Slow Social Learning: Innovation Adoption under Network Externalities

work in progress

Click here for the latest version

Tuomas Laiho, Pauli Murto, Julia Salmi*

September 4, 2023

Abstract

We analyze the consequences of payoff externalities in social learning. The leading example is innovation adoption under network effects. When choosing whether to adopt a new innovation, an individual agent faces uncertainty about the quality of the innovation and about its popularity in the future. When the experiences of early adopters work as signals about the quality, a combination of informational and payoff externalities arises. We show that this combination has a surprising effect when the payoff externality is positive: a higher potential for social learning, e.g. due to high media coverage, slows down learning in equilibrium.

JEL classification: C61, C73, D82, D83

Keywords: social learning, experimentation, optimal stopping, network effects, innovation adoption, industry equilibrium

*Laiho: Aalto University and Finnish Ministry of Finance (tuomas.s.laiho@gmail.com), Murto: Aalto University School of Business (pauli.murto@aalto.fi), Salmi: University of Copenhagen (julia.salmi@econ.ku.dk). We thank Yeon-Koo Che, Francesc Dilme, Alex Frug, Yingni Guo, Daniel Hauser, Jan Knoepfle, Pau Milán, Peter Norman Sørensen, Juuso Välimäki, and the participants in the following seminars, conferences, and workshops for their helpful comments: BU 2021, UT Austin 2021, VSET 2022, Incentives and Dynamic Games Workshop in Paris 2022, Pompeu Fabra 2023, CBS Microeconomics Workshop 2023, Essex Workshop on the Economics of Dynamic Incentives 2023, UB Microeconomics 2023, and ESSET 2023. This project has received financial support from the Academy of Finland (Murto).

1 Introduction

Adopting a new innovation is often more valuable if others start using it too. Network externalities are significant in many industries, including computer software, social media, commercial platforms, and electric vehicles. It is also widely accepted that the green transition requires coordinated actions as novel green technologies cannot operate without large investments in the infrastructure. Another characteristic feature of new innovations is that there is common uncertainty about their quality. Hence, by trying out a new product, early adopters cause an informational externality on those still considering whether to start using it. How does the concurrent existence of informational and payoff externalities affect the adoption patterns of new innovations and experimentation more broadly? How do these two externalities interact?

We approach these questions from the perspective of *learning potential.* Learning potential is high when adopters generate a lot of public information, i.e. when past experience predicts future performance well and when communication frictions are not too severe. To fix ideas, one can think of learning potential as the expected media coverage: if a new innovation attracts a lot of media attention, potential adopters are likely to hear about the drawbacks and successes of early adopters, and hence the potential of social learning is high. Our main result shows that whether learning increases or decreases as a consequence of an improved learning potential is determined by the payoff externality: the equilibrium learning is slower under positive payoff externalities, faster under negative payoff externalities, and is independent of the learning potential under no payoff externalities.

To analyze the effect of payoff externalities on experimentation, we build on the classic Brownian bandit model in Bolton and Harris (1999) where players are initially uncertain whether the state-of-the-world is high or low and receive normally distributed signals about it based on the aggregate level of experimentation. However, we make two important changes: 1) players choose when to irreversibly stop (e.g. to adopt an innovation), and 2) the payoff after stopping depends on the fraction of how many other players have stopped. Payoff externalities are positive (negative) when the payoff is increasing (decreasing) in the fraction of adopters. As actions are irreversible, also information generation is persistent. Furthermore, we focus on experimentation in large games where each player is infinitesimally small.

An individual player solves an optimal stopping problem with two state variables: the stock of stopped agents and the belief that the state is high. We show that both the equilibrium with the fastest and the equilibrium with the slowest adoption can be characterized by a cutoff belief that depends on the current stock of stopped agents. We call such equilibria the maximal equilibrium and the minimal equilibrium, respectively.

The first step to understanding why higher learning potential may decrease equilibrium learning is to consider an individual player's problem when keeping the behavior of everyone else constant. Clearly, an increase in learning potential decreases the individual's willingness to adopt due to the increased value of waiting. This endogenous response drives a wedge between equilibrium learning and learning potential, irrespective of payoff externalities. However, it is surprising that this effect is so strong that equilibrium learning may actually slow down, despite the direct effect of a higher learning potential that speeds up learning.

The equilibrium rate of learning adjusts to the level that makes players indifferent between stopping and waiting. Without payoff externalities, the rate must hence be independent of the learning potential. The innovation adoption slows down under a better learning technology just so much that the rate of learning stays at the same level.

We show that under positive payoff externalities, a higher learning potential leads to strictly slower learning and lower welfare in equilibrium. To get the intuition, start again with the innovation pattern that keeps the rate of learning constant when the learning potential improves. However, now slower innovation adoption implies lower stopping payoffs as the early adopters get less of the positive payoff externality. This effect makes players even less eager to adopt, and hence learning slows down even further. As both innovation adoption and learning slow down, the expected welfare decreases to all players. The comparison holds between the sets of equilibria and is strict under a large class of parameter values.

The result has important implications for detecting bottlenecks in innovation adoption. The adoption is slow exactly when the social value of fast adoption is the highest, i.e. when early adopters generate a lot of information and when there is a large positive payoff externality among the adopters. This suggests that the need to subsidize the adoption of new innovations with network effects is especially pronounced when there is also social learning. The payoff externality strengthens the informational free-riding caused by the informational externality.

If payoff externalities are negative instead, e.g. if the innovation adopters are firms operating in the same market, better learning technology always leads to faster learning and higher welfare in equilibrium. In that case, the equilibrium innovation adoption does not need to slow down as much as under positive externalities because the stopping payoffs increase as the innovation adoption slows down, creating a counteracting force for informational free-riding.

The key step in showing that the above intuition holds in our dynamic environment with forward-looking players is to consider a simplified problem where players ignore future changes in the stock. We prove an equivalence result between the simplified problem and the optimal stopping in an equilibrium. This observation enables solving the maximal and minimal equilibria in closed form. We show the equivalence by using the iterative elimination of strictly dominated strategies. The elimination process rules out innovation adoption that is optimal only if other players behave suboptimally.

Our equilibrium characterization allows for any form of payoff externality, including externalities that are non-monotonic. For example, there can be a saturation point so that the externality is positive until the stock reaches a certain level and is negative thereafter.

Related literature

The paper is related to the literature on social learning and experimentation, see Bolton and Harris (1999), Keller, Rady and Cripps (2005), and Keller and Rady (2015). The main differences are that actions are irreversible, and hence information generation is persistent, and that the players' payoffs depend on the stopping decisions of other players. The former feature is also present in our earlier paper Laiho, Murto and Salmi (2023) but without payoff externalities.

Frick and Ishii (2020) studies experimentation by small players under Poisson learning. In both Frick and Ishii (2020) and Laiho, Murto and Salmi (2023), informational free-riding mitigates the gains of better learning technology. However, because there are no payoff externalities, the larger potential of learning does not lead to slower learning or lower welfare. Frick and Ishii (2020) provide an extension where a better learning technology may lead to lower aggregate welfare if the opportunities to adopt the innovation arrive infrequently and there is heterogeneity in the players' patience. In the present paper, players are homogeneous and there are no frictions in innovation adoption.

There is a large literature on network effects in innovation adoption, for early papers see Katz and Shapiro (1986), Jovanovic and Lach (1989), and Farrell and Saloner (1986). We contribute to this literature by analyzing the effects of payoff externalities together with endogenous learning.

The combination of informational and payoff externalities has been studied in the context of war of attrition games: see Décamps and Mariotti (2004) and Kwon, Xu, Agrawal and Muthulingam (2016) for duopoly markets under private costs of investment. Margaria (2020) studies a related setup without private costs but with private information about the profitability of the investment. Different strategic concerns arise in duopoly models of Bergemann and Välimäki (1997) and (2000) where learning about the product quality affects how close substitutes the products are and hence how fierce the competition between the firms is. See also Strulovici (2010), Halac, Kartik and Liu (2017), and Thomas (2021) for other (less related) environments with joint informational and payoff externalities.

The way we characterize the equilibria is close to Leahy (1993) and Baldursson and Karatzas (1996) that use an equivalence result to solve a competitive industry equilibrium under fluctuating demand. We also use a similar approach in our earlier experimentation paper Laiho, Murto and Salmi (2023) without payoff externalities. The proof techniques in the above papers do not extend to the environment of the present paper. Therefore, the present paper uses a different approach to rule out equilibria where the equivalence between the simplified problem and equilibrium optimization would fail by applying the iterative elimination of strictly dominated strategies.

The distinguishing feature between the present paper and any of the papers listed above is that we focus on the combination of informational and payoff externalities in a large game between symmetrically informed homogeneous players. Importantly, the players are truly forward-looking and their payoffs change over time as a consequence of other players' actions.

2 Motivating example

Here, we present our main result in the simplest possible setup: a two-period model with linear (affine) payoffs.

Consider a game between a unit mass of identical players in two periods, $t \in \{1, 2\}$. In each period, the players who have not stopped before, choose simultaneously whether to stop in the current period. All players who have stopped receive a per-period payoff $\pi_\omega(q_t)$ where $q_t \in [0, 1]$ is the fraction of players who have stopped in period t or earlier and $\omega \in \{H, L\}$ is an unknown state-of-the world. We assume that $\pi_H(q) = a + bq$ and $\pi_L = -c$ with $a, b, c > 0$. A player who has not stopped gets a payoff 0. These assumptions imply that there is a positive payoff externality among players who have stopped. After the first period, a fully revealing signal arrives with probability λq_1 where $\lambda \in [0, 1]$ is a parameter that captures the learning potential in the game. The players maximize the sum of the expected payoffs over the two periods. Suppose that the initial belief that the state is high, x_1, is high enough that it would be optimal to stop in the first period regardless of what other players do if there was no learning: $x_1 a > (1 - x_1)c$.

We can solve the equilibrium backward. In the second period, all remaining players stop both absent news and if a signal revealing that the state is high (good news) arrives. Suppose that fraction q_1 of all players stops in the first period. Then, it is optimal for an individual player to stop in the first period if

$$\begin{aligned}&(x_1(a + bq_1) - (1 - x_1)c) + (1 - \lambda q_1)(x_1(a + b) - (1 - x_1)c) + \lambda q_1 x_1(a + b) \\ &- \lambda q_1(1 - x_1)c \geq (1 - \lambda q_1)(x_1(a + b) - (1 - x_1)c) + \lambda q_1 x_1(a + b),\end{aligned}$$

where the first term is the expected payoff in the first period and the remaining terms on the left-hand side are the second-period payoffs absent news, after good news, and after bad news, correspondingly. The right-hand side is the player's expected payoff if he does not stop in the first period and then stops in the second period if there is no bad news. The optimality condition simplifies to

$$x_1(a + bq_1) - (1 - x_1)c \geq \lambda q_1(1 - x_1)c. \tag{1}$$

Notice that large q_1 makes stopping in the first period more appealing if $x_1 b > (1 - x_1)\lambda c$. Then there is an equilibrium where everyone stops in the first period. This is true when learning potential is low, i.e. when λ is small. If learning potential is high instead, there exists a unique $\overline{q}_1$ such that all players are indifferent between stopping and waiting in the first period when the stock $\overline{q}_1$ of other players is expected to stop. Then, the unique equilibrium is such that exactly mass $\overline{q}_1$ stops in the first period. We get the stock $\overline{q}_1$ from (1):

$$\overline{q}_1 = \frac{x_1 a - (1 - x_1)c}{\lambda(1 - x_1)c - x_1 b}. \tag{2}$$

Let us now look at what happens to the probability of learning the true state, $\lambda\overline{q}_1 = \frac{x_1 a - (1 - x_1)c}{(1 - x_1)c - \lambda^{-1} x_1 b}$, when learning potential improves. In principle, there are two forces going to opposite directions: the direct effect of high λ is positive, but there is also a negative effect through the decrease in $\overline{q}_1$. We see from (2) that the latter effect dominates whenever $\overline{q}_1 < 1$ that is if $x_1 b < (1 - x_1)\lambda c$. This means that a lower learning potential paradoxically leads to more learning in the equilibrium up to the point when the learning potential is so low that everyone wants to stop in the first period anyway.

Notice that the argument above hinges on positive payoff externalities, $b > 0$. If there are no payoff externalilities ($b = 0$), the probability of receiving news is independent of the learning potential λ. Under negative externalities ($b < 0$), high learning potential increases learning in equilibrium. We see in Section 5 that these findings generalize to our full model.

3 Model

3.1 Actions and payoffs

Time is continuous and goes to infinity, $t \in [0, \infty)$. A unit mass of identical players chooses when to stop. When a player stops, he starts receiving a flow payoff $\pi_\omega(q_t)$ where $q_t \in [0, 1]$ is the fraction of players who have stopped by time t and $\omega \in \{H, L\}$ is an unknown state-of-the-world that can be either high or low. Players' payoffs are normalized to 0 before stopping, and they discount with rate r. Throughout, we assume that $\pi_H(q) > 0$ and that $\pi_L(q) < 0$ for all q, which implies that all players want to stop if the state is known to be high and no one wants to stop if the state is known to be low. We further assume that π_ω are absolutely continuous in q.[1]

3.2 Learning

Players observe the following public signal process:

$$dY_t = \lambda(q_t)\mu_\omega dt + dW_t, \tag{3}$$

where $\mu_H = 0.5$, $\mu_L = -0.5$, and dW_t is a standard Wiener process. The signal-to-noise ratio of the signal process (3) is $\lambda(q_t)$, and hence the function λ captures *learning potential* (or learning technology) in the game.

[1]Notice that a lump-sum cost of stopping can also be incorporated into the model: in addition to the flow payoffs above, let there be lump-sum cost C that each player must pay at the time of stopping. Then, we can find flow payoffs $\hat{\pi}_H(q) = \pi_H(q) - rC$ and $\hat{\pi}_L(q) = \pi_L(q) - rC$ that give exactly the same discounted stopping payoff as the combination of the flow payoff and the lump-sum cost.

Let the initial belief that the state-of-the-world is high be $x_0 \in (0,1)$. The players update their belief based on (3). We denote the realized public belief at time t by x_t. Conditional on the information available at time t, the change in the public belief, dX_t, is distributed normally with mean 0 and variance $(\lambda(q_t)x_t(1-x_t))^2dt$.

Throughout the paper, we assume that the signal-to-noise ratio is differentiable and increasing in the stock and that there is no exogenous source of information:

Assumption 1. $\lambda'(q) > 0$ *and* $\lambda(0) = 0$.

Our main interest is to analyze how higher learning potential affects equilibrium behavior under payoff externalities. We use the following definition to (partially) rank learning technologies:

Definition 1. *λ has (strictly) higher learning potential than $\hat{\lambda}$ if $\lambda(q) > \hat{\lambda}(q)$ for all $q > 0$.*

We provide a micro foundation for Process (3) as aggregate realized payoffs of the stopped agents in Appendix A.1. Under that interpretation, high learning potential means little noise in realized payoffs, conditional on the true state ω, together with good observability of other players' payoffs.

3.3 Equilibrium

Suppose that the stock of stopped agents follows an increasing process Q_t adapted to the natural filtration generated by Y_t in Equation (3). State (q,x) is then a sufficient history for players to choose the optimal stopping time. When fixing the stock process Q_t, an individual's problem is:

$$v(q_t,x_t) = \sup_{\tau} \mathbb{E}_{Q_t}\left[\int_{\tau}^{\infty} e^{-r(s-t)}\Big(x_s\pi_H(q_s) + (1-x_s)\pi_L(q_s)\Big)ds|(q_t,x_t)\right]. \quad (4)$$

We define the equilibrium as a stock process that is consistent with individual optimization:

Definition 2. *An increasing process Q_t adapted to the natural filtration generated by Y_t is an equilibrium if* *(i)* $v(q_t, x_t) = \mathbb{E}_{Q_t}\left[\int_t^\infty e^{-r(s-t)}\Big(x_s\pi_H(q_s) + (1-x_s)\pi_L(q_s)\Big)ds|(q_t, x_t)\right]$ *whenever* $dQ_t > 0$,

(ii) $v(q_t, x_t) \geq \mathbb{E}_{Q_t}\left[\int_t^\infty e^{-r(s-t)}\Big(x_s\pi_H(q_s) + (1-x_s)\pi_L(q_s)\Big)ds|(q_t, x_t)\right]$ *whenever* $dQ_t = 0$.

4 Benchmark without payoff externalities

In this section, we assume that the stopping payoffs are independent of the stock: $\pi_\omega(q) \equiv \pi_\omega$. Instead of solving for the equilibrium, we focus on the implications of learning potential. In Section 4.1, we show that the speed of equilibrium learning is essentially independent of the learning potential. In Section 4.2, we discuss informally how payoff externalities change this result and the additional steps needed in order to incorporate payoff externalities into the formal analysis.

4.1 Equilibrium learning is independent of learning potential

Here, we show that the equilibrium rate of learning is independent of the learning technology when there are no payoff externalities. More precisely, we argue that if Q_t is an equilibrium under learning potential λ, then stock process $\hat{Q}_t$ that yields exactly the same flow of information under learning potential $\hat{\lambda}$ is an equilibrium under that learning potential. Qualitatively, this observation is shared with some earlier papers using different models for experimentation, e.g. Frick and Ishii (2020). For us, the main purpose of analyzing the case without payoff externalities first is that it helps to build intuition for our main result with payoff externalities.

We seek to compare the equilibria under two different learning technologies, λ and $\hat{\lambda}$ such that λ has a strictly higher learning potential. It is useful to define *information equivalent stock* $z(q; \lambda, \hat{\lambda})$, which is the stock that produces the same amount of information under λ as stock q produces under $\hat{\lambda}$. Formally, $z(q; \lambda, \hat{\lambda}) := \hat{\lambda}^{-1}(\lambda(q))$.

Suppose that the stock process Q_t is an equilibrium under the higher learning potential λ. We suggest that then process $\hat{Q}_t$, defined as $\hat{q}_t = z(q_t; \lambda, \hat{\lambda})$ if $z(q_t; \lambda, \hat{\lambda}) < 1$ and $\hat{q}_t = 1$ otherwise, is an equilibrium under lower learning potential $\hat{\lambda}$. To see this, notice that the two processes imply an identical belief process X_t for all t such that $\hat{q}_t < 1$. As players' optimal actions depend on other players' behavior only through the effect on the belief dynamics, also the optimal stopping rules coincide when they face the same flow of information. Appendix A.2 provides a formal argument along these lines so that we obtain:

Proposition 1. *Let Q_t be an equilibrium under learning potential λ. Then, there exists an equilibrium $\hat{Q}_t$ under a lower learning potential $\hat{\lambda}$ such that $\hat{q}_t = \min\{z(q_t; \lambda, \hat{\lambda}), 1\}$ for all t.*

Proposition 1 implies that for all realizations of the Wiener process W_t in (3), the belief is exactly the same in equilibrium Q_t under λ as it is in equilibrium $\hat{Q}_t$ under $\hat{\lambda}$ as long as some players are still waiting in equilibrium under $\hat{\lambda}$. As q_t is still below 1 when (and if) $\hat{q}_t$ reaches 1, the two belief processes diverge after that point. However, that happens only when there are no decisions left to any player under $\hat{\lambda}$, and therefore we say that learning is essentially independent of learning potential. The expected welfare of all players is the same.

4.2 Payoff externalities: discussion

In the previous subsection, we used a guess-and-verify approach: we guessed that the information equivalent stock process $\hat{Q}_t$ is an equilibrium under the lower learning potential $\hat{\lambda}$ and verified individual optimality, without deriving the equilibrium properties. We cannot use the same method with payoff externalities because we lack an immediate guess for an equilibrium. Why does not the same guess work with payoff externalities? The easiest way to see this is to consider the case where one tries to construct an equilibrium under a higher learning potential based on an equilibrium under a lower potential when payoff externalities are positive.

Assume strictly positive payoff externalities such that $\pi'_\omega(q) > 0$ for both $\omega \in \{H, L\}$. Suppose that stock process Q_t is an equilibrium under learning potential

λ, and define stock process $\tilde{Q}_t$ such that the stock is information equivalent under a *higher* potential $\tilde{\lambda}$: $\tilde{q}_t := z(q_t; \lambda, \hat{\tilde{\lambda}})$ for all t such that $q_t < 1$. Let (x_t, q_t) be such that some players stop, i.e. $dQ_t > 0$. This implies that in order for $\tilde{Q}_t$ be an equilibrium, some players must be willing to stop at $(x_t, \tilde{q}_t)$. Recall that the players face the same process for information in $\tilde{Q}_t$ under $\tilde{\lambda}$ as they do in Q_t under λ. However, the expected stopping payoffs are not the same because $\tilde{q}_t < q_t$ and hence also $x_t\pi_H(\tilde{q}_t) + (1-x_t)\pi_L(\tilde{q}_t) < x_t\pi_H(q_t) + (1-x_t)\pi_L(q_t)$. Because the stock evolves slower in $\tilde{Q}_t$ than in Q_t, the positive payoff externality is lower. Therefore, $\tilde{Q}_t$ cannot be an equilibrium under learning technology $\tilde{\lambda}$ if Q_t is an equilibrium under learning technology λ.

The above reasoning seems to suggest that the actual equilibrium stock should evolve even slower than $\tilde{Q}_t$ because it would be optimal to wait if others behaved according to $\tilde{Q}_t$. We show in the next section that this is indeed true in equilibrium: higher learning potential slows down expansions so much that also learning slows down under positive payoff externalities. However, we cannot show the result directly because one needs to know how the stock is expected to evolve in the future in order to evaluate the expected stopping payoff when there are payoff externalities. Therefore, we first characterize the sets of equilibria (in Proposition 2 and Proposition 3) before proving the counterpart of Proposition 1 with payoff externalities (in Proposition 4).

5 Equilibrium with payoff externalities

5.1 Characterization

Recall that all players would like to stop if the state was known to be high and no one would like to stop if it was known to be low. Therefore, to characterize any equilibrium, it is a reasonable first step to find a stock-dependent cutoff belief such that more players stop whenever the belief is above the cutoff.The interesting question is how the cutoff depends on the stock of stopped agents. The stock affects the profitability of stopping through two channels: it controls the flow of new information and affects the payoff externality.

We proceed to solve for the set of equilibria by first considering a simple problem where the stock of stopped players stays constant over time. Then we show how we can use the solution of the simple problem in equilibrium characterization.

The first step is to solve the optimal stopping problem for fixed q. This is a standard one-dimensional optimal stopping problem and its solution is fully characterized by cutoff $\bar{x}$ (see e.g. Chapter 5 in Dixit and Pindyck (1994) and team problem in Bolton and Harris (1999)):

$$\bar{x}(q) := \frac{-\beta(q)\pi_L(q)}{(\beta(q)-1)\,\pi_H(q) - \beta(q)\pi_L(q)}, \tag{5}$$

where $\beta(q) := \frac{1}{2}\left(1 + \sqrt{1 + \frac{8r}{\lambda(q)}}\right)$. Next, we define

$$\hat{x}(q) := \max\{x \in [0,1] : x \leq \bar{x}(q') \quad \forall q' \geq q\}. \tag{6}$$

The function $\hat{x}$ is the largest monotone function that has values below $\bar{x}$. If $\bar{x}$ is monotone, $\bar{x}$ and $\hat{x}$ coincide.

We show in Appendix A.4 that the equilibrium cutoff in the game where Q_t changes endogenously coincides with the cutoff with fixed stock:[2]

Proposition 2. *In any equilibrium, $dQ_t > 0$ if $x_t > \bar{x}(q)$ and $dQ_t = 0$ if $x_t < \hat{x}(q)$, where the cutoff beliefs are defined in* (5) *and* (6)*.*

When $\bar{x}$ is monotone, Proposition 2 implies that the (q, x) state space can be divided into two regions: there is a stopping region above $\bar{x}(q)$ and a waiting region below it. If the current state (q, x) is in the stopping region, i.e. $x > \bar{x}(q)$, more players stop so that we reach a cutoff state $(q + dQ, \bar{x}(q + dQ))$ where dQ solves $\bar{x}(q + dQ) = x$. If $\bar{x}$ is non-monotone as in Figure 1, there are some states between the curves $\bar{x}$ and $\hat{x}$ and for those Proposition 2 does not imply if the players should stop or wait. This ambiguity is an integral feature of the model with payoff externalities because the profitability of stopping depends on other players' stopping decisions. Hence, multiple equilibria naturally arise. In the

[2] This result extends the characterization of the decentralized equilibrium in our earlier paper Laiho, Murto and Salmi (2023) without payoff externalities and with heterogeneous players.

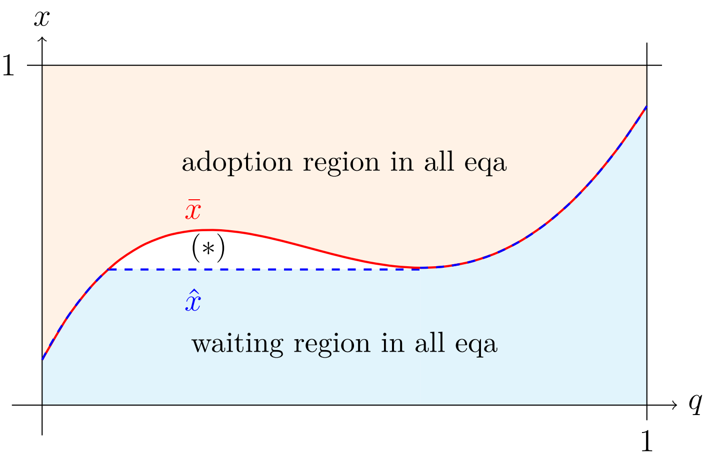

Figure 1: Illustration of Proposition 2: in any equilibrium, players prefer stopping in all states above $\bar{x}(q)$ and prefer waiting in all states below $\hat{x}(q)$. The behavior differs across equilibria in the middle region marked by $(*)$.

main text, we focus on the equilibrium with the fastest and the equilibrium with the slowest expansions, which we define in the next subsection. For completeness, we discuss other equilibria in Appendix A.6.

5.2 Maximal and minimal equilibria

In order to compare the sets of equilibria under different learning potentials, we focus on the equilibria with the fastest and slowest innovation adoption:

Proposition 3.

- *Let stock process* $\overline{Q}_t$ *be characterized by the cutoff rule* $\hat{x}$ *in* (6)*: in state* (q, x)*,* $d\overline{Q} = \max\{q' - q, 0\}$ *where* $q' = \max\{s \in [0,1] : \hat{x}(s) = x\}$*. Then,* $\overline{Q}_t$ *is an equilibrium and we call it the maximal equilibrium.*
- *Let stock process* $\underline{Q}_t$ *be characterized by the cutoff rule* $\bar{x}$ *in* (5)*: in state* (q, x)*,* $d\underline{Q} = \max\{q' - q, 0\}$ *where* $q' = \min\{s \in [q,1] : \bar{x}(s) \geq x\}$*. Then,* $\underline{Q}_t$ *is an equilibrium and we call it the minimal equilibrium.*

Proposition 3 guarantees the existence of an equilibrium by characterizing the maximal and minimal equilibria. In Appendix A.5, we check that no player has a profitable deviation under either $\overline{Q}_t$ or $\underline{Q}_t$, which verifies that they are indeed equilibria. To correctly interpret the maximal and minimal equilibria, recall that a lower cutoff belief implies a larger stock and faster learning. By Proposition 2, no equilibrium can have faster expansion in the stock than the maximal equilibrium. Similarly, no equilibrium can have slower expansion in the stock than the minimal equilibrium.

When it comes to the specific form of the maximal and minimal equilibria, notice that they are both essentially pinned down by $\bar{x}(q)$ in (5). Cutoff $\bar{x}(q)$ coincides with the optimal stopping threshold of an individual agent when the stock stays constant forever. This result means that the players can ignore both faster learning and the payoff consequences of future stopping decisions. Our earlier paper Laiho, Murto and Salmi (2023) finds a similar result but without payoff externalities. The intuition behind the result is that, in an equilibrium, the stock expands only when the player considering stopping now would be willing to stop, too. Hence, the player would never strictly prefer waiting until the stock has actually increased. Importantly, the equivalence between optimal stopping under constant and increasing stock processes is an equilibrium property. One can always find a non-equilibrium stock process such that a player's best response depends on future expansions and hence violates the equivalence.

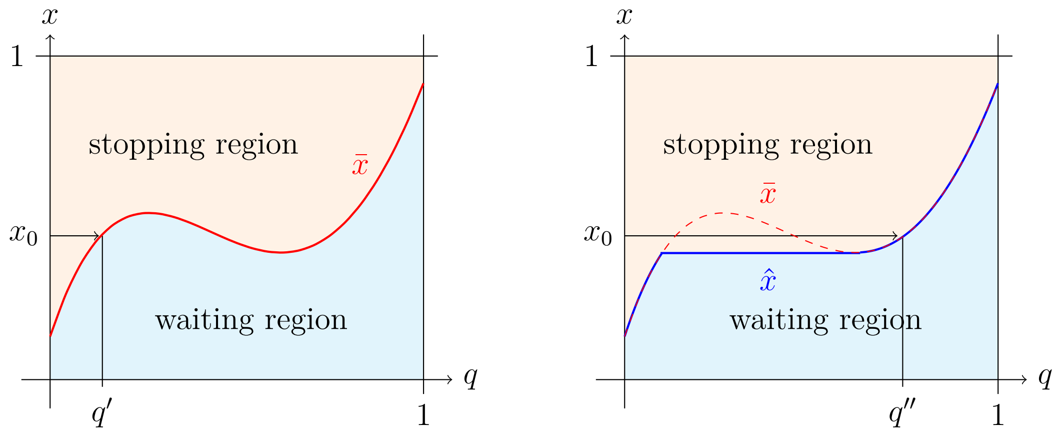

Figure 2: The left panel shows the minimal and the right panel the maximal equilibrium. With the same initial belief x_0, the equilibria differ in the mass of players who stop immediately: the stock jumps to q' in the minimal and to q'' in the maximal equilibrium.

Notice that $\bar{x}$ in (5) may be everywhere increasing, and in that case, the maximal and the minimal equilibria coincide. By Proposition 2, the equilibrium is then unique. One can verify from Formula (5) that $\bar{x}$ is necessarily increasing when there are no payoff externalities or when the externalities are negative. Hence, the multiplicity of equilibria arises only under positive payoff externalities. This is a natural feature as only positive payoff externalities create a coordination problem in the game. Notice that a coordination problem arises even without learning: suppose that $\lambda(q) = 0$ for all q and that the belief is such that $x_0\pi_H(0) + (1 - x_0)\pi_L(0) < 0$ but $x_0\pi_H(1) + (1 - x_0)\pi_L(1) > 0$. Then, there is an equilibrium where no one ever adopts the innovation and another equilibrium where everyone adopts the innovation immediately.

5.3 Main result

We now show that improvements in the learning potential may backfire: under positive payoff externalities, higher learning potential slows down learning and reduces total welfare.

We seek to compare the set of equilibria under different learning potentials. To make the comparison precise, we consider the maximal and minimal equilibria under different learning technologies. Let $\overline{q}_{t,\lambda}$ and $\underline{q}_{t,\lambda}$ denote a realization of the stock in the maximal and minimal equilibria under learning potential λ. We use $q_{t,\lambda}$ for a generic equilibrium stock under λ. Similarly, we use $\bar{x}_\lambda(q)$ and $\hat{x}_\lambda(q)$ for the cutoff beliefs (5) and (6) under learning potential λ.

As an intermediate step toward our main result that shows that learning potential may decrease the expected welfare of all players, we study how learning potential affects equilibrium learning. Recall *information equivalent stock* from Section 4, $z(q; \lambda, \hat{\lambda}) = \hat{\lambda}^{-1}(\lambda(q))$, which yields exactly the same amount of information under $\hat{\lambda}$ as stock q does under λ. The following comparison follows directly from the functional forms of (5) and (6):

Corollary 1. *Let $\hat{\lambda}$ has strictly higher learning potential than λ.*

(i) Suppose strictly positive payoff externalities. Then, $\hat{x}_{\hat{\lambda}}(z(q; \lambda, \hat{\lambda})) > \hat{x}_\lambda(q)$ and $\bar{x}_{\hat{\lambda}}(z(q; \lambda, \hat{\lambda})) > \bar{x}_\lambda(q)$ for all $q \in (0, 1]$.

(ii) Suppose strictly negative payoff externalities. Then, $\hat{x}_{\hat{\lambda}}(z(q; \lambda, \hat{\lambda})) = \bar{x}_{\hat{\lambda}}(z(q; \lambda, \hat{\lambda})) < \bar{x}_\lambda(q) = \hat{x}_\lambda(q)$ for all $q \in (0, 1]$.

Corollary 1 is the key to why we may have faster learning under lower learning potential. To correctly interpret it, recall that if $q_{t,\hat{\lambda}} < z(q_{t,\lambda}; \lambda, \hat{\lambda})$, the equilibrium under learning technology $\hat{\lambda}$ yields less information at time t than the equilibrium under learning technology λ. Corollary 1 compares the beliefs needed for new players to be willing to stop under different learning potentials when keeping the flows of information the same, i.e. when the stock is q under λ and $z(q; \lambda, \hat{\lambda})$ under $\hat{\lambda}$. This comparison shows how the comparable information stocks evolve across different learning potentials: if a lower belief is everywhere needed for the stock to expand, learning is necessarily faster. For completeness, we provide an exact comparison for the equilibrium belief processes in Appendix A.7.

Part (ii) of Corollary 1 implies that the comparison between equilibria under different learning potentials is as expected under negative payoff externalities: higher learning potential leads to a larger flow of information in equilibrium.

Part (i) of Corollary 1 implies that under positive payoff externalities, learning is paradoxically slower under better learning technology. The comparison holds in the set order: information stock evolves strictly slower in all equilibria under the better learning technology than in the maximal equilibrium under the worse technology, and it evolves faster in all equilibria under the worse learning technology than in the minimal equilibrium under the better learning technology. Next, we show that this logic implies that high learning potential decreases the welfare of all players under positive payoff externalities.

To make the welfare comparison precise, we use $v_{\lambda}^{\max}(x_0)$ to denote the expected ex-ante payoff of the players in the maximal equilibrium with a given learning potential λ, and similarly $v_{\lambda}^{\min}(x_0)$ for the minimal equilibrium. Our main result shows that the learning potential influences the equilibrium payoffs differently under positive and negative payoff externalities:

Proposition 4. *Suppose λ has strictly higher learning potential than λ' and let the initial belief x_0 satisfy $x_0 \pi_H(0) + (1 - x_0)\pi_L(0) > 0$. Then:*

(i) With strictly positive payoff externalities, $v_{\lambda'}^{\max}(x_0) > v_{\lambda}^{\max}(x_0)$ and $v_{\lambda'}^{\min}(x_0) > v_{\lambda}^{\min}(x_0)$.

(ii) With strictly negative payoff externalities, $v_{\lambda}^{\max}(x_0) = v_{\lambda}^{\min}(x_0) > v_{\lambda'}^{\max}(x_0) = v_{\lambda'}^{\min}(x_0)$.

The intuition behind the first part of Proposition 4 is that players face both a slower increase in the stock and a slower arrival of information under higher learning potential when there are positive payoff externalities. In Appendix A.8, we show the result formally by showing that the expected time until the maximal equilibrium reaches belief $\hat{x}_{\lambda'}(1)$ under λ' is shorter than the time until the maximal equilibrium reaches belief $\hat{x}_{\lambda'}(z(1;\lambda',\lambda))$ under λ' (in expectation). Then, because the stopping payoff is lower when the stock is $z(1;\lambda',\lambda)$ than when it is 1 under positive payoff externalities, the value must be lower, too. The equivalent argument applies to the comparison between the minimal equilibria. To show that the converse is true under negative payoff externalities, we argue that the equilibrium under λ reaches belief $x_{\lambda'}(1)$ sooner than the equilibrium under λ' (in expectation). The value at that belief must also be higher under λ because the stock is smaller and payoff externalities are negative. In both cases, our argument utilizes the fact that the expected payoff of all players is the same as the expected payoff of the last player.

Proposition 4 has important implications in technology adoption under positive network effects. When potential adopters can learn from each other, the option value of waiting for more information is strong and gets stronger if innovation adoption is fast. Therefore, no one wants to be the first to adopt the innovation, which makes the innovation adoption even less lucrative under positive payoff externalities. Our results suggest that subsidizing early adoption may be vital for overall efficiency if the potential for learning is high. Otherwise, the equilibrium may get stuck to very slow expansion of the new technology as adopters need a higher belief to be willing to adopt and the low fraction of early adopters makes the belief move slowly.

6 Concluding remarks

The fundamental problem in social learning is that the players do not internalize the informational externality. This affects the effectiveness of various policies trying to facilitate information generation and better informed decision making. This paper points out that technical solutions to facilitate the aggregation and diffusion of information may even backfire if there are positive network effects. Therefore, having obstacles in information transmission may actually improve welfare. An indirect implication of our analysis is that small subsidies for the use of new technologies may increase the total welfare significantly even if they are temporary. When network effects are not important or when they are negative (e.g. in the case of products and services with congestion), the backfiring effect of improvements in the learning technology does not appear.

A Appendix

A.1 Micro foundation for learning process (3)

We provide a micro foundation for Process (3) as coming from the realized payoffs of the stopped agents:

Start with N players, and let n_t be the number of players who have stopped by time t. When a player stops, he starts receiving a flow payoff

$$du_t^i = \frac{\pi_\omega(q_t)dt}{N} + \frac{\sigma}{N}dW_t^i \tag{7}$$

that depends on the fraction of other agents who have stopped, $q_t = n_t/N$, and an unknown state-of-the-world that can be either high or low, $\omega \in \{H, L\}$. dW_t^i are independently and normally distributed with mean 0 and variance dt and capture noise in payoffs. Furthermore, assume that other agents observe a noisy signal of the realized payoff, $d\hat{u}_t^i$:

$$d\hat{u}_t^i = du_t^i + \frac{\sigma'}{N}dW'^i_t, \tag{8}$$

where dW'^i_t are independently and normally distributed with mean 0 and variance dt and capture noise in communication.

The total informativeness of the experiment coincides with observing the sum of flow payoffs of the agents who have stopped:

$$dU_t = \sum_{i=1}^{n_t} d\hat{u}_t^i = \sum_{i=1}^{n_t}\left(\frac{\pi(q_t,\omega)}{N} + \sqrt{\frac{n_t}{N}}\sigma dW_t^i + \sqrt{\frac{n_t}{N}}\sigma' dW'^i_t\right). \tag{9}$$

We take the limit as $N \to \infty$. Then, by using that $q_t = \frac{n_t}{N}$, dU_t is normally distributed with unknown mean $q_t\pi(q_t,\omega)dt$ and variance $q_t(\sigma^2 + \sigma'^2)dt$. Signal dU_t is a special case of (3) with $\lambda(q_t) := \frac{\sqrt{q_t}(\pi_H(q_t)-\pi_L(q_t))}{\sqrt{\sigma^2+\sigma'^2}}$.

Here, parameters σ and σ' determine the learning potential. The interpretation of σ is that low σ means that past performance is predictive of future performance; the interpretation of σ' is that low σ' means little friction in communication.

A.2 Proof of Proposition 1

Proof. Suppose everyone follows $\hat{Q}_t$ under $\hat{\lambda}$. Suppose that an individual has a strictly profitable deviation, which means that either an individual has a strict preference for waiting when $d\hat{Q}_t > 0$ or that he has a strict preference for stopping when $d\hat{Q}_t = 0$. Consider the first case: an individual has a strict preference for waiting in some state $(z(q_t, \lambda, \hat{\lambda}), x_t)$ such that $dQ_t > 0$ in state (q_t, x_t). This contradicts that Q_t under λ because the only difference between an individual's problem in state $(z(q_t, \lambda, \hat{\lambda}), x_t)$ under $(\hat{Q}_t, \hat{\lambda})$ and in state (q_t, x_t) under (Q_t, λ) is that there is more information available under the latter scheme after $\hat{Q}_t$ reaches 1.

Consider the second case: an individual has a strict preference for stopping in some state $(z(q_t, \lambda, \hat{\lambda}), x_t)$ such that $dQ_t = 0$ in state (q_t, x_t). Again, this contradicts the optimality to wait in state (q_t, x_t) under (Q_t, λ). To see this notice there must exist x'_t such that the individual's expected payoff is weakly greater if she waits for state (q_t, x') rather than stops immediately at state (q_t, x_t) under (Q_t, λ) (otherwise Q_t would not be an equilibrium).[3] When fixing the realization of W_t and the state-of-the-world ω in (3), the implied belief processes are identical when beginning from state $(z(q_t, \lambda, \hat{\lambda}), x_t)$ under $(\hat{Q}_t, \hat{\lambda})$ and when beginning from state (q_t, x_t) under (Q_t, λ). Therefore, it must be weakly optimal to wait until the belief reaches x' also from state $(z(q_t, \lambda, \hat{\lambda}), x_t)$ under $(\hat{Q}_t, \hat{\lambda})$. □

A.3 Preliminaries

We will utilize the following lemma that states the optimal stopping problem of an individual agent as a minimization problem of the *opportunity cost of delay*:

Lemma 1. *Fix arbitrary stock process Q. Then a player should not stop at state (q, x) if there exists a stopping time τ such that*

$$\mathbb{E}\left[\int_0^\tau e^{-rt}\left(x_t\pi_H(q_t) + (1 - x_t)\,\pi_L(q_t)\right)dt\,|(q, x; Q)\right] < 0. \tag{10}$$

[3]Notice that it is not possible that $dQ_t = 0$ for all (q_t, x') because $\pi_H(q_t) > 0$ for all q_t.

If no such stopping time exists, then it is optimal for a player to stop immediately at state (q, x).

Proof. Given process Q, the optimal stopping time τ^* must satisfy:

$$\tau^* \in \arg\sup_{\tau} \mathbb{E}\left[\int_{\tau}^{\infty} e^{-r(t)} \left(x_t \pi_H(q_t) + (1 - x_t)\, \pi_L(q_t)\right) dt \,|(q, x; Q)\right], \tag{11}$$

where maximization is over all stopping times τ adapted to the natural filtration generated by X_t under Q with initial value $(q_0, x_0) = (x, q)$. Rewriting the expectation as

$$\mathbb{E}\left[\int_{0}^{\infty} e^{-rt} \left(x_t \pi_H(q_t) + (1 - x_t)\, \pi_L(q_t)\right) dt \right.$$
$$\left. - \int_{0}^{\tau} e^{-rt} \left(x_t \pi_H(q_t) + (1 - x_t)\, \pi_L(q_t)\right) dt \,|(q, x; Q)\right],$$

we note that (11) is equivalent to

$$\tau_\theta^* \in \arg\inf_{\tau} \mathbb{E}\left[\int_{0}^{\tau} e^{-rt} \left(x_t \pi_H(q_t) + (1 - x_t)\, \pi_L(q_t)\right) dt \,|(q, x; Q)\right],$$

and so $\tau_\theta^* > 0$ only if (10) holds. □

A.4 Proof of Proposition 2

We show here that any equilibrium must satisfy the conditions in Proposition 2. We will do that in two steps. In the first step, we will show that the region $x > \bar{x}(q)$ of the state space must be a stopping region in any equilibrium, and in the second step we will show that there is no equilibrium where some player wants to stop for $x < \hat{x}(q)$.

Step 1: expansion when $x > \bar{x}(q)$.

Suppose that at time t the state is $(q_t, x_t) = (q, x)$, where $x > \bar{x}(q)$. If the stock were to be fixed at $q_t \equiv q$ forever, then it would be optimal to stop immediately. Combining this with Lemma 1 implies that for fixed $q_t \equiv q$, we have

$$\mathbb{E}\left[\int_{0}^{\tau} e^{-rt} \left(x_t \pi_H(q) + (1 - x_t)\, \pi_L(q)\right) dt \,|(q, x; q_t \equiv q)\right] \geq 0 \tag{12}$$

for all stopping times τ (given initial value (q, x) and $q_t \equiv q$ fixed).

Contrary to the claim, assume that no one stops at (q,x). But as long as no one stops, the optimality conditions (10) and (12) are equivalent, and hence a player prefers stopping immediately, contradicting that no one stops at (q,x) in an equilibrium.

Step 2: no stopping when $x < \hat{x}(q)$

We will now show that it cannot be optimal for any player to stop for $x < \hat{x}(q)$ in equilibrium. We will prove this through iterated deletion of dominated stopping strategies. Define

$$\overline{\pi}_{\omega}(q', q'') := \max_{q \in [q', q'']} \pi_{\omega}(q),$$

and define a mapping $\widetilde{x}_q : [q,1] \to [0,1]$ as

$$\widetilde{x}_q(q') := \sup \Big\{ x : \text{ there exists a stopping time } \tau \text{ such that}$$
$$\mathbb{E}\left(\int_0^{\tau} e^{-rt} \left(x \overline{\pi}_H(q,q') + (1-x)\, \overline{\pi}_L(q,q') \right) dt \, | q(\theta), x; q_t \equiv q(\theta) \right) < 0 \Big\}.$$

The interpretation is that $\widetilde{x}_q(q')$ would be the optimal stopping threshold if belief x_t follows process with q_t fixed at q, but with the flow payoff given in each state by $\max_{s \in [q,q']} \pi_{\omega}(s)$. It can be solved in closed form:

$$\widetilde{x}_q\left(q'\right) \;=\; \frac{-\beta(q)\overline{\pi}_L\left(q,q'\right)}{\left(\beta(q)-1\right)\overline{\pi}_H(q,q') - \beta(q)\overline{\pi}_L(q,q')}. \tag{13}$$

We see directly from (13) that $\widetilde{x}_q(q')$ is Lipschitz continuous in q and q', strictly increasing in q, weakly decreasing in q', and $\widetilde{x}_q(q) = \bar{x}(q)$.

We next define iteratively a sequence $\{x_n\}_{n=0}^{\infty}$ of functions $x_n : [0,1] \to [0,1]$. We first set $x_0(q) := \widetilde{x}_q\left(1\right)$. The assumption that q_t is fixed at q sets the learning rate at a lower bound for what it can be in the future and the assumption that the payoff is given by $\max_{s \in [q,1]} \pi_{\omega}(s)$ in state ω sets the flow payoff at an upper bound for what it can be in the future. Since an increase in the learning rate or a decrease in the payoff flow can only make stopping less profitable, stopping at (q,x) with $x < x_0\left(q\right)$ is a dominated strategy.

For $n \geq 1$, assume that $x_{n-1}(q)$ is a Lipschitz continuous and strictly increasing function with $x_{n-1}(q) \leq \hat{x}(q)$ for all q (note that these properties hold for $x_0(q)$ defined above). Define

$$x_n(q) := \widetilde{x}_q(q'), \text{ where } q' = \min\{s : \widetilde{x}_q\left(s\right) = x_{n-1}(s)\}. \tag{14}$$

The interpretation is that $x_n(q)$ would be the optimal stopping threshold under the assumption that q_t is fixed at q and the flow payoff is given by $\max_{s\in[q,q']}\pi_\omega(s)$ in state ω. Here $q' \in (q,1)$ is chosen in such a way that if no player ever stops below $x_{n-1}(q)$, then q' is an upper bound for q_t as long as $x_t \leq x_n(q)$. Hence $x_n(q)$ is a lower bound for the optimal stopping threshold when the current stock is q and under the assumption that no player stops below $x_{n-1}(q)$. It follows from (14) that if $x_{n-1}(q) < \hat{x}(q)$, then $x_{n-1}(q) < x_n(q) \leq \hat{x}(q)$, and if $x_{n-1}(q) = \hat{x}(q)$, then $x_n(q) = \hat{x}(q)$. Further, $x_n(q)$ inherits from $x_{n-1}(q)$ the properties of being Lipschitz continuous and strictly increasing function.

We have now shown by iterated deletion of dominated strategies that no player wants to stop at any (q,x) such that $x < x_n(q)$ for any n. The final step is to show that $x_n(q) \to \hat{x}(q)$ as $n \to \infty$. By the construction, $\{x_n\}$ is a family of equicontinuous and strictly increasing functions. Therefore, the sequence $\{x_n\}_{n=0}^{\infty}$ converges uniformly to some function $x(q)$, which is increasing and continuous. It also satisfies $x(q) \leq \hat{x}(q)$ for all q.

It remains to show that $x(q) = \hat{x}(q)$. First, notice that $|x_n(q) - x_{n-1}(q)| < \epsilon$ implies that $|x_{n-1}(q) - x_{n-1}(q')| < \epsilon$ for $q' = \min\{s : \widetilde{x}_q(s) = x_{n-1}(s)\}$ by construction. This further implies that either $|q - q'|$ is small or x_{n-1} is almost flat. When $n \to \infty$, this implies that the limiting x must satisfy that either $x(q) = \widetilde{x}_q(q) = \bar{x}(q)$ or $x'(q) = 0$ for all q. Because we know that $x_n(1) = \bar{x}(1)$ and that $x_n(q) \leq \bar{x}(q)$ for all n, it cannot be that x is constant and hence the only possibility is that $x(q) = \hat{x}(q)$ for all q.

A.5 Proof of Proposition 3

Proof. We prove that the stock process $\overline{Q}_t$ is an equilibrium (part i). The proof for $\underline{Q}_t$ follows the same steps and is omitted.

We need to verify that the best response to stock process $\overline{Q}_t$ is to stop when the belief reaches $\hat{x}(q)$. We will proceed through two main steps. In step 1, we will show that whenever $x_t < \hat{x}(q)$, it is optimal to wait. In step 2, we will show that if $x_t \geq \hat{x}(q)$, it is optimal to stop.

Step 1: Optimal to wait when $x_t < \hat{x}(q)$

It is easy to verify that

$$\mathbb{E}\left[\int_0^{\tau^*} e^{-rt}\left(x_t\pi_H(q) + (1-x_t)\,\pi_L(q)\right)dt\,|(q,x;q_t \equiv q)\right] < 0,$$

where τ^* is the stopping rule that commends to stop when x_t hits $\hat{x}(q)$. But then τ^* delivers a strictly negative opportunity cost of delay also against $\overline{Q}_t$ because the stock stays constant until hitting $\hat{x}(q)$:

$$\mathbb{E}\left[\int_0^{\tau^*} e^{-rt}\left(x_t\pi_H(q) + (1-x_t)\,\pi_L(q)\right)dt\,\Big|\left(q,x;\overline{Q}_t\right)\right] < 0,$$

and hence by Lemma 1 it is non-optimal to stop when $x < \hat{x}(q)$.

Step 2: Optimal to stop when $x_t = \hat{x}(q)$

Define

$$F(q,x) := \inf_{\tau} \mathbb{E}\left[\int_0^{\tau} e^{-rt}\left(x_t\pi_H(q) + (1-x_t)\,\pi_L(q)\right)dt\,\Big|\left(q,x;\overline{Q}_t\right)\right].$$

For a contradiction, suppose that there is some q such that $F\left(q,\hat{x}(q)\right) < 0$ so that by Lemma 1 it is not optimal to stop at that boundary point.

The general solution to the HJB-equation for $F(q,x)$ in the waiting region is

$$F(q,x) = \frac{x\pi_H\left(\theta,q\right) + (1-x)\,\pi_L\left(\theta,q\right)}{r} + A(q)\cdot\Phi(q,x) + B(q)\cdot\widetilde{\Phi}(q,x),$$

where $A(q)$ and $B(q)$ are some constants and $\Phi(q,x) := x^{\beta(q)}\left(1-x\right)^{1-\beta(q)}$ and $\widetilde{\Phi}(q,x) := x^{1-\beta(q)}\left(1-x\right)^{\beta(q)}$. We will first show that our assumption $F\left(q,\hat{x}(q)\right) < 0$ implies $F_x\left(q,\hat{x}(q)\right) \leq 0$. We will do that separately for potential cases $B(q) \geq 0$ and $B(q) < 0$.

Suppose that $B(q) \geq 0$. Then $F\left(\hat{x}(q),q\right) < 0$ implies that

$$A(q)\Phi\left(q,\hat{x}(q)\right) < -\frac{\hat{x}(q)\pi_H(q) + \left(1-\hat{x}(q)\right)\pi_L(q)}{r}. \tag{15}$$

But then,

$$
\begin{aligned}
&F_x\left(q,\hat{x}(q)\right)=\frac{\pi_H(q)-\pi_L(q)}{r}+\frac{\beta(q)-\hat{x}(q)}{\hat{x}(q)\left(1-\hat{x}(q)\right)}A(q)\Phi(q,\hat{x}(q))\\
&+\frac{1-\hat{x}(q)-\beta(q)}{\hat{x}(q)\left(1-\hat{x}(q)\right)}B(q)\Phi\left(\hat{x}(q),q\right)\\
&\leq\frac{\pi_H(q)-\pi_L(q)}{r}+\frac{\beta(q)-\hat{x}(q)}{\hat{x}(q)\left(1-\hat{x}(q)\right)}A(q)\Phi\left(q,\hat{x}(q)\right)\\
&<\frac{\pi_H(q)-\pi_L(q)}{r}\\
&-\frac{\beta(q)-\hat{x}(q)}{\hat{x}(q)\left(1-\hat{x}(q)\right)}\frac{\hat{x}(q)\pi_H(q)+\left(1-\hat{x}(q)\right)\pi_L(q)}{r}.
\end{aligned}
\tag{16}
$$

where the first inequality is implied by $B(q)\geq 0$ and $\beta(q)>1$, and the second inequality is implied by (15). Noting that

$$
\left(\beta(q)-1\right)\bar{x}(q)\pi_H(q)+\beta(q)\left(1-\bar{x}(q)\right)\pi_L(q)=0,
$$

which is equivalent to

$$
\frac{\pi_H(q)-\pi_L(q)}{r}=\frac{\left(\beta(q)-\bar{x}(q)\right)\left(\bar{x}(q)\pi_H(q)+\left(1-\bar{x}(q)\right)\pi_L(q)\right)}{r\bar{x}(q)\left(1-\bar{x}(q)\right)}
$$

and so it follows from (16) that $F_x\left(q,\hat{x}(q)\right)<0$ when $\hat{x}(q)=\bar{x}(q)$. When $\hat{x}(q)<\bar{x}(q)$, notice that the state jumps immediately to some $(q',\bar{x}(q'))$ s.t. $q'>q$ when hitting $(q,\hat{x}(q))$. Therefore, the optimality of stopping at $(q,\hat{x}(q))$ is equivalent to the optimality of stopping at $(q',\bar{x}(q'))$ and we ignore the case.

Suppose next that $B(q)<0$. This is only possible if there is some $x'<\bar{x}(q)$ where $F(q,x')\geq 0$ (otherwise it would be optimal to continue for all values $x<\bar{x}(q)$ and $B(q)<0$ would imply $\lim_{x\downarrow 0}F(q,x)=B(q)\lim_{x\downarrow 0}\widetilde{\Phi}(q,x)=-\infty$ contradicting the boundary condition at $x=0$). If $A(q)\geq 0$, then $F(q,x)$ is increasing in x, but then $F\left(q,x'\right)\geq 0$ contradicts our running assumption $F\left(q,\bar{x}(q)\right)<0$. The only remaining possibility is that both $B(q)$ and $A(q)$ are strictly negative. In that case $F(q,x)$ is concave, and so $F\left(q,x'\right)=0$ and $F\left(q,\bar{x}(q)\right)<0$ together imply $F_x\left(q,\bar{x}(q)\right)<0$.

We have now shown that $F\left(q,\bar{x}(q)\right)<0$ implies $F_x\left(q,\bar{x}(q)\right)<0$. But then the derivative of $F\left(q,\bar{x}(q)\right)$ along the boundary is also negative:

$$
\frac{d}{dq}F\left(q,\bar{x}(q)\right)=F_x\left(q,\bar{x}(q)\right)\frac{d}{dq}\bar{x}(q)<0.
$$

Iterating the argument along the boundary implies that $F(1, \bar{x}(1)) < 0$. But since $q_t = 1$ is an absorbing value for q_t, we can treat $q_t \equiv 1$ as fixed. Lemma 1 and that $\bar{x}(q)$ is the solution to the problem with fixed q imply that $F(1, \bar{x}(1)) < 0$ is a contradiction. Hence it is (weakly) optimal to stop in state $(q, \hat{x}(q))$.

□

A.6 Other equilibria

For our main result, it is enough to focus on the maximal and minimal equilibria as they imply dominance in the strong set order. Here, we provide additional material on what kind of other equilibria there exists.

Proposition 2 does not give a unique prediction for the stopping belief when $\bar{x}(q) > \hat{x}(q)$. Also, the size of the jump dQ_t is ambiguous above $\bar{x}(q_t)$. One can show by following the same argument as in Proof of Proposition 3 that in any equilibrium, either the stock stays constant or it jumps up to some $q' > q$ such that $\bar{x}(q') = x$ when the initial state is (q, x).In fact, one can support any such q' as the new equilibrium state if $\bar{x}(q') = x$ and $\bar{x}'(q') > 0$. For example, the maximal equilibrium jumps to the highest $q > q_t$ such that $\bar{x}(q) = x_t$ (or all the way to 1 if $\hat{x}(1) \leq x$) and the minimal equilibrium jumps to the lowest such q.

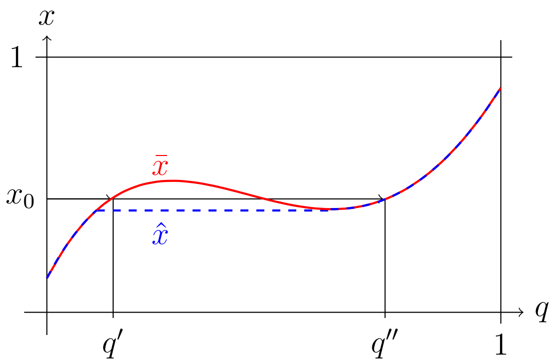

Figure 3: If the initial belief is x_0, both $dQ_0 = q'$ and $dQ_0 = q''$ are possible initial jumps in the stock.

Figure 3 illustrates why the non-monotonicity of $\bar{x}$ necessarily leads to multiple equilibria. In the figure, both q' and q'' satisfy $\bar{x}(q) = x_0$. Hence both the stock process that jumps to q' ($dQ_0 = q'$) and the stock process that jumps all the way to

q'' ($dQ_0 = q''$) are consistent with Proposition 2. Furthermore, one can construct equilibria where the jump depends on the fine details of the current belief. For instance in the situation of Figure 3, one could construct an equilibrium where the jump in the stock was (just below) q'' for the initial beliefs just below x_0 and where it was (just above) q' for the initial beliefs just above x_0. In such an equilibrium, the effect of the initial belief is not monotone. This happens because the belief works as a coordination device.

Importantly, all realizations of other equilibrium stock processes are between those from the maximal and minimal equilibria.

A.7 Belief dynamics

Here, we verify that Corollary 1 indeed implies an equivalent comparison for the equilibrium belief dynamics. We present the result only for the maximal equilibria under positive payoff externalities. The equivalent result holds for the minimal equilibria under positive payoff externalities. The result is reversed under negative externalities.

Let $\overline{x}_{t,\lambda}$ denote the belief at time t in the maximal equilibrium under learning technology λ. We get that the equilibrium belief process has larger expected volatility (more learning) when λ is low:

Proposition 5. *Suppose strictly positive payoff externalities and let λ' has strictly higher learning potential than λ. Then, $Pr(\overline{x}_{s,\lambda} \in [a,b]$ for all $s \leq t) < Pr(\overline{x}_{s,\lambda'} \in [a,b]$ for all $s \leq t)$ for all $t > 0$ and all $a < x_0$ and $b \in (x_0, \hat{x}_\lambda(1)]$.*

Proof idea.

The unconditional law-of-motion for the belief is

$$dX_t = \lambda(q_t)x_t(1 - x_t)dW'_t, \tag{17}$$

where W'_t is a standard Wiener process. We prove Proposition 5 by showing that the claim holds for all realization of W'_t:[4]

[4] Notice that W'_t is not the same Wiener process as W_t in (3). Therefore, Proposition 5 is written in terms of probabilities instead of realizations.

Lemma 2. *Fix the realization of W_t'. For each learning potential λ, define the joint belief and stock process $(\hat{X}_{t,\lambda}, \hat{Q}_{t,\lambda})$ such that the belief follows* (17) *and the stock is characterized by the cutoff rule $\hat{x}$ in* (6)*: in state (q, x), $d\hat{Q} = \max\{q' - q, 0\}$ where $q' = \max\{s \in [0, 1] : \hat{x}(s) = x\}$. Suppose that $x_{t',\lambda'} = x \in (0, x_0) \cup (x_0, \hat{x}_\lambda(1)]$ and that $x_{s,\lambda'} < \hat{x}_\lambda(1)$ for all $s \leq t'$. Then, $x_{t,\lambda} = x$ for some $t < t'$.*

Lemma 2 implies Proposition 5 because $\hat{X}_{t,\lambda}$ and $\hat{Q}_{t,\lambda}$ are distributed identically as the belief and stock processes in the maximal equilibrium.

Preliminaries.

Let x_t follow

$$dx_t = \lambda_t x_t (1 - x_t)\, dw_t' \tag{18}$$

with initial value x_0 and where w_t' is a standard Wiener process.

Lemma 3. *Let $\lambda_t := g(\max_{s \leq t}(x_s))$ and $\lambda_t' := h(\max_{s \leq t}(x_s'))$ with some strictly positive functions g and h, and let x_t and x_t' follow* (18) *with signal-to-noise-ratios λ_t and λ_t', respectively. Fix a realization of w_t' be such that $\lambda_s > \lambda_s'$ for all $s < t$. Then, $\max_{s \leq t}(x_s) > \max_{s \leq t}(x_s')$ if $\max_{s \leq t}(x_s') > x_0$ and $\min_{s \leq t}(x_s) > \min_{s \leq t}(x_s')$ if $\min_{s \leq t}(x_s') < x_0$.*

Proof. We use notation $m_t := \max_{s \leq t}(x_s)$ and $m_t' := \max_{s \leq t}(x_s')$. Since we have fixed a realization of w_t', m_t and m_t' are deterministic paths. Notice that $dm_t = 0$ if and only if $dm_t' = 0$ because λ_t and λ_t' stay constant whenever $x_t \leq m_t$ and $x_t' \leq m_t'$. Therefore both paths x_t and x_t' reach their earlier maximum again exactly at the same time.

Next, define path $\tilde{w}_t$ as $d\tilde{w}_t = dw_t'$ if $dm_t > 0$ and $d\tilde{w}_t = 0$ otherwise. Next, we use it to define path $\tilde{m}_t$ such that $d\tilde{m}_t = \lambda_t \tilde{m}_t (1 - \tilde{m}_t)\, d\tilde{w}_t$ and $\tilde{m}_0 = x_0$. Similarly, define $d\tilde{m}_t' = \lambda_t' \tilde{m}_t' (1 - \tilde{m}_t')\, d\tilde{w}_t$ with $\tilde{m}_0' = x_0$.

Now, since $\tilde{w}_t$ is an increasing path and $\lambda_s > \lambda_s'$ for all $s < t$, it must be that $\tilde{m}_t > \tilde{m}_t'$ unless $d\tilde{w}_s = 0$ everywhere. In the former case, $\tilde{m}_t = \tilde{m}_t' = x_0$. By construction, $\tilde{m}_t = m_t$ and $\tilde{m}_t' = m_t'$ for all t and hence $m_t > m_t'$ whenever $m_t' > x_0$.

The argument for the minimal paths $\min_{s\leq t}(x_s)$ and $\min_{s\leq t}(x'_s)$ is the same and is hence omitted. $\square$

Proof of Lemma 2.

Let the path of the stock defined in Lemma 2 under λ be $(\hat{q}_{t,\lambda})$ and let the path of the stock under λ' be $(\hat{q}_{t,\lambda'})$.

From Lemma 3, we get that the statement immediately holds if $\lambda(\hat{q}_{s,\lambda}) > \lambda'(\hat{q}_{s,\lambda'})$ for all $s \leq t$ (notice that the belief path is continuous). Therefore, we focus on the case where there exists s such that $\lambda(\hat{q}_{s,\lambda}) \leq \lambda'(\hat{q}_{s,\lambda'})$ and let t' be the smallest such time s.

Notice that $\lambda(\hat{q}_{t',\lambda}) \leq \lambda'(\hat{q}_{t',\lambda'})$ is equivalent to $\hat{q}_{t',\lambda'} \geq z(\hat{q}_{t',\lambda}; \lambda, \lambda')$. Because cutoff $\hat{x}_{\lambda'}(z(q;\lambda,\lambda')) > \hat{x}_{\lambda}(q)$ for all $q \in (0,1]$ (from Corollary 1, Part (i)), this further implies that we must have $\max_{s\leq t'} \hat{x}_{s,\lambda} =: z_{t'} < \max_{s\leq t'} \hat{x}_{s,\lambda'} =: z'_{t'}$ for all $z'_{t'} \leq \hat{x}_{\lambda}(1)$. But by Lemma 3, $z_{t'} < z'_{t'}$ is possible only if $\lambda(\hat{q}_{s,\lambda}) \leq \lambda'(\hat{q}_{s,\lambda'})$ for some $s < t'$, contradicting that t' is the smallest such s. Therefore, we conclude that $z_t > z'_t$ for all t such that $z'_t \in (x_0, \hat{x}_{\lambda}(1)]$, and then Lemma 3 implies the claim. $\square$

A.8 Proof of Proposition 4

Note that in equilibrium, all agents get the same expected payoff. The plan in the proof is to pick one individual agent and compare the expected discounted payoff computed at time zero across the two equilibria involving different learning potentials.

It is convenient to undertake a transformation of the state variable q_t. We denote by k_t the *information stock* at time t (as opposed to the *physical stock* q_t) and define it to be the signal-to-noise ratio at that moment:

$$k_t := \lambda\left(q_t\right).$$

Since $\lambda\left(\cdot\right)$ is strictly increasing, there is a one-to-one mapping between q_t and z_t, and therefore we may use (k_t, x_t) to summarize the state at time t. Notice that function $\beta(q)$ defined in the main text depends on q only through λ. Therefore, with some abuse of notation it is convenient to write:

$$\beta(k) := \beta\left(\lambda^{-1}(k)\right) = \frac{1}{2}\left(1+\sqrt{1+\frac{8r}{k}}\right).$$

Similarly, we may write:

$$\begin{aligned}\overline{x}(k) \quad &: \quad = \overline{x}\left(\lambda^{-1}(k)\right) = \frac{-\beta(k)\pi_L\left(\lambda^{-1}(k)\right)}{\left(\beta(k)-1\right)\pi_H\left(\lambda^{-1}(k)\right) - \beta(k)\pi_L\left(\lambda^{-1}(k)\right)}, \\ \widehat{x}(k) \quad &: \quad = \widehat{x}\left(\lambda^{-1}(k)\right) = \max\left\{x \in [0,1] : x \leq \overline{x}(k')\forall k' \in \left[k, \overline{k}_\lambda\right]\right\},\end{aligned}$$

where $\overline{k}_\lambda$ is the maximal signal-to-noise ratio $\overline{k}_\lambda := \lambda(1)$.

Notice that one can characterize a given equilibrium information stock process K_t in the (k,x)-space. We define *strictly divisive boundaries* for the maximal and minimal equilibria:

$$\begin{aligned}x^{\max}_{\lambda,x_0}(k) \quad &: \quad = \max\{x_0, \widehat{x}(k)\}, \\ x^{\min}_{\lambda,x_0}(k) \quad &: \quad = \max\{x_0, \max_{0\leq k'\leq k}\left(\overline{x}(k')\right)\}.\end{aligned}$$

The strictly divisive boundaries are defined for $k \in [0, \overline{k}_\lambda]$. Let $x(k)$ be the divisive boundary for either maximal or minimal equilibrium. The interpretation of a strictly divisive boundary is that k_t jumps up to $x(k_t)$ whenever x_t hits $x(k)$. (Notice that $x(k)$ may take a constant value within some interval, which means that k_t jumps instantaneously up by a discrete amount.) The equilibrium never reaches states above the divisive boundary. Divisive boundaries are continuous and increasing.

The payoff comparisons rely on $\mathbb{E}\left(e^{-r\tau}\right)$ for different paths in (k,x)-space, where τ is the time of reaching some state point. The following lemma shows how to compute this expectation:

Lemma 4. *Let $x(k)$ be the strictly divisive boundary for either the maximal or the minimal equilibrium, defined for $k \in \left[0, \overline{k}_\lambda\right]$ with initial state (k_0, x_0) such that $x_0 \leq x(k_0)$. Let $\overline{\tau}$ denote the time of reaching state $\left(\overline{k}_\lambda, x'\right)$ with $x' \geq x\left(\overline{k}_\lambda\right)$:*

$$\overline{\tau} := \inf\left(t : x_t = x',\ k_t = \overline{k}_\lambda\right).$$

Then

$$\mathbb{E}\left(e^{-r\overline{\tau}}\left|k_0, x_0\right.\right) = \frac{\Phi\left(k_0, x_0\right)}{\Phi\left(\overline{k}_\lambda, x'\right)} \exp\left(\int_{k_0}^{\overline{k}_\lambda} \beta'\left(s\right) \log\left(\frac{x\left(s\right)}{1 - x\left(s\right)}\right) ds\right),$$

where $\Phi\left(k, x\right) = x^{\beta(k)}\left(1 - x\right)^{1-\beta(k)}$.

Proof. For an arbitrary current state $(k_t, x_t) = (k, x)$, we can write

$$\mathbb{E}\left(e^{-r\overline{\tau}}\left|k, x\right.\right) = A\left(k\right)\Phi\left(k, x\right),$$

where $A\left(k\right)$ is a function to be determined. We will derive a differential equation for $A\left(k\right)$. Note that k increases at the boundary $x\left(k\right)$, and therefore the partial of $\mathbb{E}\left(e^{-r\overline{\tau}}\left|k, x\right.\right)$ w.r.t. k must be zero at the boundary:

$$\frac{\partial}{\partial k}\mathbb{E}\left(e^{-r\overline{\tau}}|k, x(k)\right) = A'\left(k\right)\Phi\left(k, x\left(k\right)\right) + A\left(k\right)\Phi_k\left(k, x\left(k\right)\right) = 0.$$

Noting that

$$\Phi_k\left(k, x\right) = \beta'\left(k\right)\log\left(\frac{x}{1 - x}\right)\Phi\left(k, x\right),$$

we get the following differential equation for $A\left(k\right)$:

$$A'\left(k\right) = -A\left(k\right)\beta'\left(k\right)\log\left(\frac{x\left(k\right)}{1 - x\left(k\right)}\right). \tag{19}$$

From $\mathbb{E}\left(e^{-r\overline{\tau}}\left|x', \overline{k}_\lambda\right.\right) = A\left(\overline{k}_\lambda\right)\Phi\left(\overline{k}_\lambda, x'\right) = 1$, we get the following boundary condition:

$$A\left(\overline{k}_\lambda\right) = \frac{1}{\Phi\left(\overline{k}_\lambda, x'\right)},$$

which gives us the solution to (19) as:

$$A\left(k\right) = \frac{1}{\Phi(\overline{k}_\lambda, x')}\exp\left(\int_{k}^{\overline{k}_\lambda} \beta'\left(s\right)\log\left(\frac{x\left(s\right)}{1 - x\left(s\right)}\right) ds\right).$$

This gives

$$\mathbb{E}\left(e^{-r\overline{\tau}}\left|k_0, x_0\right.\right) = A\left(k_0\right)\Phi\left(k_0, x_0\right) = \frac{\Phi\left(k_0, x_0\right)}{\Phi\left(\overline{k}_\lambda, x'\right)}\exp\left(\int_{k_0}^{\overline{k}_\lambda} \beta'\left(s\right)\log\left(\frac{x\left(s\right)}{1 - x\left(s\right)}\right) ds\right).$$

□

We now proceed to the proof of the statements in the proposition. Consider first the case of strictly positive payoff externalities and compare the maximal equilibrium $\overline{Q}_t$ and $\overline{Q}_t'$ corresponding to λ and λ', respectively, where λ has a strictly higher learning potential than λ'. For notational simplicity, we use $x(k)$ for $x^{\max}_{\lambda,x_0}(k)$ and $x'(k)$ for $x^{\max}_{\lambda',x_0}(k)$. It follows from (1) that $x(k) \geq x'(k)$ for all $0 \leq z \leq \overline{k}_{\lambda'}$ and where $x(k) > x'(k)$ whenever $x'(k) > x_0$. Figure 4 depicts the divisive boundaries, $x'(k)$ in red and $x(k)$ in black.

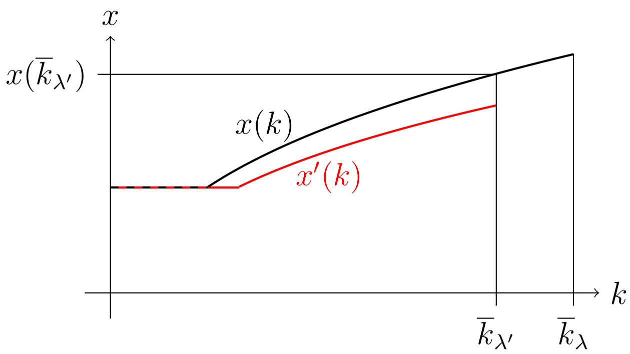

Figure 4: Divisive boundaries for learning potentials λ and λ' under positive payoff externalities: $x(k) = x^{\max}_{\lambda,x_0}(k)$ and $x'(k) = x^{\max}_{\lambda',x_0}(k)$.

We derive a lower bound for $v^{\max}_{\lambda'}(x_0)$, i.e. the ex-ante payoff of an individual agent under λ'. To do this, note first that $v^{\max}_{\lambda'}(x_0)$ is equal to the expected payoff of the agent who is the last one ever to stop, i.e. the player who stops at $\tau' := \inf\left(t : x_t = x'\left(\overline{k}_{\lambda'}\right)\right)$:

$$v^{\max}_{\lambda'}(x_0) = \mathbb{E}\left(e^{-r\tau'}\frac{x'\left(\overline{k}_{\lambda'}\right)\pi_H(1) + \left(1 - x'\left(\overline{k}_{\lambda'}\right)\right)\pi_L(1)}{r}\right).$$

Suppose that this agent makes a deviation and decides to stop only when x_t hits $x(\overline{k}_{\lambda'})$, i.e. at the belief when the information stock reaches $\overline{k}_{\lambda'}$ under the other learning potential. Let $\overline{\tau}_{\lambda'}$ denote the corresponding (stochastic) stopping time. This deviation gives a lower a payoff than equilibrium payoff $v^{\max}_{\lambda'}(x_0)$, and so we get

$$v^{\max}_{\lambda'}(x_0) > \mathbb{E}\left(e^{-r\overline{\tau}}\frac{x(\overline{k}_{\lambda'})\pi_H(1) + \left(1 - x(\overline{k}_{\lambda'})\right)\pi_L(1)}{r}\right) \tag{20}$$

$$= \frac{\Phi(k_0, x_0)}{\Phi(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'}))}\exp\left(\int_0^{\overline{k}_{\lambda'}}\beta'(s)\log\left(\frac{x'(s)}{1 - x'(s)}\right)ds\right)\frac{x(\overline{k}_{\lambda'})\pi_H(1) + \left(1 - x(\overline{k}_{\lambda'})\right)\pi_L(1)}{r}.$$

Let us now consider the corresponding payoff for the learning potential λ, denoted $v_{\lambda}^{\max}(x_0)$. Since $\left(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'})\right)$ is on the path of the maximal equilibrium (the black curve in Figure 4), we can compute $v_{\lambda}^{\max}(x_0)$ by considering the agent that plans to stop at $\left(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'})\right)$, i.e.

$$v_{\lambda}^{\max}(x_0) = \mathbb{E}\left(e^{-r\overline{\tau}_{\lambda}} u\left(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'})\right)\right), \tag{21}$$

where $u\left(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'})\right)$ is the stopping value at state $\left(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'})\right)$. Since we have strictly positive payoff externalities, the stopping value at $\left(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'})\right)$, where not all the agents have stopped yet, must be strictly lower than the stopping value assuming that all the agents have already stopped, i.e.

$$u\left(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'})\right) < \frac{x(\overline{k}_{\lambda'}\pi_H(1) + \left(1 - x(\overline{k}_{\lambda'}\right)\pi_L(1)}{r}. \tag{22}$$

Combining (21) and (22), we get

$$\begin{aligned} v_{\lambda}^{\max}(x_0) &< \mathbb{E}\left(e^{-r\overline{\tau}_{\lambda}} \frac{x(\overline{k}_{\lambda'}\pi_H(1) + \left(1 - x(\overline{k}_{\lambda'}\right)\pi_L(1)}{r}\right) \qquad (23) \\ &= \frac{\Phi(k_0, x_0)}{\Phi(\overline{k}_{\lambda'}, x(\overline{k}_{\lambda'}))} \exp\left(\int_0^{\overline{k}_{\lambda'}} \beta'(s) \log\left(\frac{x(s)}{1 - x(s)}\right) ds\right) \frac{x(\overline{k}_{\lambda'})\pi_H(1) + \left(1 - x(\overline{k}_{\lambda'})\right)\pi_L(1)}{r}, \end{aligned}$$

where the only difference to expression (21) is that we have the divisive boundary $x(s)$ inside the integral instead of $x'(s)$ (since we move along the black curve instead of the red curve in Figure 4). Since $x(k) \geq x'(k)$ for all k and $x(k) > x'(k)$ for some $k < \overline{k}_{\lambda'}$, direct comparison of the expressions (21) and (24) imply that

$$v_{\lambda}^{\max}(x_0) < v_{\lambda'}^{\max}(x_0).$$

The argument for the payoff comparison under minimal equilibrium is identical and gives

$$v_{\lambda}^{\min}(x_0) < v_{\lambda'}^{\min}(x_0).$$

Let us then consider the case of strictly negative payoff externalities. Note that we have then a unique equilibrium and therefore $v_{\lambda}^{\min}(x_0) = v_{\lambda}^{\max}(x_0) =: v_{\lambda}(x_0)$ for any learning potential λ. Let us again compare λ and λ', where λ has a strictly higher learning potential. In this case, we have $x(k) \leq x'(k)$ for all $k \in \left[0, \overline{k}_{\lambda'}\right]$ and $x(k) > x'(k)$ whenever $x(k) > x_0$.

To compare the payoffs across λ and λ', take an agent who stops at $\overline{\tau} := \inf\left(t : x_t = x'(\overline{k}_{\lambda'})\right)$. There are now two cases for the path under λ depending on whether $x\left(\overline{k}_\lambda\right) < x'(\overline{k}_\lambda)$ (the black curve in the left panel of Figure 5) or $x\left(\overline{k}_\lambda\right) \geq x'(\overline{k}_\lambda)$(the blue curve in the right panel of Figure 5). Next, we argue that both lead to the same conclusion (26).

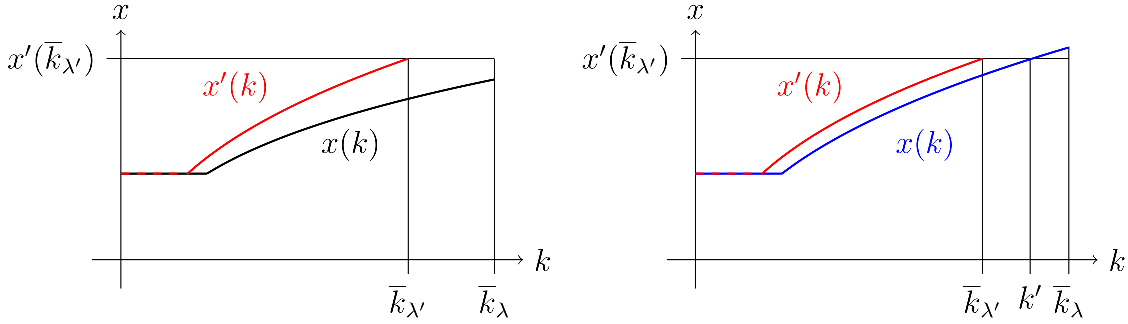

Figure 5: Two cases under negative payoff externalities.

In the latter case (blue curve in right), stopping at $x'(\overline{k}_{\lambda'})$ is optimal for some level of the information stock under learning potential λ. Let k' denote that stock: $x(k') = x'(\overline{k}_{\lambda'})$. Therefore, the stopping payoff can be written as

$$v_\lambda\left(x_0\right) = \mathbb{E}\left(e^{-r\overline{\tau}}u\left(k', x'(\overline{k}_{\lambda'})\right)\right), \tag{24}$$

where the stopping payoff satisfies

$$u\left(k', x'(\overline{k}_{\lambda'})\right) > \frac{x'(\overline{k}_{\lambda'})\pi_H\left(1\right) + \left(1 - x'(\overline{k}_{\lambda'})\right)\pi_L\left(1\right)}{r} \tag{25}$$

since not all the agents have yet stopped and payoff externalities are strictly negative. Combining (24) and (25), we get

$$v_\lambda\left(x_0\right) > \mathbb{E}\left(e^{-r\overline{\tau}}\frac{x'(\overline{k}_{\lambda'})\pi_H\left(1\right) + \left(1 - x'(\overline{k}_{\lambda'})\right)\pi_L\left(1\right)}{r}\right). \tag{26}$$

In the other case (black curve in left), stopping at $x'(\overline{k}_{\lambda'})$ is sub-optimal (since $\left(k, x'(\overline{k}_{\lambda'})\right)$ is not on the path for any k), and we have therefore

$$v_\lambda\left(x_0\right) > \mathbb{E}\left(e^{-r\overline{\tau}}u\left(\overline{k}_\lambda, x'(\overline{k}_{\lambda'})\right)\right),$$

but since in this case all the agents have stopped at $(\overline{k}_\lambda, x'(\overline{k}_{\lambda'}))$, we have

$$u\left(\overline{k}_\lambda, x'(\overline{k}_{\lambda'})\right) = \frac{x'(\overline{k}_{\lambda'})\pi_H\left(1\right) + \left(1 - x'(\overline{k}_{\lambda'})\right)\pi_L\left(1\right)}{r},$$

and so (26) holds in this case too. Combining (26) with Lemma 4 gives us:

$$v_\lambda(x_0) > \frac{\Phi(k_0, x_0)}{\Phi(\overline{k}_\lambda, x'(\overline{k}_{\lambda'}))} \exp\left(\int_{k_0}^{\overline{k}_\lambda} \beta'(s) \log\left(\frac{x(s)}{1-x(s)}\right) ds\right) \frac{x'(\overline{k}_{\lambda'})\pi_H(1) + \left(1 - x'(\overline{k}_{\lambda'})\right)\pi_L(1)}{r}. \tag{27}$$

Consider now the learning technology λ' (the red curve in both parts of Figure 5) and again consider an agent who stops at $\overline{\tau}' := \inf\left(t : x_t = x'(\overline{k}_{\lambda'})\right)$. This is now the last agent to stop in equilibrium and gives us:

$$v_{\lambda'}(x_0) = \mathbb{E}\left(e^{-r\overline{\tau}'} \frac{x'(\overline{k}_{\lambda'})\pi_H(1) + \left(1 - x'(\overline{k}_{\lambda'})\right)\pi_L(1)}{r}\right) \tag{28}$$

$$= \frac{\Phi(k_0, x_0)}{\Phi(\overline{k}_{\lambda'}, x'(\overline{k}_{\lambda'}))} \exp\left(\int_{k_0}^{\overline{k}_{\lambda'}} \beta'(s) \log\left(\frac{x'(s)}{1-x'(s)}\right) ds\right) \frac{x'(\overline{k}_{\lambda'})\pi_H(1) + \left(1 - x'(\overline{k}_{\lambda'})\right)\pi_L(1)}{r}.$$

Notice that in contrast to the expression in (27), we now have $\overline{k}_\lambda$ instead of $\overline{k}_{\lambda'}$ as the first argument in $\Phi(\cdot,\cdot)$ as well as in the upper bound of the integration. Using $\Phi(k, x) = x^{\beta(k)}(1-x)^{1-\beta(k)}$, we may write

$$\begin{aligned}
\frac{\Phi(k_0, x_0)}{\Phi(\overline{k}_{\lambda'}, x'(\overline{k}_{\lambda'}))} &= \frac{\Phi(k_0, x_0)}{\Phi(\overline{k}_\lambda, x'(\overline{k}_\lambda))} \left(\frac{x'(\overline{k}_{\lambda'})}{1 - x'(\overline{k}_{\lambda'})}\right)^{\beta(\overline{k}_\lambda) - \beta(\overline{k}_{\lambda'})} \\
&= \frac{\Phi(k_0, x_0)}{\Phi\left(\overline{k}_\lambda, x'(\overline{k}_{\lambda'})\right)} \exp\left(\int_{\overline{k}_{\lambda'}}^{\overline{k}_\lambda} \beta'(s) \log\left(\frac{x'(\overline{k}_{\lambda'})}{1 - x'(\overline{k}_{\lambda'})}\right) ds\right),
\end{aligned}$$

and so we may write (29) as

$$v_{\lambda'}(x_0) = \frac{\Phi(k_0, x_0)}{\Phi\left(\overline{k}_\lambda, x'(\overline{k}_{\lambda'})\right)} \exp\left(\int_{k_0}^{\overline{k}_{\lambda'}} \beta'(s) \log\left(\frac{x(s)}{1-x(s)}\right) ds + \int_{\overline{k}_{\lambda'}}^{\overline{k}_\lambda} \beta'(s) \log\left(\frac{x'(\overline{k}_{\lambda'})}{1 - x'(\overline{k}_{\lambda'})}\right) ds\right)$$
$$\cdot \frac{x'(\overline{k}_{\lambda'})\pi_H(1) + \left(1 - x'(\overline{k}_{\lambda'})\right)\pi_L(1)}{r}.$$

Comparing this expression with (27) and noting $x(k) \leq \min\{x'(k), x'(\overline{k}_{\lambda'})\}$ for $k \leq \overline{k}_{\lambda'}$ (with strict inequality for a range of values), we get $v_\lambda(x_0) > v_{\lambda'}(x_0)$. □

References

Baldursson, Fridrik M. and Ioannis Karatzas, "Irreversible investment and industry equilibrium," *Finance and Stochastics*, Dec 1996, *1* (1), 69–89.

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.

Bolton, Patrick and Christopher Harris, "Strategic Experimentation," *Econometrica*, 1999, *67* (2), 349–374.

Dixit, Avinash K. and Robert S. Pindyck, *Investment under Uncertainty*, Princeton: Princeton University Press, 1994.

Décamps, Jean-Paul and Thomas Mariotti, "Investment timing and learning externalities," *Journal of Economic Theory*, 2004, *118* (1), 80–102.

Farrell, Joseph and Garth Saloner, "Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation," *The American Economic Review*, 1986, *76* (5), 940–955.

Frick, Mira and Yuhta Ishii, "Innovation Adoption by Forward-Looking Social Learners," *Working Paper*, 2020.

Halac, Marina, Navin Kartik, and Qingmin Liu, "Contests for Experimentation," *Journal of Political Economy*, 2017, *125* (5), 1523–1569.

Jovanovic, Boyan and Saul Lach, "Entry, Exit, and Diffusion with Learning by Doing," *The American Economic Review*, 1989, *79* (4), 690–699.

Katz, Michael L. and Carl Shapiro, "Technology Adoption in the Presence of Network Externalities," *Journal of Political Economy*, 1986, *94* (4), 822–841.

Keller, Godfrey and Sven Rady, "Breakdowns," *Theoretical Economics*, 2015, *10* (1), 175–202.

_ , _ , and Martin Cripps, "Strategic Experimentation with Exponential Bandits," *Econometrica*, 2005, *73* (1), 39–68.

Kwon, H. Dharma, Wenxin Xu, Anupam Agrawal, and Suresh Muthulingam, "Impact of Bayesian Learning and Externalities on Strategic Investment," *Management Science*, February 2016, *62* (2), 550–570.

Laiho, Tuomas, Pauli Murto, and Julia Salmi, "Gradual Learning from Incremental Actions," *Working paper*, 2023.

Leahy, John V, "Investment in competitive equilibrium: The optimality of myopic behavior," *The Quarterly Journal of Economics*, 1993, *108* (4), 1105–1133.

Margaria, Chiara, "Learning and payoff externalities in an investment game," *Games and Economic Behavior*, 2020, *119*, 234–250.

Strulovici, Bruno, "Learning while Voting: Determinants of Collective Experimentation," *Econometrica*, 2010, *78* (3), 933–971.

Thomas, Caroline D., "Strategic Experimentation with Congestion," *American Economic Journal: Microeconomics*, February 2021, *13* (1), 1–82.