Persuading the Principal to Wait

Dmitry Orlov

University of Wisconsin–Madison

Andrzej Skrzypacz

Stanford University

Pavel Zryumov

University of Rochester

A principal decides when to exercise a real option. A biased agent influences this decision by strategically disclosing information. Committing to disclose all information with a delay is the optimal way to persuade the principal to wait. Without dynamic commitment, this promise is credible only if the agent's bias is small; otherwise, he pipets information, probabilistically delaying the principal's action. When the agent is biased toward early exercise, his lack of commitment to remain quiet leads to immediate disclosure, hurting him. Our model applies to pharmaceutical companies conducting clinical trials to influence the Food and Drug Administration or equipment manufacturers testing their products.

I. Introduction

Decision-makers commonly rely on interested parties to provide relevant information. In turn, agents use strategic communication to influence decision-makers. Conflicts arise when preferences over decisions are not

We thank Mike Harrison, Johannes Hörner, Emir Kamenica, Ron Kaniel, Erik Madsen, Heikki Rantakari, Nina Varsava, and seminar participants at the University of Rochester, the University of Pennsylvania, New York University, the University of Chicago, Chicago Booth, the Wharton School, the Fuqua School of Business, HEC (Faculté des hautes études

Electronically published June 2, 2020
[*Journal of Political Economy*, 2020, vol. 128, no. 7]

aligned. We find that agents can maintain strategic ignorance and delay information acquisition and disclosure to their advantage. For example, a pharmaceutical company can successfully manipulate its regulator (the Food and Drug Administration [FDA]) to keep their drug on the market by strategically designing and timing clinical trials. We show how spreading out such trials across time is valuable for the firm and how delegating such trials to a third party, in order to obtain a commitment to future testing, can further increase profits. Our model also applies to managers deciding what evidence to acquire to convince headquarters to launch a product or keep an existing one on the market, and to equipment manufacturers performing repeated safety tests to influence buyers.

We model such strategic interaction as a game of dynamic persuasion of a principal (she, receiver) by a biased agent (he, sender) in the context of real options. We contrast the equilibrium of the dynamic persuasion game, in which the agent can commit to test design within a period but not across periods, with the optimal dynamic persuasion mechanism, in which the agent can also commit to future test design. As we show, when the agent is biased toward late exercise, he is able to beneficially persuade the principal to wait even if he cannot commit to future information disclosures. Our first result is that when this bias is large, in equilibrium the agent reveals information gradually ("pipets" information), as opposed to revealing all information at once but with a delay, as is optimal under commitment. Our second result is that when the bias is small, the equilibrium of the game coincides with the optimal mechanism. The principal benefits from information only in the small-bias case. The agent benefits in both cases. Finally, we show how the direction of the conflict affects equilibrium persuasion. When the agent is biased toward early exercise, in equilibrium, he reveals all information immediately. Not only is the agent's noncommitment payoff smaller than in the optimal mechanism, but his equilibrium payoff can be even lower than in a game with no information disclosure. The ability to control information can hurt the agent.

In a real-option problem, a principal decides in every period whether or not to take an irreversible action; that is, she decides when to exercise the option. The payoff from taking the action in our model depends on two states. The first state, which we denote by X_t, is public and evolves exogenously via a geometric Brownian motion. The second state, θ, is binary and initially unobserved by either player. The realized payoff from exercising the option is $\theta X_t - I$.[1] We extend this classic single-player setting

commerciales) Lausanne and EPFL (École polytechnique fédérale de Lausanne), Boston University, the Midwest Economic Theory meeting, the North American Econometric Society meeting, SITE (Stanford Institute for Theoretical Economics), and the European Econometric Society meeting for useful comments and suggestions.

[1] The principal's problem, without any information from the agent, is based on models such as McDonald and Siegel (1986), Stokey (2008), and Dixit and Pindyck (2012).

by adding a strategic agent who can acquire and disclose additional verifiable information about the second state, θ. The agent's incentives are not fully aligned with those of the principal: he may prefer to either wait longer or act sooner than the principal. The agent strategically chooses what information to acquire and disclose in order to influence the principal's decision, that is, to persuade her when to exercise the option.

A. *Persuasion to Wait*

We first characterize a unique Markov perfect equilibrium (MPE) when the agent is biased toward late exercise, that is, when he wants to persuade the principal to wait, as is often the case when firms persuade their regulators to delay interventions. Information disclosed by the agent has a dual effect on the principal's behavior. On the one hand, new information affects her contemporaneous beliefs and, as a consequence, her immediate action. On the other hand, the anticipation of future information acts as an incentive device because it increases the benefit of waiting.[2]

In order to develop the intuition for our main results, consider a two-period example.[3] The first period begins with the agent disclosing some information that affects beliefs of the principal about θ, and it concludes with the principal deciding whether to exercise the option or to wait until the second period. If the principal chooses to wait in the first period, the public state stochastically changes from X_1 to X_2. Upon observing X_2, the agent can disclose more information, and after that, the principal makes the final decision of whether to exercise the option or let it expire. Suppose that the agent is strongly biased and wants the principal to wait regardless of the state. When the game reaches the second period, the agent designs a signal about θ to minimize the probability of the principal exercising the option, which is very similar to the judge-prosecutor example presented in Kamenica and Gentzkow (2011). Specifically, if the principal's expected exercise payoff is negative or zero, the agent stays quiet. Otherwise, he conducts a partially informative test, which either increases the expected exercise value or reduces it to zero. In the latter case, the principal is indifferent between letting the option expire and exercising it and does the former in equilibrium. As in Kamenica and Gentzkow (2011), such persuasion does not affect the principal's expected payoff entering the second period. Going back to the first period, the principal anticipates that her equilibrium continuation payoff is the same as if she were to receive no information from the agent. Her optimal first-period policy then takes into account the option value of waiting for the public information but not the possibility of receiving additional information

[2] Such dual use of information also plays a central role in Ely and Szydlowski (2020).
[3] See online app. A.6 for a formal treatment of the two-period model.

from the agent (we refer to it as the "autarky strategy"). In turn, the agent's optimal first-period strategy minimizes the probability that the principal exercises the option in the first period. The agent pipets information: when the principal's beliefs are such that she does not exercise the option, the agent stays quiet. If, however, principal's beliefs are such that she is about to act, the agent induces a binary distribution of posteriors with the property that after a negative signal the principal is indifferent between waiting and acting. Again, the principal does not benefit ex ante from such information disclosure.

This dynamic behavior extends to any number of periods by induction: in every period the principal follows her autarky strategy (which depends on the number of remaining periods and the current beliefs) and the agent pipets information.[4] The equilibrium payoff of the principal is equal to her autarky payoff, while the agent's payoff is strictly higher. Even though a promise of future information disclosure by the agent could serve as a carrot in motivating the principal to wait, the agent is not able to utilize this incentive device in equilibrium because of the lack of dynamic commitment. If the agent could commit to future disclosure of information, he would optimally incentivize the principal to wait in the first period by promising to disclose more information in the second period. In this setting, pipetting is dominated by delayed disclosure when the latter is credible.

Now reconsider the two-period example, but suppose that the preferences of the principal and the agent are not entirely misaligned. That is, for some states, the agent and the principal agree on whether the option should be exercised. Then, for those states, the agent will find it optimal to disclose all available information in the second period. When entering period 2, such information strictly increases the principal's expected payoff over the autarky benchmark. As a consequence, she is less willing to exercise the option in the first period. This reduces the need for the agent to reveal information in the first period. Under partial alignment of preferences, delayed disclosure of valuable information becomes credible and can even replicate the full-commitment outcome.

In the infinite-horizon model,[5] the magnitude of the conflict between the agent and the principal endogenously depends on beliefs about θ, and the (unique) equilibrium is characterized by a combination of pipetting and delayed disclosure. As a result, the principal's equilibrium policy is often nonmonotone in the expected value of θX_t, unlike in the two-period model. The intuition is that under partial alignment of preferences,

[4] This intuition also extends to the MPE of our infinite-horizon continuous-time game.

[5] The model features an infinite time horizon to ensure that the equilibrium behavior is not affected by the approaching deadline of the option. Additionally, we choose time to be continuous for tractability, as is standard in the real-options literature.

the agent finds it optimal to disclose all information about θ when X_t is sufficiently high. In the vicinity of such states, it is optimal for the principal to wait to learn θ. Similarly, when X_t is sufficiently low, the principal waits, since she never acts sooner than her autarky threshold. However, for intermediate expected exercise payoffs, near the principal's autarky threshold, the expected cost of waiting for the agent's information disclosure is too high, and she prefers to act. The agent's best response is to pipet information in order to minimize the likelihood of early option exercise. On the equilibrium path, the agent may start by pipetting information for a while to prevent early exercise, gradually reducing the belief about θ conditional on waiting. As this happens, the principal becomes increasingly more pessimistic about exercising the option and eventually chooses to wait for the agent to disclose θ fully. The agent can, thus, credibly delay full disclosure of information until his preferred exercise threshold.

B. *Persuasion to Act*

We contrast these equilibrium outcomes to the case when the agent is biased toward early exercise. Then, there exists an MPE in which the agent immediately acquires and discloses all information. This equilibrium outcome is unique under certain conditions. The intuition behind such a stark result is that there exists a region of beliefs close to the principal's acting region such that the agent would want to release some information to speed up the action. In turn, in anticipation of this information, the principal wants to wait longer. That makes the agent want to disclose the information even sooner, and so on.[6] If the agent deviates and does not disclose information, the principal rationally anticipates that the agent will be too tempted to reveal it in the future, and thus she waits for that disclosure.

Such equilibrium behavior implies that for some priors, the agent's expected payoff is not only strictly less than that with full commitment but even worse than that if the agent could not disclose any information at all. Finally, for some model parameters, we show that, under commitment, the agent discloses only imprecise information that maximizes the probability of option exercise at time 0 and does not communicate afterward.

C. *Applications*

Among other contexts, our model captures information dynamics around product recalls. A concrete example is the postmarket surveillance of drugs

[6] Full information disclosure in equilibrium requires a continuous-time limit, similar to the Coase conjecture.

and medical devices conducted by the FDA. When a product is first introduced to the market, its efficacy and safety are somewhat uncertain. As increasingly more patients use it, information about its effects—both positive and negative—is gradually revealed. The regulator monitors this (exogenous) news and can recall the product, that is, remove it from the market if evidence indicates that it is dangerous. The firm producing the drug can affect the FDA's decision by providing additional information or tests.[7] We should expect a partial misalignment between the firm and the regulator over when to exercise the real option of recall. If the firm does not fully internalize all the costs of a bad drug, it would prefer to wait longer for stronger evidence of side effects than the regulator would. The firm cannot pay the regulator to postpone a recall but can persuade it to wait by designing trials and optimally timing them.[8] We show that without a long-term commitment, it is initially optimal for the firm to engage in strategic ignorance, that is, design noisy testing procedures that have a low chance of uncovering negative effects, in order to persuade the regulator that the drug is sufficiently safe for the market.[9] However, if bad news about the drug accumulates, the firm eventually conducts a highly informative test and, conditional on results confirming the problems with the drug, voluntarily recalls the product without pressure from the FDA. These findings are supported by anecdotal evidence about ongoing trials and subsequent recalls of certain medical drugs (see sec. A.5 of the online appendix). The value of commitment to future trials creates an incentive for firms to delegate the trials to a third party, such as contract research organizations—an organizational structure that is documented in the pharmaceutical industry.[10] Our model predicts that the outsourced trials are, on average, more informative than the in-house ones but are conducted with a significant delay.

Another example, from organizational economics, is the problem of a new-product launch or cancellation of an existing product. An executive in the firm (the principal) has the final decision power. She assesses the size of the market and the consumer's willingness to pay (WTP). Public

[7] See "Postmarket Surveillance Under Section 522 of the Federal Food, Drug, and Cosmetic Act: Guidance for Industry and Food and Drug Administration Staff," issued on May 16, 2016 (https://www.fda.gov/media/81015/download). Under these FDA guidelines, the manufacturer has the opportunity to provide additional information and identify specific surveillance methodologies before the FDA issues a recall.

[8] The FDA does not require controlled clinical trials to address its concerns, but asks for "the minimum amount of information necessary to adequately address a relevant regulatory question or issue through the most efficient manner at the right time." See "The Least Burdensome Provisions: Concept and Principles: Guidance for Industry and Food and Drug Administration Staff" (https://www.fda.gov/media/73188/download) for details.

[9] In sec. VI, we consider a version of the model where the agent (drug producer) is privately informed but can still conduct credible tests. We show that the equilibrium information sharing of the Bayesian persuasion game remains an equilibrium in this alternative setting under reasonable off-equilibrium beliefs.

[10] See Mirowski and Van Horn (2005) for an overview of these institutions and their role in medical innovation.

news arrives over time about the size of the market (we model the evolution of the market size as a geometric Brownian motion). The consumer's WTP, high or low, is unobservable, and both players share a common prior over it. The agent is a product manager whose preferences might not be fully aligned with the principal's (because the agent gets private benefits from managing the product or because he does not fully internalize the cost of launching, etc.). The product manager strategically designs marketing studies that are informative about the WTP and chooses when to conduct them. He understands that good news speeds up the launch, but bad news postpones it and cannot be hidden once acquired. Our results imply that equilibrium communication depends on the direction of the agent's bias. If the product manager wishes to delay the launch, he slowly pipets negative information about the project until a partial agreement with the principal is reached. If the goal is to speed up the product launch, then, absent commitment power, he cannot make use of his superior access to information and ends up acquiring and disclosing all (good and bad) information immediately.

Another application that fits our theory is voluntary testing by firms to persuade potential customers to buy their products. Pricing alone is often of limited use, especially when buyers have concerns about the safety of the product. In such situations, product tests can be used to support sales. For example, the manufacturer of the Taser electrical stun gun conducted partially informative tests to persuade police departments that its products are sufficiently safe and should continue to be used. The common theme in these and many other applications is that both the principal and the agent are rational and forward looking and cannot commit to future actions; transfers are either not allowed or do not fully align incentives;[11] agents can strategically decide to remain ignorant about certain facts, but any information they acquire must be disclosed.

D. *Related Literature*

We study the dynamic interaction of exogenous news and endogenous communication. We model within-period communication as the management of public information by an agent, also known as Bayesian persuasion, first introduced by Aumann and Maschler (1995) and Kamenica and Gentzkow (2011). Several papers study Bayesian persuasion in dynamic settings. Ely (2017) shows that when the agent (sender) communicates with a sequence of short-lived principals (receivers), long-term commitment is essentially not valuable. Renault, Solan, and Vieille (2017) study a similar problem and obtain conditions under which optimal persuasion

[11] In lemma A.1 (sec. A.1 of the online appendix), we show that even if the principal could align preferences by paying the agent a bonus for exercising the option, she optimally chooses to set it to zero.

takes into account only short-term optimality for the agent. In contrast, we show that when both players are long-lived, the commitment and noncommitment solutions are qualitatively different.

A second significant difference is that in both Ely (2017) and Renault, Solan, and Vieille (2017) the optimal dynamic persuasion mechanism is greedy; that is, it maximizes the agent's instantaneous payoff in every period and features gradual revelation of information. In our model, because the receiver is long-lived, the optimal commitment policy uses future information disclosure as an incentive device and hence is neither gradual nor greedy. We show that pipetting of information at the receiver's autarky threshold (which is reminiscent of the greedy policies) can be an equilibrium feature of a dynamic game without commitment. While equilibrium communication under noncommitment resembles the policies in these papers, the economic intuition for optimality of the pipetting strategy is different. In those papers, information cannot be used as a carrot because the principal is short-lived. In our paper, the principal is long-lived, so she could be persuaded to wait for future information. However, the agent may not be able to credibly promise to deliver such information in the future, and hence, as we explained above in the two-period example, the principal ends up following the autarky policy in equilibrium.

Our paper is related to Smolin (2019) and Ely and Szydlowski (2020), since they also study dynamic persuasion between two long-lived players. Their focus is on the commitment solution, and their findings are consistent with ours: when the sender wants to delay the action of the receiver, he can benefit by promising future information disclosure. We complement their analysis by contrasting the commitment and noncommitment solutions, by studying a different type of dynamic game, and by allowing partial alignment of preferences between the principal and the agent. In the paper, we show that the commitment and noncommitment solutions coincide if and only if players' preferences are sufficiently aligned.

Equilibrium unraveling in persuasion to act is reminiscent of Au (2015), where the principal is privately informed about her preferences. Our unraveling result in proposition 5 holds even though the preferences of the principal are known, but the lack of commitment plays an important role in both results. Bizzotto, Rüdiger, and Vigier (2019) is also a closely related paper that studies a finite-horizon persuasion game without commitment. The main finding of that paper is that delayed persuasion may be optimal since, as the deadline of the game approaches, the principal's exercise region expands, benefiting the agent. In our paper, the environment is stationary, and hence the difficulty of persuading the principal depends only on her expectations of equilibrium communication. Henry and Ottaviani (2019) analyze a setting in which the agent collects and reveals costly Brownian signals to influence the principal's rejection or approval decision. Their focus is on providing incentives to the sender, and in their model, the receiver commits to stopping at a specific threshold. In our model, the receiver cannot commit to exercising the option suboptimally, making such arrangements infeasible. Hörner and Skrzypacz (2016) analyze a dynamic persuasion game with noncontractable monetary transfers. They find that gradual persuasion is optimal to resolve the ex post holdup problem, so the economic mechanism for pipetting is different than in our paper.

Our analysis sheds light on the role of verifiable information in dynamic decision-making. Grenadier, Malenko, and Malenko (2016) analyze a closely related model in which the agent has access to information but can use only cheap talk. We show that when the informed agent's bias toward delayed exercise is small, both verifiable and nonverifiable communication lead to delegating the execution of the real option to the agent. However, when the bias is large, the agent cannot credibly convince the principal to wait via cheap talk. With verifiable information, the agent can probabilistically delay the principal's action. Similarly, if the agent is biased toward early exercise, Grenadier, Malenko, and Malenko (2016) show that there does not exist a revealing equilibrium. This is in contrast to our findings, which show that if the agent is biased toward early exercise, it leads to full and immediate information sharing. These results highlight the distinction between communication of soft (unverifiable) information and that of hard (verifiable) information.

More broadly, our paper is related to the literature on agency conflicts in the context of real options, including Grenadier and Wang (2005), Board (2007), Kruse and Strack (2015), and Gryglewicz and Hartman-Glaser (2020). These papers study the role of incentive contracts with monetary transfers in managing conflicts of interest. We are the first to analyze the limits of strategic management of hard information in the context of real options when monetary transfers are not feasible.

The rest of the paper is organized as follows. In section II, we present our main model. In section III, we characterize the unique equilibrium in case the agent is biased toward late exercise and persuades the principal to wait. Propositions 1 and 2 are the main results of the paper. In section IV, we consider the case where the agent is biased toward early exercise and persuades the principal to act. In section V, we extend the model beyond the binary θ case. In section VI, we allow the agent to be privately informed and show that equilibria constructed in sections III and IV are robust to the introduction of private information. Section VII concludes.

II. Model

A. *Basic Setup*

We start with an informal description of the model. The principal chooses the time to make an irreversible decision. The agent strategizes over when and what kind of additional information to publicly acquire and disclose in order to influence the timing of the decision. The players share the costs and benefits of the decision differently; hence, their preferences over the option exercise time are misaligned. There are two reasons why the principal may wish to wait: exogenous innovations in the underlying state and new information about the project that the agent endogenously decides when to acquire and disclose publicly. We cast our model in continuous time to use well-established tools and intuitions from single-agent real-option problems (see Dixit and Pindyck 2012).

1. Players and Payoffs

Time is continuous and infinite, $t \in [0, +\infty)$. There are two long-lived players, a principal (she, receiver) and an agent (he, sender), who discount future payoffs at a rate r. The principal has an irreversible decision to make and chooses the optimal timing of this decision; that is, she faces a real option. The payoff from exercising the option depends on two states. The first state is given by a publicly observable process $X = (X_t)_{t\geq 0}$, which follows a geometric Brownian motion,

$$dX_t = \mu X_t \, dt + \phi X_t \, dB_t,$$

with $\mu < r$ and $\phi \geq 0$. The second state is the underlying quality of the project $\theta \in \{\theta_L, \theta_H\}$, with $\theta_H > \theta_L \geq 0$.[12] Neither party initially observes the realization of θ, and they share a common prior,

$$Y_{0-} = \mathrm{P}(\theta = \theta_H)\,.$$

We model the real-option problem using the classic approach described in Dixit and Pindyck (2012), assuming that X and θ are independent.

The flow payoffs of both players are zero before option exercise. If the principal takes the action at time t, then time 0 discounted payoffs of the agent and the principal conditional on X_t and θ are

$$v_A = e^{-rt}(\theta X_t - I_A) \quad \text{and} \quad v_P = e^{-rt}(\theta X_t - I_P).$$

Parameters I_P and I_A capture the costs of exercising the option, and we assume $I_P > 0$, $I_A \geq 0$. One interpretation is that the option is to launch a product, X, is the observed potential market size and θ is the unobserved willingness of consumers to pay, so that $\theta \cdot X$ is a measure of profits from the launch.

The disagreement between the agent and the principal is driven by the difference in costs of exercising the option, $I_P \neq I_A$. If $I_P < I_A$, then

[12] The case of $\theta_L < 0$ is qualitatively similar to the case of $\theta_L = 0$, since neither player wishes to exercise the option in this state regardless of X_t.

for any given θ, the agent's optimal timing of exercise is later than the principal's. In this case, the agent would like the principal to delay exercise time—he would like to persuade the principal to wait. If $I_P > I_A$, then the direction of the conflict is reversed, and the agent would like to accelerate exercise time—he would like to persuade the principal to act.

2. Remark about Payoffs

We model preferences in terms of the call option, with zero flow costs and only terminal payoffs. In section III.B, we analyze an equivalent formulation of our model, in which the players face flow payoffs and consider a put option to stop. Alternatively, one could model the conflict between the parties by assuming that the discount rates of the principal and agent differ. One can show that our equilibrium results depend only on the direction of the disagreement between the principal and agent and are robust to such model perturbations.

B. Strategies and Equilibrium Concept

Within every "period," the innovation to X_t is first realized, and then the agent provides an informative signal about θ by conducting a test. The test induces a posterior over θ from some distribution, subject to the martingale constraint that the average posterior belief has to be equal to the prior. The agent can commit within a period to an arbitrary distribution, but he cannot commit to future signals. After observing the signal generated by the agent, the principal decides whether to exercise the option or to continue waiting (e.g., whether or not to launch the product). She also cannot commit to future actions. Heuristically, the sequence of events in a short period of time dt is shown in figure 1. The agent's strategy is a function of the past history of the state process X and information learned about θ up to time t. The principal's strategy is, additionally, a function of the information learned about θ at time t.

In our model, the agent controls information flow about θ, which affects the principal's posterior beliefs over time. Denote by $\mathcal{F}_t$ all information available to the players at time t, which includes the path of the process

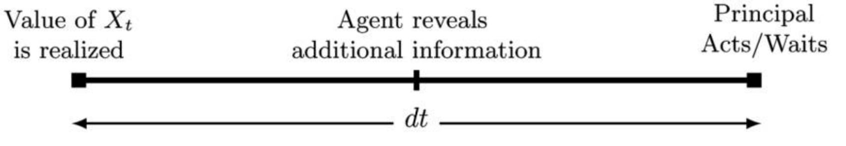

FIG. 1.—Timing of events in a short interval of time.

X as well as any signals communicated by the agent up to time t.[13] Denote by Y_t the posterior belief about θ given by

$$Y_t = \mathrm{P}(\theta = \theta_{\mathrm{H}} | \mathcal{F}_t) .$$

We define a *(Markov) state* of the game to be the pair (X_t, Y_t). That is, the state contains the posterior belief Y_t about the quality of the project θ and the current level of the process X_t. A player's strategy is Markov if it depends on history only via the current levels of state X and belief Y. We allow the agent to continuously generate informative signals whose distributions are contingent on the history of X, realization of θ, and past messages. We require only that information disclosed by the agent at time t is independent of future increments of X to ensure that the belief process Y is "not forward looking"; that is, it does not foresee the future evolution of X.

Definition (Admissible belief process). An *admissible belief process* $Y = (Y_t)_{t \geq 0-}$ is a right-continuous-with-left-limits martingale, with respect to the natural filtration of (X, Y), that takes values in $[0, 1]$ such that Y_t is independent of X's future innovation paths $\{X_{t+s}/X_t\}_{s \geq 0}$ for every $t \geq 0$. We denote by $\mathcal{Y}$ the set of all admissible belief processes.

The agent's strategy is an admissible belief process (i.e., instead of modeling messages the agent sends, we represent the strategy directly in terms of the posterior beliefs). We require the belief process Y to be right-continuous to capture the idea that the agent moves first in every period, as illustrated in figure 1. In particular, we allow the agent to generate a discrete signal right at the beginning of the game; that is, time 0 posterior Y_0 may be different from the initial prior Y_{0-}, before the principal has the first opportunity to exercise the option. The strategy is Markov if the information the agent discloses about θ depends on the history only through the current state (X_t, Y_{t-}).

Definition (Markov strategy of the agent). *The agent's Markov strategy* is an admissible belief process $Y \in \mathcal{Y}$ such that $(X_t, Y_t)_{t \geq 0}$ is a Markov process.

A class of agent's Markov strategies that is important in the equilibrium analysis is disclosure strategies, which reveal whether $\theta = \theta_{\mathrm{H}}$ over time. Define D_t to be the cumulative probability of having disclosed $\theta = \theta_{\mathrm{H}}$ up to time t. Then a disclosure strategy induces posterior belief Y_t, which is either equal to 1 if $\theta = \theta_{\mathrm{H}}$ has been disclosed or equal to Y_t^{ND} given by the Bayes rule if $\theta = \theta_{\mathrm{H}}$ has not been disclosed,

$$Y_t^{\mathrm{ND}} = \frac{Y_{0-}(1 - D_t)}{Y_{0-}(1 - D_t) + 1 - Y_{0-}} . \tag{1}$$

[13] Formally, $\mathcal{F}_t$ is the σ algebra generated by $(X_s)_{s \leq t}$ and the signals disclosed by the agent. For technical reasons, we require filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ to be complete and right-continuous. See app. E.1 of Pollard (2002) for details.

In this class of strategies, we define a pipetting strategy that minimizes the likelihood of posterior beliefs exceeding a particular threshold.

Definition (Pipetting strategy of the agent). A disclosure strategy Y is a *pipetting strategy* with respect to an (upper) boundary $b(x)$ if (i) conditional on lack of disclosure the posterior belief Y_t^{ND} is weakly below the boundary $b(X_t)$ and (ii) process $D = (D_t)_{t \geq 0}$ is the minimal process satisfying condition i.

A pipetting strategy is well defined, since for the set of disclosure strategies $\mathcal{D}$ satisfying condition i, process $D_t = \inf_{\hat{D} \in \mathcal{D}} \hat{D}_t$ satisfies condition i and is, by construction, minimal. Beliefs under the pipetting strategy behave as follows. If the prior Y_{0-} is above $b(X_0)$, then this strategy induces an immediate jump such that the posterior Y_0 is either 1 or $b(X_0)$. Afterward, posterior beliefs change only at the boundary $b(x)$, either reflecting off it or jumping to 1. In other words, a pipetting strategy delays the disclosure of $\theta = \theta_{\mathrm{H}}$ as much as possible without crossing the boundary $b(x)$.

We define a Markov strategy of the principal as an action plan for every state of the game (x, y).[14] Since her action is binary (stop/wait), a pure Markov strategy can be identified with a stopping set $\mathbb{T}$. Intuitively, the principal's strategy is Markov if her decision to stop depends only on the current level of X_t and the current belief Y_t.

Definition (Markov strategy of the principal). *The principal's Markov strategy* is a Borel set $T \subseteq [0, +\infty) \times [0, 1]$, such that the principal exercises the option at the first hitting time of $\mathbb{T}$,

$$\tau = \inf\{t \geq 0 : (X_t, Y_t) \in T\} .$$

We define an MPE (Markov perfect equilibrium) as a pair of Markov strategies that are mutual best responses.

Definition. An MPE is a pair of Markov strategies of the agent and principal $(Y^*, \mathbb{T}^*)$ such that the following conditions hold.

i (Principal's optimality). At every state $(x, y) \in [0, +\infty) \times [0, 1]$, the first hitting time τ^* of the set $\mathbb{T}^*$ is optimal, given the anticipated belief process Y^*:

$$\tau^* \in \operatorname*{arg\,max}_{\tau \in \mathcal{M}(Y^*)} \mathrm{E}\left[e^{-r\tau}\left\{\left[Y_\tau^* \cdot \theta_{\mathrm{H}} + (1 - Y_\tau^*) \cdot \theta_{\mathrm{L}}\right] X_\tau - I_{\mathrm{P}}\right\} \middle| X_0 = x, Y_{0-} = y\right], \quad (2)$$

where the maximum is taken over the set of all stopping times $\mathcal{M}(Y^*)$ with respect to the process (X, Y^*).

ii (Agent's optimality). At every state $(x, y) \in [0, +\infty) \times [0, 1]$, the posterior belief process Y^* is optimal, given the anticipated stopping rule of the principal τ^*:

[14] A general strategy for the principal could specify a stopping time as a function of the whole history, not just contemporaneous state (x, y).

$$Y^* \in \underset{\tilde{Y} \in \mathcal{Y}}{\arg\max} \mathrm{E}\left[e^{-r\tau^*}\left\{\left[\tilde{Y}_{\tau^*} \cdot \theta_{\mathrm{H}} + (1 - \tilde{Y}_{\tau^*}) \cdot \theta_{\mathrm{L}}\right] X_{\tau^*} - I_{\mathrm{A}}\right\} | X_0 = x, Y_{0-} = y\right]. \quad (3)$$

While any MPE $(Y^*, \mathbb{T}^*)$ is a pair of Markov strategies, the maximum in equations (2) and (3) is taken over all Markov and non-Markov strategies. Even though the set of all available stopping times of the principal $\mathcal{M}(Y^*)$ depends on the entire strategy of the agent Y^*, the action of the principal at time t, that is, whether $\tau \le t$, depends only on the history of (X, Y^*) up to time t. The word "perfect" in the definition of MPE emphasizes that the strategies of the principal and the agent are time consistent. The principal's strategy is time consistent because her best response given by equation (2) is a solution to an optimal stopping problem in a Markov environment. The agent's strategy is time consistent because it is Markov and equation (3) ensures that it is the agent's best response for any initial state (x, y).

C. *Autarky Thresholds*

For a given belief $Y_t = y$, the optimal stopping decision of the principal (or the agent, if he were given control rights) that is based only on exogenous evolution of X—that is, if no additional information about θ were available—can be characterized by a first entry time into the set $\{x \ge x_{\mathrm{P}}(y)\}$ (or $\{x \ge x_{\mathrm{A}}(y)\}$ for the agent), with

$$x_i(y) = \frac{\beta}{\beta - 1} \cdot \frac{I_i}{y \cdot \theta_{\mathrm{H}} + (1 - y) \cdot \theta_{\mathrm{L}}}, \quad i \in \{\mathrm{A}, \mathrm{P}\},$$

where $\beta > 1$ is the positive root of $[\phi^2\beta(\beta - 1)]/2 + \mu\beta = r$.[15] We refer to $x_i(y)$ as the *autarky threshold* of player i. When a player holds a higher belief y, they would like to exercise the option earlier; that is, $x_i(y)$ is decreasing in y.

III. Persuasion to Wait

In this section, we characterize the unique MPE in the case $I_{\mathrm{A}} > I_{\mathrm{P}}$, that is, when the agent prefers to exercise the option later (at a higher threshold) than the principal: $x_{\mathrm{A}}(y) > x_{\mathrm{P}}(y)$. Define $y_{\mathrm{P}}(x) = x_{\mathrm{P}}^{-1}(x)$ to be the lowest belief about θ at which the principal is willing to exercise the option under autarky, given $X_t = x$.

[15] See Dixit and Pindyck (2012) for details.

PROPOSITION 1. There exists an essentially unique MPE.[16] Equilibrium strategies are characterized by a threshold $x^* \in (x_P(1), x_P(0)]$ and two boundaries $a(x)$ and $b(x)$.

i (Agent). For $x \le x^*$, the agent follows the pipetting strategy with respect to the principal's autarky boundary $y_P(x)$. For $x \in (x^*, x_A(1))$, the agent follows a pipetting strategy with respect to $b(x)$. Finally, for $x \ge x_A(1)$, he fully discloses θ.

ii (Principal). For $x \le x^*$, the principal follows her autarky strategy; that is, she exercises the option only if $y > y_P(x)$. For $x \in (x^*, x_P(0))$, the principal exercises the option if and only if $y > a(x)$ or $y = 1$. For $x \ge x_P(0)$, she exercises the option regardless of y.

Threshold x^* is the point at which the principal is indifferent between following her autarky strategy and waiting for the agent to disclose all information at $x_A(1)$ (see eq. [4] below). Figure 2 depicts the equilibrium strategies of the principal and the agent graphically when, in equilibrium, $x^* \in [[r/(r-\mu)](I_A/\theta_H), x_P(0))$.[17] In this case, the equilibrium boundaries are given by $a(x) = 1$ and $b(x) = y_P(x^*)$. If $x^* < [r/(r-\mu)](I_A/\theta_H)$, then the equilibrium construction is qualitatively similar and is presented in the appendix. If $x^* = x_P(0)$, then it is necessarily the case that $x_A(1) \ge x_P(0)$ and the equilibrium features only pipetting.

In the equilibrium of proposition 1, the principal always knows θ when she exercises the option. Even though her stopping set includes states for which $y < 1$, as illustrated by the shaded set in figure 2*A*, the agent discloses $\theta = \theta_H$ such that the option is exercised only at $y = 1$ or $y = 0$. If the game starts in the interior of the shaded set of figure 2*B*, the agent immediately reveals information to move beliefs to either 1 or $y_P(x)$ for $x \le x^*$ and to $b(x)$ for $x \in (x^*, x_A(1))$. Along the equilibrium path, the agent probabilistically discloses $\theta = \theta_H$ to either induce immediate action or elicit additional waiting from the principal.

The equilibrium features four distinct regions $\{R_i\}_{i=1}^4$, depicted in figure 3. In $R_1 = \{x \ge x_A(1)\}$, the agent fully discloses θ and the principal exercises the option immediately if $\theta = \theta_H$ or, possibly, waits until $x_P(0)$ if $\theta = \theta_L$. This information disclosure is valuable for the principal if $x < x_P(0)$ and thus introduces an incentive for her to wait for it and postpone option exercise. As a result, in $R_2 = \{x < x_A(1), y < y_P(x^*)\}$, the principal is willing to wait past her autarky threshold until region R_1 is reached. The agent does not need to communicate in region R_2 and still obtains his first-best option exercise conditional on $\theta = \theta_H$ at $x_A(1)$.

[16] *Essential uniqueness* means uniqueness of the outcome of the game, i.e., distribution of (X_τ, τ) conditional on θ.

[17] This condition is satisfied whenever I_P is sufficiently close to I_A. See eq. (12), in the appendix, for the exact condition.

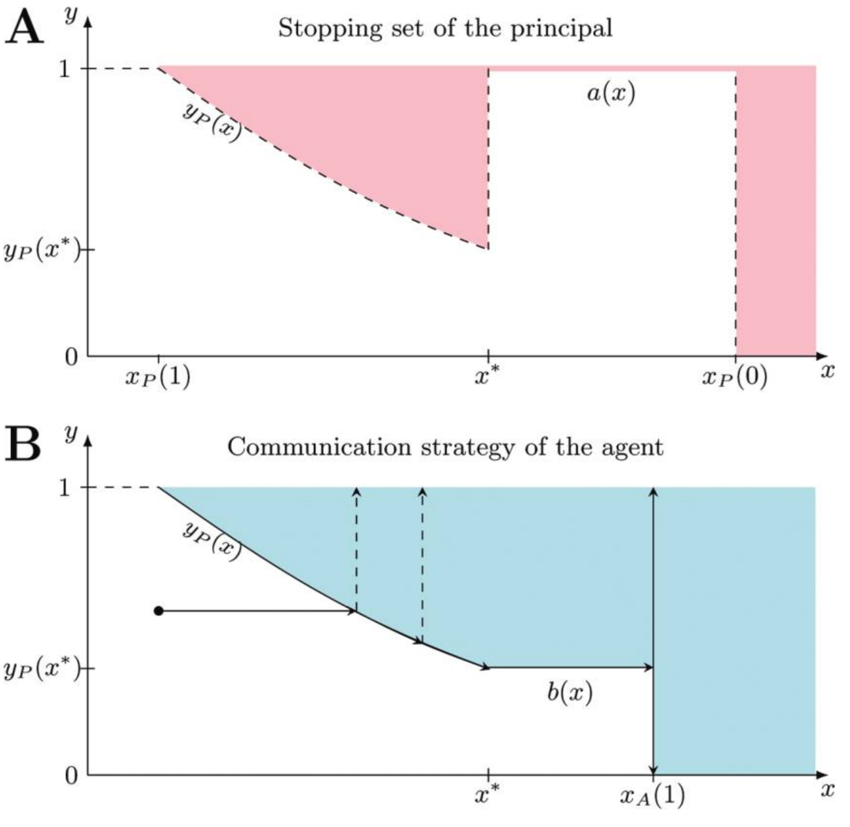

FIG. 2.—Equilibrium strategies with $[r/(r-\mu)] \cdot (I_A/\theta_H) \le x^* < x_P(0)$.

When the starting beliefs are high and the initial X is low, that is, when the game starts in $R_3 = \{x \le x^*, y > y_P(x^*)\}$, the principal (absent pipetting) prefers to exercise the option at her autarky threshold $y_P(x)$, rather than wait until region R_1 is reached. The agent's pipetting makes the principal indifferent between waiting and exercising the option (in equilibrium, she chooses the latter). When (x, y) is below the principal's autarky threshold, the state moves only horizontally, because the agent does not communicate. When x increases to $x_P(y)$, the agent begins pipetting information, which results in posterior y either jumping up to 1 or sliding down along the principal's autarky threshold $y_P(x)$, as depicted in figure 2B. Once the posterior reaches $y_P(x^*)$, the agent stops providing information until region R_1 is reached. By pipetting information in R_3, the agent minimizes the probability of early option exercise. Finally, in the remaining region, R_4, the agent sends an immediate discrete signal

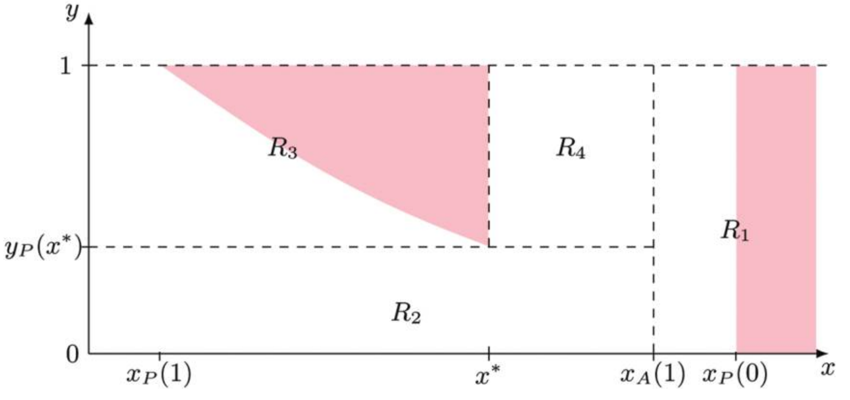

FIG. 3.—Partition of the state space into four regions, R_1–R_4.

to the principal that either reveals $\theta = \theta_H$ and results in instantaneous option exercise or lowers beliefs to $y_P(x^*)$, in which case she waits for disclosure in region R_1.

We now provide intuition and sketch a proof of why the described persuasion strategy of the agent and the stopping set of the principal are best responses to each other, region by region.

Equilibrium in region R_1.—When the principal knows θ, then she exercises the option at her autarky thresholds $x_P(1)$ and $x_P(0)$. Moreover, when $X_t \geq x_P(0)$, then the principal is past her optimal exercise thresholds for both θ_H and θ_L, and hence she exercises the option immediately regardless of beliefs. Because of this, the agent understands that it is impossible to incentivize the principal to wait beyond $x_P(0)$. By disclosing θ when $X_t \geq x_A(1)$, the agent achieves his first-best timing of exercise conditional on $\theta = \theta_H$ and simultaneously delays option exercise conditional on $\theta = \theta_L$ as much as possible, that is, until $x_P(0)$. In other words, two rounds of eliminating dominated strategies imply that it is optimal for the agent to disclose θ when X_t reaches $x_A(1)$.

Equilibrium in region R_2.—The expectation of learning θ at $x_A(1)$ increases the option value of waiting for the principal relative to her autarky strategy. If the principal exercises the option at $(x, y_P(y))$, her payoff is

$$(y_P(x) \cdot \theta_H + (1 - y_P(x)) \cdot \theta_L)x - I_P = \frac{I_P}{\beta - 1}.$$

We define x^* as the point at which the principal is indifferent between exercising the option at her autarky threshold $(x^*, y_P(x^*))$ and waiting for the agent to disclose θ at $x_A(1)$. Formally, x^* is the unique solution[18] to

[18] We show uniqueness of the solution to eq. (4) in lemma A.3 in the online appendix.

$$\frac{I_{\mathrm{P}}}{\beta - 1} = y_{\mathrm{P}}(x^*)\mathrm{E}_{x^*}[e^{-rT(x_{\mathrm{A}}(1))}(\theta_{\mathrm{H}} x_{\mathrm{A}}(1) - I_{\mathrm{P}})] \\ + (1 - y_{\mathrm{P}}(x^*))\mathrm{E}_{x^*}[e^{-rT(x_{\mathrm{P}}(0))}(\theta_{\mathrm{L}} x_{\mathrm{P}}(0) - I_{\mathrm{P}})], \tag{4}$$

where $\mathrm{E}_{x^*}[\cdot]$ is the expectation conditional on $X_0 = x^*$ and $T(x) = \inf\{t > 0 : X_t \geq x\}$ is the first time X_t crosses a given threshold x. The right-hand side of equation (4) is the expected payoff from waiting to learn θ, which results in exercise either at $x_{\mathrm{A}}(1)$, if $\theta = \theta_{\mathrm{H}}$, or at $x_{\mathrm{P}}(0)$ if $\theta = \theta_{\mathrm{L}}$. For $y < y_{\mathrm{P}}(x^*)$, the principal strictly prefers to wait for the agent's information disclosure. Anticipating this, the agent does not communicate any information for $y < y_{\mathrm{P}}(x^*)$ and waits until X_t reaches $x_{\mathrm{A}}(1)$.[19]

Equilibrium in region R_3.—A higher belief about θ makes the principal more likely to regret waiting for the agent's disclosure at $x_{\mathrm{A}}(1)$. Consequently, in region R_3 the principal would rather exercise the option in her autarky set $\{x > x_{\mathrm{P}}(y)\}$, absent pipetting. In equilibrium, the agent communicates additional information through pipetting with respect to $y_{\mathrm{P}}(x)$. Such pipetting generates posterior beliefs $Y_t \in \{y_{\mathrm{P}}(X_t), 1\}$. We claim that this information disclosure does not generate value for the principal. The intuition is as follows. Information is disclosed only when the beliefs are $y > y_{\mathrm{P}}(x)$ (i.e., in the shaded part of region R_3). If the posterior belief is 1, then it is optimal to exercise the option. If the posterior belief is $y_{\mathrm{P}}(x)$, then the principal is indifferent between exercising the option and waiting. So option exercise is one of the optimal actions in both cases. Stopping is also an optimal action under autarky (i.e., absent any persuasion), and hence the expected equilibrium payoff is the same as the autarky payoff. Since such pipetting of information provides no value to the principal, exercising the option for $y > y_{\mathrm{P}}(x)$ in region R_3 is still optimal for her.

In order to show that pipetting information at the principal's stopping boundary is optimal for the agent, we show that not disclosing $\theta = \theta_{\mathrm{H}}$ until the principal is just about to act is strictly optimal. To do so, we construct an upper bound on the equilibrium expected value of the agent, given $y_{\mathrm{P}}(x)$, and show that Y^* delivers this upper bound. Note that the agent always benefits from additional delay if $\theta = \theta_{\mathrm{L}}$, and if, along the equilibrium path, the option is exercised for some intermediate belief $Y_\tau \in (0, 1)$, the agent can strictly improve his payoff by fully disclosing θ at τ. Such disclosure does not affect option exercise if $\theta = \theta_{\mathrm{H}}$ and weakly delays it if $\theta = \theta_{\mathrm{L}}$. As a result, the agent's expected payoff can be written as

$$Y_{0-} \cdot \mathrm{E}[e^{-r\tau}(\theta X_\tau - I_{\mathrm{A}})|\theta = \theta_{\mathrm{H}}] + (1 - Y_{0-}) \\ \cdot \mathrm{E}[e^{-rT(x_{\mathrm{P}}(0))}(\theta x_{\mathrm{P}}(0) - I_{\mathrm{A}})|\theta = \theta_{\mathrm{L}}].$$

[19] The equilibrium could feature only "inconsequential" communication in this region that does not allow the posterior belief to exceed $y_{\mathrm{P}}(x^*)$ before fully disclosing θ at $x_{\mathrm{A}}(1)$.

For $x < x_P(1)$, the agent cannot persuade the principal to exercise the option, and hence any communication can be delayed until $x_P(1)$ without loss. For $x \geq x_P(1)$, the agent can incentivize the principal to exercise the option immediately by probabilistically disclosing whether $\theta = \theta_H$. We can identify the stopping time τ with the first time the agent discloses $\theta = \theta_H$ and consider only disclosure strategies.[20] Recall that, for a disclosure strategy, D_t is the cumulative, history-dependent, conditional probability of disclosing $\theta = \theta_H$ up to time t. The agent's payoff conditional on $\theta = \theta_H$ can be written as

$$\mathrm{E}[e^{-r\tau}(\theta_H X_\tau - I_A)|\theta = \theta_H] = \mathrm{E}\left[\int_0^\infty e^{-rt}(\theta_H X_t - I_A)\, dD_t\right]. \quad (5)$$

Next, we perform a change of time in the above integral to reduce the agent's problem to a point-wise optimization. Cumulative probability of disclosure D_t uniquely pins down the posterior belief Y_t^{ND} conditional on staying on the path of no disclosure via equation (1). Since this (conditional) posterior belief is monotonically decreasing over time, instead of integrating equation (5) over calendar time t, we can integrate it over levels of Y_t^{ND} when $\theta = \theta_H$ is disclosed. Define $\eta(y) = \inf\{t : Y_t^{\mathrm{ND}} \leq y\}$ to be the first time the posterior belief conditional on no disclosure falls below y. Using a change of variable $t = \eta(y)$, it follows that $D_{\eta(y)} = 1 - [(1 - Y_{0-})/Y_{0-}] \cdot [1/(1 - y)]$, and we can rewrite equation (5) as[21]

$$\mathrm{E}\left[\int_0^\infty e^{-rt}(\theta_H X_t - I_A)\, dD_t\right] = \frac{1 - Y_{0-}}{Y_{0-}} \cdot \int_0^{Y_{0-}} \left(\mathrm{E}\left[e^{-r\eta(y)}(\theta_H X_{\eta(y)} - I_A)\right] \cdot \frac{1}{(1 - y)^2}\right) dy. \quad (6)$$

Intuitively, $\eta(y)$ represents the random time at which $\theta = \theta_H$ is disclosed, given level of beliefs y. The shape of the principal's stopping set puts a constraint on what stopping times $\eta(y)$ the agent can induce. Suppose that $Y_{0-} < y_P(X_0)$. Then in order for $\eta(y)$ to be feasible, it must be the case that the path of X_t stays below $x_P(y)$ for all $t < \eta(y)$. The agent would like to exercise the option as close to his first-best threshold $x_A(1)$ as possible, that is, at the highest possible level of X_t in the principal's waiting region, which is equal to $x_P(y)$. Therefore, the optimal $\eta(y)$ is the first hitting time of $x_P(y)$,[22] and we can bound the agent's expected payoff as

$$\mathrm{E}\left[e^{-r\eta(y)}(\theta_H X_{\eta(y)} - I_A)\right] \leq \mathrm{E}\left[e^{-rT(x_P(y))}(\theta_H x_P(y) - I_A)\right] \ \forall\, y \in (y_P(x^*), Y_{0-}]. \quad (7)$$

[20] We derive this result formally in lemma 1, in the appendix. The revelation principle stated in Myerson (1986) does not apply in our setting because of a lack of agent's long-term commitment.

[21] Representation (6) holds if D_t is continuous and goes through with a minor modification if D_t has jumps.

[22] In the case $Y_{0-} > y_P(X_0)$, the optimal $\eta(y) = T(\max(X_0, x_P(y)))$.

Inequality (7) highlights that it is optimal for the agent to disclose $\theta = \theta_H$ only along the principal's boundary $y_P(x)$. This outcome is uniquely achieved by the pipetting strategy against boundary $y_P(x)$. Also note that for $y \le y_P(x^*)$, the first-best stopping time $\eta(y) = T(x_A(1))$ is feasible because the pair (X_t, y) remains below the principal's stopping set.

Equilibrium in region R_4.—For all $(X_0, Y_{0-}) \in R_4$, the agent generates a discrete signal at time 0 such that the posterior Y_0 is either 1 or $y_P(x^*)$. This policy is optimal, despite the principal being willing to wait in this region, by the following logic. Since X_t is stochastic, the agent faces a risk that X_t declines and the state (X_t, Y_{0-}) enters region R_3, where the option may be exercised by the principal. This possibility lowers the agent's option value of waiting for all $X_t > x^*$, and, as a result, he finds it optimal to speed up option exercise conditional on $\theta = \theta_H$, rather than wait for the possibility of X_t reaching his first-best threshold $x_A(1)$. Direct computation shows that, in the parametric case when $x^* \ge [r/(r-\mu)] \cdot (I_A/\theta_H)$, immediate option exercise is optimal, as the agent loses more on discounting of $\theta_H X_t$ than he gains by discounting the investment cost I_A. Formally, it implies that $\eta(y) = 0$ maximizes expression (6) for $(x, y) \in R_4$.[23]

The principal has no incentives to exercise the option for $y \in (y_P(x^*), 1)$, since she expects the agent to disclose valuable information the next instant, resulting in an expected value of

$$
\begin{aligned}
V(x, y) &= \frac{y - y_P(x^*)}{1 - y_P(x^*)} \cdot (\theta_H x - I_P) + \frac{1 - y}{1 - y_P(x^*)} \cdot V(x, y_P(x^*)) \\
&> [y\theta_H + (1 - y)\theta_L]x - I_P.
\end{aligned}
$$

Therefore, the principal strictly prefers to wait for the agent's information in state $(x, y_P(x^*))$.

Note that the agent's behavior close to x^* is very similar to the right and to the left of x^*: in both cases, the agent discloses information so that the posterior belief Y_t is either 1 or $y_P(x^*)$ (approximately for $X_t < x^*$). Yet the principal's optimal action is different. In particular, if the agent deviates and does not disclose information, the principal stops if beliefs are to the left of x^* and waits if beliefs are to the right of x^*. The intuition is that for $x < x^*$ it is optimal to stop after both belief realizations. Therefore, waiting does not create value, and the principal stops if the information does not come. In contrast, to the right of x^*, the optimal action depends on the information disclosed, and hence information has a positive value. It is thus optimal for the principal to wait if the agent deviates in region R_4. The value of information depends on the distance from x^* and converges

[23] When $x^* < [r/(r-\mu)] \cdot (I_A/\theta_H)$, the agent's option value of waiting is not completely eliminated. For these parameters, the optimal $\eta(y)$ may be positive, but the option is still exercised before the agent's first-best threshold $x_A(1)$. We characterize the equilibrium boundaries $(a(x), b(x))$ for $x \in [x^*, x_A(1)]$ in the appendix.

to zero as we approach x^* from the right. So not waiting to the left of x^* in the hope for X_t to pass from region R_3 to region R_4 is optimal.

Equilibrium uniqueness.—We show the uniqueness of the equilibrium outcome by sequentially eliminating dominated strategies of the agent and the principal. The previous logic uniquely pins down equilibrium behavior in regions R_1 and R_2. Establishing uniqueness in regions R_3 and R_4 is more involved. The agent can, potentially, induce waiting past the principal's autarky threshold by promising, for example, to reveal θ before X_t reaches $x_A(1)$. However, lack of commitment makes such promises not credible. Instead, the agent knows that the principal never stops before $x_P(y)$, and he can thus delay all communication until $x_P(y)$ is reached. Even more so, by pipetting information, the agent can improve his payoff as described in inequality (7). We show that the best payoff the principal can achieve is equal to that of the equilibrium in proposition 1. This pins down the stopping set of the principal and the boundary $a(x)$. Once this is established, the previous analysis shows that Y^* is the agent's essentially unique best response.

Equilibrium properties.—We begin by comparing payoffs obtained by the players in equilibrium with their autarky payoffs, that is, if there was no persuasion.

Proposition 2. Let $V_P(x, y)$ and $V_A(x, y)$ be the MPE payoffs of the principal and the agent, respectively. Let $V_P^{\text{Aut}}(x, y)$ and $V_A^{\text{Aut}}(x, y)$ be the payoffs of the players when the principal exercises the option at the autarky threshold $x_P(y)$ and the agent provides no information. Then for the agent,

$$V_A(x, y) > V_A^{\text{Aut}}(x, y) \text{ if } x < x_P(0) \text{ and } y \notin \{0, 1\},$$

and $V_A(x, y) = V_A^{\text{Aut}}(x, y)$ otherwise. And for the principal,

$$\begin{aligned} V_P(x, y) = V_P^{\text{Aut}}(x, y) \text{ if } &(x, y) \in R_3, \text{or } (x, y) \in R_4 \text{ and } y \geq a(x), \\ &\text{or } x \geq x_P(0), \text{ or } y \in \{0, 1\}, \end{aligned}$$

and $V_P(x, y) > V_P^{\text{Aut}}(x, y)$ otherwise.

This proposition shows that there is an asymmetry in the split of gains from information about θ in equilibrium, relative to players' autarky payoffs. The agent always benefits from information control whenever his information can have any effect on the principal's actions. There are three sources for these gains. First, the agent benefits from immediate persuasion to minimize the probability of entering the principal's stopping set. Second, the agent benefits from better decision-making when he reveals information at $x_A(1)$. Third, since information disclosure is credible at x_A (1), for some beliefs, the principal finds it optimal to wait past her autarky threshold, and that benefits the agent even further.

However, not all persuasion is valuable to the principal. Specifically, pipetting of information along the principal's stopping boundary $y_P(x)$ has no value for her, since she is locally indifferent between immediate exercise and waiting. When the game starts with a prior above $y_P(x)$, then the time 0 persuasion leaves the posterior on the linear part of principal's value function and, hence, also does not benefit her. Anticipation of future pipetting at $y_P(x)$ does not benefit the principal because of the optimality of $y_P(x)$, absent pipetting.[24] Therefore, for all priors in region R_3, the principal's payoff is equal to her autarky payoff. For priors in region R_2, it is strictly higher (the other two regions vary).

Proposition 1 states that when the belief about θ is sufficiently low—that is, $Y_{0-} \le y_P(x^*)$—the option conditional on $\theta = \theta_H$ is exercised at the agent's first-best threshold $x_A(1)$. In this case, we say that the conflict between the principal and the agent is small. Next, we show that the size of this conflict depends not only on the difference in preferences but also on the value of information about θ and on the volatility of public information.

Proposition 3. The range of beliefs $[0, y_P(x^*)]$ for which the principal is willing to wait for the agent to fully disclose θ at $x_A(1)$ is

i. larger when preferences are more aligned, that is, $\partial y_P(x^*)/\partial I_P > 0 > \partial y_P(x^*)/\partial I_A$;
ii. larger when information about θ is more valuable, that is, $\partial y_P(x^*)/\partial \theta_H > 0 > \partial y_P(x^*)/\partial \theta_L$;
iii. larger when X is more volatile and $\theta_L = 0$, that is, $(\partial y_P(x^*)/\partial \phi)|_{\theta_L = 0} > 0$.

First, when preferences of players are more aligned—that is, when $I_P - I_A$ is low—the principal's cost of waiting until $x_A(1)$ is small relative to the benefit of learning θ. As a result, the stopping region $y > y_P(x)$ shrinks and $y_P(x^*)$ goes up. Second, an increase in θ_H reduces the information revelation threshold $x_A(1)$ since, conditional on $\theta = \theta_H$, the agent's benefit of exercising the option goes up. A lower $x_A(1)$, in turn, corresponds to a smaller principal's cost of waiting for the discrete information revelation. A decrease in θ_L does not affect the incentives of the agent to fully disclose θ at $x_A(1)$. However, it makes the information disclosed at $x_A(1)$ more valuable, since the principal avoids premature option exercise conditional on $\theta = \theta_L$. As a result, waiting is more attractive for the principal when $\theta_H - \theta_L$ is high. Finally, note that the ratio $x_P(y)/x_A(1)$ does not depend on the volatility of X. Hence, for a fixed belief y, the process X has to cover the same distance from the autarky threshold of the principal to the point when the agent discloses θ. A more volatile process X moves faster, thus reducing the expected cost of waiting for information and increasing $y_P(x^*)$.

[24] This result is the dynamic extension of the property that the principal (receiver) does not gain from equilibrium persuasion in the binary judge-prosecutor example of Kamenica and Gentzkow (2011) and the two-period model discussed in sec. I.

A. *Value of Dynamic Commitment*

The equilibrium constructed in proposition 1 features pipetting of information along the boundaries $y_{\mathrm{P}}(x)$ and $a(x)$ in order to delay the option exercise. In this section, we show that this is a feature of limited commitment in the dynamic persuasion game, rather than the optimal way to induce a delayed action from the principal. The key distinction between the dynamic game and the dynamic mechanism in our setting is the ability of the agent to commit to delayed persuasion.

For an arbitrary (non-Markov) admissible belief process $Y \in \mathcal{Y}$, let $\mathcal{M}^*(Y)$ denote the set of principal's best responses to Y, that is,

$$\mathcal{M}^*(Y) = \underset{\tau \in \mathcal{M}(Y)}{\arg\max}\ \mathrm{E}[e^{-r\tau}\{[Y_\tau \cdot \theta_{\mathrm{H}} + (1 - Y_\tau) \cdot \theta_{\mathrm{L}}]X_\tau - I_{\mathrm{P}}\}].$$

A *feasible dynamic persuasion mechanism* is a strategy of the agent Y for which the set $\mathcal{M}^*(Y)$ is nonempty. Many mechanisms satisfy this requirement; for example, any Markov posterior belief process Y leads to a nonempty $\mathcal{M}^*(Y)$. Denote by $\mathcal{Y}^{\mathrm{M}}$ the set of feasible mechanisms.

A feasible dynamic persuasion mechanism is optimal if it maximizes the agent's ex ante payoff while taking into account the best response of the principal.

Definition. Optimal dynamic persuasion mechanism Y^{M} is a feasible dynamic mechanism, which maximizes the agent's ex ante payoff:

$$Y^{\mathrm{M}} \in \underset{Y \in \mathcal{Y}^{\mathrm{M}}}{\arg\max}\left[\sup_{\tau \in \mathcal{M}^*(Y)} \mathrm{E}[e^{-r\tau}\{[Y_\tau \cdot \theta_{\mathrm{H}} + (1 - Y_\tau) \cdot \theta_{\mathrm{L}}]X_\tau - I_{\mathrm{A}}\}]\right]. \quad (8)$$

There is an important difference between the optimal persuasion mechanism and the MPE formulated in section II.B. In the mechanism, the agent commits to the full dynamic strategy at time 0, influencing the best response of the principal. In contrast, in the equilibrium of the dynamic game, the strategies are mutual best responses at every time t or, equivalently, in any state (x, y). Proposition 4 characterizes the optimal dynamic persuasion mechanism and illustrates that it requires making ex post suboptimal promises.

Proposition 4. Suppose that $X_0 \le x_{\mathrm{P}}(Y_{0-})$. Then the optimal dynamic persuasion mechanism is characterized by a disclosure threshold $\bar{x}(Y_{0-})$. The agent remains quiet, and the principal waits for information when $X_t < \bar{x}(Y_{0-})$. When X_t reaches $\bar{x}(Y_{0-})$, the agent fully reveals θ. Disclosure threshold $\bar{x}(Y_{0-})$ is strictly decreasing for $Y_{0-} > y_{\mathrm{P}}(x^*)$ and coincides with $x_{\mathrm{A}}(1)$ for $Y_{0-} \le y_{\mathrm{P}}(x^*)$. There is no pipetting of information under long-term commitment.

When $Y_{0-} \le y_{\mathrm{P}}(x^*)$ the agent obtains first-best payoff for $\theta = \theta_{\mathrm{H}}$ in the equilibrium of proposition 1, and commitment is not valuable. For

$Y_{0-} > y_P(x^*)$, we first evaluate a constrained Pareto efficient exercise policy, in which, conditional on $\theta = \theta_L$, the option is exercised at $x_P(0)$, and conditional on $\theta = \theta_H$ the option is exercised at a threshold $\bar{x}$. Next, we choose $\bar{x} = \bar{x}(Y_{0-})$ to make the principal indifferent at $t = 0$ between such exercise policy and her expected value under autarky.[25] The payoff from this exercise policy is the upper bound on the agent's expected value in any long-term mechanism. The agent can implement this outcome by committing to acquire and disclose θ at $\bar{x}(Y_{0-})$, since, upon observing that $\theta = \theta_H$, the principal immediately exercises the option, while, upon observing $\theta = \theta_L$, she waits until $x_P(0)$.

Proposition 4 implies that slow pipetting of information occurring at the boundary $y_P(x)$ in figure 2 arises from the inability of the agent to commit to a delayed information-sharing rule and that lack of such commitment is costly for the agent.

B. *Product Recalls and Abandonment Options*

In section I, we discussed the relevance of our results to product recalls. In particular, our findings apply to recalls of approved drugs from the market by the FDA.[26] In this section, we map the equilibrium constructed above to this environment and, more generally, to abandonment options.

State θ is binary and is either 0 or 1. In our notation, $\theta = \theta_H$ corresponds to a harmful (high-risk) drug that should be recalled. The exogenous news process X reflects public information about the drug and its interaction effects with other medications. The combined risk of the drug is $\theta \cdot X$.[27]

As long as the drug is on the market, it generates the firm (agent) an expected profit flow that depends on the state: $F_\theta \in \{F_0, F_1\}$, with $F_0 \geq F_1$ and $F_0 > 0$. It also generates expected welfare flow for the FDA (principal), $W_\theta \in \{W_0, W_1\}$, with $W_0 > 0 > W_1$. If the product is recalled from the market at time τ, then the expected payoffs of the agent and the principal conditional on θ are, respectively,

$$\begin{aligned} v_A &= \int_0^\tau e^{-rt}(F_0 - F_1 \cdot \theta X_t)\, dt, \text{ and} \\ v_P &= \int_0^\tau e^{-rt}(W_0 - W_1 \cdot \theta X_t)\, dt. \end{aligned} \tag{9}$$

[25] This is why $\bar{x}(Y_{0-})$ depends on Y_{0-} and why we focus on $Y_{0-} > y_P(x^*)$, as it implies that $\bar{x}(Y_{0-}) < x_A(1)$.

[26] Henry and Ottaviani (2019) analyze firms' incentives to quit experimentation when the initial approval threshold is set by the FDA in equilibrium.

[27] In sec. A.7 of the online appendix, we discuss the case when X_t is a public news process about θ and refer to the option as a Wald option. The analysis is qualitatively unchanged.

If $F_1 < 0$, then under complete information about θ and a high level of X, the two players agree on a recall policy (since their flow payoffs conditional on θ have the same sign). The conflict is caused by the principal's lower tolerance for the risky drug, captured by $W_0/W_1 < F_0/F_1$.

The payoff structure (eq. [9]) is equivalent to the one described in section II; that is, it gives rise to the same equilibria once I_A and I_P are properly defined. In particular, the expected payoff of the agent can be written as

$$\mathrm{E}[v_A] = \text{const.} + \frac{F_0}{r-\mu} \cdot \mathrm{E}\left[e^{-r\tau}\left(\theta X_\tau - \frac{r-\mu}{r} \cdot \frac{F_1}{F_0}\right)\right],$$

which corresponds to $I_A = [(r-\mu)/r] \cdot (F_1/F_0)$ in the terminology of section II. Similarly, the principal's expected payoff is a linear function of $\mathrm{E}[e^{-r\tau}(\theta X_\tau - I_P)]$, with $I_P = [(r-\mu)/r] \cdot (W_1/W_0)$. Hence, the analysis in section III for $I_P < I_A$ covers the case in which the FDA has a lower tolerance for the risky drug than the drug producer.

We highlight two features of the equilibrium in the context of the postmarket surveillance program conducted by the FDA. If the prior beliefs and parameters are such that $Y_{0-} \leq y_P(x^*)$, then the timing of the action has a compromise property: in case the results of the test run by the agent at $x_A(1)$ reveal problems with the drug, the principal acts at the agent's optimal point; when the results are encouraging, the principal further delays action until her optimal threshold $x_P(0)$. As a result, anytime $\theta = \theta_H$ is revealed, the firm recalls the drug voluntarily. That is not true in the pipetting region $y > y_P(x)$, where, upon observing that $\theta = \theta_H$, the FDA recalls the drug against the wishes of the firm. This equilibrium behavior corresponds to the FDA requiring additional tests for drugs that it considers dangerous (involuntary recalls), while recalls of seemingly safe drugs are delegated to the manufacturer (voluntary recalls).

IV. Persuasion to Act

We now turn to the case when it is the agent who would like to exercise the real option sooner than the principal, which is captured by the parametric case $I_A < I_P$. This can correspond to a situation where the agent works for the principal and they learn jointly from public news about a potential project the agent would like to start. The agent is biased toward early option exercise either because the project gives him private benefits or because he does not fully internalize the fixed cost of starting it. The results in this section show that, in equilibrium, the agent can be worse off than under autarky, that is, if he never communicated with the principal.[28]

[28] The agent's autarky payoff is his expected value if the option is exercised at the principal's threshold and no information about θ is disclosed.

When the agent is able to provide credible information and cannot commit to staying quiet, the principal sets the exercise threshold higher in anticipation of this information. This hurts the agent.

PROPOSITION 5. There exists an MPE in which the agent fully discloses θ at time 0 for all (X_0, Y_{0-}).

In such an equilibrium, the principal achieves first-best payoff by adopting a very "strong bargaining position": she threatens not to exercise the option until either θ is fully known or the level of X_t is so high that θ is irrelevant for the option exercise decision, that is, $X_t \geq x_P(0)$. Given the equilibrium strategy of the agent, such a threat is credible, since waiting for the information to be released in the next instant is costless in continuous time. The agent faces a tough choice: either to fully disclose θ immediately or to wait until $X_t = x_P(0)$; any partial information acquisition does not affect the timing of option exercise. If θ is known, the option is exercised either at $x_P(1)$ (if $\theta = \theta_H$) or at $x_P(0)$ (if $\theta = \theta_L$). Since the agent prefers the option to be exercised earlier and $x_P(1) < x_P(0)$, he is better off fully disclosing θ at time 0.

PROPOSITION 6. If $X_0 \leq x_P(1)$, then the equilibrium outcome of proposition 5 is essentially unique. Moreover, if $\phi > 0$, in any MPE $(Y^*, \mathbb{T}^*)$ the equilibrium stopping set $\mathbb{T}^*$ is strictly smaller than the autarky stopping set; that is, there exists a boundary $\underline{a}(x)$ satisfying $\underline{a}(x) > y_P(x)$ for $x \in (x_P(1), x_P(0))$ such that $T^* \subseteq \{(x, y) : y \geq \underline{a}(x)\}$.

We prove the first part of proposition 6 by constructing an upper bound on the agent's payoff. Suppose that in equilibrium, the option is exercised at (X_τ, Y_τ). The principal always waits at least until her autarky threshold, and hence $X_\tau \geq x_P(Y_\tau)$. Since the agent prefers earlier option exercise, he would rather the option be exercised exactly at principal's autarky threshold, that is, at $(x_P(Y_\tau), Y_\tau)$. This would deliver the agent an expected payoff $\mathrm{E}[e^{-rT(x_P(Y_\tau))}] \cdot \{[\beta/(\beta - 1)]I_P - I_A\}$, which is weakly higher than his equilibrium payoff, yet it may not be feasible if the principal's equilibrium stopping set is smaller than her autarky stopping set. As a result of discounting, $\mathrm{E}[e^{-rT(x_P(y))}]$ is a convex function in Y. By disclosing θ immediately, the agent can achieve exercise at either $x_P(0)$ or $x_P(1)$ in any equilibrium. Thus, immediate disclosure increases the agent's payoff both because it generates option exercise at the principal's autarky threshold and also because of the convexity stemming from discounting. Figure 4*B* illustrates the information the agent discloses in any equilibrium.

The second part of proposition 6 states that in any equilibrium the principal delays her action relative to autarky. The intuition is that even if $X_0 > x_P(1)$, there is a positive probability that X_t will decline below $x_P(1)$. If that happens, the previous argument shows that the agent fully discloses θ. As a result, the principal is willing to wait beyond her autarky threshold $x_P(y)$. We define $\underline{a}(x)$ as an indifference point: at $(x, \underline{a}(x))$, the benefit of waiting for information at $x_P(1)$ is exactly offset by the associated delay

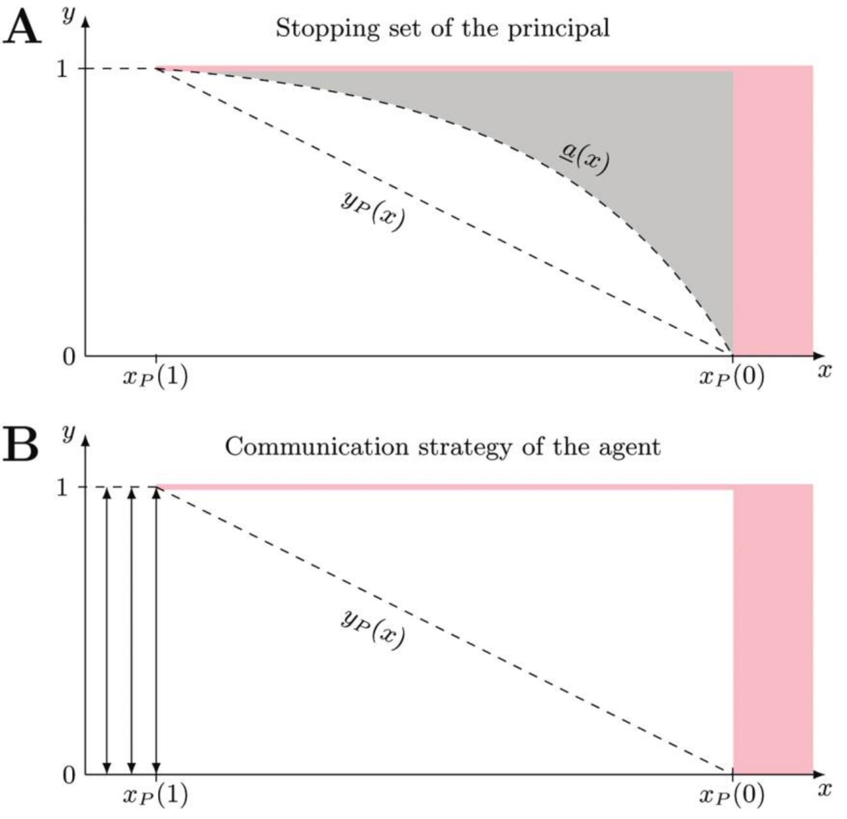

FIG. 4.—Minimal equilibrium communication and implied waiting region.

cost. Two rounds of eliminating dominated strategies imply that in any equilibrium the principal waits in the region $y_P(x) < y < \underline{a}(x)$, as shown in figure 4*A*. Since additional communication can only increase principal's value of waiting for more information, the equilibrium stopping set $\mathbb{T}^*$ is weakly above $\underline{a}(x)$.

COROLLARY 1. Suppose that $Y_{0-} \in (y_P(X_0), \underline{a}(X_0))$, $x_A(0) < X_0 \le x_P(0)$, and $\phi > 0$. Then the agent's expected payoff in any MPE is strictly lower than his value if the principal acted at her autarky threshold (i.e., if no information were available).

Absent persuasion, the principal would have exercised the real option immediately. However, in equilibrium she expects the agent to disclose θ when $X_t \le x_P(1)$ and no longer exercises the option at $t = 0$. By providing additional information about θ, the agent accelerates the option exercise conditional on $\theta = \theta_H$ but delays it conditional on $\theta = \theta_L$. The parametric condition of $x_A(0) < X_0$ implies that the agent would have preferred immediate exercise for both $\theta = \theta_H$ and $\theta = \theta_L$. In this case, delay past $x_P(Y_{0-})$ hurts the agent regardless of the equilibrium communication strategy Y^*. In fact, he would benefit from being able to commit not to disclose θ at all.

Next, we characterize the optimal dynamic persuasion mechanism in the special case of $\theta_L = 0$. We show that it is precisely the inability to commit to remaining quiet that differentiates the equilibrium of the dynamic game from the optimal dynamic mechanism.

PROPOSITION 7. Suppose that $\theta_L = 0$ and $I_A < 0$. The optimal dynamic persuasion mechanism sends a single message at time $t = 0$ that induces a posterior belief of either 0 or $y_P(X_0)$ whenever $X_0 \in [x_P(1), x_P(0)]$. When $X_0 < x_P(1)$, the agent fully reveals θ immediately. There is no additional communication for any $t > 0$.[29]

The optimal long-term mechanism involves pooling information to incentivize the principal to act. The solution resembles the trial example in Kamenica and Gentzkow (2011), in which either $\theta = \theta_L$ is perfectly revealed or posterior beliefs increase to the principal's action threshold.

Such an outcome cannot be sustained in an equilibrium of a dynamic game, as it would unravel. Suppose that the agent was to send a pooling message as in proposition 7; however, unlike in the persuasion mechanism, he cannot commit to staying quiet after such message. Because X_t might fall below $x_P(Y_t)$ "in the next instant," the principal expects to learn additional information from the agent in the future, which renders her autarky stopping rule $x_P(Y_t)$ suboptimal. We can see this by noting that the agent's message concavifies both his and the principal's value function in the waiting region $y < y_P(x)$. This creates a positive kink in the principal's equilibrium value function V_P at the action threshold $y_P(y)$. In the presence of Brownian increments dX_t, such a kink makes it optimal for the principal to wait beyond her autarky threshold and pushes the exercise threshold higher than $y_P(x)$. The inability of the agent to commit to being quiet after the first message renders the outcome of the dynamic persuasion mechanism unattainable in any MPE of the dynamic game.

V. Uniform θ

In this section, we show that the equilibria described by propositions 1 and 5 do not depend on the assumption about the binary nature of the state θ.

Suppose that $\theta \sim U[\underline{\theta}, \overline{\theta}]$, with $0 \leq \underline{\theta} < \overline{\theta}$ and $I_A, I_P > 0$. Define $x_A(y) = [\beta/(\beta - 1)] \cdot (I_A/y)$ to be the optimal threshold at which the agent would exercise the real option if he knew that $\theta = y$. Since $x_A(\theta)$

[29] The case of $\theta_L = 0$ is special, as the value of the option in the low state is fixed at 0. If this were not the case, then there are dynamic considerations associated with the optimality of immediate information pooling.

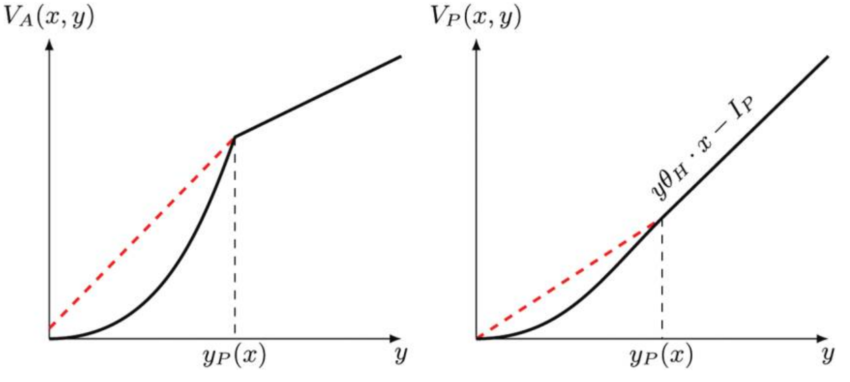

FIG. 5.—Effect of concavification on the agent's and the principal's value functions.

is monotone in θ, a natural equilibrium outcome is one in which the agent discloses a monotone sequence of θs over time. To capture this intuition in a tractable way, we restrict the agent to conduct only monotone tests. Specifically, we assume that there exist a weakly increasing process $a = (a_t)_{t\geq 0}$ and a weakly decreasing process $A = (A_t)_{t\geq 0}$ such that the posterior distribution of θ at time t is $U[a_t, A_t]$.

Given such strategies of the agent, the history of the game can be summarized by the posterior support of θ. The principal's Markov strategy can be described as a stopping set $\mathbb{T}$ in the state space (X, a, A), and the autarky threshold of the principal when the posterior distribution θ is $U[a, A]$ is given by

$$x_{\mathrm{P}}(a, A) = \frac{\beta}{\beta - 1}\cdot\frac{I_{\mathrm{P}}}{E[\theta|\theta \in [a, A]]} = \frac{\beta}{\beta - 1}\cdot\frac{2I_{\mathrm{P}}}{a + A}.$$

The notion of an MPE from section II can be easily adjusted for the newly defined strategies, and we will not repeat it here. An equilibrium consists of the principal's and the agent's strategies, given any (a, A). As we formally show in section A.3 of the online appendix, the agent discloses only the highest remaining values of θ, implying that $a_t = \underline{\theta}$ for $t < \tau$.

A. *Persuasion to Wait*

Suppose that $I_{\mathrm{A}} > I_{\mathrm{P}}$, that is, the agent prefers to exercise the option later than the principal. Figure 6*A* illustrates the principal's stopping strategy along the equilibrium path, and figure 6*B* describes the agent's persuasion strategy along the equilibrium path for the case $[\beta/(\beta - 1)]I_{\mathrm{P}} \geq [r/(r - \mu)]I_{\mathrm{A}}$. The agent discloses $\theta > \theta^*$ only if $X_t > x_{\mathrm{P}}(\underline{\theta}, \theta)$. On the other

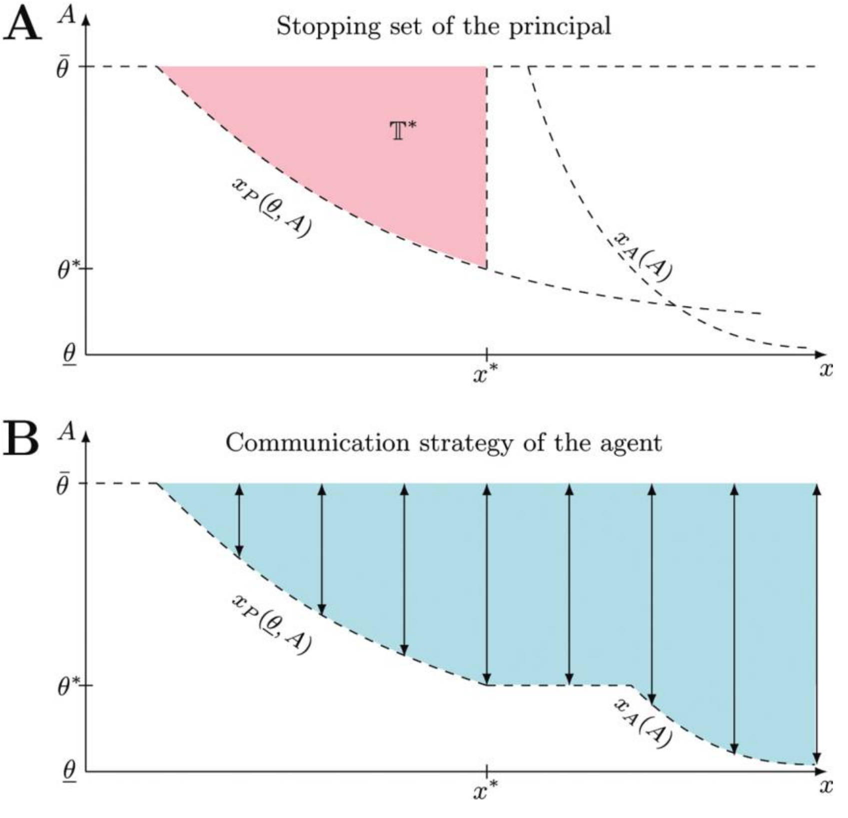

FIG. 6.—Equilibrium strategies.

hand, if $\theta < \theta^*$, then the agent discloses it at his first-best level $x_A(\theta)$. The principal is willing to wait until that threshold only for $\theta \leq \theta^*$, while for $\theta > \theta^*$ she acts at her autarky threshold. As the highest remaining θ, that is, A_t, decreases as a result of past disclosures, the distance between exercising the option at the principal's autarky threshold and the agent's first-best threshold declines. Eventually, the principal delegates option exercise to the agent, as it minimizes the risk of premature option exercise.

B. *Persuasion to Act*

If $I_A < I_P$, then full and immediate information revelation, as described in proposition 5, is still an MPE of this game. It can also be shown that in every other MPE, the principal acts past her autarky threshold, as waiting is more valuable because of the agent's information disclosure.

VI. Persuasion by an Informed Agent

So far, we have assumed that the agent has no private information and learns about θ together with the principal by publicly acquiring and disclosing information. In this section, we relax this assumption to accommodate a broader range of applications in which the agent may be privately informed. For example, a product manager can have superior information about the customer's WTP before he conducts a credible test. Such information may affect the type of test he is willing to run. Hence, the principal can potentially learn about θ not only from the outcome of the test but also from the structure of the test itself. In what follows, we informally argue[30] that equilibrium outcomes of propositions 1 and 5 can be supported as pure-strategy equilibrium outcomes of a dynamic persuasion model in which the agent knows θ but the principal does not.[31]

When the agent is privately informed about θ, he can convey information to the principal either through conducting informative tests (Bayesian persuasion) or through his choice of such tests (signaling). Note that in pure-strategy equilibria the signaling component is not needed on the equilibrium path. Instead of choosing a specific test that reveals his type, the agent might as well reveal θ by running a fully informative test. Hence, without loss, the principal's beliefs are fully determined by the hard information generated by the agent on the equilibrium path. We define the principal's beliefs to be *passive* if such property holds off path as well. Passive beliefs, among many others, support the equilibrium outcome of proposition 1.

Corollary 2 (Persuasion to wait). Suppose that $I_{\mathrm{P}} < I_{\mathrm{A}}$, and let $(Y^*, \mathbb{T}^*)$ denote the equilibrium strategies from proposition 1. Then both types of agents choosing cumulative disclosure probability process D^* that induces principal's beliefs Y^* and the principal holding passive beliefs and choosing to exercise the option in $\mathbb{T}^*$ constitute an MPE.

Corollary 2 describes a pooling equilibrium in which the agent picks the cumulative disclosure probability process D^* from proposition 1 regardless of his knowledge of θ. The principal observes the information generated by the agent and updates her beliefs, which follow a martingale Y^*. However, from the agent's perspective, principal's beliefs have a drift: the low type induces a decreasing process of beliefs Y_t^{ND}, while the high type expects beliefs to jump to 1 with a positive intensity at the boundary $y_{\mathrm{P}}(x)$.

To see why this is an equilibrium, note that by following Y^*, the low type guarantees that the posterior belief of the principal is always given by Y_t^{ND}, and hence the option is exercised at $x_{\mathrm{P}}(0)$. Because $x_{\mathrm{P}}(0)$ is the highest

[30] Formal treatment is relegated to sec. A.4 of the online appendix in the interest of space.

[31] An alternative way to model strategic communication with a privately informed agent is via disclosure of hard evidence, similar to Hörner and Skrzypacz (2016). Equilibrium outcomes of propositions 1 and 5 can be supported in such framework as well.

exercise threshold in any equilibrium, the low type has no incentives to deviate. The high type also has no incentives to deviate from Y^* for $x > x_A(1)$, since Y^* calls for a fully informative test and results in immediate option exercise, the first-best outcome for the high-type agent. Recall that for $x < x_A(1)$, belief process Y^* maximizes the expectation of $e^{-r\tau}(\theta_H X_\tau - I_A)$ conditional on $\theta = \theta_H$, given the principal's exercise region $\mathbb{T}^*$ as shown in (5)–(7). Thus, the high type has no incentives to deviate either.

Next, we show that the fully informative equilibrium outcome in proposition 5 can also be sustained when the agent is privately informed.

Corollary 3 (Persuasion to act). Suppose that $I_P > I_A$, and let $(Y^*, \mathbb{T}^*)$ denote the equilibrium strategies from proposition 5. Then both types of agents choosing belief process Y^* and the principal holding passive beliefs and choosing to exercise the option in $\mathbb{T}^*$ constitute an MPE.

By revealing himself immediately, the high type speeds up option exercise as much as possible, given the strategy of the principal. Given that the $\theta = \theta_H$ agent chooses strategy Y^* and principal's passive beliefs, the low type is indifferent between all information-sharing structures since, in equilibrium, for any belief $y < 1$, the option is exercised at $x_P(0)$ conditional on $\theta = \theta_L$. Thus, neither type has a strict incentive to deviate. In fact, this equilibrium is unique, given passive beliefs, since the high type strictly prefers to be separated from the low type in order to speed up option exercise.

VII. Conclusion

We present a theory of dynamic persuasion in the context of real options. The principal has full authority over the exercise of a real option, while the agent can disclose information to influence her decision. We show that the agent's ability to persuade always benefits the agent if he is biased toward late exercise but may hurt him if he is biased toward early exercise. We also highlight the value of dynamic commitment by comparing the equilibrium of the dynamic persuasion game to the optimal dynamic persuasion mechanism.

When the agent is biased toward late exercise, the outcomes of the equilibrium and the dynamic persuasion mechanism coincide if the conflict of interest is small. If the conflict of interest is large, the lack of commitment is costly to the agent and leads to pipetting of information. On the other hand, when the agent is biased toward early exercise, the optimal mechanism prescribes pooling information at the principal's autarky threshold. Absent dynamic commitment, however, pooling of information unravels as a result of the option of the principal to wait and obtain more information from the agent in the immediate future. The agent's inability to stop persuading undermines him, and there always exists an equilibrium in which the agent discloses all information at time 0.

Our paper is a step toward understanding the potential for manipulating information even when disclosure requirements are heavily enforced.

We show that agents can maintain strategic ignorance and delay information acquisition to their advantage. Our results suggest that despite compulsory disclosure, principals do not always gain from information disclosed in equilibrium when the agent's bias toward late exercise is large. Certain observable features of equilibrium disclosure, such as pipetting, that is, many inconclusive tests, can be used as proxies for identifying potential problems.

Appendix

We provide here additional details for proposition 1. First, we verify that the strategies in proposition 1 are mutual best responses (when $x^* \geq \hat{x}$). Second, we discuss how the model parameters determine whether $x^* \geq \hat{x} \stackrel{\text{def}}{=} [r/(r-\mu)] \cdot (I_A/\theta_H)$, corresponding to the case depicted in figure 2. Finally, we describe the equilibrium boundaries $a(x)$, $b(x)$ when $x^* < \hat{x}$ (see fig. 7). The formal proofs are in the online appendix.

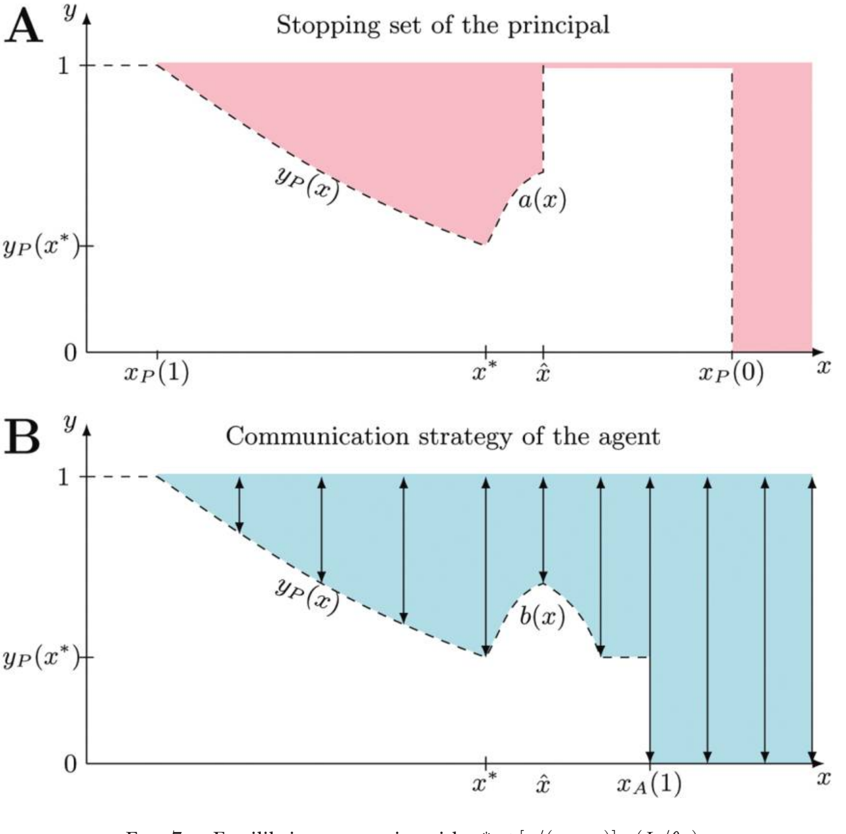

FIG. 7.—Equilibrium strategies with $x^* < [r/(r-\mu)] \cdot (I_A/\theta_H)$.

A1. *Proof of Proposition 1*

We show here that the strategies in proposition 1 are mutual best responses for the case of $\hat{x} \le x^* < x_P(0)$. This also covers the case $x^* = x_P(0)$, since then the equilibrium features only region R_3. Equilibrium verification when $x^* < \hat{x}$ and proof of uniqueness are relegated to the online appendix.

A1.1. Principal's Best Response

We show that pipetting against $y_P(x)$ does not add value for the principal, and so it is optimal for her to exercise the option for $X_t > x_P(Y_t)$ in region R_3. Conjecture that $V_P(x, y) = [y\theta_H + (1 - y)\theta_L]x - I_P$ for $y > y_P(x)$. It satisfies the boundary condition

$$\frac{\partial}{\partial y} V_P(x, y_P(x)) = \frac{V_P(x, 1) - V_P(x, y_P(x))}{1 - y_P(x)},$$

implied by the reflective nature of the belief process Y at $y_P(x)$ (see Harrison 2013).

Given this conjecture,

$$\begin{aligned}\lim_{\varepsilon \downarrow 0} \frac{\partial}{\partial y} V_P(x, y + \varepsilon)|_{y=y_P(x)} &= \frac{V_P(x, 1) - V_P(x, y_P(x))}{1 - y_P(x)} \\ &= (\theta_H - \theta_L)x.\end{aligned}$$

For $y < y_P(x)$ function $V_P(x, y)$ is given by the autarky solution; that is, it solves the waiting ordinary differential equation (ODE; eq. [13]) with appropriate value-matching and smooth-pasting conditions. Differentiating the value-matching condition at $y_P(x)$, we get

$$\lim_{\varepsilon \uparrow 0} \frac{\partial}{\partial y} V_P(x, y + \varepsilon)|_{y=y_P(x)} = (\theta_H - \theta_L)x.$$

Hence, the autarky value function of the principal satisfies the waiting ODE (eq. [13]), together with the boundary condition capturing the pipetting belief process at the boundary $y_P(x)$. Thus, the equilibrium value function of the principal is her autarky value, and the agent's pipetting against $y_P(x)$ alone does not add value to the principal.

A1.2. Agent's Best Response

To validate the approach discussed in the main text, we first verify that it is without loss to consider only disclosure strategies in the construction of the upper bound of the agent's equilibrium payoff.

LEMMA 1. If the principal's stopping set is monotone in beliefs—that is, if $(x, y) \in \mathbb{T}$, then $(x, y') \in \mathbb{T}$ for all $y' > y$—then for any best response of the agent Y there exists a disclosure strategy, $\tilde{Y}$, that is also a best response, and the two strategies induce the same outcomes almost surely.

Proof. Define D_t to be the cumulative probability of stopping (before reaching $x_A(1)$) under the belief process Y conditional on the path of X, that is, $D_t = P(Y_t = 1 | \theta = \theta_H, (X_s)_{s \le t})$ and let the belief process $\tilde{Y}_t$ to be a martingale with values in $\{1, Y_t^{ND}\}$, with Y_t^{ND} defined in equation (1).

Both belief processes induce the same distribution of stopping times (conditional on the path of X), since

$$\mathrm{P}(Y_t = 1|\theta = \theta_{\mathrm{H}}, (X_s)_{s\le t}) = \mathrm{P}(\tilde{Y}_t = 1|\theta = \theta_{\mathrm{H}}, (X_s)_{s\le t}),$$

and $Y_t^{\mathrm{ND}} = E[Y_t|Y_t < 1, (X_s)_{s\le t}]$. Hence, whenever $Y_t \le a(X_t)$, we have $Y_t^{\mathrm{ND}} \le a(X_t)$, and the strategy $\tilde{Y}$ generates the same payoffs and distribution of outcomes as Y. QED

The remainder of the proof relies on constructing the upper bound in inequality (7). In region R_3, we have $\eta(y) \le T(x_{\mathrm{P}}(y))$ for $y > y_{\mathrm{P}}(x^*)$, and thus

$$\mathrm{E}_x\left[e^{-r\eta(y)}(\theta_{\mathrm{H}}X_{\eta(y)} - I_{\mathrm{A}})\right] \le \sup_\tau \mathrm{E}_x\left[e^{-r\cdot\tau\wedge T(x_{\mathrm{P}}(y))}(\theta_{\mathrm{H}}X_{\tau\wedge T(x_{\mathrm{P}}(y))} - I_{\mathrm{A}})\right],$$

where the supremum is taken over all stopping times τ. Define $f(x, y) \stackrel{\mathrm{def}}{=} (x/x_{\mathrm{P}}(y))^{\beta_1}(\theta_{\mathrm{H}}x_{\mathrm{P}}(y) - I_{\mathrm{A}})$, and note that $f(x, y) > \theta_{\mathrm{H}}x - I_{\mathrm{A}}$ for all $x < x_{\mathrm{P}}(y)$. Then

$$\mathrm{E}_x\left[e^{-r\cdot\tau\wedge T(x_{\mathrm{P}}(y))}(\theta_{\mathrm{H}}X_{\tau\wedge T(x_{\mathrm{P}}(y))} - I_{\mathrm{A}})\right] \le \mathrm{E}_x\left[e^{-r\cdot\tau\wedge T(x_{\mathrm{P}}(y))}\cdot f\left(X_{\tau\wedge T(x_{\mathrm{P}}(y))}, y\right)\right] \stackrel{(i)}{\le} f(x, y), \quad (10)$$

where inequality (i) holds because the process $e^{-r\cdot t\wedge T(x_{\mathrm{P}}(y))}\cdot f(X_{t\wedge T(x_{\mathrm{P}}(y))}, y)$ is a positive supermartingale. Finally, note that $f(x, y)$ is the expected payoff if $\eta(y) = T(x_{\mathrm{P}}(y))$. Hence, the upper bound is achieved by exercising the option at the boundary $x_{\mathrm{P}}(y)$.

In region R_4, we have $\eta(y) \le \underline{T}(x^*)$ for $y > a(x^*)$, and thus,

$$\mathrm{E}_x\left[e^{-r\eta(y)}(\theta_{\mathrm{H}}X_{\eta(y)} - I_{\mathrm{A}})\right] \le \sup_\tau \mathrm{E}_x\left[e^{-r\cdot\tau\wedge\underline{T}(x^*)}(\theta_{\mathrm{H}}X_{\tau\wedge\underline{T}(x^*)} - I_{\mathrm{A}})\right]. \quad (11)$$

When $x^* \ge \hat{x}$, the process $e^{-r\cdot t\wedge\underline{T}(x^*)}(\theta_{\mathrm{H}}X_{t\wedge\underline{T}(x^*)} - I_{\mathrm{A}})$ is a positive supermartingale; hence, immediate stopping $\eta(y) = 0$ is optimal. QED

A2. Parametric Condition for $x^* \ge \hat{x}$

The cutoff x^* is defined by equation (4), where the right-hand side is increasing in x. In order for $x^* \ge \hat{x}$, it is necessary and sufficient to have

$$\frac{I_{\mathrm{P}}}{\beta - 1} \ge y_{\mathrm{P}}(\hat{x})\mathrm{E}_{\hat{x}}[e^{-rT(x_{\mathrm{A}}(1))}(\theta_{\mathrm{H}}x_{\mathrm{A}}(1) - I_{\mathrm{P}})] + (1 - y_{\mathrm{P}}(\hat{x}))\mathrm{E}_{\hat{x}}[e^{-rT(x_{\mathrm{P}}(0))}(\theta_{\mathrm{L}}x_{\mathrm{P}}(0) - I_{\mathrm{P}})]. \quad (12)$$

When $\theta_{\mathrm{L}} = 0$, this condition is equivalent to $[\beta/(\beta - 1)]^{\beta-1} \ge [r/(r - \mu)]^{\beta-1}[\beta - (\beta - 1)I_{\mathrm{P}}/I_{\mathrm{A}}]$, which trivially holds for I_{P} close to I_{A}. For $\theta_{\mathrm{L}} > 0$, I_{P} close to I_{A} is also sufficient for equation (12) to hold.

A3. Equilibrium Construction for $x^* < \hat{x}$

For $x \in (\hat{x}, x_{\mathrm{A}}(1))$ we set $a(x) = 1$ and denote by $x_{\mathrm{R}}(y)$ the inverse of $a(x)$ for $x \in [x^*, \hat{x}]$.[32] Similarly, for $x \in [x^*, \hat{x}]$ we set $b(x) = a(x)$ and denote by $x_{\mathrm{S}}(y)$ the inverse of $b(x)$ for $x \in (\hat{x}, x_{\mathrm{A}}(1))$.

[32] Below, we show that $x_{\mathrm{R}}(y)$ is strictly increasing; hence, its inverse $a(x)$ is well defined.

First, consider the agent. Define $S(x)$ to be the agent's optimal threshold to exercise the option conditional on $\theta = \theta_{\mathrm{H}}$ if the principal is exercising the option when $X_t \leq x$. Given an arbitrary stopping boundary $x_{\mathrm{R}}(y)$ of the principal, define $x_{\mathrm{S}}(y) \overset{\mathrm{def}}{=} S(x_{\mathrm{R}}(y))$. Such a definition of $x_{\mathrm{S}}(y)$ is motivated by the fact that pipetting information at $x_{\mathrm{S}}(y)$ triggers option exercise when $\theta = \theta_{\mathrm{H}}$. Lemma 2 formally defines $S(x)$. Its proof is contained in the online appendix.

LEMMA 2. For any $x_{\mathrm{R}} \in [x^*, \hat{x}]$, there exists a unique $S(x_{\mathrm{R}}) \in [\hat{x}, x_{\mathrm{A}}(1))$ such that

$$\underline{T}(x_{\mathrm{R}}) \wedge T(S(x_{\mathrm{R}})) = \underset{\tau \in \mathcal{M}(x_{\mathrm{R}})}{\arg\max} \mathrm{E}[e^{-r\tau}(\theta_{\mathrm{H}} X_\tau - I_{\mathrm{A}})],$$

where $\mathcal{M}(x_{\mathrm{R}})$ are all the stopping times τ such that $\tau \leq \underline{T}(x_{\mathrm{R}})$, with $\underline{T}(x_{\mathrm{R}}) \overset{\mathrm{def}}{=} \inf\{t > 0 : X_t \leq x_{\mathrm{R}}\}$ and $T(S(x_{\mathrm{R}})) = \inf\{t > 0 : X_t \geq S(x_{\mathrm{R}})\}$.

Next, consider the principal. The boundary $x_{\mathrm{R}}(y)$ is the unique solution to the principal's optimal stopping problem, given the expectation of the agent pipetting information at $x_{\mathrm{S}}(y)$. Pipetting of information at $x_{\mathrm{S}}(y)$ is valuable for the principal, while pipetting at $x_{\mathrm{R}}(y)$ is not. In the waiting region of R_4, that is, $x \in (x_{\mathrm{R}}(y), x_{\mathrm{S}}(y))$, the equilibrium value function $V_{\mathrm{P}}(x, y)$ of the principal satisfies the waiting ODE

$$rV_{\mathrm{P}}(x, y) = \mu x \frac{\partial}{\partial x} V_{\mathrm{P}}(x, y) + \frac{1}{2}\phi^2 x^2 \frac{\partial^2}{\partial x^2} V_{\mathrm{P}}(x, y), \tag{13}$$

subject to the boundary conditions given by

$$V_{\mathrm{P}}(x, y) = (y\theta_{\mathrm{H}} + (1 - y)\theta_{\mathrm{L}})x_{\mathrm{R}}(y) - I_{\mathrm{P}} \quad \text{(value matching at } x_{\mathrm{R}}(y)\text{)}, \tag{14}$$

$$\frac{\partial}{\partial x} V_{\mathrm{P}}(x_{\mathrm{R}}(y), y) = y\theta_{\mathrm{H}} + (1 - y)\theta_{\mathrm{L}} \quad \text{(smooth pasting at } x_{\mathrm{R}}(y)\text{)}, \tag{15}$$

$$\frac{\partial}{\partial y} V_{\mathrm{P}}(x_{\mathrm{S}}(y), y) = \frac{\theta_{\mathrm{H}} x_{\mathrm{S}}(y) - I_{\mathrm{P}} - V_{\mathrm{P}}(x_{\mathrm{S}}(y), y)}{1 - y} \quad \text{(pipetting at } x_{\mathrm{S}}(y)\text{)}. \tag{16}$$

Boundary condition (16) captures the effect of the agent pipetting information at $x_{\mathrm{S}}(y)$, implying that $V_{\mathrm{P}}(x, y)$ must satisfy smooth pasting with respect to y at $x_{\mathrm{S}}(y)$. Differential equation (13), together with boundary conditions (14)–(16), pins down $x_{\mathrm{R}}(y)$ as a solution to the differential equation

$$x_{\mathrm{R}}'(y) = \frac{\theta_{\mathrm{H}} S(x_{\mathrm{R}}) - I_{\mathrm{P}} - \frac{\theta_{\mathrm{H}} x_{\mathrm{R}}(1-\beta_2)+\beta_2 I_{\mathrm{P}}}{\beta_1 - \beta_2}\left(\frac{S(x_{\mathrm{R}})}{x_{\mathrm{R}}}\right)^{\beta_1} - \frac{\theta_{\mathrm{H}} x_{\mathrm{R}}(\beta_1 - 1)-\beta_1 I_{\mathrm{P}}}{\beta_1 - \beta_2}\left(\frac{S(x_{\mathrm{R}})}{x_{\mathrm{R}}}\right)^{\beta_2}}{\frac{[y\theta_{\mathrm{H}}+(1-y)\theta_{\mathrm{L}}](1-\beta_2)(1-\beta_1)-\beta_1\beta_2 I_{\mathrm{P}} x_{\mathrm{R}}^{-1}}{\beta_1 - \beta_2}\left[\left(\frac{S(x_{\mathrm{R}})}{x_{\mathrm{R}}}\right)^{\beta_1} - \left(\frac{S(x_{\mathrm{R}})}{x_{\mathrm{R}}}\right)^{\beta_2}\right](1 - y)} \tag{17}$$

and the initial condition $x_{\mathrm{R}}(y_{\mathrm{P}}(x^*)) = x^*$, where $\beta_1 > 1 > 0 > \beta_2$ are the two roots of $(1/2)\beta(\beta - 1) + \mu\beta = r$. The ODE (17) defines an increasing function in $(x^*, \hat{x})$ because the numerator is negative, because $S(x_{\mathrm{R}}) > \hat{x} > x_{\mathrm{R}}$, and the denominator is also negative, because $x_{\mathrm{R}} \geq x^* = x_{\mathrm{P}}(y_{\mathrm{P}}(x^*)) > x_{\mathrm{P}}(y) > [-\beta_2/(1 - \beta_2)]x_{\mathrm{P}}(y)$.

By construction, $x_{\mathrm{S}}(y_{\mathrm{P}}(x^*)) = S(x_{\mathrm{R}}(y_{\mathrm{P}}(x^*))) = S(x^*) < x_{\mathrm{A}}(1)$. To complete the definition of the agent's pipetting boundary, we set $b(x) = y_{\mathrm{P}}(x^*)$ for $x \in (x_{\mathrm{S}}(y_{\mathrm{P}}(x^*)), x_{\mathrm{A}}(1))$.

References

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

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

Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier. 2019. "Dynamic Persuasion with Outside Information." Working paper.

Board, Simon. 2007. "Selling Options." *J. Econ. Theory* 136 (1): 324–40.

Dixit, Robert K., and Robert S. Pindyck. 2012. *Investment under Uncertainty.* Princeton, NJ: Princeton Univ. Press.

Ely, Jeffrey C. 2017. "Beeps." *A.E.R.* 107 (1): 31–53.

Ely, Jeffrey C., and Martin Szydlowski. 2020. "Moving the Goalposts." *J.P.E.* 128 (2): 468–506.

Grenadier, Steven R., Andrey Malenko, and Nadya Malenko. 2016. "Timing Decisions in Organizations: Communication and Authority in a Dynamic Environment." *A.E.R.* 106 (9): 2552–81.

Grenadier, Steven R., and Neng Wang. 2005. "Investment Timing, Agency, and Information." *J. Financial Econ.* 75 (3): 493–533.

Gryglewicz, Sebastian, and Barney Hartman-Glaser. 2020. "Investment Timing and Incentive Costs." *Rev. Financial Studies* 33 (1): 309–57

Harrison, J. Michael. 2013. "Further Examples." In *Brownian Models of Performance and Control*, 155–70. New York: Cambridge Univ. Press.

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

Hörner, Johannes, and Andrzej Skrzypacz. 2016. "Selling Information." *J.P.E.* 124 (6): 1515–62.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *A.E.R.* 101 (6): 2590–615.

Kruse, Thomas, and Philipp Strack. 2015. "Optimal Stopping with Private Information." *J. Econ. Theory* 159:702–27.

McDonald, Robert, and Daniel Siegel. 1986. "The Value of Waiting to Invest." *Q.J.E.* 101 (4): 707–27.

Mirowski, Philip, and Robert Van Horn. 2005. "The Contract Research Organization and the Commercialization of Scientific Research." *Soc. Studies Sci.* 35 (4): 503–48.

Myerson, Roger B. 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323–58.

Pollard, David. 2002. *A User's Guide to Measure Theoretic Probability.* Cambridge: Cambridge Univ. Press.

Renault, Jérôme, Eilon Solan, and Nicolas Vieille. 2017. "Optimal Dynamic Information Provision." *Games and Econ. Behavior* 104:329–49.

Smolin, Alex. 2019. "Dynamic Evaluation Design." Working paper.

Stokey, Nancy L. 2008. *The Economics of Inaction: Stochastic Control Models with Fixed Costs.* Princeton, NJ: Princeton Univ. Press.