frac{d^2\mathbb{P}(s_t = s^i | q_t)}{dq_t^2} \phi(p^i) + \left(2 \frac{d\mathbb{P}(s_t = s^i | q_t)}{dq_t} \frac{dp^i}{dq_t} + \mathbb{P}(s_t = s^i | q_t) \frac{d^2 p^i}{dq_t^2} \right) \frac{d\phi(p^i)}{dp^i} + \right. \\ &\quad \left. + \mathbb{P}(s_t = s^i | q_t) \frac{d^2 \phi(p^i)}{d(p^i)^2} \left(\frac{dp^i}{dq_t} \right)^2 \right). \end{aligned} \quad (2)$$

Notice that: (i) $\mathbb{P}(s_t = s^i | q_t) = q_t \gamma_{Gi} + (1 - q_t) \gamma_{Bi}$, thus $\frac{d^2 \mathbb{P}(s_t = s^i | q_t)}{dq_t^2} = 0$; moreover, (ii) $\frac{dp^i}{dq_t} = \frac{\gamma_{Gi} \gamma_{Bi}}{(\mathbb{P}(s_t = s^i | q_t))^2}$ and $\frac{d^2 p^i}{dq_t^2} = -2 \frac{\gamma_{Gi} \gamma_{Bi}}{(\mathbb{P}(s_t = s^i | q_t))^3} \frac{d\mathbb{P}(s_t = s^i | q_t)}{dq_t}$, implying $2 \frac{d\mathbb{P}(s_t = s^i | q_t)}{dq_t} \frac{dp^i}{dq_t} + \mathbb{P}(s_t = s^i | q_t) \frac{d^2 p^i}{dq_t^2} = 0$. In light of (i) and (ii), (2) reduces to

$$\frac{d^2\mathbb{E}_{s_t} [\phi(p_{t+1})|q_t]}{dq_t^2} = \sum_{i=1}^K \left(\mathbb{P}(s_t = s^i | q_t) \frac{d^2 \phi(p^i)}{d(p^i)^2} \left(\frac{dp^i}{dq_t} \right)^2 \right).$$

As $\mathbb{P}(s_t = s^i | q_t) \left(\frac{dp^i}{dq_t} \right)^2 \geq 0$, if ϕ is convex (respectively concave) then $\mathbb{E}_{s_t} [\phi(p_{t+1})|q_t]$ is convex (resp. concave) as well. ■

The continuation game starting in period T with beginning-of-period- T belief p_T is identical to the static Bayesian persuasion game of Kamenica and Gentzkow (2011). We summarize some of their main results in the following lemma.

Lemma A.1. *In equilibrium, at $t=T$, the agent accepts if $q_T \geq \underline{b}$ and rejects otherwise. The principal designs the experiment*

$$M_T = \begin{cases} \{0, \underline{b}\} & \text{if } p_T \in (0, \underline{b}); \\ \{p_T\} & \text{otherwise.} \end{cases}$$

The agent does not benefit from the period- T experiment, hence his equilibrium continuation payoff is convex in p_T .²²

Lemma A.2. *Let $t < T$. Suppose that functions $\hat{g}_{t+1}(p_{t+1})$ and $\hat{f}_{t+1}(p_{t+1})$ uniquely determine the agent's (resp. the principal's) equilibrium continuation payoffs in period $t+1$. If \hat{g}_{t+1} is convex, then:*

1. *in equilibrium, the principal's period- t experiment and the agent's period- t decision are both uniquely determined; the former is a function of p_t only and the latter is a function of q_t only;*
2. *functions $\hat{g}_t(p_t)$ and $\hat{f}_t(p_t)$ uniquely determine the equilibrium continuation payoffs in period t , and \hat{g}_t is convex.*

Proof: Define

$$\tilde{g}_t(q_t) := \mathbb{E}_{s_t} [\hat{g}_{t+1}(p_{t+1}) | q_t], \quad (3)$$

and

$$\tilde{f}_t(q_t) := \mathbb{E}_{s_t} [\hat{f}_{t+1}(p_{t+1}) | q_t]. \quad (4)$$

Then the agent's equilibrium continuation payoff given q_t can be written as

$$g_t(q_t) = \max \{V_R, \delta \tilde{g}_t(q_t), V_B + q_t(V_G - V_B)\}. \quad (5)$$

²²The equilibrium period- T experiment generates information that has no value for the agent, since rejecting is an optimal choice for $q_T=0$ as well as for $q_T=\underline{b}$. The agent's equilibrium continuation payoffs at the beginning of period T are thus given by the convex function $\hat{g}_T(p_T) = \max\{V_R, V_B + p_T(V_G - V_B)\}$.

As \hat{g}_{t+1} is convex by assumption, Proposition A.1 shows that \tilde{g}_t is convex as well. Moreover,

$$\begin{cases} \delta\tilde{g}_t(0) = \delta V_R < V_R; \\ \delta\tilde{g}_t(1) = \delta V_G < V_G. \end{cases} \quad (6)$$

Then (5), (6) and convexity of \tilde{g}_t give unique a_t and b_t , with $a_t \leq \underline{b} \leq b_t$, such that

$$\begin{cases} g_t(q_t) = V_R > \max \{ \delta\tilde{g}_t(q_t), V_B + q_t(V_G - V_B) \} & \text{if } q_t < a_t; \\ g_t(q_t) = \delta\tilde{g}_t(q_t) > \max \{ V_R, V_B + q_t(V_G - V_B) \} & \text{if } q_t \in (a_t, b_t); \\ g_t(q_t) = V_B + q_t(V_G - V_B) > \max \{ V_R, \delta\tilde{g}_t(q_t) \} & \text{if } q_t > b_t. \end{cases}$$

Hence, in equilibrium, the agent rejects if $q_t < a_t$, waits if $q_t \in (a_t, b_t)$, and accepts if $q_t > b_t$. Moreover, since in equilibrium whenever indifferent the agent makes the decision preferred by the principal, the agent waits if $q_t = a_t < b_t$ and accepts if $q_t = b_t$. Hence:

$$f_t(q_t) = \begin{cases} 0 & \text{if } q_t < a_t; \\ \delta\tilde{f}_t(q_t) & \text{if } q_t \in [a_t, b_t); \\ 1 & \text{if } q_t \geq b_t. \end{cases} \quad (7)$$

Standard arguments yield

$$\hat{f}_t = \text{cav } f_t. \quad (8)$$

Furthermore, since in equilibrium whenever indifferent the principal picks the least informative experiment, the principal's experiment in period t is uniquely determined by the belief p_t at the beginning of period t . Lastly, letting $\tau_t(p_t)$ denote the principal's equilibrium experiment given p_t yields $\hat{g}_t(p_t) = \mathbb{E}_{\tau_t(p_t)}[g_t(q_t)|p_t]$.

Finally, since \tilde{g}_t is convex, (5) shows that g_t is as well which, in turn, implies

$$\hat{g}_t(p_t) = \mathbb{E}_{\tau_t(p_t)}[g_t(q_t)|p_t]. \quad (9)$$

The properties of $\tau_t(\cdot)$ implied by (8) finish to establish that \hat{g}_t is convex, since \hat{g}_t is given by (9) and g_t is convex. ■

Proof of Proposition 1: The proposition follows from Lemmata A.1 and A.2. ■

Appendix B: Proof of Theorem 1

In this appendix we prove the steps leading to Theorem 1, including all lemmas of Section 3 except for Lemma 4, whose proof was kept in the text. The order in which we prove the results is as follows: Lemma 1, 5, 6, 7, 2, 3 and 8.

Proof of Lemma 1: For $t = T$, the result follows from Lemma A.1. For $t < T$, the result was shown within the proof of Lemma A.2. ■

Lemma B.1. *At all $t \leq T - 1$, the threshold of acceptance is at least as large as the threshold of acceptance at $t = T - 1$, which itself is at least as large as the static threshold of acceptance: $b_t \geq b_{T-1} \geq b_T = \underline{b}$.*

Proof: The result follows from the arguments in the text above the statement of Lemma 2. ■

Lemma B.2. *Each period, in equilibrium $M_t = \{p_t\}$ for all $p_t \in \{0\} \cup [b_t, 1]$. Moreover:*

- (i) *either $M_t = \{0, b_t\}$ for all $p_t \in (0, b_t)$,*
- (ii) *or $a_t < b_t$ and: $M_t = \{0, a_t\}$ for $p_t \in (0, a_t)$ and there exists $c_t \in [a_t, b_t)$ such that $M_t = \{p_t\}$ for $p_t \in [a_t, c_t]$, while $M_t = \{c_t, b_t\}$ for $p_t \in (c_t, b_t)$.*

Proof: Recall (8). If $a_t = b_t$ (so that the set of beliefs for which in equilibrium the agent waits in period t is empty) then in equilibrium $M_t = \{0, b_t\}$ for all $p_t \in (0, b_t)$. So assume $a_t < b_t$. Observe that:

- (A) $\tilde{f}_t(\cdot)$ (defined by (4)) is concave,
- (B) $f_t(q_t) = \delta \tilde{f}_t(q_t)$ for all $q_t \in \{0\} \cup [a_t, b_t)$.

(A) follows from Proposition A.1 while (B) is obtained from (7). In view of (A)-(B), either (i) in the statement of the lemma holds, or (ii) does. ■

Proof of Lemma 5: If $a_{T-1} = b_{T-1} = b_T$, the claim of the lemma is straightforward.²³ Assume now $b_{T-1} > b_T$. At $q_{T-1} = b_{T-1}$, in equilibrium the agent is indifferent between waiting and accepting. The agent's expected payoff from accepting is $b_{T-1}V_G + (1 - b_{T-1})V_B$. On the other hand, using Lemmata 1 and B.1, the agent's expected payoff from waiting can be written as $\delta[b_{T-1}V_G + (1 - b_{T-1})(\gamma V_R + (1 - \gamma)V_B)]$. So b_{T-1} is the unique solution of

$$xV_G + (1 - x)V_B = \delta[xV_G + (1 - x)(\gamma V_R + (1 - \gamma)V_B)]. \quad (10)$$

Next, consider $t < T - 1$ such that $b_{t+1} = b_{T-1}$. Suppose $q_t = b_t$, so that, by definition, in equilibrium the agent is indifferent between waiting and accepting. The agent's expected payoff from accepting is $b_tV_G + (1 - b_t)V_B$. On the other hand, using Lemma B.1, $q_t = b_t \geq b_{T-1} = b_{t+1}$. Hence, conditional on $s_t = g$, the agent optimally accepts in the next period. It ensues that b_t solves (10) and, therefore, that $b_t = b_{T-1}$. A recursive argument then yields $b_t = b_{T-1}$ for all $t < T$. ■

Proof of Lemma 6: Suppose that in equilibrium the principal is aggressive in period $1 < t + 1 < T$. If $a_t = b_t$ the statement of the lemma is straightforward. Assume therefore that $a_t < b_t$. By virtue of Lemma B.2, in order to establish that the principal is also aggressive in period t it is enough to show that, at $p_t = a_t$, the principal strictly prefers the experiment $M_t = \{0, b_t\}$ over the uninformative experiment. On one hand, the principal's expected payoff from designing $M_t = \{0, b_t\}$ is $\frac{a_t}{b_t}$. On the other hand, her expected payoff from designing the uninformative experiment is given by $\delta \mathbb{E}_{s_t}[\hat{f}_{t+1}(p_{t+1}) | q_t = a_t]$. The next sequence of inequalities therefore concludes the proof:

$$\delta \mathbb{E}_{s_t}[\hat{f}_{t+1}(p_{t+1}) | q_t = a_t] \leq \delta \hat{f}_{t+1}(a_t) = \delta \frac{a_t}{b_{t+1}} < \frac{a_t}{b_t}.$$

The first inequality follows from noting that \hat{f}_{t+1} is concave; the equality follows from the assumption that the principal is aggressive in period $t + 1$, and the second inequality is due to Lemma 5. ■

²³ $a_{T-1} = b_{T-1} = b_T = \underline{b}$ implies $a_{t-1} = b_{t-1} = b_t$ whenever $a_t = b_t = \underline{b}$. Hence, a recursive argument yields the result in this case.

Proof of Lemma 7: Note that, in view of Lemma 6, it is enough to show that in equilibrium the principal is aggressive at $t=1$ when T is sufficiently large. Next, Lemma 5 shows that any benefit to the principal from being conservative at $t=1$ must come from persuading the agent to accept at $t=T$ when $\omega=B$. So these benefits are bounded from above by δ^{T-1} , which tends to 0 as $T \rightarrow \infty$. On the other hand, as $b_1 < 1$, the corresponding loss to the principal is bounded away from zero since by being aggressive at $t=1$ the principal obtains acceptance with strictly positive probability conditional on $\omega=B$. We conclude that, for T sufficiently large, in equilibrium the principal is aggressive at $t=1$. ■

Proof of Lemma 2: In the perfect bad news case, the result follows from Lemma 5. In the perfect good news case, the result is easily obtained by induction using Lemma 4. ■

Lemma B.3. *Consider a period $t < T$ and let a_t^+ denote the beginning-of-period- $t+1$ belief given $q_t = a_t$ and $s_t = g$. Then $a_t^+ > a_{t+1}$.*

Proof: The result is trivial if $a_t = b_t$, so suppose $a_t < b_t$. Assume by way of contradiction that $a_t^+ \leq a_{t+1}$. By definition of a_t , in equilibrium, at $q_t = a_t$ the agent is indifferent between waiting and rejecting; the corresponding remark applies to period $t+1$. Therefore $a_t^+ \leq a_{t+1}$ implies that, by waiting at $q_t = a_t$, the agent's expected continuation payoff is as if the agent rejected with probability 1 in period $t+1$. As $V_R > 0$, rejecting in period t thus yields the agent a strictly higher expected continuation payoff than waiting. This remark contradicts the definition of a_t . ■

Lemma B.4. *Consider a period $t < T-1$ such that, in equilibrium, in period t the principal is not aggressive. Let c_t^+ denote the beginning-of-period- $t+1$ belief given $q_t = c_t$ and $s_t = g$. With perfect bad news, $c_t^+ < b_{t+1}$.*

Proof: Suppose by way of contradiction that $c_t^+ \geq b_{t+1}$. Then, using Lemma 5, given $p_t = c_t$, in equilibrium the experiment $M_t = \{0, b_t\}$ gives the principal a strictly larger expected continuation payoff than the uninformative experiment. This cannot be, since by definition of c_t (Lemma B.2) the uninformative experiment has to be optimal at $p_t = c_t$. ■

Lemma B.5. *Consider a period $t < T - 1$ such that, in equilibrium, in period t the principal is not aggressive. With perfect bad news, $c_{t+1} = a_{t+1}$ implies $c_t = a_t$.*

Proof: Assume that the conditions in the statement of the lemma hold. Recall to begin that $c_t \geq a_t$, by definition. Suppose by way of contradiction that $c_t > a_t$. Our goal will be to show that, given $p_t = c_t$, in equilibrium the experiment $M_t = \{a_t, b_t\}$ yields the principal strictly larger expected continuation payoff than the uninformative experiment, contradicting the definition of c_t .

As a preliminary step, notice that, since in equilibrium the principal is not aggressive in period $t+1$ (Lemma 6), at $p_{t+1} = a_{t+1}$ the principal must weakly prefer the uninformative experiment over $M_{t+1} = \{0, b_{t+1}\}$ (Lemma B.2), that is,

$$f_{t+1}(a_{t+1}) \geq \frac{a_{t+1}}{b_{t+1}} f_{t+1}(b_{t+1}) + \left(1 - \frac{a_{t+1}}{b_{t+1}}\right) f_{t+1}(0). \quad (11)$$

Next, Lemmata B.3 and B.4 together imply $c_t^+ \in (a_{t+1}, b_{t+1})$. Thus, using Lemma B.2, at $p_t = c_t$, in equilibrium the principal's expected continuation payoff from the uninformative experiment can be expressed as $\delta \mathbb{E}[f_{t+1}(X)]$, where X is a random variable with mean c_t and support $\{0, a_{t+1}, b_{t+1}\}$, and $\mathbb{P}(X=0) = (1 - c_t)\gamma$. Call this remark A.

Now, at $p_t = c_t$, reasoning similarly as above and using $b_{t+1} = b_t$ (Lemma 5) establishes that the principal's expected continuation payoff from designing the experiment $M_t = \{a_t, b_t\}$ can be bounded from below (strictly) by $\delta \mathbb{E}[f_{t+1}(Y)]$, where Y is a random variable with mean c_t and support $\{0, a_{t+1}, b_{t+1}\}$, and $\mathbb{P}(Y=0) = \frac{b_t - c_t}{b_t - a_t} (1 - a_t)\gamma$. Call this remark B.

The last step of the proof is as follows. First, straightforward algebra shows $1 - c_t > \frac{b_t - c_t}{b_t - a_t} (1 - a_t)$. Hence, $\mathbb{P}(X=0) > \mathbb{P}(Y=0)$. As X and Y have the same mean and are both supported on $\{0, a_{t+1}, b_{t+1}\}$, we conclude that X is a mean-preserving spread of Y . Inequality (11) then implies $\delta \mathbb{E}[f_{t+1}(X)] \leq \delta \mathbb{E}[f_{t+1}(Y)]$. Hence, combining remarks A and B, given $p_t = c_t$, in equilibrium the experiment $M_t = \{a_t, b_t\}$ yields the principal strictly larger expected continuation payoff than the uninformative experiment, contradicting the definition of c_t . ■

Lemma B.6. *With perfect bad news, in any period such that in equilibrium the principal is not aggressive, $c_t = a_t$.*

Proof: Consider t such that, in equilibrium, the principal is not aggressive in period t . Then, by Lemma 6, the principal is not aggressive in period $T-1$. Moreover, using Lemma B.3, for any $q_{T-1} \in [a_{T-1}, b_{T-1})$, the agent waits and: (i) accepts following $s_{T-1} = g$, (ii) rejects following $s_{T-1} = b$. Thus f_{T-1} is linear in q_{T-1} over the belief interval $[a_{T-1}, b_{T-1})$. It ensues that $c_{T-1} = a_{T-1}$. Reasoning by induction using Lemma B.5 then establishes $c_t = a_t$. ■

Proof of Lemma 3: By the definitions of aggressive and conservative, Lemma 3 follows from Lemma B.2 if we can show that, in equilibrium, each period either (a) the principal is aggressive or (b) $c_t = a_t$. Lemma 4 shows that (a) always holds in the perfect good news case. In the perfect bad news case, by Lemma B.6, either (a) holds or (b) does. ■

Lemma B.7. *In equilibrium, conditional on $\omega = G$ the agent accepts with probability 1.*

Proof: Lemma 1 ensures that in equilibrium the agent rejects in period t if and only if $q_t \in [0, a_t)$. Lemma 3 ensures that in equilibrium $q_t \notin (0, a_t)$. Thus if the agent rejects in period t , it must be the case that $q_t = 0$. ■

Lemma B.8. *Let $\tilde{\gamma}(\delta) := \left(\frac{1-\delta}{\delta}\right) \frac{V_R(V_G - V_B)}{(V_G - V_R)(V_R - V_B)}$. With perfect bad news, $\gamma > \tilde{\gamma}(\delta)$ if and only if $b_{T-1} > \underline{b}$; in this case,*

$$b_{T-1} = \frac{\delta(\gamma V_R + (1-\gamma)V_B) - V_B}{(V_G - V_B)(1-\delta) + \delta\gamma(V_R - V_B)}. \quad (12)$$

Proof: The arguments in the proof of Lemma A.2 show that $b_{T-1} > \underline{b}$ if and only if given $q_{T-1} = \underline{b}$ the agent strictly prefers waiting over rejecting, that is, if and only if

$$\delta [\underline{b}V_G + (1-\underline{b})(\gamma V_R + (1-\gamma)V_B)] > V_R,$$

which, upon rearrangement, yields $\gamma > \tilde{\gamma}(\delta)$. Solving (10) gives (12). ■

Proof of Lemma 8: By Lemmata 5 and B.8, if $\gamma \leq \tilde{\gamma}(\delta)$ then in equilibrium $a_t = b_t = \underline{b}$ in every period, and so the principal is aggressive in period $T-1$. Since $\tilde{\gamma}(\delta) > 1$ for δ sufficiently

small, we conclude that, in equilibrium, for δ small enough the principal is aggressive in period $T - 1$ regardless of γ .

Suppose next that $\gamma > \tilde{\gamma}(\delta)$. Then at $q_{T-1} = a_{T-1}$, in equilibrium the agent is indifferent between rejecting and waiting. Hence, $V_R = \delta[a_{T-1}V_G + (1 - a_{T-1})(\gamma V_R + (1 - \gamma)V_B)]$, giving

$$a_{T-1} = \frac{V_R - \delta(\gamma V_R + (1 - \gamma)V_B)}{\delta(V_G - \gamma V_R - (1 - \gamma)V_B)} \quad (13)$$

after rearrangement. Now, using Lemma B.2, the necessary and sufficient condition for the principal to not be aggressive in period $T - 1$ in equilibrium is $f_{T-1}(a_{T-1}) \geq \frac{a_{T-1}}{b_{T-1}}$. Noting that $f_{T-1}(a_{T-1}) = \delta[a_{T-1} + (1 - \gamma)(1 - a_{T-1})]$ and substituting for a_{T-1} and b_{T-1} using (12) and (13), the former inequality becomes

$$\begin{aligned} & V_B\delta(\gamma - 1) + \delta V_G(1 - \gamma) + V_R\gamma(1 - \delta) \\ & \geq \frac{[V_B\delta(\gamma - 1) + V_R(1 - \delta\gamma)][\delta\gamma(V_B - V_R) + (V_B - V_G)(1 - \delta)]}{\delta[(V_B - V_R)\delta\gamma + V_B(1 - \delta)]}. \end{aligned} \quad (14)$$

One checks that for $\delta = 1$ the quadratic equation in γ obtained from (14) has roots $\gamma = 0$ and $\gamma = 1$. On the other hand, for $\delta < 1$, (14) is violated when either $\gamma = 1$ or $\gamma = \tilde{\gamma}(\delta)$. So (14) holds for all values of γ in between the roots of the corresponding quadratic equation. Letting $\underline{\gamma}(\delta)$ and $\bar{\gamma}(\delta)$ denote the real roots, the previous remarks yield $\tilde{\gamma}(\delta) < \underline{\gamma}(\delta) \leq \bar{\gamma}(\delta) < 1$ and show that these roots only exist for $\delta > \underline{\delta}$, where $\underline{\delta} > 0$ is defined implicitly by $\underline{\gamma}(\underline{\delta}) = \bar{\gamma}(\underline{\delta})$. Noting that, by Lemma 3, whenever the principal is not aggressive she is conservative concludes the proof. \blacksquare

Proposition B.1. *Let b_t^G (respectively b_t^B) denote the period- t threshold of acceptance under perfect good news (resp. perfect bad news). Then $b_t^G \leq b_t^B$ for all t , with $b_T^G = b_T^B = \underline{b}$.*

Proof: First, notice that, under perfect bad news, b_{T-1}^B satisfies the following fixed-point property: at $q_{T-1} = b_{T-1}^B$ the agent is indifferent between (a) accepting and (b) making his final decision next period given that next period the principal designs $M_T = \{0, b_{T-1}^B\}$ if $p_T \in (0, b_{T-1}^B)$ and $M_T = \{b_{T-1}^B, 1\}$ if $p_T \in (b_{T-1}^B, 1)$. Next, let X denote the random variable representing the belief at which, under perfect bad news, the agent makes his final decision following scenario (b). Then $\text{supp}(X) = \{0, b_{T-1}^B, 1\}$, $\mathbb{E}[X] = b_{T-1}^B$ and $\mathbb{P}(X = 0) = (1 - b_{T-1}^B)\gamma$. Straightforward algebra therefore yields $\mathbb{P}(X = 1) = b_{T-1}^B\gamma$ and $\mathbb{P}(X = b_{T-1}^B) = 1 - \gamma$.

We claim that, under perfect good news, b_{T-1}^G satisfies the same fixed-point property as above. To see this, consider $q_{T-1} = b_{T-1}^G$ and let Y denote the random variable representing the belief at which, under perfect good news, the agent makes his final decision following scenario (b). Then $\text{supp}(Y) = \{0, b_{T-1}^G, 1\}$, $\mathbb{E}[Y] = b_{T-1}^G$ and $\mathbb{P}(Y = 1) = b_{T-1}^G \gamma$. Straightforward algebra yields $Y \sim X$. The claim ensues.

We now show by induction that $b_t^G \leq b_t^B$ for all $t < T$ (the case $t = T$ is obvious). For $t = T - 1$ the result immediately follows from the claim above. Next, suppose $b_t^G \leq b_t^B$ for $1 < t < T$. By Lemma 5: $b_t^B = b_{T-1}^B$. Hence $b_t^G \leq b_{T-1}^B$. But then the claim above implies that, under perfect good news, at $q_{t-1} = b_{T-1}^B$ in equilibrium the agent must weakly prefer accepting over waiting. So $b_{t-1}^G \leq b_{T-1}^B = b_{t-1}^B$. ■

Appendix C: Technical Appendix for Subsection 3.2

Pareto Efficiency

Define an outcome Z by two probability distributions over $\{\text{accept}, \text{reject}\} \times \{1, \dots, T\}$, specifying, respectively, the probability $Z(x, t | G)$ that the agent makes the final decision x in period t when $\omega = G$, and the corresponding probability $Z(x, t | B)$ when $\omega = B$.

Proposition C.1. *An outcome Z is Pareto efficient if and only if (i) $Z(\text{accept}, 1 | G) = 1$, and (ii) $Z(\text{accept}, 1 | B) + Z(\text{reject}, 1 | B) = 1$.*

Proof: Let \mathcal{Z}^* denote the set of Pareto efficient outcomes and \mathcal{Z}^\dagger the set of outcomes satisfying conditions (i) and (ii) in the statement of the lemma.

We start by showing that $\mathcal{Z}^* \subseteq \mathcal{Z}^\dagger$. First, as (a) $V_R > 0$ and (b) the principal gets 0 from rejection:

$$Z \in \mathcal{Z}^* \Rightarrow Z(\text{reject}, t | \omega) = 0, \forall t > 1. \quad (15)$$

Second, as (a) the principal prefers acceptance over rejection and (b) $V_G > V_R$:

$$Z \in \mathcal{Z}^* \Rightarrow Z(\text{reject}, t | G) = 0, \forall t. \quad (16)$$

Third, since (a) $V_G > 0$ and (b) $\delta < 1$:

$$Z \in \mathcal{Z}^* \Rightarrow Z(\text{accept}, 1 | G) = 1. \quad (17)$$

Fourth, we claim that

$$Z \in \mathcal{Z}^* \Rightarrow Z(\text{accept}, t | B) = 0, \forall t > 1. \quad (18)$$

Suppose for a contradiction that this is not the case and that we can find $Z \in \mathcal{Z}^*$ and $\hat{t} > 1$ such that $Z(\text{accept}, \hat{t} | B) > 0$. Define \tilde{Z} as follows:

- (i) $\tilde{Z}(\cdot, \cdot | G) = Z(\cdot, \cdot | G)$,
- (ii) $\tilde{Z}(\cdot, t | B) = Z(\cdot, t | B), \forall t \notin \{1, \hat{t}\}$,
- (iii) $\tilde{Z}(\text{accept}, \hat{t} | B) = 0$,
- (iv) $\tilde{Z}(\text{accept}, 1 | B) = Z(\text{accept}, 1 | B) + \delta^{\hat{t}-1} Z(\text{accept}, \hat{t} | B)$,
- (v) $\tilde{Z}(\text{reject}, 1 | B) = Z(\text{reject}, 1 | B) + (1 - \delta^{\hat{t}-1}) Z(\text{accept}, \hat{t} | B)$.

Applying (15)-(17), the principal's expected payoff under \tilde{Z} can be written as

$$\begin{aligned}
& p_1 \tilde{Z}(\text{accept}, 1|G) + (1-p_1) \sum_{t=1}^T \delta^{t-1} \tilde{Z}(\text{accept}, t|B) \\
&= p_1 Z(\text{accept}, 1|G) + (1-p_1) \sum_{t \notin \{1, \hat{t}\}} \delta^{t-1} Z(\text{accept}, t|B) \\
&+ (1-p_1) \left[Z(\text{accept}, 1|B) + \delta^{\hat{t}-1} Z(\text{accept}, \hat{t}|B) \right] \\
&= p_1 + (1-p_1) \sum_{t=1}^T \delta^{t-1} Z(\text{accept}, t|B).
\end{aligned}$$

Thus \tilde{Z} and Z give the same expected payoff to the principal. On the other hand, the agent's expected payoff under \tilde{Z} can be written as

$$\begin{aligned}
& p_1 V_G \tilde{Z}(\text{accept}, 1|G) + (1-p_1) \left[V_R \tilde{Z}(\text{reject}, 1|B) + V_B \sum_{t=1}^T \delta^{t-1} \tilde{Z}(\text{accept}, t|B) \right] \\
&= p_1 V_G Z(\text{accept}, 1|G) + (1-p_1) V_B \sum_{t \notin \{1, \hat{t}\}} \delta^{t-1} Z(\text{accept}, t|B) \\
&+ (1-p_1) \left[V_R \left(Z(\text{reject}, 1|B) + (1-\delta^{\hat{t}-1}) Z(\text{accept}, \hat{t}|B) \right) + V_B \left(Z(\text{accept}, 1|B) + \delta^{\hat{t}-1} Z(\text{accept}, \hat{t}|B) \right) \right] \\
&= p_1 V_G + (1-p_1) \left[V_R Z(\text{reject}, 1|B) + V_B \sum_{t=1}^T \delta^{t-1} Z(\text{accept}, t|B) \right] \\
&+ (1-p_1) V_R (1-\delta^{\hat{t}-1}) Z(\text{accept}, \hat{t}|B),
\end{aligned}$$

where the first two terms in the final sum represent the agent's expected payoff under Z . As $(1-p_1)V_R(1-\delta^{\hat{t}-1})Z(\text{accept}, \hat{t}|B) > 0$, we find that the agent's expected payoff under \tilde{Z} is greater than it is under Z . Thus, \tilde{Z} and Z give the same expected payoff to the principal, but the agent's expected payoff is strictly greater under \tilde{Z} than it is under Z , contradicting the initial assumption that $Z \in \mathcal{Z}^*$. We conclude that (18) holds. Combining (15)-(18) shows that $\mathcal{Z}^* \subseteq \mathcal{Z}^\dagger$.

We next show that $\mathcal{Z}^\dagger \subseteq \mathcal{Z}^*$. Let $Z \in \mathcal{Z}^\dagger$. If $Z \notin \mathcal{Z}^*$, we can find Z' which Pareto dominates Z . Either $Z' \in \mathcal{Z}^\dagger$ or, by the first part of the proof, we can find $Z'' \in \mathcal{Z}^\dagger$ which Pareto dominates Z' , in which case Z'' Pareto dominates Z , by transitivity. Hence, assume without loss of generality that $Z' \in \mathcal{Z}^\dagger$. Since both Z and Z' belong to \mathcal{Z}^\dagger we have $Z'(x, t|\omega) = Z(x, t|\omega)$

unless $t = 1$ and $\omega = B$. But then either $Z'(\text{accept}, 1|B) < Z(\text{accept}, 1|B)$ and then the principal strictly prefers Z over Z' , or $Z'(\text{accept}, 1|B) > Z(\text{accept}, 1|B)$ and then the agent strictly prefers Z over Z' . Therefore, Z' does not Pareto dominate Z , contradicting the definition of Z' . This shows that $Z \in \mathcal{Z}^*$. ■

Comparison with the Single-Player Setting

We show here that increasing γ may increase the probability of type II errors made and lower the agent's expected payoff. We start with the following useful lemma.

Lemma C.1. *Either $a_t = b_t = \underline{b}$ in all periods or the interval of beliefs at which the agent waits is constant for the first $x \leq T - 1$ periods, and then strictly decreasing over the remaining periods.*

Proof: In the perfect good news case, the result follows from Lemma 4. Below, we focus on the perfect bad news case. If in equilibrium the principal is aggressive at $t = T - 1$ then the result is a consequence of Lemmata 5, 6 and A.1. Therefore, suppose henceforth that the principal is conservative at $t = T - 1$.

Notice to begin with that by Lemma 5 all we need to show is that the sequence a_t increases with t . Reasoning as in Lemma B.4 establishes that $a_{T-1}^+ \in (a_{T-1}, b_{T-1})$. Since $\hat{g}_{T-1} > \hat{g}_T$ over the belief interval (a_{T-1}, b_{T-1}) , we obtain, using definition (3),

$$\begin{aligned} \tilde{g}_{T-2}(a_{T-1}) &= \mathbb{P}(s_{T-2} = b | q_{T-2} = a_{T-1})V_R + \mathbb{P}(s_{T-2} = g | q_{T-2} = a_{T-1})\hat{g}_{T-1}(a_{T-1}^+) \\ &> \mathbb{P}(s_{T-1} = b | q_{T-1} = a_{T-1})V_R + \mathbb{P}(s_{T-1} = g | q_{T-1} = a_{T-1})\hat{g}_T(a_{T-1}^+) \\ &= \tilde{g}_{T-1}(a_{T-1}). \end{aligned}$$

We conclude from the arguments in the proof of Lemma A.2 that $a_{T-2} < a_{T-1}$.

Now, if the principal is aggressive in period $T - 2$ then Lemmata 5 and 6, immediately give $a_1 = \dots = a_{T-3} < a_{T-2} < a_{T-1}$. So suppose that the principal is conservative in period $T - 2$. Since $a_{T-2} < a_{T-1}$ and $b_{T-2} = b_{T-1}$, Lemmata B.2 and B.6 establish that $\hat{g}_{T-2} > \hat{g}_{T-1}$ over the

belief interval (a_{T-2}, b_{T-2}) . Moreover, $a_{T-2}^+ \in (a_{T-2}, b_{T-2})$. Hence:

$$\begin{aligned}\tilde{g}_{T-3}(a_{T-2}) &= \mathbb{P}(s_{T-3} = b | q_{T-3} = a_{T-2})V_R + \mathbb{P}(s_{T-3} = g | q_{T-3} = a_{T-2})\hat{g}_{T-2}(a_{T-2}^+) \\ &> \mathbb{P}(s_{T-2} = b | q_{T-2} = a_{T-2})V_R + \mathbb{P}(s_{T-2} = g | q_{T-2} = a_{T-2})\hat{g}_{T-1}(a_{T-2}^+) \\ &= \tilde{g}_{T-2}(a_{T-2}).\end{aligned}$$

We conclude from the arguments in the proof of Lemma A.2 that $a_{T-3} < a_{T-2}$. Pursuing the recursion completes the proof. \blacksquare

Section 3 revealed the existence of two possible equilibrium regimes. In one regime, the principal is aggressive at $t=1$, and triggers the agent's final decision in the first period. In the other regime, the principal is conservative, and seeks to sustain uncertainty until $t=T$. We next establish that, as long as no regime switch occurs, increasing the amount of exogenous outside information weakly increases the welfare of the agent.

Lemma C.2. *Let $(T'', \gamma'') \geq (T', \gamma')$. Assume that, in equilibrium, at $t=1$, either the principal is aggressive given $(T, \gamma) = (T', \gamma')$ as well as given $(T, \gamma) = (T'', \gamma'')$, or the principal is conservative in both cases. Then the agent's equilibrium expected payoff is greater given $(T, \gamma) = (T'', \gamma'')$ than given $(T, \gamma) = (T', \gamma')$.*

Proof: We show the proof for the perfect bad news case (the proof for the perfect good news case is similar but easier). Throughout the proof primes will be used for all objects corresponding to the situation in which $(T, \gamma) = (T', \gamma')$. Similarly, double primes will be used for all objects corresponding to the situation in which $(T, \gamma) = (T'', \gamma'')$.

In the case in which at $t=1$ the principal is aggressive given $(T, \gamma) = (T', \gamma')$ as well as given $(T, \gamma) = (T'', \gamma'')$, the result immediately follows from Lemma 5 and noting that $b''_{T''-1} \geq b'_{T'-1}$ (which, in turn, follows from Lemma B.8). Below we deal with the other case.

Since $\gamma'' \geq \gamma'$ notice first that, by Lemma A.1, $a''_{T''-1} \leq a'_{T'-1}$ and $b''_{T''-1} \geq b'_{T'-1}$. Moreover, since at $t=1$ the principal is conservative, the same must be true at all $t < T$ (Lemma 6). Hence, $\hat{g}''_{T''-1}(\cdot)$ is piecewise linear with kinks at $a''_{T''-1}$ and $b''_{T''-1}$, $\hat{g}''_{T''-1}(a''_{T''-1}) = V_R$ and $\hat{g}''_{T''-1}(b''_{T''-1}) = V_B + b''_{T''-1}(V_G - V_B)$. A similar remark applies to $\hat{g}'_{T'-1}(\cdot)$. We conclude that

$\hat{g}''_{T''-1}(\cdot) \geq \hat{g}'_{T'-1}(\cdot)$. If $T' = 2$ then Lemma C.1 finishes the proof. Otherwise,

$$\begin{aligned} \tilde{g}''_{T''-2}(a'_{T'-2}) &= \mathbb{E}_{s''_{T''-2}}[\hat{g}''_{T''-1}(p_{T''-1})|q_{T''-2} = a'_{T'-2}] \\ &\geq \mathbb{E}_{s'_{T''-2}}[\hat{g}''_{T''-1}(p_{T''-1})|q_{T''-2} = a'_{T'-2}] \\ &\geq \mathbb{E}_{s'_{T''-2}}[\hat{g}'_{T''-1}(p_{T''-1})|q_{T''-2} = a'_{T'-2}] \\ &= \tilde{g}'_{T''-2}(a'_{T'-2}). \end{aligned}$$

The first inequality follows from convexity of $\hat{g}''_{T''-1}(\cdot)$ and the fact that, since $\gamma'' \geq \gamma'$, $s''_{T''-2}$ is Blackwell-more-informative than $s'_{T''-2}$. The second inequality follows from the previously established inequality $\hat{g}''_{T''-1}(\cdot) \geq \hat{g}'_{T''-1}(\cdot)$. Hence, $\delta \tilde{g}''_{T''-2}(a'_{T'-2}) \geq \delta \tilde{g}'_{T''-2}(a'_{T'-2}) = V_R$ which, in turn, implies $a''_{T''-2} \leq a'_{T'-2}$ and, reasoning as above, $\hat{g}''_{T''-2}(\cdot) \geq \hat{g}'_{T'-2}(\cdot)$. If $T' = 3$ then Lemma C.1 finishes the proof. Otherwise, we can repeat the last step. ■

Lemma C.3. *Fix $\gamma \in (0, 1)$. There exists $\underline{\delta}(\gamma, T) < 1$ such that, in equilibrium, whenever $\delta > \underline{\delta}(\gamma, T)$, at $t = 1$ the principal is conservative.*

Proof: Fix $\gamma \in (0, 1)$. First, notice that

$$\lim_{\delta \rightarrow 1} b_t = 1 \tag{19}$$

for all $t < T$. Next, let each element of the sequence $\{x_t\}_{t=1}^{T-1}$ be defined implicitly as follows:

$$\mathbb{P}(\omega = G | p_t = x_t, s_t = g, \dots, s_{T-1} = g) = \underline{b}.$$

Thus, $x_1 < x_2 < \dots < x_{T-1} < \underline{b}$. Moreover notice that, for all $t < T$:

$$\limsup_{\delta \rightarrow 1} a_t \leq x_t. \tag{20}$$

Otherwise, we could find a δ sufficiently close to 1 such that given $q_t = a_t$ the agent would strictly prefer waiting until the deadline over rejection (contradicting the definition of a_t). Let $1 - x_1 > \epsilon > 0$. Applying (20), we can find $\underline{\delta} < 1$ such that $\delta > \underline{\delta}$ implies

$$\hat{f}_1(x_1 + \epsilon) \geq \delta^{T-1} \left[1 - (1 - (x_1 + \epsilon))(1 - (1 - \gamma)^{T-1}) \right]. \tag{21}$$

Noting that

$$1 - (1 - (x_1 + \epsilon))(1 - (1 - \gamma)^{T-1}) = (1 - \gamma)^{T-1} + (x_1 + \epsilon)(1 - (1 - \gamma)^{T-1}) > x_1 + \epsilon,$$

combining (19)-(21) yields, for δ sufficiently large:

$$\hat{f}_1(x_1 + \epsilon) > \frac{x_1 + \epsilon}{b_1}.$$

If in equilibrium the principal were aggressive in period 1 we would have $\hat{f}_1(x_1 + \epsilon) = \frac{x_1 + \epsilon}{b_1}$. ■

Proposition C.2. *With perfect good news, the agent's equilibrium expected payoff is monotonically increasing in T and γ . With perfect bad news, the agent's equilibrium expected payoff is monotonically increasing in T and, if players are sufficiently impatient, also monotonically increasing in γ ; however, if players are patient enough, the agent's equilibrium expected payoff is non-monotonic in γ .*

Proof: For the perfect good news case, the result follows from Lemma C.2. We show the proof of the result for the perfect bad news case. We start with three observations:

- *Observation 1:* if in equilibrium at $t=1$ the principal is aggressive then \hat{g}_1 is piecewise linear with a kink at b_1 , $\hat{g}_1(0) = V_R$, $\hat{g}_1(b_1) = V_B + b_1(V_G - V_B)$ and $\hat{g}_1(1) = V_G$.
- *Observation 2:* if in equilibrium at $t=1$ the principal is conservative then \hat{g}_1 is piecewise linear with kinks at a_1 and b_1 , $\hat{g}_1(0) = \hat{g}_1(a_1) = V_R$, $\hat{g}_1(b_1) = V_B + b_1(V_G - V_B)$ and $\hat{g}_1(1) = V_G$.
- *Observation 3:* b_1 is both non-decreasing, and continuous in γ .

Observations 1 and 2 immediately follow from the definitions of a_1 , b_1 , and the experiments designed by the principal when she is aggressive and conservative. Observation 3 follows from Lemmata 5 and B.8.

Now, let $T'' > T'$. We want to show that the agent's equilibrium expected payoff is at least as large in the game of length $T = T''$ as in the game of length $T = T'$. If in equilibrium the principal is aggressive at $t=1$ given $T = T''$ and given $T = T'$, the result then follows from Lemma C.2, and similarly if in equilibrium the principal is conservative at $t=1$ given both game lengths. Hence, by Lemma 6, the only case left to consider is when in equilibrium the

principal is aggressive at $t=1$ given $T=T''$, but conservative at $t=1$ given $T=T'$. In the latter case, the result follows from Observations 1-2 combined with Lemma 5.

Next, let $\gamma'' > \gamma'$. We first want to show that, if players are sufficiently impatient, then the agent's equilibrium expected payoff is at least as large under $\gamma=\gamma''$ than under $\gamma=\gamma'$. We know from Lemmata 6 and 8 that, for $\delta < \underline{\delta}$, in equilibrium the principal is aggressive at $t=1$ regardless of γ . So the result follows from Observations 1 and 3.

Finally, we want to show that, if players are patient enough, then the agent's equilibrium expected payoff is non-monotone in γ . This result follows from Lemmata B.8 and C.3, combined with Observations 1-2 showing that an equilibrium switch from aggressive at $t=1$ to conservative at $t=1$ triggers a drop of the agent's equilibrium expected payoff. ■

The rest of this appendix considers a hypothetical planner with payoffs W_{aG} from acceptance in state G , $W_{rG} < W_{aG}$ from rejection in state G , W_{rB} from rejection in state B , and $W_{aB} < W_{rB}$ from acceptance in state B . We are interested in this planner's equilibrium expected payoff, Q . For concreteness, we henceforth refer to Q as the (equilibrium) quality of the agent's final decision.²⁴

The planner's welfare differs from the agent's in two ways: first, while the planner cares about errors made by the agent, the planner is indifferent about the timing of said errors; second the planner and the agent may weigh type I and type II errors differently. Notwithstanding these differences, the effect of exogenous outside information on Q mirrors its effect on the welfare of the agent (Proposition C.2).

Proposition C.3. *With perfect good news, the quality of the agent's final decision is monotonically increasing in T and γ . With perfect bad news, Q is monotonically increasing in T and, if players are sufficiently impatient, also monotonically increasing in γ . However, if players are patient enough, then Q is non-monotonic in γ .*

Proof: We focus as usual on the perfect bad news case (the perfect good news case being similar and easier). Let X denote the random variable representing the belief at which in equilibrium the agent makes his final decision. Let $\phi: [0, 1] \rightarrow \mathbb{R}$ denote the piecewise linear function with a kink at \underline{b} such that $\phi(0) = W_{rB}$, $\phi(\underline{b}) = W_{aB} + \underline{b}(W_{aG} - W_{aB})$ and $\phi(1) = W_{aG}$. Then:

²⁴If $W_{aG} = W_{rB} = 1$ and $W_{rG} = W_{aB} = 0$, then $1 - Q$ represents the equilibrium probability that the agent makes a type I or type II error.

- (a) $\phi(\cdot)$ is convex
- (b) $\mathbb{E}[X] = p_1$;
- (c) if in equilibrium the principal is aggressive at $t = 1$ then $\text{supp}(X) \subseteq \{0, b_1\}$;
- (d) if in equilibrium the principal is conservative at $t = 1$ then $\text{supp}(X) = \{0, a_{T-1}^+, b_1\}$, where $a_{T-1}^+ > \underline{b}$ denotes the beginning of period- T belief given $q_{T-1} = a_{T-1}$ and $s_{T-1} = g$;
- (e) $Q = \mathbb{E}[\phi(X)]$.

We are now ready to prove the various parts of the proposition. First, we know from Lemmata 6 and 8 that, for $\delta < \underline{\delta}$, in equilibrium the principal is aggressive at $t = 1$ regardless of γ . Hence, suppose $\delta < \underline{\delta}$. Let $\gamma'' > \gamma'$. Then, by Lemmata 5 and B.8, $b_1'' \geq b_1'$. That $Q'' \geq Q'$ now follows from remarks (a), (b), (c) and (e) above.

Next, if players are patient enough, Lemmata B.8 and C.3 establish that, starting from $\gamma = 0$ and increasing γ , in equilibrium, at $t = 1$ the principal is aggressive at first but then switches to being conservative. Since b_1 is continuous in γ , remarks (a)-(e) establish that this equilibrium switch induces a drop in Q .

Lastly, let $T'' > T'$. If in equilibrium the principal is aggressive at $t = 1$ given $T = T''$ and given $T = T'$ then, by Lemma 5, $Q'' = Q'$. If in equilibrium the principal is aggressive at $t = 1$ given $T = T''$ but conservative at $t = 1$ given $T = T'$ then, by Lemma 5 and remarks (a)-(e), $Q'' > Q'$. By Lemma 6, the last case remaining is when in equilibrium the principal is conservative at $t = 1$ given $T = T''$ and given $T = T'$. A simple recursive argument based on Proposition A.1 then establishes $Q'' > Q'$. ■

References

- Au, P. H. (2015) Dynamic information disclosure, *The RAND Journal of Economics*, **46**, 791–823.
- Aumann, R. J., Maschler, M. and Stearns, R. E. (1995) *Repeated games with incomplete information*, MIT press.
- Bizzotto, J., Rüdiger, J. and Vigier, A. (2018) Dynamic persuasion with outside information, *Mimeo*.
- Board, S. and Lu, J. (2018) Competitive information disclosure in search markets.
- Bolton, P. and Harris, C. (1999) Strategic experimentation, *Econometrica*, **67**, 349–374.
- Bonatti, A. and Hörner, J. (2017) Learning to disagree in a game of experimentation, *Journal of Economic Theory*, **169**, 234–269.
- Che, Y.-K., Kim, K. and Mierendorff, K. (2020) Keeping the listener engaged: a dynamic model of bayesian persuasion, *mimeo*.
- Ely, J. and Szydlowski, M. (2020) Moving the goalposts, *forthcoming in Journal of Political Economy*.
- Frick, M. and Ishii, Y. (2018) Innovation adoption by forward-looking social learners, *Working Paper*.
- Gentzkow, M. and Kamenica, E. (2016) Competition in persuasion, *The Review of Economic Studies*, **84**, 300–322.
- Gratton, G., Holden, R. and Kolotilin, A. (2017) When to drop a bombshell, *The Review of Economic Studies*, **85**, 2139–2172.
- Henry, E. and Ottaviani, M. (2019) Research and the approval process: the organization of persuasion, *American Economic Review*, **109**, 911–55.
- Honryo, T. (2018) Dynamic persuasion, *Journal of Economic Theory*, **178**, 36–58.
- Kamenica, E. and Gentzkow, M. (2011) Bayesian Persuasion, *American Economic Review*, **101**, 2590–2615.

- Keller, G. and Rady, S. (2015) Breakdowns, *Theoretical Economics*, **10**, 175–202.
- Keller, G., Rady, S. and Cripps, M. (2005) Strategic experimentation with exponential bandits, *Econometrica*, **73**, 39–68.
- Kolotilin, A. (2018) Optimal information disclosure: A linear programming approach, *Theoretical Economics*, **13**, 607–635.
- Li, F. and Norman, P. (2018) Sequential persuasion, *Working Paper*.
- Orlov, D., Skrzypacz, A. and Zryumov, P. (2020) Persuading the principal to wait.
- Rothschild, M. (1974) A two-armed bandit theory of market pricing, *Journal of Economic Theory*, **9**, 185–202.
- Smolin, A. (2020) Dynamic evaluation design, *Mimeo*.
- Wald, A. (1947) Sequential analysis, *John Wiley and Sons*.
- Zhao, W., Renou, L. and Tomala, T. (2020) Contracting over persistent information.

Online Appendix to “Dynamic Persuasion with
Outside Information” by Jacopo Bizzotto,
Jesper Rüdiger and Adrien Vigier

Online Appendix 1: Frequent signals

Proof of Proposition 2 of the main text: We will prove part (ii) of the proposition (the proof of part (i) uses arguments similar to those used in the proof of Lemma 7 of the main text and is therefore omitted). Let $b_{T_n-1,n}$ denote the agent's period- (T_n-1) standard of approval given signal frequency n . We now show that $b_{T_n-1,n}$ is monotonically increasing in n . Fix n . By Lemma A.1 of Appendix A, in equilibrium the agent does not benefit from the period- T_n experiment. So $b_{T_n-1,n}$ is independent of information generated by the principal. However, with signal frequency $n+1$ the amount of exogenous outside information that the agent obtains by waiting two periods is the same as what he obtains with signal frequency n by waiting one period. These remarks imply

$$b_{T_{n+1}-2,n+1} \geq b_{T_n-1,n}.$$

By Lemma 5,

$$b_{T_{n+1}-2,n+1} = b_{T_{n+1}-1,n+1}.$$

Hence,

$$b_{T_{n+1}-1,n+1} \geq b_{T_n-1,n}.$$

In what follows, let $\bar{b} := \lim_{n \rightarrow \infty} b_{T_n-1,n}$. The condition $\lambda > \varphi(r)$ is equivalent to $b_1 > \underline{b}$. So whenever this condition holds, $\bar{b} > \underline{b}$. But then, for L sufficiently small, for all sufficiently large n being aggressive in the first period cannot be optimal for the principal: keeping uncertainty high until the last period enables the principal to benefit (at a very small cost) from a strictly lower standard of approval. ■

Online Appendix 2: Different Discount Factors

In this appendix we prove Proposition 3 of the main text. All the results in this appendix refer to the setting in which the players' discount rates are δ_P and δ_A .

Lemma OA.1. *In equilibrium, at $t=T$, the agent accepts if $q_T \geq \underline{b}$ and rejects otherwise. The principal designs the experiment*

$$M_T = \begin{cases} \{0, \underline{b}\} & \text{if } p_T \in (0, \underline{b}); \\ \{p_T\} & \text{otherwise.} \end{cases}$$

The agent does not benefit from the period- T experiment, hence his equilibrium continuation payoff is convex in p_T .

Proof: See the proof of Lemma A.1 in Appendix A. ■

Lemma OA.2. *Let $t < T$. Suppose that functions $\hat{g}_{t+1}(p_{t+1})$ and $\hat{f}_{t+1}(p_{t+1})$ uniquely determine the agent's (resp. the principal's) equilibrium continuation payoffs in period $t+1$. If \hat{g}_{t+1} is convex, then:*

1. *in equilibrium, the principal's period- t experiment and the agent's period- t decision are both uniquely determined; the former is a function of p_t only and the latter is a function of q_t only;*
2. *functions $\hat{g}_t(p_t)$ and $\hat{f}_t(p_t)$ uniquely determine the equilibrium continuation payoffs in period t , and \hat{g}_t is convex.*

Proof: Let $\tilde{g}_t(q_t)$ and $\tilde{f}_t(q_t)$ be defined as in equations (3) and (4) of the main text, respectively. Then the agent's equilibrium continuation payoff given q_t can be written as

$$g_t(q_t) = \max \{V_R, \delta_A \tilde{g}_t(q_t), V_B + q_t(V_G - V_B)\}. \quad (\text{OA.1})$$

As \hat{g}_{t+1} is convex by assumption, Proposition A.1 shows that \tilde{g}_t is convex as well. Moreover,

$$\begin{cases} \delta_A \tilde{g}_t(0) = \delta_A V_R < V_R; \\ \delta_A \tilde{g}_t(1) = \delta_A V_G < V_G. \end{cases} \quad (\text{OA.2})$$

Then (OA.1), (OA.2) and convexity of \tilde{g}_t give unique a_t and b_t , with $a_t \leq \underline{b} \leq b_t$, such that

$$\begin{cases} g_t(q_t) = V_R > \max \{ \delta_A \tilde{g}_t(q_t), V_B + q_t(V_G - V_B) \} & \text{if } q_t < a_t; \\ g_t(q_t) = \delta_A \tilde{g}_t(q_t) > \max \{ V_R, V_B + q_t(V_G - V_B) \} & \text{if } q_t \in (a_t, b_t); \\ g_t(q_t) = V_B + q_t(V_G - V_B) > \max \{ V_R, \delta_A \tilde{g}_t(q_t) \} & \text{if } q_t > b_t. \end{cases}$$

Hence, in equilibrium, the agent rejects if $q_t < a_t$, waits if $q_t \in (a_t, b_t)$, and accepts if $q_t > b_t$. Moreover, since in equilibrium whenever indifferent the agent makes the decision preferred by the principal, it ensues that the agent waits if $q_t = a_t < b_t$ and accepts if $q_t = b_t$. This gives

$$f_t(q_t) = \begin{cases} 0 & \text{if } q_t < a_t; \\ \delta_P \tilde{f}_t(q_t) & \text{if } q_t \in [a_t, b_t); \\ 1 & \text{if } q_t \geq b_t. \end{cases} \quad (\text{OA.3})$$

Standard arguments yield $\hat{f}_t = \text{cav } f_t$. Since in equilibrium whenever indifferent the principal picks the least informative experiment, the principal's experiment in period t is uniquely determined by the belief p_t at the beginning of period t . Lastly, letting $\tau_t(p_t)$ denote the principal's equilibrium experiment given p_t yields $\hat{g}_t(p_t) = \mathbb{E}_{\tau_t(p_t)}[g_t(q_t)|p_t]$.

Finally, since \tilde{g}_t is convex, (OA.1) shows that g_t is as well. Since $\hat{g}_t(p_t) = \mathbb{E}_{\tau_t(p_t)}[g_t(q_t)|p_t]$, convexity of g_t together with the properties of $\tau_t(\cdot)$ establish that \hat{g}_t is convex. ■

Proposition OA.1. *There exists a unique equilibrium.*

Proof: The proposition follows from Lemmata OA.1 and OA.2. ■

Lemma OA.3. *Each period, cutoffs $0 < a_t \leq b_t < 1$ exist such that, in equilibrium the agent rejects if $q_t < a_t$, waits if $q_t \in [a_t, b_t)$, and accepts if $q_t \geq b_t$.*

Proof: For $t = T$, the result follows from Lemma OA.1. For $t < T$, the result was shown within the proof of Lemma OA.2. ■

Lemma OA.4. *At all $t \leq T-1$, the threshold of acceptance is at least as large as the threshold of acceptance at $t=T-1$, which itself is at least as large as the static threshold of acceptance: $b_t \geq b_{T-1} \geq b_T = \underline{b}$.*

Proof: See the proof of Lemma B.1 in Appendix B. ■

Lemma OA.5. *Each period, in equilibrium $M_t = \{p_t\}$ for all $p_t \in \{0\} \cup [b_t, 1]$. Moreover:*

(i) *either $M_t = \{0, b_t\}$ for all $p_t \in (0, b_t)$,*

(ii) *or $a_t < b_t$ and: $M_t = \{0, a_t\}$ for $p_t \in (0, a_t)$ and there exists $c_t \in [a_t, b_t)$ such that $M_t = \{p_t\}$ for $p_t \in [a_t, c_t]$, while $M_t = \{c_t, b_t\}$ for $p_t \in (c_t, b_t)$.*

Proof: Recall $\hat{f}_t = \text{cav } f_t$. If $a_t = b_t$ (so that the set of beliefs for which in equilibrium the agent waits in period t is empty) then in equilibrium $M_t = \{0, b_t\}$ for all $p_t \in (0, b_t)$. So assume $a_t < b_t$. Observe that:

(A) $\tilde{f}_t(\cdot)$ (defined by equation (4) in Appendix A) is concave,

(B) $f_t(q_t) = \delta_P \tilde{f}_t(q_t)$ for all $q_t \in \{0\} \cup [a_t, b_t)$.

(A) follows from Proposition A.1 while (B) is obtained from (OA.3). In view of (A)-(B), either (i) in the statement of the lemma holds, or (ii) does. ■

Lemma OA.6. *With perfect good news, in equilibrium the principal is aggressive at $t=1$.*

Proof: The proof follows the same steps of the proof of Lemma 4 of the main text. ■

Lemma OA.7. *With perfect bad news, either $a_t = b_t = \underline{b}$ in all periods, or for all $t < T$: $b_t = b_{T-1} > b_T = \underline{b}$.*

Proof: If $a_{T-1} = b_{T-1} = b_T$, the claim of the lemma is straightforward. Assume now $b_{T-1} > b_T$. At $q_{T-1} = b_{T-1}$, the agent is indifferent between waiting and accepting. The agent's expected payoff from accepting is $b_{T-1}V_G + (1 - b_{T-1})V_B$. On the other hand, using Lemmata OA.3 and OA.4 (and noting that in equilibrium, in period T , the agent accepts for $p_T \geq b_T$), the agent's

expected payoff from waiting can be written as $\delta_A[b_{T-1}V_G + (1-b_{T-1})(\gamma V_R + (1-\gamma)V_B)]$. So b_{T-1} is the unique solution of

$$xV_G + (1-x)V_B = \delta_A[xV_G + (1-x)(\gamma V_R + (1-\gamma)V_B)]. \quad (\text{OA.4})$$

Next, consider $t < T-1$ such that $b_{t+1} = b_{T-1}$, and $q_t = b_t$, so that, by definition, the agent is indifferent between waiting and accepting. The agent's expected payoff from accepting is $b_t V_G + (1-b_t)V_B$. On the other hand, using Lemma OA.4, $q_t = b_t \geq b_{T-1} = b_{t+1}$. Hence, conditional on $s_t = g$, the agent optimally accepts in the next period. It ensues that b_t solves (OA.4) and, therefore, that $b_t = b_{T-1}$. A recursive argument then yields $b_t = b_{T-1}$ for all $t < T$. ■

Lemma OA.8. *Let $\tilde{\gamma}(\delta_A) := \left(\frac{1-\delta_A}{\delta_A}\right) \frac{V_R(V_G-V_B)}{(V_G-V_R)(V_R-V_B)}$. Then $\gamma > \tilde{\gamma}(\delta_A)$ if and only if $b_{T-1} > \underline{b}$, and either condition implies*

$$b_{T-1} = \frac{\delta_A(\gamma V_R + (1-\gamma)V_B) - V_B}{(V_G - V_B)(1 - \delta_A) + \delta_A \gamma (V_R - V_B)}. \quad (\text{OA.5})$$

Proof: The arguments in the proof of Lemma OA.2 show that $b_{T-1} > \underline{b}$ if and only if given $q_{T-1} = \underline{b}$ the agent strictly prefers waiting over rejection, that is, if and only if

$$\delta_A [\underline{b}V_G + (1-\underline{b})(\gamma V_R + (1-\gamma)V_B)] > V_R,$$

which, upon rearrangement, yields $\gamma > \tilde{\gamma}(\delta)$. Solving (OA.4) yields (OA.5). ■

Lemma OA.9. *Let $t < T-1$. In equilibrium, if the principal is aggressive in period $t+1$, then the principal is also aggressive in period t .*

Proof: Suppose that in equilibrium the principal is aggressive in period $1 < t+1 < T$. If $a_t = b_t$ the statement of the lemma is straightforward. Assume therefore that $a_t < b_t$. By virtue of Lemma OA.5, in order to establish that the principal is also aggressive in period t it is enough to show that when $p_t = a_t$ the principal strictly prefers the experiment $M_t = \{0, b_t\}$ over the uninformative experiment. On one hand, the principal's expected payoff from designing $M_t = \{0, b_t\}$ is $\frac{a_t}{b_t}$. On the other hand, her expected payoff from designing the uninformative

experiment is given by $\delta_P \mathbb{E}_{s_t}[\hat{f}_{t+1}(p_{t+1}) | q_t = a_t]$. The next sequence of inequalities therefore concludes the proof:

$$\delta_P \mathbb{E}_{s_t}[\hat{f}_{t+1}(p_{t+1}) | q_t = a_t] \leq \delta_P \hat{f}_{t+1}(a_t) = \delta_P \frac{a_t}{b_{t+1}} < \frac{a_t}{b_t}.$$

The first inequality follows from noting that \hat{f}_{t+1} is concave (which we show formally in the appendix); the equality follows from the assumption that the principal is aggressive in period $t+1$, and the second inequality is due to Lemma OA.7. ■

Lemma OA.10. *There exists $\tilde{T}(\gamma, \delta_A, \delta_P) < \infty$ such that, in equilibrium, the principal is aggressive in period 1 if and only if $T > \tilde{T}(\gamma, \delta_A, \delta_P)$.*

Proof: Note that in view of Lemma OA.9 it is enough to show that, for T sufficiently large, in equilibrium the principal is aggressive at $t=1$. Next, part (ii) of Lemma OA.7 shows that any benefit to the principal from not being aggressive at $t=1$ must come from persuading the agent to accept at $t=T$ when $\omega = B$. So these benefits are bounded from above by δ_P^{T-1} , which tends to 0 as $T \rightarrow \infty$. On the other hand, as $b_1 < 1$, the corresponding loss to the principal is bounded away from zero since by being aggressive at $t=1$ the principal obtains acceptance with strictly positive probability conditional on $\omega = B$. We conclude that, for T sufficiently large, in equilibrium the principal is aggressive at $t=1$. ■

Lemma OA.11. *In equilibrium, each period either the principal is aggressive, or the principal is conservative.*

Proof: The proof follows the same steps as the proof of Lemma 3 of the main text. ■

Lemma OA.12. *There exist cutoffs $\underline{\delta}_A \in (0, 1)$ and $\underline{\delta}_P(\delta_A) \in (0, 1]$ and, for $\delta_A > \underline{\delta}_A$ and $\delta_P > \underline{\delta}_P(\delta_A)$, functions $0 < \underline{\gamma}(\delta_A, \delta_P) < \bar{\gamma}(\delta_A, \delta_P) < 1$ such that, in equilibrium, the principal is conservative in period $T-1$ if and only if $\delta_A > \underline{\delta}_A$, $\delta_P > \underline{\delta}_P(\delta_A)$, and $\gamma \in (\underline{\gamma}(\delta_A, \delta_P), \bar{\gamma}(\delta_A, \delta_P))$.*

Proof: We saw in the proof of Lemma OA.8 that $\gamma \leq \tilde{\gamma}(\delta_A)$ implies that in equilibrium the agent never waits. So whenever $\gamma \leq \tilde{\gamma}(\delta_A)$, in equilibrium the principal has to be aggressive in

period $T - 1$. In particular, since $\tilde{\gamma}(\delta_A) > 1$ for δ_A sufficiently small, we find that for δ_A small enough the principal is aggressive in period $T - 1$ irrespective of γ and of δ_P .

Suppose next that $\gamma > \tilde{\gamma}(\delta_A)$. Then for $q_{T-1} = a_{T-1}$ in equilibrium the agent is indifferent between waiting and rejection. The agent's expected payoff from rejection is given by V_R . His expected payoff from waiting is on the other hand given by $\delta_A[a_{T-1}V_G + (1 - a_{T-1})(\gamma V_R + (1 - \gamma)V_B)]$, where we deduced from Lemma OA.1 that $s_{T-1} = g$ induces $p_T > b_T = \underline{b}$. We therefore obtain $V_R = \delta_A[a_{T-1}V_G + (1 - a_{T-1})(\gamma V_R + (1 - \gamma)V_B)]$, giving

$$a_{T-1} = \frac{V_R - \delta_A(\gamma V_R + (1 - \gamma)V_B)}{\delta_A(V_G - \gamma V_R - (1 - \gamma)V_B)}. \quad (\text{OA.6})$$

Now, using Lemma OA.5, the necessary and sufficient condition for the principal not to be aggressive in period $T - 1$ in equilibrium is $f_{T-1}(a_{T-1}) \geq \frac{a_{T-1}}{b_{T-1}}$.²⁵ Noting that $f_{T-1}(a_{T-1}) = \delta_P [a_{T-1} + (1 - \gamma)(1 - a_{T-1})]$ and substituting for a_{T-1} and b_{T-1} using (OA.5) and (OA.6), the former inequality becomes

$$\begin{aligned} & V_B \delta_A (\gamma - 1) + \delta_A V_G (1 - \gamma) + V_R \gamma (1 - \delta_A) \\ & \geq \frac{[V_B \delta_A (\gamma - 1) + V_R (1 - \delta_A \gamma)] [\delta_A \gamma (V_B - V_R) + (V_B - V_G) (1 - \delta_A)]}{\delta_P [(V_B - V_R) \delta_A \gamma + V_B (1 - \delta_A)]}. \end{aligned} \quad (\text{OA.7})$$

One checks that if (OA.7) holds for some δ'_P , it must hold for $\delta''_P > \delta'_P$: either the right-hand side is positive, and therefore decreasing in δ_P , or it is negative, but the left-hand side is always positive,²⁶ so in this case the inequality is always satisfied. Moreover, for $\delta_A = 1$ the quadratic equation in γ obtained from (OA.7) has roots $\gamma = 0$ and $\gamma = 1$. On the other hand, for $\delta_A < 1$, (OA.7) is violated whenever either $\gamma = 1$, or $\gamma = \tilde{\gamma}(\delta_A)$. So (OA.7) holds for all values of γ in between the roots of the corresponding quadratic equation. Letting $\underline{\gamma}(\delta_A, \delta_P)$ and $\bar{\gamma}(\delta_A, \delta_P)$ denote the real roots, the previous remarks yield $\tilde{\gamma}(\delta_A, \delta_P) < \underline{\gamma}(\delta_A, \delta_P) \leq \bar{\gamma}(\delta_A, \delta_P) < 1$ and show that these roots only exist for $\delta_A > \underline{\delta}_A$ and $\delta_P > \underline{\delta}_P(\delta_A)$, where (i) $\underline{\delta}_A$ is defined implicitly by $\underline{\gamma}(\underline{\delta}_A, 1) = \bar{\gamma}(\underline{\delta}_A, 1)$ and (ii) $\underline{\delta}_P(\delta_A)$ is defined implicitly for $\delta_A > \underline{\delta}_A$ by $\underline{\gamma}(\underline{\delta}_P(\delta_A), \delta_A) = \bar{\gamma}(\underline{\delta}_P(\delta_A), \delta_A)$. Noting that, by Lemma OA.11, whenever the principal is not aggressive she is conservative concludes the proof. ■

²⁵That is, at $p_{T-1} = a_{T-1}$ the principal must prefer the uninformative experiment over $M_{T-1} = \{0, b_{T-1}\}$.

²⁶Since $V_G > V_B$ and $V_R > 0$ imply $V_B \delta_A (\gamma - 1) + \delta_A V_G (1 - \gamma) + V_R \gamma (1 - \delta_A) > \delta_A (1 - \gamma) (V_G - V_B) > 0$.