# Persuading the Principal To Wait*

Dmitry Orlov† Andrzej Skrzypacz‡ Pavel Zryumov†

First Draft: $13^{\text{th}}$ February, 2016

Current Version: $23^{\text{rd}}$ August, 2018

**Abstract**

A principal decides when to exercise a real option. A biased agent influences this decision by strategically disclosing relevant information. To persuade the principal to wait it is optimal to commit to delayed disclosure of all information. Without long-term commitment, this promise is credible only if the agent's bias towards delayed exercise is small; otherwise, the agent pipets information, probabilistically delaying the principal's action. When the agent is biased towards early exercise, his lack of commitment to remain quiet leads to immediate disclosure, hurting the agent. Our model applies to pharmaceutical companies conducting post-market clinical trials to influence the FDA or equipment manufacturers testing their products to attract customers.

## 1 Introduction

Decision makers commonly rely on interested parties to provide relevant information. In turn, agents use strategic communication to influence the decision makers. Conflicts arise when preferences over decisions are not 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 FDA) to keep their drug on the market

*We thank Mike Harrison, Johannes Hörner, Emir Kamenica, Ron Kaniel, Erik Madsen, Heikki Rantakari, Nina Varsava, and seminar participants at the University of Rochester, University of Pennsylvania, New York University, University of Chicago, Chicago Booth, Wharton School, Fuqua School of Business, HEC Lausanne and EPFL, Boston University, Midwest Economic Theory meeting, North American Econometric Society meeting, SITE, and European Econometric Society meeting for useful comments and suggestions.

†Rochester Simon School of Business, dmitry.orlov@simon.rochester.edu, pavel.zryumov@simon.rochester.edu.

‡Stanford Graduate School of Business, andy@gsb.stanford.edu.

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 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 with the solution to a dynamic persuasion mechanism (i.e., under commitment). As we show, when the agent is biased towards 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 towards early exercise, in equilibrium, he reveals all information immediately. Not only is the agent's non-commitment payoff smaller than in the optimal mechanism, but his equilibrium payoff can be even lower than in a game with no information acquisition. The ability to control information can hurt the agent.

In our real option model, a principal decides in every period whether or not to take an irreversible action, i.e., she decides when to exercise the option. The payoff from taking the action depends on two states, one of which is public and evolves exogenously via a geometric Brownian motion.[1] We extend this classic single-player setting by adding a strategic agent who can acquire and disclose additional verifiable information about the second state which is binary and is initially unobserved by either player. 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, i.e., to persuade her when to exercise the option.

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

**Persuasion to Wait.** We first characterize a (unique Markov-perfect) equilibrium when the agent is biased towards late exercise, i.e., 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.

In order to develop the intuition for our main results, consider a two-period example.[2] Suppose the agent is strongly biased and wants the principal to wait regardless of the state. When the game reaches the second period, optimal persuasion of the agent minimizes the probability of option exercise and 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. Note that 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 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 logic 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.[3] The equilibrium payoff of the principal is equal to her autarky payoff while the agent's payoff is strictly higher. Even though future information from the agent

[2]See Online Appendix A.4 for a formal treatment of the two-period model.

[3]This logic also extends to the Markov perfect equilibrium of our infinite-horizon continuous-time game.

could act as a carrot for the principal to wait, the agent is not able to utilize that incentive device in equilibrium due to the lack of dynamic commitment. If the agent could commit to future disclosure of information, he would incentivize the principal to wait in the first period by promising to disclose more information in period two. Pipetting is sub-optimal with long-term commitment in this setting.

Now reconsider the two-period example, but suppose 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. Such information strictly increases the principal's expected payoff when entering period two over the autarky benchmark. As a consequence, in the first period, she is less willing to exercise the option since the value of reaching the second period is now higher relative to her autarky payoff. 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, both of these intuitions affect equilibrium strategies, and the equilibrium behavior is more complex. For example, unlike in the two-period model, the principal's equilibrium policy is often non-monotone in the expected exercise payoff. The intuition is that under partial alignment of preferences, the agent finds it optimal to disclose all information when the expected exercise payoff is sufficiently high. In the vicinity of such states, it is optimal for the principal to wait for that information. Similarly, when the expected exercise payoff 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 best responds with pipetting 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. However, after beliefs change substantially, he can credibly delay full disclosure of information until his preferred exercise threshold.

**Persuasion to Act.** We contrast these equilibrium outcomes to the case when the agent is biased towards early exercise. Then, there exists a Markov perfect equilibrium 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 information even sooner and so on.[4] 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 with full commitment but is even worse than if the agent could not acquire any information at all. Finally, for some conditions, we show that, under commitment, the agent only discloses imprecise information that maximizes the probability of option exercise at time zero.

**Applications.** Among other contexts, our model captures information dynamics around product recalls. A concrete example is the post-market surveillance of drugs 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, i.e., remove it from the market if evidence indicates it is dangerous. The firm producing the drug can affect the FDA's decision by providing additional information or tests.[5] 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[6] and optimally timing them. We show that without long-term commitment, it is initially optimal for the firm to engage in strategic ignorance, i.e., 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.[7] 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

[4] Full information disclosure in equilibrium requires many periods. For a discussion of the two-period model see Online Appendix A.4.

[5] See "Guidance for Industry and FDA Staff. Post-market Surveillance Under Section 522 of the Federal Food, Drug, and Cosmetic Act" issued on May 16, 2016. Under these FDA guidelines, the manufacturer has the opportunity to provide additional information and identify specific surveillance methodologies before the FDA issues a recall.

[6] The FDA does not require controlled clinical trials to address its concerns, but "intend(s) that manufacturers use the most practical least burdensome approach to produce a scientifically sound answer."

[7] In Section 5 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.

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 supplementary material in Section A.3 of the Online Appendix). The value of commitment to future trials creates an incentive for firms to delegate the trials to a third party – an organizational structure that is documented in the pharmaceutical industry. Our model predicts that the outsourced trials are on average more informative than the in-house ones, but are conducted with a significant delay. We show that our assumptions fit this environment, and the equilibrium dynamics are consistent with the observed behavior of large pharmaceutical companies.

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. Over time public news arrives about the size of the market (we model the evolution of the market size as a geometric Brownian motion). The consumer's willingness to pay (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. We show 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 product testing by firms to persuade potential customers to buy. 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 Tasers 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, forward-looking and cannot commit to future actions. Transfers are either not allowed or do not fully resolve the conflict of interest.[8] Agents can strategically decide to remain ignorant about certain facts, but any information they acquire must be disclosed.

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

A second significant difference is that in both Ely (2017) and Renault et al. (2017) the optimal commitment policy is greedy (i.e., can be found by solving myopic problems) 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 non-commitment resembles the policies in these papers, the economic intuition for optimality of that strategy is different. In those papers, information cannot be used as a carrot because the receiver 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 Ely and Szydlowski (2017) and Smolin (2018) since they also study dynamic persuasion between two long-lived players. Their focus is on the commitment problem, 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 add to these papers by contrasting the

[8] In Lemma 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.

commitment and non-commitment solutions, by studying a different type of dynamic game, and by allowing partial alignment of preferences between the principal and the agent. One of our results is the conditions under which the commitment and non-commitment solutions coincide (they require sufficient agreement between the players).

Equilibrium unraveling in persuasion to act is reminiscent of Au (2015) where the principal is privately informed about her preferences. Our unraveling result in Section 4 holds even though the preferences of the principal are known, but the lack of commitment plays an important role in both results. Bizzotto, Rudiger, and Vigier (2016) is also a closely related paper that studies a finite-horizon persuasion game without commitment. Their main finding is that delayed persuasion may be optimal since, as the deadline of the game approaches, the principal's exercise region increases, benefiting the agent. In our paper, the environment is stationary and the difficulty of persuading the principal depends only on her expectations of equilibrium communication. Henry and Ottaviani (2018) analyze a similar setting to ours, 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 exercise the option suboptimally making such arrangements infeasible. Hörner and Skrzypacz (2016) analyze a dynamic persuasion game with non-contractable monetary transfers. They find that gradual persuasion is optimal to resolve the ex-post holdup problem, so the economic mechanism 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 only use cheap talk. We show that when the informed agent's bias towards delayed exercise is small, both verifiable and non-verifiable 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 prefers to exercise early, 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 impatient, it leads to full and immediate information sharing. These results highlight the distinction between communication of soft (unverifiable) versus 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), Gryglewicz and Hartman-Glaser (2018), and Kruse and Strack (2015). These papers study the role of incentive contracts with monetary transfers in managing conflicts of interest. To the best of our knowledge, we are the first to analyze the limits of strategic management of hard information for aligning incentives in the context of real options when monetary transfers are not feasible.

The rest of the paper is organized as follows. In Section 2 we present our main model. In Section 3 we characterize the unique equilibrium in case the agent is biased towards early exercise and persuades the principal to wait. Propositions 1 and 2 are the main results of the paper. In Section 4 we consider the case where the agent is biased towards early exercise and persuades the principal to act. In Section 5 we allow the agent to be privately informed and show that equilibria constructed in Sections 3 and 4 are robust to the introduction of private information. Section 6 concludes.

# 2 Model

## 2.1 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 publicly acquire and disclose. 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)).

**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 decides on the optimal timing of this decision, i.e., 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 $\theta_H > \max(\theta_L, 0)$. Neither party initially observes the realization of $\theta$ 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 $\theta$ are independent. Both $X$ and $\theta$ are defined on an underlying probability space $(\Omega, \mathcal{F}, \mathrm{P})$.

The flow payoffs of both players are zero prior to 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 $\theta$ are

$$v_A = e^{-rt}\,(\theta X_t - I_A), \qquad v_P = e^{-rt}\,(\theta X_t - I_P).$$

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 $\theta$ 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 different cost of exercising the option $I_P \neq I_A$. Intuitively, if $I_P < I_A$, then for any given $\theta$ 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.

**Remark about payoffs.** We model preferences in terms of the call option with zero flow costs and only terminal payoffs. In Section 3.2 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 of interest between the parties by assuming 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. Our model is relevant to any such situation that involves misalignment of preferences between a principal who has decision rights over the timing of an action and an agent who has access to additional information.

## 2.2 Strategies and Equilibrium Concept

Within every "period" the innovation to $X_t$ is first realized, then the agent provides an informative signal about $\theta$ by conducting a test. The test induces a posterior over $\theta$ 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 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 past history of the state process $X$ and information learned about $\theta$ *up to* time $t$. The principal's strategy is, additionally, a function of the information learned about $\theta$ *at* time $t$.

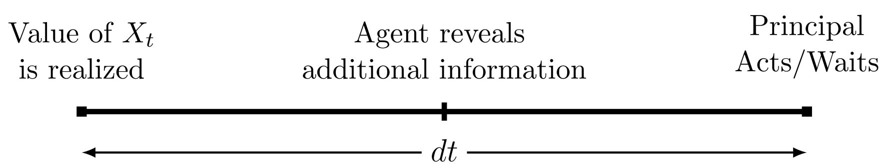


Figure 1: Timing of Events in a Short Interval of Time

In our model the agent controls information flow about $\theta$ 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 $X$ as well as any signals communicated by the agent up to time $t$.[9] Denote by $Y_t$ the posterior belief about $\theta$ given by

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

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 $\theta$ and the current level of process $X_t$. A player's strategy is Markov if it depends on past 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 $\theta$, and past messages. We require 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", i.e., it does not foresee the future evolution of $X$.

**Definition.** *An admissible belief process $Y = (Y_t)_{t\geq 0-}$ is a right-continuous with left limits martin-*

[9] Formally, $\mathcal{F}_t$ is the $\sigma$-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 Pollard (2002), Appendix E.1 for details.

*gale that takes values in* $[0,1]$ *such that* $Y_t$ *is independent of* $X$*'s future innovations* $\{X_{t+s}/X_t\}_{s\geq 0}$ *for every* $t \geq 0$*. We denote by* $\mathcal{Y}$ *the set of all admissible belief processes.*

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 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, i.e., time 0 posterior $Y_0$ may be different from the initial prior $Y_{0-}$, before the principal had the first opportunity to exercise the option. The strategy is Markov if the information the agent discloses about $\theta$ depends on the history only through the current state $(X_t, Y_t)$.

**Definition** (**Markov strategy of the agent**). *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 $\theta = \theta_H$ over time. For a given such strategy define $D_t$ to be the cumulative probability of having disclosed $\theta = \theta_H$ up to time $t$. Then the induced posterior belief $Y_t$ is either equal to one if $\theta = \theta_H$ has been disclosed or is equal to $Y_t^{ND}$ given by Bayes rule

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

The set of disclosure strategies with first-order stochastic dominance (FOSD) over the corresponding cumulative probabilities of disclosure as a partial order forms a lattice. 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**). *An admissible belief process* $Y$ *is a pipetting strategy with respect to an (upper) boundary* $b(x)$ *if there exists a cumulative probability of disclosure* $D = (D_t)_{t\geq 0}$ *such that at every time* $t \geq 0$

*(i) the posterior belief* $Y_t$ *is equal either to* 1 *or to* $Y_t^{ND}$ *given by* (1)*,*

*(ii) conditional on lack of disclosure the posterior belief* $Y_t^{ND}$ *is weakly below the boundary* $b(X_t)$*,*

*(iii) process* $D = (D_t)_{t\geq 0}$ *is the minimal process satisfying (i) and (ii).*

The ranking in condition (iii) is based on the FOSD order of the cdfs. Notice that disclosure strategies that satisfy (ii) constitute a meet-semilattice since for any two $D^1$ and $D^2$ satisfying (ii) process $D_t = \min(D_t^1, D_t^2)$ also satisfies $Y_t^{ND} \leq b(X_t)$ for all $t \geq 0$. It is then easy to verify that there exists a unique minimal process satisfying $(i)$ and $(ii)$ and thus a pipetting strategy is well-defined with respect to any upper boundary $b(x)$. If beliefs start at $Y_{0-} > b(X_0)$ then the pipetting strategy induces a jump at $t = 0$ such that the posterior is either 1 or $b(X_0)$. Afterward, posterior beliefs change only at the boundary $b(x)$, either reflecting off this boundary or jumping to 1. In other words, a pipetting strategy delays the disclosure of information 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 $(x, y)$.[10] 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).** *Principal's Markov strategy is a Borel set $\mathbb{T} \subseteq [0, +\infty) \times [0, 1]$, such that the principal exercises the option at the first hitting time of $\mathbb{T}$*

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

We define a Markov equilibrium as a pair of Markov strategies that are mutual best responses.

**Definition.** *A Markov perfect equilibrium (MPE) is a pair of Markov strategies of the agent and principal $(Y^*, \mathbb{T}^*)$ such that*

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

$$\tau^* \in \underset{\tau \in \mathcal{M}(Y^*)}{\arg\max} \, \mathrm{E}\left[e^{-r\tau}\Big((Y_\tau^* \cdot \theta_H + (1 - Y_\tau^*) \cdot \theta_L)X_\tau - I_P\Big) \,\Big|\, X_0 = x, Y_{0-} = y\right] \tag{2}$$

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

2. **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 $\tau^*$:*

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

$$Y^* \in \underset{\tilde{Y} \in \mathcal{Y}}{\arg\max}\, \mathrm{E}\left[e^{-r\tau^*}\left((\tilde{Y}_{\tau^*}\theta_H + (1-\tilde{Y}_{\tau^*})\theta_L)X_{\tau^*} - I_A\right) \,\Big|\, X_0 = x, Y_{0-} = y\right]. \tag{3}$$

While any MPE $(Y^*, \mathbb{T}^*)$ is a pair of Markov strategies, the maximum in (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 $t$, i.e., whether $\tau \leq t$, 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 dynamically consistent. Principal's strategy is dynamically consistent because her best response given by (2) is a solution to an optimal stopping problem in a Markov environment. The agent's strategy is dynamically consistent because it is Markov and equation (3) ensures that it maximizes agent's utility for any initial state $(x, y)$.

### 2.3 Autarky Thresholds

It is helpful to define two important thresholds at which the principal and the agent would like to exercise the real option in case additional information about $\theta$ is unavailable, that is, in case they can only observe the evolution of $X_t$. For a given belief $Y_t \equiv y$ the optimal stopping decision of the principal (or the agent if he were given control rights) that is based *only on exogenous evolution of* $X$ can be characterized by a first entry time into the set $\{x \geq x_P(y)\}$ (or $\{x \geq x_A(y)\}$ for the agent) with[11]

$$x_i(y) = \frac{\beta}{\beta - 1} \cdot \frac{I_i}{y\theta_H + (1-y)\theta_L}, \qquad i \in \{A, P\}$$

where $\beta > 1$ is the positive root of the equation $\frac{1}{2}\phi^2\beta(\beta-1) + \mu\beta = r$ (see Dixit and Pindyck (2012) for details). When the principal holds a higher belief $y$ she would like to exercise the option earlier, i.e., $x_P(y)$ is decreasing in $y$. The autarky thresholds $x_P(y)$ and $x_A(y)$ summarize the conflict between the principal and the agent. If $I_P < I_A$ then principal wants to exercise the option earlier than the agent. Player $i$'s payoff at their autarky threshold (if the latter is finite) is given by

$$(y\theta_H + (1-y)\theta_L)x_i(y) - I_i = \frac{I_i}{\beta - 1}, \qquad i \in \{A, P\}.$$

An equivalent characterization of the principal's autarky stopping set is to exercise the option if $Y_t > y_P(X_t) = x_P^{-1}(X_t)$. This is a useful characterization when $Y_t$ is dynamically changing as the

[11] This expression is valid only for $y > \frac{-\theta_L}{\theta_H - \theta_L}$. For $y \leq \frac{-\theta_L}{\theta_H - \theta_L}$ we set the autarky threshold $x_i(y) = +\infty$ since neither player wants to exercise the option for any level of $x$.

agent discloses information about $\theta$ over time.

## 3 Persuasion to Wait

In this section, we characterize the unique Markov perfect equilibrium in case $I_P < I_A$, that is when the principal prefers to exercise the option earlier (at a lower threshold) than the agent: $x_P(y) < x_A(y)$.

**Proposition 1.** *There exists an essentially*[12] *unique Markov Perfect Equilibrium. Equilibrium strategies are characterized by two boundaries $a(x)$ and $b(x)$.*

- *For $x < x_A(1)$ the agent follows the pipetting strategy against the boundary $b(x)$. For $x \geq x_A(1)$ he fully discloses $\theta$.*
- *For $x \in (x_P(1), x_P(0))$ the principal exercises the option either when $y > a(x)$ or when $y = 1$. For $x \geq x_P(0)$ she exercises the option regardless of $y$ and for $x \leq x_P(1)$ she waits regardless of $y$.*

*Moreover, $a(x)$ is a quasi-convex function weakly higher than $b(x)$, and both boundaries coincide with the autarky threshold $y_P(x)$ to the left of the vertex of $a(x)$.*

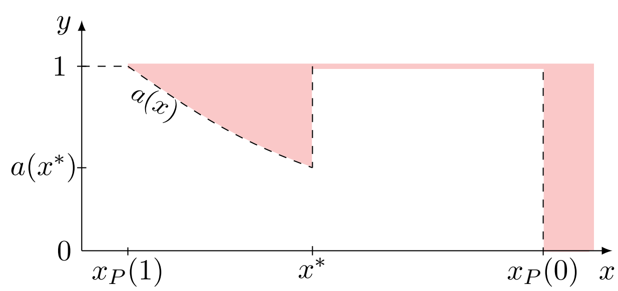


(a) Stopping set of the principal.

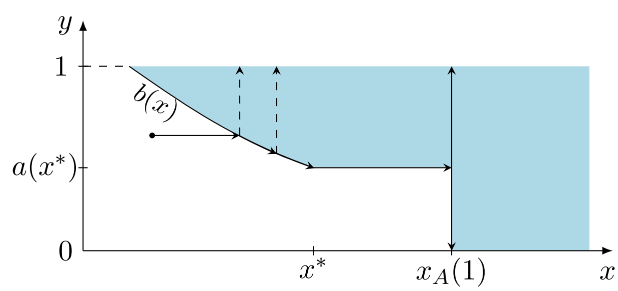


(b) Communication strategy of the agent.

Figure 2: Equilibrium strategies with $x^* \geq \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$ and $x_P(0) > x_A(1)$.

Figure 2 depicts the equilibrium strategies of the principal and the agent graphically. We denote $x^*$ to be the vertex of the principal's action boundary $a(x)$ and formally derive it in (4). To the left of $x^*$ both $a(x)$ and $b(x)$ are equal to principal's autarky threshold $y_P(x)$. If $x_P(0) < x_A(1)$ then $x^* = x_P(0)$ and we are in a "large conflict case". For such parameters $a(x) = y_P(x)$ for all $x \leq x_P(0)$ and the equilibrium only features pipetting.

[12]Essential uniqueness means uniqueness of the outcome of the game, i.e., distribution of $(X_\tau, \tau)$ conditional on $\theta$.

The equilibrium behavior is more complex when $x^* < x_P(0)$. Figure 2 is drawn for the case $x^* \geq \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$ and $x_P(0) > x_A(1)$.[13] In this case, the equilibrium boundaries are explicitly given by

$$a(x) = \begin{cases} y_P(x) & \text{if} \quad x < x^*, \\ 1 & \text{if} \quad x \geq x^*, \end{cases} \qquad b(x) = \begin{cases} y_P(x) & \text{if} \quad x < x^*, \\ y_P(x^*) & \text{if} \quad x \geq x^*. \end{cases}$$

In this case, the equilibrium has the following properties. First, along the equilibrium path the principal only stops before $x_P(0)$ if she knows that $\theta = \theta_H$. Even though her stopping set includes other states, as illustrated in Figure 2a, the agent discloses $\theta = \theta_H$ such that the option is never exercised in its interior. If the game starts in the interior of the shaded set, the agent immediately reveals information to move beliefs either to 1 or to the boundary $b(x)$. Moreover, in equilibrium, the agent reveals $\theta$ partially to either induce immediate action or elicit waiting until $x_A(1)$ is reached.

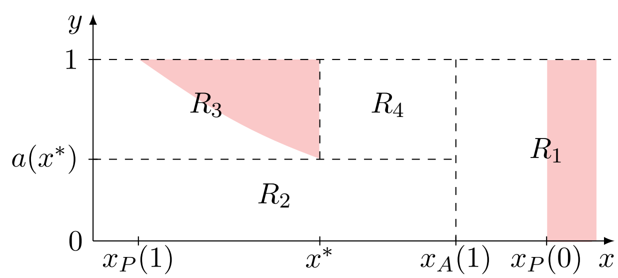


Figure 3: Partition of the state space $[0, +\infty) \times [0, 1]$ into four regions $R_1, R_2, R_3, R_4$.

The equilibrium features four distinct regions $\{R_i\}_{i=1}^4$ depicted in Figure 3. In $R_1 = \{x \geq x_A(1)\}$ the agent fully discloses $\theta$ 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 $x_A(1)$. If the starting beliefs are sufficiently high and the initial $x$ is low, i.e. the game starts in $R_3 = \{x \leq x^*, y > y_P(x^*)\}$, the principal prefers to exercise the option at her autarky threshold $a(x) = y_P(x)$, rather than wait until region $R_1$ is reached. When $(x, y)$ is below the principal's autarky threshold the state moves only horizontally

[13] To simplify exposition, the equilibrium characterization of the remaining case $x^* < \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$ and $x_P(0) > x_A(1)$ is relegated to the Appendix.

because the agent does not communicate. When $x$ increases up 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 $a(x^*)$ the agent stops providing information until the 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 to the principal that either reveals $\theta = \theta_H$ and results in instantaneous option exercise, or lowers beliefs to $a(x^*)$, 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 $\mathbf{R}_1$.** When the principal knows $\theta$ then she exercises the option at her autarky thresholds $x_P(1)$ and $x_P(0)$, respectively. Moreover, when $X_t \geq x_P(0)$ then the principal is past her optimal exercise thresholds for both $\theta_H$ and $\theta_L$, 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 $\theta$ when $X_t \geq x_A(1)$ he 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, i.e., until $x_P(0)$. In other words, two rounds of eliminating dominated strategies imply that it is optimal for the agent to disclose $\theta$ when $X_t$ reaches $x_A(1)$.

**Equilibrium in region $\mathbf{R}_2$.** The expectation of learning $\theta$ at $x_A(1)$ increases the option value of waiting for the principal relative to her autarky strategy. 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 $\theta$ at $x_A(1)$. Formally $x^*$ is the unique solution to

$$\frac{I_P}{\beta - 1} = y_P(x^*)\mathrm{E}_{x^*}\left[e^{-rT(x_A(1))}(\theta_H x_A(1) - I_P)\right] + (1 - y_P(x^*))\mathrm{E}_{x^*}\left[e^{-rT(x_P(0))}(\theta_L x_P(0) - I_P)\right], \quad (4)$$

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 left hand side of (4) is the payoff from immediate exercise at $(x^*, y_P(x^*))$ and the right hand side of (4) is the expected payoff from waiting to learn $\theta$ which results either in exercise at $x_A(1)$ if $\theta = \theta_H$ or at $x_P(0)$ if $\theta = \theta_L$. For $y < y_P(x^*)$ the principal strictly prefers to wait for the agent's information disclosure. Anticipating this, the agent does not communicate any information for $y < y_P(x^*)$ and waits until $X_t$ reaches $x_A(1)$.[14]

**Equilibrium in region $\mathbf{R}_3$.** A higher belief about $\theta$ makes the principal more likely to regret waiting for agent's disclosure at $x_A(1)$. Consequently, in region $R_3$ the principal would rather exercise the option in her autarky set $\{x > x_P(y)\}$. In equilibrium the agent communicates additional information through pipetting against boundary $b(x)$, which is equal to $y_P(x)$ in $R_3$. Such pipetting generates posterior beliefs $Y_t \in \{y_P(X_t), 1\}$ and does not affect the principal's expected value at the boundary $a(x) = y_P(x)$ due to linearity of her payoff in $y$ upon exercise. Since $x_P(y)$ is an optimal option exercise threshold, principal's value for $y < y_P(x)$ is also equal to her autarky payoff, i.e., her expected value absent any persuasion. Thus stopping at $y > y_P(x)$ in region $R_3$ is still optimal for the principal despite the agent pipetting information at that boundary.

In order to show that pipetting information at the principal's stopping boundary is optimal for the agent, we show that disclosing $\theta = \theta_H$ 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 $a(x)$ and show that $Y^*$ delivers this upper bound. Notice that the agent always benefits from additional delay if $\theta = \theta_L$ 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 $\theta$ at $\tau$. Such disclosure does does not affect option exercise if $\theta = \theta_H$ and weakly delays it if $\theta = \theta_L$. As a result, the agent's expected payoff can be written as

$$Y_{0-} \cdot \mathrm{E}\left[e^{-r\tau}\big(\theta X_\tau - I_A\big) \,\big|\, \theta = \theta_H\right] + (1 - Y_{0-}) \cdot \mathrm{E}\left[e^{-rT(x_P(0))}\big(\theta x_P(0) - I_A\big) \,\big|\, \theta = \theta_L\right].$$

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 $\theta = \theta_H$. This implies that we can identify the stopping time $\tau$ with the first time the agent discloses $\theta = \theta_H$.[15] We denote by $D_t$ the cumulative, history dependent, conditional probability of disclosing $\theta = \theta_H$ up to time $t$. Intuitively, $\Delta D_t$ is the probability of exercising the option at time $t$ knowing that $\theta = \theta_H$. The agent's payoff conditional on $\theta = \theta_H$ can be written as

[14] The equilibrium could feature only "inconsequential" communication in this region that does allow the posterior belief to exceed $y_P(x^*)$ prior to fully disclosing $\theta$ at $x_A(1)$.

[15] We derive this result formally in Lemma 4 in the Appendix. The revelation principle stated in Myerson (1986) does not apply in our setting due to lack of agent's long-term commitment.

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

Next, we perform a change of time in the above integral to reduce the agent's problem to a point-wise optimization. Given (1) process $D_t$ uniquely pins down the decreasing posterior belief $Y_t^{ND}$ conditional on staying on the path of no disclosure, and vice versa. Since this (conditional) posterior belief is monotonically decreasing over time, instead of integrating (5) over calendar time $t$ we can integrate it over levels of $Y_t^{ND}$ when $\theta=\theta_H$ is disclosed. Define $\eta(y)=\inf\{t : Y_t^{ND}\le y\}$ 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 $D_{\eta(y)} = 1-\frac{1-Y_{0-}}{Y_{0-}}\cdot\frac{1}{1-y}$ and we can rewrite (5) as[16]

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

Intuitively, $\eta(y)$ represents the random time at which $\theta=\theta_H$ is disclosed for a 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 $Y_{0-}<a(X_0)$.[17] Then in order for $\eta(y)$ to be feasible it must be the case that the path of $X_t$ stays below $a^{-1}(y)$ 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, i.e., at the highest possible level of $X_t$ in the principal's waiting region, which is equal to $a^{-1}(y)$. This implies that the optimal $\eta(y)$ is the first hitting time of $a^{-1}(y)$:

$$\mathrm{E}\left[e^{-r\eta(y)}\big(\theta_H X_{\eta(y)} - I_A\big)\right] \le \mathrm{E}\left[e^{-rT(a^{-1}(y))}\big(\theta_H a^{-1}(y) - I_A\big)\right] \qquad \forall y\in(a(x^*),Y_{0-}]. \tag{7}$$

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

**Equilibrium in region $\mathbf{R}_4$.** 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 $a(x^*)$. This policy is optimal despite the principal willing to wait in this region. The intuition is as follows. In this region, since $X_t$ is stochastic, the agent faces a risk of hitting threshold $x^*$ from the right (moving to $R_3$). This possibility lowers his option value of waiting for all $X_t>x^*$ and, as a result, he finds it optimal to speed up option exercise

[16]This representation holds if $D_t$ is continuous and goes through with a minor modification if $D_t$ has jumps.
[17]This argument goes through with minor modifications if $Y_{0-}>a(X_0)$.

conditional on $\theta = \theta_H$ relative to his first best threshold $x_A(1)$. If there were no risk of $X_t$ dropping back to the principal's action region, i.e., if $\phi = 0$ and $\mu > 0$, then he would wait in this region until $x_A(1)$.

Direct computation shows that in the parametric case when $x^* \geq \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$ immediate 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 $\eta(y) = 0$ in (7) when for $(x, y) \in R_4$.[18] Finally, the principal has no incentives to exercise the option for $y < 1$ In this region since she expects the agent to disclose valuable information the next instant.

**Equilibrium Uniqueness.** We show uniqueness of the equilibrium outcome by iteratively eliminating dominated strategies of the agent and the principal. By the previous logic this 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 $\theta$ 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 prior to $x_P(y)$ 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 (7). We can 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.

With this characterization of the equilibrium, we can also compare equilibrium payoffs with the autarky payoffs.

**Proposition 2.** *Let $V_P(x, y)$ and $V_A(x, y)$ be the Markov perfect equilibrium payoffs of the principal and agent respectively. Let $V_P^{Aut}(x, y)$ and $V_A^{Aut}(x, y)$ 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*

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

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

$$V_P(x, y) = V_P^{Aut}(x, y) \quad \text{if} \quad y \geq a(\max(x, x^*)) \text{ or } x \geq x_P(0) \text{ or } y \in \{0, 1\},$$

[18]When $x^* < \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$ 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 prior to the agent's first best. We characterize the equilibrium boundaries $(a(x), b(x))$ for $x \in [x^*, x_A(1)]$ in the Appendix.

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

This proposition shows that there is an asymmetry in the split of gains from information about $\theta$ 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 $a(x)$ has no value for her since she is locally indifferent between immediate exercise and waiting. When the game starts above $a(x)$, 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 $a(x)$ does not benefit the principal due to the optimality of $a(x)$ absent pipetting.[19] Therefore, for all priors in region $R_3$ principal's payoffs are equal to her autarky payoffs. For priors in region $R_2$ they are strictly higher (the other two regions vary).

### 3.1 Value of Dynamic Commitment

The equilibrium constructed in Proposition 1 features pipetting of information along the boundary $a(x)$ 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 $Y$, i.e.,

$$\mathcal{M}^*(Y) = \text{argmax}_{\tau \in \mathcal{M}(Y)} \mathrm{E}\left[e^{-r\tau}((Y_\tau \theta_H + (1-Y_\tau)\theta_L)X_\tau - I_P)\right].$$

A *feasible dynamic mechanism* is a strategy of the agent $Y$ for which the set $\mathcal{M}^*(Y)$ is nonempty.

[19]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 the Introduction.

Many mechanisms satisfy this requirement, e.g., any Markov posterior belief process $Y$ leads to a nonempty $\mathcal{M}^*(Y)$. Denote by $\mathcal{Y}^M$ 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:*

$$Y^M \in \arg\max_{\hat{Y} \in \mathcal{Y}^M} \left[ \sup_{\tau \in \mathcal{M}^*(\hat{Y})} \mathrm{E}\left[e^{-r\tau} \cdot ((Y_\tau \theta_H + (1 - Y_\tau)\theta_L) X_\tau - I_A)\right] \right]. \tag{8}$$

There is an important difference between the optimal persuasion mechanism and the Markov perfect equilibrium formulated in Section 2.2. In the mechanism, the agent commits to the full dynamic strategy at time zero, 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 $(x, y)$. This difference is analogous to the difference between the Stackelberg and Cournot equilibria in a standard duopoly game. Proposition 3 characterizes the optimal dynamic persuasion mechanism and illustrates that it requires making ex-post suboptimal promises.

**Proposition 3.** *Suppose that $X_0 \leq x_P(Y_{0-})$. 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 $\theta$. Disclosure threshold $\bar{x}(Y_{0-})$ is strictly decreasing for $Y_{0-} > a(x^*)$ and coincides with $x_A(1)$ for $Y_{0-} \leq a(x^*)$. There is no pipetting of information under long-term commitment.*

When $Y_{0-} \leq a(x^*)$ the agent obtains first best for $\theta = \theta_H$ in the equilibrium of Proposition 1 and commitment is not valuable. For $Y_0 > a(x^*)$ 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 $\bar{x}$. 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.[20] 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 $\theta$ 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 till $x_P(0)$. Proposition 3 implies that slow pipetting of information occurring

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

at the boundary $b(x)$ 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.

Proposition 3 shows that when the belief about $\theta$ is sufficiently low, i.e., $Y_{0-} \leq a(x^*)$, the solutions to the dynamic game and the dynamic mechanism coincide: the principal waits for the agent to fully disclose $\theta$ at $x_A(1)$. 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 b$\theta$, and the volatility of public information.

**Lemma 1.** *The range of beliefs $[0, a(x^*)]$ for which the principal is willing to wait for the agent to fully disclose $\theta$ at $x_A(1)$ is*

*(i) larger when preferences are more aligned, i.e., $\frac{\partial y^*}{\partial I_P} > 0 > \frac{\partial y^*}{\partial I_A}$;*

*(ii) larger when information aout $\theta$ is more valuable, i.e., $\frac{\partial y^*}{\partial \theta_H} > 0 > \frac{\partial y^*}{\partial \theta_L}$.*

*(iii) larger when $X$ is more volatile and $\theta_L = 0$, i.e., $\left.\frac{\partial y^*}{\partial \phi}\right|_{\theta_L=0} > 0$.*

First, when preferences of players are more aligned, i.e., $I_P - I_A$ is low, the principal's cost of waiting until $x_A(1)$ is small relative to the benefit of learning $\theta$. As a result the stopping region $y > a(x)$ shrinks and $a(x^*)$ goes up. Second, an increase in $\theta_H$ 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 $\theta_L$ has no effect on incentives of the agent to fully disclose $\theta$ at $x_A(1)$. However, it makes the information disclosed at $x_A(1)$ of better quality since the principal avoids premature option exercise conditional on $\theta = \theta_L$. As a result, waiting is a more attractive option when $\theta_H - \theta_L$ is high. Finally, notice 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 $X$ needs to cover the same distance from the autarky threshold of the principal to the point when the agent discloses $\theta$. A more volatile process $X$ moves faster thus reducing the expected cost of waiting for information and increasing $a(x^*)$.

## 3.2 Product Recalls and Abandonment Options

In the Introduction, 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.[21] In this section, we

[21] Henry and Ottaviani (2018) analyze firm's incentives to quit experimentation when the initial approval threshold is set by the FDA in equilibrium.

map the equilibrium constructed above to this environment, and more generally to abandonment options. We show that the equilibrium corresponds to the FDA requiring additional tests for drugs which it considers dangerous (involuntary recalls), while recalls of seemingly safe drugs are delegated to the manufacturer (voluntary recalls).

State $\theta$ is binary and is either $\theta_L$ or $\theta_H$. In our notation $\theta = \theta_H$ corresponds to 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 $\theta \cdot X$.[22]

As long as the drug is on the market, it generates the firm (agent) an expected profit flow which 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 $\tau$ then the expected payoffs of the agent and the principal conditional on $\theta$ are

$$v_A = \int_0^\tau e^{-rt}(F_0 - F_1 \cdot \theta X_t)\, dt, \qquad v_P = \int_0^\tau e^{-rt}(W_0 - W_1 \cdot \theta X_t)\, dt. \tag{9}$$

If $F_1 < 0$, then under complete information about $\theta$ and high level of $X$ the two players agree on a recall policy (since their flow payoffs conditional on $\theta$ have the same sign). The friction/disagreement is caused by uncertainty about the state through the misalignment between the principal's and the agent's option values of waiting for additional public information. The principal has a lower tolerance for the risky drug if $\frac{W_0}{W_1} < \frac{F_0}{F_1}$.

The payoff structure (9) is equivalent to the one described in Section (2), i.e., 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

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

which corresponds to $I_A = \frac{r-\mu}{r} \cdot \frac{F_1}{F_0}$ in the terminology of Section 2. Similarly the principal's expected payoff is an affine function of $\mathrm{E}\left[e^{-r\tau}(\theta X_\tau - I_P)\right]$ with $I_P = \frac{r-\mu}{r} \cdot \frac{W_1}{W_0}$. Hence, the analysis in Section 3 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 post-market surveillance program conducted by the FDA. If the prior beliefs and parameters are such that $Y_{0-} \leq a(x^*)$, then the timing of the action has a compromise property: in case the results of the test run by the agent

[22] In Section A.5 of the Online Appendix we discuss the case when $X_t$ is a public news process about $\theta$ and refer to the option as a Wald option. The analysis is qualitatively unchanged.

at $x_A(1)$ are good, the principal acts at the agent's optimal point; when the results are bad, the principal further delays action until her optimal threshold $x_P(0)$. This has two consequences. First, anytime $\theta = \theta_H$ is revealed, the firm does not need to be compelled to recall the drug – it is happy to do so voluntarily. That is not true in the pipetting region $y > b(x)$, whereupon observing that $\theta = \theta_H$ the FDA decides to recall the drug and that is against the wishes of the firm. Second, the compromise property implies that the equilibrium is robust to replacing the verifiable information about $\theta$ with cheap talk.[23] The reason is that at $x_A(1)$ the incentives of the firm and the regulator are aligned. If the firm does not perform a voluntary recall past $x_A(1)$, the FDA can credibly infer that $\theta = \theta_L$ and stop accordingly. This is, again, not true in the pipetting region $y > b(x)$: in this region the firm would prefer to mislead (at least temporarily) the FDA that $\theta = \theta_L$, rather than truthfully reveal that $\theta = \theta_H$.

# 4 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 towards early option exercise either because the project gives him private benefits or 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, i.e., if he never communicated with the principal. [24] 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 4.** *There exists a Markov Perfect Equilibrium in which the agent fully discloses $\theta$ at time 0 for all $(X_0, Y_{0-})$.*

In such equilibrium the principal achieves first best by adopting a very "strong bargaining position": she threatens not to exercise the option until either $\theta$ is fully known, or the level of $X_t$ is so high that $\theta$ is irrelevant for the option exercise decision, i.e., $X_t \geq x_P(0)$. Given the equilibrium strategy

[23]The equilibrium outcome of our model coincides with Grenadier, Malenko, and Malenko (2016) only when $Y_{0-} < a(x^*)$.

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

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 $\theta$ immediately, or to wait until $X_t = x_P(0)$; any partial information acquisition does not affect the timing of option exercise. If $\theta$ 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 $\theta$ at time 0.

Next we provide sufficient conditions under which fully disclosing $\theta$ at time 0 is the unique equilibrium outcome.

**Proposition 5.** *If $X_0 \leq x_P(1)$, then the equilibrium outcome of Proposition 4 is essentially unique. Moreover if $\phi > 0$ there exists a boundary $\underline{a}(x)$ such that in any MPE $(Y^*, \mathbb{T}^*)$ the stopping set $\mathbb{T}^* \subseteq \{(x, y) : y \geq \underline{a}(x)\}$ and $\underline{a}(x) > y_P(x)$ for $x \in (x_P(1), x_P(0))$.*

We prove the first part of Proposition 5 by noticing that the agent's payoff can be improved by conducting all equilibrium communication at time 0. This induces the same distribution of posteriors $Y_\tau$ as in equilibrium but the principal, anticipating no further communication, acts sooner than in equilibrium, namely at $x_P(y)$. Furthermore, since $\mathrm{E}\left[e^{-rT(x_P(y))}\right]$ is convex in $y$ the agent's payoff is maximized when $Y_\tau \in \{0, 1\}$. By disclosing $\theta$ immediately the agent guarantees that the principal stops at either $x_P(1)$ or $x_P(0)$ and achieves the upper bound of his equilibrium payoff. Figure 4b illustrates the information the agent shares in any equilibrium.

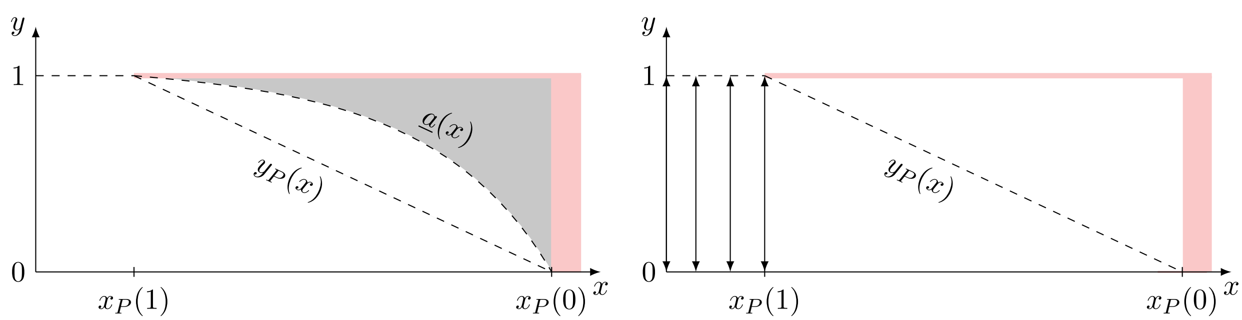


(a) Stopping set of the principal.



(b) Communication strategy of the agent.

Figure 4: Minimal equilibrium communication and implied waiting region.

The second part of Proposition 5 states 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 proves that the agent fully discloses $\theta$. 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 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 4a. 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 $Y_{0-} \in (y_P(x), \underline{a}(x))$ and $x \leq x_P(0)$. 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 $\theta$ when $X_t \leq x_P(1)$ and no longer exercises the option at $t = 0$. By providing additional information about $\theta$ the agent accelerates the option exercise conditional on $\theta = \theta_H$ but delays it conditional on $\theta = \theta_L$. Parametric condition of $x_A(0) < X_0$ implies that the 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 $\theta$ at all.

Next, we characterize the optimal dynamic persuasion mechanism in the special case of $\theta_L = 0$ and illustrate why it cannot be sustained in equilibrium without commitment. 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.

**Lemma 2.** *Suppose $\theta_L = 0$ and $I_A < 0$. The optimal dynamic persuasion mechanism sends a single message at time $t = 0$ such that the induced posterior belief is 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 $\theta$ immediately. There is no additional communication for any $t > 0$.*[25]

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 the agent was to send a pooling message as in Lemma 2, however, unlike in the persuasion

[25] 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.

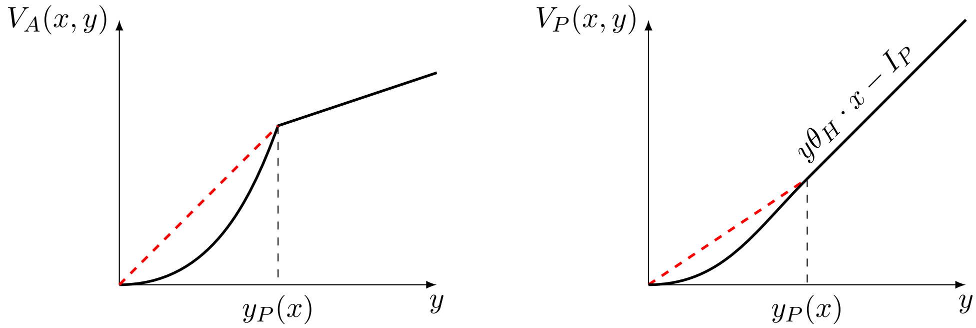


Figure 5: Effect of concavification on the agent's and the principal's value functions

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 $x_P(y)$ 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 an MPE of the dynamic game.[26]

## 5 Persuasion by an Informed Agent

So far we have assumed that the agent has no private information and learns about $\theta$ 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 willingness to pay 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 $\theta$ not only from the outcome of the test but also from the structure of the test itself. In what follows, we informally argue[27] that equilibrium outcomes of Propositions 1 and 4 can be supported as pure strategy equilibrium

[26] Persuasion using a pooling message implies that the value function of the principal has to satisfy both the "reflecting" (see Harrison (2013), ch.6) and smooth pasting (see Dixit and Pindyck (2012), ch.4) boundary conditions at the optimal exercise threshold. Since persuasion occurs over a strictly convex part of principal's value function, the slope implied by the reflecting boundary condition is always less than the smooth pasting slope $\theta_H \cdot x$. The two conditions can not be satisfied simultaneously.

[27] Formal treatment is relegated to Section A.2 of the Online Appendix in the interest of space.

outcomes of a dynamic persuasion model in which the agent knows $\theta$ but the principal does not.[28]

When the agent is privately informed about $\theta$ he can convey information to the principal either through conducting informative tests (Bayesian persuasion) or through his choice of such tests (signaling). Notice 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 $\theta$ 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 principal's beliefs to be *passive* if such property holds off path as well. That is, with passive beliefs the principal completely ignores the signaling component of the game. Passive beliefs, among many others, support the equilibrium outcome of Proposition 1.

**Corollary 2 (Persuasion to Wait).** *Suppose $I_P < I_A$ and let $(Y^*, \mathbb{T}^*)$ denote the equilibrium strategies from Proposition 1. 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 a Markov perfect equilibrium.*

Corollary 2 describes a pooling equilibrium in which the agent picks the pipetting belief process $Y^*$ from Proposition 1 regardless of his knowledge of $\theta$. 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 $b(x)$.

To see why this is an equilibrium notice 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 $x_P(0)$. Because $x_P(0)$ is the highest 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 $x < x_A(1)$ the belief process $Y^*$ maximizes $\mathrm{E}\left[e^{-r\tau}(\theta_H X_\tau - I_A) \,|\, \theta = \theta_H\right]$ given the principal's exercise region $\mathbb{T}^*$ as shown in (5) – (7). Thus, the high type has no incentives to deviate either.

[28] 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 4 can be supported in such framework as well.

Next, we show that the fully informative equilibrium outcome in Proposition 4 can also be sustained when the agent is privately informed.

**Corollary 3 (Persuasion to Act).** *Suppose $I_P > I_A$ and let $(Y^*, \mathbb{T}^*)$ denote the equilibrium strategies from Proposition 4. 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 a Markov perfect equilibrium.*

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.

# 6 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 towards late exercise, but may hurt him if he is biased towards 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 towards late exercise, the outcomes of the equilibrium and dynamic persuasion mechanism coincide if the conflict of interest is small. If the conflict of interest is large, lack of commitment is costly to the agent and leads to pipetting of information. On the other hand, when the agent is biased towards early exercise, the optimal mechanism prescribes pooling information at the principal's autarky threshold. Absent dynamic commitment, however, pooling of information unravels due to the option of the principal to wait and obtain more information from the agent in the immediate future. Agent's inability to stop persuading undermines him and there always exists an equilibrium in which the agent discloses all information at time zero.

Our paper is a step towards 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 gain from information disclosed in equilibrium when the agent's bias towards late exercise is large. Certain observable features of equilibrium disclosure, such as pipetting, i.e. many inconclusive tests, can be used as proxies for identifying potential problems.

## References


Pak Hung Au. Dynamic information disclosure. *The RAND Journal of Economics*, 46(4):791–823, 2015. ISSN 1756-2171.

Robert J. Aumann and Michael Maschler. *Repeated games with incomplete information.* MIT press, Cambridge University Press, 1995.

Jacopo Bizzotto, Jesper Rudiger, and Adrien Vigier. Dynamic bayesian persuasion with public news. *Working Paper*, 2016.

Robert K Dixit and Robert S Pindyck. *Investment under Uncertainty.* Princeton University Press, 2012.

Jeffrey C Ely. Beeps. *The American Economic Review*, 107(1):31–53, 2017.

Jeffrey C Ely and Martin Szydlowski. Moving the goalposts. *Working Paper*, 2017.

Steven R Grenadier and Neng Wang. Investment timing, agency, and information. *Journal of Financial Economics*, 75(3):493–533, 2005.

Steven R Grenadier, Andrey Malenko, and Nadya Malenko. Timing decisions in organizations: Communication and authority in a dynamic environment. *The American Economic Review*, 106 (9):2552–2581, 2016.

Sebastian Gryglewicz and Barney Hartman-Glaser. Dynamic agency and real options. *Working Paper*, 2018.

J Michael Harrison. *Brownian Models of Performance and Control.* Cambridge University Press, 2013.

Emeric Henry and Marco Ottaviani. Research and approval process: The organization of persuasion. *Working Paper*, 2018.

Johannes Hörner and Andrzej Skrzypacz. Selling information. *Journal of Political Economy*, 124 (6):1515–1562, 2016.

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

Thomas Kruse and Philipp Strack. Optimal stopping with private information. *Journal of Economic Theory*, 159:702–727, 2015.

Robert McDonald and Daniel Siegel. The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–727, 1986.

Roger B. Myerson. Multistage games with communication. *Econometrica*, 54(2):323–358, 1986.

David Pollard. *A user's guide to measure theoretic probability*, volume 8. Cambridge University Press, 2002.

Jérôme Renault, Eilon Solan, and Nicolas Vieille. Optimal dynamic information provision. *Games and Economic Behavior*, 104:329–349, 2017.

Alex Smolin. Dynamic evaluation design. *Working Paper*, 2018.

Nancy L Stokey. *The Economics of Inaction: Stochastic Control models with fixed costs*. Princeton University Press, 2008.

# Appendix

## Equilibrium Construction

**Case $x_P(0) \leq x_A(1)$.** The stopping boundary of the principal $a(x)$ and the pipetting boundary of the agent $b(x)$ coincide with the autarky boundary of the principal, i.e.

$$a(x) = b(x) = y_P(x) \quad \text{ for } x_P(1) \leq x \leq x_P(0).$$

**Case $x_P(0) > x_A(1)$ and $x^* \geq \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$.** The stopping boundary of the principal $a(x)$ and the pipetting boundary of the agent $b(x)$ are given by

$$a(x) = \begin{cases} y_P(x), & \text{if } x_P(1) \leq x \leq x^* \\ 1, & \text{if } x \notin [x_P(1), x^*] \end{cases} \qquad b(x) = \begin{cases} y_P(x), & \text{if } x_P(1) \leq x \leq x^* \\ a(x^*), & \text{if } x^* < x < x_A(1) \end{cases}$$

**Case $x_P(0) > x_A(1)$ and $x^* < \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$.** Put $\hat{x} = \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$. First, for every $x_R \in (0, \hat{x})$ we define an auxiliary function $S(x_R)$ as a local best response of the agent.

**Lemma 3.** *For any $x_R \in (0, \hat{x})$ there exists a unique $x_S = S(x_R) \in (\hat{x}, x_A(1))$ such that*

$$\underline{T}(x_R) \wedge T(x_S) = \arg\max_{\tau \in \mathcal{M}(x_R)} \mathrm{E}\left[e^{-r\tau}(\theta_H X_\tau - I_A)\right]$$

*where* $\mathcal{M}(x_R)$ *are all the stopping times* $\tau$ *w.r.t. process* $X$ *such that* $\tau \leq \underline{T}(x_R)$ *with* $\underline{T}(x) = \inf\{t > 0 : X_t \leq x\}$ *and* $T(x) = \inf\{t > 0 : X_t \geq x\}$.

The proof of this lemma is in the Online Appendix (see Lemma (A.3) )

We construct the right boundary of the principal's stopping test $x_R(y)$ as the unique solution of

$$x_R'(y) = \frac{\theta_H x_S - I_P - \dfrac{\theta_H x_R(1-\beta_2) + \beta_2 I_P}{\beta_1 - \beta_2}\left(\dfrac{x_S}{x_R}\right)^{\beta_1} - \dfrac{\theta_H x_R(\beta_1 - 1) - \beta_1 I_P}{\beta_1 - \beta_2}\left(\dfrac{x_S}{x_R}\right)^{\beta_2}}{\dfrac{[y\theta_H + (1-y)\theta_L](1-\beta_2)(1-\beta_1) - \beta_1\beta_2 I_P x_R^{-1}}{(\beta_1 - \beta_2)}\left[\left(\dfrac{x_S}{x_R}\right)^{\beta_1} - \left(\dfrac{x_S}{x_R}\right)^{\beta_2}\right](1-y)} \tag{10}$$

with $x_P(y^*) = y^*$, where $x_S = S(x_R(y))$ and $\beta_1 > 1 > 0 > \beta_2$ are the two solutions of $\frac{1}{2}\beta(\beta - 1) + \mu\beta = r$.

The ODE (10) defines an increasing function in $(x^*, \hat{x})$ because the numerator is negative due to $x_S > \hat{x} > x_R$ and the denominator is also negative due to $x_R \geq x^* = x_P(y^*) > x_P(y) > \frac{-\beta_2}{1-\beta_2} x_P(y)$. Denote by $\hat{y}$ the solution to the equation $x_R(\hat{y}) = \hat{x}$ and let $x_S(y) = S(x_R(y))$. Then the stopping boundary of the principal $a(x)$ and the pipetting boundary of the agent $b(x)$ are given by

$$a(x) = \begin{cases} y_P(x), & \text{if } x_P(1) \leq x \leq x^* \\ x_R^{-1}(x), & \text{if } x^* \leq x \leq \hat{x} \\ 1, & \text{if } x \notin [x_P(1), \hat{x}] \end{cases} \qquad b(x) = \begin{cases} y_P(x), & \text{if } x_P(1) \leq x \leq x^* \\ x_R^{-1}(x), & \text{if } x^* \leq x \leq \hat{x} \\ x_S^{-1}(x), & \text{if } \hat{x} \leq x \leq x_S(y^*) \\ y^*, & \text{if } x_S(y^*) \leq x < x_A(1) \end{cases}$$

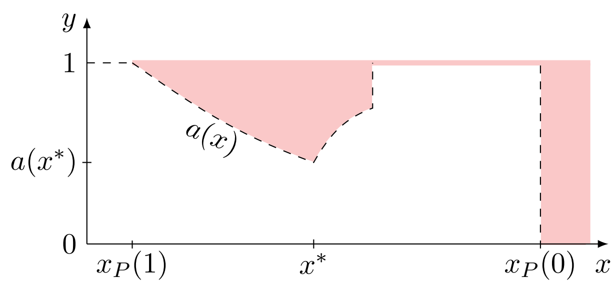

(a) Stopping set of the principal.

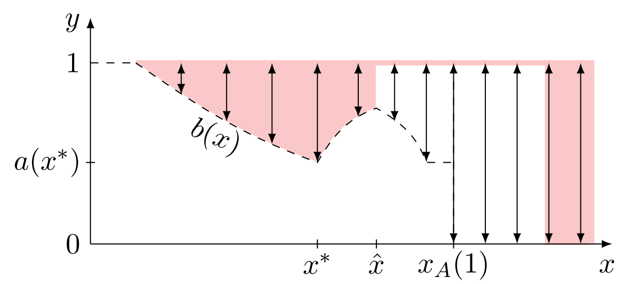

(b) Communication strategy of the agent.

Figure 6: Equilibrium strategies with $x^* < \frac{r}{r-\mu} \cdot \frac{I_A}{\theta_H}$ and $x_P(0) > x_A(1)$.

In the waiting region $x \in (x_R(y), x_S(y))$ equilibrium value function $V_P(x, y)$ of the principal satisfies

$$rV_P(x,y) = \mu x \frac{\partial}{\partial x} V_P(x,y) + \frac{1}{2}\phi^2 x^2 \frac{\partial^2}{\partial x^2} V_P(x,y)$$

with boundary conditions

$$\begin{aligned} V_P(x_R(y),y) &= (y\theta_H + (1-y)\theta_L)x_R(y) - I_P \\ \frac{\partial}{\partial x} V_P(x_R(y),y) &= y\theta_H + (1-y)\theta_L \\ \frac{\partial}{\partial y} V_P(x_S(y),y) &= \frac{\theta_H x_S(y) - I_P - V_P(x_S(y),y)}{1-y}. \end{aligned}$$

It is the last, reflecting, boundary condition at $x_S(y)$ that gives rise to the ODE (10).

## Proof of Proposition 1.

Below we formally prove the equilibrium existence for the case $x^* \geq \hat{x}$ and $x_P(0) > x_A(1)$. This also covers the case $x_P(0) \leq x_A(1)$ since the equilibrium features only region $R_3$. Formal verification of equilibrium existence for the case $x^* < \hat{x}$ and $x_P(0) > x_A(1)$ and proof of uniqueness in all three cases is relegated to Online Appendix.

*Principal's Best Response.* The only part of the principle's best response that is not obvious is the optimality of the stopping boundary $a(x) = y_P(x)$ in presence of pipetting. Below we formally show that pipetting against $b(x) = y_P(x)$ does not add value for the principal. We do so by conjecturing that $V_P(x,y) = [y\theta_H + (1-y)\theta_L]\, x - I_P$ for $y > a(x)$ and then check that 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), Chapter 9). Given the conjecture about the principal's value above $a(x)$ we get

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

Below $y = a(x)$ function $V_P$ is given by the autarky solution, i.e., it solves a waiting ODE with value matching and smooth-pasting condition at $y_P(x)$. Differentiate the value matching condition at $y_P(x)$ to get:

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

Hence the autarky value function of the principal satisfies the waiting ODE together with the boundary conditions of the equilibrium value function with reflected belief process at the boundary

$y_P(x)$. Thus, the equilibrium value function of the principal is her autarky value and pipetting alone does not add value.

*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.

**Lemma 4.** *If the agent's stopping set is monotone in beliefs, i.e. 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 best response $\tilde{Y}$ under which he stochastically reveals $\theta = \theta_H$ with the same outcome.*

*Proof.* Define $D_t$ to be the cumulative probability of stopping (prior to reaching $x_A(1)$) under the belief process $Y$ conditional on the path of $X$, i.e. $D_t = \mathrm{P}\big(Y_t = 1 \,|\, \theta = \theta_H, \mathcal{F}_t^X\big)$ 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 (1).

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

$$\mathrm{P}\big(Y_t = 1 \,|\, \theta = \theta_H, \mathcal{F}_t^X\big) = \mathrm{P}\Big(\tilde{Y}_t = 1 \,|\, \theta = \theta_H, \mathcal{F}_t^X\Big),$$

and $Y_t^{ND} = \mathrm{E}\left[Y_t \,|\, Y_t < 1, \mathcal{F}_t^X\right]$. Hence whenever $Y_t \leq a(X_t)$ we have $Y_t^{ND} \leq a(X_t)$, and the strategy $\tilde{Y}$ generates the same payoffs and distribution of outcomes as $Y$. □

The remainder of the proof relies on constructing an upper bound for (7). In region $R_3$ we have $\eta(y) \leq T(x_P(y))$ for $y > a(x^*)$ and thus

$$\mathrm{E}_x\left[e^{-r\eta(y)}\big(\theta_H X_{\eta(y)} - I_A\big)\right] \leq \sup_\tau \mathrm{E}_x\left[e^{-r\cdot\tau\wedge T(x_P(y))}\big(\theta_H X_{\tau\wedge T(x_P(y))} - I_A\big)\right]$$

where supremum is taken over all stopping times $\tau$. Define $f(x) = (x/x_P(y))^{\beta_1}(\theta_H x_P(y) - I_A)$ and notice that $f(x) > \theta_H x - I_A$ for all $x < x_P(y)$. Then

$$\mathrm{E}_x\left[e^{-r\cdot\tau\wedge T(x_P(y))}\big(\theta_H X_{\tau\wedge T(x_P(y))} - I_A\big)\right] \leq \mathrm{E}_x\left[e^{-r\cdot\tau\wedge T(x_P(y))}\cdot f\big(X_{\tau\wedge T(x_P(y))}\big)\right] \leq f(x)$$

where the last inequality holds because the process $e^{-r\cdot t\wedge T(x_P(y))}\cdot f\big(X_{t\wedge T(x_P(y))}\big)$ is a positive super-martingale. Finally notice that $f(x)$ is the expected payoff when $\eta(y) = T(x_P(y))$, hence, the upper bound is achieved by exercising the option at the boundary $x_P(y) = a^{-1}(y)$.

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

$$\mathrm{E}_x\left[e^{-r\eta(y)}\big(\theta_H X_{\eta(y)} - I_A\big)\right] \leq \sup_\tau \mathrm{E}_x\left[e^{-r\cdot\tau\wedge\underline{T}(x^*)}\big(\theta_H X_{\tau\wedge\underline{T}(x^*)} - I_A\big)\right].$$

When $x^* \geq \hat{x}$ the process $e^{-r\cdot t\wedge\underline{T}(x^*)}\big(\theta_H X_{t\wedge\underline{T}(x^*)} - I_A\big)$ is a positive super-martingale, hence immediate stopping $\eta(y) = 0$ is optimal. □