Timing of Irreversible Actions under Informational and Payoff Externalities

work in progress

Tuomas Laiho, Pauli Murto, Julia Salmi*

May 11, 2022

Abstract

This paper studies 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 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 rather surprising effects: a technological improvement in learning may in fact slow down learning in equilibrium.

JEL classification: C61, C73, D82, D83

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

*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). We thank Yeon-Koo Che, Francesc Dilme, Yingni Guo, Daniel Hauser, Jan Knoepfle, Vasiliki Skreta, Peter Norman Sørensen, Juuso Välimäki, and the seminar audiences at BU, UT Austin, VSET, and Incentives and dynamic games workshop in Paris for their helpful comments. This project has received financial support from 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 and commercial platforms, and electric vehicles. Another characteristic feature of new innovations is that there is common uncertainty about their quality. Hence, by experimenting with 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 widely?

We show that if experimentation by others is very informative about the overall quality, innovation adoption may slow down in equilibrium. Surprisingly, also learning may be slower if the learning technology is better. It turns out that payoff externalities play a major role in the result: equilibrium learning slows down only if there are positive 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-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 make irreversible stopping decisions (e.g. when to adopt an innovation), and 2) the payoff after stopping depends on the fraction of how many other players have stopped. Furthermore, we focus on experimentation in large games where the number of players goes to infinity.

Any Markov perfect equilibrium can essentially be characterized by a cutoff belief that depends on the current stock of stopped agents. More players are willing to stop when the belief that the state is high gets above the cutoff. The profitability of stopping depends on other players through two channels: the players who have already stopped generate information and affect the current payoff externality. The players who are still waiting may stop later and hence change the payoff of other players. Finding the cutoff belief is therefore potentially a complicated problem that depends on, among other things, expectations for the future. We extend the observation made in our earlier paper Laiho, Murto and Salmi (2022) that in the equilibrium, players may ignore future changes in the stock when choosing their optimal stopping time. This observation enables solving the cutoff in closed form.

Laiho, Murto and Salmi (2022) does not include payoff externalities. From a technical perspective, payoff externalities make the stopping payoffs depend on the expectations of what other players will do in the future. To eliminate other kind of equilibria in the present paper, we need a different proof that is based on the iterative elimination of strictly dominated strategies. Payoff externalities complicate the analysis also through the shape of the stock-dependent cutoff belief. Absent payoff externalities, players require a higher belief to be willing to stop if the stock is high because then there is more information arriving. This need not be the case with positive payoff externalities because a high stock has a direct impact on the stopping payoff. The non-monotonicity requires special treatment in the proofs and leads to the multiplicity of equilibria. We characterize the equilibrium with the fastest learning and the equilibrium with the slowest learning, and show when they coincide.

Payoff externalities change the results also qualitatively. Most importantly, a better learning technology (less noise in signals) may cause the learning to be strictly slower in the equilibrium. This happens when the payoff externality is positive. The intuition behind the result is that the stock must increase slower under a better learning technology because faster learning makes waiting more profitable. This effect brings the stock down as long as the flow of new information is higher under the better learning technology. Positive payoff externalities then imply that no player is willing to stop under the better learning technology when the flows of information are equal because then the stock, and hence the stopping payoff, is lower under the better technology. This result is striking because it means that precisely when a fast adoption of a new innovation would be the most important for the aggregate welfare – both because of experimentation creating a lot of information and because of positive network effects – it is the hardest to achieve.

If payoff externalities are negative instead, better learning technology always leads to faster learning in the equilibrium. Furthermore, the Markov perfect equilibrium is unique. Negative payoff externalities naturally arise between firms operating in a same market. A firm investing in production capacity imposes a negative externality on other producers. We solve an industry equilibrium in a market of unknown size and derive new economic insights, including the effects of taxes and subsidies on investments in productive capital.

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 also information generation is persistent, and that the player's payoffs depend on the stopping decisions of other players. The former feature is present also in our earlier paper Laiho, Murto and Salmi (2022) 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 (2022), informational free-riding mitigates the gains of better learning technology. However, because there are no payoff externalities, larger potential of learning does not lead to lower welfare. Frick and Ishii (2020) provide an extension where a better learning technology may lead to a 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 e.g. 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.

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 Model

2.1 Actions and payoffs

Time is continuous and goes to infinity, $t \in [0, \infty)$. N identical agents choose when to stop. When an agent stops, he starts receiving a flow payoff

$$du_t^i = \frac{\pi(q_t, \omega)dt}{N} + \frac{\sigma}{N} dW_t^i \tag{1}$$

that depends on the fraction of other agents who have stopped, $q_t = N(stopped)/N$, 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.

Agents discount with rate r. Before stopping, all agents receive a flow payoff 0. Throughout, we assume that $\pi(q, H) > 0$ and that $\pi(q, L) < 0$ for all q, which implies that all agents 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. Throughout, we assume that π is absolutly continuous in q. We focus on the case of a large game and take $N \to \infty$.

2.2 Learning

Let x_t denote the public belief that the state-of-the-world is high. The realized payoffs are informative about the state of the world. Let n_t be the number of agents who have stopped by time t. The total informativeness of the experiment coincides with observing the sum of flow payoffs of the agents who have stopped:

$$dU_t = \sum_{i=1}^{n_t} \left(\frac{\pi(q_t, \omega)}{N} + \sqrt{\frac{n_t}{N}} \sigma dW_t^i \right). \tag{2}$$

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 $q_t \sigma^2 dt$. The implied signal-to-noise ratio of the signal process U_t is $\lambda(q_t) = \frac{\sqrt{q_t}(\pi(q_t,H) - \pi(q_t,L))}{\sigma}$. Throughout the paper, we assume that the signal-to-noise ratio is increasing in q_t.[1]

The public belief, x_t, then follows the Brownian motion with zero drift and variance $(\lambda(q_t) x_t (1 - x_t))^2$.[2]

3 Equilibrium

3.1 Definition

Let $N = \infty$. A Markov strategy for player i is a mapping $\xi_i : [0,1] \times [0,1] \to [0,1]$ from the current belief and the fraction of stopped agents to a probability of stopping. When all players follow Markov strategies, the stock of stopped agents follows an increasing process Q_t adopted to the natural filtration generated by U_t. Then, a player's value satisfies

$$v(q_t, x_t) = \sup_{\tau} \mathbb{E}_{Q_t} \left[\int_{\tau}^{\infty} e^{-r(s-t)} \Big(x_s \pi(q_s, H) + (1 - x_s) \pi(q_s, L) \Big) ds | (q_t, x_t) \right]. \tag{3}$$

Definition 1. *A stock process Q_t is an equilibrium if*

(i) $v(q_t, x_t) = \mathbb{E}_{Q_t} \left[\int_t^{\infty} e^{-r(s-t)} \Big(x_s \pi(q_s, H) + (1 - x_s) \pi(q_s, L) \Big) ds | (q_t, x_t) \right]$ *whenever* $dQ_t > 0$,

[1] A sufficient condition is that $\pi(q, H) - \pi(q, L)$ is non-decreasing in q, i.e. having more players stopping does not diminish the effect of the state-of-the-world on payoffs.

[2] This is a standard results in Brownian learning, for derivations see e.g. Bolton and Harris (1999).

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

3.2 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.

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 $\bar{x}$:

$$\bar{x}(q) := \frac{-\beta(q)\pi(q, L)}{(\beta(q) - 1)\,\pi(q, H) - \beta(q)\pi(q, L)}, \tag{4}$$

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

$$\hat{x}(q) := \max\{x \in [0, 1] : x \leq \bar{x}(q') \quad \forall q' \geq q\}. \tag{5}$$

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.2 that the equilibrium cutoff in the game where Q_t changes endogenously coincides with the cutoff with fixed stock: [3]

Proposition 1. *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* (4) *and* (5).

Let us first comment briefly the implications of the cutoff structure generally and then the specific form of the equilibrium cutoff. If $\bar{x}$ is monotone, Proposition 1

[3]This result extends the characterization of the decentralized equilibrium in Laiho, Murto and Salmi (2022) without payoff externalities and with heterogeneous players.

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 agents 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, there are some states between the curves $\bar{x}$ and $\hat{x}$ and for those Proposition 1 does not imply if the players should stop or wait. This 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. We will discuss possible non-monotonicity more in Sections 4 and 5, separately for positive and negative payoff externalities.

When it comes to the specific form of $\bar{x}(q)$ in (4), it coincides with the optimal stopping threshold of an individual agent when the stock stays constant forever. This result is rather surprising as it means that the players can ignore both faster learning and the payoff consequences of future stopping decisions. 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 some other stock process Q_t such that Proposition 1 does not characterize the optimal stopping rule for individual players.

4 Positive payoff externalities

Positive payoff externalities feature innovation adoption with network effects. Consider for instance electric cars, new piece of software, or certification standards. The more users they have, the better services are available for them and the more aware of other people are about them.

4.1 Multiplicity of equilibria

As discussed in the previous section, there can be multiple equilibria. With positive payoff externalities, there is a coordination game between the players, which causes multiplicity of equilibria even without learning. If players expect many other players to stop, they want to stop, too.

Technically, the multiplicity is related to the non-monotonicity of $\bar{x}$. By looking at the formula in (4), one sees that $\bar{x}$ is non-monotone if there are sufficiently strong positive payoff externalities, i.e. $\bar{x}'(q, \omega)$ are large enough. How can one see that non-monotone $\bar{x}$ leads to multiplicity of equilibria? The first thing to notice is that Proposition 1 does not give a unique prediction for the stopping belief when $\bar{x}(q) > \hat{x}(q)$, which happens precisely when $\bar{x}$ is non-monotone. Second, the size of the jump dQ_t is ambiguous: when there are multiple $q > q_t$ that satisfy $\bar{x}(q) = x_t$, the stock process that jumps to any of them from state (q_t, x_t) is an equilibrium.

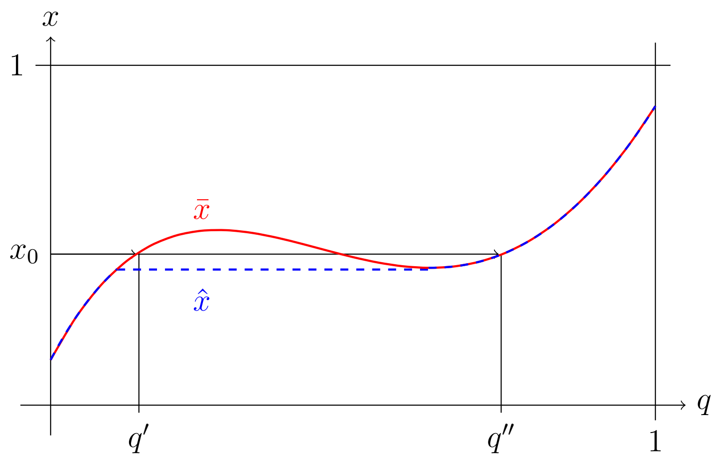

Figure 1: If the initial belief is x_0, both $dQ_0 = q'$ and $dQ_0 = q''$ are possible initial jumps in the stock.

Figure 1 illustrates the consequences of non-monotone $\bar{x}$. 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 1.

The multiplicity of cutoff equilibria where more players stop if the belief is high gives rise to the existence of equilibria where the effect of the belief is not monotone. For instance in the situation of Figure 1, 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. We do not analyze this kind of non-monotonicity further in this paper because it is somewhat artificial and the realizations of such stock processes are anyway between those from the maximal and minimal cutoff equilibria, which we define next.

Definition 2.

- *The maximal cutoff equilibrium is characterized by the cutoff rule $\hat{x}$: for any (x, q), $dQ = q' - q$ where $q' = \max\{s \geq q : \hat{x}(q) = x\}$.*
- *The minimal cutoff equilibrium is characterized by the cutoff rule $\bar{x}$: for any (x, q), $dQ = q' - q$ where $q' = \min\{s \geq q : \bar{x}(q) = x\}$.*

To correctly interpret the maximal and minimal equilibria, recall that a lower cutoff belief implies a larger stock and faster learning. By Proposition 1, 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. We verify in Appendix A.3 that the stock processes defined in Definition 2 are indeed equilibria.

4.2 Better learning technology may harm learning and welfare

We now show that improvements in the learning technology may backfire: under positive payoff externalities, an improvement in the learning technology (lower σ) slows down learning and reduces total welfare. This shows that the interaction of informational and payoff externalities may lead to quite unexpected results.

To make the comparison precise, we consider the maximal equilibria under different learning technologies. Let $\hat{x}_\sigma$ denote the stopping cutoff in the maximal equilibrium under learning technology σ. We have the following welfare result that holds independent of any properties of the solution (proof in Appendix A.4):

Proposition 2. *Let $\sigma' < \sigma$. Then, the players are strictly better off in the maximal equilibrium under σ than under σ' for all initial beliefs $x_0 \in (\hat{x}_\sigma(0), \hat{x}_\sigma(1)]$.*

The key step to prove Proposition 2 is to show that learning is paradoxically slower under the better learning technology. The intuition is that a player does not want to stop when learning is equally fast because it means that the stock is lower under the better learning technology. A lower stock implies a lower stopping payoff due to positive payoff externalities. As a result, players are worse off because they face both a slower increase in the stock and slower arrival of information.

Proposition 2 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.

5 Negative payoff externalities

Negative payoff externalities naturally arise when firms choose about investments in productive capital and operate in the same market. We first show some general results that hold under negative payoff externalities that enable comparison with the case with positive payoff externalities. Then, we analyze industry equilibrium more thoroughly and discuss how endogenous learning interacts with subsidy and tax policies.

5.1 Equilibrium under negative payoff externalities

We start with a proposition that fully characterizes the equilibrium when there are no payoff externalities or when they are negative:[4]

Proposition 3. *When there are no payoff externalities or when they are negative, there is a unique equilibrium characterized by $\bar{x}$ in (4). Furthermore, a better learning technology leads to faster learning under negative payoff externalities and has no effect on learning without payoff externalities.*

The first part of Proposition 3 is almost a direct consequence of Proposition 1. One needs to show that the cutoff belief $\bar{x}$ in (4) is increasing under weakly negative payoff externalities, which further implies that there is a unique candidate for the equilibrium as $\hat{x}(q) = \bar{x}(q)$ for all q. Then, the result follows once one verifies that $\bar{x}$ characterizes an equilibrium. The second part is a counter part of Proposition 2. The intuition is pretty much the same but reversed. See Appendix A.5 for the complete proof.

Why do positive and negative payoff externalities lead to so different equilibrium predictions? Notice that endogenous learning causes a positive informational externality that goes to the players who have not stopped yet. When comparing optimal stopping decisions under better and worse learning technologies, both the informational and the payoff effects of a high stock go to the same direction under negative payoff externalities: a player is less willing to stop if the stock is high. Therefore, if the learning technology improves, the stock does not need to adjust down very much to make the player willing to stop again. Under positive payoff externalities, the decrease in the stock affects a player's incentives less because the informational effect and the payoff externality work in opposite directions. Therefore, a large decrease in the stock is needed to make a player willing to stop if the learning technology improves.

[4]The part of the result that the equilibrium rate of learning is independent of the learning technology without payoff externalities is not unique to this paper. See e.g. Frick and Ishii (2020) and Laiho, Murto and Salmi (2022).

5.2 Capital investments in a competitive industry

In this section, we consider an important application with negative payoff externalities: entry to a market with unknown demand. In markets for new products and services, firms investing in productive capital face uncertainty about market demand. More firms entering the market means a larger capacity and hence more sales and faster learning in the market. Furthermore, firms impose a negative payoff externality on each other because the larger the production capacity, the fiercer the competition and lower the price.

5.2.1 Model of investments in a new market

Consider an investment problem where identical small firms choose when to make an irreversible investment to a market with unknown demand. Let the market use capital as the only input and assume a firm can invest in a unit of capital at lump sum cost c. Equivalently, cost of investment amounts to a perpetual stream of rental cost flow rc. We assume that all capital is in production so that we can normalize the production level to equal the current capital stock q_t.[5] Let the inverse demand be

$$p_t = x_t D_H (q_t) + (1 - x_t) D_L (q_t),$$

where x_t is the belief that the state of the world is high and D_ω is the inverse demand in state $\omega \in \{H, L\}$. We assume $D'_\omega (q) \leq 0$ for all $q \geq 0$.

Let the belief follow the diffusion process in Section 2.2 where the payoff flow is $\pi(q, \omega) = D_\omega (q) - rc$. The interpretation of the learning process is that new purchases generate information and that the level of purchases is pinned down by investment decisions in the past. An investing firm imposes a negative payoff externality on other firms as an investment pushes the price down.[6]

[5]This arises, for example, when the market is competitive and the marginal cost of using capacity is zero.

[6]Notice that there is no payoff externality on the social level when we take into account consumer surplus. However, to solve for the competitive outcome, we need to take into account the effect between the firms.

5.2.2 Competitive industry equilibrium

The equilibrium, which in this context we call the competitive industry equilibrium, can be easily computed using Proposition 1:

Corollary 1. *The investment cutoff in the competitive industry equilibrium is*

$$x^D(q) = \frac{-\beta(q)\left(D_L(q) - rc\right)}{(\beta(q) - 1)\, D_H(q) - \beta(q)\, D_L(q) + rc}.$$

The competitive industry equilibrium suffers from insufficient investments when there is a lot of uncertainty because the potential entry of other firms eliminates all gains from information generation.

Our model is essentially the same as in Leahy (1993), except in our model information generation is *endogenous*, which drives a wedge between competitive equilibrium and social optimum. Notice that if learning were *exogenous*, the competitive equilibrium corresponds to the social optimum, as established by Leahy (1993). With endogenous learning, this correspondence is no longer true because of the informational externality.

5.2.3 Taxes and subsidies

When knowing that the industry equilibrium is inefficient, a natural question is how different tax and subsidy policies could improve it. Consider a production subsidy $S(q)$ that is allowed to depend on the total capital/ production in the market. The subsidy essentially adds a new stock dependent term into the flow payoff of the firms. The new industry equilibrium would then be

$$x^D(q) = \frac{-\beta(q)\left(D_L(q) - rc + S(q)\right)}{(\beta(q) - 1)\, D_H(q) - \beta(q)\, D_L(q) + rc - S(q)}. \tag{6}$$

By changing S, the regulator could implement any cutoff rule, including the one that maximizes the social welfare, in the equilibrium.[7]

It is hardly surprising that subsidies improve welfare. However, when carefully interpreting (6), one notices that significant improvements can be made even with

[7] For how to solve the socially optimal expansion path under a similar kind of endogenous learning process, see Laiho, Murto and Salmi (2022).

subsidies that are *ex ante* budget balanced, i.e. involve both subsidies and taxes. Consider a decreasing scheme that starts with strictly positive $S(q)$ when q is low and has strictly negative $S(q)$ when q is high. Despite early investors understanding that their revenue will be taxed in the future if q increases, it does not affect their decision to invest today when q is low and there is a subsidy in place. They ignore the unfavorable change in the subsidy/ tax scheme similarly as they ignore the effect through the negative payoff externality. When there is a lot of initial uncertainty (when the initial belief is not too high), this kind of joint subsidy and tax scheme improves welfare because it speeds up information generation, which further speeds up the growth in the investments and hence shortens the time when the regulator needs to subsidize production.

6 Concluding remarks

This paper shows that even in an environment with informational and payoff externalities, forward-looking players can ignore the effect from the stopping decisions in the future. Therefore, we get an easy-to-interpret closed-form solution for the equilibrium.

The most surprising feature of the equilibrium is that more informative signals actually lead to slower learning in equilibrium under positive payoff externalities. Furthermore, for a large class of stopping payoffs, the equilibrium is not unique. In contrast, the equilibrium is always unique and a better learning technology always speeds up learning under negative payoff externalities.

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. In addition, small subsidies in 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 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 (x,q) if there exists a stopping time τ such that*

$$\mathbb{E}\left[\int_0^\tau e^{-rt}\left(x_t\pi(q_t,H)+(1-x_t)\,\pi(q_t,L)\right)dt\,|(x,q;Q)\right]<0. \tag{7}$$

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

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

$$\tau^*\in\arg\sup_\tau\mathbb{E}\left[\int_\tau^\infty e^{-r(t)}\left(x_t\pi(q_t,H)+(1-x_t)\,\pi(q_t,L)\right)dt\,|(x,q;Q)\right], \tag{8}$$

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

$$\begin{aligned}\mathbb{E}\Big[&\int_0^\infty e^{-rt}\left(x_t\pi(q_t,H)+(1-x_t)\,\pi(q_t,L)\right)dt\\ &-\int_0^\tau e^{-rt}\left(x_t\pi(q_t,H)+(1-x_t)\,\pi(q_t,L)\right)dt\,|(x,q;Q)\Big],\end{aligned}$$

we note that (8) is equivalent to

$$\tau^*_\theta\in\arg\inf_\tau\mathbb{E}\left[\int_0^\tau e^{-rt}\left(x_t\pi(q_t,H)+(1-x_t)\,\pi(q_t,L)\right)dt\,|(x,q;Q)\right],$$

and so $\tau^*_\theta>0$ only if (7) holds. □

A.2 Proof of Proposition 1

We show here that any equilibrium must satisfy the conditions in Proposition 1. 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

$$\mathbb{E}\left[\int_0^\tau e^{-rt}\left(x_t\pi(q,H) + (1-x_t)\,\pi(q,H)\right)dt\,|(q,x;q_t\equiv q)\right] \geq 0 \tag{9}$$

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

Contrary to the claim, assume that no one stops at (x,q). But as long as no one stops, the optimality conditions (7) and (9) are equivalent, and hence a player prefers stopping immediately, contradicting that no one stops at (x,q) 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

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

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

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

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 $\max_{s\in[q,q']}\pi(s,\omega)$. It can be solved in closed form:

$$\widetilde{x}_q\left(q'\right) \quad = \quad \frac{-\beta(q)\overline{\pi}\left(q,q';L\right)}{\left(\beta(q)-1\right)\overline{\pi}(q,q';H) - \beta(q)\overline{\pi}(q,q';L)}. \tag{10}$$

We see directly from (10) 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 $\max_{s\in[q,1]}\pi(s,\omega)$ 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 $(x, q(\theta))$ with $x < x_0(q(\theta))$ 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

$$x_n(q) := \widetilde{x}_q(q'), \text{ where } q' = \min\{s : \widetilde{x}_q(s) = x_{n-1}(s)\}. \tag{11}$$

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 $\max_{s\in[q,q']}\pi(s,\omega)$ 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 (11) 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 some 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 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 $x_n(q) = \bar{x}(q)$ 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.3 Existence of the maximal and minimal equilibria

Lemma 2. *Cutoff rule $\hat{x}$ in* (5) *characterizes an equilibrium in Definition 2.*

Proof. We need to verify that the best response to stock process $\hat{Q}_t$ implied by $\hat{x}$ 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

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

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 $\hat{Q}_t$ because the stock stays constant until hitting $\hat{x}(q)$:

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

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

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

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 $F\left(\hat{x}(q),q\right) <$

0 implies $F_x(q, \hat{x}(q)) \leq 0$. We will do that separately for potential cases $B(q) \geq 0$ and $B(q) < 0$.

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

$$A(q)\Phi(q, \hat{x}(q)) < -\frac{\hat{x}(q)\pi(q, H) + (1 - \hat{x}(q))\,\pi(q, L)}{r}. \tag{12}$$

But then,

$$\begin{aligned}
F_x(q, \hat{x}(q)) &= \frac{\pi(q, H) - \pi(q, L)}{r} + \frac{\beta(q) - \hat{x}(q)}{\hat{x}(q)\,(1 - \hat{x}(q))} A(q)\Phi(q, \hat{x}(q)) \\
&+ \frac{1 - \hat{x}(q) - \beta(q)}{\hat{x}(q)\,(1 - \hat{x}(q))} B(q)\Phi(\hat{x}(q), q) \\
&\leq \frac{\pi(q, H) - \pi(q, L)}{r} + \frac{\beta(q) - \hat{x}(q)}{\hat{x}(q)\,(1 - \hat{x}(q))} A(q)\Phi(q, \hat{x}(q)) \\
&< \frac{\pi(q, H) - \pi(q, L)}{r} \\
&- \frac{\beta(q) - \hat{x}(q)}{\hat{x}(q)\,(1 - \hat{x}(q))} \frac{\hat{x}(q)\pi(q, H) + (1 - \hat{x}(q))\,\pi(q, L)}{r}.
\end{aligned} \tag{13}$$

where the first inequality is implied by $B(q) \geq 0$ and $\beta(q) > 1$, and the second inequality is implied by (12). Noting that

$$(\beta(q) - 1)\,\bar{x}(q)\pi(q, H) + \beta(q)\,(1 - \bar{x}(q))\,\pi(q, L) = 0,$$

which is equivalent to

$$\frac{\pi(q, H) - \pi(q, L)}{r} = \frac{(\beta(q) - \bar{x}(q))\,(\bar{x}(q)\pi(q, H) + (1 - \bar{x}(q))\,\pi(q, L))}{r\bar{x}(q)\,(1 - \bar{x}(q))}$$

and so it follows from (13) that $F_x(q, \hat{x}(q)) < 0$ 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 $\lim_{x \downarrow 0} F(x, q) = B(q) \lim_{x \downarrow 0} \widetilde{\Phi}(q, x) = -\infty$ contradicting the boundary condition at $x = 0$). If $A(q) \geq 0$, then $F(q, x)$ is increasing in x, but then $F(q, x') \geq 0$ contradicts our running assumption $F(q, \bar{x}(q)) < 0$. 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 $F(q, x') = 0$ and $F(q, \bar{x}(q)) < 0$ together imply $F_x(q, \bar{x}(q)) < 0$.

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

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

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 $(q, \hat{x}(q))$.

□

Lemma 3. *Cutoff rule $\bar{x}$ in* (5) *characterizes an equilibrium in Definition 2.*

Proof. The only difference to the previous lemma and is that now we need to show that players do not want to stop between $\hat{x}$ and $\bar{x}$ when expecting that other players do not stop either. But when no one is stopping, the same argument as used in Step 1 of the proof of Lemma 2 carries over to the case where the cutoff is $\bar{x}$ instead of $\hat{x}$. Hence the proof for this case is identical. □

A.4 Proof of Proposition 2

The maximal equilibria exist by Lemma 2.

Define $\tilde{q}(q)$ to be such that the informativeness of stock $\tilde{q}(q)$ under learning technology σ' is the same as the the informativeness of stock q under learning technology σ: $\tilde{q}(q) := s : \lambda_\sigma(q) = \lambda_{\sigma'}(s)$. Similarly, let $\tilde{x} : [0, 1] \to [0, 1]$ be $\tilde{x}(\tilde{q}(q)) := \hat{x}_\sigma(q)$. The interpretation of $\tilde{x}$ is that if $\hat{x}_{\sigma'}(q) > \tilde{x}(q)$, then the agents require a higher belief to be willing to stop under σ' than under σ, given the same arrival rate of information.

Notice that the information flow under stopping cutoff $\tilde{x}$ and technology σ' is exactly the same as the information flow under stopping cutoff $\hat{x}_\sigma$ and technology σ. Next, use (4) to see that $\hat{x}_{\sigma'}(\tilde{q}(q)) > \hat{x}_\sigma(q)$ for all q: for all beliefs, the stock is lower under σ' than under σ. Now, because $\tilde{x}(\tilde{q}(q)) = \hat{x}_\sigma(q)$, we also have $\hat{x}_{\sigma'}(\tilde{q}(q)) > \tilde{x}(\tilde{q}(q))$, which implies that learning is slower under stopping cutoff $\hat{x}_{\sigma'}$ and technology σ' than under stopping cutoff $\hat{x}_\sigma$ and technology σ for all $x \leq \hat{x}_\sigma(1)$.

This means that agents who stop before the stock reaches $\tilde{q}(1)$ are worse off because they face both a slower increasing stock and slower arrival of information. As agents are homogeneous, this further implies that all agents are worse off because they are indifferent between stopping and waiting, and therefore the claim follows. □

A.5 Proof of Proposition 3

Part 1. Uniqueness: As pointed out in the main text, when $\bar{x}$ in (4) is monotone, Proposition 1 gives a unique candidate for the equilibrium, which is the minimal equilibrium in Definition 2. By Lemma 3, the candidate is indeed an equilibrium. Hence, our only task is to verify that (4) is monotone when $\pi(\cdot, L)$ and $\pi(\cdot, H)$ are everywhere weakly decreasing. To see that, take the derivative of (4) to get

$$\bar{x}'(q) = \frac{\beta'(q)\pi(q,L)\pi(q,H) + \beta(q)(\beta(q)-1)(\pi_q(q,L))\pi(q,H) - \pi_q(q,H))\pi(q,L))}{((\beta(q)-1)\pi(q,H) - \beta(q)\pi(q,L))^2},$$

which is always strictly positive if $\pi_q(q, L) \leq 0$ and $\pi_q(\cdot, H) \leq 0$ for all q as $\beta'(q) > 0$,$\pi(q, L) < 0$ and $\pi(q, H) > 0$.

Part 2. Learning: For the case where $\pi(q, L) = \pi_L$ and $\pi(q, H) = \pi_L$ are constants, the equilibrium stock of stopped players satisfies (from (4)) $\beta(q) = \frac{\bar{x}(q)\pi_H}{\bar{x}(q)\pi_H + (1-\bar{x}(q))\pi_L}$ for any value of parameter σ. Hence, if $(\bar{x}_\sigma(q), q)$ and $(\bar{x}_{\sigma'}(q'), q')$ such that $\bar{x}_\sigma(q) = \bar{x}_{\sigma'}(q')$ are points at the cutoffs for σ and σ' respectably, $\beta_\sigma(q)$ under σ must equal $\beta_{\sigma'}(q')$ under σ' and hence the flows of information are equal.[8]

Next, suppose $\pi_q(q, L), \pi_q(q, H) \leq 0$ and at least one with strict inequality. We use the same but reversed argument as in the proof of Proposition 2: suppose $\sigma' < \sigma$ and define $\tilde{q}(q) := \frac{\sigma'}{\sigma} q$ and $\tilde{x}(\tilde{q}(q)) := \bar{x}_\sigma(q)$. Use (4) to see that $\bar{x}_{\sigma'}(\tilde{q}(q)) < \bar{x}_\sigma(q)$ (this inequality is reversed in the proof of Proposition 2). Now, because

[8]To be precise, the result only holds for initial beliefs such that not everyone stops immediately.

$\tilde{x}(\tilde{q}(q)) = \hat{x}_\sigma(q)$, we also have $\bar{x}_{\sigma'}(\tilde{q}(q)) < \tilde{x}(\tilde{q}(q))$, which implies that learning is faster under cutoff $\bar{x}\sigma'$ and technology σ' than under cutoff $\bar{x}_\sigma$ and technology σ. $\square$

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.

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

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.

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*, 2022.

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.