Persuading the Regulator To Wait*

Dmitry Orlov† Andrzej Skrzypacz‡ Pavel Zryumov§

First Draft: 13th February, 2016

Abstract

We study a Bayesian persuasion game in the context of real options. The Sender (firm) chooses signals to reveal to the Receiver (regulator) each period but has no long-term commitment power. The Receiver chooses when to exercise the option, affecting welfare of both parties. When the Sender favors late exercise relative to the Receiver but their disagreement is small, in the unique equilibrium information is disclosed discretely with a delay and the Receiver sometimes ex-post regrets waiting too long. When disagreement is large, the Sender, instead of acquiring information once and fully, pipets good information over time. These results shed light on the post-market surveillance practices of medical drugs and instruments by the FDA, and the role of pharmaceutical companies in keeping marginally beneficial drugs in the market. When the Sender favors early exercise, the 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 Sender.

1 Introduction

Consider a pharmaceutical firm selling an FDA-regulated drug or medical device. When the drug is first introduced, 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 (FDA), monitors this news and can recall the drug (i.e. remove it from the market) if it turns out to be unsafe. While the firm does not want to sell a bad drug, it is natural to expect a partial misalignment between the firm and the regulator over when to exercise the real option of recalling the drug. For example, if the firm does not fully internalize all the costs of a bad drug, it would prefer to wait longer for stronger

*We thank Mike Harrison, Johannes Hörner, Ron Kaniel, Eric Madsen, and seminar participants at the University of Rochester for useful comments and suggestions.

†University or 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

evidence of side effects than the regulator would. Similarly, if the regulator discounts future more than the firm, it would choose to exercise the real option sooner than the firm would want it to. The firm cannot pay the regulator to postpone a recall, but can persuade it to wait by providing additional information/tests.[1]

In this paper we analyze a model of dynamic persuasion of the regulator by the firm and characterize equilibrium dynamics of information provision. The firm and the regulator learn jointly from public news about quality of the product produced by the firm (e.g. drug efficacy). The regulator has the power to stop the marketing of the product at any time (i.e. he has the decision rights over the execution of the real option) and we assume it is rational and forward-looking.

In addition to the public news, the firm can at any time run additional tests that provide a noisy signal of product quality. The firm's decision is what information to acquire and when. In most of the paper we assume that the firm cannot perfectly learn the state via tests. This assumption is motivated by the observation that short-term tests are not perfectly informative about long-term effects of a drug and the regulator cares both about the short and long-term effects.[2] The design and timing of the tests is chosen by the firm - the regulator interprets the results of the tests, but does not have the power to compel the firm to take them. The firm does not have private information about the quality of the product and we assume that when it runs a test it cannot hide the results from the regulator.[3]

Beyond pharmaceutical firms and their regulators, our model applies to many economic situations in which one agent has decision rights over a real option and another agent controls some information acquisition (and there is misalignment of incentives). One example is environmental regulation: for instance, information about price of oil or local risks of a given oil well comes gradually to the market, and the firm can let engineers perform additional tests about the environmental impact of the well. Another example is decision making within organizations: for instance, a product manager and a CEO learn over time about profitability of a product and the product manager can perform market tests. While the manager may be paid for performance, as long as the contract does not fully align incentives, the manager will have incentives to manipulate information acquisition to influence CEO's

[1] For example, see "Guidance for Industry and FDA Staff. Postmarket Surveillance Under Section 522 of the Federal Food, Drug and Cosmetic Act" issued on April 27, 2006 and available at http://www.fda.gov/downloads/medicaldevices/deviceregulationandguidance/guidancedocuments/ucm072564.pdf. Under these FDA guidelines, the manufacturer has the opportunity to provide additional information and identify specific surveillance methodologies before the FDA issues postmarket surveillance. The FDA does not require controlled clinical trials to address its concerns, but "intend(s) that manufacturers use the most practical least burdensome approach to produce a scientifically sound answer."

[2] Additionally, very accurate tests may be prohibitively costly, but we abstract away from any direct costs of testing.

[3] We do not allow the firm to secretly run tests about the drug and hide unfavorable results and we assume that the information is verifiable since it comes from scientific 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 regulator knows that the firm has verifiable information then due to standard unraveling arguments the firm would reveal it). 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.

decision to terminate the project. Finally, the model can also apply to stakeholders' decisions when to sell off firm's assets: for example, some stakeholders may have access to sources of information while others may have the majority required to make the decision to sell assets and their preferences need not be perfectly aligned. The key assumptions of the model are that the regulator cannot be paid by the firm to delay a recall, the regulator cannot pay the firm to disclose information and both parties have no commitment to future actions. For tractability, we also assume that the underlying state and signal are binary, time is continuous and information comes to the market via a Brownian motion with a drift that depends on the true state.

The incentives to acquire information depend on the misalignment of preferences between the firm and the regulator. On one hand, information improves decisions so if there was no misalignment, the firm would acquire information immediately. On the other hand, information affects when the regulator recalls the drug - it can either speed up or delay it.[4] For example, if the regulator expects to reveal information in near future, it may find it optimal to delay the recall.

Persuasion to Wait. In the main part of the paper we assume that the firm/Sender is "more patient" than the regulator/Receiver in the sense that absent any information from the firm, the regulator would exercise the real option sooner than the firm would like to. So it is in the firm's interest to persuade the regulator to wait. In this case, we show that this game has a unique Markov-perfect equilibrium (Markov in beliefs) in which the Sender delays information disclosure and the Receiver sometimes ex-post regrets waiting too long. For a given prior, we show that the pattern of information disclosure depends on the amount of disagreement between the firm and the regulator when to recall the drug.

When the misalignment is small, then in equilibrium the firm postpones information acquisition until its beliefs reach a point that in the event the noisy signal would deliver bad news about the drug, the firm would want to recall the drug itself. In that case the firm acquires information once and disclosure is fully revealing. The regulator, not wanting to take a potentially effective drug off the market, and expecting the firm to provide comprehensive information when beliefs become sufficiently pessimistic, waits in equilibrium for the firm to provide information. This can mean that the regulator waits past a belief threshold at which it would stop with only public news. Finally, the timing of the recall has a "compromise" property: in case the test results are bad, the regulator recalls the drug at the firm's optimal belief; when the test results are good, the regulator further delays the recall but if he eventually does recall, it is at the regulator's optimal threshold. This seems to correspond to the two types of recalls in the market: in some instances the regulator suggest the firm a voluntary

[4] In the special case when the firm can perfectly learn the quality of the drug, as long as the misalignment of preferences is not so extreme that the firm would prefer to market even a bad drug, the firm would reveal that information immediately. Signal imprecision creates more interesting equilibrium dynamics.

recall and the firm decides to follow the recommendation and in some cases the regulator recalls the drug in opposition to the firm. This "compromise" property implies also that even if the information about test results was not verifiable, the equilibrium would be robust - when the news is bad it is in the firm's self-interest to recall the drug, so seeing the firm not recalling the drug, the regulator would rationally infer that the news is good and optimize accordingly. When the news is bad, the regulator has ex-post regret that it waited, but since it wants to avoid a false-positive of recalling a good drug, ex-ante it is willing to wait for the information.

When the misalignment is large, equilibrium information provision is radically different. We show that instead of acquiring information once and fully, the firm "pipets" good information. That is, once the regulator reaches its recall threshold, which in equilibrium is the same as if the regulator acted on the public news only, the firm reveals just enough information to either slightly reduce the belief or to fully reveal bad information. In the jargon of real options, the regulator's belief threshold becomes a reflecting boundary (and the reflection is accompanied by a jump of beliefs that the drug is bad, so that beliefs are martingales).[5] Somewhat surprisingly, while this strategy helps the firm obtain higher payoffs than with no access to additional information, for the regulator additional information pipetted this way has no value.

Persuasion to Act. We also discuss the case where the Sender is less patient than the Receiver, in the sense that its threshold for exercising the real option is lower than the Receiver's. For example, if the real option is to invest into project that the Sender likes, or stop a project that the Sender does not like but cannot stop on its own.[6] We call this case *persuasion to act*.

In that case there always exists a Markov-perfect equilibrium in which the Sender immediately acquires all information about the noisy signal. While in some cases this is beneficial for the Sender who wants to speed-up the execution of the real option, it can be also harmful. For example, if the regulator's prior is above its threshold for execution based solely on the public news, and the posterior belief after bad news lies in between the firm's and regulator's thresholds, the firm is hurt by equilibrium disclosure. Unlike the case of persuasion to wait, where having access to the noisy signal is always beneficial for the Sender, in the case of persuasion to act, this access can be detrimental for the Sender. Moreover, we show that for some parameter values if the Sender is less patient than the Receiver, in all Markov equilibria the Sender is worse off as a result of having access to additional information.

[5] If 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 the bad news do not arrive, changes to the "waiting for full information" mode.

[6] In case the drug manufacturer faces triple penalties if the drug turns out to be bad, it is possible that it would like to recall the drug even sooner than the regulator; since firms have the right to self-recall a product that situation does not create frictions. Therefore our analysis of persuasion to act applies more to regulator's approval of a product or to a product manager who would like to start a new project and hopes to persuade the CEO to re-assign her.

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 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 Sender'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 technically cumbersome in continuous time. Instead, we propose to define the Sender'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 Sender can in continuous time provide information about the noisy signal in response to the public news. While continuous time implies that for some Markov strategies of the Receiver (mappings from beliefs about the state and noisy signal to the decision to exercise the option or not) there does not exist a Sender's best response (because the supremum over responses is not attainable), the problem does not lead to non-existence of Markov equilibria.

1.1 Related Literature.

Our technology of information acquisition and disclosure is the same as in the literature on Bayesian persuasion, as in Aumann and Maschler (1995) and Kamenica and Gentzkow (2011). That is, the agent does not have private information, chooses how to "split" a prior belief and cannot hide results. In that literature there are a few papers that study dynamic information disclosure, for example, Renault, Solan, and Vieille (2014) and Ely (2015). There are many differences between these two papers and our model. First, in their models the Sender cares only about the Receiver's actions, while in our paper the Sender's preferences over actions depend on the underlying state: if the state is bad, the firm wants to recall the drug, but if the state is good, it does not want it to do so. The second is that we have two long-lived players and consider non-commitment equilibria rather than design of dynamic information acquisition with dynamic commitment. Despite these important modeling differences, somewhat surprisingly we find that in some instances the equilibrium information provision features "pipetting" which in continuous time is equivalent to the "greedy" strategy from their results: releasing the least information possible to prevent the recall. While pipetting is sometimes an equilibrium outcome, we have also identified important cases when information is instead revealed only once and fully.[7]

[7] Another difference is that absent information acquisition, beliefs in those two papers drift because the underlying state changes, while in our model the state does not change but beliefs respond to public news. Therefore, in our paper

Our paper is also somewhat related to literature on voluntary dynamic information disclosure with verifiable information, as in Acharya, DeMarzo, and Kremer (2011) or Guttman, Kremer, and Skrzypacz (2014). The main difference between these models is that in our model the firm decides when to acquire information and has to reveal all news it obtains, while in those models the firm learns exogenously and privately: the "Receiver" in those paper is the market who does not know if the agent obtained the news. Moreover, in those papers the agent can decide to hide bad information, something we do not allow motivated by markets where such withholding of information would be illegal (or where due to unraveling argument such hiding would not happen in equilibrium). We think our model applies better to either markets with regulation compelling firms to disclose certain types of information, or to internal-organization situations where agents may not be able to hide results of tests from their managers.

Grenadier, Malenko, and Malenko (2015) and Guo (2014) consider two variants of dynamic models of experts providing information to a principal with decision rights and partial misalignment of incentives. The difference from our model is that their Senders are privately informed and provide information via cheap talk. This leads to qualitatively different information disclosure in equilibrium. The main similarity is that like in our paper, the dynamics of information disclosure depend on the direction of misalignment of incentives between the Sender and Receiver. For example, in Grenadier, Malenko and Malenko (2015) when the Sender is more patient, in equilibrium information is revealed fully but the Receiver takes action ex-post too late; when the Receiver is more patient, information is partially pooled in equilibrium.

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.

The rest of the paper is organized as follows. In the next section we present the model. In Section 3 we provide some additional notation and preliminary results. In Section 4 we discuss persuasion to wait. In Section 5, persuasion to act. We conclude in Section 6. The Appendix contains proofs of all our formal results.

beliefs are martingales not only within a period but also across periods, which is not true in their papers.

2 Model

2.1 Basic Setup

We start with an informal description of the model. There are two long-lived players, a firm and a regulator, who we also refer to as the Sender (she) and Receiver (he). The firm/Sender sells a product, for example a drug or a medical device, that is either safe or unsafe, which we refer to as the state. The players share a common prior about the state. Over (continuous) time public news arrives about the state - information can come from payoffs or other sources, for example reports on side effects or long-term effects of the product. We model this information as a Brownian motion with a drift that depends on the state. Additional information can be acquired by the firm about a noisy indicator of the state. As long as the drug is on the market, the players earn payoff flows that on average depend on the state. Expected payoff flows are positive if the drug is safe and negative if not (we also consider the case when they are always positive for the firm/Sender). Profits are noisy and do not immediately reveal the state, especially that some consequences of the sales today are delayed with long-term negative effects on the patients. The regulator has control rights over product recall and faces a real-option problem of deciding when to recall the product from the market - his Markov strategy is stopping time that depends on beliefs and he wants to recall unsafe drugs but protect safe ones. We model the firm's actions as the option to acquire information about a binary random variable that is correlated with the state. This information is noisy because some long-term effects of the product cannot be perfectly discovered with short-term tests or because running fully revealing tests is prohibitively expensive. The firm's Markov strategy is a mapping from beliefs to an information acquisition policy that we model as a filtration. We assume that the Sender can decide not to learn certain information, but once she obtains it, she has to disclose it. As a result, at all times the players have the same beliefs about the state. The Receiver makes his decisions based on the exogenous public news and the additional information chosen by the Sender; his option value of waiting for more information depends on both the variance of the public news and the expectation of the Sender's future disclosures. We characterize Markov Perfect Equilibria in which the strategies of both players depend on the current beliefs about the state of the drug and the realization of the additional random variable the Sender can disclose.

Players and Payoffs. There are two long-lived players: a firm/Sender/she and a regulator/Receiver/he. The firm sells a product that can be either safe $\theta = 0$ or unsafe $\theta = 1$.

Time is continuous and infinite. As long as the product is on the market, it generates for the firm expected profit flow which depends on the state: $F_\theta \in \{F_0, F_1\}$, with $F_0 \geq F_1$ and $F_0 > 0$. It also generates expected welfare flow for the Regulator, $W_\theta \in \{W_0, W_1\}$, with $W_0 > 0 > W_1$.

If the product is recalled from the market at time τ then the expected payoffs of the Sender and the Receiver conditional on θ are

$$u_S(\tau,\theta) = \int_0^{\tau} e^{-r_S t}(F_1\theta + (1-\theta)F_0)dt \qquad u_R(\tau,\theta) = \int_0^{\tau} e^{-r_R t}(W_1\theta + (1-\theta)W_0)dt, \tag{1}$$

where r_S and r_R are the discount rates of the two players. We allow the discount rates and payoff flows of the two players to be different. If $F_1 < 0$, then under complete information the two players agree on recall policy (since their payoff flows conditional on θ have the same sign). The friction/disagreement is caused by uncertainty about the state and hence misalignment between the Sender and the Receiver about the option value of waiting for additional information.

Denote by $\mathcal{H}_\tau$ the sigma-algebra reflecting all the information available to the parties at the time of the drug recall τ, and let $p_\tau = \mathrm{P}(\theta = 1 \,|\, \mathcal{H}_\tau)$. Then the expected unconditional payoffs can be written as

$$\begin{aligned}
\mathrm{E}\left[u_S(\tau,\theta)\right] &= \mathrm{E}\left[\int_0^{\tau} e^{-r_S t}(F_1\theta + (1-\theta)F_0)dt\right] \\
&= \frac{1}{r_S}\mathrm{E}\left[F_1\theta + (1-\theta)F_0 - e^{-r_S\tau}(F_1\theta + (1-\theta)F_0)\right] \\
&= \frac{1}{r_S}\mathrm{E}\left[\mathrm{E}\left[F_1\theta + (1-\theta)F_0 - e^{-r_S\tau}(F_1\theta + (1-\theta)F_0)\Big|\, h_\tau\right]\right] \\
&= \frac{1}{r_S}\mathrm{E}\left[F_1 p_\tau + (1-p_\tau)F_0 - e^{-r_S\tau}(F_1 p_\tau + (1-p_\tau)F_0)\right] \\
&= \frac{1}{r_S}\mathrm{E}\left[F_1 p_\tau + (1-p_\tau)F_0\right] + \frac{1}{r_S}\mathrm{E}\left[v_S(\tau,p_\tau)\right] \\
&= \frac{1}{r_S}\left[F_1 p_0 + (1-p_0)F_0\right] + \frac{1}{r_S}\mathrm{E}\left[v_S(\tau,p_\tau)\right],
\end{aligned}$$

where

$$v_S(\tau,p) = e^{-r_S\tau}\Big[(F_0 - F_1)p - F_0\Big], \tag{2}$$

Analogously for the Receiver, define $v_R(\tau,p)$ as

$$v_R(\tau,p) = e^{-r_R\tau}\Big[(W_0 - W_1)p - W_0\Big]. \tag{3}$$

Note that expected values of v_S and v_R are affine transformations of the expected values of u_S and u_R with a positive slope. Therefore, they define the same preferences over stopping times. For convenience, in the rest of the paper we work with v_R and v_S instead of the original utility functions given by (1).

Remark about Payoffs This last observation allows us to apply our model to a wider class of real-option problems. Motivated by the drug recall application we derived preferences of the players v_S and v_R from the model primitives (F_θ, W_θ). Yet, our model is applicable to any situation in which the players' utilities have a form of

$$v_i(\tau, p) = e^{-r_i \tau}\big[a_i p - b_i\big] \quad a_i \geq 0, b_i > 0 \quad i \in \{R, S\}.$$

For example, parameter b_i can incorporate a fixed cost incurred at the time the option is exercised, as in models of timing of the start of a project. Similarly, in a model of a product manager and a CEO, it can represent the foregone private benefits of the Sender from running a project when the Receiver stops it. The real option model has been applied to a wide array of economic decisions. Our model is relevant to any such situation if it involves a conflict of interest between an agent who has decision rights over stopping and an agent who has access to additional information.

Information: Public News. The state of nature θ is a random variable defined on a probability space $(\Omega, \mathcal{F}, \mathrm{P})$ with $\mathrm{P}(\theta = 1) = p_0$ being the common prior. Neither the Sender nor the Receiver observe at any time the realization of θ.

Over time public news about the state θ arrive via stochastic process X:

$$dX_t = \theta\, dt + \sigma\, dB_t$$

where $B = (B_t)_{t \geq 0}$ is a Standard Brownian Motion with respect to its natural filtration $\mathcal{F}^B$ under measure P. The interpretation of the news process is that it can reflect sales data as well as information from patients and medical care providers about the effects of the product. The process of learning 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.[8]

Information: Sender's Additional Information. Sender can choose to learn about a (typically imperfect) signal ξ independent of the Brownian Motion B. The signal ξ is binary and has conditional distribution matrix $Q_0 = (q_0^{ij})$ with $q_0^{ij} = \mathrm{P}(\xi = j | \theta = i)$. Since the columns of Q_0 add up to one, the whole matrix can be summarized by two numbers $q_0^i = \mathrm{P}(\xi = 1 \,|\, \theta = i)$ for $i = 0, 1$. For our purposes, precision of signal ξ about θ is better summarized by the beliefs this signal can induce in the absence

[8] To be concrete, our model encompasses the following situation: the cumulative payoffs of the two players are $dF(t) = F_\theta dt + dZ_t^F$ and $dW(t) = W_\theta dt + dZ_t^W$ where Z^F and Z^W are standard Brownian Motions. In this case public information X is a sufficient statistic for θ given public realizations of $F(t)$ and $W(t)$.

of public news. This is given by

$$z_0^i = \mathrm{P}(\theta = 1 \,|\, \xi = i) \qquad i = 0, 1$$

We normalize that $z_0^0 < z_0^1$, i.e. the signal $\xi = i$ is more informative about the fundamental state being $\theta = i$. The precision of the signal ξ determines how much information in total is available to the Sender. For example, if $z_0^0 = 0$ and $z_0^1 = 1$ then the Sender has a perfect information technology and she can potentially choose information structure that reveals the state to both agents, however, when $z_0^0 > 0$ ($z_0^1 < 1$) then even perfect knowledge of the signal $\xi = 0$ ($\xi = 1$) leaves both parties uncertain about the true state θ. We describe the information acquisition technology next.

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 public news is first realized, the Sender can then provide information about ξ by "splitting" the prior about ξ. That is, the Sender induces some posterior distribution over ξ subject to the martingale constraint that the average posterior belief has to be equal to the prior. The Sender can commit within a period to an arbitrary distribution, but she cannot commit to future actions. After that, the Receiver makes the decision to either stop or continue the project (that is, to either recall or not the product) and he also cannot commit to future actions. The Sender's strategy is a function of past history of public news and information revealed about ξ up to time t. The Receiver's strategy is additionally a function of the information disclosed about ξ at time t. Markov strategies depend on the history only insofar that they affect joint beliefs about θ and ξ.

Sender's Strategy. Information available to the Sender can be organized in three groups:

1. *Public Information*: $\mathbb{F}^X = (\mathcal{F}_t^X)_{t \geq 0}$ with $\mathcal{F}_t^X = \sigma(X_s, s \leq t)$
2. *Private Signal*: $\mathcal{F}^\xi = \sigma(\xi)$
3. *Randomization Device*: A sufficiently rich[9] filtration $\mathbb{F}^R = (\mathcal{F}_t^R)_{t \geq 0}$ such that $\mathcal{F}_\infty^R = \bigvee_{t=0}^\infty \mathcal{F}_t^R$, $\mathcal{F}_\infty^B = \bigvee_{t=0}^\infty \mathcal{F}_t^B$, and $\mathcal{F}^\theta = \sigma(\theta)$ are mutually independent sigma algebras[10].

[9] For our purposes it will be sufficient to require that this filtration contains sigma-algebras generated by a countable and independent uniform random variables and Poisson processes. In the construction of equilibria, we first assume that this filtration is sufficiently rich and then show by construction what is sufficient.

[10] We also require the original probability space to be rich enough to accommodate Brownian news and independent randomization, i.e. $\mathcal{F}_\infty^R \subset \mathcal{F}$ and $\mathcal{F}_\infty^B \subset \mathcal{F}$.

For technical reasons we require filtrations $\mathbb{F}^X$ and $\mathbb{F}^R$ to satisfy standard conditions, i.e., to contain all P-null sets and to be right-continuous.

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

$$\forall t \quad \mathcal{F}_t^X \subseteq \mathcal{H}_t \subseteq \sigma(\mathcal{F}_t^X, \mathcal{F}^\xi, \mathcal{F}_t^R). \tag{4}$$

Let $\mathscr{H}$ *denote the set of all feasible action profiles.*

Such definition of action profiles allows the Sender to generate informative messages at every time t whose distribution can be contingent on the path of X, realization of ξ, and realization of past messages. Standard restrictions on the filtration $\mathbb{H}$ guarantee that all martingales with respect to this filtration have cádlág versions and first hitting times of such martingales are stopping times with respect to $\mathbb{H}$.[11] That assures that expected payoffs from our strategies are well-defined. The set of all action profiles corresponds to the set of all possible histories of the game in case the Receiver never stopped.[12].

Note that according to this definition an action profile specifies not just 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

$$\mathbb{H}^t = \{\mathcal{H}_s,\, s < t\}.$$

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

$$\mathscr{H}(\mathbb{H}^t) = \left\{\tilde{\mathbb{H}} \in \mathscr{H} \,:\, \tilde{\mathcal{H}}_s = \mathcal{H}_s \quad \forall\, s < t\right\}. \tag{5}$$

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

$$S(\mathbb{H}^t) \in \mathscr{H}(\mathbb{H}^t). \tag{6}$$

[11] For further details see Pollard (2002).

[12] We chose $\mathscr{H}$ 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

At any time and after any history of the game, the Sender'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.

Note that in discrete-time games it is more common to define a strategy as a mapping from all possible histories to current period information disclosure. Iteratively one can then recover the whole sequence of actions in the future periods and specify the information filtration chosen by the Sender. 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

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

In words, if the Sender 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) she should not change her mind given the history $S(\mathbb{H}^t)^{t'} = \{S_s(\mathbb{H}^t),\, s < t'\}$. We will not restrict the Sender to time consistent strategies, allowing her to deviate to non-time consistent strategy after any time (yet, 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, Sender'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 Sender chooses filtration/sigma-algebras of the set Ω and they induce posterior beliefs, so that in discrete time the definitions are equivalent. While it may be possible to define a strategy in continuous time using posterior beliefs, we find it more convenient to define actions in terms of the filtrations as they are a more fundamental object. Also recall that we assume that the filtration chosen by the Sender up to time t is public, so that at any time the Sender and Receiver's beliefs coincide. This captures our assumption that the firm can decide not to learn about some aspects of its drug, but if it learns something, it has to disclose it to the regulator.

We define a *(Markov) state* of the game, π, to be the joint distribution of a pair (θ, ξ). That is, π contains information about the posterior belief about the pay-off relevant state θ, about the realization of the Receiver's signal ξ, and about the correlation between the two. Denote by Π the set of all possible joint distributions of the pair (θ, ξ).

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

$$\pi_{t-} = Law(\theta, \xi \,|\, \mathcal{H}_{t-}).$$

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

$$\pi_{t+s} = Law(\theta, \xi \,|\, S_{t+s}(\mathbb{H}^t))$$

is a Markov process.

Verbally, we define the strategy to be Markov, if the future evolution of posterior beliefs about (θ, ξ), that are induced jointly by public news and the Sender's information disclosure, depend only on current beliefs. Also note that since in our game the Sender never has private information, posterior beliefs are uniquely pinned down for any history $\mathbb{H}^t$, both on and off the equilibrium path.

Receiver's Strategy. We now turn to the Receiver's strategy.

Definition. *A* strategy of the Receiver*, $\mathcal{T}$, is a collection of stopping times. For any action profile of the Sender $\mathbb{H}$ 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 Receiver'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 $\mathbb{H}$.
Similarly to the Sender's strategy, we now define the strategy of the Receiver which is Markov in the state π.

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

$$\forall\, \mathbb{H} \in \mathscr{H},\, t \geq 0 \qquad \mathcal{T}(\mathbb{H}, t) = \inf\{s \geq t \,:\, \pi_s \in \mathbf{T}\}, \tag{7}$$

where $\pi_s = Law(\theta, \xi \,|\, \mathcal{H}_s)$

In words, the Receiver's strategy is Markov if the decision to stop depends only on the beliefs induced by $\mathcal{H}_t$ and nothing else. In particular, if the Sender deviates from the equilibrium path but later follows disclosure that would bring beliefs back to the equilibrium path, the Receiver 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 public news $X_t(w)$ is realized, (ii) second, additional information from $S_t(\mathbb{H}^t)$ is announced (iii) finally, the Receiver upon observing everything prior decides whether to exercise the option or not.

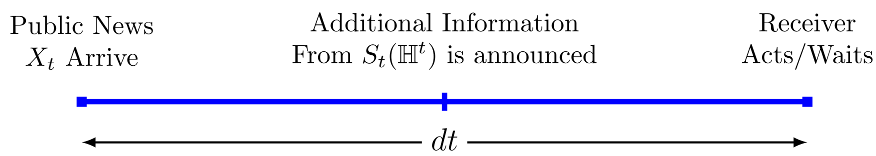

Figure 1: Timing of Events in a Short Interval of Time

2.3 Equilibrium.

Now, we define equilibrium that is Markov in posterior beliefs, π.

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

1. *The Receiver's Strategy* $\mathcal{T}^*$ *is optimal given the anticipated flow of information* S^**, i.e.*

$$\forall\, \mathbb{H} \in \mathscr{H}, \forall\, t \geq 0 \qquad \mathcal{T}^*(\mathbb{H}, t) \in \underset{\tau \in \mathcal{M}(S^*(\mathbb{H}^t)), \tau \geq t}{\arg\max} \mathrm{E}\Big[v_R(\tau, p_\tau)\Big] \tag{8}$$

where $p_s = \mathrm{P}(\theta = 1 \,|\, S^*_s(\mathbb{H}^t))$ *and the maximum is taken over the set of all stopping times* $\mathcal{M}(S^*(\mathbb{H}^t))$ *with respect to filtration* $S^*(\mathbb{H}^t)$*.*

2. *The Sender's strategy* S^* *is optimal given the anticipated stopping rule of the Receiver* $\mathcal{T}^*$*, i.e.*

$$\forall\, \mathbb{H} \in \mathscr{H}, \forall\, t \geq 0 \qquad S^*(\mathbb{H}^t) \in \underset{\tilde{\mathbb{H}} \in \mathscr{H}(\mathbb{H}^t)}{\arg\max} \mathrm{E}\Big[v_S\left(\mathcal{T}^*(\tilde{\mathbb{H}}, t), p_{\mathcal{T}^*(\tilde{\mathbb{H}}, t)}\right)\Big] \tag{9}$$

Notice that in the spirit of subgame perfection, we require both the strategy of the Sender S^* and the strategy of the Receiver $\mathcal{T}^*$ to be optimal after arbitrary histories of the game, even those that do not occur along the equilibrium path.

3 General Analysis

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

No Information Acquisition Thresholds We start with defining two variables: belief thresholds the agents would like to use in case ξ was unavailable, that is, in case they could use only the public news. The optimal stopping decision of the Receiver (or the Sender if she is given control rights) that is based *only on public news* can be characterized by a single threshold[13] p_R (p_S) and a corresponding value function V_R^{NI} (V_S^{NI}) that solve

$$r_i V_i^{NI}(p) = \frac{1}{2} p^2 (1-p)^2 \phi^2 \frac{d^2}{dp^2} V_i^{NI}(p), \tag{10}$$

$$V_i^{NI}(0) = 0, \tag{11}$$

$$V_i^{NI}(p_i) = v_i(0, p_i), \tag{12}$$

$$\frac{d}{dp} V_i^{NI}(p_i) = \frac{d}{dp} v_i(0, p_i) \tag{13}$$

where $\phi = \frac{1}{\sigma}$ and $i \in \{R, S\}$. The above is a second order differential equation with an unknown boundary. Hence it requires three initial conditions to be correctly specified. The smooth pasting condition is a requirement for the threshold p_R to indeed be optimal.

The thresholds p_R and p_S summarize the conflict of interest between the Sender and the Receiver. If $p_R < p_S$ then Receiver wants to undertake the option too early relative to the Sender and vice versa. With imperfectly aligned incentives between the two agents, the Sender searches for ways to structure information sharing to implement an exercise policy more inline with her preferences.

State Variables. While we have introduced above π as the Markov state of the game, in our analysis below we are going to use a sufficient statistic that describes the history of the game. To that end, define

$$x_t = \frac{\phi}{\sigma} X_t - \frac{\phi}{2\sigma} t$$

$$y_t = \mathrm{P}(\xi = 1 \,|\, \mathcal{H}_t)\,.$$

Process x_t is a normalized measure of the public news and up to a linear transformation is equal to the log-likelihood ratio $ln(p_t/(1-p_t))$.

[13] Similar to the standard real options approach in Dixit and Pindyck (2012)

Next, define the most extreme posteriors about θ that the Sender can induce at any time t as

$$z_t^1 = \mathrm{P}(\theta = 1 \,|\, \xi = 1, \mathcal{H}_t)\,, \quad z_t^0 = \mathrm{P}(\theta = 1 \,|\, \xi = 0, \mathcal{H}_t)\,.$$

Lemma 1. *Given prior (p_0, z_0^0, z_0^1), pair (x_t, y_t) is a sufficient statistic for posterior beliefs (p_t, z_t^0, z_t^1).*

Proofs of all formal results are in the Appendix. For intuition, notice that z_t^1 and z_t^0 evolve based only on the new information revealed by the public news. Thus, we can write $z_t^1 = z^1(x_t)$ and $z_t^0 = z^0(x_t)$. The remaining parameter y_t allows to recover belief about θ as $p_t = y_t z^1(x_t) + (1 - y_t) z^0(x_t)$. Thus, the state (p_t, z_t^0, z_t^1) (or, equivalently, state (p_t, q_t^0, q_t^1)) maps naturally into $(z^0(x_t), z^1(x_t), y_t)$ and we can treat the pair (x_t, y_t) as the new state variable, which we do in the rest of the paper.

The benefit of using (x, y) instead of (p_t, q_t^0, q_t^1) is that when the Sender discloses information about ξ, it changes both the beliefs about θ and ξ, while in the (x, y) space, only y changes. That facilitates characterization of the Sender's value function. Revelation of public new over time, on the other hand, affects both x_t and y_t. Define $\hat{y}_t$ to be the non-strategic component of y_t, i.e. the belief about ξ conditional only on the history of *public* news. Then the pair $(x_t, \hat{y}_t)$ evolves as:

$$dx_t = \phi^2 \left(p(x_t, y_t) - \frac{1}{2} \right) dt + \phi dB_t,$$
$$d\hat{y}_t = (z^1(x_t) - z^0(x_t)) \hat{y}_t (1 - \hat{y}_t) \phi dB_t,$$

where $p(x, y) = y z^1(x) + (1 - y) z^0(x)$. Absent any persuasion the evolution of (x_t, y_t) is exactly the same as $(x_t, \hat{y}_t)$. Persuasion allows the Sender to overlay an "arbitrary" martingale on top of $\hat{y}_t$ and introduce discrete jumps in posterior y_t.

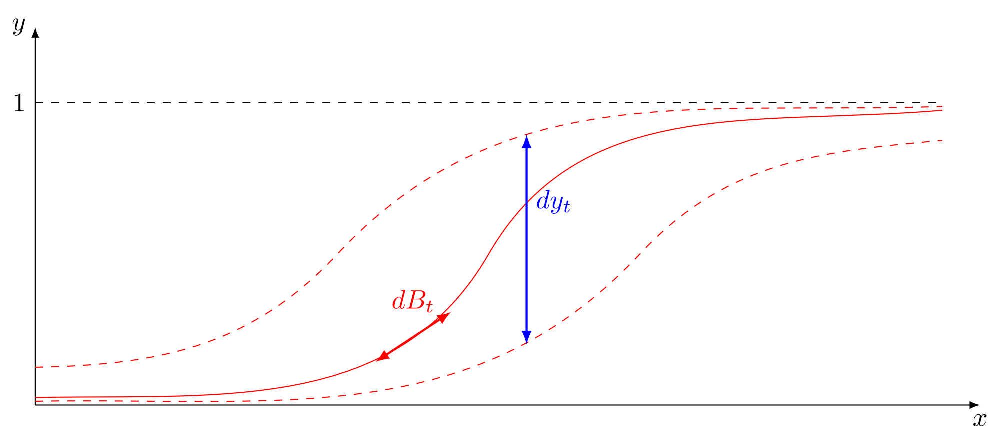

Figure 2: Evolution of beliefs in (x, y) the space due to public news dB_t and persuasion dy_t.

We are looking at Markov Perfect Equilibria in posterior beliefs. Any posterior is summarized by the pair (x_t, y_t). In a given MPE denote the value function of the Receiver and Sender by $V_R(x, y)$ and $V_S(x, y)$ respectively. We now provide characterization of these value functions for given strategies of the opponent.

Sender's Best Response. We start with the Sender's value function and show that it solves an HJB equation and a fixed point of a concavification operator (familiar from Bayesian persuasion literature). Which of the two applies depends on whether beliefs (x, y) are the extreme point of a lower contour set of $V_S(x, y)$. Then, given a value function, we construct a strategy that achieves it. It has the property that in those points where HJB is satisfied the Sender keeps quiet and in points where the concavification applies, there is randomization between two posterior beliefs about y. $V_S(x, y)$ is weakly concave in y since the Sender can always induce any posterior distribution over y subject to the martingale constraint. At points where $V_S(x, y)$ is strictly concave in y it has to be optimal to not reveal any information.
The first step is to describe beliefs (x, y) at which $V_S(x, y)$ is "strictly concave" in y. Given function $V_S(x, y)$ define a correspondence $B : \mathbb{R} \to \mathcal{B}(\mathbb{R}^2)$

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

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 $V_S(x, y)$ for a fixed x. For an equilibrium value function $V_S(x, y)$ define[14]

$$E(x) = \{y : (y, V_S(x, y)) \text{ is an extreme point of } B(x)\}.$$

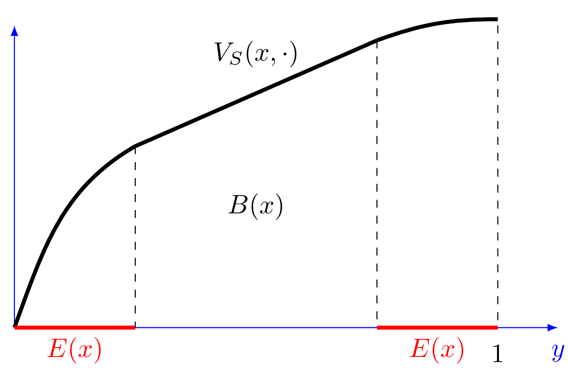

Figure 3: Lower contour set $B(x)$ and the set of extreme points $E(x)$.

For $y \in E(x)$, the Sender's value function is locally strictly concave in y at (x, y) and any signal that

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

she would generate about ξ would decrease her total payoff. Thus for $y \in E(x)$ the Sender must prefer waiting and not obtaining any additional information. For $y \notin E(x)$, the Sender's value function is locally linear in y. Hence she can send the Receiver a signal about ξ and her value function remains unchanged.

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

We show that the if $y \in E(x)$ then Sender's value function satisfies a "waiting" Hamilton-Jacobi-Bellman equation in (x, y).[15] When $y \notin E(x)$ then Sender's value function is a convex combination between her value function when she waits, and her value function when she induces the Receiver to exercise the option:

Lemma 2 (Dynamic Concavification)**.** *In any MPE Sender's value function $V_S(x, y)$ is concave in y and continuous in x. For every $y \in E(x)$ and $(x, y) \notin \mathbf{T}^*$ function $V_S(x, y)$ satisfies*

$$\begin{aligned} r_S V_S(x,y) =& \frac{\partial V_S}{\partial x} \cdot \phi^2 \left(p(x,y) - \frac{1}{2} \right) + \frac{1}{2} \frac{\partial^2 V_S}{\partial^2 x} \phi^2 + \frac{1}{2} \frac{\partial^2 V_S(x,y)}{\partial^2 y} \cdot (z^1(x) - z^0(x))^2 y^2 (1-y)^2 \phi^2 \\ &+ \frac{\partial^2 V_S(x,y)}{\partial x \partial y} \cdot (z^1(x) - z^0(x)) y (1-y) \phi^2 \end{aligned} \tag{14}$$

Otherwise it satisfies

$$V_S(x,y) = cav_\mu \Big[V_S(x,\mu) \cdot \mathbb{1}\{(x,\mu) \notin \mathbf{T}^*\} + v_S(0, z^1(x)\mu + z^0(x)(1-\mu)) \cdot \mathbb{1}\{(x,\mu) \in \mathbf{T}^*\} \Big](y). \tag{15}$$

Notice that equation (15) is not tautological since for in the region $(x, y) \in \mathbf{T}^*$ it might be possible for the Sender to split the belief y into $\{\underline{y}, \overline{y}\}$ and obtain a pay-off $V_S(x, y) > v_S(0, z^1(x)y + z^0(x)(1 - y))$.

Sender Strategy Implementation. Given a function $V_S(x, y)$ that satisfies the conditions outlined in Lemma 2 the Sender has a well-defined strategy attaining this payoff. If $y_t \in E(x_t)$, then the Sender does not disclose information. If $y_t \notin E(x_t)$ define

$$\underline{y}(x_t, y_t) = \sup\{y : y < y_t \text{ and } y \in E(x_t)\}$$
$$\overline{y}(x_t, y_t) = \inf\{y : y > y_t \text{ and } y \in E(x_t)\}$$

The Sender uses a public lottery between beliefs $\underline{y}_t$ and $\overline{y}_t$ with probabilities that satisfy that beliefs about y are a martingale.

[15] In fact, it is an ODE w.r.t. public belief about the fundamental p.

Receiver's Best Response. We now turn to characterizing the Receiver's equilibrium value function and strategy. We show that any MPE is outcome-equivalent to a MPE in which the Receiver follows a threshold strategy for every x: if beliefs about ξ are below the threshold he waits and if they are above the threshold, he stops:

Lemma 3 (Threshold Strategy). *In any MPE, Receiver's value function $V_R(x, y)$ is convex in y. For any x the set where Receiver strictly prefers to wait is*

$$\mathcal{W}(x) = \{y \,:\, V_R(x, y) > v_R(0, z^1(x)y + z^0(x)(1 - y))\}$$

is an interval $[0, l(x))$. Moreover, in any MPE the Receiver without loss of generality follows a threshold strategy: wait when $y < l(x)$, and act when $y > l(x)$.

While the convexity of $V_R(x, y)$ in y is non-trivial, the characterization of the set $\mathcal{W}(x)$ follows immediately from linearity of $v_R(0, z^1(x)y + z^0(x)(1-y))$ in y. The second part of the Lemma 3 shows that regardless of the Receiver's action in the region $\{(x, y) \,:\, y > l(x)\}$ the beliefs spend no time there, thus, one can without loss resolve the Receiver's indifference and focus only on threshold strategies.

Recommendation Principle. The last step in this section is to simplify the relevant set of equilibria further by noticing that without loss of generality we can consider equilibria in which the Sender discloses information about ξ only to affect Receiver's immediate action.

Definition. *An MPE $(\mathcal{T}^*, \mathcal{H}^*)$ is a Recommendation Equilibrium if there exists a message process $m = (m_t)_{t\geq 0}$ with values in $\{0, 1\}$ such that*

1. $\mathcal{S}^*_s(\mathbb{H}^t) = \sigma(\mathcal{H}_{t-}, X_s, m_s; s \geq t)$ *for all* $s \geq t$, $t \geq 0, \mathbb{H} \in \mathscr{H}$,
2. $\mathcal{T}^*(\mathbb{H}, t) = \inf\{s \geq t : m_s = 1\}$ *for all* $t \geq 0$, $\mathbb{H} \in \mathscr{H}$.

That is, in a Recommendation Equilibrium one can think of messages m_t as recommended actions to Receiver with $m_t = 0$ being the recommendation to wait and $m_t = 1$ being the recommendation to exercise the option. Whether such class of equilibria is a restrictive may be a priori unclear. The subtlety arises because a richer sigma algebra $\mathcal{H}_t$ might not necessarily change the instantaneous action of the Receiver but induce a variety of posterior beliefs affecting the continuation game. Nevertheless, we show that it is in Sender's best interest to simply recommend the action to the Receiver and not reveal any additional information. The intuition behind this result it twofold: on one hand, similar to Myerson (1986), making information set of the Receiver coarser relaxes incentive constraints. On the other hand, coarser history up to time t expands the set of feasible actions of the Sender allowing to secure a higher pay-off.

Lemma 4 (Markov Recommendation Principle)**.** *For any MPE* $(\mathcal{T}^*, \mathcal{H}^*)$ *there exists a Recommendation Equilibrium* $(\hat{\mathcal{T}}, \hat{\mathcal{H}})$ *such that* $\hat{\mathcal{T}} = \mathcal{T}^*$.

4 Impatient Receiver $p_S > p_R$

In this section we analyze the game in case $p_S > p_R$, that is, the conflict of interest is such that with only public information, the threshold belief at which the Receiver would choose to stop is lower than the optimal stopping point for the Sender. This case is motivated by our leading example of a drug or a medical device and the assumption that the firm would like to further delay the decision to recall the product at the time the regulator is just indifferent.

For any level of public news x_t the posterior belief p_t about θ is always between

$$z^0(x_t) \leq p_t = y_t z^1(x_t) + (1 - y_t) z^0(x_t) \leq z^1(x_t).$$

By revealing ξ, the Sender could "speed up" or "delay" exercise of the real option depending on the realization of ξ, since it would move the belief closer to $z^1(x_t)$ or $z^0(x_t)$. Consider the incentives of the Sender state by state: if $\xi = 0$ then she would prefer this information to become public as it lowers public beliefs and delays option exercise. If $\xi = 1$, then the Sender wishes to exercise the option when $z^1(x) \geq p_S > p_R$. This suggests that when $z^1(x_t) \geq p_S$, the Sender should prefer to reveal ξ because in this case if it turns out that $\xi = 1$ then the Receiver would stop which is in the Sender's best interest, and if $\xi = 0$ the Receiver will delay as much as it is possible. In other words, two rounds of eliminating dominated strategies imply that it is optimal for the Sender to reveal ξ when $z^1(x)$ reaches p_S. The following lemma established this result formally.

Lemma 5. *In any MPE, if after some history the posterior* $z^1(x_t) \geq p_S$*, then the Sender immediately fully reveals* ξ.

If $z^1_0 = 1$, then Lemma 5 fully characterizes equilibrium outcome: immediate disclosure of ξ at $t = 0$ followed by option exercise as soon as $p_t \geq p_R$. Otherwise, define x^1_S to be the level[16] of the (normalized) public signal at which, conditional on $\xi = 1$, beliefs are p_S:

$$z^1(x^1_S) = p_S.$$

Lemma 5 implies as soon as the (normalized) public signal x reaches in equilibrium level $x_t = x^1_S$ then the Sender immediately reveals ξ and the resulting exercise of the real option has the "compromise"

[16] Equation $z^1(x^1_S) = p_S$ has a unique solution since when $z^1_0 < 1$, then $z^1(\cdot)$ is a strictly increasing function with 0 and 1 being the limits at $-\infty$ and $+\infty$ respectively.

property that we described in the Introduction: Conditional on $\xi = 1$, the option is exercised at the Sender's first-best threshold, p_S; conditional $\xi = 0$, the option is exercised at the Receiver's first-best threshold, p_R.

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

Receiver's Equilibrium Strategy. When deciding whether to exercise the real option the Receiver weights the benefits of obtaining more precise information in the future against the costs of waiting. There are two sources of information for the Receiver: (i) public news and (ii) Sender's persuasion. While the amount of information the Sender shares is an equilibrium object, Lemma 5 states that there is always full information revelation by the Sender when beliefs (x, y) are such that $x \geq x_S^1$. For now, assume that public information and discrete revelation when x_S^1 are *the only* two sources of valuable information for the Receiver.[17]

The value of learning from public news alone is given by[18] $V_R^{NI}(x, y)$ which has been defined in Section 3. If the Receiver additionally learns ξ when public news reach x_S^1, we denote the corresponding value function by $V_R^I(x, y)$. When $x \geq x_S^1$ it is simply given by

$$V_R^I(x, y) = y \cdot v_R(0, z^1(x)) + (1 - y) \cdot V_R^{NI}(x, 0).$$

When $x < x_S^1$ we can write $V_R^I(x, y) = f(p(x, y); x, y)$, where, given the initial pair (x, y), it solves the usual ODE:

$$r_R f(p) = \frac{1}{2}\phi^2 p^2 (1 - p)^2 f''(p) \quad p \in (0, p_R) \cup (\underline{p}(x, y), \overline{p}(x, y)) \tag{16}$$

subject to standard boundary conditions:

$$f(0) = 0, \tag{17}$$

$$f(p_R) = v_R(0, p_R) \quad f'(p_R) = \frac{\partial}{\partial p} v_R(0, p_R) \tag{18}$$

$$f(p) = v_R(0, p) \quad p \in (p_R, \underline{p}(x, y)) \tag{19}$$

that guarantee optimality of the exercise threshold p_R when the beliefs approach the stopping region

[17] In equilibrium communication (about ξ) occurs sometimes in the region $x < x_S^1$ as well. However, as we show later, it does not generate any value for the Receiver, so this fictitious calculation allows us to compute the Receiver's equilibrium payoff.

[18] With a slight abuse of notation we define $V_R^{NI}(x, y) \equiv V_R^{NI}(p(x, y)) = V_R^{NI}(z^1(x)y + z^0(x)(1 - y))$.

$(p_R, \underline{p}(x,y))$ from below, and

$$f(\underline{p}(x,y)) = v_R(0, \underline{p}(x,y)) \quad f'(\underline{p}(x,y)) = \frac{\partial}{\partial p} v_R(0, \underline{p}(x,y)) \tag{20}$$

$$\tag{21}$$

that guarantee optimality of the exercise threshold $\underline{p}(x,y)$ when the beliefs approach the stopping region $(p_R, \underline{p}(x,y))$ from above. The last boundary condition

$$f(\overline{p}(x,y)) = y \cdot v_R(0, p_S) + (1-y) \cdot V_R^{NI}(x_S^1, 0). \tag{22}$$

pins down from the value that Receiver derives from full disclosure of ξ: if $\xi = 1$ beliefs increase to p_S which leads to immediate exercise, if $\xi = 0$, beliefs decrease to $z^0(x_S^1)$ and the Receiver subsequently exercises the real option when his beliefs reach p_R. The threshold $\overline{p}(x,y)$ is an implied belief about θ when $x_t = x_S^1$, formally it is defied by

$$\begin{aligned}
\overline{p}(x,y) &= \frac{p_S \cdot q^0(x,y)}{q^1(x,y)(1-p_S) + q^0(x,y)p_S}, \\
q^0(x,y) &= (1 - z^1(x)) \cdot \frac{y}{1 - yz^1(x) - (1-y)z^0(x)}, \\
q^1(x,y) &= z^1(x) \cdot \frac{y}{yz^1(x) + (1-y)z^0(x)}.
\end{aligned}$$

A solution to this system is illustrated in Figure 4. There are three regions of beliefs: between $(0, p_R)$ the Receiver waits optimally because of his option to wait for public news. In $(p_R, \underline{p}(x,y))$ the Receiver stops immediately. Finally, in $(\underline{p}(x,y), \overline{p}(x,y))$ the Receiver waits because of the option value that ξ will be discretely revealed.

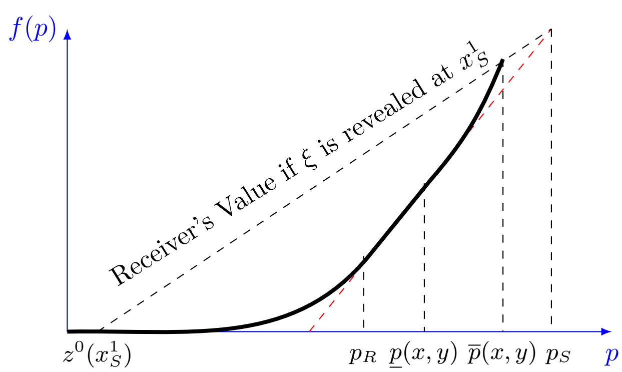

Figure 4: Receiver's expected value $f(p)$ from acting only on public news and the discrete message from the Sender at x_S^1.

When the signal ξ is very informative, or when disagreement between Sender and Receiver is small, then $p_R \geq \underline{p}(x,y)$, so instead of two waiting regions as in Figure 4, there is only one. In that case the Receiver prefers to wait everywhere between $p = 0$ and $\overline{p}(x,y)$ until the Sender reveals ξ. In this case f is given as a solution of the ODE (16) subject to boundary conditions (17) and (22) only. The equilibrium is quite simple – once belief $\overline{p}(x,y)$ is reached, the Sender reveals ξ and Receiver acts accordingly.

When ξ is not that informative or the disagreement is large, so that $p_R < \underline{p}$, the equilibrium is more complicated. In that case, for beliefs in $(p_R, \underline{p})$, the Receiver would prefer to exercise the option, rather than wait for the x_S^1 threshold. For $x \leq x_S^1$ define $l(x)$ as the minimal solution to the equation

$$l(x) = \inf \left\{ y \in [0,1] : v_R(0, z^1(x)y + z^0(x)(1-y)) = V_R^I(x,y) \right\}.$$ [19,20]

Define a candidate Markov strategy of the Receiver, $\mathcal{T}^*$, via the following stopping set $\mathbf{T}^*$:

$$\begin{aligned}\mathbf{T}^* = &\{(x,y) : p_R \leq z^1(x) \leq p_S,\, y > l(x)\} \cup \\ &\{(x,y) : z^1(x) > p_S,\, z^0(x) \leq p_R,\, y = 1\} \cup \\ &\{(x,y) : z^0(x) > p_R\}\end{aligned} \tag{23}$$

Intuitively, according to (23) the Receiver waits if the value of waiting $V_R(x,y)$ is strictly above the value of exercising the option $v_R(0, p(x,y))$. This happens when $z^1(x) < p_R$ or when $p_R \leq z^1(x) \leq p_S$ and $y < l(x)$. Additionally, when $z^1(x) > p_S$ but $z^0(x) \leq p_R$ the Receiver anticipates immediate revelation of ξ, thus, waiting everywhere but $y = 1$ is optimal. Finally, when $z^0(x) > p_R$ the Receiver exercises the option regardless of the realization of ξ. The distinction from a standard real option is that the Receiver now has an endogenous option of waiting for the Sender to reveal ξ.[21] This strategy is depicted in Figure 5.

Sender's Equilibrium Strategy. For $x < x_S^1$ even if the signal is $\xi = 1$ the Sender would prefer to wait longer to observe public news. This means that, intuitively, the Sender should try to minimize the probability of option exercise as much as possible when $x < x_S^1$. We heuristically define the strategy of the Sender via the messages communicated to the Receiver:

- $z^1(x) \geq p_S$. Sender reveals ξ immediately. As we saw before, Sender obtains her highest possible equilibrium payoff this way (she gets her most-preferred stopping when $\xi = 1$ and she delays

[19] We interpret $\inf\{\emptyset\} \equiv +\infty$.

[20] Notice that $p(x, l(x))$ is either p_R or $\underline{p}(x, l(x))$

[21] Additionally the Receiver waits in a measure zero subset of all states in which $V_R(x,y) = v_R(0, p(x,y))$ to make sure that the best response of the Sender is well-defined.

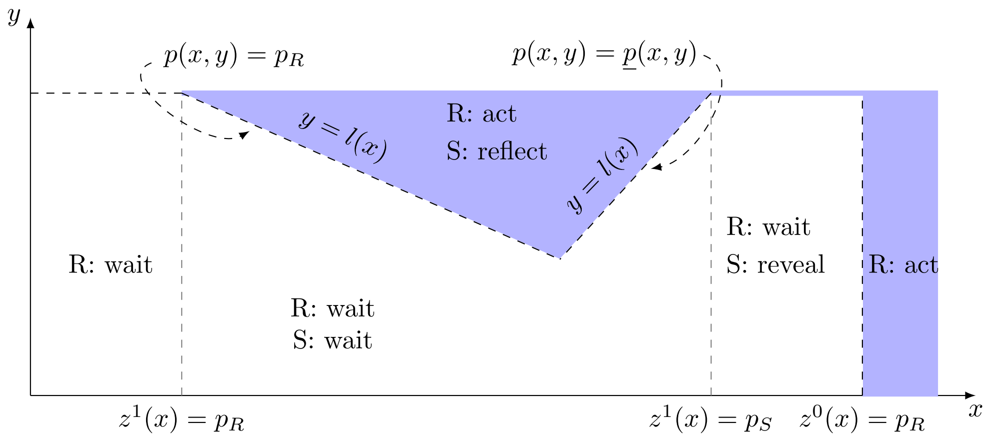

Figure 5: Equilibrium Strategies for Receiver (R) and Sender (S).

the Receiver as much as possible in any equilibrium when $\xi = 0$).

- $z^1(x) < p_S$, $y \le l(x)$. Sender communicates nothing since both parties prefer to wait for more public news.

- $z^1(x) < p_S$, $y > l(x)$. Receiver is about to act with certainty absent additional information. Sender reveals $\xi = 1$ with conditional probability $\alpha(x,y) \in (0,1)$ such that

$$y \cdot \alpha(x,y) + l(x) \cdot (1 - y \cdot \alpha(x,y)) = y \quad \Rightarrow \quad \alpha(x,y) = \frac{y - l(x)}{y(1 - l(x))}$$

leading beliefs about ξ to either jump to 1 or to $l(x)$.

This strategy is described in Figure 5: in words, the Sender reveals ξ completely for any $x \ge x_S^1$. For smaller x the Sender waits until beliefs hit the "triangle" defined by the $y = l(x)$ curve. When they do, at the boundary, he reveals information about ξ that makes the beliefs either reflect at the boundary or jump to $y = 1$. The "left side of the triangle" is defined by $p(x,y) = p_R$ and the "right side" is defined by $p(x,y) = \underline{p}$. As public news arrives, say increasing x, the state of the game moves to the left and up, since public news affect indirectly beliefs about ξ. When the Sender reveals information about ξ, beliefs move vertically in Figure 5.

In the Appendix we explicitly construct the strategy of the Sender S^* that induces a Markov process of beliefs $(x_t, y_t)_{t \ge 0}$ 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, 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 (**Impatient Receiver Equilibrium**). *If $p_S > p_R$, then the pair of strategies $(S^*, \mathcal{T}^*)$ described above constitute a Markov Perfect Equilibrium. Moreover, in this equilibrium $V_R(x,y) \equiv V_R^I(x,y)$.*

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 Sender S^*, Receiver's strategy $\mathcal{T}^*$ is a best response. When $z^1(x) \geq p_S$ the Sender fully discloses ξ, thus, $\mathcal{T}^*$ is optimal as long as the Receiver acts in the set $\{y = 0, z^0(x) > p_R\}$ and $\{y = 1, z^1(x) > p_R\}$. Both of these sets are in the action region $\mathbf{T}^*$.

If $z^1(x) < p_S$ then waiting below $l(x)$ is optimal for the Receiver. Since the conjectured equilibrium involves communication in the region $x < x_S^1$ the value function corresponding to any best response on S^* is weakly above V_R^I defined above. Intuitively, the Receiver can do only better with more information. But $V_R^I(x,y) > v_R(0, p(x,y))$ in this region, thus, waiting is optimal. Next we check that acting is optimal for $y > l(x)$. First, Lemma 9 in the Appendix establishes that for $y > l(x)$ value function $V_R^I(x,y)$ equals to $v_R(0, p(x,y))$, thus, $V_R^I(x,y)$ is linear in y for $y > l(x)$. Since Sender's strategy randomizes over $y = 1$ and $y = l(x)$ in this region, such randomization does not generate any value for the Receiver. In other words

$$
\begin{aligned}
V_R^I(x,y) &= \alpha V_R^I(x,1) + (1-\alpha) V_R^I(x, l(x)) \\
&= \alpha v_R(0, p(x,1)) + (1-\alpha) v_R(0, p(x, l(x)) \\
&= v_R(0, p(x,y)),
\end{aligned}
$$

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

Second, we check whether Sender behaves optimally, given $\mathcal{T}^*$. Lemma 5 establishes that it is optimal for the Sender to reveal ξ immediately when $z^1(x) > p_S$, consistent with S^*. Next, consider the region $\{z^1(x) < p_S, y < l(x)\}$. Recommendation principle implies that any message generated by the Sender must have a realization that occurs with positive probability after which the Receiver stops. In the region $\{x < x_S^1, y < l(x)\}$ the Sender prefers waiting regardless of the realization of ξ, thus it is optimal to stay quiet.

Finally, consider a state $(\hat{x}, \hat{y}) \in \{(x,y) : z^1(x) < p_S, y > l(x)\}$. Absent any information from the Sender, the Receiver acts immediately in $(\hat{x}, \hat{y})$. If the Sender does nothing, she gets her terminal value $v_S(\hat{x}, \hat{y})$. This is strictly dominated by a fully revealing message. Hence persuasion is optimal for the Sender at $(\hat{x}, \hat{y})$. According to the recommendation principle, the Sender will send a binary lottery $\underline{y}, \overline{y}$ such that $1 \geq \overline{y} > l(\hat{x}) \geq \underline{y}$. Suppose $\overline{y} < 1$. State $(\hat{x}, \overline{y})$ has similar problems to $(\hat{x}, \hat{y})$ in the sense that at $(\hat{x}, \overline{y})$ the Sender would reveal all of the information rather than stay quiet. Hence it must be the case that $\overline{y} = 1$. Suppose that $\underline{y} < l(\hat{x})$. This information sharing can be implemented via a compound message. First, the Sender sends beliefs either to $l(\hat{x})$ or to 1. Then, conditional on being at $l(\hat{x})$, the Sender sends beliefs either to $\underline{y}$ or to 1. We know, however, that at $(\hat{x}, l(\hat{x}))$ the Sender prefers to stay quiet since the Receiver is not acting, and hence $\underline{y} < l(\hat{x})$ cannot be optimal. By setting $\underline{y} = l(x)$ the Sender 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 Sender at the boundary $y = l(x)$. We show that the Sender generates a messages that reveals $\xi = 1$ with positive *intensity* making the belief process reflect from the boundary of the action region conditional on sending a negative "message". □

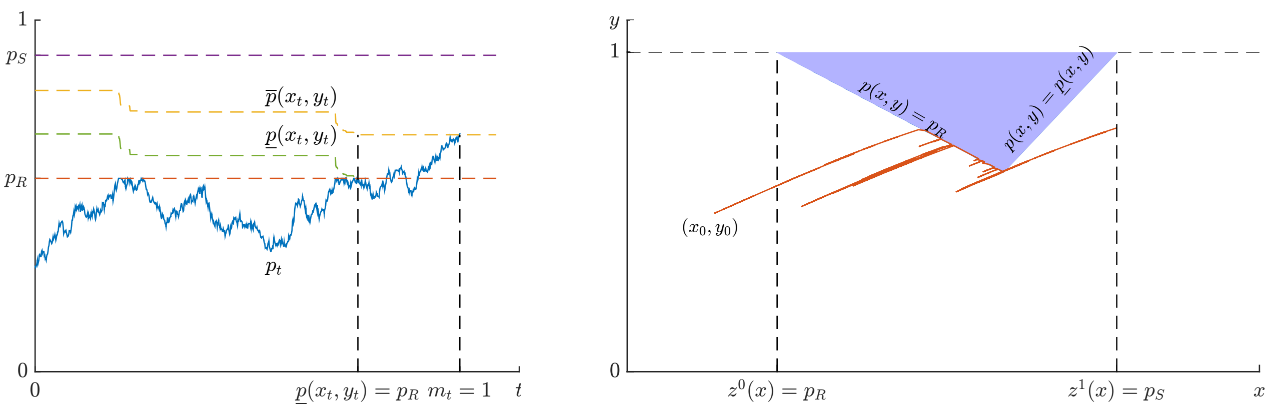

(a) Successful Persuasion in p space

(b) Successful Persuasion in (x, y) space

Figure 6: Equilibrium dynamics with $\xi = 1$ in the lower waiting region.

The strategies $(S^*, \mathcal{T}^*)$ generate the dynamics show on Figure 6. In this example we start with $\underline{p}(x_0, y_0) > p_R$ meaning that the Receiver would rather exercise the option at p_R, rather than wait for the Sender's information to be revealed at x_S^1. Whenever the belief p_t is below the Receiver's "autarky" threshold, p_R, the Sender stays quiet and the belief is updated based on the public news only. At the threshold $p(x, y) = p_R$ the beliefs, conditional on a "negative" message reflect down inducing the Receiver to wait further (or they jump to $y = 1$ with intensity to satisfy the Bayes rule). As y_t decreases due to persuasion the belief threshold at which the Sender is willing to reveal ξ approaches p_R from above. A decrease in the revelation threshold increases the value of the Receiver to wait for the full information revelation. When this value is high enough to incentivize the Receiver to wait past p_R, the Sender stops persuasion (this point corresponds to the "tip of the triangle" in

Figure 6(b)). Further waiting ensues until the public news reach a new higher threshold x_S^1 at which the Sender fully reveals ξ.

The next proposition shows that what we just constructed is essentially a unique Markov Perfect Equilibrium of our game when the Sender is more patient than the Receiver.

Proposition 2 (**Uniqueness**). *If $p_S > p_R$, then $(S^*, \mathcal{T}^*)$ is the essentially*[22] *unique MPE.*

We establish uniqueness by first eliminating dominated strategies of the players. Lemma 5 shows that in any equilibrium the Sender fully discloses ξ as soon as $z^1(x) > p_S$. This implies that it is dominant for the Receiver to wait when $z^1(x) < p_S$ and $y < l(x)$ in any equilibrium. Next, we notice that in the region $\{z^1(x) < p_S, y < l(x)\}$ the Sender prefers to wait regardless of ξ, thus, in any equilibrium she stays quiet.

The argument above pins down the strategies of the Sender and Receiver everywhere except the region $\{z^1(x) < p_S, y \geq l(y)\}$. The proof in the Appendix shows that in any equilibrium the strategies in the 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 Receiver $\mathbf{A} = \{(x, y) \in \mathbf{T} : z^1(x) < p_S\}$ is strictly inside of the "triangle" $\{z^1(x) < p_S, y \geq l(y)\}$. In this case it is optimal (similar to the argument in the proof of Proposition 1) for the Sender to stay quiet outside of $\mathbf{A}$. Since the Receiver is waiting outside of $\mathbf{A}$ his payoff is weakly above $v_R(0, p(x, y))$, at the same time, since he is receiving less information compared to S^*, his payoff should be weakly below $V_R(x, y)$. Thus, in the region $\{z^1(x) < p_S, y \geq l(y)\} \setminus \mathbf{A}$ the payoff of the Receiver should be equal to $V_R(x, y) = v_R(0, p(x.y))$. This payoff is linear in belief p, which is incompatible with strictly positive waiting time in this region (since the payoff from waiting it strictly convex in p). Since in any equilibrium the beliefs spend no time in the region $\{z^1(x) < p_S, y \geq l(y)\}$, it is outcome-equivalent to $(S^*, \mathcal{T}^*)$.

Remark. Earlier in this section we have constructed value function $V_R^I(x, y)$ fictitiously assuming that the Sender reveals information only at x_S^1, while in the equilibrium we just described the Sender actually reveals also information at the boundary $l(x)$. At that boundary, beliefs are either p_R or $\underline{p}$ (at the left and right sides of the "triangle," respectively). Even though typically additional information increases payoff of a decision maker who faces a real option problem, in our equilibrium the information release by the firm is timed in a way that it affects the regulator's actions but does not affect his payoff. Technically, even though the value of the Receiver is convex in y, it is locally linear at the boundaries where information is revealed (recall the smooth-pasting conditions). Since

[22]By "essentially" we mean that for any initial conditions (x_0, y_0) the distribution of outcome posterior beliefs of any MPE coincides with $\text{Law}(x_\tau, y_\tau, \tau)$ implied by $(S^*, \mathcal{T}^*)$.

the equilibrium information disclosure reflects beliefs only locally, it keeps the Receiver indifferent and hence does not create any value for him.

Remark. We want to highlight two features of the equilibrium. If the prior beliefs and parameters are such that the beliefs on the equilibrium path do not cross the $l(x)$ "triangle" in Figure 5, then the outcome has the compromise property we discussed. That has two consequences. First, anytime $\xi = 1$ is revealed, the firm does not need to be compelled to recall the drug - it is happy to do so voluntarily. That is not true in the "pipetting region" - there upon revealing $\xi = 1$ the regulator decides to recall the draft and that is against the wishes of the firm. Second, the "compromise" property implies that the equilibrium is robust to replacing the verifiable information about ξ with cheap talk. The reason is that at x_S^1 the firm has aligned incentives with the regulator to reveal the truth. If the firm does not perform self-recall past x_S^1, the regulator can credibly infer that $\xi = 0$ and stop accordingly. That is again not true in the "pipetting" region: there if the firm is about to reveal $\xi = 1$, it would prefer to mislead (at least temporarily) the regulator that $\xi = 0$.

4.1 Comparative Statics

We finish this section discussing some comparative statics of the equilibrium.

Comparative static on the size of the triangle. Keeping the assumption $p_S > p_R$, define by (x^*, y^*) to be the unique solution of $p_R = \underline{p}(x^*, y^*)$. Visually, point (x^*, y^*) is the lower tip of the Receiver's action "triangle" depicted in Figure 5. This point is important because it affects equilibrium dynamics: arrival of public news shifts beliefs along non-intersecting curves in the (x, y) space as shown in Figure 2. If such curve passing through the prior (x_0, y_0) does not intersect the triangle, then along the equilibrium path information is revealed once and fully when $x_t = x_S^1$ and the timing of the option exercise features a "compromise" property. If the initial curve crosses the "triangle", along the equilibrium path the Sender initially "pipets" information until either bad news arrives revealing $\xi = 1$ or beliefs reach the tip of the triangle, at which point both players wait until x_S^1.

Dependence on Sender's Information Precision. The next lemma shows that the amount of delay in equilibrium is increasing with the precision of Sender's signal ξ.

Lemma 6. *The Receiver is willing to wait longer when the Sender's signal ξ is more precise:*

$$\frac{\partial y^*}{\partial z_0^1} > 0, \qquad \frac{\partial y^*}{\partial z_0^0} < 0.$$

As z_0^1 goes up Sender's signal ξ becomes more precise about the state $\theta = 1$. As a result, the information revelation threshold x_S^1 decreases since, conditional on $\xi = 1$, the Sender is requires less positive public news to reach a belief p_S. A lower x_S^1, in turn, corresponds to a smaller Receiver's cost of waiting for the discrete information revelation. At (x^*, y^*) the Receiver is exactly indifferent between acting immediately and waiting for x_S^1. As z_0^1 increases, Receiver's option value to wait at (x^*, y^*) increasing leading him to prefer waiting rather than acting. This shifts the tip of the triangle up.

As z_0^0 decreases, Sender's signal ξ becomes more precise about the state $\theta = 0$. This has no effect on her incentives to reveal information at x_S^1. However it makes the information disclosed that is discretely disclosed of higher quality. Because Receiver's value function is strictly convex in beliefs, this increases the payoff of waiting for the Receiver. Graphically, as z_0^0 decreases, the tip of the triangle (x^*, y^*) increases.

Level of Disagreement. In our model disagreement between the Sender and Receiver is summarized by the distance between p_R and p_S, the no-additional-information thresholds.

Lemma 7. *As the disagreement of preferences increases the Receiver is willing to wait less before acting:*

$$\frac{\partial y^*}{\partial p_S} < 0$$

The intuition is simply that a higher p_S implies a higher x_S^1 since the firm is going to wait longer to self-recall the drug. The Receiver now has to wait longer for the discrete revelation of information. Naturally, this decreases the value of waiting to learn ξ and expands the range of beliefs for which the Receiver is willing to exercise the option in equilibrium.

5 Impatient Sender $p_S \leq p_R$

We now turn to the opposite case, when it is the Receiver who would like to exercise the real option sooner than the Sender, which is captured by the assumption $p_S \leq p_R$. 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 that either because the project gives the agent private benefits or that the agent does not fully internalize the fixed cost of starting it, she would like to start her project sooner than the principal.

5.1 Immediate Revelation Equilibrium

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

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

In such equilibrium the Receiver achieves the first best by adopting a very "strong bargaining position": he threatens not to exercise the option until either all information about ξ is revealed, or the public news render additional information that could be obtained by the Sender irrelevant, i.e. $z^0(x) \geq p_R$. Given the equilibrium strategy of the Sender such threat is credible, since waiting for the information to be released at the next instance is costless in continuous time. The Sender faces a tough choice: either to reveal ξ immediately, or to wait until $z^0(x) = p_R$; any partial information revelation does not affect the timing of option exercise. If the Sender reveals ξ, then the option is exercised either at $z^1(x) = p_R$ (if $\xi = 1$) or at $z^0(x) = p_R$ (if $\xi = 0$). Since she prefers the option to be exercised earlier and $z^0(x) < z^1(x)$ the Sender is better off revealing ξ at time 0.
Note that the same reasoning does not apply to the previous case of a patient Sender. In that case, if the Receiver could commit to a strategy "I stop in the next second unless you reveal all information about ξ," the Sender's best response would be to reveal such additional information. But that threat is not credible: if the Sender does not provide information, it is not credible for the Receiver to stop before beliefs reach p_R.
In the next sections we show that under some parametric assumptions this is the unique Markov Perfect Equilibrium and that in general, Receiver's decision threshold is always higher than absent persuasion, which implies that in these cases the Sender is worse off as a result of having access to information about ξ.

5.2 Equilibrium Uniqueness

First we show that in any Recommendation Equilibrium if $m_t = 1$ is an informative message, then it fully reveals $\xi = 1$. As a consequence, in equilibrium the option is exercised either when $y_t = 1$ or after a period of no persuasion.

Lemma 8. *In any Recommendation Equilibrium if* $\mathrm{P}_t(\xi = 1|m_t = 1) > \mathrm{P}_t(\xi = 1|m_t = 0)$ *then* $\mathrm{P}_t(\xi = 1|m_t = 1) = 1$.

The above lemma highlights the impact of public news on the equilibrium behavior. The proof constructs an unraveling argument that we illustrate with an example in which signal ξ is perfectly informative of θ. Suppose the Receiver's strategy is to wait when belief is below some threshold $\overline{p}$ and act otherwise. We show that $\overline{p} < 1$ cannot be part of a Markov Perfect Equilibrium. Similar to the judge example of Kamenica and Gentzkow (2011) the Sender might want to design a lottery which sends Receiver's beliefs from $p < \overline{p}$ to either 0 or $\overline{p}$. In order for this to constitute an equilibrium, Receiver's decision to exercise the option at $\overline{p}$ must be optimal, even though for $p < \overline{p}$ he expects to receive additional information from the Sender. If the receiver deviates and waits dt he obtains public news. If the news push beliefs up, then delayed exercise would result in a *second order loss* compared to immediate exercise, since $\bar{p}$ was optimally chosen in the first place. However, if the news push beliefs down, the Sender would be tempted to persuade him 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_R and the Receiver does not end up exercising the option when $\theta = 0$. Hence it cannot be an equilibrium where Receiver acts at a threshold $\overline{p} < 1$.

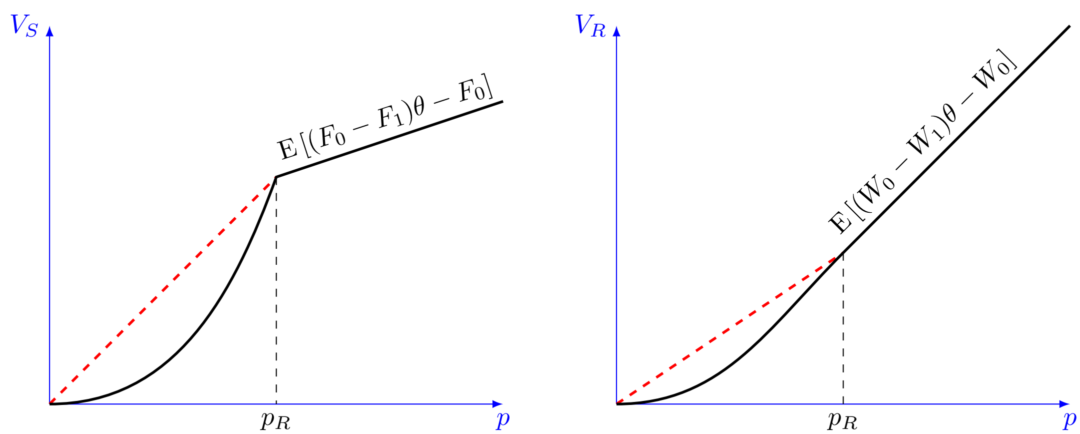

Figure 7: Effect of concavification both on Sender's and Receiver's value functions expressed as functions of public belief p when Sender has perfect information $\xi = \theta$.

Figure 7 provides a graphical illustration of the core idea behind Lemma 8. It shows the effect of one-step concavification of the Sender's value function that pools information about $\xi = 1$. As a result Receiver's value function, that previously satisfied a smooth pasting condition at p_R, obtains a kink. This kink in presence of Brownian news makes it optimal for the Receiver to wait and pushes the exercise threshold higher than p_R. Technically speaking, persuasion using a pooling message implies that the value function of the Receiver has to satisfy both the "reflecting"[23] and smooth pasting[24] boundary conditions at the optimal exercise threshold which is impossible.

Next we show that the for sufficiently low prior $z^0(x_0) < p_R$ the outcome of any MPE coincides with the outcome of equilibrium of Proposition 3:

[23]See Harrison (2013) for details.
[24]See Dixit and Pindyck (2012) for details.

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

We prove the first part of the proposition by showing that given the initial condition $z^1(x_0) \leq p_R$ in any MPE the option is exercised only when ξ is fully revealed. Since the Sender prefers early exercise immediate revelation of ξ guarantees that the Receiver stops exactly when $z^1(x) = p_R$ or $z^0(x) = p_R$. To see why the second part of Proposition 4 holds, one has to notice that full revelation of ξ at x_R^1 such that $z^1(x_R^1) = p_R$ creates additional option value for the Receiver when the beliefs approach this threshold from above. Hence he will no longer prefer to stop at $p(x, y) = p_R$ because his value from waiting is now strictly higher.

Suppose now that Receiver's prior beliefs at $t = 0$ are p_R and the Sender would prefer him to exercise the option. Then, absent persuasion, this would be the outcome of the game. The Receiver, however, expects the Sender to reveal more information when public beliefs drop. Hence he will no longer wish to exercise the option at p_R. This hurts the value of the Sender since, under some parameters, unconditional exercise would result in a first best outcome had they started with sufficiently high beliefs.

Corollary 1. *Suppose $V_S^{NI}(z^0(x_0)) < (F_0 - F_1)z^0(x_0) - F_1$ at $t = 0$ and $y_0 \in (0, 1)$. If initial beliefs are $p_0 = p_R$ then Sender's value in any MPE is strictly lower than his value if the Receiver acted just based on public news.*

Proof. Absent persuasion the Receiver would have exercised the real option at p_R which would have been the best outcome for the Sender. Due to persuasion, however, Receiver's decision threshold is strictly higher than p_R. □

5.3 Special Case: $z^0(x) \equiv 0$

In this section we prove the uniqueness of the equilibrium for a special case when the Receiver can obtain a perfect information about $\theta = 0$:

Proposition 5. *Suppose $z^0(x) \equiv 0$ then the equilibrium of Proposition 3 is essentially unique.*

When $z^0 = 0$ and the Sender learns that $\xi = 0$, then, the Sender knows that the state of nature is $\theta = 0$. If $F_0 < 0$, the Sender does not wish to exercise the option. In that case it is a dominant strategy for her to disclose all of information immediately: if $\xi = 0$ she gets her preferred action and when $\xi = 1$ she gets the Receiver speed up exercise as much as possible.

If, however, she prefers pooling ($F_0 > 0$), then for any interior decision threshold $y = l(x)$ of the Receiver there is always a point (x^*, y^*), such that she always prefers to send a pooling message to induce action. Given such a best response of the Sender to threshold $l(\cdot)$ we have a contradiction with Lemma 8. The Receiver will expect such a pooling message at x^* and his threshold will be strictly higher than $l(x^*)$. Hence a threshold $0 < l(x) < 1$ cannot be part of any Markov Perfect Equilibrium.

Remark. This is a stark difference to a result of persuasion with commitment. It can be shown that when $z_0^0 = 0$ the Sender engages in immediate persuasion by doing a lottery between z_0^0 and $\min(z_0^1, p_R)$.

6 Conclusion

In this paper we study a dynamic Bayesian persuasion game in which the Sender and the Receiver learn jointly about the state of the nature from a stream of public news. The Sender can acquire additional informative signals; she decides what signals and when to acquire but has no long-term commitment power. The Receiver, upon observing public news and additional information generated by the Sender, decides when to take an irreversible action (exercise a real option) and also have no commitment power to future stopping decisions. This analysis is motivated by many problems where one agent has decision power over a real option (e.g. a regulator) and another agent, with partially misaligned preferences (e.g. a medical device supplier), can acquire additional information to influence the decision-maker.

We show that there exists a unique Markov Perfect equilibrium and characterize equilibrium dynamics of information provision when the Sender prefers later exercise. If misalignment of incentives is small, in equilibrium the Sender acquires information once and it is fully revealing. Anticipating a discrete information disclosure, the Receiver is willing to wait past a belief threshold at which he would stop with only public news. Depending on the information disclosed, the option is exercised at either the Senders optimal threshold, in which case the Receiver ex-post regrets waiting for too long, or at the Receivers optimal threshold. Thus, the timing of option exercise features a compromise property. If misalignment of incentives is large, Sender [initially] resorts to pipetting of information over time. Pipetting results in belief process that either reflects from the action threshold of the Receiver or discretely jumps, triggering immediate exercise of the option. Such information provision changes the timing of the option exercise and allows the Sender to reach the compromise region with positive probability, however, it does not generate any value for the Receiver.

When the Sender prefers earlier exercise her inability to stop persuading undermines her ability to speed up the option exercise time. We show that there always exists an equilibrium in which the Sender discloses all information at time zero, and in several cases this equilibrium is unique. In contrast to the static environment of Kamenica and Gentzkow (2011) pooling of information unravels due to the option of the Receiver to wait and obtain more information from the Sender in the immediate future. Moreover, we show that for some parameter values the Sender is worse off as a result of having access to additional information in all Markov equilibria.

Appendix

Proof of Lemma 1. The most extreme beliefs that the sender can induce at any time t can be expressed as a function of initial conditions and x_t:

$$z_t^i = \mathrm{P}_t(\theta = 1|\xi = i) = \frac{z_0^i e^{x_t}}{z_0^i e^{x_t} + (1 - z_0^i)} = z^i(x_t)$$

The triplet (p_t, q_t^0, q_t^1) maps naturally into $(z^0(x_t), z^1(x_t), y_t)$ via

$$\begin{aligned} p_t &= y_t \cdot z^1(x_t) + (1 - y_t) \cdot z^0(x_t), \\ q_t^1 &= \frac{z^1(x_t) y_t}{p_t,} \\ q_t^0 &= \frac{(1 - z^1(x_t)) y_t}{1 - p_t}. \end{aligned}$$

Public news and persuasion are subsequently captured by the pair (x_t, y_t). □

Proof of Lemma 2 . Suppose $V_S(x, y)$ corresponds to Sender's equilibrium value function. First, it must satisfy

$$V_S(x, y) = \mathrm{cav}_y V_S(x, y) \tag{A.1}$$

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

Second, the Sender has the option to choose the time when he starts persuading the receiver:

$$V_S(x, y) = \sup_{\kappa} \mathrm{E}_{(x,y)} \left[e^{-r\tau} \cdot V_S(x_{\tau\wedge\kappa}, y_{\tau\wedge\kappa}) \right] \tag{A.2}$$

where

$$\tau = \inf\{t : (x_t, y_t) \in T^*\}$$

Moreover because the simple solution of $\kappa = 0$ achieves the optimum, we have

$$V_S(x, y) = \max_{\kappa} \mathrm{E}_{(x,y)} \left[e^{-r\tau} \cdot V_S(x_{\tau\wedge\kappa}), y_{\tau\wedge\kappa} \right]$$

The solution must be continuous along the path of no persuasion $(x, y(x))$. Because it is continuous in y it implies that it is continuous in x. Suppose $y \in E(x)$. It implies that there is no action and no persuasion being conducted by the Sender (because of strict concavity).

Suppose that $y \in E(x)$. This implies that in the vicinity of (x, y) there is no persuasion. Hence for

$$\hat{\kappa} = \inf\{t : y_t \notin E(x_t)\}$$

we have

$$V_S(x, y) = \mathrm{E}_{(x,y)}\left[e^{-r\kappa\wedge\tau} \cdot V_S(x_{\tau\wedge\hat{\kappa}}), y_{\tau\wedge\hat{\kappa}}\right]$$

Hence V_S must satisfy (14). □

Proof of Lemma 3. First we show a property of the Sender's best response to any strategy. Suppose $y_0^L < y_0^H$ and $y_0^L, y_0^H \in E(x)$. There exists an information sharing strategy of the Sender delivering him his equilibrium value such that $y_t^H \geq y_t^L$. From the contrary define (we can define it with some small gap Δ just as way to make it a non-strict inequality)

$$\kappa = \inf\{t < \kappa : Y_t^L > Y_t^H\}$$

Denote the random set $\tilde{E} = E(x_\kappa)$. Then

$$\lim_{\varepsilon \to 0, \varepsilon > 0} Y_{\kappa-\varepsilon}^L = Y_{\kappa-}^L \notin \tilde{E}$$

Process Y_t^L is continuous when $Y_t^L \in E(x_t)$. Suppose that $Y_{\kappa-}^H \in \tilde{E}$. Then according to the Sender's strategy it will be the case that $Y_\kappa^L \leq Y_\kappa^H$ and we have a contradiction. Hence it must be the case that $Y_{\kappa-}^H \notin \tilde{E}$. This implies that both along both belief paths L and H in state x_τ the Sender creates signals. It is clear that we must focus on the case when $\forall y \in [Y_{\kappa-}^L, Y_{\kappa-}^H]$ we have $y \notin \tilde{E}$. Otherwise Sender's signals would not alter the ranking of beliefs.

The lottery is going to be between two points $\underline{y}$ and $\overline{y}$. The probability that message $\overline{y}$ is sent by the Sender is given by

$$\alpha^L = \frac{Y_{\kappa-}^L - \underline{y}}{\overline{y} - \underline{y}}$$
$$\alpha^H = \frac{Y_{\kappa-}^H - \underline{y}}{\overline{y} - \underline{y}}$$

The Sender can implement his strategy by delivering two messages. One that sends messages from Y^L to $\underline{y}$ and Y^H. It does not hurt the Sender and the Receiver's continuation value does not change. Hence $Y_t^H \geq Y_t^L$. This implies that if the Receiver has a signal about y_0 at $t = 0$, this leads to a lottery over all subsequent paths. Because he can ignore this information when making the decision, this leads to convexity in his value function in equilibrium.

It is not obvious that the above constitutes a threshold rule. Suppose $V_R(x,1) > v_R(0, z^1(x))$. This implies that the receiver prefers waiting even if all of the information is revealed. The Receiver solves his real option problem and the above is equivalent to $z^1(x) < p_R$. In this case $z^1(x)y + z^0(x)(1-y) < p_R$ for $y < 1$ as well. Hence $V_R(x,y) > v_R(0, z^1(x)y + z^0(x)(1-y))$ for $y \in [0,1]$. If $V_R(x,1) \leq v_R(0, z^1(x))$, then $l(x) = \sup\{y : V_R(x,y) > v_R(0, z^1(x)y + z^0(x)(1-y))\}$ is well defined and

$$V_R(x,y) > v_R(0, z^1(x)y + z^0(x)(1-y)) \quad \text{for } y < l(x),$$
$$V_R(x,y) = v_R(0, z^1(x)y + z^0(x)(1-y)) \quad \text{for } y \geq l(x).$$

Clearly the strategy of the Receiver "wait if $y_t \geq l(x_t)$ and act otherwise" is a best response, however, there are many others since in the region $y > l(x)$ the Receiver is indifferent between acting and not. We now show that in any equilibrium the beliefs spend measure zero time in the region $1 > y > l(x)$, thus, it is without loss of generality to consider a threshold best response $\tilde{\mathbf{T}}$.

Suppose the contrary, and consider a given (x_0, y_0) with $y_0 > l(x_0)$ and $T(x_0, y_0) = 0$. Define E to be the set of pairs (x,y) that correspond to the extreme point of V_S, i.e.

$$E = \{(x,y) : y \in E(x)\}.$$

If $(x_0, y_0) \in \text{int}(E)$, then starting at (x_0, y_0) a positive waiting ensues with only the public news arriving. In this region V_S is linear in p but has to satisfy a second order ODE and be strictly convex - a contradiction. If $(x_0, y_0) \notin \text{int}E$, then $y_0 \in [\underline{y}(x_0, y_0), \overline{y}(x_0, y_0)]$. Notice that $\underline{y}(x_0, y_0) \leq l(x_0)$ otherwise, there would be some positive quiet time starting at some (x_0, y_0') with $y_0' \in (l(x_0), \underline{y}(x_0, y_0))$ which we have shown to be a contradiction. Similarly $\overline{y}(x_0, y_0) = 1$, otherwise, there would is some quite time starting at some (x_0, y_0') with $y_0' > \overline{y}(x_0, y_0)$. Thus, starting at (x_0, y_0) the sender immediately sends beliefs out of the region $1 > y > l(x)$, and we can put $\tilde{\mathbf{T}}(x_0, y_0) = 1$ without changing the outcome of the equilibrium.

□

Proof of Lemma 4. In the existing equilibrium $(\mathcal{T}^*, \mathcal{H})$ process (x_t, y_t) is Markov. Define

$$m_t = \max_{s \leq t} \mathbf{T}(x_s, y_s)$$

Define

$$\tau = \inf\{t : m_t = 1\}$$

The new belief process is given by

$$\hat{y}_t = \mathrm{P}(\xi = 1 | x_t, m_t = 0)$$

It satisfies the SDE

$$d\hat{y}_t = (z^1(x_t) - z^0(x_t))\hat{y}_t(1-\hat{y}_t)\phi dB_t + (\hat{y}(t, x_t, m_t = 1) - \hat{y}_t)(dN_t - \lambda(t, x_t, m_t = 0)\, dt)$$

This implies process $(x_t, \hat{y}_t, t)$ is Markov. This equilibrium is not MPE in beliefs yet because of the time dependency of the intensity $\lambda(t, x_t, m_t = 0)$. Define the corresponding value function $\hat{V}_S(t, \hat{y}, x)$. Suppose

$$\max_t V_S(t, \hat{y}, x) > \min_t V_S(t, \hat{y}, x)$$

Then the Sender has a profitable deviation in the original equilibrium. Hence $\hat{V}_S(t, \hat{y}, x) = \hat{V}_S(0, \hat{y}, x)$. Define the optimal strategy as

$$\lambda^*(x, m_t = 0) = \lambda(0, x, m = 0)$$

This leads to the same expected value for the Sender and a Markov recommendation strategy. □

Proof of Lemma 5. If $\xi(\omega) = 1$, then disclosing it results in immediate exercise of the option by the receiver, which is the first best outcome for the sender conditional on $\xi(\omega) = 1$. Any delay in such a state ω decreases sender's pay-off conditional on $\xi = 1$.
Consider the case with $\xi(\omega) = 0$. Let p_τ denote the equilibrium level of beliefs about θ at the moment of option exercise. It must be that $p_\tau \geq \mathrm{P}(\theta = 1 | p_\tau, \xi = 0)$. Without loss, this inequality is strict in some states of the world. In those states revealing $\xi = 0$ at $t = \tau$ lowers the belief of the receiver and weakly delays option exercise. If $\mathrm{P}(\theta = 1 | p_\tau, \xi = 0) > p_R$, then the option is exercised by the receiver regardless of the sender's information, hence, revealing is weakly optimal. If $\mathrm{P}(\theta = 1 | p_\tau, \xi = 0) \leq p_R < p_S$, delaying option exercise in this region improves sender's conditional payoff pay-off – a contradiction. Thus, $p_\tau = \mathrm{P}(\theta = 1 | p_\tau, \xi = 0)$ i.e. negative signal $\xi = 0$ is fully disclosed in any equilibrium. Revealing $\xi = 0$ at $t = 0$ path-by-path results in the same exercise time τ and induces the same pay-off to both receiver and the sender, conditional on $\xi = 0$. □

Lemma 9. *For any fixed x there exists $y^H(x)$ such that $V_R^{NI}(x, y) \geq V_R^I(x, y)$ when $y \geq y^H(x)$.*

Proof of Lemma 9. It is sufficient to show that

$$\frac{\partial}{\partial y}\left[V_R^I(x, y) - V_R^{NI}(x, y)\right]\Big|_{p(x,y) = \underline{p}} = 0.$$

Consider a function f that solves

$$
\begin{aligned}
r_R f(p) &= \frac{1}{2}\phi^2 p^2 (1-p)^2 f''(p) \quad p \in (a,b) \\
f(a) &= v_R(0,a) \\
f(b) &= y \cdot v_R(0,p_S) + (1-y) V_R^{NI}(\bar{x},0)
\end{aligned}
$$

To show the dependence on the boundary conditions can write $f(p;a,b)$. Notice that $f(p;\underline{p},\overline{p})$ coincides with V_R^I as a function of p in the relevant region. Recall that

$$
\frac{\partial}{\partial p} f(\underline{p};\underline{p},\overline{p}) = \frac{\partial}{\partial p} v_R(0,\underline{p}).
$$

Next increase y to $y+dy$, and consider the new triplet $p' = p + dp(y)$ $\underline{p}' = \underline{p} + dp(y)$ and $\overline{p}' = \overline{p} + d\overline{p}(y)$. We can check whether the smoothpasting condition holds and the new boundary $\underline{p}'$:

$$
\begin{aligned}
\frac{\partial}{\partial p} f(\underline{p}';\underline{p}',\overline{p}') &= \frac{\partial}{\partial p}\left(f(\underline{p};\underline{p},\overline{p}) + \frac{\partial}{\partial p} f(\underline{p};\underline{p},\overline{p}) dp(y) + \frac{\partial}{\partial \underline{p}} f(\underline{p};\underline{p},\overline{p}) dp(y) + \frac{\partial}{\partial \overline{p}} f(\underline{p};\underline{p},\overline{p}) d\overline{p}(y) \right) \\
&= \frac{\partial}{\partial p}\left(f(\underline{p};\underline{p},\overline{p}) + \frac{\partial}{\partial \overline{p}} f(\underline{p};\underline{p},\overline{p}) d\overline{p}(y) \right) \\
&< \frac{\partial}{\partial p} f(\underline{p};\underline{p},\overline{p}) \\
&= \frac{\partial}{\partial p} v_R(0,\underline{p})
\end{aligned}
$$

where the second equality holds since

$$
\frac{\partial^2}{\partial p^2} f(\underline{p};\underline{p},\overline{p}) = -\frac{\partial^2}{\partial p \partial \underline{p}} f(\underline{p};\underline{p},\overline{p}) = \frac{2 r_R f(p)}{\phi^2 \underline{p}^2 (1-\underline{p})^2}.
$$

Thus, the optimal exercise threshold $\underline{p}(y)$ increases faster than the current belief $p(y)$ and an increase in y keeps the beliefs in the linear part of $V_R^I = V_R^{NI}$.

□

Explicit Construction of the Reflected Process First, define $\hat{y}_t$ to be the naive posterior about ξ that is updated based only on public news, i.e.

$$
\hat{y}_t = \mathrm{P}(\xi = 1 \mid X_t) = g(x_t; x_0, y_0).
$$

For any belief process y induced by the Sender define $a(x_t, y_t)$ to be

$$a(x_t, y_t) = g(x_0; x_t, y_t). \tag{A.3}$$

Intuitively, $a(x_t, y_t)$ is the belief about ξ that is updated only on the information generated by the sender above and beyond public news. Clearly the state of the game can be described by a pair (x, a) instead of (x, y) and it is natural to think of a as a pure control variable of the Sender since it is not confounded by the public news. Given a path x_t define a_t^* to be

$$a_t^* = \min\Big[y_0, \inf_{s \le t} g(x_0; x_s, l(x_s)) \Big]. \tag{A.4}$$

Given the process a^* we now define the belief y simply as

$$y_t^* = g(x_t; x_0, a_t^*). \tag{A.5}$$

Lemma 10. *Process a_t^* is the unique process such that $(x_t, y_t^*) \notin \mathbf{T}$ and it changes if and only if $y_t^* = l(x_t)$.*

Proof. **Step 1: reflection.** Suppose there exists some realization $\omega \in \Omega$ and time t such that $y_t^*(\omega) > l(x_t(\omega))$. This implies that $a(x_t(\omega), y_t^*(\omega)) > a_t^*(\omega)$ and gives rise to a contradiction

$$y_t^*(\omega) \equiv g(x_t(\omega); x_0, a(x_t(\omega), y_t^*(\omega))) > g(x_t(\omega); x_0, a_t^*(\omega)) = y_t^*(\omega).$$

Moreover, notice that y_t^* satisfies

$$y_t^* = \int_0^t \Big[(z^1(x_s) - z^0(x_s)) y_s^* (1 - y_s^*) \phi dB_s + \frac{\partial}{\partial y} g(x_s; x_0, a_s^*) da_s \Big].$$

Step 2: uniqueness. Suppose there is an alternative process $\hat{y}_t$ such that in the region $(x, y) \notin \mathbf{T}$ it behaves like the standard belief process, while $(x_t, \hat{y}_t) \notin \mathbf{T}$. The new process $(x_t, \hat{y}_t)$ corresponds to the process $\hat{a}_t$ given by

$$\hat{a}_t = \min\Big[y_0, \inf_{s \le t} g(x_0; x_s, \hat{y}_s) \Big]$$

Then for the same paths we must have $a_t^* \ge \hat{a}_t$. This implies that $y_t^* \ge \hat{y}_t$. It is possible to show that the Sender's payoff is greater if path by path process a_t^* is greater. $\square$

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

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

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

Denote by $\hat{V}(x,y)$ the value obtained by the receiver in the alternative equilibrium. Then $\hat{V}(x,y) \leq V(x,y)$. Otherwise we can add the external news process from the original equilibrium and implement the value function $V(x,y)$ but the receiver will have strictly more information. On the other hand, in the original equilibrium belief reflection does not add value. This implies that $\hat{V}(x,y) \geq V(x,y)$. This keeps the possibility that $V(x,y) = \hat{V}(x,y)$.

Note that for $y < \frac{l(x)+\hat{l}(x)}{2}$ there is no persuasion in the alternative equilibrium. Convexity of the ODE solution in p in the waiting for these y and the result that $V(x,y) = \hat{V}(x,y)$ once $y < l(x)$ and $y > \hat{l}(x)$ it implies that $\hat{V}(x,y) < V(x,y)$ in that region. This contradicts $V(x,y) = \hat{V}(x,y)$ for $y < \hat{l}(x)$. □

Proof of Lemma 6 . Holding y fixed, when z^1 is increased two things happen. First, the Sender discloses information sooner. Second, Receiver's belief about the fundamental $p = yz^1 + (1-y)z^0$ also increases. At the point (x^*, y^*) the Receiver is indifferent between acting and waiting for the discrete piece of news

$$V_R^{NI}(x^*, y^*) = V_R^I(x^*, y^*)$$

Moreover these functions just depend on p

$$V_R^{NI}(x^*, y^*) = f_R^{NI}(y^* z^1(x^*) + (1-y^*)z^0(x^*))$$
$$V_R^I(x^*, y^*) = f_R^I(y^* z^1(x^*) + (1-y^*)z^0(x^*))$$

Note that $\frac{df_R^I}{dp}(p_R) \geq \frac{df_R^I}{dp}(p_R)$ when at $z^1(x_\tau) = p_S$ the Sender reveals everything. Sensitivity of $V_R^{NI}(x,y)$ with respect to z^1 is equal to y at the smooth pasting point. Now consider the representation of V_R^I.

$$V_R^I(x,y) = f_R(p_S(z^1), yz^1 + (1-y)z^0)$$

The Receiver strictly prefers observing the signal z^1 earlier, rather than later. Hence

$$\frac{\partial f_R}{\partial p_S} < 0$$

Then

$$\frac{\partial V_R^I(x^*, y^*)}{\partial z^1} = \frac{\partial f_R}{\partial p_S}\frac{\partial p_S}{\partial z^1} + \frac{\partial f_R}{\partial p} \cdot y^* > y^* = \frac{\partial V_R^{NI}}{\partial z^1}(x^*, y^*)$$

Hence when z^1 is increased, V_R^I increases by more than V_R^{NI}, and which implies that $l(x^*)$ increases.

Next, suppose z^0 decreases. This implies that conditional on $\xi = 0$ the Sender has more precise information. The value function of the Receiver upon obtaining this information is

$$V_R^I(x,y) = \mathrm{E}\left[e^{-r\tau_S}(y_{\tau_S}((W^0 - W^1)p_S - W^0) + (1 - y_{\tau_S})(V_R^{NI}(z^0_{\tau_S})))\right]$$

The derivative wrt z^0 is given by

$$V_R^I(x,y) = (1 - y_{\tau_S}) \cdot \frac{\partial f_R^{NI}}{\partial p} \cdot (z^0_{\tau_S})\mathrm{E}\left[e^{-r\tau_S}\right] < 1 - y$$

Hence when z^0 is decreased the value of V_R^I decreases at a smaller rate than V_R^{NI}. This implies that when z^0 declines, $l(x^*)$ has to increase. □

Proof of Proposition 3. Consider the following pair of Markov strategies:

$\mathcal{T}^*$: Receiver acts if and only if

$$v_R(0, z^1(x)y + z^0(x)(1-y)) \geq y \cdot V_R^{NI}(1,x) + (1-y) \cdot V_R^{NI}(0,x)$$

$\mathcal{H}^*$: Sender reveals ξ after every history.

Now we need to check that these strategies are a best response to one another after every history. Receiver's best response is easy: if he believes a positive amount of information is about to be revealed the next instance he chooses to wait. The only case when he does not is when this information is irrelevant for decision making $z^0 \geq p_R$. In that region, however, the Sender is indifferent between disclosing information and not since it does not affect option exercise.

To compute the Sender's best response we must show that if all information must be revealed for the action to take place, she chooses to reveal it immediately (when $z^0 \geq p_R$ the information is irrelevant for Sender's decision as well and she might as well reveal everything). Denote by τ the moment when ξ is fully revealed. Then

$$p^i_\tau = z^i_\tau \qquad i = 0, 1$$

The Sender solves

$$\max_\tau \mathrm{E}\left[e^{-r\tau} \cdot y_\tau \cdot V_S^{NI}\left(z^1(x_\tau)\right) + e^{-r\tau} \cdot (1 - y_\tau) \cdot V_S^{NI}\left(z^0(x_\tau)\right)\right]$$

In order for immediate disclosure to be the best response, the solution to the above stopping problem must be $\tau^* = 0$. We can simplify by aligning incentives of the Sender conditional on ξ. Type 1 Sender prefers immediate disclosure after any history. Type 0 Sender also prefers immediate disclosure since the Receiver will not act at a threshold lower than p_R after any history. □

Proof of Lemma 8. Fix the strategy of the Sender. It specified a public beliefs process p_t. It is a combination of public news and Sender's messages. The belief p_t can be decomposed into (x_t, y_t). In the region where the Sender chooses to send an informative message, she sends a signal leading beliefs in y_t to be either $y_H(x)$ and $y_L(x)$.

$$V_R(x,y) = \frac{y - y_L(x)}{y_H(x) - y_L(x)} \cdot V_R(x, y_H(x)) + \frac{y_H(x) - y}{y_H(x) - y_L(x)} \cdot V_R\left(x, y_L(x)\right)$$

In particular

$$V_R(x, y_H(x)) = v_R\left(0, y_H(x) z^1(x) + (1 - y_H(x)) z^0(x)\right)$$

The optimal stopping problem of the receiver implies that $\forall (x,y)$

$$V_R(x,y) \geq v_R(0, z^1(x) y + z^0(x)(1-y))$$

Function $z^1(x)y + z^0(x)(1-y)$ is increasing in x. The above inequality implies

$$\begin{aligned}
\lim_{\varepsilon \to 0} \frac{V_R(x, y_H(x)) - V_R(x - \varepsilon, y_H(x))}{\varepsilon} &\leq \frac{dz^1(x)}{dx} \cdot y_H(x) + \frac{dz^0(x)}{dx} \cdot (1 - y_H(x)) \\
\lim_{\varepsilon \to 0} \frac{V_R(x + \varepsilon, y_H(x)) - V_R(x, y_H(x))}{\varepsilon} &\geq \frac{dz^1(x)}{dx} \cdot y_H(x) + \frac{dz^0(x)}{dx} \cdot (1 - y_H(x))
\end{aligned}$$

If the above is not true, then there exists x^* and $y^* = y_H(x^*)$ such that

$$V_R(x^*, y^*) < v_R(0, z^1(x^*) y^* + z^0(x^*) \cdot (1 - y^*))$$

which contradicts the optimal stopping of the Receiver. The derivative with respect to y is given by

$$\begin{aligned}
\frac{\partial V_R}{\partial y}(x,y) &= \frac{1}{y_H(x) - y_L(x)} \cdot \left(v_R\left(0, z^1(x) y_H(x) + z^0(x)(1 - y_H(x))\right) - V_R(x, y_L(x))\right) \\
&= -\frac{V_R(x, y_L(x)) - v_R\left(0, z^1(x) y_L(x) - z^0(x)(1 - y_L(x))\right)}{y_H(x) - y_L(x)} + (z^1(x) - z^0(x)) \cdot (W^0 - W^1) \\
&< (z_1(x) - z_0(x)) \cdot (W^1 - W^0)
\end{aligned}$$

The calculations above show that $V_R(x,y)$ has a non-negative kink in x and strictly positive kink in y at every point of $(x, y_H(x))$. Consider one such point $(\hat{x}, \hat{y}) = (\hat{x}, y_H(\hat{x}))$ with a corresponding belief

$\hat{p} = \mathrm{P}(\theta = 1|\hat{x}, \hat{y})$ and suppose that $\hat{y} < 1$. Then the function $f_S(p)$ defined as

$$f_R(p) = V_R(x(p), y(p)) \quad \text{with} \quad x'(p), y'(p) > 0$$

has a positive kink at $\hat{p}$. $\square$

Proof of Proposition 4 . Recall that for all $p(x, y) < p_R$ it is dominant strategy for the receiver to wait, thus, without loss of generality, we can assume that for all x such that $z^1(x) < p_R$ sender does not reveal any information.

Define

$$\tau = \inf\Big\{t : y_t \geq y_H(x_t) \text{ or } y_t \notin E(x_t)\Big\}$$

This is the first moment where there will be either action of information sharing. We know that $p_t = p(x_t, y(x_t))$ evolves according to the public news process for $t \leq \tau$ and no action is undertaken by the Receiver.

Suppose $y_\tau \geq y_H(x_\tau)$. Since the sender is not disclosing any information up to (and including) the time τ it has to be that $p(x_\tau, y_\tau) = p_R$. Thus, $V_S(x_0, y_0) = V_S^{NI}(p(x_0, y_0))$. If the sender fully reveals information at time 0 (or, pay-off equivalently at time $\tau' = \inf\{t > 0 : z^1(x_t) \geq p_R\}$) her pay-off is

$$y_0 V_S^{NI}(z^1(x_0)) + (1 - y_0) V_S^{NI}(z^0(x_0)) > V_S^{NI}(p(x_0, y_0)).$$

The inequality holds due to strict convexity of $V_S^{NI}(p)$ for all $p \leq p_R$.

Profitable deviation outlined in the previous paragraph implies that $y_\tau < y_H(x_\tau)$ for all (y_0, x_0) such that $z^1(x_0) < p_R$. Consider

$$\underline{y} = \inf_{y_0 \in [0,1]} \{y_\tau = y_\tau(x_0, y_0)\},$$

and suppose that $\underline{y} > 0$. Then there exists y_0 small enough that path of (x_t, y_t) updated only on public news satisfies $y(x_t) \in E(x_t)$ for all $t \geq 0$, i.e., the sender does not reveal any information along this path. This can not be the case, since the receiver would act at $p(x_t, y(x_t)) = p_R$ which contradicts $y_\tau < y_H(x_\tau)$. Thus, $\underline{y} = 0$, implying that for all (x_0, y_0) with $z^1(x_0) < p_R$ in any MPE the action is taken only after ξ is fully revealed and

$$V_S(x_0, y_0) \leq y_0 V_S^{NI}(z^1(x_0)) + (1 - y_0) V_S^{NI}(z^0(x_0)).$$

Clearly, revealing information at $t = 0$ (or at τ') generates

$$V_S(x_0, y_0) = y_0 V_S^{NI}(z^1(x_0)) + (1 - y_0) V_S^{NI}(z^0(x_0))$$

and achieves the upper bound on the equilibrium value function of the Sender. $\square$

Proof of Proposition 5 . After any history Sender's value function cannot exceed:

$$\frac{p}{p_R} \cdot ((F_0 - F_1)p_R - F_0)$$

Suppose $z^1(x) = p_R$. Sender's beliefs are given by $p = y \cdot z^1(x) = p_R$. Then full information revelation implies

$$y \cdot ((F_0 - F_1)p_R - F_0) = \frac{p}{p_R} \cdot ((F_0 - F_1)p_R - F_0)$$

For a given equilibrium payoff $V_S(x, y)$ along the particular path of beliefs define $\overline{W}(p)$ to be the upper bound. It is constructed as follows. Given the path of (x, y) between the waiting and the persuasion region, denote by $\underline{p}$ the first time when the Sender sends the message $\{m = act\} = \{\xi = 1\}$. At $\overline{p}$ beliefs reach Receiver's decision threshold and the latter acts. Then

$$V_S(x, y) \leq \frac{\underline{p}}{p_R} \cdot \frac{\overline{p} - p}{\overline{p} - \underline{p}} \cdot ((F_0 - F_1)p_R - F_0) + \frac{p - \underline{p}}{\overline{p} - \underline{p}} \cdot ((F_0 - F_1)\overline{p} - F_0) = W(p)$$

where, to the benefit of the Sender, Receiver is willing to act at p_R when public beliefs reach $\underline{p}$ along this trajectory.

For a given p define by $u(p)$ the lowest current belief at which the Receiver is willing to act. This is coming from the threshold strategy of the Receiver in the (x, y) space – depending on the current level of p the persuasion threshold will be different. Compare $W(p)$ with immediate persuasion. The value from the latter is

$$w(p) = \frac{p}{u(p)} \cdot ((F_0 - F_1)u(p) - F_0) = (F_0 - F_1)p - F_0 \frac{p}{u(p)}$$

At $p = \overline{p}$ by definition $u(p) = \overline{p}$ and hence $w(\overline{p}) = W(\overline{p})$. I show that $W'(p) > w(p)$ implying that for $p < \overline{p}$ it is better to persuade immediately, rather than wait.

$$\left[-\frac{\underline{p}}{p_R} \cdot ((F_0 - F_1)p_R - F_0) \cdot \frac{1}{\overline{p} - \underline{p}} + \frac{1}{\overline{p} - \underline{p}} \cdot ((F_0 - F_1)\overline{p} - F_0)\right] vs. \left[F_0 - F_1 - F_0 \cdot \frac{1}{u(p)} + F_0 \frac{p}{u(p)^2} u'(p)\right]$$

$$F_0 \frac{\underline{p}}{p_R} \cdot \frac{1}{\overline{p} - \underline{p}} - F_0 \cdot \frac{1}{\overline{p} - \underline{p}} \, vs. \, - F_0 \cdot \frac{1}{\overline{p}}$$

$$\frac{\underline{p}}{p_R} - 1 \, vs. \, - \frac{\overline{p} - \underline{p}}{\overline{p}}$$

$$\frac{\underline{p}}{p_R} \, vs. \, \frac{\underline{p}}{\overline{p}}$$

$$\overline{p} > p_R$$

For the inequalities above to hold it is sufficient that $u'(p) < 0$. Hence locally around $\overline{p}$ immediate persuasion dominates the upper bound on the value that the Sender can obtain if she delays persuasion. Q.E.D. □

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.

Ricardo Alonso and Odilon Câmara. Persuading skeptics and reaffirming believers. *mimeo*.

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

Isabelle Brocas and Juan D Carrillo. Influence through ignorance. *The RAND Journal of Economics*, 38(4):931–947, 2007.

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. *mimeo*, 2015.

Péter Eso and Yuk-Fai Fong. Wait and see: A theory of communication over time. *mimeo*, 2010.

Péter Eso, Balázs Szentes, et al. The one who controls the information appropriates its rents. *mimeo*.

Matthew Gentzkow and Emir Kamenica. Competition in persuasion. *mimeo*, 2016.

Steven R Grenadier, Andrey Malenko, and Nadya Malenko. Timing decisions in organizations: Communication and authority in a dynamic environment. *mimeo*, 2015.

Faruk Gul and Wolfgang Pesendorfer. The war of information. *The Review of Economic Studies*, 79 (2):707–734, 2012.

Yingni Guo. Dynamic delegation of experimentation. *mimeo*, 2014.

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.

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

Aaron M. Kolb. Reputation and strategic adoption. *mimeo*, 2015.

Ilan Kremer and Andrzej Skrzypacz. Dynamic signaling and market breakdown. *Journal of Economic Theory*, 133(1):58–82, 2007.

Erik Madsen. Optimal project termination with an informed agent. *mimeo*, 2015.

George J Mailath and Larry Samuelson. *Repeated games and reputations*, volume 2. Oxford university press Oxford, 2006.

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

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

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.

Alex Smolin. Optimal feedback design. *mimeo*, 2015.