American Economic Journal: Microeconomics 2021, 13(1): 179–194
https://doi.org/10.1257/mic.20180141

Dynamic Persuasion with Outside Information†

By Jacopo Bizzotto, Jesper Rüdiger, and Adrien Vigier*

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* C73, D82, D83)

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. On the other hand, the agent's payoff from accepting the offer depends on an unknown

*Bizzotto: Oslo Business School, Oslo Metropolitan University, Pilestredet 35, 0166 Oslo, Norway (email: jacopobi@oslomet.no); Rüdiger: Department of Business Administration, Universidad Carlos III de Madrid, Calle Madrid 126, 28903 Getafe, Madrid, Spain (email: jrudiger@emp.uc3m.es); Vigier: Department of Economics, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway (email: a.h.vigier@gmail.com). John Asker was coeditor for this article. 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 errors are our own.
†Go to https://doi.org/10.1257/mic.20180141 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

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 follows. 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 persuade 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; in

[1] We show that waiting is socially wasteful even if the agent's payoff from a type-II error is negative.
[2] The "belief" refers to the probability attached to the state in which accepting the offer is optimal for the agent.
[3] The lower the threshold of acceptance, the more type-II errors the principal can induce the agent to make.

this case, we say 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; in this case, we say 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 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, which 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 I. The core of the analysis is in Section II. Several extensions of the model are examined in Section III. Section IV 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 strands of research.

The first strand of research examines the case in which multiple senders compete to persuade the agent (Gentzkow and Kamenica 2017, Li and Norman 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, Skrzypacz, and Zryumov (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 (2017); and Zhao et al. (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 nonmonotonic function of signal accuracy is linked in spirit to a related result in Kolotilin (2018). Gratton, Holden, and Kolotilin (2018) 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 ours, analyzes 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.

I. 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]

Timing.—The state of the world is drawn by nature according to $\Pr(\omega = G) = p_1$. We suppose for expository purposes that the agent initially leans toward rejection, that is, $p_1 \in (0, \underline{b})$, where $\underline{b} := (V_R - V_B)/(V_G - V_B)$ 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.

[4] Section IIIB examines the case in which players discount time differently.

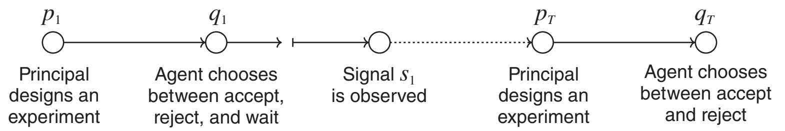

FIGURE 1. TIMING

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: $E_{\tau_t}[q_t] = p_t$. 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 upward 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 downward 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.

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 $\{\text{accept}, \text{reject}, \text{wait}\}$ for $t < T$, and to a decision in $\{\text{accept}, \text{reject}\}$ for $t = T$. The equilibrium concept is Perfect Bayesian Equilibrium (PBE): the player at each decision node maximizes her/his expected payoff conditional on (i) the other player's strategy and (ii) the belief obtained using Bayes' rule.

II. Analysis

Section IIA characterizes the equilibrium of our game. A general discussion of the main theorem is provided in Section IIB, and a sketch of its proof is presented in Section IIC. All omitted proofs of this section are in online Appendices A, B, and C.

[5] See, for instance, Keller and Rady (2015), Bonatti and Hörner (2017), or Frick and Ishii (2020).

A. *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.*

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

[6] For 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.

[7] For 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 $q_t = x$.

[8] If $q_t \geq b_t$ (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.

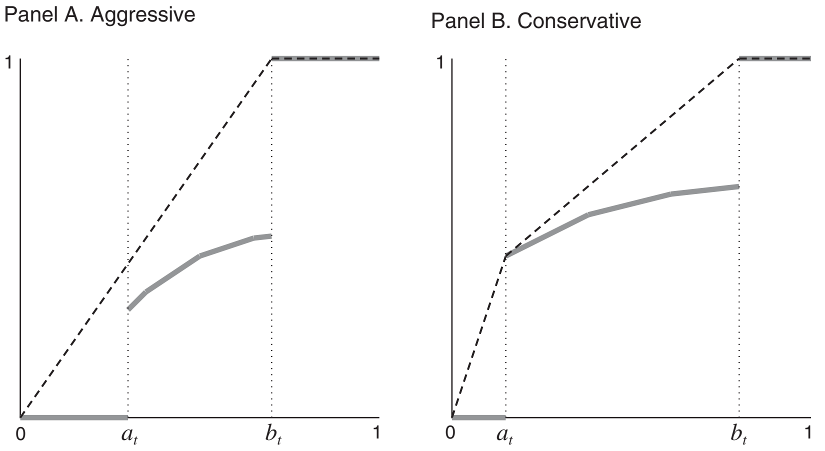

FIGURE 2. PRINCIPAL'S PERSUASION STRATEGY

and Maschler 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 A. 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 B. In this case, $a_t < b_t$ and

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

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*

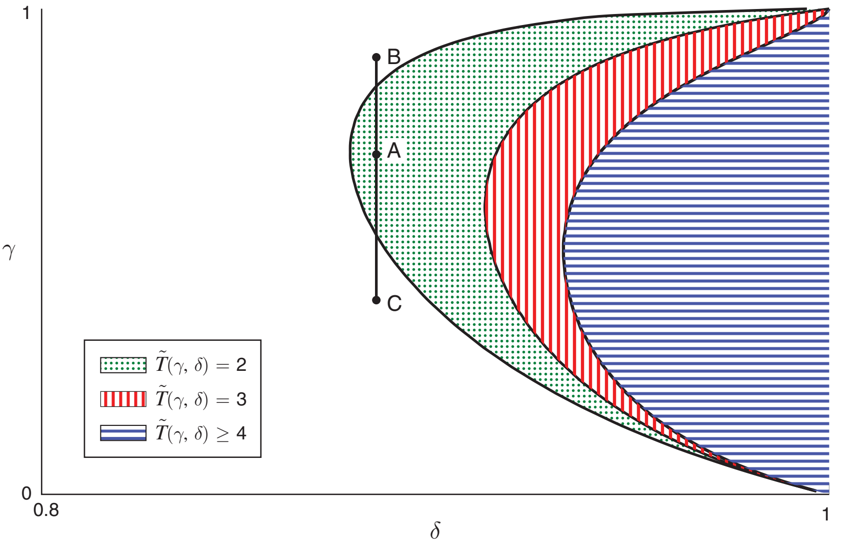

FIGURE 3. THE PERFECT BAD NEWS CASE

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* $\gamma \in (\underline{\gamma}(\delta), \overline{\gamma}(\delta))$, *where* $0 < \underline{\gamma}(\delta) < \overline{\gamma}(\delta) < 1$.

The perfect bad news case is illustrated in Figure 3.[9] For example, each parameter pair (δ, γ) that belongs to the vertically striped 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 tension at the heart of our model and how this tension explains Theorem 1. In Section IIB, 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 section.

Information that 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.

[9]The figure is drawn for $V_G = 2$, $V_R = 1$, and $V_B = 0$. The code is available at Bizzotto, Rüdiger, and Vigier (2021).

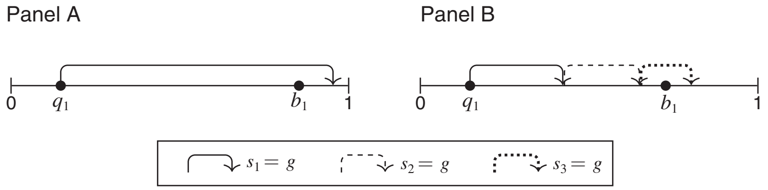

FIGURE 4. THE OVERSHOOTING EFFECT

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.

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 A and panel B 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.

B. *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 (i) accepts with probability 1 conditional on state G and (ii) makes his final decision at $t = 1$ with probability 1.[13] By Lemmas 1 and 3, condition (i) 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 (ii) holds in this case (by Lemma 1). On the

[10] In panel A, to prevent the overshooting, the principal must generate information inducing the agent to accept in the first period with positive probability. In panel B, the principal has more freedom, but to avert overshooting the principal must generate information inducing the agent to accept with positive probability before $t = T$.

[11] To 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 $t = 1$.

[12] Moreover, we show in online 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.

[13] See Proposition C.1 in online Appendix C.

[14] 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 $q_t = 0$.

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, and so in this case (ii) 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 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 $\delta_A = \delta_C$).

C. *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

[15] See Propositions C.2 and C.3 in online Appendix C.

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 OF LEMMA 4:

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

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

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 $E[\phi(X)]$. On the other hand, as $\delta < 1$, her expected payoff from designing the uninformative experiment is bounded from above by $E[\phi(Y)]$. Since ϕ is concave, $E[\phi(Y)] \leq E[\phi(X)]$. 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 $t < T$.

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

[17] The 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.

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]

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

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

III. Extensions

A. *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

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

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 $\Delta_n = 1/2^{n-1}$ 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 I) is obtained by setting $\gamma_n = 1 - e^{-\lambda\Delta_n}$, and $\delta_n = e^{-r\Delta_n}$. Relabeling appropriately, the analysis in Section II 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 ($\lambda_n > \overline{\lambda}_n$), and 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 in online Appendix C).[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 $\lim_{n\to\infty} \tilde{L}_n = 0$ (respectively, $\lim_{n\to\infty} \tilde{L}_n = \infty$) 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 D) 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

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

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* $L > L^\dagger$;

[19] Keeping L fixed, increasing n also increases T_n.

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

B. *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,* $\gamma \in (\underline{\gamma}(\delta_A, \delta_P), \bar{\gamma}(\delta_A, \delta_P))$; *and*
- *both players are sufficiently patient, that is,* $\delta_A > \underline{\delta}_A$, *and* $\delta_P > \underline{\delta}_P(\delta_A)$,

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 E. 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 $t = 1$.

C. *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 $U_R := C\delta/(1-\delta)$ 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]

IV. 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 nonstationary environment in which the agent's threshold of acceptance evolves over time, providing 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 in which, in equilibrium, the agent makes his final decision in the first period, and those in which the agent sometimes waits until the deadline. We 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.

REFERENCES

Au, Pak Hung. 2015. "Dynamic Information Disclosure." *RAND Journal of Economics* 46 (4): 791–823.

Aumann, Robert J., and Michael B. Maschler. 1995. *Repeated Games with Incomplete Information*. Cambridge: MIT Press.

Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier. 2021. "Replication data for: Dynamic Persuasion with Outside Information." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E117901V1.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica* 67 (2): 349–74.

Bonatti, Alessandro, and Johannes Hörner. 2017. "Learning to Disagree in a Game of Experimentation." *Journal of Economic Theory* 169: 234–69.

Che, Yeon-Koo, Kyungmin Kim, and Konrad Mierendorff. 2020. "Keeping the Listener Engaged: A Dynamic Model of Bayesian Persuasion." https://arxiv.org/abs/2003.07338.

Ely, Jeffrey C., and Martin Szydlowski. 2020. "Moving the Goalposts." *Journal of Political Economy* 128 (2): 468–506.

Frick, Mira, and Yuhta Ishii. 2020. "Innovation Adoption by Forward-Looking Social Learners." Unpublished.

Gentzkow, Matthew, and Emir Kamenica. 2017. "Competition in Persuasion." *Review of Economic Studies* 84 (1): 300–22.

Gratton, Gabriele, Richard Holden, and Anton Kolotilin. 2018. "When to Drop a Bombshell." *Review of Economic Studies* 85 (4): 2139–72.

Henry, Emeric, and Marco Ottaviani. 2019. "Research and the Approval Process: The Organization of Persuasion." *American Economic Review* 109 (3): 911–55.

Honryo, Takakazu. 2018. "Dynamic Persuasion." *Journal of Economic Theory* 178: 36–58.

[20] The 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 (available upon request).

[21] Without 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.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (6): 2590–2615.

Keller, Godfrey, and Sven Rady. 2015. "Breakdowns." *Theoretical Economics* 10 (1): 175–202.

Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73 (1): 39–68.

Kolotilin, Anton. 2018. "Optimal Information Disclosure: A Linear Programming Approach." *Theoretical Economics* 13 (2): 607–35.

Li, Fei, and Peter Norman. 2018. "Sequential Persuasion." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2952606.

Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov. 2020. "Persuading the Principal to Wait." *Journal of Political Economy* 128 (7): 2502–41.

Rothschild, Michael. 1974. "A Two-Armed Bandit Theory of Market Pricing." *Journal of Economic Theory* 9 (2): 185–202.

Smolin, Alex. 2017. "Dynamic Evaluation Design." https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3051703.

Wald, Abraham. 1947. *Sequential Analysis*. Hoboken, NJ: John Wiley and Sons.

Zhao, Wei, Claudio Mezzetti, Ludovic Renou, and Tristan Tomala. 2020. "Contracting over Persistent Information." Unpublished.