

Persuading the Principal To Wait*

Dmitry Orlov[†] Andrzej Skrzypacz[‡] Pavel Zryumov[§]

First Draft: 13th February, 2016

Current Version: 9th March, 2017

Abstract

We study strategic communication of verifiable information in a dynamic setting. The principal chooses when to exercise a real option. The agent affects the timing of option exercise by collecting and revealing additional information to the principal over time. When the agent favors late exercise relative to the principal and the conflict of interest is small, all information is disclosed discretely but with a delay. The principal exercises the option at the agent's preferred threshold. When the conflict is large, the agent "pipets" information over time. Credible threat of option exercise is crucial for incentives to provide new information. When the agent favors early exercise, his lack of commitment not to persuade in the future leads to unraveling; in equilibrium all information is disclosed immediately. In sharp contrast to static environments, the ability to persuade might hurt the agent.

1 Introduction

Few decisions are made without doubt. Decision makers commonly rely on agents to provide relevant information. In turn, agents use strategic communication of verifiable information to influence the decision makers. For example, firms choose when to collect and disclose hard information to the regulators in an effort to keep their products in the market. Managers decide on what evidence to acquire to convince headquarters to launch a product. Conflicts of interest arise when preferences over decisions are not aligned. We model strategic communication as a game of dynamic persuasion of the principal by the agent in the context of real options. We characterize equilibrium decisions, dynamics of information provision, implications for value of authority and access to information.

*We thank Mike Harrison, Johannes Hörner, Emir Kamenica, Ron Kaniel, Eric Madsen, Heikki Rantakari and seminar participants at the University of Rochester, University of Pennsylvania, University of Chicago, Chicago Booth, Wharton School, 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.

[†]University of Rochester, Simon School of Business. dmitry.orlov@simon.rochester.edu

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

[§]University of Pennsylvania, Wharton School. pzryumov@wharton.upenn.edu

In a real option problem, a principal decides on the optimal timing of a certain action, i.e., when to exercise an option. Innovations, that change the expected payoffs of the action, arrive over time from multiple sources. In a classic¹ real-option model innovations arrive exogenously. We extend the standard model by adding a strategic agent who can acquire (and disclose) additional verifiable² information about the expected payoffs. Although the agent has superior access to information, his incentives are not fully aligned with the principal's: he may prefer to either wait longer or act sooner than the principal. The agent chooses strategically when and what to learn in order to influence the principals decision, i.e., to persuade him.

One example of such setting is the post market surveillance of drugs and medical devices conducted by FDA. When the product is first introduced to the market, its efficacy is somewhat uncertain. As more and more patients use it, information about its effects, the good ones and the bad ones, is gradually revealed. The regulator monitors this (exogenous) news and can recall the product (i.e. remove it from the market) if it turns out to be unsafe. The firm producing the drug can affect FDA's decision by providing additional information/tests³. It is natural to expect a partial misalignment between the firm and the regulator over when to exercise the real option of recall. For example, 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 the tests⁴ and optimally timing them. Similar strategic communication takes place between firms and their regulators in many other industries (e.g. environmental regulation and firms obtaining and disclosing results on environmental impact that go beyond the required minimum testing) and countries.

A more stylized example from organizational economics is the problem of a new product launch or of canceling an existing product. Senior executives in the company have the ultimate control over the timing of such decisions, which are made based on exogenous innovations, e.g., estimates of the market size, and additional information provided by the product manager. While the manager can design and run additional tests about, for example, the willingness to pay of potential customers, he might be biased toward an earlier product launch or favor later product shutdown⁵. We assume that the product manager cannot pay the executives to change the timing of the decision, instead, he can influence the decision making of the executives by strategically timing information acquisition.

Several other applications fit our theory: it can be applied to decisions when to sell off firm's assets when some stakeholders have access to sources of information while others have the majority required to make the decision to sell assets. The common theme across these examples are that both the principal and the

¹See Stokey (2008), Dixit and Pindyck (2012), McDonald and Siegel (1986).

²As opposed to cheap talk of Grenadier, Malenko, and Malenko (2016).

³For example, 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.

⁴The 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."

⁵We are agnostic about the contracts present between the principal and the agent as long as they do not perfectly resolve the conflict of interest. We study the equilibrium information provision under both the assumptions: that the agent would like to persuade the principal to wait or to act sooner.

agent are rational, forward-looking and cannot commit to future actions. Transfers are either not allowed or do not fully resolve the conflict of interest⁶. Agents can strategically decide not to learn certain facts, but anything they decide to learn they have to disclose⁷.

Under these assumptions we ask when and what information is obtained in equilibrium, who benefits from agent's access to information, and how successful is dynamic persuasion. The answers to these questions turn out to depend on the direction and strength of the conflict. In the remainder of the paper we refer to information, or news, being good (bad) when they increase (decrease) the payoff of the option upon exercise.

Persuasion to Wait. In the first part of the paper we analyze the case where the agent is “more patient” than the principal, in the sense that absent any information from the agent, the principal would exercise the real option sooner than the agent. So, it is in the agent's interest to persuade the principal to wait. We show this game admits a unique Markov-perfect equilibrium in which the agent delays information disclosure and the principal sometimes regrets waiting too long ex-post. We show that the pattern of information disclosure is determined by the strength of the conflict of interest between the principal and the agent, which in turn depends on the size of the agent's bias and the value of information that the agent could provide.

When the conflict is small, i.e., either the agent's bias is small or the additional information is sufficiently valuable, then in equilibrium the agent postpones information acquisition until a point when he acquires all available information at once. If this information turns out to be good, the principal acts immediately. If the information turns out to be bad, the principal continues to wait. The principal waits past her autarky threshold (the state at which she would act if there was no agent in the game) because with a small conflict of interest, the cost of additional waiting is small compared to the value of information obtained later. Even though ex-post the principal regrets waiting when the information is good, she is happy to have waited if the news is bad.

The timing of the action has a “compromise” property: in case the results are good, the principal acts at the agent's optimal point; when the results are bad, the principal further delays but if she eventually does act, it is at her optimal threshold. Equilibrium outcome can be implemented via a time-contingent delegation policy that initially delegates all authority to the agent, however, if the agent continues to wait past a certain threshold, the principal reclaims the authority and exercises the option herself. Finally, in this case both parties benefit from the agent's access to information and the equilibrium outcome would be the same if information was soft (i.e. cheap talk) instead of hard (i.e. persuasion or disclosure).

⁶We consider transfers in form of a bonus to the agent for exercising the real option. We show that, while the principal could perfectly align preferences by choosing the appropriate bonus, he optimally chooses to set the bonus to zero.

⁷We do not allow the firm/agent to secretly run tests and hide unfavorable results and we assume that the information is verifiable since it comes from physical tests. While voluntary non-disclosure is legal in some markets, we focus on markets in which either full disclosure is mandated or the agent cannot hide having information (if the executive knows that the manager has run a test then due to standard unraveling arguments the agent would reveal its results). The analysis of fraud related to disclosure or fabrication of results is of its own interest, especially given the famous examples of firms knowingly hiding negative effects of their products. While our paper can help understand the incentives to engage in such fraud, a proper analysis of fraud and fraud prevention is beyond the scope of this paper.

When the conflict of interest is large, equilibrium information provision is radically different. We show that instead of acquiring information once and fully, the agent “pipets” bad information. That is, once the principal reaches her action threshold, which in equilibrium is the same as if the principal acted on the exogenous innovation only, the agent reveals just enough information to either slightly decrease the belief or to fully reveal good information. In the jargon of real options, the principal’s belief threshold becomes a reflecting boundary (and the reflection is accompanied by a jump of beliefs when good news are revealed, so that beliefs are martingales)⁸. Authority of the principal and a credible threat of option exercise is crucial for incentivizing the agent to provide information, rendering meaningful delegation strictly suboptimal. Somewhat surprisingly, while pipetting bad information helps the agent to obtain a higher payoff than with no access to additional information, for the principal additional information provided by the agent provides just enough value to compensate for the extra cost of waiting.

In Section 6 we consider a version of the model where the agent is privately informed, but can still conduct credible tests. In this game of asymmetric information the agent has private information about the outcomes of such tests, relative to the principal. We show that the equilibrium information sharing of the Bayesian persuasion game remains an equilibrium in this alternative setting.

Persuasion to Act. We also explore the case where the agent is less patient than the principal, in the sense that his threshold for exercising the real option is lower than the principal’s. We call this case *persuasion to act*. We show that there always exists a Markov-perfect equilibrium in which the agent immediately acquires and shares all information, moreover, under certain conditions, it is unique. The intuition behind such stark result is similar to the one of Coase conjecture: given an option exercise threshold of the principal, the agent provides additional information in hope of changing principal’s instantaneous action and speeding up the option exercise time. However, information provided by the agent today increases principal’s value of waiting for more information in the future. As a result, the action threshold of the principal becomes more stringent, which, in turn, requires the agent to provide more precise information resulting in even more stringent threshold and so on. While in some cases acquiring and sharing all information at the initial date is beneficial for the agent who wants to speed-up the execution of the real option, it can be also harmful to him. For example, if absent additional information from the agent the principal were to act, but the bad news induce excessive waiting, then the agent is hurt by having access to information. This result is in stark contrast with persuasion to wait, where having access to information is always beneficial for the agent.

Two Models of Real Options. We cast our model as the classical Dixit and Pindyck (2012) real-option problem. The profitability of the action depends on two uncorrelated states. One of the states is public and

⁸If the misalignment is large enough, this is the only type of information disclosure on the equilibrium path. When misalignment is intermediate, the equilibrium starts in the “pipetting” mode and, if good news do not arrive, eventually changes to the “waiting for full information” mode.

evolves exogenously, and the other one is unobserved by either party directly but can be learned through the tests conducted by the agent.

In the end of the paper we discuss an alternative Wald option model, in which there is a single unobserved underlying state of nature that affects payoffs. The players observe a exogenous news process that provides public information about the state over time.⁹ In addition, the agent can collect and share information that is imperfectly correlated with underlying state. Such extension allows us to better capture the firm-regulator applications: the product is either safe or not safe, the market learns about the safety from the reports of consumer experiences (the exogenous news) and the firm can provide additional tests that are informative but not 100% conclusive (for example, because safety may depend on long-term effects that are hard to test for)¹⁰.

In this model the Markov state of the game is the joint belief over the fundamental state and the additional information that the agent could learn. There are several differences between those models; for example, the latter model captures the possibility that as consumer reports reach the regulator, its beliefs about what the firm may disclose change. Also, in the Dixit and Pindyck (2012) model the exogenous state can have a positive drift, while in the Wald model it is a martingale (since the state is a belief). Despite these differences, we show that all qualitative results continue to hold for the Wald option.

On the technical side, we have decided to write the model in continuous time to use the standard tools of real-option analysis. The drawback of this approach is that we have a game between two strategic agents instead of a single-decision maker, and the actions are perfectly observable (unlike in Sannikov (2007)). In general, modeling such games in continuous time is notoriously difficult (see for example, Simon and Stinchcombe (1989)). In discrete-time communications games, it is customary to define the agent's strategy as a choice of message sequences from some abstract message space as a function of the history. This can be simplified in Bayesian persuasion problems since it is without loss of generality to describe communication by posterior beliefs induced by messages and the set of all possible communications as the set of all distributions that satisfy the martingale property of beliefs. Even that turns out to be somewhat technically cumbersome in continuous time. Instead, we define the agent's strategy to be a choice of information partition about the additional signal, and we put constraints on the information partitions to be a well-defined filtration. Such filtrations capture a general way the agent can in continuous time provide information in response to the public news. While continuous time implies that for some Markov strategies of the principal (mappings from the payoff-relevant state to the decision to act or not) there does not exist an agent's best response (because the supremum over responses is not attainable), the problem does not lead to non-existence of Markov equilibria.

⁹This learning process is explored in Wald (1973) and the corresponding real option model is studied in Stokey (2008).

¹⁰In the special case when the firm is able to acquire information that perfectly reveals the state, equilibrium features either only pipetting, or full disclosure at time 0. Signal imprecision creates transition dynamics, similar to those in Dixit and Pindyck (2012) and Stokey (2008) option models.

1.1 Related Literature.

This work is related to papers on dynamic decision making with communication such as Grenadier, Malenko, and Malenko (2016) and Guo (2016). We contribute to this literature by considering effects of managing verifiable information for dynamic decision making. Grenadier, Malenko, and Malenko (2016) show that when an informed agent has a bias towards delayed exercise and that bias is not very large, the principal achieves his best outcome by delegating the execution of the real option to the agent. When the bias is large, the principal exercises the option at the uninformed threshold. We show that when the agent strategically manages verifiable information, there is the middle scenario: when the conflict is large, there is slow release of information in the beginning with the principal retaining authority. The conflict eventually diminishes and the principal delegates option exercise to the agent.¹¹ If the agent prefers early exercise, Grenadier, Malenko, and Malenko (2016) show that there does not exist a revealing equilibrium. This is in contrast to our findings, which show that if the agent is impatient, in many cases it leads to full and immediate information sharing. These results highlight the distinctions between considering soft (unverifiable) versus hard (verifiable) information communication.

Guo (2016) studies optimal delegation in a dynamic experimentation setting. When the agent has a preference for more experimentation relative to the principal, the optimal policy is a sliding deadline: principal commits to terminate the agent unless success has been observed. We show in this paper that in order to incentivize the principal to delay exercise of the real option (akin to stopping experimentation) the agent produce negative information about the value of the real option. When the agent has a preference to stop experimentation sooner, Guo (2016) finds it optimal for the principal to set a floor on experimentation a solution which requires dynamic commitment from the principal. We show that when the agent wants to stop sooner than the principal, the equilibrium features full and immediate information revelation. This is the first best outcome for the principal and does not require dynamic commitment on her behalf.

More broadly, our paper is applicable to vertical authority and communication in organizations. An excellent overview of the literature can be found in Bolton and Dewatripont (2013).

We model communication as management of public information by the agent, also known as Bayesian persuasion, first introduced by Aumann and Maschler (1995) and Kamenica and Gentzkow (2011). The agent does not have private information, chooses what information to acquire (how to “split” a prior belief), and cannot hide results. There are a few papers that study dynamic Bayesian persuasion, for example, Renault, Solan, and Vieille (2014) and Ely (2017). There are many differences between these two papers and our model. First, in their models the agent cares only about the principal’s actions, while in our paper the agent’s preferences over actions depend on the underlying state: if the state is good, the agent wants to exercise the option, while if the state is bad, he prefers to wait. The second is that we have two long-lived players and consider non-commitment equilibria rather than design of dynamic persuasion with long-term

¹¹This is exactly true in our model when $\theta_L \leq 0$.

commitment. Despite these important modeling differences, somewhat surprisingly we find that in some instances the equilibrium information provision features “pipetting”. Such slow revelation of information is the continuous time equivalent of the “greedy” strategy from Renault, Solan, and Vieille (2014): releasing the least information possible to delay the exercise of the real option. While pipetting is sometimes an equilibrium outcome, we also identify important cases when information is revealed only once and fully.

A closely related paper is Bizzotto, Rudiger, and Vigier (2016). It studies dynamic persuasion in the presence of public news and the principal facing a single stopping decision. The main focus of this paper is on the effect of a deadline: authors show that if public news was stationary over time, then the unique Perfect Bayesian Equilibrium in their model involves the agent persuading at $t = 0$ and the principal acting. Our analysis applies more broadly to real option problems (see Sections 2, 4, and 5) and Bayesian persuasion in the presence of public news (see Section 7). We show that when agent’s information is imperfect, then delays occur when news arrives continuously over time. Moreover, in Section 7 we show that in continuous time perfect unraveling occurs in which the agent reveals a perfectly informative signal about the state.

Our paper is related to literature on voluntary dynamic information disclosure with verifiable information, as in Acharya, DeMarzo, and Kremer (2011), Guttman, Kremer, and Skrzypacz (2014), and Hörner and Skrzypacz (2016). The main part of this paper focuses on managing public information and thus does not allow the firm to hide the information it obtains. To make sure our findings are robust to alternative specifications of communication of hard information, we extend the model to allow for the agent to be privately informed. In this setting he still shares verifiable information by conducting credible tests, which allows him to reveal private information in small pieces. While such a game can be rewritten as a game of dynamic disclosure, we, instead, extend the game-theoretic framework developed in Section 2 to analyze the expanded game. We show that equilibria constructed in Sections 4 and 5 remain equilibria when the agent has private information.

The observation that in some of our equilibria the regulator waits for the additional news to arrive is somewhat reminiscent of the results in Kremer and Skrzypacz (2007) and Daley and Green (2012) where the expectation of news arriving to the market leads to a market breakdown due to adverse selection. In our model news cause the agent to delay communication. Unlike Daley and Green (2012), the agent never plays a mixed strategy and the randomness is embedded in the optimal message structure.

Our paper is related to literature on agency conflicts in real options such as Grenadier and Wang (2005), Gryglewicz and Hartman-Glaser (2017), and Kruse and Strack (2015). These papers study delegation in the presence of incentive contracts. In our model, real option decision rights belong to the principal and we study the effect of information choice by the agent in persuading the principal. We begin by not allowing incentive contracts between the principal and the agent. Our findings are robust to the introduction of simple contracts: if the principal could offer the agent a positive bonus for exercising the option, he would prefer to set it to zero.

The rest of the paper is organized as follows. Section 2 we present our main model. In Section 3 we provide some additional notation and preliminary results. In Sections 4 and 5 we discuss persuasion to wait and persuasion to act respectively. In Section 6 we allow the agent to be privately informed and show that equilibria constructed in Sections 4 and 5 are robust to introduction of private information. In Section 7 we present an alternative specification of real options and discuss the role of commitment in Section 8. We conclude in Section 9. The supplementary Appendix contains proofs of all our formal results.

2 Model

2.1 Basic Setup

We start with an informal description of the model. There are two long-lived players, a principal (she) and an agent (he). The principal faces a real-option problem: she chooses the time to make an irreversible (at least partially) decision. There are two reasons why the principal may wish to wait: exogenous innovations in the underlying state; and new information about the project that the agent endogenously decides when to acquire. A conflict of interest arises because the players share differently the costs and benefits of the decision. Once the agent learns some information he has to disclose it, so at all times the players have the same beliefs about the unknown state. The agent strategizes over what kind of additional information to gather and when to do it in order to persuade the principal. Neither player can commit to future actions; both are strategic and forward looking. We are interested in understanding when and what kind of information is shared in equilibrium. We are also interested in whether the ability to persuade hurts or benefits the agent.

There are many economic situations that fit this broad description. Consider a firm deciding when to launch a new product. An executive in the firm (the principal) has the final decision power. She assesses the size of the market and the consumers' willingness to pay. Over time public news arrive about the size of the market (we model the evolution of the market size as a geometric Brownian motion). Consumers' willingness to pay (WTP), high or low, is unobservable and both players share a common prior over it. The agent is a product manager. He strategically designs marketing studies that are informative about the WTP (e.g., how informative they are) and chooses when to conduct those. His preferences are not fully aligned with the principal's (because he gets private benefits from managing the product or because he does not fully internalize the cost of launching, etc.). He understands that good news speed up the launch but bad news postpone it and cannot be hidden once acquired.

A different example captured by our model is that of information sharing between producers and regulators who have the power to recall products. For example, the FDA conducts a Post market Surveillance program that collects information from patients and health care providers about potential side effects of drugs and medical devices. The FDA also allows manufacturers to submit results of additional experiments. The

regulators and the firms are likely to have different preferences over the optimal timing of recalls in response to the news.¹²

Players and Payoffs. We proceed to define the model formally. There are two long-lived players: a principal (she) and an agent (he). 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 the observable state of the market X , that evolves over time, and the underlying quality of the project, $\theta \in \{\theta_H, \theta_L\}$, with $\theta_H > \max(\theta_L, 0)$. Neither party observes the realization of θ , they share a common prior

$$Y_{0-} = P(\theta = \theta_H). \quad (1)$$

We model the real option problem using the classic approach described in Dixit and Pindyck (2012), assuming that X and θ are independent.¹³ We discuss a different approach, where X is a news process about the payoff-relevant θ in Section 7 (the Wald option model, that may better describe some applications), and argue that the results do not change.

Time is continuous and infinite, $t \in [0, +\infty)$, both players discount future payoffs at a rate $r > 0$. If the option is exercised at time t then time 0 discounted payoffs of the agent and the principal conditional on θ are

$$v_A(t, \theta) = e^{-rt} (\theta X_t - I_A), \quad v_P(t, \theta) = e^{-rt} (\theta X_t - I_P), \quad (2)$$

where the publicly observed process $X = (X_t)_{t \geq 0}$, follows a geometric Brownian motion

$$dX_t = \mu X_t dt + \phi X_t dB_t \quad (3)$$

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

The disagreement between the agent and the principal is driven by differential cost of exercising the option $I_P \neq I_A$. Intuitively, if $I_P < I_A$, then for any given θ the agent's optimal timing of exercise is later than the principal's. In this case the agent would like the principal to delay exercise time – he would like to persuade the principal to wait. If $I_P > I_A$ then direction of the conflict is reversed and the agent would like to accelerate exercise time – he would like to persuade the principal to act.

¹²We propose to model these two applications slightly differently. We start with the product launch application and return to the regulator application in Section 7. The main difference is that in the product case we assume that the size of the market and the willingness to pay are independent. In the regulator recall case it is more natural to assume that the news from customers is correlated with the news that the firm can obtain by running additional tests.

¹³Grenadier, Malenko, and Malenko (2016) use identical payoff structure, however, in their model the agent is privately informed about the realization of θ , and can send only unverifiable messages, while in our model there is no private information, the agent decides when and what to publicly learn about θ , and the information is hard.

Remark about payoffs. Although we model preferences in terms of the call option with zero flow costs and only terminal payoffs, an equivalent formulation of our model is to consider a put option to stop, e.g.:

$$\tilde{v}_A(t, \theta) = \int_0^t e^{-rs}(\theta X_s - I_A)ds, \quad \tilde{v}_P(t, \theta) = \int_0^t e^{-rs}(\theta X_s - I_P)ds. \quad (4)$$

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 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 a conflict of interest between a principal who has decision rights over timing of action and an agent who has access to additional information. The details of the real option do not affect our results.

2.2 Strategies

We cast our model in continuous time to use well-established tools and intuitions for single-agent real option problems (see for example Dixit and Pindyck (2012)). However, since our problem is a game and not a single decision-maker problem, continuous time requires special care in defining strategies. While this subsection is highly technical, it is designed to capture the following idea from discrete time: within every “period” the innovation to X_t is first realized, the agent can then provide an informative signal about θ by “splitting” the prior about θ . That is, the agent induces a posterior over θ from some distribution subject to the martingale constraint that the average posterior belief has to be equal to the prior. The agent can commit within a period to an arbitrary distribution, but he cannot commit to future actions. After that, the principal decides whether to exercise the option or continue waiting (for example, whether or not to launch the product), she also cannot commit to future actions. The agent’s strategy is a function of past history of the state process X and information revealed about θ up to time t . The principal’s strategy is additionally a function of the information disclosed about θ at time t . Markov strategies depend on the history only insofar that it contains information about current level of X and current belief about θ .

Sources of Information. Information in this game comes from three sources:

1. *Exogenous Evolution of X* : $\mathbb{F}^X = (\mathcal{F}_t^X)_{t \geq 0}$ with $\mathcal{F}_t^X = \sigma(X_s, s \leq t)$
2. *Quality of the Project*: $\mathcal{F}^\theta = \sigma(\theta)$
3. *Randomization Device*: A sufficiently rich¹⁴ sigma-algebra \mathcal{F}^R

¹⁴For our purposes it will be sufficient to require that this sigma-algebra contains sigma-algebras generated by two independent uniform random variables. In the construction of equilibria, we first assume that this sigma-algebra is sufficiently rich and then show by construction what is sufficient.

such that $\mathcal{F}_\infty^X = \bigvee_{t=0}^\infty \mathcal{F}_t^X$, \mathcal{F}^R , and \mathcal{F}^θ are mutually independent sigma algebras¹⁵. For technical reasons we require filtration \mathbb{F}^X to satisfy standard conditions, i.e., to contain all P-null sets and to be right-continuous.

Agent's strategy. We start by defining agent's action profile.

Definition. A feasible action profile of the agent is a filtration $\mathbb{H} = (\mathcal{H}_t)_{t \geq 0}$ satisfying standard conditions such that¹⁶

$$\forall t \quad \mathcal{F}_t^X \subseteq \mathcal{H}_t \subseteq \sigma(\mathcal{F}_t^X, \mathcal{F}^\theta, \mathcal{F}^R) \quad \text{and} \quad \mathcal{H}_{0-} \subseteq \mathcal{H}_0. \quad (5)$$

Let \mathcal{H} denote the set of all feasible action profiles.

Such definition of action profiles allows the agent to generate informative messages at every time t whose distribution can be contingent on the path of X , realization of θ , and realization of past messages. Standard restrictions on the filtration \mathbb{H} guarantee that all martingales with respect to this filtration have càdlàg versions and first hitting times of such martingales are stopping times with respect to \mathbb{H} .¹⁷ That assures the expected payoffs from our strategies are well-defined. The set of all action profiles \mathcal{H} corresponds to the set of all possible histories of the game in case the principal never stopped.¹⁸

According to this definition an action profile specifies not just the current information disclosure but a whole plan of contingent dynamic information disclosures starting at time 0. Intuitively, we are capturing an action in a normal form of the dynamic game. To introduce a notion of perfection, we next define histories and strategies.

Define the *history* generated by a feasible action profile \mathbb{H} up to time t to be

$$\mathbb{H}^t = (\mathcal{H}_{0-}, \mathcal{H}_s; s < t).$$

Notice that such definition of history in a way encompasses all paths that could happen up to time t under the action profile \mathbb{H} . Thus, any decisions made conditional on history up to time t would have to specify the plan of action for every possible realization of stochastic uncertainty.

Denote by $\mathcal{H}(\mathbb{H}^t)$ the set of all feasible action profiles that agree with \mathbb{H} up to time t , i.e.

$$\mathcal{H}(\mathbb{H}^t) = \left\{ \tilde{\mathbb{H}} \in \mathcal{H} : \tilde{\mathcal{H}}_s = \mathcal{H}_s \quad \forall s < t \right\}. \quad (6)$$

¹⁵We also require the original probability space to be rich enough to accommodate process X and independent randomization, i.e. $\mathcal{F}^R \subset \mathcal{F}$ and $\mathcal{F}_\infty^X \subset \mathcal{F}$.

¹⁶We formally put $\mathcal{H}_{0-} = \sigma(X_0, Y_{0-}, \cdot)$, i.e. sigma algebra \mathcal{H}_{0-} contains information only about starting point of the game: the prior Y_{0-} and the level of the state process X_0 before agent generated news are revealed.

¹⁷For further details see Pollard (2002).

¹⁸We chose \mathcal{H} to denote the set of all action profiles motivated by Mailath and Samuelson (2006) who use that symbol to denote the set of all histories of a repeated game

Definition. A strategy of the agent, S , is a mapping from any possible history \mathbb{H}^t into the set of all feasible action profiles $\mathcal{H}(\mathbb{H}^t)$ that agree with \mathbb{H}^t up to time t , i.e.

$$S(\mathbb{H}^t) \in \mathcal{H}(\mathbb{H}^t). \quad (7)$$

At any time and after any history of the game, the agent's strategy S is a choice of an action profile, i.e. it specifies the whole structure of past and future information sharing, $S(\mathbb{H}^t)$. Since the history states what already happened and an action profile specifies information sharing starting at time 0, the strategy at time t can only choose actions that are consistent with the realized history.

In discrete-time persuasion games it is more common to define a strategy as a mapping from all possible histories to current period information disclosure (Ely (2017), Renault, Solan, and Vieille (2014)). Iteratively one can then recover the whole sequence of actions in future periods and specify the information filtration chosen by the agent. We define the strategy to be the whole contingent plan of disclosures in current and future periods. Such definition allows us to avoid some technical issues related to continuous time modeling.

A strategy of the seller is *time consistent* if

$$S(\mathbb{H}^{t'}) = S\left(S(\mathbb{H}^t)^{t'}\right) \quad \forall t' > t \geq 0 \text{ and } \forall \mathbb{H} \in \mathcal{H}.$$

If the agent decides on the whole structure of information sharing at time t after arbitrary history \mathbb{H}^t and follows it up to some future time t' , then at t' (for a strategy to be time consistent) he should not change his mind given the history $S(\mathbb{H}^t)^{t'} = \{S_s(\mathbb{H}^t), s < t'\}$. We do not restrict the agent to time consistent strategies, allowing him to deviate to non-time consistent strategy after any time (time consistency will be a feature of the equilibrium due to sequential rationality).

Remark. Our model can be described as a dynamic Bayesian persuasion model *without* commitment. In Bayesian persuasion models, agent's strategy is typically defined either as a choice of messages, or, more commonly, as a choice of posterior distribution of beliefs subject to a martingale constraint. In our definition, instead of choosing posteriors, the agent chooses filtration/sigma-algebras of the set Ω and they induce posterior beliefs¹⁹, so that in discrete time the definitions are equivalent. An additional benefit of defining actions in terms of filtrations is that it allows us to easily incorporate constraints like (5). Also recall that we assume that the filtration chosen by the Sender up to time t is public, so that at any time the agent and principal's beliefs coincide both on and off equilibrium path. In our examples, this corresponds to the firm having the option not to perform additional tests about safety of its product, but if it learns

¹⁹One can, of course, recover the filtration given the belief process as the natural filtration of this process. All other filtrations that give rise to the same belief process can only include additional information that is independent of the θ conditional on the belief process itself.

anything, it is legally obliged to disclose it to the regulator. Similarly, a product manager can decide not to collect some additional information, but he is compelled to reveal any news to his boss whenever he learns anything.²⁰

We define a (*Markov*) *state* of the game, π , to be the pair $\pi_t = (X_t, Y_t)$. That is, π_t contains the posterior belief Y_t about the quality of the option θ and the current level of state process X_t . Denote by Π the set of all possible pairs (x, y) , i.e., $\Pi = [0, +\infty) \times [0, 1]$.

For any feasible action profile of the agent $\mathbb{H} = (\mathcal{H}_s)_{s \geq 0} \in \mathcal{H}$ and arbitrary time t define the posterior belief about the state θ to be²¹:

$$Y_{t-} = P(\theta = \theta_H | \mathbb{H}^t) = P(\theta = \theta_H | \mathcal{H}_{t-}).$$

Definition (Markov strategy of the agent). *A strategy of the agent, S , is Markov in state π if for any feasible action profile $\mathbb{H} \in \mathcal{H}$ and any time $t \geq 0$ the induced process $(X_{t+s}, Y_{t+s})_{s \geq 0}$ with*

$$Y_{t+s} = P(\theta = \theta_H | S_{t+s}(\mathbb{H}^t))$$

is Markov.

Verbally, we define the strategy to be Markov, if the future evolution of posterior belief about θ , depends only on the current values of X_t and belief Y_{t-} . Also note that since in our game the agent never has private information, posterior belief is uniquely pinned down for any history \mathbb{H}^t , both on and off the equilibrium path²².

Principal's strategy. We now turn to the principal's option exercise strategy.

Definition. *A strategy of the principal, \mathcal{T} , is a collection of stopping times. For any action profile of the agent \mathbb{H} and any time $t \geq 0$, $\mathcal{T}(\mathbb{H}, t)$ is a stopping time with respect to \mathbb{H} that takes values in $[t, +\infty)$.*

Intuitively, $\mathcal{T}(\mathbb{H}, t)$ is the principal's optimal time of exercising the option calculated from the time t stand point given the history \mathbb{H}^t up time t , the latest information \mathcal{H}_t available at time t , and expectations about future information sharing given by \mathbb{H} .

Next we define a Markov strategy of the principal as an action plan for every state of the game π . Since the action of the principal is binary (stop/wait), a Markov strategy can be identified with a stopping set \mathbf{T} .

²⁰In Section 6 we allow the agent to be privately informed about the state.

²¹As usual sigma-algebra \mathcal{H}_{t-} is defined as $\mathcal{H}_{t-} = \sigma\left(\bigcup_{s < t} \mathcal{H}_s\right)$. Recall that \mathbb{H} satisfies standard conditions, hence Y_{t-} can be also defined as a left limit of π_t , i.e., $Y_{t-} = \lim_{s \uparrow t} Y_s$ with $Y_s = P(\theta = \theta_H | \mathcal{H}_s)$.

²²Formally we restrict the agent's belief after any history \mathbb{H}^t to be Y_{t-} and the principal's belief at time t to be Y_t . That is, we do not allow for belief divergence off path.

Definition (Markov strategy of the principal). A strategy of the principal \mathcal{T} is Markov in state π if there exists a set $\mathbf{T} \subseteq \Pi$ such that

$$\forall \mathbb{H} \in \mathcal{H}, t \geq 0 \quad \mathcal{T}(\mathbb{H}, t) = \inf\{s \geq t : (X_s, Y_s) \in \mathbf{T}\}, \quad (8)$$

where $Y_s = P(\theta = \theta_H | \mathcal{H}_s)$.

The principal's strategy is Markov if the decision to stop depends only on the current level of X and the beliefs induced by \mathcal{H}_t . In particular, if the agent deviates from the equilibrium path but later acquires and shares information that brings beliefs back to the equilibrium path, the principal cannot punish such behavior and has to behave the same way as if that belief was reached on-path.

Heuristically, the sequence of actions in a short period of time dt is shown in Figure 1 and can be described as follows: for a given $\omega \in \Omega$ (i) first, the value of $X_t(\omega)$ is realized, (ii) second, additional information about θ from $S_t(\mathbb{H}^t)$ is announced (iii) finally, the principal updates her belief upon observing everything and decides whether to exercise the option or not.

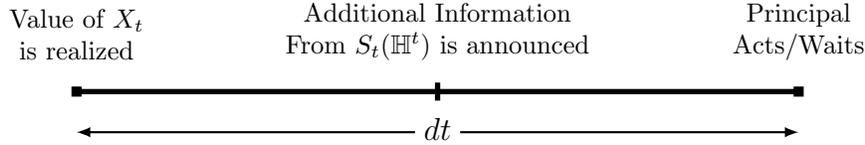


Figure 1: Timing of Events in a Short Interval of Time

2.3 Equilibrium.

We define Markov equilibrium in posterior beliefs, π , as a pair of Markov strategies that are mutual best responses. Lack of agent's commitment manifests itself through condition 3 (Consistency) which requires the future evolution of beliefs to depend only on the current state.

Definition. A Markov Perfect Equilibrium in posterior beliefs, π , is a collection of Markov strategies of the agent and principal (S^*, \mathcal{T}^*) such that

1. **Principal's optimality:** The principal's strategy \mathcal{T}^* is optimal given the anticipated flow of information $\mathbb{H}^* = S^*(\mathbb{H}^0)$, i.e.

$$\forall \pi \in \Pi \quad \mathcal{T}^*(\mathbb{H}^*, 0)|_{\pi_{0-}=\pi} \in \arg \max_{\tau \in \mathcal{M}(\mathbb{H}^*)} \mathbb{E} \left[e^{-r\tau} \left((Y_\tau \theta_H + (1 - Y_\tau) \theta_L) X_\tau - I_P \right) | \pi_{0-} = \pi \right] \quad (9)$$

where $Y_\tau = P(\theta = \theta_H | S_\tau^*(\mathbb{H}^0))$ and the maximum is taken over the set of all stopping times $\mathcal{M}(\mathbb{H}^*)$ with respect to filtration \mathbb{H}^* .

2. **Agent's optimality:** *The agent's strategy S^* is optimal given the anticipated stopping rule of the principal \mathcal{T}^* , i.e.*

$$\forall \pi \in \Pi \quad S^*(\mathbb{H}^0)|_{\pi_{0-}=\pi} \in \arg \max_{\mathbb{H} \in \mathcal{H}} \mathbb{E} \left[e^{-r\tau^*} \left((Y_{\tau^*} \theta_H + (1 - Y_{\tau^*}) \theta_L) X_{\tau^*} - I_A \right) \mid \pi_{0-} = \pi \right] \quad (10)$$

where $\tau^* = \inf\{t > 0 : (X_t, Y_t) \in \mathbf{T}^*\}$.

3. **Consistency:** *After any history \mathbb{H}^t the evolution of the state π_s for $s \geq t$, induced by the strategy S^* , depends only on π_{t-} , i.e.*

$$\forall \mathbb{H} \in \mathcal{H} \quad t \geq 0, \pi \in \Pi \quad Law\left((X_{t+s}, Y_{t+s})_{s \geq 0} \mid \pi_{t-} = \pi\right) = Law\left((X_s, Y_s)_{s \geq 0} \mid \pi_{0-} = \pi\right), \quad (11)$$

where $Y_{t+s} = P(\theta = \theta_H \mid S_{t+s}^*(\mathbb{H}^t))$ and $Y_s = P(\theta = \theta_H \mid S_s^*(\mathbb{H}^0))$.

In order to capture the spirit of subgame perfection we require the strategy of the agent S^* to be consistent. Notice that consistency implies that continuation strategy $S^*(\mathbb{H}^t)$ is optimal after arbitrary histories of the game \mathbb{H}^t , even those that do not occur along the equilibrium path. Since the payoff of the principal depends only on the evolution of the belief process Y and state process X , the strategy of the principal \mathcal{T}^* is also optimal after arbitrary histories of the game.

3 General Analysis

In this section we provide some preliminary analysis that facilitates delivery of our main results in the next two sections.

Autarky thresholds. We start with defining two variables: the thresholds at which the principal and the agent would like to exercise the real option in case additional information about θ was unavailable, that is, in case they could only observe the evolution of X_t . For any given belief y about the underlying quality θ , the optimal stopping decision of the principal (or the agent if he is given control rights) that is based *only on exogenous evolution of X* can be characterized by a single threshold²³ $x_P(y)$ ($x_A(y)$) and a corresponding

²³As in the real options analysis in Dixit and Pindyck (2012)

value function V_P^{NI} (V_A^{NI}) that solves

$$rV_i^{NI}(x, y) = \mu x \frac{\partial}{\partial x} V_R^{NI}(x, y) + \frac{1}{2} \phi^2 x^2 \frac{\partial^2}{\partial x^2} V_i^{NI}(x, y), \quad (12)$$

$$V_i^{NI}(0, y) = 0, \quad (13)$$

$$V_i^{NI}(x_i(y), y) = (y\theta_H + (1 - y)\theta_L)x_i(y) - I_i, \quad (14)$$

$$\frac{\partial}{\partial x} V_i^{NI}(x_i(y), y) = y\theta_H + (1 - y)\theta_L \quad (15)$$

Equation (12) is a second order differential equation with an unknown boundary. Hence it requires three initial conditions (13) - (15) to be correctly specified. The smooth pasting condition (15) is a requirement for the thresholds $x_i(y)$ to indeed be optimal. The equations above have closed form solutions, namely for $i \in \{P, A\}$

$$V_i^{NI}(x, y) = \begin{cases} \left(\frac{x}{x_i(y)}\right)^{\beta_1} [(y\theta_H + (1 - y)\theta_L)x_i(y) - I_i], & \text{if } x \leq x_i(y) \\ (y\theta_H + (1 - y)\theta_L)x - I_i, & \text{if } x > x_i(y) \end{cases} \quad (16)$$

$$x_i(y) = \frac{\beta_1}{\beta_1 - 1} \frac{I_i}{y\theta_H + (1 - y)\theta_L}, \quad (17)$$

where $\beta_1 > 1$ is the positive root of the equation $\frac{1}{2}\phi^2\beta(\beta - 1) + \mu\beta = r$.

The autarky thresholds $x_P(y)$ and $x_A(y)$ summarize the conflict of interest between the principal and the agent. If $x_P(y) < x_A(y)$ then principal wants to exercise the option earlier than the agent. With imperfectly aligned incentives between the two parties, the agent searches for ways to structure information sharing to implement an exercise policy more in line with his preferences.

Agent's equilibrium value function. We start with the agent's value function and show that it solves a Hamilton-Jacobi-Bellman equation and a fixed point of a concavification operator (familiar from Bayesian persuasion literature). Which of the two applies depends on whether (x, y) is an extreme point of a lower contour set of the equilibrium value function of the agent $V_A(x, y)$. At the points where the resulting value function is strictly concave in y , it is optimal for the agent to not reveal any information. More broadly, when the HJB is satisfied, the agent can achieve his equilibrium payoff by staying quiet. While we do not solve for the value function explicitly, we use the implications of dynamic concavification in characterizing the equilibrium strategies of both parties.

The first step is to describe beliefs (x, y) at which $V_A(x, y)$ is "strictly concave" in y . Given function $V_A(x, y)$ define a correspondence $B : \mathbb{R} \rightarrow \mathcal{B}(\mathbb{R}^2)$

$$B(x) = \{(x, v) : v \leq V_A(x, y)\},$$

where $\mathcal{B}(\mathbb{R}^2)$ is all the Borel subsets of \mathbb{R}^2 . In words, $B(x)$ is the lower contour set of $V_A(x, y)$ for a fixed x . For an equilibrium value function $V_A(x, y)$ define²⁴

$$E(x) = \left\{ y : (y, V_A(x, y)) \text{ is an extreme point of } B(x) \right\}. \quad (18)$$

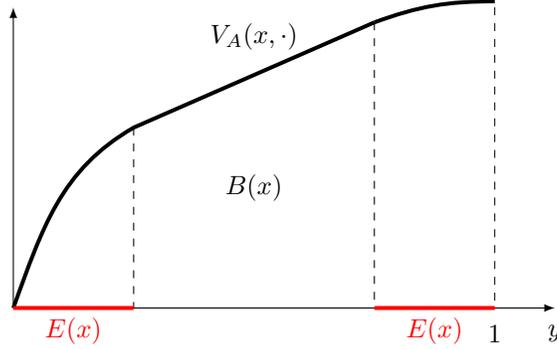


Figure 2: Lower contour set $B(x)$ and the set of extreme points $E(x)$.

For $y \in E(x)$, the agent's value function is strictly concave in y at (x, y) and any signal that he would generate about θ would decrease his total payoff. If $y \in E(x)$ the agent must prefer waiting and not obtaining any additional information. For $y \notin E(x)$, the agent's value function is locally linear in y . Hence he can send the principal a signal about θ and his value function remains unchanged. It may be the case, however, that the agent's value function is locally linear due to information shared in the future, rather than presently.

Definition. For a function $f : \mathbb{R} \rightarrow \mathbb{R}$ define $\text{cav}[f]$ to be the smallest concave function that weakly majorizes f .

We show that if $(x, y) \in \text{int}(E)$, where, $E = \cup_x E(x)$ the agent's value function satisfies a “waiting” Hamilton-Jacobi-Bellman equation in (x, y) similar to (12). When $(x, y) \notin E$ then agent's value function is a convex combination between his value function when he waits, and his value function when he induces the principal to exercise the option:

Lemma 1 (Dynamic Concavification). *In any MPE agent's value function $V_A(x, y)$ is concave in y and continuous in x . For every $(x, y) \in \text{int}(E)$ and $(x, y) \notin \mathbf{T}^*$ function $V_A(x, y)$ satisfies*

$$rV_A(x, y) = \mu x \frac{\partial}{\partial x} V_A(x, y) + \frac{1}{2} \phi^2 x^2 \frac{\partial^2}{\partial x^2} V_A(x, y) \quad (19)$$

If $(x, y) \notin E$ it satisfies

$$V_A(x, y) = \text{cav}_\mu \left[V_A(x, \mu) \mathbb{1} \{ (x, \mu) \notin \mathbf{T}^* \} + \left((\mu \theta_H + (1 - \mu) \theta_L) x - I_A \right) \mathbb{1} \{ (x, \mu) \in \mathbf{T}^* \} \right] (y). \quad (20)$$

²⁴Point $a \in A \subset \mathbb{R}^2$ is extreme if a does not lie on any open line segment joining any two points in A .

Equation (20) is not tautological since for in the region $(x, y) \in \mathbf{T}^*$ it might be possible for the agent to split the belief y into $\{\underline{y}, \bar{y}\}$ and obtain a pay-off $V_A(x, y) > (y\theta_H + (1 - y)\theta_L)x - I_A$. Equations (19) and (20) are necessary equilibrium conditions on the value functions but do not specify the strategies. If agent's value function is locally linear in y at (x_0, y_0) but at the same time satisfies the Hamilton-Jacobi-Bellman equation (19), then he is indifferent between sharing information now or later. When information is shared is relevant for the principal's payoff. Hence whether or not immediate information sharing is consistent with equilibrium behavior is determined by the strategy of the principal.²⁵

4 Persuasion to Wait

In this section we analyze the game in case $I_P < I_A$ which implies that the principal always prefers to exercise the option at a lower threshold than the agent: $x_P(y) < x_A(y)$. By learning and revealing θ fully, the agent could either “speed up” or “delay” the exercise of the real option: the principal would exercise the option at either $x_P(1)$ or $x_P(0)$, depending on the realized value θ . Consider the incentives of the agent state by state. If $\theta = \theta_L$ then he prefers this information to become public as it delays option exercise until threshold $x_P(0)$ is reached, which is the highest possible exercise threshold in any equilibrium²⁶. If $\theta = \theta_H$, then the agent wishes to exercise the option when $X_t \geq x_A(1) > x_P(1)$. This suggests that when $X_t \geq x_A(1)$, the agent prefers to learn and reveal θ : if $\theta = \theta_H$ then the principal stops immediately which is in the agent's best interest, and if $\theta = \theta_L$ the principal delays option exercise until $x_P(0)$ is reached. In other words, two rounds of eliminating dominated strategies imply that it is optimal for the agent to learn and reveal θ when X_t reaches $x_A(1)$. The following lemma established this result formally.

Lemma 2. *In any Markov Perfect Equilibrium, if after any history $X_t \geq x_A(1)$, then the agent immediately reveals full information about θ .*

Lemma 2 implies that whenever X_t reaches $x_A(1)$, the resulting exercise of the real option has a “*compromise*” property: at that time the agent reveals θ and if it turns out that $\theta = \theta_H$, the option is exercised at the agent's first-best threshold, $x_A(1)$; while if it turns out that $\theta = \theta_L$, the option is exercised at the principal's first-best threshold, $x_P(0)$ (in case $\theta = \theta_H$ the stopping is immediate; in case $\theta = \theta_L$ stopping is further delayed).

Our next step is to construct an equilibrium and discuss under what conditions $x_A(1)$ is reached. We then prove that the Markov Perfect Equilibrium is essentially unique.

Principal's equilibrium strategy. When deciding whether to exercise the real option the principal weights the benefits of obtaining more precise information and potential for X_t to increase in the future

²⁵Resolving agent's indifference over when to persuade in equilibrium is important in other models of dynamic persuasion and comes up in Renault, Solan, and Vieille (2014) and Bizzotto, Rudiger, and Vigier (2016).

²⁶At $x_P(0)$ it is optimal for the principal to exercise the option immediately regardless of θ .

against the costs of waiting. While the amount of information shared by the agent is an equilibrium object, Lemma 2 states that in any equilibrium he fully learns and communicates θ whenever $X_t \geq x_A(1)$. For now, assume that discrete revelation of θ when $X_t \geq x_A(1)$ is *the only* information communicated by the agent and consider the best response of the principal.²⁷

Definition. Let $V_P^I(x, y)$ denote the solution of the following optimization problem

$$V_P^I(x, y) = \max_{\tau \in \mathcal{M}(\hat{\mathbb{H}})} \mathbb{E} \left[e^{-r\tau} \left((Y_\tau \theta_H + (1 - Y_\tau) \theta_L) X_\tau - I_P \mid X_0 = x, Y_{0-} = y \right) \right],$$

where

$$Y_t = \begin{cases} y & \text{if } \sup_{s \leq t} X_s < x_A(1), \\ \mathbb{1} \{ \theta = \theta_H \} & \text{if } \sup_{s \leq t} X_s \geq x_A(1), \end{cases}$$

and $\hat{\mathbb{H}}$ is the sigma-algebra generated by (X, Y) .

That is $V_P^I(x, y)$ is the value that the principal obtains when she observes only evolution of X and learns θ as soon as $X_t = x_A(1)$. When $x \geq x_A(1)$ it is simply given by

$$V_P^I(x, y) = y \cdot (\theta_H x - I_P) + (1 - y) \cdot V_P^{NI}(x, 0).$$

The next lemma characterizes some useful properties of $V_P^I(x, y)$ for $x < x_A(1)$.

Lemma 3. *Given the anticipated flow of information $\hat{\mathbb{H}}$, for every y there exists $\underline{x}(y)$ such that for $x \in (x_P(y), \underline{x}(y))$ the principal prefers to act immediately and for $x \in [0, x_P(y)] \cup [\underline{x}(y), x_A(1)]$ the principal prefers to wait. Moreover there exists y^* such that $\underline{x}(y) > x_P(y)$ if and only if $y > y^*$.*

The shape of $V_P^I(x, y)$ for $y > y^*$ is illustrated in Figure 3. There are three distinct regions: between $(0, x_P(y))$ the principal waits optimally due to the standard real options logic (see Dixit and Pindyck (2012)). In $(x_P(y), \underline{x}(y))$ the principal stops immediately. Finally, in $(\underline{x}(y), x_A(1))$ the principal waits because of the option value associated with θ being discretely revealed at $x_A(1)$.

Lemma 3 divides the characterization of $V_P^I(x, y)$ into two cases, depending on how the belief y compares to the threshold y^* . When y is below y^* , it turns out that the conflict of interest between the principal and agent is small, in the sense that instead of two waiting regions as in Figure 3, there is only one. In that case the principal prefers to wait everywhere between 0 and $x_A(1)$ until the agent learns and reveals θ . Notice that this constitutes equilibrium behavior: according to Lemma 2 the agent secures second best if the principal is willing to wait until $x_A(1)$ and for $y < y^*$ it is in the principal's best interest to do so,

²⁷In equilibrium communication (about θ) sometimes occurs also in the region $x < x_A(1)$. However, as we show and explain later, it does not generate any additional value for the principal, so this fictitious calculation allows us to compute the principal's equilibrium payoff.

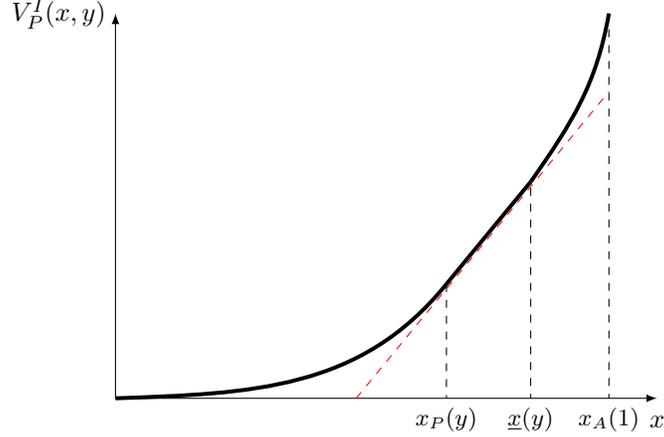


Figure 3: Principal's expected value $V_P^I(x, y)$ from acting only on public process X and the discrete message from the agent at $x_A(1)$.

thus, $V_A(x, y) = V_A^I(x, y)$. The intuition for why small y corresponds to less conflict is that in this case the principal expects to learn that $\theta = \theta_L$ with a high probability so that waiting for the revelation has a large benefit of not exercising the option early.

When $y > y^*$ the disagreement is large, in the sense that $x_P(y) < \underline{x}(y)$, which means that all three regions in Figure 3 are relevant. In this case the equilibrium is more complicated. In particular, when $X_t \in (x_P(y), \underline{x}(y))$, the principal prefers to exercise the option, rather than wait for the $x_A(1)$ threshold. But that cannot be an equilibrium since then the agent would prefer to reveal at least some information to slow down exercise at least in case $\theta = \theta_L$ (and maybe even in more cases).

Define a candidate Markov strategy of the principal, \mathcal{T}^* , based on the shape of $V_A^I(x, y)$ via the following stopping set \mathbf{T}^* :

$$\mathbf{T}^* = \{(x, y) : x_P(y) < x < \underline{x}(y)\} \cup \{(x, y) : x_P(1) < x, y = 1\} \cup \{(x, y) : x_P(0) \leq x\} \quad (21)$$

Intuitively, according to (21) the principal waits if the value of waiting $V_P^I(x, y)$ is strictly above the value of exercising the option $(y\theta_H + (1 - y)\theta_L)x - I_P$. This happens either when $x < x_P(y)$ or when $x_P(y) \leq x \leq x_A(1)$. Additionally, when $x > x_A(1)$ but $x \leq x_P(0)$ the principal anticipates immediate revelation of θ , thus, waiting everywhere but $y = 1$ is optimal. Finally, when $x > x_P(0)$ the principal exercises the option regardless of the realization of θ . The distinction from a standard real option is that the principal now has an endogenous option of waiting for the agent to reveal θ .²⁸ This strategy is depicted in Figure 4.

Agent's equilibrium strategy. For $x < x_A(1)$ even if the quality is θ_H the agent prefers to wait longer for an increase in X_t . This means that, intuitively, the agent will try to minimize the probability of option

²⁸Additionally the principal waits in a measure zero subset of all states in which $V_P(x, y) = (y\theta_H + (1 - y)\theta_L)x - I_P$ to make sure that the best response of the agent is well-defined.

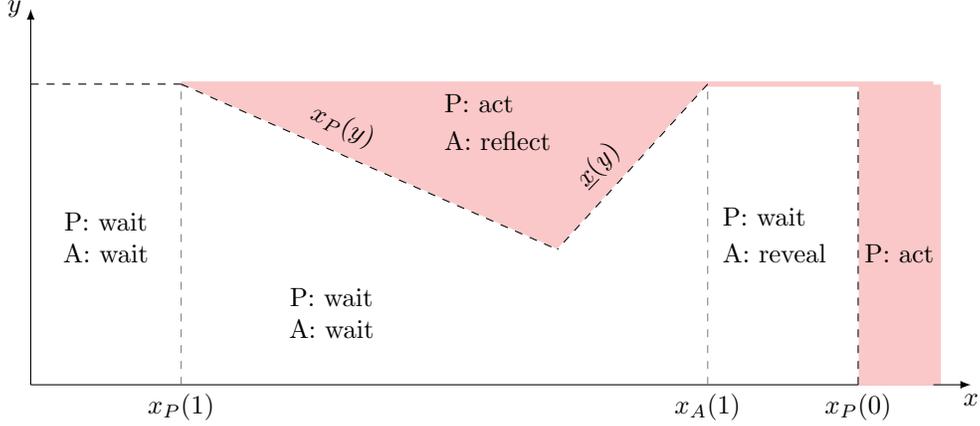


Figure 4: Equilibrium strategies for principal (P) and agent (A). Principal's strategy \mathbf{T}^* in red.

exercise as much as possible when $x < x_A(1)$. We heuristically define the strategy of the agent \mathcal{S}^* via the belief dynamics that it induces:

$$dY_t = (1 - \underline{y}(X_t, Y_{t-}))dN_t - (1 - \underline{y}(X_t, Y_{t-}))d\lambda(X_t, Y_{t-}),$$

where N_t is a pure jump process with such that whenever $dN_t = 1$ both principal and agent learn that the state is θ_H , $d\lambda(x, y)$ is the probability (or intensity) with which the parties learn the state is θ_H , and $\underline{y}(x, y)$ is the posterior belief conditional on $dN_t = 0$.

If $x \geq x_A(1)$, the agent learns and reveals θ immediately, which corresponds to $\underline{y}(x, y) = 0$ and $d\lambda(x, y) = y$. As we saw before, the agent obtains his highest possible equilibrium payoff this way (he gets his optimal stopping when $\theta = \theta_H$ and he delays the principal as much as possible in any equilibrium when $\theta = \theta_L$). If $x < x_A(1)$ and y is in the principal's waiting region, i.e., $y < l(x)$ ²⁹, the agent communicates nothing ($d\lambda(x, y) = 0$, $\underline{y}(x, y) = y$) since both parties prefer to wait. In the action region of the principal $x < x_A(1)$ and $y > l(x)$ absent additional information the principal will act with certainty. In order to delay action the agent conducts a test such that posterior either jumps up to 1 or exits the action region, i.e., jumps down to $l(x)$. Such test corresponds to $\underline{y}(x, y) = l(x)$ and $d\lambda(x, y) = (y - l(x))/(1 - l(x))$.

Finally, at the boundary of the action set $y = l(x)$, the agent pipets information such that belief either jumps to 1 with a positive *intensity* $d\lambda(x, y)$ or is marginally revised downwards, i.e., $\underline{y}(x, y) = y = l(x)$. In absence of the jump the belief process continuously *reflects* off the boundary $y = l(x)$, due to a downward drift $-(1 - Y_{t-})d\lambda(X_t, Y_{t-})$ that exactly offsets the positive increment in dX_t .

In the Appendix we explicitly construct the dynamic strategy of the agent S^* that induces a Markov process $(X_t, Y_t)_{t \geq 0}$ corresponding to the heuristic description above. It requires some technical details because in continuous time we need to construct a reflected belief process along a moving boundary $l(x)$. For now,

²⁹Where $l(x)$ is the appropriate inverse $x_P^{-1}(x)$ or $\underline{x}^{-1}(x)$.

assume that such a process exists and is well-defined. Next proposition delivers the main result of this section, it shows that the strategies (S^*, \mathcal{T}^*) constitute an equilibrium.

Proposition 1 (Persuasion to Wait). *In the case of Persuading to Wait ($I_P > I_A$), the pair of strategies (S^*, \mathcal{T}^*) described above constitute a Markov Perfect Equilibrium. Moreover, in this equilibrium $V_P(x, y) \equiv V_P^I(x, y)$.*

Below we present a sketch of the proof of the Proposition. Technical details are deferred to Appendix.

Sketch of Proof.

First, we check that given the strategy of the agent S^* , principal's strategy \mathcal{T}^* is a best response. When $x \geq x_A(1)$ the agent perfectly learns and discloses θ . Thus \mathcal{T}^* is optimal as long as the principal acts in the set $\{y = 0, x > x_P(0)\}$ and $\{y = 1, x > x_P(1)\}$. Both of these sets are in the action region \mathbf{T}^* .

If $x < x_A(1)$ then waiting below $l(x)$ is optimal for the principal. Since the conjectured equilibrium involves communication in the region $x < x_A(1)$ the value function corresponding to any best response on S^* is weakly above V_P^I defined above. Intuitively, the principal can do only better with more information. But $V_P^I(x, y) > (y\theta_H + (1 - y)\theta_L)x - I_P$ in this region, thus, waiting is optimal. Next we check that acting is optimal for $y > l(x)$. First, Lemma 8 in the Appendix establishes that for $\underline{x}(y)$ strictly increasing in y , thus, value function $V_P^I(x, y) = (y\theta_H + (1 - y)\theta_L)x - I_P$ is linear in y for $y > l(x)$. Since the agent's strategy randomizes over $y = 1$ and $y = l(x)$ in this region, such randomization does not generate any value for the principal. In other words

$$\begin{aligned} V_P^I(x, y) &= \alpha V_P^I(x, 1) + (1 - \alpha)V_P^I(x, l(x)) \\ &= \alpha[\theta_H x - I_P] + (1 - \alpha)[(l(x)\theta_H + (1 - l(x))\theta_L)x - I_P] \\ &= (y\theta_H + (1 - y)\theta_L)x - I_P, \end{aligned}$$

and the principal's value from following \mathcal{T}^* does not depend whether the agent randomizes between 1 and $l(x)$ or remains silent (in the latter case the option is exercised). Since the principal is indifferent in this region between acting and not, \mathcal{T}^* is one possible best response to S^* .

Second, we check whether the agent behaves optimally, given \mathcal{T}^* . Lemma 2 establishes that it is optimal for the agent to learn θ immediately when $x > x_A(1)$, consistent with S^* . Next, consider the region $\{x < x_A(1), y < l(x)\}$, in which the agent prefers waiting over immediate action regardless of θ . Since an informative message can only accelerate the timing of the option exercise, it is optimal for him to stay quiet.

Finally, consider a state $(\hat{x}, \hat{y}) \in \{(x, y) : x < x_A(1), y > l(x)\}$. Absent any information from the agent, the principal acts immediately in (\hat{x}, \hat{y}) . If the agent does nothing, he gets his terminal value $(\hat{y}\theta_H + (1 - \hat{y})\theta_L)\hat{x} - I_A$. This is strictly dominated by a fully revealing message. Hence persuasion is optimal for the agent at

(\hat{x}, \hat{y}) . It is sufficient to consider binary lotteries \underline{y}, \bar{y} such that $1 \geq \bar{y} > l(\hat{x}) \geq \underline{y}$. Suppose $\bar{y} < 1$. State (\hat{x}, \bar{y}) has similar problems to (\hat{x}, \hat{y}) in the sense that at (\hat{x}, \bar{y}) the agent would reveal all of the information rather than stay quiet. Hence it must be the case that $\bar{y} = 1$. Suppose that $\underline{y} < l(\hat{x})$. This information sharing can be implemented via a compound message. First, the agent sends beliefs either to $l(\hat{x})$ or to 1. Then, conditional on being at $l(\hat{x})$, the agent sends beliefs either to \underline{y} or to 1. We know, however, that at $(\hat{x}, l(\hat{x}))$ the agent prefers to stay quiet since the principal is not acting, and hence $\underline{y} < l(\hat{x})$ cannot be optimal. By setting $\underline{y} = l(x)$ the agent minimizes the probability of immediate action and keeps the option of sharing the information in subsequent periods.

The proof in the Appendix deals explicitly with the technical difficulty of defining actions of the agent at the boundary $y = l(x)$. We show that the agent generates a messages that reveals $\theta = \theta_H$ with positive *intensity* making the belief process reflect from the boundary of the action region conditional on sending a negative “message”. \square

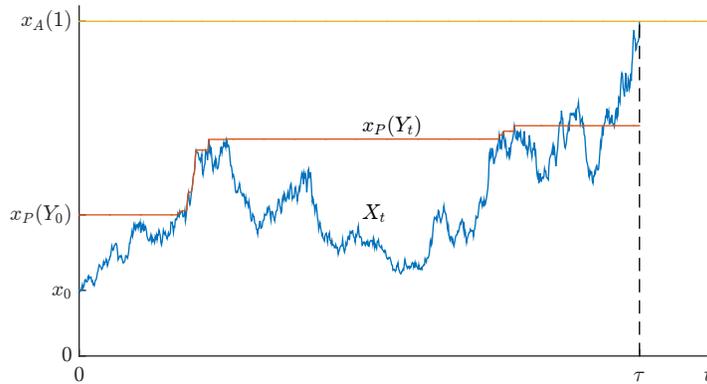


Figure 5: Equilibrium dynamics with $\theta = \theta_H$ in the lower waiting region.

The strategies (S^*, \mathcal{T}^*) generate the dynamics shown on Figure 5. In this example we start with $x_0 < x_P(y_0)$ and $y_0 > y^*$ meaning that the principal would rather exercise the option at $x_P(y_0)$ than wait for θ to be revealed at $x_A(1)$. Whenever X_t is below the principal’s “autarky” threshold, $x_P(Y_t)$, the agent stays quiet and both parties observe only the evolution of X . At the threshold $X_t = x_P(Y_t)$ the belief y_t , reflects at $x_P(y_t)$ conditional on a negative message about θ (or the belief jumps up to $y = 1$ with the intensity to satisfy the Bayes rule). A negative signal about θ (decrease in y) makes the principal more pessimistic and willing to wait for more news about X_t . As y_t decreases due to persuasion, the action threshold of the principal $x_P(y)$ increases and approaches the constant threshold $x_A(1)$. A decrease in the gap between $x_P(y)$ and $x_A(1)$ lowers the principal’s cost of waiting for the full information. When the value of information is sufficiently high relative to the cost of waiting and the principal is willing to wait until $x_A(1)$ to make a decision, the agent optimally stops persuasion (this point corresponds to the “tip of the triangle” in Figure 4). Further waiting ensues until the state process X_t reach a new higher threshold $x_A(1)$ at which the agent

fully learns and reveals θ .

Value of Information and Authority. Earlier in this section we constructed value function $V_P^I(x, y)$ fictitiously assuming that the agent reveals information only at $x_A(1)$, while in the equilibrium we just described the agent also pipets information at the boundary $l(x)$. In general, additional information weakly increases the payoff of a decision maker who faces a real option problem. In our equilibrium information release is timed in a way that affects the principal’s actions but does not affect her payoff. Technically, even though the value function of the principal is convex in y , it is locally linear at the boundaries where information is revealed (this follows from the Envelope Theorem). Economically, the equilibrium information disclosure reflects beliefs only locally: at the boundary the principal is indifferent between exercising and waiting; upon receiving information she is either indifferent again or strictly prefers to exercise. Since both exercising and waiting is optimal at those boundary points, the additional information pipetted by the agent does not create positive value for the principal. The principal exactly trades off benefits of a more informed decision with the costs of a delayed decision. Information has no value for the principal in this region, but authority is important. Delegation results in option exercise at either $x_A(1)$ or $x_A(0)$ and would yield a lower payoff to the principal.

The “compromise” property implies that if the prior beliefs and other parameters of the game are such that $Y_{0-} < y^*$, then the equilibrium is robust to replacing verifiable information about θ with cheap talk. The reason is that at $x_A(1)$ the agent has aligned incentives with the principal to reveal the truth. If the agent does not say that $\theta = \theta_H$ past $x_A(1)$, the principal can credibly infer that $\theta = \theta_L$ and stop accordingly. That is not true in the “pipetting” region: there the θ_H type agent would prefer to mislead (at least temporarily) the principal into believing that $\theta = \theta_L$ in order to delay option exercise.

When $y < y^*$ information has value for the principal. In the special case when $\theta_L = 0$, i.e. both parties agree that the option should not be exercised in the low state, the “compromise” property is equivalent to delegating the real option to the agent. In this region the principal has no real authority.

Proposition 2 (Uniqueness). *If $I_P < I_A$, then (S^*, \mathcal{T}^*) is the essentially³⁰ unique MPE.*

We establish uniqueness by first eliminating dominated strategies of the players. Lemma 2 shows that in any equilibrium the agent fully learns θ as soon as $x \geq x_A(1)$. This implies that it is dominant for the principal to wait when $x < x_A(1)$ and $y < l(x)$ in any equilibrium. Next, we notice that in the region $\{x < x_A(1), y < l(x)\}$ the agent prefers to not to learn any new information regardless of θ , thus, in any equilibrium she stays quiet.

The argument above pins down the strategies of the principal and the agent everywhere except the region $\{x < x_A(1), y \geq l(y)\}$. The proof in the Appendix shows that in any equilibrium the strategies in the

³⁰By “essentially” we mean that for any initial conditions (X_0, Y_{0-}) the distribution over outcomes (X_τ, Y_τ, τ) of any MPE coincides with the one induced by (S^*, \mathcal{T}^*) .

remaining region have to coincide with (S^*, \mathcal{T}^*) . Below we give a sketch of the proof to convey the intuition.

Suppose that the action region of the principal $\mathbf{A} = \{(x, y) \in \mathbf{T} : x < x_A(1)\}$ is strictly inside of the “triangle” $\{x < x_A(1), y > l(y)\}$. In this case it is optimal (similar to the argument in the proof of Proposition 1) for the agent to stay quiet outside of \mathbf{A} . Since the principal is waiting outside of \mathbf{A} his payoff is weakly above $(y\theta_H + (1 - y)\theta_L)x - I_P$, at the same time, since he is receiving less information compared to S^* , his payoff should be weakly below $V_P^I(x, y)$. Thus, in the region $\{x < x_A(1), y \geq l(y)\} \setminus \mathbf{A}$ the payoff of the principal should be equal to $V_P(x, y) = (y\theta_H + (1 - y)\theta_L)x - I_P$. This payoff is linear in the state variable x , which is incompatible with strictly positive waiting time in this region (since the linear payoff from waiting can not satisfy the second order ODE (19)). Since in any equilibrium the beliefs spend no time in the region $\{x < x_A(1), y \geq l(y)\}$, it is outcome-equivalent to (S^*, \mathcal{T}^*) .

4.1 Comparative Statics

We finish this section discussing comparative statics of the equilibrium. Recall that belief level y^* is defined as the unique solution of $x_P(y^*) = \underline{x}(y^*)$, i.e. at $(X_t, Y_t) = (x_P(y^*), y^*)$ the principal is indifferent between exercising the option immediately and waiting till $x_A(1)$ to learn θ . Visually, point (x^*, y^*) is the lower tip of the principal’s action “triangle” depicted in Figure 4. This point is important because it affects equilibrium dynamics: if $Y_{0-} \leq y^*$ then along the equilibrium path information is acquired and revealed once and fully when $X_t = x_A(1)$ and the timing of the option exercise features a “compromise” property. If, instead, $Y_{0-} > y^*$, then along the equilibrium path when X_t reaches $x_P(Y_{0-})$ the agent engages in persuasion and initially “pipets” information in order to prevent the principal from exercising the option.

Role of agent’s information. The next lemma shows that the amount of persuasion in equilibrium is decreasing with the importance of the agent’s information, i.e. higher θ_H or lower θ_L . Consequently, it is easier to reach the “compromise” outcome if the information that the agent can acquire is of high importance.

Lemma 4. *The range of beliefs y for which the agent does not engage in persuasion, $[0, y^*]$*

(i) *expands with an increase in θ_H , i.e. $\frac{\partial y^*}{\partial \theta_H} > 0$,*

(ii) *shrinks with an increase in θ_L , i.e. $\frac{\partial y^*}{\partial \theta_L} < 0$.*

An increase in θ_H affects incentives to wait through two channels. First, it reduces the information revelation threshold $x_A(1)$ since, conditional on $\theta = \theta_H$ the Sender’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. At $(x_P(y^*), y^*)$ the principal is exactly indifferent between acting immediately and waiting until $X_t = x_A(1)$. Second, as θ_H increases, the principal’s option value to wait at $(x_P(y^*), y^*)$ goes up as well, leading her to

prefer waiting rather than acting, thus, the action region of the principal shrinks. In response, the agent persuades less and the compromise is more likely to be reached along the equilibrium path.

An increase in θ_L has no effect on incentives of the agent to learn and reveal θ perfectly at $x_A(1)$. However, it makes the information disclosed at $x_A(1)$ of worse quality since the gap in value between acting (suboptimally) and waiting (optimally) at a fixed x^* is decreasing in θ_L . As a result acting immediately becomes a more attractive option and the action region of the principal increases. Larger action region of the principal implies more persuasion by the agent and smaller probability of compromise along the equilibrium path.

Level of disagreement. In our model disagreement between the principal and the agent is summarized by the distance between the autarky thresholds, $x_P(y)$ and $x_A(y)$, that in turn depends on magnitude of $I_P - I_A$. The next lemma shows that with higher disagreement, equilibrium features more persuasion and compromise is less likely to be reached. Moreover, information is acquired earlier (in the FOSD sense) along the equilibrium path when the principal is more impatient, i.e. has lower $x_P(y)$.

Lemma 5. *The range of beliefs y for which the agent does not engage in persuasion, $[0, y^*]$*

(i) *expands with an increase in I_P , i.e. $\frac{\partial y^*}{\partial I_P} > 0$,*

(ii) *shrinks with an increase in I_A , i.e. $\frac{\partial y^*}{\partial I_A} < 0$.*

Let $\tau_1 = \inf\{t > 0 : Y_t = 1\}$. Then for $I_P > I'_P$ the τ_1 first order stochastically dominates τ'_1 .

The level of disagreement between the parties affects equilibrium dynamics through the distance between $x_P(y)$ and $x_A(1)$. The *nature* of equilibrium information sharing depends on whether the principal finds it worthwhile to wait for a discrete piece of information, obtained at $x_A(1)$, when $X_t = x_P(Y_{t-})$. More disagreement implies longer wait times between $x_P(y)$ and $x_A(1)$, which makes it suboptimal for the principal to delay option exercise and increases the action region. Larger action region of the principal implies more persuasion by the agent and smaller probability of compromise along the equilibrium path.

In addition, the level of disagreement between the agent and the principal affects the *timing* of information acquisition in equilibrium. A decrease in $x_P(y)$ forces the agent to start persuading earlier (in the FOSD sense) in order to prevent the principal from exercising the option at the autarky threshold. At the same time any change in x_P has no effect on the incentives of the agent to perfectly learn θ at $x_A(1)$, thus, overall information sharing occurs earlier when the principal is more impatient.

An increase in x_A has no clear effect on the timing of equilibrium information acquisition due to two opposing forces. On the one hand, higher level of disagreement increases the persuasion region, which accelerates equilibrium learning. On the other hand, higher threshold $x_A(1)$, at which the agent is willing to perfectly learn θ , implies delayed learning in equilibrium.

5 Persuasion to Act

We now turn to the opposite case, when it is the principal who would like to exercise the real option sooner than the agent, which is captured by the assumption $I_A < I_P$. As we discussed above, that can correspond to a situation where an agent works for a principal and they learn jointly from public news about a potential project the agent would like to start. The conflict of interest is either because the project gives the agent private benefits or the agent does not fully internalize the fixed cost of starting it, he would like to start his project sooner than the principal.

Our first result is that there exists a MPE in which the agent reveals information immediately, after every history.

Proposition 3. *There exists a MPE in which signal θ is fully revealed by the agent after all histories.*

In such equilibrium the principal achieves first best by adopting a very “strong bargaining position”: she threatens not to exercise the option until either all information about θ is revealed, or the level of X_t renders additional information that could be obtained by the agent irrelevant, i.e. $X_t \geq x_P(0)$. Given the equilibrium strategy of the agent such threat is credible, since waiting for the information to be released at the next instance is costless in continuous time. The agent faces a tough choice: either to reveal θ immediately, or to wait until $X_t = x_P(0)$; any partial information revelation does not affect the timing of option exercise. If the agent reveals θ , then the option is exercised either at $x_P(1)$ (if $\theta = \theta_H$) or at $x_P(0)$ (if $\theta = \theta_L$). Since she prefers the option to be exercised earlier and $x_P(1) < x_P(0)$ the agent is better off revealing θ at time 0.

Note that the same reasoning does not apply to the previous case of a patient agent. In that case, if the principal could commit to a strategy “I stop in the next second unless you reveal all information about θ ,” the agent’s best response would be to reveal such additional information. But that threat is not credible: if the agent does not provide information, the principal will not stop before state X_t reaches $x_P(y)$.

Despite a similar outcome, the standard unraveling argument of Grossman (1981) and Milgrom (1981) has no bite here since the agent has no private information. Instead this result is based on an intuition similar to the Coase conjecture, formalized in Gul, Sonnenschein, and Wilson (1986). Suppose that the principal employs a less stringent decision rule, i.e., she is willing to exercise the option prior to learning all information. In this case the agent has incentives to provide just enough information to induce posterior beliefs that would culminate in immediate action. However, the information generated by the agent does not only affect the instantaneous level of principal’s belief, due to lack of agent’s commitment it also increases the temptation of the principal to wait for more information *in the future*. As a result of persuasion, the principal shifts the action threshold, which prompts the agent to generate more information, which shifts the action threshold even further, and so on.

In the rest of this section we show sufficient conditions under which this is the unique Markov Perfect Equilibrium. We also show that in general, the principal’s decision threshold is always higher than absent

persuasion, which implies that in these cases the agent could be worse off as a result of having access to information about θ .

Proposition 4. *If $x_0 \leq x_P(1)$, then the equilibrium outcome of Proposition 3 is essentially unique. Moreover, regardless of x_0 in any MPE (\mathcal{T}^*, S^*) the stopping set \mathbf{T}^* is strictly inside of the autarky stopping set $\{(x, y) : x \geq x_P(y)\}$.*

We prove the first part of the proposition by showing that given the initial condition $x_0 \leq x_P(y)$ in any MPE the option is exercised only when θ is fully revealed. Since the agent prefers early exercise immediate revelation of θ guarantees that the principal stops exactly when $X_t = x_P(1)$ or $X_t = x_P(0)$.

To see why the second part of Proposition 4 holds, one has to notice that full revelation of θ at $x_P(1)$ creates additional option value for the principal when $x_0 > x_P(1)$. Hence he will no longer prefer to stop at $x_P(y)$ for $0 < y < 1$ because his value from waiting is now strictly higher.

Suppose now that the principal's prior belief at $t = 0$, Y_{0-} , is such that $X_0 = x_R(Y_{0-})$. Then, absent persuasion, the principal would immediately stop. However, in equilibrium he expects the agent to reveal more information when X_t decreases. Hence he will no longer wish to exercise the option at $x_P(Y_{0-})$. This hurts the payoff of the agent since, since she prefers immediate exercise at $x_P(Y_{0-})$ due to $x_A(y) < x_P(y)$ for all y .

Corollary 1. *Suppose $V_A^{NI}(x_0, 0) < \theta_L x_0 - I_A$ at $t = 0$ and $y_0 \in (0, 1)$. If initial belief is such that $X_0 = x_P(Y_{0-})$ then agent's (principal's) value in any MPE is strictly lower (higher) than his value if the principal acted just based on evolution of X_t .*

Proof. Absent persuasion the principal would have exercised the real option at $x_P(Y_{0-})$ which would have been the best outcome for the agent. Due to persuasion, however, principal's decision threshold is strictly higher than $x_P(Y_{0-})$. \square

When the agent is significantly more impatient than the principal both authority and information generate strictly positive value for the principal. Full delegation results in the agent learning the state θ but exercising the option too early either at $x_A(1)$ or $x_A(0)$. Authority absent any communication improves principal's payoff and results in exercise at $x_P(y)$. Although the principal does not learn the state, she benefits from the ability to delay the option exercise past x_A . Finally, communication delivers first best for the principal: she gets to both learn the state and exercise the option at her optimal threshold $x_P(1)$ or $x_P(0)$.

6 Persuasion by an Informed Agent

So far we assumed that the agent has no private information and learns about θ together with the principal. In this section we relax this assumption to allow for a broader range of applications in which such assumption

might fail. For example, a product manager could have superior information about the customer's willingness to pay *before* he conducts a credible test, such information might affect the type of test he is willing to run. Hence the principal could potentially learn about θ not only from the outcome of test, but also from the structure of the test itself. In what follows, we show that equilibrium outcomes of Propositions 1 and 3 are consistent with the dynamic persuasion model in which the agent is privately informed about θ .³¹ We proceed by extending the definitions of agent's and principal's strategies and incorporate the possibility of dynamic signaling into the equilibrium definition.

The stopping strategy of the principal carries over from the main model Section 2. In order to incorporate agent's private information we extend the definition of agent's strategy to a pair (S^H, S^L) . Similar to Section 2 for every \mathbb{H}^t the strategy of the type $i \in \{H, L\}$ is a mapping into the continuation filtration, i.e. $S^i(\mathbb{H}^t) \in \mathcal{H}(\mathbb{H}^t)$.

Next we define the *naive* beliefs of the principal for a given filtration \mathbb{H} .

Definition (Naive Beliefs). *For any filtration $\mathbb{H} = (\mathcal{H}_t)_{t \geq 0}$ the naive belief process $\hat{Y} = (\hat{Y}_t)_{t \geq 0}$ associated with that filtration is defined as $\hat{Y}_t = P(\theta = \theta_H | \mathcal{H}_t)$.*

In other words, the naive belief process \hat{Y} updates only on the hard information about θ contained in filtration \mathbb{H} . In a standard Bayesian persuasion framework beliefs of the principal are naive, since the signaling component is absent - the agent either does not know θ himself or commits to a persuasion policy ex-ante. With private information and lack of commitment on the side of the agent the principal could learn not only from the content of the information generated by the agent but also from the structure of such information, thus, principal's equilibrium belief process might differ from the naive one.

We call the strategy of the agent Markov if it induces Markov process of naive beliefs.

Definition (Markov strategy of the informed agent). *A strategy of the agent, (S^H, S^L) , is Markov if, for any feasible action profile $\mathbb{H} \in \mathcal{H}$ and any time $t \geq 0$, the induced process $(X_{t+s}, \hat{Y}_{t+s}^i)_{s \geq 0}$ with*

$$\hat{Y}_{t+s}^i = P(\theta = \theta_H | S_{t+s}^i(\mathbb{H}^t)) \quad i \in \{H, L\}$$

is Markov.

For any filtration $\mathbb{H} = (\mathcal{H}_t)_{t \geq 0}$ played along the game path (on or off equilibrium) we can define two stopping times τ^H and τ^L , such that the observed filtration becomes inconsistent with conjectured equilibrium

³¹An alternative way to model strategic communication with privately informed agent is via a disclosure of hard evidence similar to Hörner and Skrzypacz (2016). Equilibrium outcomes of Propositions 1 and 3 can be supported in such framework as well.

behavior of high and low type agent respectively³². Formally, put

$$\tau^H = \sup\{\tau : \mathcal{H}_\tau = S_\tau^H(\mathbb{H}^0)\} \quad \text{and} \quad \tau^L = \sup\{\tau : \mathcal{H}_\tau = S_\tau^L(\mathbb{H}^0)\}, \quad (22)$$

where supremum is taken over all stopping times τ with respect to \mathbb{H} . The supremum is well defined since the stopping times that satisfy $\mathcal{H}_\tau = S_\tau^H(\mathbb{H}^0)$ (or $\mathcal{H}_\tau = S_\tau^L(\mathbb{H}^0)$) form a lattice as shown in Lemma 10 in the Appendix.

Intuitively, the principal does not observe the whole sigma-algebra \mathcal{H}_t at time t , since this object contains information about the distribution of \hat{Y}_t for all possible realizations of X_t . She only observes the relevant part of \mathcal{H}_t which is conditioned on the observed path $(X_s, \hat{Y}_s)_{s \leq t}$. From principal's point of view the observed distribution and path of naive beliefs up to time t are consistent with the equilibrium strategy of the high (low) type agent when $t < \tau_H$ ($t < \tau_L$).

Equilibrium. We define a (pure) Markov Perfect Equilibrium as a triple: a Markov strategy of the agent (S^H, S^L) , a Markov strategy of the principal \mathcal{T}^* and a collection principal's belief processes $Y = \{Y(\mathbb{H}), \mathbb{H} \in \mathcal{H}\}$ ³³ that satisfy (i) principal's optimality, (ii) agent's optimality, (iii) consistency and (iv) belief consistency. The first three equilibrium conditions are direct analogs of those in Section 2. Belief consistency captures the Bayesian updating of principal's beliefs from the content of shared information, as well as the structure of it.

Definition (Belief Consistency). *Principal's beliefs Y are consistent with the strategy of the agent (S^H, S^L) if for any filtration \mathbb{H} and any time t the belief process $Y(\mathbb{H})$ satisfies*

1. $Y_t(\mathbb{H}) = \hat{Y}_t$ for $t < \min(\tau^H, \tau^L)$,
2. $Y_t(\mathbb{H}) = 1$ for $\tau^L \leq t < \tau^H$, $Y_t(\mathbb{H}) = 0$ for $\tau^H \leq t < \tau^L$,
3. $Y(\mathbb{H})$ is a càdlàg and adapted to \mathbb{H} process for $t \geq \max(\tau^H, \tau^L)$.

Condition 1 above requires the principal's belief to be naive as long as the played filtration \mathbb{H} is consistent with $S^H(\mathbb{H}^0)$ and $S^L(\mathbb{H}^0)$ conditional on the observed path up to time t . Condition 2 pins down principal's belief along the equilibrium path if either type of the agent reveals himself. Finally, condition 3 puts no economic restrictions on the out of equilibrium beliefs.

Along the equilibrium path the agent is able to signal his type by picking $S^H \neq S^L$. Upon observing the difference the principal immediately learns the type of the agent and then acts accordingly. The same

³²The definition of agent's strategy adopted from Section 2 essentially allows only pure strategies, thus, along the equilibrium path principal's beliefs are either naive (pooling) or degenerate (separating). If one were to allow the agent to privately mix between different filtrations then there could potentially exist semi-pooling equilibria as well.

³³We endow the principal with a collection of belief processes to capture beliefs for all possible out of equilibrium filtrations $\mathbb{H} \in \mathcal{H}$.

outcome can be achieved by modifying the original strategies of the agent to include a fully informative test at $t = \min(\tau^H, \tau^L)$. Thus, it is without loss to have the principal's beliefs be naive on path in any equilibrium³⁴. Next we introduce specific off path beliefs of the principal that ignore the signaling component of the game and update only on the hard information about θ .

Definition (Passive Beliefs). *Principal's beliefs Y are passive if for every $\mathbb{H} \in \mathcal{H}$ and every $t \geq 0$*

$$Y_t(\mathbb{H}) = \hat{Y}_t, \quad (23)$$

where \hat{Y} is the naive belief process associated with \mathbb{H} .

Next we show that equilibrium outcomes of Propositions 1 and 3 can be supported by pooling equilibria in which the agent chooses the same strategy regardless of his type and the principal holds passive beliefs.

Corollary 2 (Persuasion to Wait). *Suppose $I_P < I_A$ and let (S^*, \mathcal{T}^*) be the equilibrium strategies from Proposition 1. Then the strategy of the agent (S^*, S^*) , the strategy of the principal \mathcal{T}^* and passive beliefs Y constitute a Markov Perfect Equilibrium.*

The intuition behind the proof of Proposition 2 rests on the indifference of the θ_L type agent as long as the action is not undertaken prior to $x_P(0)$. By following S^* the low type guarantees himself exercise at $x_P(0)$, since $x_P(0)$ is the highest exercise threshold in any equilibrium, the agent has no incentives to deviate. The high type has no incentives to deviate from S^* for $X_t > x_A(1)$, since S^* calls for a fully informative test and results in immediate option exercise, first-best outcome for the high type agent. For $X_t < x_A(1)$ the high type prefers the principal to wait and has incentives to conduct only the least informative tests to delay the action. As argued in Proposition 1, S^* induces maximal pooling of both types and minimizes the probability of option exercise prior to $x_A(1)$, thus, the high type has no incentives to deviate. The principal observes the information generated by the agent and updates her beliefs, which follow a martingale. However, from the agent's perspective principal's beliefs have a drift: the low type expects the process to never jump to 1, while the high type expects a positive intensity of a jump at the boundary of the principal's action region. The logic is similar to Hörner and Skrzypacz (2016).

Next we show that the equilibrium outcome of the Proposition 3 can be sustained with privately informed agent.

Corollary 3 (Persuasion to Act). *Suppose $I_P > I_A$ and let (S^*, \mathcal{T}^*) be the equilibrium strategies from Proposition 3. Then the strategy of the agent (S^*, S^*) , the strategy of the principal \mathcal{T}^* and passive beliefs Y constitute a Markov Perfect Equilibrium.*

³⁴For any equilibrium for which $S^H \neq S^L$ there exists an outcome equivalent equilibrium with $S^H = S^L$ in which the beliefs are naive on path.

By revealing himself at time 0 the high type agent speeds up option exercise as much as possible, given the strategy of the principal, hence, he has not incentives to deviate. With passive beliefs the low type is indifferent between all possible information sharing structures since for any belief below 1 the option is exercised at $x_P(0)$, thus, he has no strict incentives to deviate either. 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.

This dynamic signaling game unsurprisingly features equilibrium multiplicity. For example, in case of persuasion to wait full revelation at time 0 is an equilibrium outcome supported by harsh out of equilibrium beliefs that assign any deviation to the high type agent.

7 Wald Options

In many natural applications the innovations to principal’s expected payoff are correlated with the fundamental state of nature. For example, reports from patients and doctors provide information to the FDA on the likelihood of possible outcomes of the laboratory tests run by the drug sponsor. Similarly, prices and analysts’ forecasts that provide the creditors with information on the financial stability of a particular bank are informative about the outcomes of the stress tests of the same institution. In this section we relax the earlier assumption that the payoff relevant state is a product of two orthogonal components θ and X_t . As in the applications outlined above, evolution of X_t may provide information about θ .

The principal and the agent observe public news about a payoff relevant fundamental state of nature. The real option, colloquially referred to as the Wald option³⁵, faced by the principal is to decide at what belief threshold it is optimal to stop. This structure introduces correlation between the exogenously moving state (exogenous news) and the information that the agent can acquire (endogenous news). We argue below, that qualitative features of the Wald option model, strategies, equilibrium concept, and equilibrium are identical to our main model specification in Section 2. We choose to work with a Wald option as it is a direct continuous time extension of Kamenica and Gentzkow (2011) utilizing the methodology introduced in Sections 2.2 and 2.3.

7.1 Model

Consider the FDA example. Suppose the drug, currently present in the market, can be either safe $\xi = 0$ or unsafe $\xi = 1$. Neither the Regulator (principal = FDA) nor the drug producer (agent) know the state ξ and share a common prior $q_{0-} = P(\xi = 1)$. The Regulator has the option to withdraw the product from the market at any time. If the Regulator withdraws a safe drug from the market she suffers a welfare loss

³⁵Bayesian inference is explored in the seminal contribution Wald (1973), and applied to real option problems in Stokey (2008).

$W_0 > 0$. However, if a harmful drug is withdrawn the regulator increases welfare by $W_1 > 0$.³⁶ The expected welfare gain from withdrawing the product is

$$v_P(q) = E[W_1 \cdot \mathbb{1}\{\xi = 1\} - W_0 \cdot \mathbb{1}\{\xi = 0\}] = (W_1 + W_0)q - W_0, \quad (24)$$

where q is the principal's posterior probability of the drug being unsafe $\xi = 1$. Preferences of the drug firm (agent) can be written in a similar way

$$v_A(q) = (F_1 + F_0)q - F_0, \quad (25)$$

where F_0, F_1 are losses / gains in profit from removing the drug from the market. If $F_1 \geq 0$ then the agent prefers the option to be exercised when $q = 1$.

Over time exogenous news arrive via stochastic process X :

$$dX_t = \xi dt + \frac{1}{\phi} dB_t, \quad (26)$$

where $B = (B_t)_{t \geq 0}$ is a Standard Brownian Motion. The interpretation of the news process is that it reflects information from patients and medical care providers about the side effects of the product. The public news process is gradual because it takes time to learn about all effects of a drug and we assume outcomes of any particular patient do not change beliefs discontinuously. As in Section 3 define q_P and q_A the autarky exercise thresholds by the principal and the agent respectively.

7.2 Perfect Information

The agent (drug producer) can conduct studies and learn additional information about ξ . The following lemma characterizes the equilibrium outcomes when the principal's and agent's objectives are positively collinear.

Proposition 5. *Suppose the agent can conduct a perfectly informative test about ξ . If $F_1 \geq 0$ the unique Markov Perfect Equilibrium is equivalent to a fully informative test about ξ conducted at $t = 0$. Principal acts only when he knows that $\xi = 1$.*

The intuition for the above proposition is two-fold. If $F_0 > 0$ the agent does not wish to exercise the option if $\xi = 0$, but wishes to exercise it if $\xi = 1$. Under complete information principal's and agent's preferences agree and hence the agent can implement first best by conducting a fully informative test at $t = 0$. If $F_0 < 0$, then the agent would like to exercise the option for both $\xi = 0$ and $\xi = 1$. In this case immediate and complete information provision is a result of a dynamic unraveling similar to Coase (1972). Suppose for a

³⁶This is without loss of generality and we choose this parametrization to fit the FDA application closer.

moment that the principal follows the autarky threshold strategy: act if $q_t \geq q_P$ and wait otherwise. The agent's best response is to send beliefs either to 0 or to the decision threshold of the principal q_P . In order for this to constitute an equilibrium, the principal's decision to exercise the option at $q = q_P$ must remain optimal, even though for $q < q_P$ he expects to receive additional information from the agent. If the principal deviates and waits dt he observes innovation dX_t . If dX_t is positive, then delayed exercise would result in a *second order loss* compared to immediate exercise, since q_P was optimally chosen in the first place. However, if dX_t is negative, the agent would be tempted to persuade her again. Such persuasion results in a *first order gain* compared to immediate exercise, since persuasion is done in the region of strict convexity of V_P and the principal does not end up exercising the option when $\xi = 0$. Hence it cannot be an equilibrium where principal acts at a threshold q_P (and similarly at any other candidate stopping threshold $\bar{q} < 1$).

Consider the judge example from Kamenica and Gentzkow (2011) with the following preferences: the agent gets 1 if the option is exercised (e.g. the defendant is convicted) and 0 otherwise, while the principal gets 1 if the option exercised in state $\xi = 1$ (e.g. the defendant is guilty) and -1 if the option is exercised in the state $\xi = 0$ (e.g. the defendant is innocent). This setup corresponds to utility functions

$$v_A(t, q) = e^{-rt}, \quad v_P(t, q) = e^{-rt}(2q - 1).$$

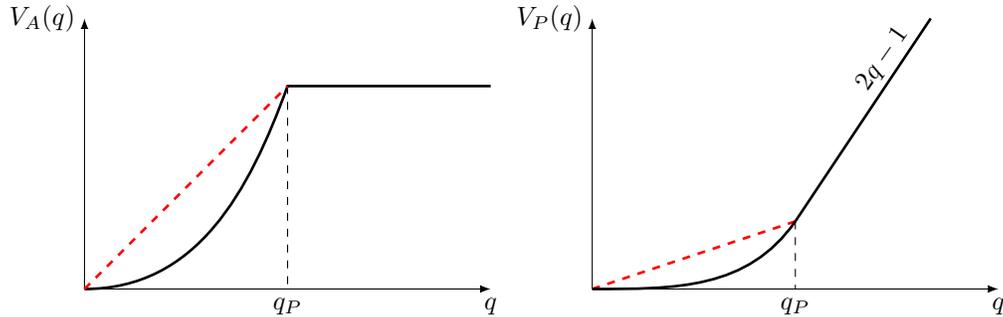


Figure 6: Effect of concavification on principal's value function, $v_A(t, q) = e^{-rt}$, $v_P(t, q) = e^{-rt}(2q - 1)$.

Figure 6 provides a graphical illustration of the core idea behind the unraveling argument in this case. It shows the effect of one-step concavification of the agent's value function that pools information about $\xi = 1$. As a result, principal's value function, that previously satisfied a smooth pasting condition at q_P , obtains a kink. This kink in presence of Brownian increments dX_t makes it optimal for the principal to wait and pushes the exercise threshold higher than q_P . Technically speaking, persuasion using a pooling message implies that the value function of the principal has to satisfy both the "reflecting"³⁷ and smooth pasting³⁸ boundary conditions at the optimal exercise threshold.

In stark contrast to the static solution of Kamenica and Gentzkow (2011), in which the agent is able to

³⁷See Harrison (2013) for details.

³⁸See Dixit and Pindyck (2012) for details.

increase the chances of conviction by designing a test that pools information, the unique equilibrium in our environment is a fully informative test at time 0. The option of the principal to wait for exogenous news benefits her directly, since the autarky action threshold q_P is higher than the static one $1/2$, and more importantly it benefits her indirectly - by incentivizing the agent to produce information earlier, which leads to complete unraveling in equilibrium. Proposition 5 implies that the ability to persuade itself hurts the agent for initial beliefs $q \in [q_P, 1]$ where the principal would have acted immediately absent endogenous information.

7.3 Imperfect Information

It may be the case that the scope of the information that the agent can obtain by a binary random variable $\theta \in \{0, 1\}$ is not perfect and is only correlated with ξ . We interpret θ as the information contained in the most informative test that the agent can conduct, which might not perfectly reveal ξ . In the FDA example this could be due to the presence of long term effects which can not be detected during a short medical test. Both parties have a common prior $y_{0-} = P(\theta = 1)$ and the informativeness of θ is given by

$$z_{0-}^i = P(\xi = 1 | \theta = i), \quad i \in \{0, 1\}.$$

Similar to Section 2 we allow the agent to conduct arbitrary tests on θ every instance of time, but do not allow her to commit to future tests.

After an arbitrary history \mathbb{H}^t the evolution of posteriors in such model can be summarized by two values:

$$\begin{cases} x_t = \phi^2 \left(X_t - \frac{t}{2} \right) \\ y_{t-} = P(\theta = 1 | \mathbb{H}^t) \end{cases}$$

since the posterior $q_{t-} = P(\xi = 1 | \mathbb{H}^t)$ can be decomposed as

$$q_{t-} = q(x_t, y_{t-}) = y_{t-} \cdot z^1(x_t) + (1 - y_{t-}) \cdot z^0(x_t)$$

where $z^i(x) = P(\xi = 1 | \theta = i, \mathbb{H}^t) = P(\xi = 1 | \theta = i, x_t = x)$ have closed form solutions:

$$z^i(x) = \frac{z_{0-}^i \cdot e^x}{z_{0-}^i \cdot e^x + 1 - z_{0-}^i} \quad i = 0, 1$$

Notice that in the Wald option setting $z^1(x)$ and $z^0(x)$ are the highest and the lowest value of payoff relevant belief $q(x, y)$ that the agent can induce through persuasion over y conditional on exogenous level of x . This corresponds to $\theta_H x$ and $\theta_L x$ respectively in the setup of Section 2.

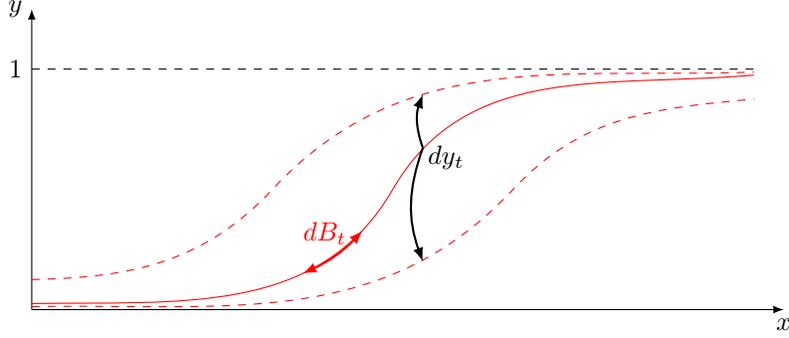


Figure 7: Evolution of beliefs in (x, y) the space due to public news dx_t and persuasion dy_t .

Autarky thresholds. Absent additional information from the agent, exogenous variation in x_t translates one-to-one into the variation in $q_t = q(y_{0-}, x_t)$ and gives rise to a standard real option problem. The autarky solution for agent $i \in \{A, P\}$ is characterized by a threshold q_i such that when $q_t \leq q_i$ the agent i prefers to wait, and when $q_t > q_i$ the agent i would like the option to be exercised. The autarky value function f_i^{NI} with $i \in \{A, P\}$, which we for convenience write as a function of q , satisfies a second order ODE with boundary conditions that pin the optimal exercise threshold q_i :

$$r f_i^{NI}(q) = \frac{1}{2} \phi^2 q^2 (1-q)^2 \frac{\partial^2}{\partial q^2} f_i^{NI}(q), \quad (27)$$

$$f_i^{NI}(0) = 0, \quad (28)$$

$$f_i^{NI}(q_i) = v_i(q_i), \quad (29)$$

$$\frac{d}{dq} f_i^{NI}(q_i) = \frac{d}{dq} v_i(q_i). \quad (30)$$

Equations (27) - (30) are the adaptations of equations (12) - (15) to the Wald option setting. By defining $V_i^{NI}(x, y) = f_i(q(x, y))$ and $x_i^w(y)$ as the unique solution of

$$q(x_i^w(y), y) = q_i \quad \text{for } i \in \{A, P\} \quad (31)$$

we obtain the primitives of the model similar to those of Section 3 that serve as building blocks for the equilibrium. From here one can proceed in a manner similar to Sections 4 and 5 and construct equilibria that are analogous to those of Propositions 1 and 3.³⁹

Persuasion to Wait. Suppose the principal is more impatient: $q_P < q_A$. Analogous to Section 4 we make the observation that at $x_A^w(1)$ the agent discloses all available information. Iterated elimination of dominated strategies results in a threshold $l(x)$ such that principal rationally acts whenever y is sufficiently large. The equilibrium structure is depicted in Figure 7.

³⁹For more details see an earlier draft of this paper "Persuading the Regulator to Wait".

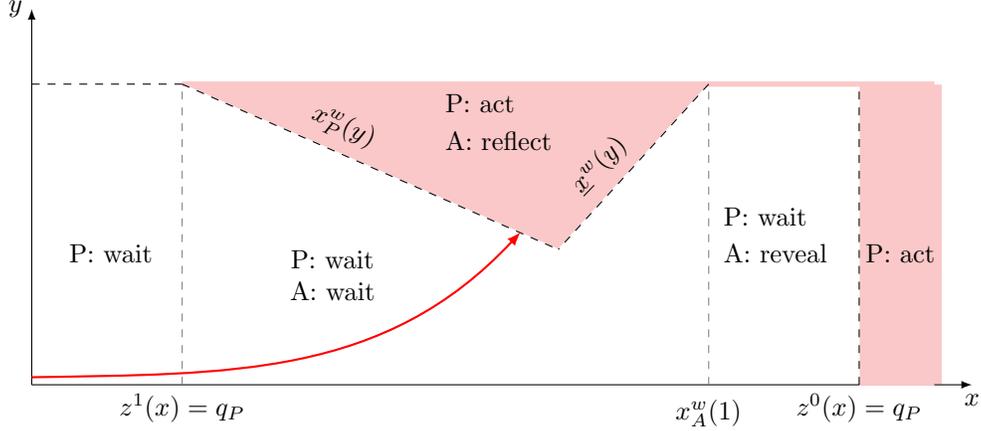


Figure 8: Equilibrium strategies for the principal (P) and the agent (A) for wald options.

If initial belief y_{0-} is sufficiently high, both parties wait until the exogenous news push the beliefs up to principal's autarky threshold $x_P^w(y)$. In order to minimize the probability of instantaneous action, the agent communicates a partially informative message which either reveals $\theta = \theta_H$ with an endogenous intensity, or pools the remaining states. This causes the belief trajectory (x_t, y_t) reflect against the principal's autarky threshold $x_P^w(y)$. This belief dynamic is akin to the one shown in Figure 4. Over time, if $\theta = \theta_H$ is not revealed via gradual persuasion, the principal becomes sufficiently pessimistic about θ (lower y) and she becomes willing to wait for the agent to reveal θ perfectly at \bar{x}^w .

8 Discussion of Dynamics and Commitment

8.1 Static Persuasion

Suppose the agent can persuade the principal only once at $t = 0$. This is the static counterpart of the game analyzed in Sections 2-7. Solving for static persuasion in the presence of a real option subgame amounts to computing the agent's no-persuasion value function and concavifying it subject to a constraint that the posterior expectation of θ is constrained by the support of the distribution

$$\theta_L \leq \mathbb{E}[\theta] \leq \theta_H$$

Lemma 6. *Suppose $I_A \geq I_P$. The agent conducts a perfectly informative test at $t = 0$. If $I_A < I_P$, then two cases are possible:*

- $V_A^{NI}(x, 0) \geq \theta_L x - I_A$. The agent fully reveals θ at $t = 0$.
- $V_A^{NI}(x, 0) < \theta_L x - I_A$. The agent communicates a partially informative signal such which makes the principal either act on pooled information, or reveals that $\theta = \theta_L$.

Lemma 6 highlights the differences between static and dynamic persuasion. In the case of *persuasion to wait* the agent eventually reveals θ in both static and dynamic cases. The latter generates more value for the agent since he can delay option exercise more efficiently by pooling information over time. In the case of *persuasion to act* the inability to communicate in the future serves as a commitment device for the agent to stop persuading: he pools information at the principal’s action threshold. Static persuasion generates more value for the agent compared to the dynamic game.

8.2 Agent’s commitment

Persuasion to Wait. The equilibrium in case of “Persuasion to Wait” features a compromise region. Namely, as soon as X_t is sufficiently high (above $x_A(1)$) the agent conducts a test that reveals the state θ to both parties. Moreover, when the conflict of interest between the parties is small, i.e. I_P is close to I_A , and the prior on θ is sufficiently low, full revelation of θ at $x_A(1)$ is the only information shared along the equilibrium path. Notice that in this case dynamic commitment has no value to the agent. The intuition behind result is similar to the one of Lemma 2: if $\theta = \theta_H$ the equilibrium delivers the first best outcome for the agent. The only room to improve the agent’s payoff is through delaying the exercise time when $\theta = \theta_L$. However, that is impossible since it is dominant for the principal to exercise the option immediately for $x > x_P(0)$ regardless of the strategy of the agent. Commitment to within period tests can also be relaxed. In particular, the outcome of the equilibrium remains unchanged if one considers a disclosure (in which the agent *learns* θ perfectly at time 0 and decides when to credibly disclose it) or even a cheap talk (in which the agent *learns* θ perfectly at time 0 and but has no means to credibly communicate it) setting. This is a consequence of perfect misalignment of incentives between the θ_H and θ_L type agents when X is sufficiently large: at $x_A(1)$ the agent with $\theta = \theta_H$ effectively separates from the low type agent with $\theta = \theta_L$ in order to induce immediate action.

Dynamic commitment might be valuable when the conflict of interests between the parties is sufficiently large. For example, consider the case in which in addition to having higher I_A the agent has a slightly higher discount rate $r_A > r_P$ that does not change the direction of the conflict. Instead of pipeting information continuously at $x_P(y)$ the agent is better off by committing to disclose a discrete chunk of information in the future. Such commitment incentivizes the principal to wait via higher continuation value waiting and is cheaper than instantaneous information provision due to differences in discount rates. The exact shape of the optimal information sharing policy in a general setup is hard to characterize and we leave it to future research.

Persuasion to Act. The equilibrium of the “persuasion to act” case features unraveling and the outcome is a fully informative test conducted at time 0. Incentives of the agent to provide information in order to induce instantaneous action backfire due to the increase of the option value of the principal to wait for

more information in the future. Dynamic commitment allows the agent to disentangle instantaneous effect of information provision from the continuation value side-effect for the principal. In particular, the agent benefits from even a very weak form dynamic commitment – commitment to stop providing information after a certain date. The optimal policy in this class features communication only at $t = 0$ and employs information pooling in order to speed up and maximize the probability of the action taken by the principal in a manner similar to Kamenica and Gentzkow (2011): at time 0 the agent conducts a test that either reveals $\theta = \theta_L$ or increases the belief of the principal up to the autarky threshold $y = x_P^{-1}(X_0)$, conditional on the increase of beliefs the option is exercised immediately. This policy is optimal in a broader class of persuasion strategies with dynamic commitment for certain parameters of our model, and we leave the full characterization of the optimal policy to future research.

9 Conclusion

We present a theory of communication with hard information in the context of real options. Even though the principal has full authority over the exercise of the real option, when the conflict of interest between the parties is large, she can do no better than his outside option. The agent is then able to extract all of the rents associated with his access to better information. When the conflict of interest is small, we observe state-contingent delegation. The agent is allowed to exercise the real option at his optimal threshold when his information is good, and the principal takes back control when the agent’s information is bad. Importantly, the correct measure of conflict of interest is the net value of the agent’s bias and the value of information he has access to. In equilibrium, the value of agent’s information changes and the principal may start by keeping all authority, but then relinquish it to the agent as the latter’s information value endogenously increases.

We show the equilibrium outcome can be sustained when the agent is privately informed (Section 6) and is robust to alternative specifications of the real options (Section 7). We believe the equilibrium structure presented in Section 4 is important for understanding implications of strategic communication of hard information for decision making and delegation of authority in organizations.

When the agent prefers earlier exercise his inability to stop persuading undermines him. There always exists an equilibrium in which the agent discloses all information at time zero, and in a large set of parameters this equilibrium is unique. In sharp contrast to static environments, pooling of information unravels due to the option of the principal to wait and obtain more information from the agent in the immediate future. Moreover, there is always an initial prior over agent’s information, for which the agent does strictly worse, than if he could not communicate at all. The robust economic prediction in is that the principal acts at a strictly higher confidence level in the project, than in the case of static communication.

References

- Viral V. Acharya, Peter DeMarzo, and Ilan Kremer. Endogenous information flows and the clustering of announcements. *The American Economic Review*, 101(7):2955–2979, 2011.
- 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. *Mimeo*, 2016.
- Patrick Bolton and Mathias Dewatripont. Authority in organizations. *Handbook of Organizational Economics*, pages 342–372, 2013.
- Ronald H Coase. Durability and monopoly. *The Journal of Law and Economics*, 15(1):143–149, 1972.
- Brendan Daley and Brett Green. Waiting for news in the market for lemons. *Econometrica*, 80(4):1433–1504, 2012.
- 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.
- 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.
- Sanford J Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483, 1981.
- Sebastian Gryglewicz and Barney Hartman-Glaser. Dynamic agency and real options. *mimeo*, 2017.
- Faruk Gul, Hugo Sonnenschein, and Robert Wilson. Foundations of dynamic monopoly and the coase conjecture. *Journal of Economic Theory*, 39(1):155–190, 1986.
- Yingni Guo. Dynamic delegation of experimentation. *The American Economic Review*, 106(8):1969–2008, 2016.
- Ilan Guttman, Ilan Kremer, and Andrzej Skrzypacz. Not only what but also when: A theory of dynamic voluntary disclosure. *The American Economic Review*, 104(8):2400–2420, 2014.
- J Michael Harrison. *Brownian Models of Performance and Control*. Cambridge University Press, 2013.

- 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.
- Ilan Kremer and Andrzej Skrzypacz. Dynamic signaling and market breakdown. *Journal of Economic Theory*, 133(1):58–82, 2007.
- Thomas Kruse and Philipp Strack. Optimal stopping with private information. *Journal of Economic Theory*, 159:702–727, 2015.
- George J Mailath and Larry Samuelson. *Repeated games and reputations*, volume 2. Oxford university press Oxford, 2006.
- Robert McDonald and Daniel Siegel. The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–727, 1986.
- Paul R Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.
- David Pollard. *A user’s guide to measure theoretic probability*, volume 8. Cambridge University Press, 2002.
- Jerome Renault, Eilon Solan, and Nicolas Vieille. Optimal dynamic information provision. *arXiv preprint arXiv:1407.5649*, 2014.
- Yuliy Sannikov. Games with imperfectly observable actions in continuous time. *Econometrica*, 75(5):1285–1329, 2007.
- Leo Simon and Maxwell Stinchcombe. Extensive form games in continuous time: Pure strategies. *Econometrica*, 57(5):1171–1214, 1989.
- Nancy L Stokey. *The Economics of Inaction: Stochastic Control models with fixed costs*. Princeton University Press, 2008.
- Abraham Wald. *Sequential analysis*. Courier Corporation, 1973.

Appendix (Supplementary Material)

Proof of Lemma 1 . Suppose $V_A(x, y)$ corresponds to agent's equilibrium value function. First, it must satisfy

$$V_A(x, \cdot) = \text{cav}_\mu [V_A(x, \mu)] \quad (\text{A.1})$$

where x is taken as constant and the concavification operator is applied only with respect to y . If (A.1) did not hold for any y , then the agent can improve his value by sending a single informative signal and then returning to the original strategy. This proves global concavity in y , and by extension, continuity.

Second, the agent is free to choose when to send the first message:

$$V_A(x, y) = \sup_{\kappa} \mathbf{E}_{(x, y)} \left[e^{-r \cdot \tau \wedge \kappa} \cdot V_A(x_{\tau \wedge \kappa}, y) \right] \quad (\text{A.2})$$

where

$$\tau = \inf\{t : (x_t, y) \in \mathbf{T}^*\} \quad (\text{A.3})$$

and for $t < \tau \wedge \kappa$ the agent passively waits. Hence the belief about θ does not change between 0 and $\tau \wedge \kappa$. Moreover because $\kappa = 0$ achieves the sup in (A.2), we have

$$V_A(x, y) = \max_{\kappa} \mathbf{E} \left[e^{-r \cdot \tau \wedge \kappa} \cdot V_A(x_{\tau \wedge \kappa}, y) \mid (x_0, y_0) = (x, y) \right] \quad (\text{A.4})$$

A function of single-variable $V_A(\cdot, y)$ is continuous in x as a solution of the optimal stopping problem (A.4). If it were not continuous, the could wait in the vicinity of the jump. This contradicts the optimality of the candidate best response.

Lemma 7 (Auxiliary lemma.). *Suppose function of a single variable $f(y)$ is concave. Consider the set E of extreme points of $f(\cdot)$ as defined in (18). Then $f(y)$ is strictly concave in the interior of E .*

Proof. Function f is concave:

$$f(\alpha y_1 + (1 - \alpha)y_2) \geq \alpha f(y_1) + (1 - \alpha)f(y_2)$$

Since elements in the interior of E do not lie on a straight line it implies

$$f(\alpha y_1 + (1 - \alpha)y_2) > \alpha f(y_1) + (1 - \alpha)f(y_2)$$

End of Lemma 7 proof. □

Suppose (x, y) lies in the interior of $E(x)$. Due to Lemma 7 the agent does not persuade at (x, y) . Suppose there exists a ε such that $B_{\varepsilon_1}(y) \subset E$ and $\forall (x, y) \in B_{\varepsilon}(x, y)$ function $V_A(x, y)$ is strictly concave in y .

Define

$$\kappa = \inf \left\{ t : x_t \notin (x - \varepsilon, x + \varepsilon) \right\} > 0$$

Then for $t < \kappa$ neither the receiver acts, nor the agent persuades. Hence

$$V_A(x, y) = \mathbb{E} \left[e^{-r\kappa} \cdot V_A(x_\kappa, y_\kappa) \mid (x_0, y_0) = (x, y) \right]$$

Define

$$\hat{V}_t = \mathbb{E}_t \left[e^{-r(\kappa-t)} \cdot V_A(x_\kappa, y_\kappa) \mid (x_0, y_0) = (x, y) \right]$$

where the expectation is taken wrt the evolution of X up to κ , i.e. $\hat{\mathbb{F}} = \{\mathcal{F}_{t \wedge \kappa}^X\}_{t=0}^\infty$. The process \hat{V} is a Levy martingale which implies $\mathbb{E} [d\hat{V}_t] = 0$ if it admits an Ito decomposition. Moreover because the agent is not sharing information in that region $\hat{V}_t = V(x_t, y)$. It follows that $\mathbb{E} [dV_A(x_t, y)] = 0$. Applying Ito's lemma to $e^{-rt}V_S(x_t, y)$ we obtain

$$rV_A(x, y) = \mu x \frac{\partial}{\partial x} V_A(x, y) + \frac{1}{2} x^2 \phi^2 \frac{\partial^2}{\partial x^2} V_A(x, y).$$

End of Lemma 1 proof. □

Proof of Lemma 2. If $\theta(\omega) = \theta_H$, then disclosing it results in immediate exercise of the option by the principal, which is the first best outcome for the agent conditional on $\theta(\omega) = \theta_H$. Any delay in this case decreases sender's pay-off.

If $\theta(\omega) = \theta_L$, let Y_τ denote the equilibrium level of beliefs about θ at the moment of option exercise. It must be that $X_\tau \leq x_P(0)$. Since $x_P(0) < x_A(0)$ the sender strictly prefers to reveal $\theta(\omega) = \theta_L$ to guarantee exercise at $x_P(0)$ and not at X_τ . □

Proof of Lemma 3. Define y^* as a solution of an equation

$$\left(\frac{x_P(y^*)}{x_A(1)} \right)^{\beta_1} \left(y^*(\theta_H x_A(1) - I_P) + (1 - y^*)V_P^{NI}(x_A(1), 0) \right) = (y^*\theta_H + (1 - y^*)\theta_L)x_P(y^*) - I_P \quad (\text{A.5})$$

Rewrite the LHS of (A.5):

$$\left(\frac{x_P(1)}{x_A(1)\mathbb{E}[\theta]} \right)^{\beta_1} \left(y(\theta_H x_A(1) - I_P) + (1 - y)(V_P^{NI}(x_A(1), 0) - \theta_L x_A(1) + I_P) \right) \quad (\text{A.6})$$

$$= \left(\frac{x_P(1)}{x_A(1)\mathbb{E}[\theta]} \right)^{\beta_1} \left(x_A(1)\mathbb{E}[\theta] - I_P + (1 - y)(V_P^{NI}(x_A(1), 0) - \theta_L x_A(1) + I_P) \right) \quad (\text{A.7})$$

$$= \left(\frac{x_P(1)}{x_A(1)\mathbb{E}[\theta]} \right)^{\beta_1} \left(x_A(1)\mathbb{E}[\theta] - I_P \right) + \left(\frac{x_P(1)}{x_A(1)\mathbb{E}[\theta]} \right)^{\beta_1} (1 - y) \left(V_P^{NI}(x_A(1), 0) - \theta_L x_A(1) + I_P \right) \quad (\text{A.8})$$

The derivative of the LHS with respect to y is given by

$$\underbrace{(\theta_H - \theta_L) \frac{\partial}{\partial \mathbb{E}[\theta]} \left(\left(\frac{x_P(1)}{x_A(1) \mathbb{E}[\theta]} \right)^{\beta_1} (x_A(1) \mathbb{E}[\theta] - I_P) \right)}_{(A)} \quad (\text{A.9})$$

$$- \underbrace{\left(\frac{x_P(1)}{x_A(1) \mathbb{E}[\theta]} \right)^{\beta_1} \left(V_P^{NI}(x_A(1), 0) - \theta_L x_A(1) + I_P \right)}_{(B)} \quad (\text{A.10})$$

$$- \beta_1 \underbrace{\left(\frac{x_P(1)}{x_A(1) \mathbb{E}[\theta]} \right) \frac{\theta_H - \theta_L}{\mathbb{E}[\theta]} (1 - y) \left(V_P^{NI}(x_A(1), 0) - \theta_L x_A(1) + I_P \right)}_{(C)} \quad (\text{A.11})$$

Term A is negative since it would have been optimal to exercise the option at $x_A(1) \mathbb{E}[\theta] > \frac{\beta_1 I_P}{\beta_1 - 1}$. Term (B) is negative since it's better for the principal to wait when at $x_A(1)$ if $\theta = \theta_L$. Term (C) is clearly negative.

The RHS of (A.5) can be rewritten as

$$(y^* \theta_H + (1 - y^*) \theta_L) x_P(y^*) - I_P = \frac{I_P}{\beta_1 - 1} \quad (\text{A.12})$$

and is thus constant. So the LHS is monotonically decreasing in y , while the RHS is a constant. This implies that $y < y^*$ we have LHS being greater than the RHS and the principal is willing to wait. If $y > y^*$, then the LHS is smaller than the RHS and the principal is willing to act immediately.

Case $y \leq y^*$. The principal waits for additional information released at $x_A(1)$ and hence

$$V_A^I(x, y) = \left(\frac{x_P(y^*)}{x_A(1)} \right)^{\beta_1} \left(y^* (\theta_H x_A(1) - I_P) + (1 - y^*) V_P^{NI}(x_A(1), 0) \right) \quad (\text{A.13})$$

Case $y > y^*$. If $x \leq x_P(y)$ it must be the case that

$$V_P^I(x, y) = V_P^{NI}(x, y) \quad (\text{A.14})$$

since the principal optimally stops at her autarky threshold without waiting for additional information.

Suppose $x > x_P(y)$. Consider an auxiliary function $g(x; \underline{x})$ that solves:

$$\begin{aligned} r g &= \mu x g' + \frac{1}{2} \sigma^2 x^2 g'' \\ g(\underline{x}) &= (y \theta_H + (1 - y) \theta_L) \underline{x} - I_P = a \underline{x} - b \\ g'(\underline{x}) &= y \theta_H + (1 - y) \theta_L = a \end{aligned}$$

Function g can be found explicitly as $g(x) = C_1x^{\beta_1} + C_2x^{\beta_2}$. We can solve for the constants

$$\begin{cases} C_1\underline{x}^{\beta_1} + C_2\underline{x}^{\beta_2} = ax - b \\ C_1\beta_1\underline{x}^{\beta_1-1} + C_2\beta_2\underline{x}^{\beta_2-1} = a \end{cases}$$

$$\begin{cases} C_1\beta_1\underline{x}^{\beta_1} - \beta_2C_1\underline{x}^{\beta_2} = ax - \beta_2(ax - b) \\ C_2\beta_1\underline{x}^{\beta_2} - \beta_2C_2\underline{x}^{\beta_2} = \beta_1(ax - b) - ax \end{cases}$$

$$\begin{cases} C_1 = \frac{ax(1-\beta_2)+\beta_2b}{(\beta_1-\beta_2)\underline{x}^{\beta_1}} \\ C_2 = \frac{(\beta_1-1)ax-b\beta_1}{(\beta_1-\beta_2)\underline{x}^{\beta_2}} \end{cases}$$

Notice that when $\underline{x} = x_P(y)$, then $C_1 > 0$ and $C_2 = 0$. Thus for $\underline{x} > x_P(y)$ we have $C_1 > 0$ and $C_2 > 0$, which implies that $g'' > 0$.

It must be the case that for $y < y^*$:

$$g(x_A(1), x_P(y)) > y(\theta_H x_A(1) - I_P) + (1-y)V_P^{NI}(x_A(1), 0). \quad (\text{A.15})$$

Otherwise the Receiver would prefer to always (for all x) wait for the information shared at $x_S(1)$, which contradicts $y < y^*$.

Find \underline{x} as the root of an equation

$$g(x_A(1), \underline{x}) = y(\theta_H x_A(1) - I_P) + (1-y)V_P^{NI}(x_A(1), 0). \quad (\text{A.16})$$

This equation has at least one solution, since

$$g(x_A(1), x_A(1)) = (y\theta_H + (1-y)\theta_L)x_A(1) - I_P < y(\theta_H x_A(1) - I_P) + (1-y)V_P^{NI}(x_A(1), 0). \quad (\text{A.17})$$

By continuity in of $g(x, \cdot)$ in the second variabel there is a solution. Uniqueness can be seen from direct calculations below

$$\begin{aligned} (\beta_1 - \beta_2) \frac{\partial}{\partial \underline{x}} g(x; \underline{x}) &= x^{\beta_1} \cdot \frac{a(1-\beta_2)(1-\beta_1) - \beta_1\beta_2b\underline{x}^{-1}}{\underline{x}^{\beta_1}} + x^{\beta_2} \cdot \frac{a(\beta_1-1)(1-\beta_2) + b\beta_1\beta_2\underline{x}^{-1}}{\underline{x}^{\beta_2}} \\ &= \left[\left(\frac{x}{\underline{x}}\right)^{\beta_1} - \left(\frac{x}{\underline{x}}\right)^{\beta_2} \right] \cdot [a(1-\beta_2)(1-\beta_1) - \beta_1\beta_2b\underline{x}^{-1}] \end{aligned}$$

For $x > \underline{x} > x_R(y)$ the first term is clearly positive since $\beta_1 > 1 > 0 > \beta_2$. The second term can be signed

as follows:

$$\begin{aligned}
\underline{x}a(1 - \beta_2)(1 - \beta_1) - \beta_1\beta_2b &< \frac{\beta_1}{\beta_1 - 1} \frac{b}{a} a(1 - \beta_2)(1 - \beta_1) - \beta_1\beta_2b \\
&= \beta_1b(\beta_2 - 1) - \beta_1\beta_2b \\
&= -\beta_1\beta_2b \\
&< 0
\end{aligned}$$

Denote the solution of the equation above as the function of (currently a parameter) y as $\underline{x}(y)$. Finally we can write $V_P(x, y) = g(x, \underline{x}(y))$ for $x \in [\underline{x}(y), x_A(1)]$. The resulting function is convex in x and is strictly convex outside of the linear region between $x_P(y)$ and $\underline{x}(y)$ as shown in Figure 3. \square

Lemma 8. *Threshold function $\underline{x}(y)$ is increasing in y .*

Proof. Consider $g(x, \underline{x})$ and take a derivative with respect to y which determines the boundary conditions:

$$\frac{\partial}{\partial y} g(x, \underline{x}) = \frac{\underline{x}(\theta_H - \theta_L)}{\beta_1 - \beta_2} \left[(1 - \beta_2) \left(\frac{x}{\underline{x}} \right)^{\beta_1} + (\beta_1 - 1) \left(\frac{x}{\underline{x}} \right)^{\beta_2} \right]$$

If we keep \underline{x} constant and simply increase y in the boundary conditions, then at $x_A(1)$ the value of $g(x)$ goes up by more than $(\theta_H - \theta_L)x_A(1)$ which is higher than the increase in the right hand boundary condition $\theta_H x_A(1) - I_P - V_P^{NI}(x_A(1), 0)$. Hence, in order to match the boundary condition \underline{x} should increase.

The missing comparison follows from

$$\begin{aligned}
\frac{\underline{x}(\theta_H - \theta_L)}{\beta_1 - \beta_2} \left[(1 - \beta_2) \left(\frac{x}{\underline{x}} \right)^{\beta_1} + (\beta_1 - 1) \left(\frac{x}{\underline{x}} \right)^{\beta_2} \right] &vs. (\theta_H - \theta_L)x \\
(1 - \beta_2)\alpha^{\beta_1} + (\beta_1 - 1)\alpha^{\beta_2} &vs. (\beta_1 - \beta_2)\alpha
\end{aligned}$$

for $\alpha = x/\underline{x}$.

Define $h(\alpha) = (1 - \beta_2)\alpha^{\beta_1} + (\beta_1 - 1)\alpha^{\beta_2} - (\beta_1 - \beta_2)\alpha$. Notice that $h(1) = h'(1) = 0$ and

$$h''(\alpha) = (1 - \beta_2)\beta_1(\beta_1 - 1)\alpha^{\beta_1 - 2} + (\beta_1 - 1)\beta_2(\beta_2 - 1)\alpha^{\beta_2 - 2} > 0.$$

Thus $h(\alpha) > 0$ for all $\alpha > 1$, implying that $\frac{\partial}{\partial y} g(x, \underline{x}) > (\theta_H - \theta_L)x$ for all $x > \underline{x}$. \square

Lemma 9. *There exists a increasing process $M = \{M_t\}_{t \geq 0}$ measurable with respect to \mathcal{F}^X such that*

- process $(X_t, y_0 - M_t)$ lies outside the principal's action region:

$$X_t \leq x_P(y_0 - M_t).$$

- Process M is constant outside principal's action region:

$$dM_t = 0 \quad \Leftrightarrow \quad X_t < x_P(y_0 - M_t)$$

- for any other increasing process $\hat{M} = \{\hat{M}_t\}_{t \geq 0}$ satisfying the above conditions

$$M_t \leq \hat{M}_t \quad \forall \quad t \geq 0.$$

Proof. Consider the case when $X_0 < x_P(Y_{0-})$ and recall that (x^*, y^*) is the unique solution of $x^* = x_P(y^*) = \underline{x}(y^*)$. Define process $L = \{L_t\}_{t \geq 0}$ as

$$L_t = \min \left[\max_{s < t} X_s, x^* \right],$$

i.e. L_t measures how far into the action region of the principal, the process X_t would have made absent persuasion and action. Next, construct the reflected process $\hat{Y} = \{\hat{Y}_t\}_{t \geq 0}$ as follows:

$$\hat{Y}_t = \begin{cases} Y_{0-}, & \text{if } L_t < x_P(Y_{0-}) \\ x_P^{-1}(L_t), & \text{if } L_t \geq x_P(Y_{0-}) \end{cases}$$

Process \hat{Y}_t is weakly decreasing and we can define

$$M_t = Y_0 - \hat{Y}_t$$

Notice that the process $M = \{M_t\}_{t \geq 0}$ always satisfies

$$X_t \leq x_P(Y_0 - M_t).$$

Moreover, $dM_t > 0$ if and only if $X_t = x_P(Y_0 - M_t)$, thus, it is the unique process with these properties. □

Proof of Proposition 1. First, we show that under the described strategy of the agent, principal's best response stopping set \mathbf{T} remains optimal. To do so, we show is that when princial's beliefs reflect against either $x_P(y)$ or $\underline{x}(y)$, her value function does not change. In order to do this, we need to make sure that the smooth pasting condition for the value $V_P(x, y)$ is satisfied in the y dimension as well.

Consider the boundary $\underline{x}(y)$ (the other boundary $x_P(y)$ can be dealt this analogously). Above $\underline{x}(y)$ we have

$$V_P(x, y) = y(\theta_H \underline{x}(y) - I_P) + (1 - y)(\theta_L \underline{x}(y) - I_P)$$

hence

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

Below $\underline{x}(y)$ function V_P is equal to g evaluated at at the appropriate boundary. Thus

$$\begin{aligned} \lim_{\varepsilon \uparrow 0} \frac{\partial}{\partial y} V_P(x, y + \varepsilon) \Big|_{x=\underline{x}(y)} &= \frac{d}{dy} g(x, \underline{x}(y)) \Big|_{x=\underline{x}(y)} \\ &= \frac{\partial}{\partial y} g(x, \underline{x}) \Big|_{x=\underline{x}=\underline{x}(y)} \\ &= \frac{\underline{x}(\theta_H - \theta_L)}{\beta_1 - \beta_2} \left[(1 - \beta_2) \left(\frac{x}{\underline{x}} \right)^{\beta_1} + (\beta_1 - 1) \left(\frac{x}{\underline{x}} \right)^{\beta_2} \right] \Big|_{x=\underline{x}=\underline{x}(y)} \\ &= (\theta_H - \theta_L) \underline{x}(y). \end{aligned}$$

The second equality holds due to the Envelope theorem. We see that the original value function satisfies the smoothpasting condition of the value function with reflected belief process, thus, reflection does not generate any value for the principal.

Agent's equilibrium value function. Denote by $V_A(x, y)$ the agent's value function corresponding to the proposed strategy of the agent. For $x < x_P(y)$ and $y > y^*$ it takes the form

$$V_A(x, y) = c(y) \cdot x^\beta \tag{A.18}$$

where $\beta > 1$ is the positive characteristic of the agent's waiting HJB. The boundary conditions at $x_P(y) = \frac{\beta}{\beta-1} \cdot \frac{I_P}{\theta_L + y(\theta_H - \theta_L)}$

$$c'(y) x_P(y)^\beta = \frac{\theta_H x_P(y) - I_A - c(y) x_P(y)^\beta}{1 - y} \tag{A.19}$$

$$c'(y) = -\frac{c(y)}{1 - y} + \frac{\theta_H x_P(y)^{1-\beta} - I_A x_P(y)^{-\beta}}{1 - y} \tag{A.20}$$

$$-c'(y) = \frac{c(y)}{1 - y} - \frac{\theta_H x_P(y)^{1-\beta} - I_A x_P(y)^{-\beta}}{1 - y} \tag{A.21}$$

Second derivative is

$$c''(y) = -\frac{c(y)}{(1-y)^2} - \frac{c'(y)}{1-y} + \frac{(1-\beta)\theta_H x_P(y)^{-\beta} x'_P(y) + I_A \beta x_P(y)^{-\beta-1} x'_P(y)}{1-y} \quad (\text{A.22})$$

$$+ \frac{\theta_H x_P(y)^{1-\beta} - I_A x_P(y)^{-\beta}}{(1-y)^2} \quad (\text{A.23})$$

$$= x'_P(y) \cdot \frac{(1-\beta)\theta_H x_P(y)^{-\beta} + I_A \beta x_P(y)^{-\beta-1}}{1-y} \quad (\text{A.24})$$

Substitute the term for $c'(y)$ and we would like to obtain

$$c''(y) < 0 \quad (\text{A.25})$$

$$(1-\beta)\theta_H x_P(y)^{-\beta} x'_P(y) + I_A \beta x_P(y)^{-\beta-1} x'_P(y) < 0 \quad (\text{A.26})$$

$$(1-\beta)\theta_H x_P(y)^{-\beta} + I_A \beta x_P(y)^{-\beta-1} > 0 \quad (\text{A.27})$$

$$(1-\beta)\theta_H x_P(y) + I_A \beta > 0 \quad (\text{A.28})$$

$$\frac{\beta}{\beta-1} \cdot I_A > \theta_H x_P(y) \quad (\text{A.29})$$

The above is true for all $y > y^*$ by definition of y^* .

Recommendation best response. In any equilibrium the principal acts only when $P_t(\theta = \theta_H) = 1$. Otherwise the agent can weakly improve by fully revealing θ instead of letting the principal act. For a given equilibrium process (x_t, \hat{y}_t) define

$$\hat{\tau} = \inf\{t : (x_t, \hat{y}_t) \in \mathbf{T}\} \quad (\text{A.30})$$

Then define

$$\hat{m}_t = \mathbb{1}\{\hat{\tau} \geq t\} \quad (\text{A.31})$$

In that case the agent can always send a binary message about when to reveal $\theta = \theta_H$. Define by Λ_t the history dependent cumulative intensity of sending the message $\theta = \theta_H$. Consider an arbitrary process (x_t, \hat{y}_t) where

$$d\hat{y}_t = (1 - \hat{y}_t) dN_t - (1 - \hat{y}_t) d\Lambda_t \quad (\text{A.32})$$

Then suppose the agent follows strategy corresponding to (x_t, \hat{y}_t) and then switches at time t to the equilibrium strategy corresponding to reflection at $x_P(y)$ and the value function $V_A(x, y)$. This value function is given by

$$Z_t = e^{-r\tau} (\theta_H x_\tau - I_A) \mathbb{1}\{\tau < t\} + e^{-rt} V_A(x_t, \hat{y}_t) \quad (\text{A.33})$$

Then

$$dZ_t = (\theta_H x_t - I_A - V_A(x_t, \hat{y}_t)) d\Lambda_t - \frac{\partial V_A}{\partial y}(x_t, \hat{y}_t) d\Lambda_t + \mu x_t \frac{\partial V_A}{\partial x}(x_t, \hat{y}_t) \quad (\text{A.34})$$

$$+ \sigma x_t \frac{\partial V_A}{\partial x}(x_t, \hat{y}_t) dB_t + \frac{\sigma^2 x_t^2}{2} \frac{\partial^2 V_A}{\partial x^2}(x_t, \hat{y}_t) dt \quad (\text{A.35})$$

For $(x_t, \hat{y}_t) \notin \mathbf{T}$ process Z_t has a negative drift due to concavity of $V_A(x, y)$ in that region. This implies that the agent always finds it optimal to keep quiet outside area \mathbf{T} . Hence it is optimal to set $d\Lambda_t = 0$ outside \mathbf{T} . This

$$\mathbb{E}_{(x_0, y_0)} [e^{-r\hat{\tau}} (\theta_H x_{\hat{\tau}} - I_A)] \leq V_A(x_0, y_0) \quad (\text{A.36})$$

Implementation. The equilibrium belief process Y differs from $Y_0 - M_t$ in two ways:

- Y is a martingale, while $Y_0 - M_t$ is weakly decreasing
- $Y_0 - M_t$ has positive intensity of jumps at the boundary of the stopping set.

Process $Y_0 - M_t$ can be defined as follows: consider an exponentially random variable η , independently distributed from X and θ , and let κ denote

$$\kappa = \inf\{t > 0 : \ln(1 - Y_0 + M_t) - \ln(1 - Y_0) \geq \eta\},$$

finally put

$$Y_t = \begin{cases} Y_0 - M_t, & \text{if } t < \kappa \\ 1, & \text{if } t \geq \kappa. \end{cases}$$

Notice that Y_t is a martingale, since at t such that $x_P(Y_{t-}) = X_{t-}$ the process Y_t has either a jump $1 - Y_{t-}$ up with intensity $1 - e^{d \ln(1 - \hat{Y}_t)} = -\frac{d\hat{Y}_t}{1 - \hat{Y}_t}$ or simply decreases by $d\hat{Y}_t$:

$$\mathbb{E}[dY_t] = d\hat{Y}_t - (1 - Y_{t-}) \frac{d\hat{Y}_t}{1 - \hat{Y}_t} = d\hat{Y}_t - d\hat{Y}_t = 0$$

since $\hat{Y}_t = \hat{Y}_{t-}$ and the latter is equal Y_{t-} for $t < \kappa$.

□

Proof of Proposition 2. Define $l(x)$ to be the threshold in the already constructed equilibrium such that the principal stops when $y > l(x)$. Suppose there exists an alternative equilibrium with $\hat{l}(x) \neq l(x)$. The following properties apply:

- (i) Principal never finds it incentive compatible to stop before $l(x)$. This implies that $\hat{l}(x) \geq l(x)$.

(ii) Agent does not reveal any information for $y \leq l(x)$. Otherwise he can delay sending the same message.

Denote by $\hat{V}_P(x, y)$ the value obtained by the principal in the alternative equilibrium. Then $\hat{V}_P(x, y) \leq V_P(x, y)$. Otherwise we can add the external news process from the original equilibrium and implement the value function $V_A(x, y)$ but the principal will have strictly more information. On the other hand, in the original equilibrium belief reflection does not add value and $V_P(x, y)$ corresponds to a potentially suboptimal strategy of stopping in \mathbf{T} when the agent reveals θ at $x_A(1)$. This is still a valid strategy in the $\hat{l}(x)$ equilibrium. This implies that $\hat{V}_P(x, y) \geq V_P(x, y)$. Thus it must be the case that $\hat{V}_P(x, y) = V_P(x, y)$.

Suppose $\hat{l}(x) > l(x)$. Then $\forall y < \frac{l(x)+\hat{l}(x)}{2}$ there is no persuasion in the alternative equilibrium. Principal's value function $\hat{V}_P(x, y)$ is strictly convex in x for $x < \hat{l}^{-1}(y)$ since $\hat{V}_P(0, y) = 0$ to the left of the stopping region and satisfies the ODE (19). However, $V_P(x, y)$ is linear in x in the same region - a contradiction.

To the right of the stopping region (between $x \in (l^{-1}(y), x_A(1))$) the convexity argument does not work. However, by assumption $\hat{l}^{-1}(y) < \underline{x}(y)$ ($\hat{l}(x) \geq l(x)$ and $l(\underline{x}(y)) = y$ in that region). Thus function $g(x, \hat{l}^{-1}(y))$ cannot satisfy the matching condition (A.16) at $x_A(1)$. This is equivalent to $\hat{V}_P(x, y)$ not satisfying smooth-pasting condition in x at $(\hat{l}^{-1}(y), y)$.

To illustrate this point, consider a function $h(\cdot)$ defined as:

$$\begin{aligned} h(x) &= C_1 x^{\beta_1} + C_2 x^{\beta_2} \\ h(\hat{x}) &= A > 0 \\ h'(\hat{x}) &= B > 0 \end{aligned}$$

The value $h(x_A(1))$ is increasing in B if $x_A(1) > \hat{x}$:

$$\frac{\partial}{\partial B} h(x_A(1)) = \frac{x_A(1)}{\beta_1 - \beta_2} \left[\left(\frac{x_A(1)}{\hat{x}} \right)^{\beta_1} - \left(\frac{x_A(1)}{\hat{x}} \right)^{\beta_2} \right] > 0.$$

This value $\hat{V}_P(x_A(1), y)$ is increasing in $\frac{\partial \hat{V}}{\partial x}(\hat{l}^{-1}(y), x_A(1))$. Suppose the following is true:

$$\begin{cases} \hat{V}(\hat{l}^{-1}(y), y) = (y\theta_H + (1-y)\theta_L)\hat{l}^{-1}(y) - I_P & \text{value matching} \\ \frac{\partial \hat{V}}{\partial x}(\hat{l}^{-1}(y), y) \geq y\theta_H + (1-y)\theta_L \end{cases}$$

Because $x_P(y) < \hat{l}^{-1}(y) < \underline{x}(y)$ it follows that it cannot satisfy the boundary condition at $x_A(1)$:

$$\hat{V}(x_A(1), y) = y(\theta_H x_A(1) - I_P) + (1-y)V_P^{NI}(x_A(1), 0), \quad (\text{A.37})$$

since the implied $\hat{V}(x_A(1), y)$ is too high. To match the value at $x_A(1)$ the partial derivative $\frac{\partial \hat{V}}{\partial x}(\hat{l}^{-1}(y), y)$ must be below $y\theta_H + (1-y)\theta_L$, which implies that for sufficiently small ε we have $\hat{V}_P(\hat{l}^{-1}(y) + \varepsilon, y) < V_P(x, y)$.

This again contradicts with $\hat{V}_P(x, y) = V_P(x, y)$ in that region. \square

Proof of Lemma 4. We can write the value of waiting for the information as

$$(W) = y \left(\frac{x^*}{x_A(1)} \right)^{\beta_1} [\theta_H x_A(1) - I_P] + (1 - y) \left(\frac{x^*}{x_P(0)} \right)^{\beta_1} [\theta_L x_P(0) - I_P],$$

and the value of acting as

$$(A) = y[\theta_H x^* - I_P] + (1 - y)[\theta_L x^* - I_P].$$

The difference is

$$\begin{aligned} (W) - (A) &= y \left[\left(\frac{x^*}{x_A(1)} \right)^{\beta_1} [\theta_H x_A(1) - I_P] - [\theta_H x^* - I_P] \right] \\ &\quad + (1 - y) \left[\left(\frac{x^*}{x_P(0)} \right)^{\beta_1} [\theta_L x_P(0) - I_P] - [\theta_L x^* - I_P] \right]. \end{aligned}$$

(i) Notice, that θ_H appears only in the first term of the expression above, moreover, since $\theta_H x_A(1)$ does not depend on θ_H the derivative of the first term w.r.t to θ_H has the same sign as the derivative of the first term w.r.t x^* . Finally, notice that x^* is chosen optimally, hence $\frac{d}{dx^*} [(W) - (A)] = 0$, hence we can only need to sign the derivative of the second term:

$$\begin{aligned} \frac{d}{dx^*} \left[\left(\frac{x^*}{x_P(0)} \right)^{\beta_1} [\theta_L x_P(0) - I_P] - [\theta_L x^* - I_P] \right] &= \frac{\beta_1}{x^*} \left(\frac{x^*}{x_P(0)} \right)^{\beta_1} \frac{I_P}{\beta_1 - 1} - \theta_L \\ &= \theta_L ((x^*)^{\beta_1 - 1} \cdot (x_P(0))^{1 - \beta_1} - 1) < 0 \end{aligned}$$

Thus, the value of waiting grows faster than the value of exercising the option as θ_H increases and y^* needs to adjust up to compensate for that.

(ii) For a fixed y^* an increase in θ_L translates into an increase of the exercise value by $(1 - y^*)x^*$. At the same time the value of waiting goes up by

$$\frac{d}{d\theta_L} \left[\left(\frac{x^*}{x_A(1)} \right)^{\beta_1} (y^*(\theta_H x_A(1) - I_P) + (1 - y^*)V_P^{NI}(x_A(1), 0)) \right] = (1 - y^*) \left(\frac{x^*}{x_A(1)} \right)^{\beta_1} \frac{d}{d\theta_L} V_P^{NI}(x_A(1), 0),$$

where we can treat x^* as constant due to the envelope theorem. Explicit calculations yield

$$\left(\frac{x^*}{x_A(1)} \right)^{\beta_1} \frac{d}{d\theta_L} V_P^{NI}(x_A(1), 0) = (x^*)^{\beta_1} \cdot (x_P(0))^{1 - \beta_1}.$$

Since $x^* < x_P(0)$ the value of exercise goes up faster than the value of waiting, thus, y^* needs to adjust down to compensate for that. \square

Proof of Lemma 5. Start with equation (A.5) which pins down y^* :

$$\left(\frac{x_P(y^*)}{x_A(1)}\right)^{\beta_1} \left(y^*(\theta_H x_A(1) - I_P) + (1 - y^*)V_P^{NI}(x_A(1), 0)\right) - \frac{I_P}{\beta_1 - 1} = 0 \quad (\text{A.38})$$

Express everything in terms of I_P :

$$\begin{aligned} \left(\frac{I_P}{I_A \mathbb{E}_{y^*}[\theta]}\right)^{\beta_1} \left(y^* \left(\frac{\beta_1 I_A}{\beta_1 - 1} - I_P\right) + (1 - y^*) \left(\frac{x_A(1)}{x_P(0)}\right)^{\beta_1} \frac{I_P}{\beta_1 - 1}\right) - \frac{I_P}{\beta_1 - 1} &= 0 \\ y^* \left(\frac{I_P}{I_A \mathbb{E}_{y^*}[\theta]}\right)^{\beta_1} \left(\frac{\beta_1 I_A}{\beta_1 - 1} - I_P\right) + (1 - y^*) \left(\frac{x_P(y^*)}{x_P(0)}\right)^{\beta_1} \frac{I_P}{\beta_1 - 1} - \frac{I_P}{\beta_1 - 1} &= 0 \\ y^* \left(\frac{I_P}{I_A \mathbb{E}_{y^*}[\theta]}\right)^{\beta_1} \left(\frac{\beta_1 I_A}{\beta_1 - 1} - I_P\right) + (1 - y^*) \left(\frac{1}{\mathbb{E}_{y^*}[\theta]}\right)^{\beta_1} \frac{I_P}{\beta_1 - 1} - \frac{I_P}{\beta_1 - 1} &= 0 \end{aligned}$$

Define $\alpha = \frac{I_P}{I_A} < 1$. Divide the above equation by I_A . Then y^* is pinned down by

$$\begin{aligned} y^* \left(\frac{\alpha}{\mathbb{E}_{y^*}[\theta]}\right)^{\beta_1} \left(\frac{\beta_1}{\beta_1 - 1} - \alpha\right) + (1 - y^*) \left(\frac{1}{\mathbb{E}_{y^*}[\theta]}\right)^{\beta_1} \frac{\alpha}{\beta_1 - 1} - \frac{\alpha}{\beta_1 - 1} &= 0 \\ y^* \alpha^{\beta_1} \left(\frac{\beta_1}{\beta_1 - 1} - \alpha\right) + (1 - y^*) \frac{\alpha}{\beta_1 - 1} - \frac{\alpha}{\beta_1 - 1} \cdot (\mathbb{E}_{y^*}[\theta])^{\beta_1} &= 0 \\ y^* \alpha^{\beta_1 - 1} \left(\frac{\beta_1}{\beta_1 - 1} - \alpha\right) + (1 - y^*) \frac{1}{\beta_1 - 1} - \frac{1}{\beta_1 - 1} \cdot (\mathbb{E}_{y^*}[\theta])^{\beta_1} &= 0 \end{aligned}$$

Hence we can get the comparative statics with respect to both I_P and I_A by looking at the comparative statics with respect to α . The derivative with respect to α is

$$y(\beta_1 \alpha^{\beta_1 - 2} - \beta_1 \alpha^{\beta_1 - 1}) > 0 \quad (\text{A.39})$$

Hence we obtain that y^* satisfies $G(y^*, \alpha) = 0$ satisfying $\frac{\partial G}{\partial y^*} < 0$ (coming from the proof of Lemma 3) and $\frac{\partial G}{\partial \alpha} < 0$. This implies that

$$\frac{dy^*}{d\alpha} = -\frac{\frac{\partial G}{\partial \alpha}(y^*, \alpha)}{\frac{\partial G}{\partial y^*}(y^*, \alpha)} > 0 \quad (\text{A.40})$$

Since $\alpha = \frac{I_P}{I_A}$ we obtain both points of Lemma 5. □

Proof of Proposition 3. Fix strategy of the principal to be:

- if $y = 1$, then act for $x \geq x_P(1)$,
- if $y < 1$, then act for $x \geq x_P(0)$.

This is a best response to an equilibrium in which the agent reveals θ immediately after any history.

If the agent reveals information immediately, then the action will be taken at $\max(x, x_P(1))$ if $\theta = \theta_H$ and at $x_P(0)$ for $\theta = \theta_L$. Any other strategy by the agent can only increase the threshold of action conditional

on $y = 1$ (e.g. no information results in action taken at $x_P(0)$) and does not affect the threshold of action conditional on $y = 0$. Since

$$\left(\frac{x}{x^*}\right)^{\beta_1} ((y\theta_H + (1-y)\theta_L)x^* - I_A)$$

is decreasing in x^* for $x^* > x_A(y)$ and $x_A(y) < x_P(y) \leq x_P(0)$, revealing information immediately is the best response of the agent. \square

Proof of Proposition 4 . In the first step of the proof we show that the agent can always reveal information at $t = 0$. Consider an equilibrium with some stopping set of the Receiver \mathbf{T} . Denote by $y^*(x)$ the lower bound of that set, i.e.

$$y^*(x) = \inf\{y : (x, y) \in \mathbf{T}\}. \quad (\text{A.41})$$

$$x^*(y) = \inf\{x : (x, y) \in \mathbf{T}\} \quad (\text{A.42})$$

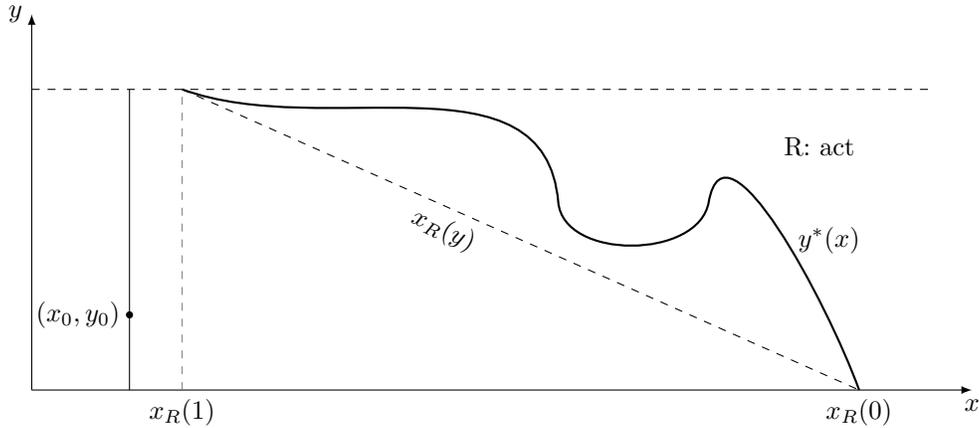


Figure 9: Illustration for the proof.

Agent's payoff is

$$V_A(x_0, y_0) = \mathbb{E}_{(x_0, y_0)} \left[e^{-r\tau(\tilde{x}, \tilde{y})} v_A(\tilde{x}, \tilde{y}) \right] \quad (\text{A.43})$$

where $v_A(x, y) = [y\theta_H + (1-y)\theta_L]x - I_A$. Note that $\tilde{x} \geq x^*(\tilde{y}) \geq x_R(\tilde{y})$.

Consider the following alternative strategy: at $t = 0$ the agent reveals the pair (\tilde{x}, \tilde{y}) . In other words, the agent conducts a partially informative test about θ generating a set of posterior signals \tilde{y} . In addition, he provides the principal with a recommendation to stop at $\tilde{x} \geq x_P(\tilde{y})$. After $t = 0$ the agent does not communicate with the principal.

Suppose the principal agrees to stop at \tilde{x} given agent's recommendation. The pair $(\tilde{x}, \tilde{y}) \in \mathbf{T}$ hence the principal can still stop then. Potentially, the principal may find it optimal to stop for a lower x , than wait until \tilde{x} , but this only benefits the agent.

Define

$$\hat{\tau} = \inf\{t : x_t \geq \tilde{x}\} \quad (\text{A.44})$$

Note that \tilde{y} and \tilde{x} are disclosed at $t = 0$. Thus $\hat{\tau}$ is a first hitting time with respect to the filtration chosen by the agent. The payoff from this alternative strategy to the agent is

$$\hat{V}_A(x_0, y_0) = \mathbb{E}_{(x_0, y_0)} [e^{-r\hat{\tau}} v_A(\tilde{x}, \tilde{y})] \quad (\text{A.45})$$

Suppose $\tau(\tilde{x}, \tilde{y}) < \hat{\tau}(\tilde{x}, \tilde{y})$ for some ω . Then it must imply that $x_t < \tilde{x}$ at $t = \tau(\tilde{x}, \tilde{y})$ which would be a contradiction with $x_{\tau(\tilde{x}, \tilde{y})} \equiv \tilde{x}$. Hence the agent weakly benefits from revealing (\tilde{x}, \tilde{y}) at $t = 0$. Because this alternative strategy delivers the same distribution over stopped outcomes (\tilde{x}, \tilde{y}) but at shorter delay we obtain:

$$\hat{V}_A(x_0, y_0) \geq V_A(x, y) \quad (\text{A.46})$$

Note that under this strategy the principal may choose to stop prior to \tilde{x} . Conditional on \tilde{y} revealed by the agent at $t = 0$ we have

$$\hat{V}_A(x_0, y_0) \leq \mathbb{E} [e^{-r\tilde{\tau}(\tilde{y})} v_A(x^*(\tilde{y}), \tilde{y})] \quad (\text{A.47})$$

where

$$\tilde{\tau}(\tilde{y}) = \inf\{t : x_t = x^*(\tilde{y})\} \quad (\text{A.48})$$

which is, again, a valid stopping time since \tilde{y} is measurable with respect to \mathcal{F}_0 . Thus by revealing \tilde{y} at $t = 0$ the agent weakly benefits in any equilibrium. This completes the first step of the proof.

In the second step, we show that the agent prefers to give the principal a fully informative signal at $t = 0$ if he can only persuade then. Specifically, suppose $\mathbb{P}(\tilde{y} \in (0, 1)) > 0$. The payoff of the agent in this case is exactly equal to

$$\left(\frac{x_0}{x^*(\tilde{y})}\right)^{\beta_1} v_A(x^*(\tilde{y}), \tilde{y}) \leq \left(\frac{x_0}{x_P(\tilde{y})}\right)^{\beta_1} v_A(x_P(\tilde{y}), \tilde{y}) \quad (\text{A.49})$$

It is easy to show that the RHS is strictly less than

$$\tilde{y} \left(\frac{x_0}{x_P(1)}\right)^{\beta_1} v_A(x_P, 1) + (1 - \tilde{y}) \left(\frac{x_0}{x_P(0)}\right)^{\beta_1} v_A(x_P, 0) \quad (\text{A.50})$$

Hence the agent can weakly improve any equilibrium payoff by simply revealing θ at $t = 0$ and forcing the

principal to act at either $x_P(1)$ or $x_P(0)$.

□

Lemma 10. Consider two filtrations $\mathbb{H} = (\mathcal{H}_t)_{t \geq 0}$ and $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ and two stopping times τ_1, τ_2 w.r.t. \mathbb{H} such that

$$\mathcal{H}_{\tau_i} = \mathcal{F}_{\tau_i} \quad i \in \{1, 2\}.$$

Then the maximum of τ_1 and τ_2 , $\tau_1 \vee \tau_2$ is also a stopping time w.r.t. \mathbb{H} and

$$\mathcal{H}_{\tau_1 \vee \tau_2} = \mathcal{F}_{\tau_1 \vee \tau_2}.$$

Proof. Obviously, $\tau_1 \vee \tau_2$ is a stopping time w.r.t. \mathbb{H} . The assertion above follows from the following representation of $\mathcal{H}_{\tau_1 \vee \tau_2}$ ($\mathcal{F}_{\tau_1 \vee \tau_2}$):

$$\mathcal{H}_{\tau_1 \vee \tau_2} = \{A \cup B, A \in \mathcal{H}_{\tau_1}, B \in \mathcal{H}_{\tau_2}\}.$$

To show \supseteq inclusion we prove that $\mathcal{H}_{\tau_1}, \mathcal{H}_{\tau_2} \subseteq \mathcal{H}_{\tau_1 \vee \tau_2}$. Consider an arbitrary $A \in \mathcal{H}_{\tau_1}$, by definition

$$A \cap \{\tau_1 \leq t\} \in \mathcal{H}_t \quad \forall t \geq 0.$$

Recall that $\{\tau_2 \leq t\} \in \mathcal{H}_t$ since τ_2 is a stopping time w.r.t. to \mathbb{H} , thus,

$$A \cap \{\tau_1 \vee \tau_2 \leq t\} = [A \cap \{\tau_1 \leq t\}] \cap \{\tau_2 \leq t\} \in \mathcal{H}_t \quad \forall t \geq 0,$$

i.e., $A \in \mathcal{H}_{\tau_1 \vee \tau_2}$. Which in turn implies $\mathcal{H}_{\tau_1} \subseteq \mathcal{H}_{\tau_1 \vee \tau_2}$. Similarly, one can show that $\mathcal{H}_{\tau_2} \subseteq \mathcal{H}_{\tau_1 \vee \tau_2}$.

Next, we show that \subseteq inclusion holds as well. Consider an arbitrary set $C \in \mathcal{H}_{\tau_1 \vee \tau_2}$. This set can be decomposed into

$$C = [C \cap \{\tau_1 \geq \tau_2\}] \cup [C \cap \{\tau_1 \leq \tau_2\}].$$

Let $A = C \cap \{\tau_1 \geq \tau_2\}$. First notice that $\{\tau_1 \geq \tau_2\} \in \mathcal{H}_{\tau_1}$ because

$$\{\tau_1 \geq \tau_2\} \cap \{\tau_1 \geq t\} = \{\tau_1 \wedge t \geq \tau_2 \wedge t\} \cap \{\tau_1 \leq t\} \cap \{\tau_2 \leq t\} \in \mathcal{H}_t.$$

Since $C \in \mathcal{H}_{\tau_1 \vee \tau_2}$ and $\{\tau_1 \geq \tau_2\} \in \mathcal{H}_{\tau_1} \subseteq \mathcal{H}_{\tau_1 \vee \tau_2}$ the set A is also in $\mathcal{H}_{\tau_1 \vee \tau_2}$. By definition we have

$$C \cap \{\tau_1 \geq \tau_2\} \cap \{\tau_1 \vee \tau_2 \leq t\} \in \mathcal{H}_t \quad \forall t \geq 0,$$

which clearly implies

$$C \cap \{\tau_1 \geq \tau_2\} \cap \{\tau_1 \leq t\} \in \mathcal{H}_t \quad \forall t \geq 0,$$

i.e., $A = C \cap \{\tau_1 \geq \tau_2\} \in \mathcal{H}_{\tau_1}$. Similarly one can show that $B = C \cap \{\tau_1 \leq \tau_2\} \in \mathcal{H}_{\tau_2}$.

□

Proof of Corollary 2. Notice that the low type agent faces an equilibrium exercise threshold $x_P(0)$, which is the longest the principal is willing to wait in any equilibrium, hence, deviations from S^* are suboptimal. Following S^* delivers the first best for the high type agent when $X_t \geq x_A(1)$ or when $Y_t \leq y^*$, thus, deviations from S^* in that region are suboptimal. For $Y_t > y^*$ and $X_t < x_A(1)$ the high type agent has no incentives to deviate from S^* due to the argument analogous to the one in Proposition 1: his value function can be written as $c(y)x^\beta$ with $c''(y) < 0$, thus, persuasion strictly outside of the principal's action set is suboptimal. Finally, any deviation from S^* at the boundary of \mathcal{T}^* results either in immediate option exercise or a posterior that lies strictly outside of action region of the principal, which is again suboptimal. □

Proof of Corollary 3. Recall that the strategy \mathcal{T}^* of the principal is to exercise the option at $x_P(1)$ if $Y_t = 1$ and at $x_P(0)$ if $Y_t < 1$. The high type agent guarantees the earliest exercise time $x_P(1)$ by following S^* , any deviation could only delay option exercise and, thus, are suboptimal. The low type is indifferent between all possible deviations, since his exercise time is $x_P(0)$ regardless of principal's beliefs (inducing $Y_t = 1$ is infeasible for the low type agent when the principal holds passive beliefs), hence, he has no strict incentives to deviate. □

Proof of Proposition 5. A general feature of perfect information available to the agent is that persuasion is immediate. Denote by \tilde{q} the posterior distribution of beliefs at which the principal acts in any Perfect Markov Equilibrium. It must be the case that $\text{supp}(\tilde{q}) \subseteq \mathbf{T}$. The payoff of the agent is

$$\mathbb{E} [e^{-r\tau}((F_1 + F_0)\tilde{q} - F_0)] \tag{A.51}$$

There are delays in equilibrium if $\mathbb{P}(\tau > 0) > 0$. The agent can generate the same posterior distribution of \tilde{q} at $t = 0$. Then there is a profitable deviation since

$$\mathbb{E} [(F_1 + F_0)\tilde{q} - F_0] > \mathbb{E} [e^{-r\tau}((F_1 + F_0)\tilde{q} - F_0)] \tag{A.52}$$

Hence immediate persuasion must be part of the equilibrium and we can focus on instant signals communicated by the agent.

Suppose $F_0 \geq 0$. The agent prefers the option to be exercised for $\xi = 1$ and not exercised at $\xi = 0$. Agent's and principal's preferences are thus aligned. So the agent can achieve his first best by disclosing ξ immediately and fully.

Suppose $F_0 < 0$. In this case the agent prefers the option to be exercised for both $\xi = 0$ and $\xi = 1$. The principal still prefers to exercise the option only when $\xi = 1$. Agent's and principal's preferences are not

aligned in this case. Because we are looking at Perfect Markov Equilibria in posterior beliefs, principal's stopping strategy is a set $\mathbf{T} \subset [0, 1]$ of beliefs about ξ at which she stops. Define by q_P her autarky threshold as before. Define

$$\bar{q} = \inf\{q \in \mathbf{T}\}$$

Since $\bar{q} \geq q_P$, it implies that $\bar{q} > 0$.

If $q_0 < \bar{q}$, the best response of the agent is to conduct a partially informative signal between 0 and \bar{q} . This is the unique best response since the agent wants to maximize the probability of the action being taken at $t = 0$.

Can $\bar{q} < 1$ be an equilibrium? Suppose it is. Define by $V_P(q)$ to be principal's value function. In order for stopping at \bar{q} to be optimal the following conditions must be satisfies:

$$V_P(\bar{q}) = (W_1 + W_0)\bar{q} - W_0 \quad \text{value matching} \quad (\text{A.53})$$

$$V'_P(\bar{q}) = W_1 + W_0 \quad \text{smooth pasting} \quad (\text{A.54})$$

On the other hand, when the agent sends the principal a partially informative message for $q_0 < \bar{q}$, principal's value function is given by

$$V_P(q) = \frac{q}{\bar{q}} ((W_1 + W_0)\bar{q} - W_0) = (W_1 + W_0)q - W_0 \frac{q}{\bar{q}} \quad (\text{A.55})$$

Then

$$\lim_{\varepsilon \uparrow 0} V'_P(q + \varepsilon) = W_1 + W_0 - \frac{W_0}{\bar{q}} < V'_P(\bar{q}) \quad (\text{A.56})$$

since $W_0 < 0$ (otherwise immediate action is optimal and the agent might as well disclose all of the information at $t = 0$). This is a contradiction since $V_P(q)$ is a solution to the principal's optimal stopping problem. This is a contradiction with $\bar{q} < 1$ because the smooth pasting condition of optimal real option exercised is an interior first order condition. When $\bar{q} = 1$ it does not apply since the principal's optimal strategy is a corner solution. Hence we rule out $\bar{q} < 1$ as equilibrium candidates, while $\bar{q} = 1$ is a Markov Perfect Equilibrium by the same arguments as in Proposition 3. \square

Lemma 11. *Suppose $I_A > I_P$, $\theta_L = 0$, and $\theta_H = 1$. If the principal can commit to pay the agent a bonus $b \geq 0$ upon exercising the option, the optimal bonus is $b = 0$.*

Proof. We are assuming, for simplicity, that the bonus is picked at $t = 0$ and kept constant for the rest of the game. Conjecture an optimal $b \geq 0$ chosen at $t = 0$ in state (x_0, y_0) . Define by $y^*(b)$ the tip of the action triangle as a function of b .

Suppose $y_0 \geq y^*(b)$. Then, in equilibrium the principal gets his autarky value corresponding to the investment cost $I_P + b$ (investment cost and the bonus paid to the agent). This implies that this candidate b cannot be optimal since the principal is better off at setting $b = 0$.

Alternatively, suppose $y_0 < y^*(b)$. For simplicity, we assumed that $\theta_L = 0$. Then the payoff of the Receiver is

$$\begin{aligned} & y_0 \cdot \left(\frac{x}{\frac{\beta_1}{1-\beta_1}(I_A - b)} \right)^{\beta_1} \cdot \left(\frac{\beta_1}{1-\beta_1}(I_A - b) - I_P - b \right) \\ &= y_0 \cdot \left(\frac{x(1-\beta_1)}{\beta_1} \right)^{\beta_1} \cdot \left(\frac{\beta_1}{1-\beta_1}(I_A - b)^{1-\beta_1} - (I_P + b)(I_A - b)^{-\beta_1} \right) \end{aligned}$$

The derivative of the above expected payoff with respect to b is

$$y_0 \left(\frac{x(1-\beta_1)}{\beta_1} \right)^{\beta_1} \left(-(1-\beta_1) \frac{\beta_1}{1-\beta_1} (I_A - b)^{-\beta_1} - (I_A - b)^{-\beta_1} - \beta_1 (I_P + b) (I_A - b)^{-\beta_1 - 1} \right) < 0$$

Hence the principal would prefer to set a lower b . Eventually $y_0 = y^*(b)$ and hence the principal optimally chooses $b = 0$. □