

Gradual Learning from Incremental Actions

Tuomas Laiho, Pauli Murto, Julia Salmi*

June 30, 2022

Abstract

We introduce a collective experimentation problem where agents choose the timing of irreversible actions under uncertainty and where public feedback from the actions arrives gradually over time. This kind of gradual learning where information is delayed and arrives over time is present in many real-life situations, such as adoption of new technologies, progressive market entry, and incremental roll-out of public policies. The socially optimal expansion path entails an informational trade-off where acting today speeds up learning but postponing capitalizes on the option value of waiting. We solve the social optimum and contrast it to the decentralized equilibrium where agents ignore the social value of information they generate. We develop mechanism design techniques to coordinate collective experimentation and illustrate the techniques in the context of pricing a new durable good.

JEL classification: C61, C73, D82, D83

Keywords: social learning, experimentation, optimal stopping, dynamic markets, innovation adoption

1 Introduction

Many actions have long-run consequences that can be observed only gradually over time. As a concrete example, consider a decision to purchase the latest model of

*Laiho: Ministry of Finance, Finland (tuomas.s.laiho@gmail.com), Murto: Aalto University School of Business (pauli.murto@aalto.fi), Salmi: University of Copenhagen (julia.salmi@econ.ku.dk).

an electric car. When using the car after purchase, a new owner observes how well the car functions in different situations and whether any technical problems emerge. When more consumers buy the new model – the fleet of existing cars expands – learning on the aggregate level becomes more precise but remains gradual because each owner continues to make observations long after the buying decision. Similarly, the profitability of capital investments, environmental damage caused by pollution, or effects of public policies are realized gradually over time, potentially long after the critical decisions were made. While such situations abound, we lack a flexible modeling framework to address the welfare implications of this kind of endogenous gradual learning.

This paper introduces a novel learning problem to address this. The key feature of the model is that an action taken today has a long-run impact on the flow of information. A continuum of small agents chooses when to stop – for example, when to adopt an innovation, make an irreversible investment, or purchase a durable good. An unknown binary state determines if stopping is profitable for the agents. Crucially, learning is *gradual*: upon stopping, each agent initiates a persistent flow of information over time. This is in contrast to the standard experimentation models where an action generates an *instantaneous* one-time signal and further actions are needed to learn more.

Our main question is how the incremental path of stopping decisions – the expansion path – is determined on one hand when agents optimize individually and on the other hand when a central planner chooses the actions. We also ask how a planner can coordinate the actions of individual agents with a posted price mechanism. The paper’s contribution is twofold. First, we develop a novel methodological approach with suitable solution techniques. Second, we derive new economic insights that are caused by endogenous gradual learning.

Gradual learning creates a new trade-off for the socially optimal expansions: *the information generation effect* calls for aggressive expansion in order to improve information for future decisions and *the option value effect* calls for cautious expansion in order to have better information for the current decisions. A social planner balances these two effects, but individual agents internalize only the latter

effect and thus the decentralized equilibrium suffers from informational free-riding.

We approach experimentation under gradual learning by modeling the path of individual actions as a stock process, which controls the speed of learning. Each agent who has stopped produces a flow of i.i.d. signals conditional on the true state. In continuous time, this leads to an aggregate signal that follows a Brownian motion with an unknown drift, determined by the true state, and a signal-to-noise ratio proportional to the stock of agents who have stopped. Each stopping decision thus affects information generation gradually over time.

The techniques to solve the decentralized equilibrium and the socially optimal policy turn out to be quite different. The common challenge is that the problems are two-dimensional as both the belief that the state is high and the stock affect the future. Furthermore, the stock and the belief processes are interlinked as the stock determines the flow of new information. We show that the decentralized equilibrium can be solved by analyzing “shortsighted” agents who optimize their stopping decisions against the assumption that no agent stops in the future. The shortsighted approach works because information arrives smoothly over time under gradual learning.

Unlike the decentralized equilibrium, the socially optimal policy takes into account the social value of faster learning. The equivalence with shortsighted optimization breaks because the value of information depends on the expected future actions. Because of the information generation effect, socially optimal policy favors earlier and more aggressive expansions than what happens in the decentralized equilibrium. The difference between the two is especially pronounced when the learning technology is good and learning could potentially be fast. Compared to the no-learning benchmark, gradual learning tends to increase the socially optimal stock for low beliefs and to decrease it for high beliefs due to the informational trade-off between information generation and the option value of waiting.

Due to its technical tractability, the model with gradual learning has a potential to work as a workhorse model that can be further extended to analyze other phenomena. We provide one such extension in the companion paper, Laiho, Murto and Salmi (2022), by including the possibility that agents’ payoffs depend directly

on the other agents' stopping behavior.

In the last part of the paper, we show how to incorporate mechanism design techniques to implement the social optimum, or any other policy, as a decentralized equilibrium. We then demonstrate our theoretical results by solving the pricing problem of a monopolist selling a new durable good under social learning. We point out the potential benefit of market power to consumers: a large player internalizes the informational externality partly to the benefit of the consumers and this may result in a higher consumer welfare than in a competitive market equilibrium. The example highlights the quality of the learning process as an important determinant of welfare implications. An improvement in the learning technology – for example improved social learning caused by better communication technologies – amplifies the welfare loss of the decentralized equilibrium and hence puts more weight on the identified positive impact of market power. More generally, a better learning technology strengthens the case for policies to correct informational distortions highlighted by our model.

1.1 Related literature

Using the framework of our paper, the previous literature on learning can be organized based on whether the information generation effect or the option value effect is present in the model. The current paper is the first to analyze the dual effect of endogenous learning.

The information generation effect is present in papers analyzing classic single agent bandit problems and experimental consumption (Gittins and Jones 1974, Rothschild 1974, Prescott 1972 and Grossman, Kihlstrom and Mirman 1977). Introducing multiple agents to these models adds an informational externality that dampens the information generation effect. Bolton and Harris (1999), Keller, Rady and Cripps (2005) and Keller and Rady (2010) analyze such models under different assumptions on the learning technology. Applications include Bergemann and Välimäki (1997, 2000) and Bonatti (2011) who analyze dynamic pricing. No option value effect exists in these papers because actions are reversible and hence

learning always increases the level of optimal quantities relative to the no-learning benchmark.

When actions are irreversible but information arrives exogenously rather than endogenously, only the option value effect is present. Seminal papers in this literature include McDonald and Siegel (1986), Pindyck (1988), and Dixit (1989) and the ensuing literature on real options is summarized in Dixit and Pindyck (1994). One can see our solution to the social planner's problem as extending the real options literature to endogenous learning.

A few papers investigate social learning with irreversible actions, which bears similarities with informational free-riding in our decentralized solution. Frick and Ishii (2020) analyze the adoption of new technologies using a Poisson process with instantaneous feedback to model learning. The adoption rate of innovations is lower than without learning in their model because of the option value effect. An early paper by Rob (1991) makes a similar observation when analyzing sequential entry into a market of unknown size. Similarly, in the models of optimal timing under observational learning, the option value creates an incentive to wait causing socially inefficient delays (Chamley and Gale 1994, Murto and Välimäki 2011).

Introducing a large player can overturn the effect of social learning and irreversibility on optimal quantities because a large player internalizes the information generation effect. Che and Hörner (2017) study how a social planner, who designs a recommendation system for consumers, can mitigate informational free-riding. Laiho and Salmi (2021) analyze monopoly pricing in a similar setup. Both in Che and Hörner (2017) and in Laiho and Salmi (2021), the presence of a social planner or a monopolist induces learning to increase quantities. The crucial difference from the present paper is that these papers model instantaneous learning from each consumption decision: the planner and the monopolist do not face the option value effect since they get more information only by attracting new consumers. More generally, there is no informational trade-off under instantaneous learning.

Our assumption that learning is gradual implies that past actions matter for the current information flow. Two contemporaneous papers share this feature

with us, although their models and key trade-offs are otherwise different from ours. Liski and Salanié (2020) analyze a single-agent problem where a decision-maker controls the accumulation of a stock that triggers a one-time catastrophe at an unknown threshold level. The novel feature in their model is a random delay between the crossing of the threshold and the onset of the catastrophe. Martimort and Guillouet (2020) analyze a model with similar features focusing on a time-inconsistency problem under their assumptions.

Finally, our other paper Laiho, Murto and Salmi (2022) builds on the modeling techniques developed in this paper to analyze the joint effect of informational and payoff externalities. Laiho, Murto and Salmi (2022) shows that agents can be strictly better off and learning can be faster in an equilibrium under a worse learning technology if there are positive payoff externalities.

2 Model

2.1 Actions and payoffs

A unit mass of small agents choose when, if ever, to take an irreversible action (to stop). We index individual agents by their type θ and assume that θ is distributed according to a continuously differentiable distribution function F with a full support on $\Theta := [\underline{\theta}, \bar{\theta}]$. Time t is continuous and goes to infinity.

An agent's stopping payoff, $v_\omega(\theta)$, depends on the state of the world $\omega \in \{H, L\}$ such that the payoff is higher in the high state of the world for all types: $v_H(\theta) \geq 0 > v_L(\theta)$.¹ Payoffs are continuously differentiable with bounded derivatives and increasing in type: for each $\theta \in \Theta$, $v'_\omega(\theta) \geq 0$ for both $\omega \in \{H, L\}$ and $v'_\omega(\theta) > 0$ for at least one $\omega = H$ or $\omega = L$. The realized payoff for an agent of type θ , who stops at time t , is $e^{-rt}v_\omega(\theta)$ where r is the common discount rate. An agent's outside option is zero and we normalize $v_H(\underline{\theta}) = 0$ so that $\underline{\theta}$ is the lowest type who would ever want to stop. The model is equivalent to a setting where agents receive

¹The analysis easily extends to the case where $v_L(\theta) > 0$ for some types. The only change is that all types, who get a positive stopping payoff in both states of the world, stop immediately.

a flow of state dependent payoffs $\pi_\omega(\theta) = rv_\omega(\theta)$ at every instant after stopping.

Agents are risk-neutral and maximize their expected discounted stopping payoffs. The agents do not know the state of the world ω but learn about it over time as we will describe next.

2.2 Learning

The key idea of *gradual* learning is that every agent who has stopped generates a flow of conditionally independent public signals. Therefore, we consider endogenous learning from the *stock* of stopped agents: let q_t denote the stock (measure) of agents who have stopped by time t .

Specifically, the public learns about the state by observing a Brownian diffusion:²

$$dy_t = q_t\mu_\omega dt + \sigma\sqrt{q_t}dw_t, \quad (1)$$

where we normalize $\mu_H = 1/2$ and $\mu_L = -1/2$, $\sigma > 0$ is the standard deviation of the process, and w_t is a standard Wiener process. Signal process (1) is the limit of a model where q_t is composed of discrete units that produce conditionally independent noisy signals over time and where the total informativeness per unit of q is normalized to stay constant. The signals can be for example interpreted as realized individual payoffs (see Appendix A).³

We denote by x_t the public posterior belief $x_t = Pr(\omega = H|\mathcal{F}_t)$, where \mathcal{F}_t is the natural filtration generated by the signal process (1). The unconditional law of motion for the public belief follows from Bayes' rule:

$$dx_t = \frac{\sqrt{q_t}}{\sigma}x_t(1-x_t)d\tilde{w}_t, \quad (2)$$

where \tilde{w}_t is a standard Wiener process. In equation (2), the term $\frac{\sqrt{q_t}}{\sigma}$ is the signal-to-noise ratio of the process (1) and determines how fast the belief converges to

²The process is otherwise equivalent to the learning processes in Bolton and Harris (1999) and in Moscarini and Smith (2001) but learning is from the stock of cumulative actions instead of being from the flow of new actions.

³See Bergemann and Välimäki (1997, 2000), Bolton and Harris (1999), Moscarini and Smith (2001), and Bonatti (2011) for other applications and further discussion. The difference to these papers is that they do not consider learning from the stock but from the flow of new actions.

the truth. Hence, the higher the stock of stopped agents, the more informative the public signals. In Online Appendix, we extend the model to allow for a more general relationship between the stock q_t and the signal-to-noise ratio.

2.3 Solution concepts

We use the term *policy* for a description of how the stock q_t evolves over time. A policy $Q = \{q_t\}_{t \geq 0}$ is an increasing stochastic process adapted to \mathcal{F}_t . Notice that the signal process itself depends on the evolution of q_t , so that in effect we are defining policy Q jointly with signal process Y .

Individual agents take the policy Q as given when they choose their stopping strategies. A strategy for an agent of type θ is a stopping time $\tau(\theta)$ adapted to \mathcal{F}_t . The payoff to type θ adopting $\tau(\theta)$ under Q is

$$\mathbb{E} \left[e^{-r\tau(\theta)} v_\omega(\theta) \middle| Q \right], \quad (3)$$

where the vertical line notation means that the expectation is for some fixed process Q .

We say that a stopping profile $\mathcal{T} = \{\tau(\theta)\}_{\theta \in \Theta}$ is *consistent with* Q if

$$Pr \left[\int_{\underline{\theta}}^{\bar{\theta}} \mathbf{1}(\tau(\theta) \leq t) dF(\theta) = q_t \middle| Q \right] = 1$$

for all t . In other words, \mathcal{T} is consistent with Q if the measure of agents that it commands to stop always matches the policy.

It is convenient to define solution concepts directly in terms of a policy rather than in terms of a stopping profile. We consider two solution concepts. In a decentralized equilibrium agents optimize individually taking the policy as given:

Definition 1. *A policy Q^E is a decentralized equilibrium if there exists a profile \mathcal{T}^E such that i) it is consistent with Q^E and ii) $\tau^E(\theta)$ maximizes (3) for each θ when $Q = Q^E$.*

The socially optimal policy maximizes the expected total welfare:

Definition 2. A policy Q^* is socially optimal if there exists a profile \mathcal{T}^* such that
i) it is consistent with Q^ and ii)*

$$\mathbb{E} \left[\int_{\underline{\theta}}^{\bar{\theta}} e^{-r\tau^*(\theta)} v_{\omega}(\theta) dF(\theta) | Q^* \right] \geq \mathbb{E} \left[\int_{\underline{\theta}}^{\bar{\theta}} e^{-r\tau(\theta)} v_{\omega}(\theta) dF(\theta) | Q \right],$$

for any policy Q and profile $\mathcal{T} = \{\tau(\theta)\}_{\theta \in \Theta}$ consistent with Q .

In section 3.5 we recast this as a control problem for the stock process $\{q_t\}$.

3 Analysis

Our objective is to analyze how gradual learning affects stopping decisions. First, we discuss some common properties that hold regardless of whether stopping times are individually or socially optimal and present the no-learning benchmark. Then, we solve both the (unique) decentralized equilibrium and the socially optimal policy. Lastly, we compare the decentralized equilibrium and the socially optimal solution to the no-learning benchmark and provide comparative statics results on the effects of learning.

3.1 Higher types stop first

In principle, one can implement a policy Q by many different stopping profiles. However, because the stopping payoffs are increasing in θ , in equilibrium higher type agents want to stop whenever a lower type agent wants to stop, which leads to monotone stopping profiles:

Lemma 1. If $\mathcal{T} = \{\tau(\theta)\}_{\theta \in \Theta}$ maximizes (3) for each θ for given process Q , then

$$\Pr \left[\tau(\theta) \leq \tau(\theta') | \mathcal{F}_t; Q \right] = 1$$

whenever $\theta > \theta'$.

Also socially optimal stopping order is monotone:

Lemma 2. Any stopping profile $\mathcal{T} = \{\tau(\theta)\}_{\theta \in \Theta}$ consistent with Q satisfies:

$$\mathbb{E} \left[\int_{\underline{\theta}}^{\bar{\theta}} e^{-r\tau(\theta)} v_{\omega}(\theta) dF(\theta) | \mathcal{F}_t; Q \right] \leq \mathbb{E} \left[\int_{\underline{\theta}}^{\bar{\theta}} e^{-r\tau^{mon}(\theta)} v_{\omega}(\theta) dF(\theta) | \mathcal{F}_t; Q \right],$$

where $\tau^{mon}(\theta) := \inf \{t : q_t \geq 1 - F(\theta)\}$.

We prove both Lemma 1 and Lemma 2 in Appendix A. The lemmas mean that it is without loss of generality to restrict attention to monotone stopping profiles for which there is a one-to-one mapping between the stock q_t and the largest type θ_t who has not stopped: $q_t = 1 - F(\theta_t)$. Throughout the paper we use notation $q(\theta) := 1 - F(\theta)$ to denote the stock as a function of the current highest type, which has an inverse (current highest type): $\theta(q) := \{\theta : 1 - F(\theta) = q\}$. With a slight notational abuse, we use $v_\omega(q)$ to denote the stopping payoff of type $\theta(q)$.

3.2 Boundary policies

This subsection discusses the dynamics in our model. It turns out that both solutions can be characterized as *boundary policies*:

Definition 3. *A policy Q is a boundary policy if there exists a continuous function $\tilde{q} : [0, 1] \rightarrow [0, 1]$ such that $q_t = \tilde{q}(\max_{s \in [0, t]} x_s)$ where \tilde{q} is strictly increasing for all x such that $\tilde{q}(x) > 0$.*

A boundary policy is Markovian: agents' stopping decisions depend only on the stock and the belief. Because stopping is irreversible, the stock at time t is determined by the highest belief reached up to t . A boundary policy hence divides the stock-belief state space into two regions: in *the expansion region*, more agents stop until the stock equals $\tilde{q}(x)$ and in *the waiting region*, everyone waits.

A boundary policy is fully characterized by the inverse of \tilde{q} , a *policy function* $\tilde{x} : [0, 1] \rightarrow [0, 1]$, which maps the stock to the cutoff belief. It turns out that it is easier to use policy functions to characterize our solutions than functions \tilde{q} . Figure 1 illustrates a boundary policy and the implied dynamics in the state space. Above the boundary, the stock increases (horizontal movement in the figure) and below it, the stock stays constant and only the belief moves (vertical movement). As soon as the belief hits the boundary from below, the quantity is pushed towards right along the boundary. The expansions in the stock are immediate (depicted by solid arrows in the figure), whereas the belief fluctuates according to the diffusion




Figure 1: Dynamics in the waiting and expansion regions of the state space.

process (2) (dashed arrows). Apart from the possible initial jump, the stock process stays below the boundary and is continuous almost surely.

It is useful to note that since a boundary policy is Markovian in the stock-belief state space, we can express an individual agent's best-response to such a policy as an optimally chosen stopping region in the state space. We utilize this in establishing the existence and uniqueness of a decentralized equilibrium.

3.3 No-learning benchmark

We start our analysis with the benchmark case without learning, which allows us to disentangle how learning affects the decentralized equilibrium and the socially optimal solution.

When there is no learning but the common belief stays constant, the agents' stopping problem is myopic. An agent stops if and only if his type is so high that the expected payoff is positive: $xv_H(\theta) + (1-x)v_L(\theta) \geq 0$. Hence, the no-learning policy is characterized by the following myopic cutoff:

$$x^{myop}(q) = \frac{-v_L(q)}{v_H(q) - v_L(q)},$$

where $v_\omega(q) := v_\omega(\theta(q))$.

Individual optimization and socially optimal policies coincide when there is no learning.

3.4 Decentralized equilibrium

We next characterize the decentralized equilibrium defined in Definition 1. An optimal stopping time for an individual agent trades off the cost of waiting with the option value of waiting. Because the belief process changes endogenously as the stock of stopped agents increases, waiting not only brings more information but also faster learning. Despite this, we show that we can solve *equilibrium* stopping times by first solving a sequence of stopping problems where each agent finds the optimal time to stop when the stock is fixed. That is, we fix $q_t = \hat{q}$ for all t and find the optimal stopping time for type $\theta(\hat{q})$ assuming that q_t is constant and equal to \hat{q} . This problem is a one-dimensional stopping problem and can be solved using standard techniques in the literature (see e.g. Dixit and Pindyck (1994) and the team problem in Bolton and Harris (1999)). We show that the equilibrium in the original problem corresponds to this “shortsighted” problem in which agents do not take future stopping decisions by other agents into account.⁴

The intuition for the equivalence between the shortsighted problem and the original problem is the following: because later expansions in the stock happen only when it is optimal for lower type agents to stop, all higher type agents strictly prefer stopping always when the stock expands (Lemma 1). Hence, higher type agents want to stop even before the expansion, which means that future expansions do not change the optimal stopping times. The equivalence with shortsighted optimization is hence an equilibrium property and may be violated against other (non-equilibrium) stock processes. In Appendix B, we formalize this argument to get the following result:

Proposition 1. *There is a unique decentralized equilibrium, which is characterized by an increasing policy function x^E :*

$$x^E(q) := \frac{-\beta(q)v_L(q)}{(\beta(q) - 1)v_H(q) - \beta(q)v_L(q)},$$

where $\beta(q) := \frac{1}{2} \left(1 + \sqrt{1 + \frac{8r\sigma^2}{q}} \right)$.

⁴Our method to solve the decentralized equilibrium is inspired by an industry investments paper by Leahy (1993) who shows that under exogenous uncertainty the competitive equilibrium behavior coincides with that of ‘myopic’ investors who ignore the effect the future investments have on the price.

According to Proposition 1, an agent of type θ waits until the belief reaches the cutoff $x^E(q(\theta))$. The decentralized equilibrium is thus a boundary policy: the policy function x^E defines a boundary so that whenever the belief is about to cross the boundary, more agents stop.

Notice that the cutoff $x^E(q)$ is increasing in the signal precision (decreasing in σ), which means that a better learning technology decreases the stock of agents who are willing to stop at any given belief. This is because the better the learning technology, the greater the option value of waiting and hence the higher the belief threshold at which an agent stops. The no-learning benchmark is a special case of the decentralized equilibrium as we take $\sigma \rightarrow \infty$, which directly gives:

Corollary 1. *For all beliefs in $(x^{stat}(0), 1)$, the stock of stopped agents is strictly smaller in the decentralized equilibrium than in the no-learning benchmark.*

The intuition behind Corollary 1 is that agents want to fully utilize the information provided by other agents. Proposition 1 implies that in the decentralized equilibrium only *past* stopping decisions affect individual agents' behavior because from an individual agent's perspective past actions determine the speed of learning. In the next section, we analyze the socially optimal policy which takes into account the informational externality between agents. Then both *past and future* stopping decisions affect the solution.

3.5 Social optimum

We now consider the problem in Definition 2 where a benevolent social planner seeks to maximize agents' expected joint payoff. The problem is identical to a problem of a single decision maker who controls a path of incremental expansions. Solving the problem is the main analytical contribution of this paper.

From Lemma 2, we know that the skimming property holds for the social optimum and hence the problem is reduced to finding the policy Q that maximizes the expected social welfare. We denote the planner's payoff in state (x, q) as

$$U(Q; x, q) = \mathbb{E} \left[\int_q^1 e^{-r\tau(s)} (x_{\tau(s)} v_H(s) + (1 - x_{\tau(s)}) v_L(s)) ds \mid x, q; Q \right]. \quad (4)$$

The planner's problem is then to find $\sup_Q U(Q; x, q)$. By applying Itô's lemma and using the properties of the Brownian motion, we have the following Hamilton-Jacobi-Bellman (HJB) equation for the planner's problem:

$$rV(x, q) = \max_{q^* \geq q} \left(r \int_q^{q^*} (xv_H(s) + (1-x)v_L(s)) ds + \frac{1}{2} V_{xx}(x, q^*) \frac{x^2(1-x)^2}{\sigma^2} q^* \right). \quad (5)$$

We solve the planner's problem by showing that the HJB equation is solved by a boundary policy that cuts the state space into an expansion region and a waiting region. A verification argument then shows that our candidate solution also maximizes the original objective (4).

The optimal policy could in principle consist of several waiting and expansion regions. We proceed by guessing that there is only one expansion and only one waiting region and then later verify this guess (in Appendix C). Let $x^* : [0, 1] \rightarrow [0, 1]$ denote our candidate for the socially optimal policy, which we derive next. Function x^* splits the state space in two so that for a given q the planner waits for beliefs $x < x^*(q)$ and expands for beliefs $x \geq x^*(q)$. Since the planner internalizes the value of information for further decisions, we should intuitively expect the socially optimal expansion region to be larger than in the case of decentralized equilibrium, i.e. $x^*(q) < x^E(q)$. We shall verify that this property indeed holds.

We start by finding the value function that solves the HJB equation (5). In the waiting region below x^* , we have $q^* = q$ and hence (5) reduces to a differential equation. The value consists of the value of potential future actions. Solving the differential equation yields:⁵

$$V(x, q) = B(q)\Phi(x, q), \quad (6)$$

where $B(q)$ is a function to be determined and

$$\Phi(x, q) := x^{\beta(q)}(1-x)^{1-\beta(q)} \text{ and } \beta(q) = \frac{1}{2} \left(1 + \sqrt{1 + \frac{8r\sigma^2}{q}} \right) \text{ as in Proposition 1.}$$

The next step is to find functions B and x^* that maximize the right-side of the HJB equation. To characterize these, we apply value matching and smooth

⁵We have discarded the other root of the characteristic equation, $\tilde{\Phi}(x, q) := x^{1-\beta(q)}(1-x)^{\beta(q)}$, as we must have that the value converges to the static solution as $x \rightarrow 0$ and $x \rightarrow 1$.

pasting conditions, which are necessary for the optimality of policy x^* . Because the planner controls the intensity of experimentation, the conditions apply to a marginal increase of the stock q . The value matching condition is thus $V_q(x^*(q), q) = -x^*(q)v_H(q) - (1 - x^*(q))v_L(q)$ and the smooth pasting condition is $V_{qx}(x^*(q), q) = -v_H(q) + v_L(q)$. Notice that the HJB equation consists of only future, not past, stopping payoffs and therefore the value matching condition equates the marginal value of increasing the stock with the lost stopping payoff.

Using Equation (6), we can write the value matching and smooth pasting conditions as

$$x^*(q)v_H(q) + (1 - x^*(q))v_L(q) + B_q(q)\Phi(x^*(q), q) + B(q)\Phi_q(x^*(q), q) = 0, \quad (7)$$

$$v_H(q) - v_L(q) + B_q(q)\Phi_x(x^*(q), q) + B(q)\Phi_{qx}(x^*(q), q) = 0. \quad (8)$$

Our candidate policy x^* must balance the direct payoff effect, the first term in both equations, against both the option value of waiting and the value of information generation. The last two show up in the latter terms of each equation as the derivatives of the value function.

We show in Appendix C that the system (7) - (8) can be transformed into a non-linear differential equation that defines our candidate policy x^* :

$$x^{*'}(q) = g(x^*(q), q), \quad (9)$$

where

$$g(x, q) = x(1 - x) \left[x \left(\beta'(q)(\beta(q) - 1)v_H'(q) - ((\beta(q) - 1)\beta''(q) - 2(\beta'(q))^2)v_H(q) \right) \right. \\ \left. + (1 - x) \left(\beta'(q)\beta(q)v_L'(q) - (\beta(q)\beta''(q) - 2(\beta'(q))^2)v_L(q) \right) \right] / \\ \left[\left(x(\beta(q) - 1)^2v_H(q) + (1 - x)(\beta(q))^2v_L(q) \right) \beta'(q) \right].$$

The appropriate initial condition for the differential equation is $x^*(1) = 1$ because the solution must equal the no-learning benchmark when the belief equals one.

The denominator of function g is zero at $(1, 1)$ and hence a potential singularity problem arises. However, we show in Appendix C that the initial value problem has a unique solution below the decentralized solution i.e. a solution satisfying $x^*(q) \leq$

$x^E(q)$ for all $q \in [0, 1]$ (proof of Lemma 5 in Appendix C.2).⁶ We then verify that together with the value function in (6) this candidate policy x^* solves the HJB equation, and we further verify that it also maximizes the original objective (4). In the process, we show that the policy function x^* is continuous and strictly increasing in q and hence satisfies the requirements for a boundary policy.

Proposition 2. *The socially optimal policy is a boundary policy. It is characterized by the unique solution to the initial value problem $x'(q) = g(x^*(q), q)$ and $x^*(1) = 1$ satisfying $x^*(q) \leq x^E(q)$ for all $q \in [0, 1]$.*

Proposition 2 confirms that we can solve the potentially complicated history-dependent problem with a simple boundary policy. However, unlike the decentralized equilibrium, we cannot solve the planner's problem in closed form because the planner is truly forward-looking. For the socially optimal policy, *both past and future* actions are relevant. The past generates information that is useful in evaluating the right decision today, whereas future decisions can be based on information generated by today's action. The socially optimal policy balances the resulting trade-off between the efficient use of information (*option value effect*) and the efficient production of information (*information generation effect*).




Figure 2: Socially optimal policy $x^*(q)$ for different σ when $v_H(q) = 1 - q$, $v_L(q) = -1/2$, and $r = 0.1$.

Figure 2 provides a numerical example of the effects of the signal precision.

⁶To see that the uniqueness can only hold in a restricted domain, note that $g(1, q) = 0$ and hence the initial value problem has a trivial solution $x(q) = 1$ for all $q \leq 1$.




Figure 3: Different policies when $v_H(q) = 1 - q$, $v_L(q) = -1/2$, $\sigma = 0.5$, and $r = 0.1$.

The smaller the noise parameter σ is, the more precise the signals are. Better learning technology decreases the cutoff belief $x^*(q)$ when the stock is small and increases it when the stock is high. This arises because improved learning amplifies both information generation and option value effects. The former dominates in the beginning, when the existing stock is low and there are many uncommitted agents who benefit from more information. Conversely, the option value effect dominates later when there are few such agents. Notice that the policies with learning (finite σ) are first below and later above the myopic policy without learning ($\sigma = \infty$). Hence, gradual learning may either increase or decrease expansions as the informational trade-off suggests. The following proposition generalizes this observation (see Appendix C.3 for the proof).

Proposition 3. *There exists $\underline{x} \in (x^{myop}(0), 1)$ and $\bar{x} \in [\underline{x}, 1)$ such that the optimal stock is strictly larger than the no-learning benchmark for all beliefs in $(x^*(0), \underline{x})$ and strictly lower for all beliefs in $(\bar{x}, 1)$.*

Figure 3 illustrates the relationship between the solutions. Compared to the no-learning benchmark, gradual learning first increases and then decreases optimal expansions over time. The decentralized policy requires a higher belief for further expansions than the other policies.

Finally, it is illuminating to look at what happens to the actual speed of learning when the learning technology improves. To do that, let $q_\sigma^*(x)$ and $q_\sigma^E(x)$ denote the socially optimal and the decentralized stocks for signal precision σ .

Proposition 4. (a) *The socially optimal signal-to-noise ratio $\sqrt{q_\sigma^*(x)}/\sigma \rightarrow \infty$ as $\sigma \rightarrow 0$ for all $x \in (0, 1)$. (b) *The decentralized signal-to-noise ratio $\sqrt{q_\sigma^E(x)}/\sigma \rightarrow a(x)$ as $\sigma \rightarrow 0$ where $a(x) = 0$ for all $x \leq x^{stat}(0)$ and $a(x) \in (0, \infty)$ for all $x \in (x^{stat}(0), 1)$.**

Learning gets arbitrarily fast in the socially optimal solution when the learning technology improves, whereas learning remains slow in the decentralized equilibrium. The latter is caused by informational free-riding: no-one wants to be the first one to stop if information arrives fast. This result suggests that the signal precision σ is an important determinant of welfare implications of the model as we will highlight in section 4.3. In Appendix C.4, we prove Proposition 4 and derive the functional form for $a(x)$.

4 Mechanism design

We have seen that the social planner's solution and the decentralized equilibrium differ from each other. In this complementary section we bridge this gap by showing how to implement a given policy in a decentralized manner with anonymous posted prices. We complement the analysis by deriving a revenue maximizing mechanism. Lastly, we demonstrate the results of this paper in the context of a market for a new durable good.

4.1 Posted prices

We now show how a designer can implement a boundary policy by using anonymous posted prices.⁷ A *posted price rule* $P : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$ defines a transfer payment that an agent has to pay to the designer if willing to stop at a given state. With such a policy in place, the stopping payoff for an agent of type θ who decides to stop in state (x, q) is

$$u_\theta(x, q) = xv_H(\theta) + (1 - x)v_L(\theta) - P(x, q). \quad (10)$$

⁷Depending on the context and sign, these can also be interpreted as taxes or subsidies.

Let Q be an arbitrary boundary policy with policy function \tilde{x} satisfying $\tilde{x}(q_1) = 1$ and its inverse $\tilde{q}(x) : [0, 1] \rightarrow [0, q_1]$ with the convention $\tilde{q}(x) = 0$ for $x \leq \tilde{x}(0)$. Note that we allow here for the possibility that the designer wants to implement a restricted maximal stock $q_1 < 1$. The following proposition characterizes a posted price rule that implements Q :

Proposition 5. *Fix a boundary policy Q with policy function \tilde{x} . Then there exists a posted price rule P such that Q is a decentralized equilibrium of the game where the stopping payoff is given by (10). For states along the stopping boundary, the posted price rule is uniquely pinned down by*

$$P(\tilde{x}(q), q) = \tilde{x}(q)(v_H(\theta(q)) + (1 - \tilde{x}(q))v_L(\theta(q))) - \mathbb{E} \left[\int_{\underline{\theta}}^{\theta(q)} e^{-r(\tau(s) - \tau(\theta(q)))} (\tilde{x}(q(s))v'_H(s) + (1 - \tilde{x}(q(s)))v'_L(s)) ds \middle| \tilde{x}(q), q \right]. \quad (11)$$

For states away from the boundary, the posted price rule is not uniquely determined. For states $x > \tilde{x}(q)$, we can assign $P(x, q) = P(x, \tilde{q}(x))$. For states $x < \tilde{x}(q)$, we can assign any sufficiently high value to $P(x, q)$.

We prove the proposition in Appendix D. The posted price at the boundary is pinned down by the envelope theorem of Milgrom and Segal (2002) and we verify that incentive compatibility holds globally along the stopping boundary. To guarantee that the agents do not have an incentive to stop at any state below the boundary, the designer only needs to make sure that the transfer payment is high enough to make the cost of stopping prohibitive. The posted price that implements a desired policy is hence unique at the boundary but not away from it.⁸

The posted price rule considered here responds to changes in both x and q . In Online Appendix, we investigate two alternative posted price rules that only depend on the stock or the belief, respectively. Such rules are of special interest because they are often easier to use in practice. For example, a seller of a new durable good can set the price based on the cumulative past sales instead of both sales and reviews or other feedback from past buyers.

⁸Above the stopping boundary, we must guarantee that those agents needed for the state to move to the boundary do indeed want to stop. Setting $P(x, q) = P(x, \tilde{q}(x))$ is one way to accomplish that.

4.2 Revenue maximizing designer

In some applications a designer maximizes an objective other than social surplus. Here we show how we can adapt our techniques to find a mechanism that maximizes the revenues from the transfers. For this subsection we assume that the type distribution F is twice continuously differentiable and has monotone hazard rate.

Consider a designer whose objective is to maximize the expected sum of transfers, $\mathbb{E}[\int_0^\infty e^{-rt} P_t dq_t]$, where P_t is the transfer that the agent pays if he stops at time t . For a given policy, the incentive compatible posted price is pinned down by (11). To back out the incentive compatible revenue, we change the order of integration to get a virtual surplus representation for the designer's payoff (see Appendix D for the proof):⁹

Lemma 3. *Incentive compatibility implies that the designer's expected revenue is:*

$$\mathbb{E}\left[\int_0^\infty e^{-rt} P_t dq_t\right] = \mathbb{E}\left[\int_{\underline{\theta}}^{\bar{\theta}} e^{-r\tau(\theta)} \phi_\omega(\theta) d\theta\right],$$

where $\phi(\theta)$ is the virtual stopping payoff: $\phi_\omega(\theta) := v_\omega(\theta) - v'_\omega(\theta) \frac{1-F(\theta)}{f(\theta)}$.

Using this result we can solve the revenue maximizing designer's problem by using the planner's solution in Proposition 2. We only need to replace the stopping payoffs with virtual stopping payoffs and use the revenue maximizing quantity under complete information as the initial value. We demonstrate this in the next subsection where we analyze the problem of a durable goods monopolist.

4.3 Application: markets for new durable goods

We now pull together our theoretical results to discuss an application. Consider a new durable good with uncertainty about the product quality. Gradual social learning naturally arises: after an individual buyer purchases the product, he starts using it and observes how well it functions over time. As a result, potential

⁹ The approach extends a result in Laiho and Salmi (2021) and shares similar features with Board (2007) who analyzes the optimal sale of options.

future buyers learn gradually from the fleet of existing users. This is in contrast to experience goods, such as movies, where consumption is one-shot and hence instantaneous learning is a natural modeling approach. The existing literature on experimentation has focused solely on experience goods.

The model is as follows. Neither the monopolist nor the buyers know the true quality of the product, $\omega \in \{H, L\}$, but they observe a public signal process (1), generated by past sales, q_t . Each buyer wants to purchase one unit and exits after purchase. Similar to the general model, a buyer's utility from consumption depends on his private type, $\theta \in [\underline{\theta}, \bar{\theta}]$, and the common quality: $\mathbb{E}[u(\theta, \omega)] = \mathbb{E}[\mathbf{1}_{\omega=H} \cdot \theta] = x_t \theta$ where x_t is the current belief that the quality is high. In addition, we assume that the type distribution F is twice continuously differentiable and has monotone hazard rate. The monopolist faces marginal cost of production $c > 0$ and commits to a pricing scheme, P_t .

The monopolist's problem is a mechanism designer's problem whose objective is to maximize the expected sum of transfers, $\mathbb{E}[\int_0^\infty e^{-rt} P_t dq_t]$. Using Lemma 3, the monopolist's objective can be written as

$$\mathbb{E}\left[\int_0^\infty e^{-rt} P_t dq_t\right] = \mathbb{E}\left[\int_{\underline{\theta}}^{\bar{\theta}} e^{-r\tau(\theta)} \phi_\omega(\theta) d\theta\right],$$

where ϕ_ω is the virtual valuation: $\phi_H(\theta) := \theta - \frac{1-F(\theta)}{f(\theta)}$ and $\phi_L(\theta) := 0$.

The monopolist seeks to maximize the expected sum of virtual valuations net of the cost of production. With this in mind, we can use the planner's solution in Proposition 2 to characterize the monopolist's solution:¹⁰

Corollary 2. *The monopolist's policy x^M is characterized by $x^M(\bar{q}^M) = 1$ and $x^M(q) = g(q, x^M(q))$, where \bar{q}^M solves $\theta(\bar{q}^M) - (1 - F(\theta(\bar{q}^M)))/f(\theta(\bar{q}^M)) = c$ and g is given in (9) for $v_H(\theta) = \theta - (1 - F(\theta))/f(\theta) - c$ and $v_L(\theta) = -c$.*

The monopolist sells more whenever the belief reaches the boundary $x^M(q)$ and waits otherwise. The posted prices that implement the solution satisfy (11)

¹⁰Notice that the monotone hazard rate condition is important as it guarantees that the virtual valuation is increasing in θ and hence the problem satisfies all our assumptions in Section 2. In addition, because the solution in Proposition 2 is a boundary policy, Proposition 5 guarantees that posted prices are without loss of generality and hence the revenue maximizing policy can be implemented.

in Proposition 5.

We contrast the monopoly solution with the planner's socially optimal solution and the competitive market equilibrium. The planner's solution can be found also by applying Proposition 2 but with different payoff functions:

Corollary 3. *The social planner's policy x^P is characterized by $x^P(\bar{q}^P) = 1$ and $x^{P'}(q) = g(q, x^P(q))$, where \bar{q}^P solves $\theta(\bar{q}^P) = c$ and g is given in (9) for $v_H(q) = \theta(q) - c$ and $v_L(q) = -c$.*

Suppose next that there are no barriers of entry to the market so that the price equals the marginal cost: $P_t = c$. An individual buyer's purchasing problem then coincides with the decentralized equilibrium with $v_H(q) = \theta(q) - c$ and $v_L(q) = -c$, and we have the following corollary to Proposition 1:

Corollary 4. *The competitive market policy is*

$$x^C(q) = \frac{\beta(q)c}{(\beta(q) - 1)\theta(q) + c}.$$

The planner's solution and the competitive equilibrium are the socially optimal and the decentralized solutions of the same problem, whereas the monopoly solution uses different stopping payoffs. This difference leads to different inefficiencies in monopoly and competitive markets.

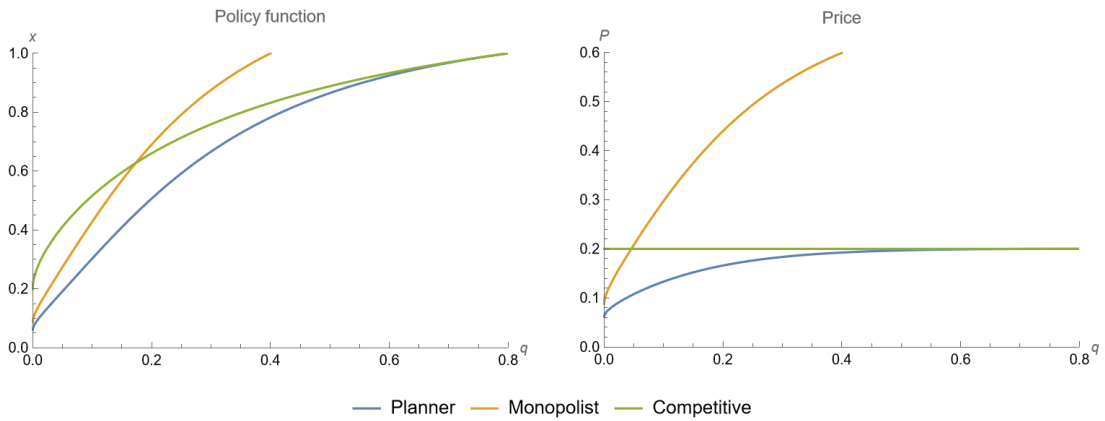


Figure 4: Different solutions for uniform $(0, 1)$ types, $c = 0.2$, $r = 0.1$, and $\sigma = 0.5$.

Figure 4 gives a numerical example of the solutions. The left panel shows the policy functions. The monopolist's and the competitive policies are everywhere

above the planner's policy: both markets require inefficiently high belief for new consumers to purchase the product. The monopoly policy is first below and then above the competitive policy. This is because the information generation effect encourages the monopolist to sell in the beginning and because early sales do not generate large information rents to other buyers. Later on, the monopolist reduces sales as the option value effect gets stronger and because later sales impose information rents to higher type buyers. The competitive market ignores the information generation effect but is otherwise efficient and therefore competitive sales are larger for high beliefs. This comparison in fact holds for all type distributions and parameter values (see Online Appendix).

The right panel shows the corresponding posted prices (11) at the stopping boundary. The monopoly starts at a price below marginal cost in order to initially boost information generation, but then increases the price steeply to extract more revenues. The social planner prices everywhere below the marginal cost in order to fully internalize the informational externality.

The monopolist's incentives to generate information are weaker than the planner's because the monopolist cannot capture all value from the buyers. This creates a distortion at the top of the type distribution, which is not present when the quality is known. In other words, a higher initial belief is needed for the monopolist to be willing to launch the product: $x^M(0) > x^P(0)$.

With endogenous learning, the quality of the learning technology is a key determinant of welfare differences between different market structures. The parameter that captures this in our model is the noise parameter σ . The left panel in Figure 5 shows the difference in total welfare between the planner's solution and the competitive solution as a function of initial belief and with different values of σ .¹¹ The better the learning technology, i.e. the lower the value of σ , the greater the welfare loss in competitive equilibrium. This is because the social planner takes full advantage of improved learning technology, whereas informational free-riding limits how much the actual speed of learning in equilibrium increases when the

¹¹ The kink at initial belief 0.2 is the threshold at which the first unit is sold in competitive equilibrium.

learning technology improves (see Proposition 4).




Figure 5: Difference in total welfare in the social planner vs. competitive solution (left panel) and in monopoly vs. competitive solution (right panel) as a function of the initial belief for uniform $(0, 1)$ types, $c = 0.2$, and $r = 0.1$.

A striking implication of endogenous learning in the current context is that market power can benefit consumers. Although a monopoly harms consumers by distorting the allocation in the usual way, it also spills benefits to the consumers from its ability to internalize the informational externality. The right panel in Figure 5 shows the difference in total welfare between the monopoly solution and the competitive equilibrium. Positive values indicate that the monopoly solution has a greater total welfare. Since consumers get a positive share of the total surplus in information rents and since consumer surplus is zero in the competitive solution for initial beliefs below 0.2, this implies that for a range of initial beliefs the consumers are better off in the monopoly solution than in the competitive solution. Again, we see that the quality of the learning technology plays a key role: the better the learning technology, the greater the scope for the consumers to benefit from market power.

As a final remark, notice that a regulator can implement the socially efficient consumption in both monopoly and competitive markets by using appropriate subsidies but there is a crucial difference between the two different competitive environments. To encourage the monopolist to sell more, the regulator should use *back-loaded* subsidy that increases over sales: $s(x, q) = x(1 - F(\theta(q)))/f(\theta(q))$. A back-loaded subsidy scheme incentivizes the monopolist to sell the socially optimal amount because she internalizes the benefits of information generation. If

the market is competitive, however, the subsidy must be *front-loaded* because competition eliminates dynamic incentives (see the difference between the social planner’s price and marginal cost in the right panel of Figure 4).

5 Concluding remarks

We conclude by summarizing a few compelling reasons why the gradual arrival of endogenous information matters. First, gradual learning enables the analysis of various real-life situations where the long-run consequences of a decision determine its profitability. Because gradual learning creates a novel informational trade-off on the social level between information generation and the option value of waiting, it dramatically shapes the incentives of experimentation.

A second important motive to model the gradual arrival of information is technical. As demonstrated in this paper, the decentralized equilibrium can be solved in a closed form under gradual learning. We further show how mechanism design techniques can be utilized to conduct policy analysis in our environment. The solution method extends to richer environments, such as to models with payoff externalities (see Laiho, Murto and Salmi 2022).

An important takeaway from the paper is that the signal precision has subtle implications for learning and welfare. We show that even if signals get arbitrarily precise, learning remains slow in the equilibrium. This contrasts with the socially optimal solution, in which the true state is learned arbitrarily fast as the learning technology improves. As a result, the equilibrium welfare loss is particularly severe if the learning technology is good.

Acknowledgements

We thank Aislinn Bohren, Martin Cripps, Rahul Deb, Francesc Dilme, Nils Christian Framstad, Dino Gerardi, Yingni Guo, Bård Harstad, Daniel Hauser, Jan Knoepfle, Matti Liski, Erik Madsen, Helene Mass, Konrad Mierendorff, Francesco Nava, Mariann Ollar, Marco Ottaviani, Sven Rady, Peter Norman Sørensen,

Roland Strausz, Juuso Toikka, Christian Traeger, Juuso Välimäki, and various seminar and conference audiences for their helpful comments.

References

- Bergemann, Dirk and Juuso Välimäki**, “Market diffusion with two-sided learning,” *RAND Journal of Economics*, 1997, 28 (4), 773–795.
- and –, “Experimentation in markets,” *Review of Economic Studies*, 2000, 67 (2), 213–234.
- Board, Simon**, “Selling options,” *Journal of Economic Theory*, 2007, 136 (1), 324–340.
- Bolton, Patrick and Christopher Harris**, “Strategic Experimentation,” *Econometrica*, 1999, 67 (2), 349–374.
- Bonatti, Alessandro**, “Menu pricing and learning,” *American Economic Journal: Microeconomics*, 2011, 3 (3), 124–163.
- Chamley, Christophe and Douglas Gale**, “Information Revelation and Strategic Delay in a Model of Investment,” *Econometrica*, 1994, 62 (5), 1065–1085.
- Che, Yeon-Koo and Johannes Hörner**, “Recommender Systems as Mechanisms for Social Learning,” *The Quarterly Journal of Economics*, 2017, pp. 871–925.
- Dixit, Avinash**, “Entry and Exit Decisions under Uncertainty,” *Journal of Political Economy*, 1989, 97 (3), 620–638.
- Dixit, Avinash K. and Robert S. Pindyck**, *Investment under Uncertainty*, Princeton: Princeton University Press, 1994.
- Fleming, Wendell H and Halil Mete Soner**, *Controlled Markov processes and viscosity solutions*, Vol. 25, Springer Science & Business Media, 2006.
- Frick, Mira and Yuhta Ishii**, “Innovation Adoption by Forward-Looking Social Learners,” *Working Paper*, 2020.
- Gittins, J. and D. Jones**, “A Dynamic Allocation Index for the Sequential Allocation of Experiments,” *Progress in Statistics*, ed. by J. Gani, 1974, pp. 241–266.
- Grossman, Sanford J., Richard E. Kihlstrom, and Leonard J. Mirman**, “A Bayesian Approach to the Production of Information and Learning by Doing,” *The Review of Economic Studies*, 1977, 44 (3), 533–547.

- Keller, Godfrey and Sven Rady**, “Strategic experimentation with Poisson bandits,” *Theoretical Economics*, 2010, 5 (2), 275–311.
- , – , and **Martin Cripps**, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 2005, 73 (1), 39–68.
- Laiho, Tuomas and Julia Salmi**, “Coasian Dynamics and Endogenous Learning,” *Working paper*, 2021.
- , **Pauli Murto, and Julia Salmi**, “Timing of Irreversible Actions under Informational and Payoff Externalities,” *Working paper*, 2022.
- Leahy, John V**, “Investment in competitive equilibrium: The optimality of myopic behavior,” *The Quarterly Journal of Economics*, 1993, 108 (4), 1105–1133.
- Liski, Matti and Francois Salanié**, “Catastrophes, delays, and learning,” *Toulouse School of Economics Working Papers No 1148*, 2020.
- Martimort, David and Louise Guillouet**, “Precaution, information and time-inconsistency: on the value of the precautionary principle,” *CEPR Discussion Paper Series DP15266*, 2020.
- McDonald, Robert and Daniel Siegel**, “The value of waiting to invest,” *The quarterly journal of economics*, 1986, 101 (4), 707–727.
- Milgrom, Paul and Ilya Segal**, “Envelope theorems for arbitrary choice sets,” *Econometrica*, 2002, 70 (2), 583–601.
- Moscarini, Giuseppe and Lones Smith**, “The Optimal Level of Experimentation,” *Econometrica*, 2001, 69 (6), 1629–1644.
- Murto, Pauli and Juuso Välimäki**, “Learning and Information Aggregation in an Exit Game,” *The Review of Economic Studies*, 2011.
- Pindyck, Robert S.**, “Irreversible Investment, Capacity Choice, and the Value of the Firm,” *The American Economic Review*, 1988, 78 (5), 969–985.
- Prescott, Edward C.**, “The Multi-Period Control Problem Under Uncertainty,” *Econometrica*, 1972, 40 (6), 1043–1058.
- Rob, Rafael**, “Learning and Capacity Expansion under Demand Uncertainty,” *The Review of Economic Studies*, 1991, 58 (4), 655–675.
- Rothschild, Michael**, “A two-armed bandit theory of market pricing,” *Journal of Economic Theory*, 1974, 9 (2), 185–202.

Appendix

A Additional material for Sections 2 and 3.1

A.1 Learning process as the continuous limit

Consider a discrete model where the number of agents is n and where the period length is dt . Let the signal process be such that in each period, each agent who has stopped generates a normally distributed conditionally iid. signal:

$$y_t^i \sim N\left(\frac{\mu_\omega dt}{n}, \frac{\sigma^2 dt}{n}\right).$$

This normalization keeps the informativeness of the aggregate signal constant while letting the number of small agents to grow as in Bergemann and Välimäki (1997).

When the number of agents who has stopped is $k \leq n$, this implies the following aggregate signal:

$$\sum_{i=1}^k y_t^i \sim N\left(\mu_\omega dt \frac{k}{n}, \sigma^2 dt \frac{k}{n}\right).$$

Let $q = k/n$ denote the fraction of agents who have stopped. Now, the signal process (1) follows once we take the limit when $n \rightarrow \infty$ (and $k \rightarrow \infty$ so that k/n stays fixed) and $dt \rightarrow 0$.

Notice that the limiting distribution for the aggregate signal depends only on the mean and the variance of y_t^i (the central limit theorem). Hence, the signal process (1) is also the limiting process for the case where y_t^i is not normally distributed, including the case where agents communicate through binary signals.

Furthermore, we can rewrite the model so that the individual signals represent realized payoffs in a model where agents start receiving a stochastic flow payoff after stopping: $\pi_t(\theta) = \pi_\omega(\theta) + \epsilon_t(\theta)$ where $\epsilon_t(\theta) \sim N(0, \sigma^2(\pi_H(\theta) - \pi_L(\theta))^2)$. The noise term is scaled so that every increment in q is equally informative. This assumption is not necessary: we analyze in Online Appendix the case of heterogeneous informativeness and show that both the analysis and the qualitative results remain unchanged if the stopping profile is monotone. When we set

$\pi_\omega(\theta) = rv_\omega(\theta)$, the expected stopping payoff is $x_t v_H(\theta) + (1 - x_t)v_L(\theta)$ just like in the main text. Since there are no further actions after stopping, it does not matter how fast the agents learn privately after they have stopped: the parameter σ can be interpreted to capture both the noise in the private learning and the noise in communication.

A.2 Proof of Lemma 1

Proof. Let policy Q be fixed. Type θ wants to stop at time t if

$$x_t v_H(\theta) + (1 - x_t)v_L(\theta) \geq \mathbb{E}[e^{-r(\tau-t)}(x_\tau v_H(\theta) + (1 - x_\tau)v_L(\theta)) | \mathcal{F}_t; Q],$$

for all stopping rules τ . Or equivalently,

$$v_L(\theta)(1 - x_t - \mathbb{E}[e^{-r(\tau-t)}(1 - x_\tau) | \mathcal{F}_t; Q]) + v_H(\theta)(x_t - \mathbb{E}[e^{-r(\tau-t)}x_\tau | \mathcal{F}_t; Q]) \geq 0.$$

The left-hand side is increasing in θ because expressions $(1 - x_t - \mathbb{E}[e^{-r(\tau-t)}(1 - x_\tau)])$ and $(x_t - \mathbb{E}[e^{-r(\tau-t)}x_\tau])$ are positive (follows from that x_τ is a martingale and $e^{-r(\tau-t)} < 1$) and v_ω is increasing. Therefore, if type θ wants to stop, type $\theta' > \theta$ wants to stop too. \square

A.3 Proof of Lemma 2

Proof. \mathcal{T} and \mathcal{T}^{mon} are both consistent with Q . We show that monotone stopping ordering maximizes *ex post* welfare for all realized paths of (X, Q) . The claim follows once we show that for all types $\theta, \theta' \in [\underline{\theta}, \bar{\theta}]$ such that $\theta > \theta'$ and for all realized stopping times $t, t' \in \mathbb{R}_+$ such that $t \leq t'$,

$$e^{-rt}v_\omega(\theta) + e^{-rt'}v_\omega(\theta') \geq e^{-rt'}v_\omega(\theta) + e^{-rt}v_\omega(\theta').$$

The above condition is equivalent with $(e^{-rt} - e^{-rt'})(v_\omega(\theta) - v_\omega(\theta')) \geq 0$, which necessarily holds as $t \leq t'$ and $v_\omega(\theta) \geq v_\omega(\theta')$ by assumption if $\theta > \theta'$. \square

B Decentralized equilibrium

B.1 Proof of Proposition 1

Throughout the proof, we will also utilize the solution to the auxiliary optimal stopping problem, where the stock is assumed to be fixed at $q_t \equiv q$ forever:

Lemma 4. *Take the policy Q , where the stock is fixed at $q_t \equiv q$ forever. Then, it is optimal for θ to stop if and only if $x_t \geq \hat{x}_\theta(q)$, where*

$$\hat{x}_\theta(q) = \frac{\beta(q)v_L(\theta)}{\beta(q)v_L(\theta) + (1 - \beta(q))v_H(\theta)}.$$

Proof. This is a standard one-dimensional optimal stopping problem. See e.g. Dixit and Pindyck (1994) or the team problem in Bolton and Harris (1999) for derivations in equivalent problems. \square

We denote by $\bar{F}_\theta(x; q)$ the value function of the problem where the stock is fixed.

We now start with the uniqueness part of Proposition 1. Assume that Q is a decentralized equilibrium. By Lemma 1 the optimized stopping times are monotone in θ . Hence, in equilibrium, θ cannot stop before q_t reaches $q(\theta)$, nor can she delay stopping until q_t has reached a level higher than $q(\theta)$. Take some history h_t such that $q_t = q(\theta)$ and denote by $F_\theta(h_t)$ the equilibrium continuation value of type θ . It follows that the optimized continuation value for θ at $q(\theta)$ cannot exceed the continuation value of the constrained problem where q_t stays fixed at q forever, i.e. we have $F_\theta(h_t) = \bar{F}_\theta(x; q(\theta))$. Any policy other than stopping at threshold $\hat{x}_\theta(q(\theta)) = x^E(q(\theta))$ would give $F_\theta(h_t) < \bar{F}_\theta(x; q)$ and so stopping at the threshold is a necessary condition in a decentralized equilibrium. Since this holds for all types θ , we can conclude that if a decentralized equilibrium exists, then it must be the cutoff policy in Proposition 1.

Next, we show that the policy in Proposition 1 is indeed a decentralized equilibrium. Fix policy Q to be the cutoff policy in Proposition 1, and consider optimal stopping for θ against it. Note that Q defines a feasible region \mathbf{X} in (x, q) -space such that $(x_t, q_t) \in \mathbf{X}$ for all $t > 0$: $\mathbf{X} = \{(x, q) : 0 \leq q \leq 1, 0 < x \leq x^E(q)\}$.




Figure 6: Optimal stopping for type θ .

Since Q is a Markovian process, an optimal stopping rule against it can be expressed as some stopping region $S_\theta \subseteq \mathbf{X}$ in the feasible region. Denote by $F_\theta(x, q)$ the value function under optimally chosen stopping region S_θ :

$$F_\theta(x, q) = \mathbb{E} \left(e^{-r\tau(S_\theta)} u_\theta(x_{\tau(S_\theta)}) \mid x, q \right),$$

where $\tau(S_\theta) = \inf(t : (x_t, q_t) \in S_\theta)$ is the time of hitting S_θ and $u_\theta(x) := xv_H(\theta) + (1-x)v_L(\theta)$ is the stopping value. Clearly, $F_\theta(x, q) = u_\theta(x)$ whenever $(x_t, q_t) \in S_\theta$. At and above the boundary $x^E(q)$, where q_t increases, we must have:

$$\frac{\partial}{\partial q} [F_\theta(x, q)]_{x \geq x^E(q)} = 0. \quad (12)$$

The plan for the rest of the proof is to show that the optimal stopping region S_θ is the shaded region in Figure 6, i.e.

$$S_\theta = \left\{ (x, q) : q \geq q(\theta), x \in [\hat{x}_\theta(q), x^E(q)] \right\}. \quad (13)$$

As a first step, we note that it can never be optimal to stop for $q < q(\theta)$. This follows from the observation that for $q < q(\theta)$ any $(x, q) \in \mathbf{X}$ satisfies $x < \hat{x}_\theta(q)$, i.e. it is not optimal to stop even assuming q_t to stay fixed forever. The continuation value in the equilibrium where q_t may increase in the future must be at least as high as the continuation value in the problem where q_t stays fixed forever. Hence, $\bar{F}_\theta(x; q) > u_\theta(x, q)$ implies $F_\theta(x, q) > u_\theta(x, q)$ for all $x < \hat{x}_\theta(q)$ and so it cannot be optimal to stop.

As a second step, we will show that when $q \geq q(\theta)$, it is always optimal to stop at the boundary of \mathbf{X} , i.e. at $x = x^E(q)$. Suppose, to the contrary, that there is some $(x, q) \notin S_\theta$, where $x = x^E(q)$ and $q \geq q(\theta)$. This amounts to assuming that $F_\theta(x^E(q), q) > u_\theta(x^E(q))$, and we will show that this leads to a contradiction. First, suppose that it is optimal to stop at some (x', q) , where $x' < x^E(q)$. In that case $F_\theta(x', q) = u_\theta(x')$. Since $F_\theta(x, q)$ is convex in x and we necessarily have $F_\theta(x, q) \geq u_\theta(x)$ for all $x \in (x', x^E(q))$, we get that

$$\frac{\partial}{\partial x} [F_\theta(x, q)]_{x=x^E(q)} \geq v_H(\theta) - v_L(\theta) = \frac{\partial}{\partial x} [u_\theta(x)]_{x=x^E(q)} \quad (14)$$

Now, suppose that it is optimal to wait for all (x, q) , where $x < x^E(q)$, in which case $F_\theta(x, q) > u_\theta(x)$ for all $x < x^E(q)$ and $F_\theta(0, q) = 0$. In that case, following the same argument as with the value function $\bar{F}_\theta(x; q)$ for fixed q , the value function $F_\theta(x, q)$ must take the form $F_\theta(x, q) = A_\theta(q) \Phi(x, q)$ for some function $A_\theta(q)$ and hence $\frac{\partial}{\partial x} [F_\theta(x, q)]_{x=x^E(q)} = A_\theta(q) \Phi_x(x^E(q), q)$. Our assumption $F_\theta(x^E(q), q) > u_\theta(x^E(q))$ is equivalent to

$$A_\theta(q) \Phi(x^E(q), q) > x^E(q) v_H(\theta) + (1 - x^E(q)) v_L(\theta),$$

which further implies

$$\begin{aligned} \frac{\partial}{\partial x} [F_\theta(x, q)]_{x=x^E(q)} &> \frac{\Phi_x(x(q), q)}{\Phi(x(q), q)} [x^E(q) v_H(\theta) + (1 - x^E(q)) v_L(\theta)] \\ &= \frac{\beta(q) - x^E(q)}{(1 - x^E(q))} v_H(\theta) + \frac{\beta(q) - x^E(q)}{x^E(q)} v_L(\theta). \end{aligned}$$

The last expression is greater than $v_H(\theta) - v_L(\theta)$ if and only if

$$x' \geq \frac{\beta(q) v_L(\theta)}{\beta(q) v_L(\theta) + (1 - \beta(q)) v_H(\theta)} = \hat{x}_\theta(q),$$

which is the case if and only if $q \geq q(\theta)$. Therefore, we may conclude that (14) holds in this case too.

Given that (14) holds, the rate of change in $F_\theta(x, q)$ along the boundary is

$$\begin{aligned} \frac{d}{dq} F_\theta(x^E(q), q) &= \frac{\partial}{\partial x} [F_\theta(x, q)]_{x=x^E(q)} \frac{d}{dq} x^E(q) + \frac{\partial}{\partial q} [F_\theta(x, q)]_{x=x^E(q)} \\ &= \frac{\partial}{\partial x} [F_\theta(x, q)]_{x=x^E(q)} \frac{d}{dq} x^E(q) \geq [v_H(\theta) - v_L(\theta)] \frac{d}{dq} x^E(q) = \frac{d}{dq} u_\theta(x^E(q)), \end{aligned}$$

where the second term of the first line disappears by (12) and where the inequality is implied by (14).

But if $F_\theta(x^E(q), q) > u_\theta(x^E(q))$ implies $\frac{d}{dq}F_\theta(x^E(q), q) \geq \frac{d}{dq}u_\theta(x^E(q))$, then it must further imply $F_\theta(x^E(q'), q') > u_\theta(x^E(q'))$ for all $q' \in [q, 1]$, and in particular $F_\theta(x^E(1), 1) > u_\theta(x^E(1))$. But $x^E(1) = 1$, so this yields $F_\theta(x^E(1), 1) > v_H(\theta)$, which is a contradiction since the value cannot be higher than the stopping payoff under certainty of state $\omega = H$.

We conclude that it is optimal to stop at all boundary points for $q > q(\theta)$. To see that this implies that it is also optimal to stop within the whole shaded region in Figure 6, i.e. $\{(x, q) : q \geq q(\theta), x \in [\hat{x}_\theta(q), x^E(q)]\} \in S_\theta$, note that q_t can only increase if x_t reaches $x^E(q)$. Since θ stops at latest when x_t reaches $x^E(q)$, the optimal continuation value $F_\theta(x, q)$ cannot exceed the corresponding value with q fixed, i.e. $\bar{F}_\theta(x; q)$. To reach that value, θ must stop at all points $[\hat{x}_\theta(q), x^E(q)]$.

We have now shown that the stopping rule defined in (13) maximizes (3) for policy Q . Since q_t can only increase at the boundary points $x^E(q)$, the first point in S_θ ever reached is $(\hat{x}_\theta(q(\theta)), q(\theta))$ and so the optimal stopping rule commands θ to stop exactly when q_t reaches $1 - F(\theta)$ and is therefore consistent with Q . We can conclude that Q is a decentralized equilibrium. \square

C Socially optimal policy

We use the derivatives of $\Phi(x, q)$ in many proofs of this section:

$$\begin{aligned}\Phi &= \left(\frac{x}{1-x}\right)^{\beta(q)} (1-x), \Phi_q = \Phi\beta'(q) \ln\left(\frac{x}{1-x}\right), \\ \Phi_x &= \Phi\frac{(\beta(q)-x)}{x(1-x)}, \Phi_{xx} = \Phi\beta(q)\frac{(\beta(q)-1)}{x^2(1-x)^2}, \\ \Phi_{qx} &= \Phi\beta'(q)x^{-1}(1-x)^{-1} \left[1 + (\beta(q)-x) \ln\left(\frac{x}{1-x}\right)\right], \\ \Phi_{xxq} &= \Phi\frac{\beta'(q)}{x^2(1-x^2)} \left[\beta(q) + (\beta(q)-1)(1 + \beta(q) \ln\left(\frac{x}{1-x}\right))\right].\end{aligned}$$

Deriving the differential equation

We first show that the value matching and smooth pasting conditions, (7) and (8), imply the differential equation in (9). Solving (7) and (8) for $B_q(q)$ and $B(q)$ yields

$$B_q(q) = A^1(x^*(q), q)x^*(q) + A^2(x^*(q), q), \quad (15)$$

$$B(q) = U^1(x^*(q), q)x^*(q) + U^2(x^*(q), q), \quad (16)$$

where

$$\begin{aligned} A^1(x, q) &: = \frac{-\Phi_{qx}(x, q)(v_H(q) - v_L(q))}{\Phi(x, q)\Phi_{qx}(x, q) - \Phi_q(x, q)\Phi_x(x, q)}, \\ A^2(x, q) &: = \frac{\Phi_{qx}(x, q)(-v_L(q)) + \Phi_q(x, q)(v_H(q) - v_L(q))}{\Phi(x, q)\Phi_{qx}(x, q) - \Phi_q(x, q)\Phi_x(x, q)}, \\ U^1(x, q) &: = \frac{\Phi_x(x, q)(v_H(q) - v_L(q))}{\Phi(x, q)\Phi_{qx}(x, q) - \Phi_q(x, q)\Phi_x(x, q)}, \\ U^2(x, q) &: = \frac{-\Phi_x(x, q)(-v_L(q)) - \Phi(x, q)(v_H(q) - v_L(q))}{\Phi(x, q)\Phi_{qx}(x, q) - \Phi_q(x, q)\Phi_x(x, q)}. \end{aligned}$$

Differentiating (16) with respect to q and using the chain rule gives

$$\begin{aligned} B_q(q) &= \left[U_x^1(x^*(q), q)x^{*'}(q) + U_q^1(x^*(q), q) \right] x^*(q) + U^1(x^*(q), q)x^{*'}(q) \\ &\quad + U_x^2(x^*(q), q)x^{*'}(q) + U_q^2(x^*(q), q) \end{aligned} \quad (17)$$

Equating (15) and (17), solving for $x^{*'}(q)$, and simplifying yields the expression (9) in the text.

Any solution that satisfies the differential equation (9) must be continuous.

C.1 Proof of Proposition 2

The proof contains three parts. In part 1, we show that the initial value problem (9) has a unique solution $x^*(q)$ with the property $x^*(q) < x^E(q)$ for all $q < 1$. We also show that $x^*(q)$ is continuous and strictly increasing and hence defines a boundary policy. In part 2, we show that our candidate policy $x^*(q)$ satisfies the HJB equation (5). In part 3, we verify that the solution to the HJB equation solves the original problem.

Part 1: solution to the initial value problem (9)

We first establish some key properties of function g in (9) (all proofs of the lemmas are in Appendix C.2):

Lemma 5. *For all (x, q) such that $q < 1$ and $x \leq x^E(q)$, function $g(x, q)$ in (9) is strictly positive and strictly increasing in x and it is Lipschitz continuous for all $q \in [0, q_1]$ if $q_1 < 1$ and for all $x \leq x^E(q)$. Furthermore, $g(x^E(q), q) > x^{E'}(q)$ for $q < 1$ and $\lim_{q \rightarrow 1} g(x^E(q), q) = x^{E'}(1)$.*

The singularity at (1,1) prevents us from directly applying the Picard-Lindelöf theorem to show the existence and uniqueness of a solution to the initial value problem (9). Instead, we note that the requirements for the Picard-Lindelöf theorem are satisfied for all initial conditions $x(q_1) = x_1$ where $x(q_1) \leq x^E(q_1)$ and $q_1 < 1$, and hence each such initial value problem defines a unique solution. Since g is increasing in x , these solutions diverge when approaching (1,1) and hence at most one path can approach (1,1) from below the decentralized policy. The fact that $\lim_{q \rightarrow 1} g(x^E(q), q) = x^{E'}(1)$ implies that there is a path that approaches (1,1) from the same direction as the decentralized policy $x^E(q)$ and the fact that $g(x^E(q), q) > x^{E'}(q)$ for $q < 1$ implies that such a path must be strictly below the decentralized solution for all $q < 1$. It follows that the initial value problem has a unique solution below the decentralized solution.

We have now shown that the initial value problem (9) has a unique solution x^* such that $x^*(q) \leq x^E(q)$ for all $x \leq q$. This solution $x^*(q)$ is continuous and strictly increasing in q , and it is our candidate policy.

Part 2: our candidate x^* solves the HJB equation

Fix $x^*(q)$ to be the candidate policy defined in the proposition and let $q^*(x)$ be its inverse with the convention $q^*(x) = 0$ for $x \leq x^*(0)$. Its associated value function is

$$V(x, q) = \begin{cases} \int_q^{q^*(x)} (xv_H(s) + (1-x)v_L(s)) ds + V(x, q^*(x)), & \text{for } q < q^*(x) \\ B(q) \Phi(x, q), & \text{for } q \geq q^*(x), \end{cases} \quad (18)$$

where $B(q)$ is given by (16). By construction, for $q \geq q^*(x)$, $V(x, q)$ satisfies

$$rV(x, q) = \frac{1}{2}V_{xx}(x, q) \frac{x^2(1-x)^2}{\sigma^2} q \quad (19)$$

and at the boundary $q = q^*(x)$, the value matching and smooth pasting conditions (7) and (8) hold:

$$V_q(x, q^*(x)) + xv_H(q^*(x)) + (1-x)v_L(q^*(x)) = 0, \quad (20)$$

$$V_{qx}(x, q^*(x)) + v_H(q^*(x)) - v_L(q^*(x)) = 0. \quad (21)$$

Differentiating (19) with respect to q , we have

$$rV_q(x, q) = \frac{1}{2}V_{xx}(x, q) \frac{x^2(1-x)^2}{\sigma^2} + \frac{1}{2}V_{xxq}(x, q) \frac{x^2(1-x)^2}{\sigma^2} q, \quad (22)$$

which allows us to re-write (20) as:

$$\begin{aligned} r[xv_H(q^*(x)) + (1-x)v_L(q^*(x))] + \frac{1}{2}V_{xx}(x, q^*(x)) \frac{x^2(1-x)^2}{\sigma^2} \\ + \frac{1}{2}V_{xxq}(x, q^*(x)) \frac{x^2(1-x)^2}{\sigma^2} q = 0. \end{aligned} \quad (23)$$

We next state three lemmas that concern the partials of the value function below, above, and at the boundary, respectively. Their proofs are in a separate section, Appendix C.2.

Lemma 6. *For $q \geq q^*(x)$, we have $V_q(x, q) + xv_H(q) + (1-x)v_L(q) \leq 0$.*

Lemma 7. *For $q < q^*(x)$, we have $V_{xx}(x, q) = V_{xx}(x, q^*(x))$, $V_{qq}(x, q) = V_{qq}(x, q^*(x))$, and $V_{xxq}(x, q) = 0$.*

Lemma 8. *For $q = q^*(x)$, we have $V_{xxq}(x, q) < 0$.*

We are now ready to show that our candidate policy satisfies the HJB-equation (5), which we re-write here using notation q' instead of q^* for the maximizer (this is to avoid confusion with boundary $q^*(x)$):

$$rV(x, q) = \max_{q' > q} \left(r \int_q^{q'} (xv_H(s) + (1-x)v_L(s)) ds + \frac{1}{2}V_{xx}(x, q') \frac{x^2(1-x)^2}{\sigma^2} q' \right). \quad (24)$$

The right-hand side of this equation is a continuous function in q' and its derivative with respect to q' is

$$r(xv_H(q') + (1-x)v_L(q')) + \frac{1}{2}V_{xx}(x, q') \frac{x^2(1-x)^2}{\sigma^2} + \frac{1}{2}V_{xxq}(x, q') \frac{x^2(1-x)^2}{\sigma^2} q'. \quad (25)$$

Let us inspect the sign of this for different values of q' . For $q' \geq q^*(x)$, we can use (22) to write (25) as

$$r(xv_H(q') + (1-x)v_L(q')) + rV_q(x, q),$$

which is negative by lemma 6. It follows that whenever $q \geq q^*(x)$, the right-hand side of (24) is maximized by choosing $q' = q$, i.e. keeping q fixed.

For $q' < q^*(x)$, we can use lemma 7 to write (25) as

$$r(xv_H(q') + (1-x)v_L(q')) + \frac{1}{2}V_{xx}(x, q^*(x)) \frac{x^2(1-x)^2}{\sigma^2},$$

which is decreasing in q' . Moreover, combining (23) and Lemma 8 we can conclude that it is positive in the limit $q' \rightarrow q^*(x)$, and hence it is positive for all $q' < q^*(x)$. Since the right-hand side of (24) is continuous, and its derivative is positive (negative) for $q' < q^*(x)$ ($q' \geq q^*(x)$), it is maximized at $q' = q^*(x)$ if $q < q^*(x)$.

We have now shown that for any $x \in (0, 1)$, the right-hand side of the HJB equation is maximized by choosing $q' = \max\{q, q^*(x)\}$. Furthermore, since $V(x, q)$ satisfies (19) for $q \geq q^*(x)$, the left- and right-hand sides of (24) coincide with this choice of q' . Hence, we have shown that $V(x, q)$ defined in (18) satisfies the HJB-equation.

Part 3: verification

The verification of the solution follows from the standard arguments in the literature (see e.g. Fleming and Soner (2006)). Let V^* solve the HJB equation (5) and let $q^*(x, q) = \max\{q, q^*(x)\}$ be the corresponding q^* . Then, let $T \geq t$ be the time at which the candidate value function is evaluated. From generalized Itô's

formula we have¹²

$$\begin{aligned}
e^{-rT}V^*(x_T, q_T) &= e^{-rt}V^*(x_t, q_t) - \int_t^T e^{-rs}rV^*(x_s, q_s)ds + \int_t^T e^{-rs}V_x^*(x_s, q_s)dx_s \\
&+ \int_t^T e^{-rs}V_q^*(x_s, q_s)dq_s + \frac{1}{2} \int_t^T e^{-rs}V_{xx}^*(x_s, q_s)d[x]_s + \frac{1}{2} \int_t^T e^{-rs}V_{qq}^*(x_s, q_s)d[q]_s \\
&+ \int_t^T e^{-rs}V_{qs}^*(x_s, q_s)d[q, x]_s
\end{aligned}$$

where $d[x]_t$ and $d[q]_t$ are the quadratic variations of x and q and $d[x, y]_t$ is their quadratic covariation. The process Q_t has bounded variation and hence $d[q]_t = d[x, y]_t = 0$. Notice also that $dx_t = x_t(1-x_t)\sigma^{-1}\sqrt{q_t}dw_t$ and $d[x]_t = x_t^2(1-x_t)^2\sigma^{-2}q_t dt$. We can further simplify the equation by noting that $V_q^*dq = -(xv_H(q) + (1-x)v_L(q))dq$. The HJB equation gives an upper bound for $\frac{q_s}{\sigma^2}x_s^2(1-x_s)^2V_{xx}^*(x_s, q_s) - rV^*(x_s, q_s) \leq \int_{q_s}^{q^*(x_s, q_s)} (xv_H(q) + (1-x)v_L(q))dq$, which equals zero for almost all s . Combining gives:

$$\begin{aligned}
e^{-rT}V^*(x_T, q_T) &\leq e^{-rt}V^*(x_t, q_t) - \int_t^T e^{-rs}(x_s v_H(q_s) + (1-x_s)v_L(q_s))dq_s \\
&+ \int_t^T e^{-rs}V_x^*(x_s, q_s)\frac{\sqrt{q_s}}{\sigma}x_s(1-x_s)dw_s.
\end{aligned}$$

Taking conditional expectations, multiplying by $-e^{rt}$ and simplifying then gives

$$V^*(x_t, q_t) \geq \mathbb{E} \left[\int_t^T e^{-r(t-s)}(x_s \pi_H(q_s) + (1-x_s)\pi_L(q_s))ds + e^{-r(T-t)}V^*(x_T, q_T) | \mathcal{F}_t \right].$$

The candidate value function is bounded and therefore clearly satisfies the transversality condition: $\lim_{T \rightarrow \infty} \mathbb{E}[e^{-r(T-t)}V^*(x_T, q_T)] = 0$. Hence, taking the limit $T \rightarrow \infty$ gives that $V^*(x, q) \geq \max_Q U(Q; x, q)$.

The last step is to use the fact that Q , induced by policy x^* , achieves the pointwise maximum of the HJB-equation and thus the inequalities above become equalities: $V^*(x, q) = \max_Q U(Q; x, q)$. Our solution solves the original problem.

¹²To see that $V \in C^2$ check V_x at the boundary. The continuity of V_{xx} and V_{qq} follows from Lemma 7 and the continuity of V_q and V_{qx} are implied by the value matching and smooth pasting conditions.

C.2 Proof of Lemmas 5, 6, 7, and 8

Proof of Lemma 5. Taking the derivative of $g(x, q)$ with respect x gives:

$$\begin{aligned}
g_x(x, q) = & - \left[\beta''(q) \left(x^2(1-2x)(\beta(q)-1)^3 v_H(q)^2 - 2(1-x)x\beta(q)(\beta(q)-1) \right. \right. \\
& \times v_H(q)v_L(q)((1-2x)\beta(q)-x) + (1-x)^2(1-2x)\beta(q)^3 v_L(q)^2 \left. \right) \\
& + \beta'(q) \left(2x^2(2x-1)(\beta(q)-1)^2 v_H(q)^2 \beta'(q) + (1-x)^2 \beta(q)^2 v_L(q) \right. \\
& \times \left(2(1-2x)v_L(q)\beta'(q) - 2x(\beta(q)-1)v_H'(q) - (1-2x)\beta(q)v_L'(q) \right) \\
& + xv_H(q) \left(4(1-x)v_L(q)\beta'(q) \left((1-2x)\beta(q)^2 + 2x\beta(q) + x \right) \right. \\
& \left. \left. - x(\beta(q)-1)^2 \left((1-2x)(\beta(q)-1)v_H'(q) + 2(1-x)\beta(q)v_L'(q) \right) \right) \right] / \\
& \left[(x(\beta(q)-1)^2 v_H(q) + (1-x)(\beta(q))^2 v_L(q))^2 \beta'(q) \right].
\end{aligned}$$

Both $g(x, q)$ and $g_x(x, q)$ are bounded if their denominators are bounded away from zero. We show that this is true if $q < 1$ and $x \leq x^E(q)$ by showing that it hold at $x = x^E$:

$$x^E(q)(\beta(q)-1)^2 v_H(q) + (1-x^E(q))(\beta(q))^2 v_L(q) < 0, \quad (26)$$

for all $q \in [0, 1)$. Notice that the left-side is increasing in x and hence (26) implies the same inequality for all lower x . The condition (26) is equivalent with

$$\frac{\beta(q)(\beta(q)v_L(q) - (\beta(q)-1)v_H(q))}{\beta(q)^2 v_L(q) - (\beta(q)-1)^2 v_H(q)} > 1 \iff (\beta(q)-1)v_H(q) > 0,$$

which holds as $\beta(q) > 1$ and $v_H(q) > 0$. We can conclude that g and g_x are bounded and continuous in both x and q for all (x, q) such that $q < 1$ and $x \leq x^E(q)$ and hence g is Lipschitz when we gap q away from 1. The denominator of $g(x, q)$ is strictly positive.

To see that $g(x, q) > 0$, it is now enough to show that the numerator of (9) is strictly positive. First notice that the second term inside the brackets is always positive but the first term can be negative.¹³ The first term is scaled by x , while the second term is scaled by $(1-x)$. Therefore, if the numerator is positive at a belief above the boundary, it must be positive for the belief at the boundary

¹³This follows from $v_L(q) < 0, v_L'(q) < 0, \beta'(q) < 0, \beta(q) > 1$ and that $\beta(q)\beta''(q) > 2(\beta'(q))^2$.

as well. Since the decentralized belief, $x^E(q)$, is always above the fully optimal boundary, we can use it to show that the numerator is positive.

Plugging in $x^E(q)$ to the numerator of (9) and dividing by $x(1-x)$ gives:

$$\frac{\beta(q)v_L(q) \left(\beta'(q) (\beta(q) - 1) v'_H(q) - ((\beta(q) - 1) \beta''(q) - 2(\beta'(q))^2) v_H(q) \right)}{\beta(q)v_L(q) + (1 - \beta(q))v_H(q)} + \frac{(1 - \beta(q)) v_H(q) \left(\beta'(q)\beta(q)v'_L(q) - (\beta(q)\beta''(q) - 2(\beta'(q))^2) v_L(q) \right)}{\beta(q)v_L(q) + (1 - \beta(q))v_H(q)}.$$

Since the denominator is negative ($v_L < 0$ and $\beta > 1$), this is proportional to

$$[v_H(q)v'_L(q) - v'_H(q)v_L(q)]\beta'(q)\beta(q)(\beta(q) - 1) - 2v_H(q)v_L(q)(\beta'(q))^2,$$

which is always positive because $v_H(q) > 0$ and $v_L(q), v'_H(q), v'_L(q) < 0$. Hence, $q(x, q) > 0$ for all $q \in [0, 1)$ and $x \leq x^E(q)$.

Similar direct calculations show that $g_x > 0$ for all (x, q) such that $q < 1$ and $x \leq x^E(q)$.

Next, insert $x^E(q)$ to (9):

$$\begin{aligned} g(x^E(q), q) &= \frac{-\beta(1-\beta)v_Lv_H}{(\beta v_L + (1-\beta)v_H)^2} \left(\frac{\beta' \beta (1-\beta)(v_L v'_H - v'_L v_H)}{\beta v_L + (1-\beta)v_H} \right. \\ &\quad \left. + \frac{\beta v_L \gamma_H (-2\beta'^2 + (\beta-1)\beta'')}{\beta v_L + (1-\beta)v_H} + \frac{(\beta-1)v_L v_H (-2\beta'^2 + \beta\beta'')}{\beta v_L + (1-\beta)v_H} \right) \\ &= \frac{v_H (2v_L \beta' - (\beta-1)\beta \gamma'_L) + (\beta-1)\beta v_L v'_H}{((\beta-1)v_H - \beta v_L)^2}. \end{aligned}$$

The derivative of the decentralized policy x^E is

$$x^{E'}(q) = \frac{v_H (v_L \beta' - (\beta-1)\beta v'_L) + (\beta-1)\beta v_L v'_H}{((\beta-1)v_H - \beta v_L)^2}.$$

By subtracting $x^{E'}(q)$ from $g(x^E(q), q)$, we get

$$g(x^E(q), q) - x^{E'}(q) = \frac{\beta'(q)v_L(q)v_H(q)}{(\beta(q)v_L(q) + (1 - \beta(q))v_H(q))^2}.$$

This expression is strictly positive for $q < 1$ and goes to zero as q goes to 1 (since $v_H(q) \rightarrow 0$). \square

Proof of Lemma 6. If the claim is not true, there must be some x and $q > q^*(x)$ such that

$$V_q(x, q) + xv_H(q) + (1-x)v_L(q) > 0. \quad (27)$$

We show that this leads to a contradiction by showing that (27) implies $V_{qx}(x, q) + v_H(q) - v_L(q) > 0$, which further implies that (27) holds also for all beliefs in $[x, x^*(q)]$, including $V_q(x^*(q), q) + x^*(q)v_H(q) + (1 - x^*(q))v_L(q) > 0$, which contradicts the value matching condition (20).

It remains to show that (27) implies $V_{qx}(x, q) + v_H(q) - v_L(q) > 0$. First notice that $V_q(x, q) = B_q(q)\Phi(x, q) + B(q)\Phi_q(x, q)$, which then together with (27) implies

$$B_q > -\frac{\Phi_q}{\Phi}B - \frac{xv_H + (1-x)v_L}{\Phi}$$

where we have left out all dependencies to simplify notation. We now get the following lower bound:

$$\begin{aligned} V_{qx} + v_H - v_L &= B_q\Phi_x + B\Phi_{qx} + v_H - v_L \\ &> -\frac{\Phi_q\Phi_x}{\Phi}B - \frac{\Phi_x}{\Phi}(xv_H + (1-x)v_L) + B\Phi_{qx} + v_H - v_L \\ &= \Phi^{-1}[B(\Phi_{qx}\Phi - \Phi_q\Phi_x) + \Phi(v_H - v_L) - \Phi_x(xv_H + (1-x)v_L)]. \end{aligned} \quad (28)$$

The first term can be simplified as

$$\begin{aligned} \Phi^{-1}B(\Phi_{qx}\Phi - \Phi_q\Phi_x) &= \frac{B\Phi\beta'}{x(1-x)} = \frac{\Phi\beta'}{x(1-x)} \frac{\Phi_x^*(x^*v_H + (1-x^*)v_L) - \Phi^*(v_H - v_L)}{\Phi_{qx}^*\Phi^* - \Phi_q^*\Phi_x^*} \\ &= \frac{x^*(1-x^*)}{x(1-x)} \frac{\Phi}{\Phi^*\Phi^*} [\Phi_x^*(x^*v_H + (1-x^*)v_L) - \Phi^*(v_H - v_L)], \end{aligned}$$

where the notation Φ^* refers to $\Phi(x^*(q), q)$.

Now, (28) becomes

$$\begin{aligned} &\frac{x^*(1-x^*)}{x(1-x)} \frac{\Phi}{\Phi^*\Phi^*} [\Phi_x^*(x^*v_H + (1-x^*)v_L) - \Phi^*(v_H - v_L)] \\ &- \frac{1}{\Phi} [\Phi_x(xv_H + (1-x)v_L) - \Phi(v_H - v_L)] \\ &= \frac{1}{x(1-x)} \left(\frac{\Phi}{\Phi^*} ((\beta-1)x^*v_H + \beta(1-x^*)v_L) - ((\beta-1)xv_H + \beta(1-x)v_L) \right), \end{aligned} \quad (29)$$

where we have used the following for both terms inside the brackets:

$$\begin{aligned} \Phi(v_H - v_L) - \Phi_x(xv_H + (1-x)v_L) &= \Phi(v_H - v_L) - \Phi \frac{\beta-x}{x(1-x)} (xv_H + (1-x)v_L) \\ &= \frac{-\Phi}{x(1-x)} ((\beta-1)xv_H + \beta(1-x)v_L). \end{aligned}$$

To conclude that (29) is larger than 0, notice first that $(\beta - 1)xv_H + \beta(1 - x)v_L < 0$ whenever $x < x^E(q)$ and that it is increasing in x . Then observe that $\Phi/\Phi^* \in (0, 1)$ and hence $(\beta - 1)xv_H + \beta(1 - x)v_L < (\Phi/\Phi^*)((\beta - 1)x^*v_H + \beta(1 - x^*)v_L)$.

We conclude that $V_q + xv_H + (1 - x)v_L > 0$ implies $V_{qx} + v_H - v_L > 0$ and the proof is complete. \square

Proof of Lemma 7. Fixing some (x, q) such that $q < q^*(x)$, differentiating (18) twice with respect to x , and simplifying gives:

$$\begin{aligned} V_{xx}(x, q) &= V_{xx}(x, q^*(x)) \\ &+ 2(q^*)'(x)(V_{qx}(x, q^*(x)) + v_H(q^*(x)) - v_L(q^*(x))) \\ &+ (q^*)''(x)(V_q(x, q^*(x)) + xv'_H(q^*(x)) + (1 - x)v'_L(q^*(x))) \\ &+ \left((q^*)'(x)\right)^2(V_{qq}(x, q^*(x)) + xv'_H(q^*(x)) + (1 - x)v'_L(q^*(x))). \end{aligned} \quad (30)$$

The second term on the right-hand side vanishes by the value-matching condition (20) and the third term vanishes by the smooth-pasting condition (21). Let us look at the last term. First, since (20) holds along the boundary $(x, q^*(x))$, we can totally differentiate it with respect to x to get:

$$\begin{aligned} 0 &= V_{qx}(x, q^*(x)) + V_{qq}(x, q^*(x))(q^*)'(x) + v_H(q^*(x)) - v_L(q^*(x)) \\ &+ [xv'_H(q^*(x)) + (1 - x)v'_L(q^*(x))](q^*)'(x). \end{aligned}$$

Applying (21), several terms disappear and this reduces to

$$V_{qq}(x, q^*(x)) + xv'_H(q^*(x)) + (1 - x)v'_L(q^*(x)) = 0.$$

The last term in (30) vanishes as well, and it follows that $V_{xx}(x, q) = V_{xx}(x, q^*(x))$. Since this holds for any $q < q^*(x)$, it immediately implies that $V_{xxq}(x, q) = 0$. \square

Proof of Lemma 8. This is by direct computation. Recall that the value function for $q \geq q^*(x)$ is $V(x, q) = B(q)\Phi(x, q)$ and hence

$$V_{xxq} = B_q(q)\Phi_{xx}(x, q) + B(q)\Phi_{xxq}(x, q).$$

Plugging in the expressions for $B_q(q)$ and $B(q)$ from (15) and (16), multiplying by r , simplifying, and evaluating at $q = q^*(x)$ gives:

$$\begin{aligned} rV_{xxq}(x, q^*(x)) &= \frac{(\beta(q^*(x)) - 1)^2 + x(1-x)}{x^2(1-x)^2} v_L(q^*(x)) \\ &\quad - \frac{\beta(q^*(x))(\beta(q^*(x)) - 1)}{x^2(1-x)^2} (v_H(q^*(x)) - v_L(q^*(x))). \end{aligned}$$

Noting that $\beta(q^*(x)) > 1$, $v_L(q^*(x)) < 0$, and $v_H(q^*(x)) - v_L(q^*(x)) > 0$, it follows that $V_{xxq}(x, q^*(x)) < 0$.

□

C.3 Proof of Proposition 3

Proof. First, recall that $x^*(0) < x^E(0) = x^{stat}(0)$ by the proof of Proposition 2. Using this together with the continuity of the policy functions we find that there exists $\underline{q} > 0$ such that $x^{stat}(q) > x^*(q)$ for all $q < \underline{q}$. As the policy functions are strictly increasing and continuous, the stocks $q^*(x)$ and $q^{stat}(x)$ are pinned down as the inverse of the policy functions for all $x \geq x^*(0)$ and $x \geq x^{stat}(0)$ respectively. In addition, $q^*(x) = 0$ for all $x \leq x^*(0)$ and $q^{stat}(x) = 0$ for all $x \leq x^{stat}(0)$, and hence q^* and q^{stat} are continuous.

Let $\underline{x} := x^{stat}(\underline{q}) > x^{stat}(0)$ where the inequality follows from x^{stat} being strictly increasing. Then, $q^{stat}(x) < q^*(x)$ for all $x \in [x^{stat}(0), \underline{x}]$ by that q^* and q^{stat} are the inverse functions of x^* and x^{stat} . Furthermore, $q^{stat}(x) = 0 < q^*(x)$ for all $x \in [x^*(0), x^{stat}(0)]$, which completes the proof.

Next, we show the other direction by showing that $x^*(1) = x^{stat}(1) = 1$ and $x_q^*(1) < x_q^{stat}(1)$. The first part is immediate. For the second part, use Lemma 5 and the uniqueness of the solution to get $x_q^*(1) = x_q^E(1)$. Now it is enough to show that the derivative of the equilibrium is smaller than of the myopic solution:

$$x_q^E(1) - x_q^{stat}(1) = \frac{(\beta - 1)\beta v_L v'_H}{(\beta v_L)^2} - \frac{v_L v'_H}{(v_L)^2} = -\frac{v_L v'_H}{\beta v_L^2} < 0.$$

The myopic and optimal solutions meet at $q = 1$ but the optimal solution reaches the point above the myopic solutions. Hence, by continuity there must

exist $\bar{q} < 1$ such that $x^*(q) > x^{stat}(q)$ for all $q \in (\bar{q}, 1)$, which then further implies the existence of $\bar{x} < 1$ by the same argument as used above for \underline{x} . \square

C.4 Proof of Proposition 4

Proof. Part (a): We show the result by contradiction. By using the solution from Proposition 2 and the value function derived in its proof, we show that $q^* = 0$ cannot maximize the HJB equation (5) in the limit as $\sigma \rightarrow 0$ unless $\sqrt{q_\sigma^*(x)}/\sigma \rightarrow \infty$. If $q^*(x)$ goes to any other value than 0, the claim immediately follows.

By taking the first order condition from (5), we get

$$xv_H(q^*) + (1-x)v_L(q^*) + \frac{1}{2} \frac{x^2(1-x)^2}{\sigma^2} (V_{xx}(x, q^*) + V_{xxq}(x, q^*)q^*).$$

The first order condition is necessarily strictly positive at $q^* = 0$ in the limit as $\sigma \rightarrow 0$ once we show that $V_{xx}(x, q) > 0$ and $V_{xxq}(x, q)$ is finite.

Recall that the value function is $V(x, q) = B(q)\Phi(x, q)$ and its derivatives are then $V_{xx} = B(q)\Phi_{xx}$ and $V_{xxq} = B_q(q)\Phi_{xx} + B(q)\Phi_{xxq}$. By plugging in the values of Φ_{xx} , we get

$$V_{xx} = B(q)\beta(q)\Phi \frac{(\beta(q) - 1)}{x^2(1-x)^2}.$$

We know that $B > 0$ for all $q < 1$ in the optimal solution and that $\Phi > 0$ for all $x \in (0, 1)$. Then, $V_{xx} > 0$ whenever $\beta > 1$ which is true whenever the signal-to-noise ration is finite.

We can write V_{xxq} as

$$V_{xxq} = \frac{(\Phi_x \Phi_{xxq} - \Phi_{qx} \Phi_{xx})(xv_H + (1-x)v_L)}{\Phi \Phi_{qx} - \Phi_q \Phi_x} + \frac{(\Phi_q \Phi_{xx} - \Phi \Phi_{xxq})(v_H - v_L)}{\Phi \Phi_{qx} - \Phi_q \Phi_x}.$$

The first term equals $\frac{(\beta-x)^2+x(1-x)}{x^2(1-x)^2}(xv_H + (1-x)v_L)$ and the second term equals $-\frac{\beta+(\beta-1)\ln(\frac{x}{1-x})}{x(1-x)}(v_H - v_L)$. Both are finite for all $x \in (0, 1)$.

Hence, we conclude that for the first order condition to be satisfied, we must have $\sqrt{q_\sigma^*(x)}/\sigma \rightarrow \infty$ as $\sigma \rightarrow 0$.

Part (b): We fix the belief to be $x \in (0, 1)$. By rearranging the solution in

Proposition 1, we get

$$\beta(q) = \frac{xv_H(q)}{xv_H(q) + (1-x)v_L(q)}.$$

We take the limit $\lim_{\sigma \rightarrow 0} \beta(q_\sigma^E(x)) = \frac{xv_H(0)}{xv_H(0) + (1-x)v_L(0)}$, which is strictly larger than 1 for all $x > x^{stat}(q)$ and hence further implying that $\lim_{\sigma \rightarrow 0} \sqrt{q_\sigma^E(x)}/\sigma < \infty$. More precisely, we get the limit of the signal-to-noise ratio as $a(x)$ satisfying $\frac{xv_H(0)}{xv_H(0) + (1-x)v_L(0)} = \frac{1}{2} \left(1 + \sqrt{1 + 8ra(x)^{-2}}\right)$. \square

D Mechanism design

D.1 Proof of Proposition 5

Proof. Fix policy Q and posted price $P(x, q)$ as in Proposition 5. To show that Q is a decentralized equilibrium, we have to show that it is optimal for type θ to stop at $\tau(\theta) := \inf \{t : q_t \geq 1 - F(\theta)\}$. We call $\tau(\theta)$ the intended stopping time for type θ .

For an individual player this is a Markovian stopping problem with fixed policy Q and stopping payoff given by (10). From the structure of the problem it is clear that it is optimal for θ to stop at the first hitting time of *some* boundary point. To see this, note that whenever the state (x, q) is above the boundary, it moves immediately to the boundary point $(x, \tilde{q}(x))$ and since we have defined $P(x, q) = P(x, \tilde{q}(x))$ above the boundary, any player is indifferent between stopping at (x, q) or $(x, \tilde{q}(x))$ (or any point between those). Furthermore, we have defined $P(x, q)$ to be so large below the boundary that no player wants to stop there. Since every boundary point is an intended stopping point for some $\theta \in [\underline{\theta}, \bar{\theta}]$, we only have to show that it is better for θ to stop at the intended stopping time $\tau(\theta)$ than at the intended stopping time for some other type $\tau(\tilde{\theta})$, $\tilde{\theta} \neq \theta$.

Denote by $U(\theta, \tilde{\theta})$ the expected value at time zero for type θ who intends to stop at $\tau(\tilde{\theta})$, $\tilde{\theta} \neq \theta$. To complete the proof, we must show that $U(\theta, \theta) \geq U(\theta, \tilde{\theta})$ for all θ and for all $\tilde{\theta}$.

We next proceed to derive the expression for $U(\theta, \tilde{\theta})$. Denote by $P_0(\tilde{\theta})$ the

expected discounted transfer payment that θ has to pay if she stops at the intended stopping time $\tau(\tilde{\theta})$:

$$P_0(\tilde{\theta}) := \mathbb{E} \left[e^{-r\tau(\tilde{\theta})} P(x_{\tau(\tilde{\theta})}, q_{\tau(\tilde{\theta})}) \right].$$

Using (11), this can be written as:

$$\begin{aligned} P_0(\tilde{\theta}) &= \mathbb{E} \left[e^{-r\tau(\tilde{\theta})} x_{\tau(\tilde{\theta})} v_H(\tilde{\theta}) + (1 - x_{\tau(\tilde{\theta})}) v_L(\tilde{\theta}) \right] \\ &\quad - \mathbb{E} \left[\int_{\underline{\theta}}^{\tilde{\theta}} e^{-r\tau(s)} (x_{\tau(s)} v'_H(s) + (1 - x_{\tau(s)}) v'_L(s)) ds \right]. \end{aligned}$$

With this, we get an expression for $U(\theta, \tilde{\theta})$:

$$U(\theta, \tilde{\theta}) = \mathbb{E} \left[e^{-r\tau(\tilde{\theta})} (x_{\tau(\tilde{\theta})} v_H(\theta) + (1 - x_{\tau(\tilde{\theta})}) v_L(\theta)) \right] - P_0(\tilde{\theta}).$$

Its partial derivative with respect to θ is

$$U_1(\theta, \tilde{\theta}) = \mathbb{E} \left[e^{-r\tau(\tilde{\theta})} (x_{\tau(\tilde{\theta})} v'_H(\theta) + (1 - x_{\tau(\tilde{\theta})}) v'_L(\theta)) \right]. \quad (31)$$

The key property that we want to prove is that $U_1(\theta, \tilde{\theta})$ is increasing in $\tilde{\theta}$. To do that, note that applying the law of iterated expectations, we can write $U_1(\theta, \tilde{\theta})$ in terms of the initial belief x_0 as

$$\begin{aligned} U_1(\theta, \tilde{\theta}) &= x_0 \mathbb{E} \left(e^{-r\tau(\tilde{\theta})} v'_H(\theta) \mid \omega = H \right) + (1 - x_0) \mathbb{E} \left(e^{-r\tau(\tilde{\theta})} v'_L(\theta) \mid \omega = L \right) \\ &= x_0 v'_H(\theta) \mathbb{E} \left(e^{-r\tau(\tilde{\theta})} \mid \omega = H \right) + (1 - x_0) v'_L(\theta) \mathbb{E} \left(e^{-r\tau(\tilde{\theta})} \mid \omega = L \right). \end{aligned}$$

Since $v'_H(\theta) \geq 0$ and $v'_L(\theta) \geq 0$ (with at least one of the inequalities strict), both terms in the above expression are positive. Type $\tilde{\theta}$ enters the expression only through the discounting terms $\mathbb{E} \left(e^{-r\tau(\tilde{\theta})} \mid \omega \right)$. Since $\theta'' > \theta'$ implies that $\tau(\theta'') := \inf \{t : q_t \geq 1 - F(\theta'')\} \leq \tau(\theta') := \inf \{t : q_t \geq 1 - F(\theta')\}$ with probability 1, it follows that $\mathbb{E} \left(e^{-r\tau(\tilde{\theta})} \mid \omega \right)$ is increasing in $\tilde{\theta}$ irrespective of state ω , and hence $U_1(\theta, \tilde{\theta})$ is increasing in $\tilde{\theta}$ as well.

We now utilize this property to complete the proof. Note first that we have:

$$\begin{aligned} U(\theta, \theta) &= \mathbb{E} \left[e^{-r\tau(\theta)} (x_{\tau(\theta)} v_H(\theta) + (1 - x_{\tau(\theta)}) v_L(\theta)) \right] - P_0(\theta) \\ &= \mathbb{E} \left[\int_{\underline{\theta}}^{\theta} e^{-r\tau(s)} (x_{\tau(s)} v'_H(s) + (1 - x_{\tau(s)}) v'_L(s)) ds \right]. \end{aligned}$$

Therefore, for arbitrary θ' and θ'' , we have

$$U(\theta'', \theta'') - U(\theta', \theta') = \mathbb{E} \left[\int_{\theta'}^{\theta''} e^{-r\tau(s)} \left(x_{\tau(s)} v'_H(s) + (1 - x_{\tau(s)}) v'_L(s) \right) ds \right]. \quad (32)$$

We can now write:

$$\begin{aligned} U(\theta, \tilde{\theta}) &= U(\tilde{\theta}, \tilde{\theta}) + \int_{\tilde{\theta}}^{\theta} U_1(s, \tilde{\theta}) ds \\ &\leq U(\tilde{\theta}, \tilde{\theta}) + \int_{\tilde{\theta}}^{\theta} U_1(s, s) ds = \\ &= U(\tilde{\theta}, \tilde{\theta}) + \mathbb{E} \left[\int_{\tilde{\theta}}^{\theta} e^{-r\tau(s)} \left(x_{\tau(s)} v'_H(s) + (1 - x_{\tau(s)}) v'_L(s) \right) ds \right] \\ &= U(\theta, \theta), \end{aligned}$$

where the inequality uses the property that $U_1(\theta, \tilde{\theta})$ is increasing in $\tilde{\theta}$, the second last equality uses (31), and the last equality uses (32).

□

D.2 Proof of Lemma 3

Proof. Here, we show how to derive the virtual valuation representation for the designer's value. Suppose transfers follow an arbitrary policy, P_τ , adapted to \mathcal{F}_t .

The designer's expected revenue can be written as

$$\begin{aligned} &\int_{\underline{\theta}}^{\bar{\theta}} \mathbb{E} \left[e^{-r\tau(\theta)} \left(x_{\tau(\theta)} v_H(\theta) + (1 - x_{\tau(\theta)}) v_L(\theta) - W(\theta, x_{\tau(\theta)}) \right) \right] f(\theta) d\theta \\ &= \int_{\underline{\theta}}^{\bar{\theta}} \left(\mathbb{E} \left[e^{-r\tau(\theta)} \left(x_{\tau(\theta)} v_H(\theta) + (1 - x_{\tau(\theta)}) v_L(\theta) \right) \right] \right) f(\theta) d\theta \\ &\quad - \int_{\underline{\theta}}^{\bar{\theta}} \mathbb{E} \left[\int_{\underline{\theta}}^{\theta(q)} e^{-r\tau(s)} \left(x_{\tau(s)} v'_H(s) + (1 - x_{\tau(s)}) v'_L(s) \right) ds \right] f(\theta) d\theta, \end{aligned}$$

where we have used the envelope theorem for the agent's value $W(\theta, x)$. We can use Fubini's theorem to change the order of integration in the second term:

$$\begin{aligned} &\int_{\underline{\theta}}^{\bar{\theta}} \mathbb{E} \left[\int_{\underline{\theta}}^{\theta(q)} e^{-r\tau(s)} \left(x_{\tau(s)} v'_H(s) + (1 - x_{\tau(s)}) v'_L(s) \right) ds \right] f(\theta) d\theta \\ &= \mathbb{E} \left[\int_{\underline{\theta}}^{\bar{\theta}} \int_s^{\bar{\theta}} e^{-r\tau(s)} \left(x_{\tau(s)} v'_H(s) + (1 - x_{\tau(s)}) v'_L(s) \right) f(\theta) d\theta ds \right] \\ &= \mathbb{E} \left[\int_{\underline{\theta}}^{\bar{\theta}} e^{-r\tau(s)} \left(x_{\tau(s)} v'_H(s) + (1 - x_{\tau(s)}) v'_L(s) \right) (1 - F(s)) ds \right]. \end{aligned}$$

The rest is simply to plug the above expression back into the designer's payoff and to write the integral over quantities rather than types (where we use $1 - F(\theta(q)) = q$). The profit maximizing designer's objective becomes

$$\mathbb{E} \left[\int_0^1 e^{-r\tau(q)} (x_{\tau(q)} \phi_H(q) + (1 - x_{\tau(q)}) \phi_L(q)) dq \right],$$

where $\tau(q)$ is the stopping time of the q highest type buyer and $\phi(q)$ is his virtual valuation:

$$\phi_\omega(q) := v_\omega(\theta(q)) - v'_\omega(\theta(q)) \frac{1 - F(\theta(q))}{f(\theta(q))}.$$

□