EUROPEAN ECONOMIC REVIEW

Available online at www.sciencedirect.com
SciVerse ScienceDirect
North-Holland

European Economic Review 56 (2012) 1220–1240

Contents lists available at SciVerse ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

Rational exuberance

Paul Heidhues [a,b], Nicolas Melissas [c,*]

[a] ESMT, European School of Management and Technology, Schlossplatz 1, 10178 Berlin, Germany
[b] CEPR, United Kingdom
[c] Instituto Tecnológico Autónomo de México, CIE-ITAM, Camino a Santa Teresa 930, 10700 D.F., Mexico

ARTICLE INFO

Article history:
Received 6 September 2011
Accepted 30 May 2012
Available online 14 June 2012

JEL classification:
D62
D83

Keywords:
Information externality
Social learning
Strategic waiting
Delay
Information cascade

ABSTRACT

We study a two-player investment game with information externalities. Necessary and sufficient conditions for a unique symmetric switching equilibrium are provided. When public news indicates that the investment opportunity is very profitable, too many types are investing early and investments should therefore be taxed. Conversely, any positive investment tax is suboptimally high if the public information is sufficiently unfavorable.

1. Introduction

In 1996 the then Chairman of the Federal Reserve Board Alan Greenspan warned investors that they were in the grip of "irrational exuberance". This catchy wording instantly made headlines and initially caused a sell-off of stock around the world. Greenspan's concern that markets were overoptimistic was based, among other things, on the—in his words—"epic" multibillion-dollar investments of telecom companies into fiber-optic cable, which he believed would lead to significant losses for most companies involved. Greenspan, however, did not think that it was up to the Fed or to the federal government to intervene in the investment decisions of other agents. In particular, 30 months after his famous speech he came to the conclusion that

> "How do you draw the line between a healthy, exciting economic boom and a …bubble …?… After thinking a great deal about this, I decided that… the Fed would not second-guess "hundreds of thousands of informed investors." Instead the Fed would position itself to protect the economy in the event of a crash." (Greenspan, 2008, p. 200-1)

Greenspan thus thought that when investors are informed policy makers should not interfere in their investment decisions—a rationale that we believe was shared broadly among policy makers in many countries. His eloquent quote is a specific example of the common argument that governments should not interfere with better-informed business decisions.

* Corresponding author. Tel.: +55 52 46 28 40 00; fax: +55 52 56 28 40 58.
E-mail addresses: Paul.Heidhues@esmt.org (P. Heidhues), nicolas.melissas@itam.mx (N. Melissas).

0014-2921/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.euroecorev.2012.05.010

We question this rationale in a basic social-learning model and illustrate that even when rational investors are better informed than a policy maker, investments should be taxed when public news is sufficiently favorable.

Formally, we develop a two-player investment game to analyze this reasoning in more detail. There is a state of the world drawn from a normal distribution whose mean equals $\overline{\theta}$. If the state of the world is positive, both players should invest. If the state of the world is negative, no one should invest. Public news regarding the realization of the state of the world is captured by the prior mean $\overline{\theta}$. Public news can be favorable for different reasons. For example, there are favorable "stories" in the public domain about the realization of the fundamental or—perhaps as a consequence of past favorable public news—an investment boom occurred in the previous period. The public news is on average correct, but can sometimes be completely wrong. In particular, it is possible that the public news depicts the investment opportunity as a "golden" one, when in reality the state of the world is negative. To capture the idea that individual investors are better informed than the policy maker, we assume that both players get an additional normally distributed private signal about the realization of the state. Players combine the public news and their private information to compute the expected returns from investing. They then simultaneously decide whether or not to invest in period one. If a player invests, her payoff equals the state of the world. A player who has not invested in period one, observes the other player's period-one decision and then reconsiders her choice in period two. Payoffs from acting late are discounted, and a player who does not invest receives her outside option.

A rational player, by delaying her investment decision, can thus learn by observing the other player's first-period investment decision. The more optimistic a player is, the less willing she is to delay as returns from investing at a later point in time are discounted. In line with this intuition, we focus on switching equilibria in which a player invests whenever her expectation of the state of the world lies above a certain cutoff. Crucially, the value of waiting in our social-learning model depends on the other player's behavior. Whenever the other player's cutoff is sufficiently low, seeing her investing comes as no surprise. In this case an investment decision contains little information, which makes waiting for further information relatively undesirable. When the other player's cutoff becomes higher, she will invest less often. An investment decision then reveals that she has good private information which, in turn, makes waiting more desirable. Whenever this force is strong enough, our game is characterized by multiple symmetric switching equilibria.[1] Section 4 therefore investigates conditions under which a laissez-faire economy has a unique symmetric switching equilibrium. We find among other conditions that if players are sufficiently patient, our game possesses only one symmetric switching equilibrium. The symmetric switching equilibrium is also unique whenever the public news is either sufficiently good or bad.

Building on this characterization, Section 5 investigates the optimal symmetric investment cutoff that a social planner would want to implement. As mentioned above, we assume that the social planner has access only to the public information regarding the investment opportunity, and has no private knowledge that is concealed from the potential private investors. We first establish that the social planner does not want to distort second-period investment cutoffs; in the final period rational players want to and should invest whenever the expected state of the world is positive. We also characterize the optimal first-period cutoff. In particular, we show that if the public news is sufficiently favorable, it is optimal to raise the first-period investment cutoff. Roughly speaking, if the public news is favorable, both players are very likely to invest early in the laissez-faire economy, which implies that the informational content of an investment decision is low. Raising the cutoff increases the informational content of the first-period investment decision and leads to a greater positive externality. We then analyze the case in which the public news is unfavorable. Consider any cutoff higher than the one that prevails in a laissez-faire economy. We show that welfare with this cutoff is lower than in the laissez-faire economy whenever the prior mean is sufficiently low. In the limit as the prior mean goes to minus infinity, thus, either investments should be subsidized or the laissez-faire policy is optimal.

Section 6 establishes that the optimal cutoff can be implemented through a period-one investment tax (or subsidy). In particular, whenever the public news is sufficiently favorable, this is achieved through taxing first-period investment activity. The implementation, however, need not be unique even if the equilibrium of the laissez-faire economy is unique. Nevertheless, we establish that taxation is strictly optimal when the public news is sufficiently favorable by showing that positive tax rates exists for which both equilibrium remains unique and welfare is higher.

Section 2 outlines the existing social-learning literature. We highlight that one needs to extend the canonical social-learning model to address our policy question: The standard model with a binary state and signal space, for example, is characterized by multiple equilibria and the optimality of taxing or subsidizing investments in this model depends on—in our view—ad-hoc assumptions about equilibrium selection. Section 3 outlines the model. Section 4 analyzes the laissez-faire economy and derives a variety of conditions that ensure uniqueness of the symmetric switching equilibrium. In Section 5 we solve the social planner's problem; we discuss the implementation through state-dependent and time-varying taxation in Section 6. Finally, in Section 7 we discuss possible extensions and variants of our model, and some shortcomings of our approach.

[1] A switching strategy is symmetric if both players use the same cutoff. In our symmetric environment, it is natural to analyze these equilibria. In the working paper version of this paper (see Heidhues and Melissas, 2010), we prove that if the public news is sufficiently favorable, every switching equilibrium is symmetric.

2. Literature review

Social learning has been intensively studied when players are assumed to move in an exogenously specified order.[2] Hendricks and Kovenock (1989) were the first to analyze a game with information externalities in which players choose whether and when to drill. They were also the first to highlight the possibility of an informational cascade: If Player one did not drill at time one, this signals unfavorable private information. In turn, this induces Player two not to drill at time two. In equilibrium, both players may end up not drilling even though—had they pooled their private information—at least one player should have drilled at time one.

Although numerous papers analyze different waiting games,[3] to the best of our knowledge only Gossner and Melissas (2006), Levin and Peck (2008) and Doyle (2010) study optimal taxation in such a game. Furthermore, no paper in this literature analyzes the relationship between public information and optimal taxation. To fix ideas, motivate our modeling choices, highlight the novel contribution as well as compare our paper to existing ones, consider the following stylized social-learning setup: N players must decide whether or not to invest in a project, the cost of which is denoted by c. The returns of the project depend on the realized state of the world. If the state of the world is "high", the investment project yields a revenue equal to one. If the state of the world is "low", the project is assumed to yield zero revenues. Players receive a binary signal concerning the realized state of the world. Call a player who received a "low" signal a low-type player, while a high-type player received a "high" signal. After receiving their signals, players compute their posteriors. Let the posterior of a low-type player be denoted by μ_l while μ_h denotes that of a high type. Suppose $c<\mu_l<\mu_h$. This parameter configuration either occurs because the investment cost is low, or because of a "favorable" prior. At time one, players simultaneously decide whether to invest or wait. If a player waits, she observes how many other players invested at time one and takes a final investment decision at time two. If a player invests at time two, however, her payoff gets discounted.

This set-up is plagued by multiple-equilibria. In one equilibrium, which is analyzed in the seminal paper of Chamley and Gale (1994), high types randomize between investing and waiting while low types wait.[4] As high types do not internalize their information externalities, Gossner and Melissas (2006) have shown that in this equilibrium investments should be *subsidized*.[5] Recall that $c<\mu_l$. As a low-type player also faces a positive gain from investing, it is a best reply for her to invest (at time one) if she expects all other $N-1$ players to invest as well. Hence, there also exists another equilibrium in which all players invest at time one. In this equilibrium, some "wrong" types are investing (i.e. the low types) and thereby reduce the informational value of overall investment activity. Gossner and Melissas highlight that a social planner can then raise welfare by *taxing* investments.[6] The intuition should be clear: Through an appropriate investment tax, a social planner can reduce the profitability of investing such that only high types face a positive gain from investing. In that case, low types wait and benefit from the information externality. A model with a binary state and signal space is thus unable to generate unambiguous economic policy recommendations when either public information is very conducive to investing—i.e. when the prior mean is "favorable"—or when the investment cost is "low". Below, we derive such unambiguous policy recommendations by replacing the unrealistic assumption of binary returns of the investment project with the, in our view, more plausible assumption that the returns of investing are normally distributed. We show that policymakers should tax investments when the public sentiment is that the investment opportunities are highly beneficial.

Chamley (2004a) analyzes a two-player continuous (and unbounded) signals version of the binary return-to-investment model and establishes the existence of multiple symmetric switching equilibria. Furthermore, Chamley (2004b) establishes that a unique symmetric switching equilibrium exists if the discount factor is sufficiently high. He does not, however, investigate the optimal tax policy, nor does he provide other sufficient conditions which guarantee a unique equilibrium within the class of the symmetric switching strategies.

Recently, a global games approach (see Morris and Shin, 2003 for a survey) has been developed to overcome the multiplicity of equilibria in various coordination settings—often with the aim of deriving policy recommendations. This approach typically consists in enriching the type and state space and assuming that some "extreme" types possess a "dominant strategy". *Some* authors then derive sufficient conditions that guarantee a unique equilibrium outcome, while others reduce the set of equilibrium outcomes. To keep the analysis tractable, many authors[7] enrich the type and state space by working with normally distributed random variables. In this paper, we also work with normally distributed

[2] For an excellent overview, see Chamley (2004b).

[3] Waiting games have, among others, also been analyzed by Chamley and Gale (1994), Gul and Lundholm (1995), Zhang (1997), Choi (1997, Section 4), Aoyagi (1998), Caplin and Leahy (1998), Frisell (2003), and Gunay (2008).

[4] Strictly speaking, Chamley and Gale do not prove that it is optimal for low types to wait. Instead, they assume that low types do not possess an "investment option" and therefore cannot invest. But giving low types the option to invest does not destroy their equilibrium.

[5] Doyle (2010) introduces idiosyncratic investment costs in such a set-up and—following Chamley and Gale (1994)—assumes that low types cannot invest. High types invest if their investment costs lie below some critical level. In his model, the government cannot commit to a future tax/subsidy scheme. Players might thus postpone their investment plans in the hope of enjoying higher subsidies in the future. Doyle also finds that investments should be subsidized.

[6] Levin and Peck (2008) introduce an idiosyncratic investment cost in the binary-return-to-investment set-up and show that a small investment subsidy can reduce welfare for the same reason: It encourages some "wrong" types (i.e. those with bad private information about investment returns but with low investment costs) to invest—thereby reducing the informational value of overall investment activity. They do not provide conditions that guarantee the switching equilibrium is unique.

[7] See, among others, Angeletos et al. (2007), Angeletos and Werning (2006), Dasgupta (2007) and Morris and Shin (1999, 2000, 2002, 2003, 2004, 2005).

random variables when enriching the state and type space with the aim of predicting a unique symmetric switching equilibrium. This enables us to derive clear policy recommendations for the case in which public information is sufficiently favorable.

So far, only Dasgupta (2007) uses a global game approach in a social learning environment with the aim of predicting a unique symmetric switching equilibrium. He considers a two-period irreversible investment model with a continuum of players, exogenous observation noise, and positive network externalities. Dasgupta's paper focuses on how the ability to wait influences the extent of coordination failures in environments with positive network externalities and private information. He does not investigate the relationship between public information and optimal tax policy.

3. The model

Two risk-neutral players have the possibility to invest in a risky project. Players can invest in two periods. If Player i invests at time one, she gets a monetary payoff of $\theta-\tau$. Henceforth, we refer to $\theta \in \mathbb{R}$ as the state of the world and $\tau \in \mathbb{R}$ as a temporary investment tax ($\tau > 0$) or subsidy ($\tau < 0$). If Player i invests at time two, she gets $\delta\theta$, where $\delta \in (0,1)$ denotes the common discount factor. Investments are irreversible. The state of the world θ is randomly drawn from a normal distribution with mean $\overline{\theta}$ and variance σ_θ^2. In our model, the prior mean $\overline{\theta}$ captures public information that is available to policy makers and investors alike. It can, for example, be high because many "stories", "studies", or "expert opinions" are circulating that depict the investment opportunity as a "golden" one, or simply because past activity indicates a boom in the industry. To model potential investors as being better informed, we suppose that Player i receives a normally distributed private signal s_i concerning θ's realization. More precisely, we assume that $s_i = \theta + \epsilon_i$, where ϵ_i is independently drawn from a normal distribution with zero mean and variance σ_ϵ^2.

The timing is as follows: At time zero, the government sets the period-one investment tax τ. Thereafter, our waiting game starts with nature drawing the state of the world and all signals. After observing the investment tax τ and their private signals, players at period 1 simultaneously decide whether to invest or wait. At the beginning of period 2, players observe past investment choices. Any player who has not invested in period 1 then decides whether or not to invest in period 2. Finally, players receive their payoffs and the game ends.

Our stylized model has some noteworthy features. First, we presume that investments are irreversible so that it applies better to real investments than to financial ones in which transaction costs are considerably lower, and where market prices may contain more information than observed past investment behavior. Second, we assume that a player who waits observes the other player's investment decision but not her investment return. In many instances, investment returns realize a long time after the original investment decision has been made, and our model applies to the interim period. In the fiber-optic case mentioned in the Introduction, for example, the returns depend on many future developments in the telecommunications industry, which do not become publicly revealed after a firm's decision to lay out these cables. Third, we assume that the policy variable τ either increases or decreases the returns from investing by some lump-sum amount. One may thus think of τ as a tax credit, or as a tax on the investment good itself. Furthermore, we will argue below that if the optimal τ is positive and small, the optimal policy can also be implemented by taxing profits.

Below, we refer to the expected state of the world conditional on a player's signal as the player's time-one posterior mean, i.e. $\mu_i \equiv E(\theta|s_i)$. Throughout we exclusively focus on equilibria in symmetric switching strategies. Player i is said to follow a switching strategy if she invests at time one whenever her time-one posterior mean exceeds a critical threshold value μ^c and refrains from investing otherwise. A pair of strategies is a symmetric switching equilibrium if, given that Player j follows a switching strategy with critical threshold μ^*, one has (E1) it is strictly optimal for Player i to invest in period one if and only if $\mu_i > \mu^*$; and (E2) if Player i did not invest at time one, she does so at time two if and only if her expectation of θ (given μ_i and given Player j's time-one decision) is positive.[8] Below, equilibrium more generally refers to Bayesian equilibrium.[9] A type is said to be active at time two if she did not invest at time one.

4. Existence and uniqueness of switching equilibria in a laissez-faire economy

In this section, we characterize equilibrium cutoffs when the government does not intervene in the economy, i.e. when $\tau = 0$. Let μ^{LF} denote the first-period *equilibrium* cutoff in a laissez-faire economy. We first establish properties of the best response to a switching strategy. To do so, it is useful to consider the expected payoff difference between investing early and delaying the investment decision. Let $\Delta(\mu_i,\mu_j^c)$ denote the difference between the gain of investing in period 1 and the gain of waiting as a function of Player i's posterior mean μ_i under the assumption that Player j follows a switching strategy characterized by μ_j^c. Thus,

$$\Delta(\mu_i,\mu_j^c) = \mu_i - \delta \Pr(\mu_j > \mu_j^c|\mu_i) \max\{0, E(\theta|\mu_i,\mu_j > \mu_j^c)\} - \delta \Pr(\mu_j < \mu_j^c|\mu_i) \max\{0, E(\theta|\mu_i,\mu_j < \mu_j^c)\}. \quad (1)$$

[8] In Heidhues and Melissas (2010), we also prove that if the prior mean $\overline{\theta}$ is high enough, no asymmetric equilibrium in switching strategies exists.

[9] In our model players with sufficiently high (low) signals strictly prefer to invest (wait) at time one, independent of the other player's strategy. Hence, there are no off-the-equilibrium-path observations and players can always apply Bayes's rule so that any Bayesian equilibrium is consistent and sequentially rational.

If $\Delta(\cdot) > 0$ Player i prefers to invest, while if $\Delta(\cdot) < 0$ she prefers to wait. We first observe that a player who is more optimistic regarding the state of the world has a bigger incentive to invest early. Formally

Lemma 1. *A player's incentive to invest early increases in her time-one posterior mean, i.e.*

$$\frac{\partial \Delta(\mu_i, \mu_j^c)}{\partial \mu_i} > 0 \quad \forall \mu_j^c.$$

Lemma 1 states a common property of waiting games studied in the literature.[10] To understand the intuition behind the lemma, suppose Player i possesses a "very high" time-one posterior mean μ_i. She then has no incentive to wait. For she would always invest at time two, even if she observes the other player waiting. Thus, if her time-one posterior mean increases from "very" to "extremely" high, this leads to an increase in her opportunity cost of delaying the investment and to an increase in the value of $\Delta(\cdot)$. If her time-one posterior mean is "moderately" high or "simply" high instead, she only refrains from investing at time two if she observes the other player waiting. If her time-one posterior mean increases from "moderately" to "simply" high, this leads to a decrease in the probability that she ascribes to the event "Player j will wait". Hence, the higher a player's time-one posterior mean, the lower her incentive to wait.

Lemma 1 implies that there exists a unique time-one posterior mean at which Player i is indifferent between investing and waiting given that Player j follows a switching strategy characterized by μ_j^c. Formally, i's cutoff $\mu_i^I(\mu_j^c)$ is implicitly defined through the equation $\Delta(\mu_i^I, \mu_j^c) = 0$. (The superscript "$I$" stands for "indifferent".)

Suppose $\mu_i > 0$ and that i expects j to always wait so that $\mu_j^c = \infty$. Then, of course, j's waiting decision bears no informational content. Thus, the difference between the gain of investing early and the gain of waiting and investing late is $\Delta(\mu_i, \infty) = (1-\delta)\mu_i > 0$. On the other hand, if $\mu_i < 0$ and Player i expects Player j to always wait, Player i prefers not to invest. Hence, in this case i invests in the first period whenever her time-one posterior mean is greater than zero and refrains otherwise. Using a similar reasoning, if Player i expects j to always invest, j's investment decision has no informational content and thus $\mu_i^I(-\infty) = \mu_i^I(\infty) = 0$. Furthermore, mere inspection of Eq. (1) reveals that i's best response cutoff μ_i^I is continuous in μ_j^c. Lemma 1 thus implies that the cutoff μ^{LF} characterizes a symmetric switching equilibrium if and only if $\mu_i^I(\mu^{LF}) = \mu^{LF}$, or equivalently, $\Delta(\mu^{LF}, \mu^{LF}) = 0$.[11] Graphically, μ^{LF} is the point at which $\mu_i^I(\mu_j^c)$ crosses the 45° line. Since $\mu_i^I(-\infty) = \mu_i^I(\infty) = 0$, and since μ_i^I is continuous in μ_j^c, a symmetric switching equilibrium exists.

We now investigate conditions that guarantee uniqueness. First, observe that a player who is indifferent between investing and waiting must face a positive gain of investing. This implies that $\mu^{LF} > 0$. Because $\mu^{LF} < E(\theta | \mu_i = \mu^{LF},\ \mu_j > \mu^{LF})$ a player with time-one posterior mean μ^{LF} invests at time two after observing her fellow player investing. We next argue that if $\mu_i = \mu^{LF}$, Player i does not invest in period two after observing that Player j waited, i.e. $E(\theta | \mu_i = \mu^{LF}, \mu_j < \mu^{LF}) < 0$. Given that j follows a switching strategy, observing him investing rather than waiting must make i more optimistic. Hence, if i wants to invest after having observed that j waited, she must also want to invest after having observed that j invested. In such a case she invests at time two independent of j's time-one action. Her expected gain of waiting therefore is $\delta\mu^{LF}$. She is then better off, however, investing at time one and receiving an expected payoff of μ^{LF}.

Given this observation, we say that Player i receives "good news" when she observes j investing. Using that a cutoff type invests in period two only when receiving good news, $\Delta(\mu^{LF}, \mu^{LF})$ simplifies to

$$\Delta(\mu^{LF}, \mu^{LF}) = \mu^{LF} - \delta \Pr(\mu_j > \mu^{LF} | \mu_i = \mu^{LF}) E(\theta | \mu_i = \mu^{LF}, \mu_j > \mu^{LF}) = 0. \tag{2}$$

Our analysis below makes use of some intuitive and well-known properties of the normal distribution (see the Appendix for proofs). First, Player i's first-period posterior mean μ_i is computed as

$$\mu_i = \alpha s_i + (1-\alpha)\overline{\theta} \quad \text{where } \alpha = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\epsilon^2}.$$

In words, μ_i is a weighted average between her private signal s_i and the prior mean $\overline{\theta}$. The more precise the prior information—i.e. the lower σ_θ^2 -the more weight Player i puts on the prior mean and the less weight she puts on her signal. Conversely, the more precise her private information—i.e. the lower σ_ϵ^2 -the more she trusts her signal as opposed to the prior mean. In particular, this implies that if the variance of the prior is infinite, or if the variance of her signal is zero, her posterior mean is equal to her signal.

Second, Player i's expectation of Player j's posterior mean μ_j is computed as

$$E(\mu_j | \mu_i) = \alpha\mu_i + (1-\alpha)\overline{\theta}.$$

Intuitively, Player i believes that j's signal is distributed around her best guess of the true state of the world—i.e. her posterior mean. Player i, however, also realizes that Player j's posterior mean is a weighted average between j's signal and the prior mean, and therefore is likely to lie between i's posterior and the prior mean. Based on this, a key fact we use below is that if Player i's posterior mean increases by one unit, her expectation about j's posterior mean increases by *less* than one unit. Hence, for example, the further her posterior mean lies above the prior mean, the more likely i thinks that j

[10] See for example Hendricks and Kovenock (1989) and Chamley (2004b, Lemma 6.1, p. 124).

[11] It follows from Lemma 1 that equilibrium condition E1 (see Section 3) is satisfied when $\Delta(\mu^{LF}, \mu^{LF}) = 0$. Equilibrium condition E2 is also satisfied because Eq. (1) prescribes Player i to make an optimal time-two choice.

is more pessimistic than herself. Closely related, if the signal is (nearly) perfect—i.e. the variance of the signal is (close to) zero—both players possess (almost) the same posterior. In that case Player *i* believes that she always (almost) lies in the "center of the world"—i.e. independent of her posterior there is a 50% chance of *j* being more optimistic than herself. A similar argument also applies with a completely uninformative prior—i.e. when the variance of the prior is infinite. In this case *j* puts zero weight on the prior mean when computing his posterior. As *i* believes *j*'s signal to be distributed around her posterior mean, she also always believes that she lies in the center of the world.

Third, conditional on having the cutoff posterior mean μ^{LF}, the probability that *j*'s posterior mean is greater than the cutoff is

$$\Pr(\mu_j > \mu^{LF} | \mu_i = \mu^{LF}) = 1 - F(\kappa_1(\mu^{LF} - \overline{\theta})), \tag{3}$$

where F denotes the cumulative distribution function of the standard normal and where κ_1 is a positive constant depending on the prior and signal variances. It follows from our second observation as well as the formula above that an increase in $\mu^{LF} - \overline{\theta}$ reduces the probability of *j* being more optimistic than the cutoff type *i*.

Fourth, we are interested in the cutoff type's expectation about the state of the world when waiting and receiving good news. In a symmetric switching equilibrium, Player *i*'s expectation will be based on her own signal, the prior mean, and the fact that *j* invested and thus had a first-period posterior mean above the common cutoff μ^{LF}. Here our distributional assumptions allow us to use known properties of the truncated normal distribution. Formally, in the Appendix we establish that

$$E(\theta | \mu_i = \mu^{LF}, \mu_j > \mu^{LF}) = \mu^{LF} + \kappa_2 h(\kappa_1(\mu^{LF} - \overline{\theta})), \tag{4}$$

where κ_2 is a positive constant which (just as κ_1) depends on σ_θ^2 and σ_ϵ^2, and where h represents the hazard rate of the standard normal distribution. Recall that the hazard rate h is defined here as: $h(x) \equiv f(x)/(1-F(x))$.[12] Recall also that the hazard rate of a normal distribution is increasing in x. Intuitively, Player *i*'s second-period expectation is the first-period expectation about the state of the world plus an upward shift that depends on the cutoff, the prior mean, as well as—through the constants—the variance of signals and the prior. We have seen above that the cutoff type's probability of getting good news decreases in the cutoff μ^{LF}. The above formula reveals that the impact of good news is also higher for higher cutoffs. Formally, this follows from the fact that the hazard rate of the standard normal distribution is increasing and thus, the upward shift is greater. The statistical intuition is as follows: Player *i*'s belief of Player *j*'s first-period posterior mean is normally distributed with—as we observed above—a mean that lies between *i*'s posterior mean and the prior mean. As the cutoff increases, the expectation of Player *j*'s posterior mean increases by less than the cutoff. Thus, if *j* invests he reveals that he lies in a higher quantile of this distribution. Since the expectation of a left-truncated normally distributed variable is increasing in the truncation point, the higher the cutoff, the better the news for the cutoff type when observing *j* investing. Consider now the case in which the variance of the prior goes to infinity. As explained above, Player *i* then believes that she is in the "center of the world", i.e. there is, independent of her posterior, a 50% chance that *j* possesses a higher posterior than herself. This implies that the upward shift does not depend on the cutoff μ^{LF}. Mathematically, in the Appendix we show that κ_1 tends to zero as the variance of the prior goes to infinity, while κ_2 converges to a positive constant. Thus in this special case the upward shift is independent of where the cutoff lies.

Using Eqs. (3) and (4), we rewrite the equilibrium condition (2) as

$$\underbrace{\mu^{LF}}_{\text{Gain of investing}} = \underbrace{\delta \underbrace{[1 - F(\kappa_1(\mu^{LF} - \overline{\theta}))]}_{\text{Prob of good news}} \left[\mu^{LF} + \underbrace{\kappa_2 h(\kappa_1(\mu^{LF} - \overline{\theta}))}_{\text{Upward shift in beliefs}} \right]}_{\text{Discounted gain of waiting}}.$$

As μ^{LF} increases, there are two countervailing forces affecting the gain of waiting. On the one hand, the probability of getting good news decreases. On the other hand, as μ^{LF} increases receiving good news leads to a greater upward shift in beliefs. Indeed, the expected upwards shift $[1-F(\cdot)]\kappa_2 h(\cdot) = \kappa_2 f(\cdot)$ and therefore is non-monotone and unimodal. Rearranging by moving the linear terms in μ^{LF} to the left-hand side and rewriting, yields

$$\mu^{LF} = \kappa_2 \mathcal{X}(\kappa_1(\mu^{LF} - \overline{\theta})) \quad \text{where } \mathcal{X}(\cdot) \equiv \frac{\delta f(\cdot)}{1 - \delta(1 - F(\cdot))}. \tag{5}$$

The left-hand side is linear in μ. The right-hand side is positive and goes to zero as μ goes to plus or minus infinity. Furthermore, Lemma 3 in the Appendix formally establishes many properties of our $\mathcal{X}$-function that are intuitive given that its numerator is the p.d.f. of a normally distributed random variable. In particular, we prove that $\mathcal{X}$ is unimodal, convex and increasing up to a critical value μ^m and thereafter concave and increasing up to its mode $\hat{\mu}$. It is also easy to see

[12] Throughout the paper, f denotes the p.d.f. of a standard normal distribution.

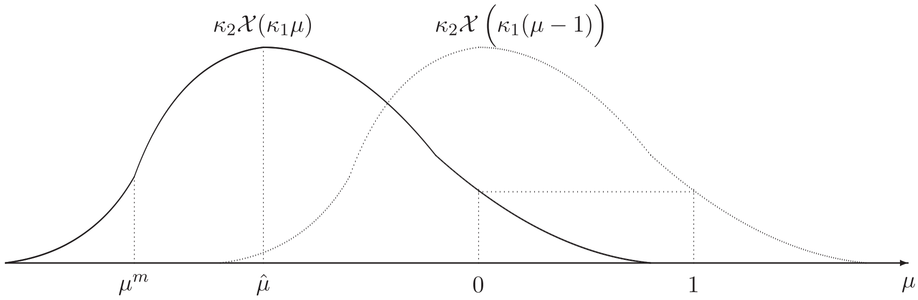

Fig. 1. Shape of $\kappa_2\mathcal{X}(\kappa_1(\mu-\overline{\theta}))$ for $\overline{\theta}=0$ and $\overline{\theta}=1$.

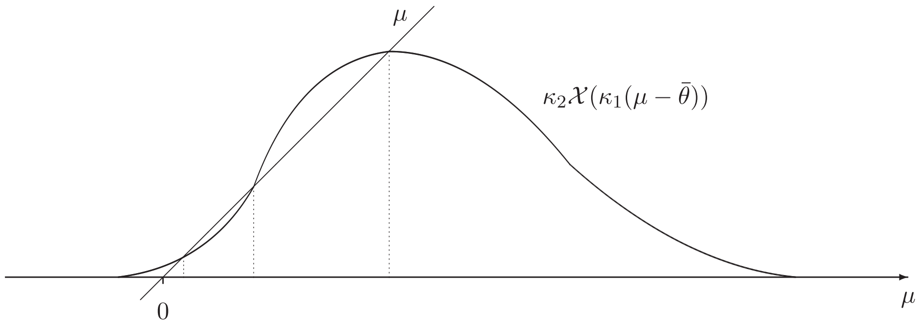

Fig. 2. Three different equilibria.

that a unit increase in $\overline{\theta}$ leads to a translation to the right of $\mathcal{X}$ by one unit. This property is easiest to check when $\overline{\theta}$ increases from zero to one. In that case $\mathcal{X}(\kappa_1(0-0))=\mathcal{X}(\kappa_1(1-1))$ as illustrated in Fig. 1.[13]

As Fig. 2 illustrates, whenever the slope of $\kappa_2\mathcal{X}$ is greater than one, multiple symmetric switching equilibria can arise. Intuitively, a low cutoff can be self-fulfilling since if μ^{LF} is low an agent's expected upward shift is also low; this makes waiting unattractive and thus induces players with low posterior means to invest early. If agents, however, expect a higher cutoff the expected upward shift can be higher, making waiting in turn more attractive.

In the Appendix, we show that—for all -μ the slope of $\kappa_2\mathcal{X}$ is less than one if and only if:

$$\left(\frac{\sigma_\theta}{\sigma_\epsilon}\right)^2 \geq \frac{1}{2}[\mathcal{X}'(\eta)-1] \quad \forall \eta \in \mathbb{R}. \tag{6}$$

Recall from the above discussion and Fig. 1 that the *maximal slope* of $\mathcal{X}$ depends only on the discount factor and not on other exogenous parameters. In the limit when players are perfectly impatient ($\delta=0$), for example, $\mathcal{X}(\cdot)=0$ everywhere and, hence, the maximal slope of $\mathcal{X}$ also equals zero. Furthermore, in Lemma 3, which can be found in the Appendix, we prove that the maximal slope of $\mathcal{X}$ tends to infinity as the discount factor δ approaches one. Inequality (6) thus implies that there exists a unique switching equilibrium if the variance of the public news σ_θ^2 is sufficiently high, or if either the variance of the private signal σ_ϵ^2 or the discount factor δ are sufficiently low.

Those three sufficient conditions are intuitive. Recall that if the variance of the prior is (infinitely) large, *i* believes *j*'s posterior mean to be equally likely to lie above or below hers—independent of her posterior mean. The cutoff type's expected upward shift in this case is thus independent of her posterior mean. Hence, as the variance of the prior becomes large, the expected upward shift tends towards a constant and therefore the slope of $\kappa_2\mathcal{X}$ tends to zero. Thus, for a high enough variance of the prior, there exists a unique symmetric switching equilibrium. Similarly, as the agent's signal becomes infinitely precise (i.e. as $\sigma_\epsilon^2\to 0$) she believes that she is in the center of the world and the expected upward shift tends to a constant. Thus, the symmetric switching equilibrium is also unique in this case. Furthermore, if the future becomes heavily discounted the gain of waiting and the slope of $\kappa_2\mathcal{X}$ tend to zero, and thus the unique symmetric equilibrium cutoff approaches zero in this case.

Of course, even if the maximal slope of $\kappa_2\mathcal{X}$ is greater than one, the symmetric switching equilibrium may be unique. For example, if the gain of investing μ crosses the function $\kappa_2\mathcal{X}$ in its right tail, i.e. when its slope is negative, the switching equilibrium is unique. Similarly, if it crosses $\kappa_2\mathcal{X}$ where its slope is positive but sufficiently low, the symmetric switching equilibrium will be unique. We have argued above that a unit increase in $\overline{\theta}$ leads to a translation by one unit to the right of $\kappa_2\mathcal{X}$. Hence, one can reduce $\overline{\theta}$ until the equilibrium condition (5) is satisfied in the decreasing part of $\kappa_2\mathcal{X}$. Similarly, we can increase $\overline{\theta}$ until μ cuts $\kappa_2\mathcal{X}(\cdot)$ "far enough" in its left tail. Thus for sufficiently high or sufficiently low $\overline{\theta}$, there exists a

[13] At the risk of stating the obvious, $\mathcal{X}(\kappa_1(0-0))$ denotes the value of $\mathcal{X}(\kappa_1(\mu-\overline{\theta}))$ when $\mu=\overline{\theta}=0$.

unique symmetric switching equilibrium. If the prior mean $\overline{\theta}$ even becomes arbitrarily large, μ cuts $\kappa_2\mathcal{X}(\cdot)$ when μ is arbitrarily close to zero. (This is illustrated in Fig. 3.) Similarly, when the prior mean becomes arbitrarily negative, μ^{LF} also tends to zero.

Finally, we argue that the symmetric switching equilibrium is unique if players are very patient. To see the intuition, consider the limit case in which $\delta=1$. In that case $\mathcal{X}(\kappa_1(\mu-\overline{\theta}))$ simplifies to the reverse hazard rate of the standard normal distribution.[14] Hence in the limit case of perfectly patient agents, μ and $\mathcal{X}(\kappa_1(\mu-\overline{\theta}))$ cross each other once, as μ is increasing and $\mathcal{X}(\kappa_1(\mu-\overline{\theta}))$ decreasing in μ. In the Appendix, we extend this argument by showing that $\mathcal{X}(\kappa_1(\mu-\overline{\theta}))$ is decreasing in the relevant range if players are sufficiently patient. The following proposition summarizes the above discussion:

Proposition 1. *There exists a unique symmetric switching equilibrium if the parameters* $(\sigma_\theta^2,\sigma_\epsilon^2,\delta)$ *satisfy Inequality* (6). *If Inequality* (6) *is violated, there exist values of* $\overline{\theta}$ *that support multiple symmetric switching equilibria. Multiplicity, however, only arises for "intermediate" values of* $\overline{\theta}$. *In particular, for any vector* $(\sigma_\theta^2,\sigma_\epsilon^2,\delta)$ *there exists a critical value* $\overline{\theta}_u<\infty$ *such that our game has a unique symmetric switching equilibrium if* $\overline{\theta}\geq\overline{\theta}_u$, *or if* $\overline{\theta}\leq 0$. *Furthermore, for any vector* $(\sigma_\theta^2,\sigma_\epsilon^2,\overline{\theta})$, *there exists a critical value* $\overline{\delta}<1$ *such that our game has a unique symmetric switching equilibrium if players are sufficiently patient* (*i.e. if* $\delta\geq\overline{\delta}$). *Finally,* $\lim_{\overline{\theta}\to-\infty}\mu^{LF}=\lim_{\overline{\theta}\to\infty}\mu^{LF}=0$.

Dasgupta (2007) also analyzes a dynamic game with social learning (see our literature review for more details) and establishes uniqueness of the symmetric switching equilibrium if $\sigma_\theta/\sigma_\epsilon$ is high enough. In our two-player model without network externalities, we identify additional conditions that yield uniqueness. Chamley (2004b) analyzes a similar set-up as ours and shows that the switching equilibrium is unique if the discount factor is sufficiently high. As we work with normally distributed random variables, we were able to identify additional sufficient conditions.

In many applications, players observe each other frequently and can relatively quickly react upon observing a player's investment decision. Hence, in such situations the discount factor is high and our model yields a unique symmetric switching equilibrium. Similarly, "boom" times are typically characterized by "stories" that depict some investment opportunities as "golden" ones. Hence, in such a situation our model also yields a unique symmetric switching equilibrium even if players observe each others actions only infrequently.

5. The social planner's problem

In this section, we consider a social planner who chooses three cutoff levels: a time-one cutoff μ^c above which a player invests at time one if (and only if) her time-one posterior mean exceeds it; a time-two cutoff $\underline{\mu}^1$ for the case in which her fellow player invested at time-one; and a time-two cutoff $\underline{\mu}^0$ for the case in which her fellow player did not invest at time one. At time two, a player invests if (and only if) her time-one posterior mean lies above the relevant time-two cutoff. The social planner aims to maximize expected welfare W, which is defined as

$$\begin{aligned}W\equiv&\int\Pr(\mu_i>\mu^c,\mu_j>\mu^c|\theta)2\theta f\left(\frac{\theta-\overline{\theta}}{\sigma_\theta}\right)d\theta\\&+2\int\Pr(\mu_i>\mu^c,\mu_j\in[\min\{\underline{\mu}^1,\mu^c\},\mu^c]|\theta)(1+\delta)\theta f\left(\frac{\theta-\overline{\theta}}{\sigma_\theta}\right)d\theta\\&+2\int\Pr(\mu_i>\mu^c,\mu_j<\min\{\underline{\mu}^1,\mu^c\}|\theta)\theta f\left(\frac{\theta-\overline{\theta}}{\sigma_\theta}\right)d\theta\\&+\delta\int\Pr(\mu_i\in[\min\{\underline{\mu}^0,\mu^c\},\mu^c],\mu_j\in[\min\{\underline{\mu}^0,\mu^c\},\mu^c]|\theta)2\theta f\left(\frac{\theta-\overline{\theta}}{\sigma_\theta}\right)d\theta\\&+2\delta\int\Pr(\mu_i\in[\min\{\underline{\mu}^0,\mu^c\},\mu^c],\mu_j<\min\{\underline{\mu}^0,\mu^c\}|\theta)\theta f\left(\frac{\theta-\overline{\theta}}{\sigma_\theta}\right)d\theta.\end{aligned}\quad(7)$$

The first integral captures the case in which both players invest at time one in which case welfare is equal to 2θ. The second integral captures the case in which one player invests at time one and thereby induces the other player to invest at time two. The player who invests at time one gets θ, the one who invests at time two $\delta\theta$. Observe that the second integral equals zero if $\mu^c\leq\underline{\mu}^1$. The third integral captures the case in which only one player invests at time one. The other player's (time-one) posterior is too low and she therefore never invests. In the fourth integral both players invest at time two. Welfare then equals $2\delta\theta$. Observe also that the fourth integral equals zero if $\mu^c\leq\underline{\mu}^0$. Finally, in the last integral only one player invests at time two. One can think of $\frac{1}{2}W\equiv U$ as the expected utility of a representative player in our model.

[14] Recall that the reverse hazard rate of a standard normal distribution, r, is defined as: $r\equiv f(x)/F(x)$. As is well known, r is decreasing in x.

It is useful to rewrite U in terms of the ex ante distribution of a player's posterior. From the planner's point of view, μ_i is normally distributed with mean $\overline{\theta}$ and with some variance denoted by σ_μ^2.[15] Using this, in the Appendix we show that the expected utility of the representative player can be rewritten as

$$U = \int_{\mu^c}^{\infty} \mu_i f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i + \delta \int_{\min\{\underline{\mu}^0, \mu^c\}}^{\mu^c} \Pr(\mu_j < \mu^c | \mu_i) E(\theta | \mu_i, \mu_j < \mu^c) f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i.$$

$$+ \delta \int_{\min\{\underline{\mu}^1, \mu^c\}}^{\mu^c} \Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c) f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i \tag{8}$$

Eq. (8) *is intuitive*: The first integral represents the weighted expected utility of all types that invest at time one. Our second integral captures Player i's payoff in the case in which both players waited at time one; in this case a player invests if her time-one posterior mean lies in the interval $[\underline{\mu}^0, \mu^c]$ and she gets an expected payoff of $\delta E(\theta | \mu_i, \mu_j < \mu^c)$ when investing. The third integral captures the case in which Player i does not invest in period 1 but does so in period 2 when receiving good news ($\mu_i \in [\underline{\mu}^1, \mu^c]$), which in turn happens with probability $\Pr(\mu_j > \mu^c | \mu_i)$.

We are now ready to analyze the optimal time-two cutoffs. Clearly, welfare cannot be raised by obliging a player to forego a profitable investment opportunity at time two or by forcing a rational player to invest in the second period when she believes this to be unprofitable. Hence, welfare-maximization implies that the critical investment type when getting good news ($\underline{\mu}^1$) is implicitly defined by setting the expected second-period investment return to zero (i.e. through $E(\theta | \mu_i = \underline{\mu}^1, \mu_j > \mu^c) = 0$). Similarly, the critical investment type when getting bad news ($\underline{\mu}^0$) is implicitly defined through $E(\theta | \mu_i = \underline{\mu}^0, \mu_j < \mu^c) = 0$.[16] With a slight abuse of notation, $\underline{\mu}^0$ and $\underline{\mu}^1$ will henceforth denote the *optimal* time-two cutoffs. Note that $\underline{\mu}^0$ and $\underline{\mu}^1$ depend on the time-one cutoff μ^c.

For which time-one cutoff μ^c is it optimal to have some active players invest in period two? Consider the cutoff type μ^c. Suppose that the expected state of the world is negative for this cutoff type even when receiving good news, i.e. that $E(\theta | \mu_i = \mu^c, \mu_j > \mu^c) \le 0$. As we established above, it is then optimal for the cutoff type to refrain from investing in period two when getting good news—and because the cutoff type is the most optimistic type who waits, no other type wants to invest in period two. Furthermore, as the expected state of the world is even lower when getting bad news, no type will want to invest in period 2 whenever $E(\theta | \mu_i = \mu^c, \mu_j > \mu^c) \le 0$. Lemma 2 in the Appendix proves that there exists a unique lower bound $\underline{\mu}$ such that $E(\theta | \mu_i = \underline{\mu}, \mu_j > \underline{\mu}) = 0$, which implies that there is no time-two investments whenever $\mu^c < \underline{\mu}$. Because players become more optimistic when getting good news, it is obvious that the lower bound is negative, i.e. $\underline{\mu} < 0$.

When $\mu^c > \underline{\mu}$, the social planner instructs the cutoff type who receives good news to invest as her expectation of the realized state of the world is then positive ($E(\theta | \mu^c, \mu_j > \mu^c) > 0$). If μ^c, however, is close to $\underline{\mu}$, the expected state of the world when getting bad news is still negative for the cutoff type. Hence, no active type will be instructed to invest when getting bad news. When μ^c is high enough, even when getting bad news the expected state of the world is positive ($E(\theta | \mu^c, \mu_j < \mu^c) > 0$). Active types close enough to μ^c will then optimally invest at time two when getting bad news. As these types are even more optimistic when getting good news, they will also invest in that case. Lemma 2 in the Appendix proves the existence of a unique upper bound $\overline{\mu}$ such that $E(\theta | \mu_i = \overline{\mu}, \mu_j < \overline{\mu}) = 0$, which implies that active types will refrain from investing when getting bad news if and only if $\mu^c < \overline{\mu}$. Since bad news makes a player more pessimistic, observe that $\overline{\mu} > 0$.

To summarize, if the social planner implements a very low cutoff (i.e. if $\mu^c < \underline{\mu}$), no one invests at time two and thus the utility of the representative agent is

$$\forall \mu^c \le \underline{\mu}, \quad U = \int_{\mu^c}^{\infty} \mu_i f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i.$$

If $\mu_i \in [\underline{\mu}, \overline{\mu}]$, on the other hand, some active types invest at time two if they receive good news but everyone refrains from investing when getting bad news. Hence

$$\forall \mu^c \in [\underline{\mu}, \overline{\mu}], \quad U = \int_{\mu^c}^{\infty} \mu_i f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i + \delta \int_{\underline{\mu}^1}^{\mu^c} \Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c) f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i. \tag{9}$$

For $\mu^c > \overline{\mu}$, some active types will invest when getting either bad or good news while others invest only when getting good news.

[15] In Section 4, we argued that i's posterior mean μ_i is a weighted average between her signal and the prior mean. Formally, $\mu_i = \alpha s_i + (1-\alpha)\overline{\theta}$ (where $\alpha \in [0,1]$ depends on the prior and signal variances). By assumption $s_i = \theta + \epsilon_i$, where $\theta \sim N(\overline{\theta}, \sigma_\theta^2)$ and $\epsilon_i \sim N(0, \sigma_\epsilon^2)$. As ϵ_i is independent of θ, from the planner's point of view $s_i \sim N(\overline{\theta}, \sigma_\theta^2 + \sigma_\epsilon^2)$. Hence, μ_i is the sum of a normally distributed random variable (multiplied by α) with mean $\overline{\theta}$ and a constant (i.e. $(1-\alpha)\overline{\theta}$). This implies that $\mu_i \sim N(\overline{\theta}, \sigma_\mu^2)$, where $\sigma_\mu^2 = \alpha^2(\sigma_\theta^2 + \sigma_\epsilon^2)$.

[16] Lemma 2 in the Appendix formally establishes that for any first-period cutoff μ^c, there exist unique second-period cutoffs $\underline{\mu}^0$ and $\underline{\mu}^1$, and that the expectations $E(\theta | \mu_i, \mu_j > \mu^c)$ and $E(\theta | \mu_i, \mu_j < \mu^c)$ are increasing in μ_i. Hence, Player i should invest at time two if (and only if) her time-one posterior mean lies above the relevant cutoff.

Let μ^{SP} denote the first-period cutoff that maximizes U. We now argue that $\mu^{SP} > \underline{\mu}$. Suppose otherwise, i.e. that $\mu^{SP} \le \underline{\mu} < 0$. In this case the expected utility of the representative player equals

$$\int_{\mu^{SP}}^{\infty} \mu_i f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i.$$

Because $\mu^{SP} < 0$, the social planner can raise welfare by setting $\mu^{SP} = 0$—a contradiction. Thus $\mu^{SP} > \underline{\mu}$. This result allows us to prove that the social planner will choose a higher than the laissez-faire cutoff when the prior mean is sufficiently high.

Recall from our previous section that, in a Bayesian equilibrium without taxes, no active type invests at time two if no one invested at time one, i.e. $\mu^{LF} < \overline{\mu}$.[17] Furthermore, $\mu^{LF} > 0$ as players with a non-positive posterior mean prefer to wait. Hence, $\mu^{LF} \in (\underline{\mu}, \overline{\mu})$. It follows from Eq. (9) that $\forall \mu^c \in [\underline{\mu}, \mu^{LF}]$

$$\frac{dU}{d\mu^c} = -[\mu^c - \delta \Pr(\mu_j > \mu^c | \mu^c) E(\theta | \mu^c, \mu_j > \mu^c)] f\left(\frac{\mu^c - \overline{\theta}}{\sigma_\mu}\right) - \delta \frac{d\underline{\mu}^1}{d\mu^c} \Pr(\mu_j > \mu^c | \underline{\mu}) \underbrace{E(\theta | \underline{\mu}^1, \mu_j > \mu^c)}_{=0} f\left(\frac{\underline{\mu}^1 - \overline{\theta}}{\sigma_\mu}\right)$$
$$+ \delta \underbrace{\int_{\underline{\mu}^1}^{\mu^c} \frac{\partial}{\partial \mu^c} (\Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c)) f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i}_{\text{Weighted change in the inframarginal types' gain of waiting}}. \tag{10}$$

For any given first-period cutoff, recall that the socially optimal cutoff $\underline{\mu}^1$ is implicitly defined through $E(\theta | \mu_i = \underline{\mu}^1, \mu_j > \mu^c) = 0$. Hence, the second term of the right-hand side equals zero. The term between square brackets represents $\Delta(\mu^c, \mu^c)$, i.e. the difference between the cutoff type's gain of investing early and her gain of waiting. Suppose that $\mu^c = \underline{\mu} < 0$. By definition of $\underline{\mu}$ this means that the cutoff type μ^c is indifferent between investing and not investing in case she gets good news. Hence, one can think of her as someone who will not invest at time two—independent of the other player's time-one decision. In that case her gain of waiting is zero and $\Delta(\underline{\mu}, \underline{\mu}) = \underline{\mu} < 0$. Suppose now that $\mu^c = \mu^{LF}$. As the cutoff player is indifferent between investing and waiting, $\Delta(\mu^{LF}, \mu^{LF}) = 0$. If equilibrium is unique—as is the case if the prior mean is sufficiently high[18]—there exists no other cutoff $\mu^c \ne \mu^{LF}$ such that $\Delta(\mu^c, \mu^c) = 0$. By continuity, the term between square brackets (i.e. *excluding* the minus sign in front) is thus negative if equilibrium is unique and if $\mu^c \in [\underline{\mu}, \mu^{LF})$. Hence, if the prior mean is sufficiently high, and if the social planner implements a cutoff $\mu^c < \mu^{LF}$, then the first term on the right-hand side (i.e. *including* the minus sign in front) is positive. Using (10), we thus conclude that $\forall \mu^c \in [\underline{\mu}, \mu^{LF}]$, $dU/d\mu^c > 0$ if the prior mean $\overline{\theta}$ is sufficiently high and if

$$\int_{\underline{\mu}^1}^{\mu^c} \frac{\partial}{\partial \mu^c} [\Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c)] f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i > 0. \tag{11}$$

We now argue that if the prior mean $\overline{\theta}$ is high enough, then the above inequality is satisfied. To build some intuition, we first explain how Player i's gain of waiting is influenced by the cutoff μ^c. Suppose $\mu_i \in [\underline{\mu}^1, \mu^c]$ in which case Player i invests at time two if the other player did so at time one. In the Appendix (see proof of Proposition 2), we prove that

$$\frac{\partial}{\partial \mu^c} [\Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c)] > 0 \Leftrightarrow (1-\alpha)\overline{\theta} - \mu_i > \mu^c. \tag{12}$$

Stated differently, Player i's gain of waiting is unimodal: It increases until $\mu^c = (1-\alpha)\overline{\theta} - \mu_i$ and decreases thereafter. To understand the unimodality of Player i's gain of waiting, observe that the above derivative is equal to

$$\frac{\partial \Pr(\mu_j > \mu^c | \mu_i)}{\partial \mu^c} E(\theta | \mu_i, \mu_j > \mu^c) + \Pr(\mu_j > \mu^c | \mu_i) \frac{\partial E(\theta | \mu_i, \mu_j > \mu^c)}{\partial \mu^c}.$$

The first term of this sum is negative, while the second one is positive. Recall from our discussion of Eq. (4) that a player's expectation of the state of the world when receiving good news is the sum of her time-one posterior mean and an upward shift. Furthermore, if the critical time-one cutoff μ^c is low, both μ^c and the upward shift are small, and hence the first term in the above sum is not very negative. In addition, if μ^c is low, it is very likely that the other player will invest early. Hence, any increase in Player i's upward shift is multiplied by a large number, so that the second term in the above sum is large. Hence, if μ^c is low the above derivative is positive. In other words, if μ^c is low, Player i wants the social planner to raise the cutoff as the decrease in the probability of getting good news is more than compensated by the increase in her upward shift. The contrary situation prevails when μ^c is high: Player i then prefers the cutoff to be lowered in order to increase her probability of getting good news.

[17] Using the same reasoning as in Section 4, it is straightforward to prove that the inequality ($\mu^{LF} < \overline{\mu}$) is strict: Suppose $\mu^{LF} = \overline{\mu}$. If the cutoff type μ^{LF} receives bad news, she is indifferent between investing and not investing. One can thus think of her as investing at time two—independent of the other player's decision. But then she is better off investing early and saving on the discounting cost.

[18] See Proposition 1.

To further understand the change in the active types' gain of waiting, let μ_i^{max} denote the cutoff (above which Player *j* invests) that maximizes type μ_i's gain of waiting. It follows from (12) that

$$\mu_i^{max} = (1-\alpha)\overline{\theta} - \mu_i. \tag{13}$$

Suppose Player *j* invests. This reveals to Player *i* that $\mu_j > \mu^c$. Now recall from the previous section that $\mu_j|\mu_i$ is normally distributed with mean $\alpha\mu_i + (1-\alpha)\overline{\theta}$. Thus the higher $\overline{\theta}$, the less "good news" is contained in the observation that Player *j* invested, and hence the lower is Player *i*'s upward shift. Similarly, the higher the type μ_i, the greater the expectation that Player *j* invests, and the smaller the upward shift when observing Player *j* investing. Nevertheless, Eq. (13) reveals that μ_i^{max} is *decreasing* in her posterior mean μ_i. To understand why, recall that Player *i* only invests at time two if Player *j* did so at time one. Hence, the higher μ_i the more Player *i* "fears" that Player *j* will not invest at time one. Stated differently, if μ_i increases Player *i* prefers the cutoff μ^c to be reduced to increase the likelihood of *j* investing even if this comes at the cost of a lower upward shift.

Now, suppose that

$$(1-\alpha)\overline{\theta} - \mu^{LF} > \mu^{LF}. \tag{14}$$

Economically, this inequality states that the cutoff-type of Player *i* in a laissez-faire economy thinks that the cutoff μ^{LF} of Player *j* is too low. Hence, if this condition holds the gain of waiting of the cutoff type $\mu_i = \mu^{LF}$ is increasing in the other player's first-period cutoff. In this case the marginal type and—because the gain of waiting is increasing in the cutoff for low types μ_i—all inframarginal types prefer a cutoff above the laissez-faire one. Hence, Inequality (14) implies that Inequality (11) must hold. Whenever this is the case, we already argued that it is optimal to raise the first-period cutoff. Finally, recall from Section 4 that the laissez-faire cutoff tends to zero as the prior mean $\overline{\theta}$ goes to infinity, which implies that Inequality (14) is indeed satisfied, and hence it is optimal for the social planer to implement a higher period-one cutoff. We summarize these insights in our next proposition.

Proposition 2. *Suppose that the prior mean $\overline{\theta}$ is high enough. The social planner's optimal period-one investment cutoff is then strictly higher than in the laissez-faire economy $\mu^{SP} > \mu^{LF}$. Any period-one investment cutoff $\mu^c < \mu^{LF}$ yields then a lower welfare than the one which prevails in a laissez-faire economy. There exists then an $\epsilon > 0$ such that for all time-one cutoffs μ^c satisfying $\mu^c - \mu^{LF} \in (0,\epsilon)$, welfare is strictly greater than in the laissez-faire economy.*

Proposition 2 extends intuition about the insufficient use of private information derived in the original herding papers (see Banerjee, 1992 and Bikhchandani et al., 1992) to an endogenous queue set-up. In an informational cascade, Player *i* gets say sufficiently good public information about the returns from investing, which arises when enough predecessors in a queue decide to invest, so that she follows the public information and invests even when possessing an unfavorable private signal. This investment decision is typically socially inefficient as it impedes subsequent movers to infer this player's signal from her action. A similar inefficiency also arises in our model: When stories about the high profitability of an investment opportunity abound (i.e. if $\overline{\theta}$ is sufficiently high), an inefficiently high mass of types end up investing early, making it harder for players who wait to confidently infer that the state of the world is indeed conducive to investing.

Unfortunately, policy recommendations are harder to derive when the prior mean $\overline{\theta}$ is not sufficiently high as inframarginal types then disagree among themselves.[19] Types with "very" negative time-one posterior means experience a huge upward shift upon observing their rival investing at time one, which induces them to invest at time two. Increasing the cutoff increases the upward shift dramatically and thus overcompensates the lower probability of getting good news for these types. A higher inframarginal type's upward shift, however, is increased by less and for such a type the lower probability of getting good news dominates. This implies that they prefer the social planner to *reduce* μ^{LF}. The social planner thus faces a tradeoff in this case: she needs to weigh the benefit of a decrease in μ^{LF} for some inframarginal types against the losses for other inframarginal types. As this exercise is analytically demanding, we have not been able to prove that a (strictly positive) subsidy is optimal when the prior mean is sufficiently low. In the Appendix (see proof of Lemma 5), however, we show that

$$\lim_{\overline{\theta}\to-\infty} \underbrace{\mu^{LF}}_{\text{Equilibrium cutoff}} = 0 = \lim_{\overline{\theta}\to-\infty} \underbrace{(1-\alpha)\overline{\theta}-\underline{\mu}^1}_{\text{Type } \underline{\mu}^{1\prime}\text{s preferred cutoff}}.$$

Hence, in the limit the inframarginal type with the lowest posterior mean (i.e. type $\underline{\mu}^1$) is perfectly happy with the cutoff μ^{LF}. It then follows from (13) that all the other inframarginal types (i.e. all types $\mu_i \in (\underline{\mu}^1, \mu^{LF}]$) think that the laissez-faire cutoff is too high. Their gain of waiting would be higher if the social planner were to implement a lower cutoff. This result implies that if the social planner implements a cutoff $\mu^c > \mu^{LF}$, there exists some critical prior mean $\overline{\theta}^c$ such that if $\overline{\theta} < \overline{\theta}^c$ the social planner could raise welfare by reducing the cutoff μ^c. This result also implies that, as $\overline{\theta}$ goes to minus infinity, $\mu^{SP} \le \mu^{LF}$. To summarize:

Proposition 3. *Any time-one cutoff $\mu^c > \mu^{LF}$ is suboptimally high if $\overline{\theta}$ lies below some threshold $\overline{\theta}^c(\mu^c)$.*

[19] This is proven in Lemma 4 for the case in which the prior mean $\overline{\theta}$ is negative.

6. Implementation

We now discuss how a social planner that selects first- and second-period taxes can implement the optimal investment cutoffs. In our previous section we argued that the time-two cutoffs $\underline{\mu}^0$ and $\underline{\mu}^1$ for any given period-one cutoff should be chosen such that no profitable investment opportunity is wasted at time two, i.e. $E(\theta|\mu_i=\underline{\mu}^0,\mu_j<\mu^c)=0$ and $E(\theta|\mu_i=\underline{\mu}^1,\mu_j>\mu^c)=0$. Hence the optimal second-period tax is zero. Thus, in the remainder of this section we analyze the optimal time-one tax. A tax τ is said to implement a first-period cutoff μ^c if there exists *a* symmetric switching equilibrium with the property that $\mu^*=\mu^c$.

We have established in Section 5 that the social planner sets a cutoff μ^c that induces some types who wait to invest when getting good news, i.e. that $\mu^{SP}>\underline{\mu}$. Intuitively, this ensures that a meaningful information externality remains. We thus restrict attention to cutoffs $\mu^c>\underline{\mu}$. Note also that $\mu^{LF}<\overline{\mu}$ since in the laissez-faire equilibrium a period-one cutoff type invests in period two if and only if she gets good news, while for cutoffs above $\overline{\mu}$ a period-one cutoff type also invests in period two when getting bad news.

Suppose the social planner wants to implement a cutoff $\mu^c\in[\underline{\mu},\overline{\mu}]$ so that the cutoff type μ^c invests at time two if and only if the other player did so at time one. Then the cutoff type is indifferent between investing and waiting if

$$\mu^c-\tau=\delta \Pr(\mu_j>\mu^c|\mu_i=\mu^c)E(\theta|\mu_i=\mu^c,\mu_j>\mu^c). \tag{15}$$

Recall from Eqs. (3) and (4) that

$$\Pr(\mu_j>\mu^c|\mu_i=\mu^c)=1-F(\kappa_1(\mu^c-\overline{\theta})) \tag{16}$$

and that

$$E(\theta|\mu_i=\mu^c,\mu_j>\mu^c)=\mu^c+\kappa_2 h(\kappa_1(\mu^c-\overline{\theta})). \tag{17}$$

Using (16) and (17), the indifference Eq. (15) can be rewritten as

$$\mu^c-\frac{\tau}{1-\delta(1-F(\kappa_1(\mu^c-\overline{\theta})))}=\kappa_2\mathcal{X}(\kappa_1(\mu^c-\overline{\theta})). \tag{18}$$

The right-hand side of the indifference equation above was analyzed in Section 4. Observe that if $\tau=0$, the indifference condition boils down to (5). Call *LHS* the left-hand side of the above equation after replacing μ^c with μ and observe that *LHS* is continuous in μ, and that it goes from minus to plus infinity as μ goes from $-\infty$ to $+\infty$. Observe also that *LHS* always lies above μ if τ is negative. If τ is positive, however, *LHS* always lies below μ. This implies that if the symmetric switching equilibrium is unique in the laissez-faire economy, and if the social planner wants to raise the equilibrium cutoff (i.e. if $\mu^{SP}\in(\mu^{LF},\overline{\mu})$), he must tax investments. (This case is illustrated in Fig. 3.) Conversely, if the symmetric switching equilibrium is unique in the laissez-faire case and the social planner wants to implement a cutoff $\mu^{SP}\in(\underline{\mu},\mu^{LF})$, he must subsidize investments.

One can rewrite the equilibrium condition (18) as

$$\tau=[1-\delta(1-F(\kappa_1(\mu^c-\overline{\theta})))]\mu^c-\kappa_2\delta f(\kappa_1(\mu^c-\overline{\theta})). \tag{19}$$

The social planner can thus implement any $\mu^c\in[\underline{\mu},\overline{\mu}]$ by setting τ equal to the right-hand side of the above equation.[20]

Using Eq. (19), it is straightforward to show that if $\overline{\theta}<0$ and if $\mu^c\in(\mu^{LF},\overline{\mu}]$, $\partial\tau/\partial\mu^c>0$.[21] Hence, if the prior mean $\overline{\theta}$ is negative and if the social planner wants to implement a cutoff lower than $\mu^c(>\mu^{LF})$, she should reduce the tax τ. Recall from Proposition 3 that for any cutoff $\mu^c>\mu^{LF}$ there exists a critical $\overline{\theta}^c(\mu^c)$ such that the social planner can raise welfare by implementing a lower cutoff if the prior mean $\overline{\theta}$ is less than $\overline{\theta}^c(\mu^c)$. Both results thus imply that for any positive tax τ there exists a critical $\overline{\theta}^c(\tau)$ such that if $\overline{\theta}<\overline{\theta}^c(\tau)$, the social planner can raise welfare by reducing the investment tax.

Suppose now that the social planner wants to implement a cutoff $\mu^c>\overline{\mu}>\mu^{LF}$. It then follows from our previous section that the cutoff type μ^c will invest at time two—independent of the other player's time-one action. Hence, in this case the cutoff type is indifferent between investing and waiting if $\mu^c-\tau=\delta\mu^c$. Any $\mu^c>\overline{\mu}$ can thus be implemented by choosing τ such that $\tau=(1-\delta)\mu^c>0$, where the inequality follows from the fact that $\mu^c>\overline{\mu}>0$.

An important caveat, however, is that a given τ need not uniquely implement μ^{SP}. To understand this, consider Fig. 3. In the figure, the prior mean $\overline{\theta}$ is implicitly assumed to be "high". (This explains why in the figure the median of $\mathcal{X}$ is drawn so much to the right and—as explained in Section 4—why μ^{LF} is close to zero.) As summarized in Proposition 2, if the prior mean is high enough, the equilibrium cutoff μ^{LF} is too low. (This is illustrated in Fig. 3 by the fact that $\mu^{LF}<\mu^{SP}$.) The figure also shows that if investments are appropriately taxed (i.e. if $\tau=\tau'$), there exists a symmetric switching equilibrium in which players coordinate on the efficient time-one cutoff μ^{SP}. In the figure, however, there also exists another symmetric switching equilibrium in which a player invests if and only if her time-one posterior mean exceeds $\tilde{\mu}$. In this case the investment tax τ' actually deters too many types from investing.

[20] It is straightforward to check that Lemma 1 remains true after introducing a lump-sum tax τ in our model. Hence, if τ is chosen such that Eq. (19) is satisfied, Player i is indifferent between investing and waiting if and only if her posterior mean is equal to μ^c, while she prefers to wait if and only if her posterior mean is less than μ^c.

[21] Implicitly, we also use our earlier result that $0<\mu^{LF}$.

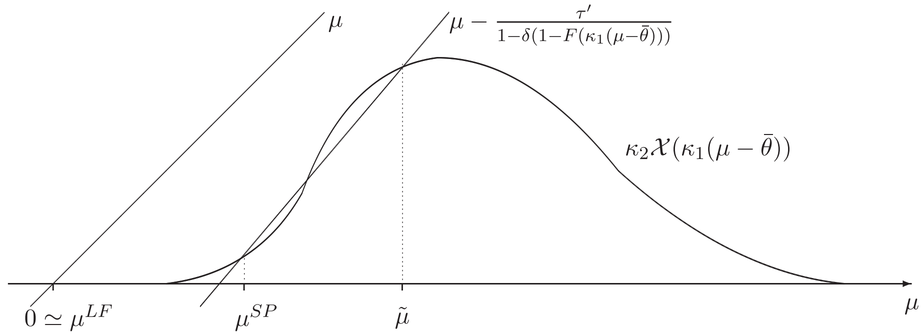

Fig. 3. Non-unique implementation of μ^{SP}.

If players focus on this latter equilibrium, the investment tax τ' may even *decrease* welfare as compared to the prevailing one in a laissez-faire economy. Recall, however, from Proposition 2 that if the prior mean is high enough, the switching equilibrium is unique and a small increase in the equilibrium cutoff increases welfare. Thus, if the prior mean is high enough (i.e. if $\kappa_2\mathcal{X}$ lies sufficiently to the right as is the case in Fig. 3) there exists a value of $\mu^c > \mu^{LF}$ such that the switching equilibrium remains unique and welfare is higher under cutoff μ^c than under cutoff μ^{LF}. Hence, even if the social planner anticipates that players will focus on cutoff $\tilde{\mu}$ instead of μ^{SP} in case $\tau = \tau'$, it is still optimal for him to tax investments. (Of course, the tax will have to be smaller than τ'.) Our main findings are summarized below.

Proposition 4. *If the prior mean is high enough, it is optimal to tax first-period investments. Suppose $\tau > 0$. Then there exists a critical value $\overline{\theta}^c(\tau)$ such that if $\overline{\theta} < \overline{\theta}^c(\tau)$, the social planner can raise welfare by reducing the investment tax τ. Furthermore, any socially optimal cutoff μ^{SP} can be implemented through an appropriate first-period investment tax/subsidy τ and a second-period investment tax of zero. Implementation, however, need not be unique.*

According to the (perhaps recent) conventional wisdom, governments should not intervene in the presence of an investment bubble as one cannot ex ante know whether it is due to fundamentals (corresponding to the case in which $\theta > 0$ in our model) or whether it is the result of incorrect stories. Alan Greenspan's quote in our introduction, for example, nicely illustrates the wisdom that prevailed in 1999. Our model questions this rationale for non-intervention: Even if policymakers in contrast to market participants receive no private signal about the state of the world, the policymakers' knowledge of $\overline{\theta}$ can still be used to improve welfare.[22] In particular, in the presence of sufficiently favorable public information, investments should be taxed.

More broadly, Propositions 2 and 3 are consistent with the idea that investment policy should be countercyclical: when $\overline{\theta}$ is high (which is likely to occur when many players have invested in the previous period(s)), investments should be taxed, while if $\overline{\theta} \to -\infty$ (i.e. when expected investment activity is zero) investment should not be taxed.

We allow the tax τ to be set conditional on the public information. In practice, the existing information in the public domain needs to be interpreted and this can be difficult.[23] The exact process of how this is done—e.g. through public hearings or a special committee—is not crucial for our results as long as policy makers do not have access to information that remains hidden from investors. If policy makers have information that is unavailable to investors, the selected tax rate could signal the state of the world, an aspect our modeling approach abstracts from.[24]

Insofar as investments into new technologies are concerned, it is natural to presume that tax or subsidy rates are technology specific. Most industrialized countries—including the US—are actively pursuing industrial policies. Policy makers try to identify and encourage investments into promising technologies or "strategically important" industries. Similarly, they often encourage the adoption of certain technologies, such as solar panels. The amount of tax breaks or subsidies investors receive depends on the presumed need or the individual profitability, which corresponds to the state of the world in our stylized set-up.

So far, we assumed that the social planner can freely change the investment tax/subsidy between the two periods. In general, one would not expect the government to frequently change investment policy on the basis of the latest investment activity. We only made this assumption, however, to simplify the analysis of the optimal cutoffs. In particular, we have shown that even if the investment tax has to be kept fixed for both periods, it remains optimal to tax investments whenever the prior mean is sufficiently high. Intuitively, if the social planner raises the investment tax τ from zero to ϵ, she raises the first-period equilibrium cutoff which increases welfare. An increase in the investment tax τ, however, also

[22] Greenspan was primarily worried about the existence of an investment boom in the U.S. stockmarket, i.e. in a context in which prices may aggregate information. As our model is void of any price mechanism, we cannot address the question whether one should interfere in the stockmarket. We feel, however, that (perhaps until recently) the vast majority of policy-makers would agree (or would have agreed) with Greenspan even in a fixed-price context. For example, it is often argued that the state should not intervene into technology adoption or development decisions as the private sector has more information regarding appropriate technologies. Our model highlights that even a less informed policy maker may want to interfere after all.

[23] Alternatively, one can try to find a real-life proxy for $\overline{\theta}$ and make τ contingent on that proxy. Finding a good proxy, however, is challenging. Past investment activity, for example, is an imperfect proxy for $\overline{\theta}$: A past investment boom may indeed indicate a high realization of θ. It may, however, also be due to the fact that players use low cutoffs.

[24] The optimality of taxing strictly better informed investors is also theoretically more surprising than it would be in an environment in which the policy maker has private information that is hidden from investors.

distorts time-two investment decisions. This welfare loss, however, is a second-order effect. Furthermore, we have also shown that the issue of non-unique implementation disappears with permanent taxes.[25] To be more specific, if the investment tax is the same in both periods, and if the prior mean is sufficiently high, the constrained-optimal symmetric time-one cutoff can be *uniquely* implemented by appropriately choosing the appropriate permanent investment tax.

We stated our results in terms of a social planner who taxes or subsidizes the investment activity itself.[26] Alternatively to taxing investment itself a social planner can also steer the incentives to invest by altering the tax rate on the resulting profits. She could, for example, commit herself to tax future profits at a certain rate $t \in (0,1)$. In principle, the rate t can depend on the investment date. A firm's tax bill is then proportional to the amount of profits she is making. If she turns a loss, however, she pays no taxes. Player i's tax bill is thus equal to $\int_0^\infty t\theta f((\theta-\mu_i)\sigma_\mu)\,d\theta$. A profit tax t implements $\mu^{SP} > \mu^{LF}$ if (i) $\Delta(\mu^{SP},\mu^{SP}) = 0$ and if (ii) Player i prefers to wait if and only if her posterior mean is less than μ^{SP}. It can be shown that requirement (i) can be satisfied: There always exists a tax rate $t \in (0,1)$ such that the cutoff type $\mu^{SP} > \mu^{LF}$ is indifferent between investing and waiting given that the other player uses the cutoff μ^{SP}. Requirement (ii), however, need not be satisfied with a profit tax. Intuitively, with a profit tax more optimistic players expect to pay ahigher tax bill if they invest at time one. By assumption they do not pay any taxes if they invest at time two. A player with a posterior mean $\mu_i > \mu^{SP}$ may therefore strictly prefer to wait. Nevertheless, it is easy to show that if the tax rate t is sufficiently small, condition (ii) is satisfied. Hence, the main implication of Proposition 4 remains unchanged: If the prior mean is sufficiently high, the introduction of a small profit tax t improves welfare.

Finally, in our model it is sometimes optimal to tax investments because the increase in the cutoff leads to an upward shift in beliefs. This upward shift is due to our assumption that Player i only observes whether or not Player j invested. If investment by Player j would immediately reveal the state of the world, the updating process would be different. We conjecture that because players do not internalize the information externalities of investing early, investments in this case should always be subsidized.[27]

7. Final remarks

We analyzed some policy implications of social learning when players are fully rational and better informed than the policymaker. Our model is particularly useful when public information is conducive to investing—which typically happens during "boom times" or when many "stories" circulate praising the profit prospects of the investment opportunity. In this case, we establish that, in a laissez-faire economy, too many types are investing early and investments should therefore be taxed.

We have chosen a two-player set-up for our model. The general N player game is difficult to analyze.[28] One "simple" alternative, however, would be to consider a model with a continuum of players. In that variation, for any given symmetric equilibrium cutoff, social learning would be perfect and hence a laissez-faire policy optimal. To circumvent this unrealistic feature, one needs to assume social learning to be imperfect. One possibility is to assume observational noise as in Chamley (2004a) or Dasgupta (2007). In such a setup, Player i's distribution about the other players' posterior means (i.e. $f(\mu_j|\mu_i)$) would still be computed in the same way as in our two-player model. Therefore, if the prior mean is "very high", an inframarginal type expects—for "many" realizations of the state of the world—a large mass of players to invest at time one. As noisy observation of past investment behavior is then expected to reveal relatively little information about the realized state of the world, we conjecture that—as in our model—the inframarginal types prefer the social planner to raise the equilibrium cutoff via taxes. One drawback of such an approach, however, is that the observational noise is completely exogenously specified. An alternative assumption is that each player can only observe some (neighboring) players' first-period decisions.[29]

In our model information can only be transmitted through actions. As there are no payoff externalities, it is natural to ask why information cannot be transmitted through words instead. If players can fully exchange their private information via cheap talk, an efficient equilibrium of course exists. We feel, however, that this simple argument is misleading as communication—even where allowed and feasible—is often imperfect. Suppose, for example, that player one is asked to reveal her type to the other player(s) prior to the waiting game. As her signal is imperfect, she also wants to learn the other player(s)' signal(s). She therefore has an incentive to send the message which maximizes her gain of waiting. In an analysis of cheap talk, Gossner and Melissas (2006) have shown that this game may—depending on the values of the parameters—be characterized by a unique monotone equilibrium in which all types send the same message, i.e. information can only be

[25] The formal result is contained in the working paper version (Heidhues and Melissas, 2010).

[26] In our model, θ should be interpreted as the gross expected profits from the investment project minus the expected investment cost. Governments often promote investments by subsidizing the investment cost. It may, for example, offer a grant that covers R&D expenses or offer a tax break proportional to the investment cost. If Player i invests, she thus gets $\mu_i - \tau$, where $\tau(<0)$ represents the subsidy on the investment cost. Conversely, a government can use a non-refundable sales tax to increase the investment cost by some lump-sum amount. And if the investment requires a particular input—such as public land, mining rights, or a permit—the government can also simply affect the investment cost directly.

[27] We also conjecture that a similar subsidization result would hold if a player's signal were to become common knowledge once she invests.

[28] We have been able to establish, however, that equilibrium is unique with N players when the state of the world θ is drawn from a Laplacian distribution. Again, the proof is available upon request.

[29] Our model, for example, can be seen as a special case in which countably many players live on a circle and each player only observes her right-hand neighbor.

revealed through actions. More generally, we believe the study of waiting games in the presence of imperfect communication to be an interesting avenue for future research.

Another noteworthy aspect of our model is that investment costs are exogenous.[30] In many applications in which policymakers are concerned about investment bubbles—such as stock market or housing market bubbles—one would expect investment costs to increase in the number of present and past investments. In an exogenous queue model with a competitive market maker, Park and Sabourian (2011) establish that herding is possible even if markets are informationally efficient. An interesting question for future research is how these results extend to an endogenous queue setting and whether it is also optimal to tax investments during boom times in such a model.

We assumed that players are fully rational. Eyster and Rabin (2010) nicely highlight some counterintuitive features of the rational learning model in an exogenous queue environment and propose a plausible alternative learning model. An interesting question is whether the introduction of inferentially naive and/or cursed players strengthens or qualifies our "taxation during booms" result in an endogenous queue environment.

Acknowledgments

Earlier versions of this paper circulated under the title "Technology Adoption, Social Learning, and Economic Policy". We are grateful to Micael Castanheira, Shurojit Chatterji, Gianni De Fraja, Shachar Kariv, Claudio Mezzetti, Matthew Rabin, Frank Riedel, Tridib Sharma, Rodney Strachan, Adam Szeidl, Glen Weyl, seminar participants in ITAM, two anonymous referees and Dan Kovenock, the associate editor, for helpful suggestions. Heidhues gratefully acknowledges financial support from the Deutsche Forschungsgemeinschaft through SFB/TR-15. Part of this paper were written while Heidhues visited the Department of Economics at UC Berkeley whose hospitality he gratefully acknowledges. Melissas gratefully acknowledges financial support from the Consejo Nacional de Ciencia y Tecnología (grant #79741) and from the Asociación Mexicana de Cultura A.C.

Appendix A

A.1. Definitions and preliminaries

Throughout the appendix F, f, h, and r represent, respectively, the c.d.f., the p.d.f., the hazard rate $(\equiv f(\cdot)/(1-F(\cdot)))$, and the reverse hazard rate $(\equiv f(\cdot)/F(\cdot))$ of the standard normal distribution. We will also use the following notations: $\alpha \equiv \sigma_\theta^2/(\sigma_\theta^2+\sigma_\epsilon^2)$, $\beta \equiv (2/\sigma_\epsilon^2)/(1/\sigma_\theta^2+2/\sigma_\epsilon^2)$, $\sigma_p^2 \equiv \sigma_\theta^2\sigma_\epsilon^2/(\sigma_\theta^2+\sigma_\epsilon^2)$, $\sigma_2^2 \equiv \sigma_p^2+\sigma_\epsilon^2$, $\sigma_o^2 \equiv \alpha^2\sigma_2^2$, $\sigma_\mu^2 \equiv \alpha^2(\sigma_\theta^2+\sigma_\epsilon^2)$, $\kappa_1 \equiv (1-\alpha)/\sigma_o$, $\kappa_2 \equiv \frac{1}{2}\beta\sigma_2$, $x(\mu^c,\mu_i) \equiv (\mu^c-\alpha\mu_i-(1-\alpha)\overline{\theta})/\sigma_o$, $\mathcal{X}(\eta) \equiv \delta f(\eta)/(1-\delta(1-F(\eta)))$, $g(\mu) \equiv \mu-\kappa_2\mathcal{X}(\kappa_1(\mu-\overline{\theta}))$, and $\phi(\mu) \equiv \mu+\kappa_2 h(x(\mu^*,\mu))$.

In our set-up (see DeGroot, 1970 for proofs) $\theta|s_i \sim N(\mu_i,\sigma_p^2)$, where

$$\mu_i = \alpha s_i + (1-\alpha)\overline{\theta}. \tag{20}$$

As ϵ_j is independent from θ and ϵ_i, $s_j|s_i = \theta|s_i+\epsilon_j$. As $\epsilon_j \sim N(0,\sigma_\epsilon^2)$, $s_j|s_i \sim N(\mu_i,\sigma_p^2+\sigma_\epsilon^2)$. Furthermore, $\mu_j = \alpha s_j+(1-\alpha)\overline{\theta}$, and thus

$$\mu_j|s_i \sim N(\alpha\mu_i+(1-\alpha)\overline{\theta},\sigma_o^2). \tag{21}$$

Hence

$$\Pr(\mu_j > \mu^*|\mu_i) = 1-F\left(\frac{\mu^*-\alpha\mu_i-(1-\alpha)\overline{\theta}}{\sigma_o}\right) \tag{22}$$

and

$$\Pr(\mu_j > \mu^*|\mu_i = \mu^*) = 1-F(\kappa_1(\mu^*-\overline{\theta})).$$

Lemma 2. *If signals and the state of the world are drawn from Normal distributions*

1. $E(\theta|\mu_i,\mu_j > \mu^c)$ *and* $E(\theta|\mu_i,\mu_j < \mu^c)$ *are increasing in* μ_i *and* μ^c.
2. *One has*

$$E(\theta|\mu_i,\mu_j > \mu^c) = \mu_i+\kappa_2 h(x(\mu^c,\mu_i)) \text{ and}$$

$$E(\theta|\mu_i,\mu_j < \mu^c) = \mu_i-\kappa_2 r(x(\mu^c,\mu_i)).$$

3. *For any first-period cutoff* μ^c, *there exist unique-second period cutoffs* $\underline{\mu}^0$ *and* $\underline{\mu}^1$ *such that* $E(\theta|\mu_i = \underline{\mu}^0,\mu_j < \mu^c) = 0$ *and* $E(\theta|\mu_i = \underline{\mu}^1,\mu_j > \mu^c) = 0$.

[30] See also the discussion in Footnote 22.

4. *There exists a unique* $\underline{\mu}<0$ *such that* $E(\theta|\underline{\mu},\mu_j>\underline{\mu})=0$. *There also exists a unique* $\overline{\mu}>0$ *such that* $E(\theta|\overline{\mu},\mu_j<\overline{\mu})=0$. *If* $\mu^c<\underline{\mu}$, *no active type invests at time two. If* $\mu^c>\overline{\mu}$, *some active types invest at time two even if no one did so at time one. If* $\mu^c\in(\underline{\mu},\overline{\mu})$, *an active player only invests at time two if she received good news.*

Proof. We first prove points 1 and 2 of the lemma. A well known statistical result (see DeGroot, 1970 for a proof) is that if $\theta\sim N(\overline{\theta},\sigma_\theta^2)$ and if $\epsilon_i\sim N(0,\sigma_\epsilon^2)$, then $\theta|s_i,s_j$ also tends to a normal and

$$E(\theta|s_i,s_j)=\beta\frac{s_i+s_j}{2}+(1-\beta)\overline{\theta}. \tag{23}$$

We first tackle the case in which $\mu_j>\mu^c$. It follows from (20) that $\mu_j>\mu^c\Leftrightarrow s_j>s^c\equiv(\mu^c-(1-\alpha)\overline{\theta})/\alpha$. One has

$$\begin{aligned}E(\theta|\mu_i,\mu_j>\mu^c)&=\int\left[\beta\frac{s_i+s_j}{2}+(1-\beta)\overline{\theta}\right]f(s_j|s_i,s_j\geq s^c)\,ds_j,\\&=\frac{\beta}{2}s_i+\frac{\beta}{2}E(s_j|s_i,s_j>s^c)+(1-\beta)\overline{\theta}.\end{aligned} \tag{24}$$

From the explanations provided after (20), we know that $s_j|s_i,s_j>s^c$ is a left-truncated normal distribution with mean μ_i, variance σ_2^2 and truncation point s^c. Using Johnson et al. (1995) to calculate the expectation of a truncated normal variable, one has

$$E(s_j|s_i,s_j>s^c)=\mu_i+h\left(\frac{s^c-\mu_i}{\sigma_2}\right)\sigma_2. \tag{25}$$

Replacing s^c by $(\mu^c-(1-\alpha)\overline{\theta})/\alpha$ and taking into account that $\sigma_o^2=\alpha^2(\sigma_p^2+\sigma_\epsilon^2)=\alpha^2\sigma_2^2$, allow us to rewrite (25) as $E(s_j|s_i,s_j>s^c)=\mu_i+h(x(\mu^c,\mu_i))\sigma_2$. Inserting this last equality into (24), and taking into account the fact that $\mu_i=\alpha s_i+(1-\alpha)\overline{\theta}$ and that $\beta(1+\alpha)=2\alpha$, yields

$$E(\theta|\mu_i,\mu_j>\mu^c)=\mu_i+\kappa_2h(x(\mu^c,\mu_i)). \tag{26}$$

Differentiating (26), and taking into account that $(\alpha/\sigma_o)\kappa_2=\frac{1}{2}\beta$, one has

$$\frac{\partial E(\theta|\mu_i,\mu_j>\mu^c)}{\partial\mu_i}=1-\frac{1}{2}\beta h'(x(\mu^c,\mu_i)). \tag{27}$$

As is well known (see, e.g. Greene, 1993, Theorem 22.2), the slope of the hazard rate of a standard normal distribution, $h'(z)\in(0,1)\forall z$. This insight, combined with the fact that $\beta\in[0,1]$, allows us to conclude that $\partial E(\theta|\mu_i,\mu_j>\mu^c)/\partial\mu_i>0$. Differentiating (26) with respect to μ^c and taking into account that $\kappa_2/\sigma_o=\beta/2\alpha$, one has

$$\frac{\partial E(\theta|\mu_i,\mu_j>\mu^c)}{\partial\mu^c}=\frac{\beta}{2\alpha}h'(x(\mu^c,\mu_i)).$$

As $h'(z)\in(0,1)$ and as both α and β are positive, we conclude that $\partial E(\theta|\mu_i,\mu_j>\mu^c)/\partial\mu^c>0$.

We now tackle the case in which $\mu_j<\mu^c$. As above

$$E(\theta|\mu_i,\mu_j<\mu^c)=\frac{\beta}{2}s_i+\frac{\beta}{2}E(s_j|\mu_i,\mu_j<\mu^c)+(1-\beta)\overline{\theta}. \tag{28}$$

From Johnson et al. (1995), we know that

$$E(s_j|\mu_i,s_j<s^c)=\mu_i-r\left(\frac{s^c-\mu_i}{\sigma_2}\right)\sigma_2. \tag{29}$$

Inserting (29) into (28), replacing s^c by $(\mu^c-(1-\alpha)\overline{\theta})/\alpha$, and taking into account that $\beta/2\alpha=1/(1+\alpha)$ and that $1-\beta=(1-\alpha)(1+\alpha)$, yields

$$E(\theta|\mu_i,\mu_j<\mu^c)=\mu_i-\kappa_2r(x(\mu^c,\mu_i)). \tag{30}$$

Differentiating this last equation, and using the fact that $\frac{\alpha}{\sigma_o}\kappa_2=\frac{1}{2}\beta$, one has

$$\frac{\partial E(\theta|\mu_i,\mu_j<\mu^c)}{\partial\mu_i}=1-\frac{1}{2}\beta r'(x(\mu^c,\mu_i)).$$

It is well known (see, e.g. Greene, 1993, Theorem 22.2) that $r'(\cdot)\in(-1,0)$. As $\beta\in[0,1]$, we conclude that $\partial E(\theta|\mu_i,\mu_j<\mu^c)/\partial\mu_i$ is positive. Differentiating (30) with respect to μ^c, one has $\partial E(\theta|\mu_i,\mu_j<\mu^c)/\partial\mu^c=-(\kappa_2/\sigma_o)r'(\cdot)$, which is positive as $r'(\cdot)<0$.

We now prove Point 3 of the lemma. From above, we know that $E(\theta|\mu,\mu_j>\mu^c)=\mu+\kappa_2h(x(\mu^c,\mu))$. At the second-period cutoff $\underline{\mu}^1$, one has $\underline{\mu}^1+\kappa_2h(x(\mu^c,\underline{\mu}^1))=0$. From what precedes, we also know that $\mu+\kappa_2h(x(\mu^c,\mu))$ is increasing in μ. Hence, if there exists a solution, it is unique. We are left to establish that a solution exists. First, observe that $\lim_{\mu\to\infty}[\mu+\kappa_2h(x(\mu^c,\mu))]>0$. Second, note that $\lim_{\mu\to-\infty}[\mu+\kappa_2h(x(\mu^c,\mu))]<0$, is equivalent to $\lim_{\mu\to-\infty}[(\mu+\kappa_2h(x(\mu^c,\mu)))/\mu]>0$, which by l'Hôpital's rule is equivalent to $\lim_{\mu\to-\infty}[1-(\beta/2)h'(x(\mu^c,\mu))]>0$.Since $h'\in(0,1)$ and $\beta<1$ this holds, which establishes the existence of $\underline{\mu}^1$.

Recall that $E(\theta|\mu,\mu_j<\mu^c)=\mu-\kappa_2 r(x(\mu^c,\mu))$. At the second-period cutoff $\underline{\mu}^0$, one has $\underline{\mu}^0-\kappa_2 r(x(\mu^c,\underline{\mu}^0))=0$. From above, we know that $\mu-\kappa_2 r(x(\mu^c,\mu))$ is increasing in μ. Hence, if a solution exists, it is unique. Note that $\lim_{\mu\to-\infty}[\mu-\kappa_2 r(x(\mu^c,\mu))]<0$. Using l'Hôpital's rule, $\lim_{\mu\to\infty}[\mu-\kappa_2 r(x(\mu^c,\mu))]>0$, is equivalent to $\lim_{\mu\to\infty}[1+\frac{1}{2}\beta r'(x(\mu^c,\mu))]>0$, which is satisfied as $r'(z)\in(-1,0)$. Hence, a solution exists.

We now prove point 4 of the lemma. Let $\hat{\psi}(\mu)\equiv\mu-\kappa_2 r(\kappa_1(\mu-\overline{\theta}))$, and observe that $\overline{\mu}$ is implicitly defined as $\hat{\psi}(\overline{\mu})=0$. Using the fact that $\kappa_1\kappa_2=(1-\alpha)/(1+\alpha)$, yields

$$\frac{\partial\hat{\psi}}{\partial\mu}=1-\frac{1-\alpha}{1+\alpha}r'(\cdot)>0$$

as $r'(\cdot)\in(-1,0)$. This insight, combined with the fact that $\lim_{\mu\to-\infty}\hat{\psi}(\mu)=-\infty$ and that $\lim_{\mu\to\infty}\hat{\psi}(\mu)=\infty$ proves the existence of a unique $\overline{\mu}$. In the paper, we prove that $0<\overline{\mu}$. Suppose $\overline{\mu}<\mu^c\le\mu_i$. Then $0\equiv E(\theta|\overline{\mu},\mu_j<\overline{\mu})<E(\theta|\overline{\mu},\mu_j<\mu^c)\le E(\theta|\mu_i,\mu_j<\mu^c)$, where all inequalities follow from Point 1 of this lemma. Hence, if $\overline{\mu}<\mu^c$, some active types will invest at time two even ifno one invested at time one. Similarly, let $\hat{\phi}(\mu)\equiv\mu+\kappa_2 h(\kappa_1(\mu-\overline{\theta}))$, and observe that $\underline{\mu}$ is implicitly defined as $\hat{\phi}(\underline{\mu})=0$. Using the fact that $\kappa_1\kappa_2=(1-\alpha)/(1+\alpha)$, yields

$$\frac{\partial\hat{\phi}}{\partial\mu}=1+\frac{1-\alpha}{1+\alpha}h'(\cdot)>0.$$

This insight, combined with the fact that $\lim_{\mu\to-\infty}\hat{\phi}(\mu)=-\infty$ and that $\lim_{\mu\to\infty}\hat{\phi}(\mu)=\infty$ proves the existence of a unique $\underline{\mu}$. In the paper, we prove that $\underline{\mu}<0$. Suppose $\mu_i\le\mu^c<\underline{\mu}$. Then $E(\theta|\mu_i,\mu_j<\mu^c)<E(\theta|\mu_i,\mu_j>\mu^c)<E(\theta|\underline{\mu},\mu_j>\underline{\mu})=0$, where the first inequality follows from the fact that observing $\mu_j>\mu^c$ is good news and where the second inequality follows from Point 1 of this lemma. Hence, if $\mu^c<\underline{\mu}$,no active type invests at time two. Finally, suppose that $\mu^c\in(\underline{\mu},\overline{\mu})$. As $\underline{\mu}<\mu^c$, $0=E(\theta|\underline{\mu},\mu_j>\underline{\mu})<E(\theta|\mu^c,\mu_j>\mu^c)$, where the inequality follows from point 1 of this lemma. By continuity, there exist values of μ_i close to (but less than) μ^c such that Player i wants to invest upon getting good news. As $\mu^c<\overline{\mu}$, it follows from point 1 of this lemma that $E(\theta|\mu^c,\mu_j<\mu^c)<E(\theta|\overline{\mu},\mu_j<\overline{\mu})=0$. Hence, no active type invests at time two if no one did so at time one. □

A.2. *Proof of Lemma 1*

Observe that for any finite μ_1 and μ_2^c, $E(\theta|\mu_1,\mu_2<\mu_2^c)<E(\theta|\mu_1,\mu_2>\mu_2^c)$. There are thus three possibilities: (i) $E(\theta|\mu_1,\mu_2<\mu_2^c)<E(\theta|\mu_1,\mu_2>\mu_2^c)\le 0$, (ii) $E(\theta|\mu_1,\mu_2<\mu_2^c)\le 0<E(\theta|\mu_1,\mu_2>\mu_2^c)$, and (iii) $0<E(\theta|\mu_1,\mu_2<\mu_2^c)<E(\theta|\mu_1,\mu_2>\mu_2^c)$.

In case (i), $\Delta(\cdot)=\mu_1$, which is increasing in μ_1.

In case (ii), $\Delta(\cdot)=\mu_1-\delta\Pr(\mu_2>\mu_2^c|\mu_1)E(\theta|\mu_1,\mu_2>\mu_2^c)$. Observe that

$$\mu_1=\Pr(\mu_2>\mu_2^c|\mu_1)E(\theta|\mu_1,\mu_2>\mu_2^c)+\Pr(\mu_2<\mu_2^c|\mu_1)E(\theta|\mu_1,\mu_2<\mu_2^c).$$

Inserting this last equality into $\Delta(\cdot)$, yields

$$\Delta(\cdot)=(1-\delta)\Pr(\mu_2>\mu_2^c|\mu_1)E(\theta|\mu_1,\mu_2>\mu_2^c)+\Pr(\mu_2<\mu_2^c|\mu_1)E(\theta|\mu_1,\mu_2<\mu_2^c).$$

Differentiating this last expression of $\Delta(\cdot)$ yields

$$\frac{\partial\Delta(\mu_1,\mu_2^c)}{\partial\mu_1}=(1-\delta)\frac{\partial\Pr(\mu_2>\mu_2^c|\mu_1)}{\partial\mu_1}E(\theta|\mu_1,\mu_2>\mu_2^c)+(1-\delta)\frac{\partial E(\theta|\mu_1,\mu_2>\mu_2^c)}{\partial\mu_1}\Pr(\mu_2>\mu_2^c|\mu_1)$$
$$+\frac{\partial\Pr(\mu_2<\mu_2^c|\mu_1)}{\partial\mu_1}E(\theta|\mu_1,\mu_2<\mu_2^c)+\frac{\partial E(\theta|\mu_1,\mu_2<\mu_2^c)}{\partial\mu_1}\Pr(\mu_2<\mu_2^c|\mu_1).\tag{31}$$

In case (ii), $E(\theta|\mu_1,\mu_2>\mu_2^c)>0$. As $\partial\Pr(\mu_2>\mu_2^c|\mu_1)/\partial\mu_1$ is also positive, the first term of the RHS of (31) is positive. Moreover, from Lemma 2 we know that both $\partial E(\theta|\mu_1,\mu_2>\mu_2^c)/\partial\mu_1$ and $\partial E(\theta|\mu_1,\mu_2<\mu_2^c)/\partial\mu_1$ are positive. Hence, the second and the fourth term of the RHS of (31) are also positive. In case (ii), $E(\theta|\mu_1,\mu_2<\mu_2^c)\le 0$. This assumption, combined with the fact that $\partial\Pr(\mu_2<\mu_2^c|\mu_1)/\partial\mu_1<0$, proves that the third term of the RHS of (31) is also positive.

Finally, in case (iii) $\Delta(\cdot)=(1-\delta)\mu_1$, which is also increasing in μ_1. □

A.3. *Proof of Proposition 1*

Recall that $\kappa_1=(1-\alpha)/\sigma_o$, $\kappa_2=\frac{1}{2}\beta\sigma_2$, $\sigma_2=\sqrt{\sigma_p^2+\sigma_\epsilon^2}$ and that $x(\mu_2^c,\mu_1)=(\mu_2^c-\alpha\mu_1-(1-\alpha)\overline{\theta})/\sigma_o$. Recall also that

$$\mathcal{X}(\eta)=\frac{\delta f(\eta)}{1-\delta(1-F(\eta))}.\tag{32}$$

As f denotes the p.d.f. of a standard normal distribution, $f(\eta)\equiv(1/\sqrt{2\pi})e^{-(1/2)\eta^2}$. Hence $f'(\eta)=-f(\eta)\eta$.

Lemma 3. *There exists a unique $\hat{\eta}<0$ such that $\mathcal{X}(\hat{\eta})=-\hat{\eta}$. $\mathcal{X}(\eta)$ increases until $\eta=\hat{\eta}$, after which it decreases. One has $\lim_{\eta\to-\infty}\mathcal{X}(\eta)=\lim_{\eta\to+\infty}\mathcal{X}(\eta)=0$; $\lim_{\eta\to-\infty}\mathcal{X}'(\eta)=\lim_{\eta\to+\infty}\mathcal{X}'(\eta)=0$; $\mathcal{X}''(\eta)>0$ if $\eta<\eta^m$ (where $\eta^m<\hat{\eta}$) and $\mathcal{X}''(\eta)<0$ if $\eta\in(\eta^m,\hat{\eta})$; $\lim_{\delta\to1}\hat{\eta}=-\infty$; $\lim_{\delta\to0}\mathcal{X}'(\eta^m)=0$; and $\lim_{\delta\to1}\mathcal{X}'(\eta^m)=\infty$.*

Proof. Observe that $\mathcal{X}(\eta)>0$ for $\delta>0$. Hence, $\mathcal{X}(\eta)>-\eta$, $\forall\eta>0$. Mere introspection of (32) reveals that for sufficiently low values of η, $\mathcal{X}(\eta)<-\eta$. By continuity, there exists at least one $\hat{\eta}<0$ such that $\mathcal{X}(\hat{\eta})=-\hat{\eta}$. Observe that the right hand side of the equality $\mathcal{X}(\hat{\eta})=-\hat{\eta}$ decreases in η. Using the fact that $f'(\eta)=-f(\eta)\eta$, one has

$$\frac{\partial\mathcal{X}(\eta)}{\partial\eta}=\mathcal{X}'(\eta)=-\mathcal{X}(\eta)[\eta+\mathcal{X}(\eta)]. \tag{33}$$

This slope is equal to zero if and only if $\mathcal{X}(\eta)=-\eta$. Hence, whenever $\mathcal{X}(\eta)=-\eta$, the right hand side of the equality strictly decreases in η, while its left hand side remains constant. As the slope of $\mathcal{X}(\eta)$ varies smoothly with changes in η, this implies that there is exactly one $\hat{\eta}<0$ such that $\mathcal{X}(\hat{\eta})=-\hat{\eta}$.

Note that if $\eta<\hat{\eta}$, $\mathcal{X}(\eta)<-\eta$, and $\mathcal{X}'(\eta)>0$. Similarly, if $\eta>\hat{\eta}$, $\mathcal{X}'(\eta)<0$. As the denominator of (32) is greater than $1-\delta$ and as $\lim_{\eta\to+\infty}f(\eta)=\lim_{\eta\to-\infty}f(\eta)=0$, one has $\lim_{\eta\to-\infty}\mathcal{X}(\eta)=\lim_{\eta\to+\infty}\mathcal{X}(\eta)=0$.

On the basis of (33), one has

$$\lim_{\eta\to\infty}\mathcal{X}'(\eta)=\lim_{\eta\to\infty}\mathcal{X}(\eta)(-\eta)-\lim_{\eta\to\infty}[\mathcal{X}(\eta)]^2.$$

Since $\lim_{\eta\to\infty}\mathcal{X}(\eta)=0$, $\lim_{\eta\to\infty}[\mathcal{X}(\eta)]^2=0$. Observe also that

$$\mathcal{X}(\eta)(-\eta)=\frac{\delta f(\eta)(-\eta)}{1-\delta(1-F(\eta))}=\frac{\delta f'(\eta)}{1-\delta(1-F(\eta))}.$$

Using the well-known fact that $\lim_{\eta\to\infty}f'(\eta)=0$, $\lim_{\eta\to\infty}\mathcal{X}(\eta)(-\eta)=0$. Hence, $\lim_{\eta\to\infty}\mathcal{X}'(\eta)=0$. By the same reasoning—and as $\delta<1$—$\lim_{\eta\to-\infty}\mathcal{X}'(\eta)$ is zero.

Differentiating Eq. (33) gives

$$\mathcal{X}''(\eta)=-\eta\mathcal{X}'(\eta)-2\mathcal{X}'(\eta)\mathcal{X}(\eta)-\mathcal{X}(\eta). \tag{34}$$

Using Eq. (33), to rewrite the above yields

$$\mathcal{X}''(\eta)=\mathcal{X}(\eta)[(\eta+\mathcal{X}(\eta))(\eta+2\mathcal{X}(\eta))-1]. \tag{35}$$

Using our earlier finding that $\lim_{\eta\to-\infty}\mathcal{X}(\eta)=0$, the term between square brackets goes to infinity as η goes to minus infinity. Recall that for any finite η, $\mathcal{X}(\eta)>0$. It thus follows from (35) that for any finite and sufficiently low value of η, $\mathcal{X}''(\eta)>0$. For $\mathcal{X}'(\hat{\eta})=0$, it follows from (34) that $\mathcal{X}''(\hat{\eta})<0$. By continuity, there exists at least one $\eta^m\in(-\infty,\hat{\eta})$ such that $\mathcal{X}''(\eta^m)=0$. Differentiating (34), and evaluating at the point $\eta=\eta^m$, one has

$$\mathcal{X}'''(\eta)\big|_{\eta=\eta^m}=-\mathcal{X}'(\eta^m)-2(\mathcal{X}'(\eta^m))^2<0,$$

where the inequality follows from the fact that $\mathcal{X}'(\eta^m)>0$, because $\eta^m<\hat{\eta}$. We conclude that η^m is unique.

Recall that $\hat{\eta}<0$. Suppose $\lim_{\delta\to1}\mathcal{X}'(\hat{\eta})=0$ for some $\hat{\eta}\in(-\infty,0)$. It follows from (33) that

$$\lim_{\delta\to1}\mathcal{X}'(\hat{\eta})=0\Leftrightarrow-\hat{\eta}=\lim_{\delta\to1}\mathcal{X}(\hat{\eta})=\frac{f(\hat{\eta})}{F(\hat{\eta})}=r(\hat{\eta}).$$

It is easy to check that $\partial r(\eta)/\partial\eta=-r(\eta)(r(\eta)+\eta)$. Hence, $r'(\hat{\eta})=0$. This, however, contradicts the fact that $r'(\eta)<0\forall\eta\in(-\infty,\infty)$ (see Greene, 1993, Theorem 22.2). Thus, $\lim_{\delta\to1}\hat{\eta}=-\infty$.

Observe that $\lim_{\delta\to0}\mathcal{X}(\eta)=0\ \forall\eta$. Hence, $\lim_{\delta\to0}\mathcal{X}'(\eta^m)=0$.

As $\eta^m<\hat{\eta}$ and $\lim_{\delta\to1}\hat{\eta}=-\infty$, $\lim_{\delta\to1}\eta^m=-\infty$. Therefore

$$\lim_{\delta\to1}\mathcal{X}(\eta^m)=\frac{\lim_{\eta^m\to-\infty}f(\eta^m)}{\lim_{\eta^m\to-\infty}F(\eta^m)}=\infty,$$

where the last equality follows from l'Hôpital's rule and the fact that $f'(\eta)=-\eta f(\eta)$. It follows from (34) that

$$\mathcal{X}''(\eta^m)=0\Leftrightarrow\eta^m=-\mathcal{X}(\eta^m)\left(\frac{1}{\mathcal{X}'(\eta^m)}+2\right). \tag{36}$$

Recall that $\mathcal{X}'(\eta^m)=-\mathcal{X}(\eta^m)[\eta^m+\mathcal{X}(\eta^m)]$. Replacing η^m on the right-hand side of this equality by the right-hand side of the last equality in (36), and rearranging, one has

$$\frac{[\mathcal{X}'(\eta^m)]^2}{1+\mathcal{X}'(\eta^m)}=[\mathcal{X}(\eta^m)]^2.$$

As $\lim_{\delta\to1}\mathcal{X}(\eta^m)=\infty$, $\lim_{\delta\to1}[\mathcal{X}(\eta^m)]^2=\infty$. Thus $\lim_{\delta\to1}[\mathcal{X}'(\eta^m)]^2/(1+\mathcal{X}'(\eta^m))=\infty$, which implies that $\lim_{\delta\to1}\mathcal{X}'(\eta^m)=\infty$. □

Call *LHS* (*RHS*) the left-hand side (respectively right-hand side) of Eq. (5) after replacing μ^{LF} by μ, and observe that

$$\frac{\partial LHS}{\partial\mu}=1$$

and that

$$\frac{\partial RHS}{\partial \mu} = \kappa_1\kappa_2\mathcal{X}'(\kappa_1(\mu-\overline{\theta})) = \frac{1-\alpha}{1+\alpha}\mathcal{X}'(\kappa_1(\mu-\overline{\theta})),$$

where the last equality follows from the fact that $\beta/2\alpha = 1/(1+\alpha)$. From Lemma 3, we know that $\mathcal{X}'(\kappa_1(\mu^{LF}-\overline{\theta}))$ is maximal when $\kappa_1(\mu^{LF}-\overline{\theta}) = \eta^m$.[31] As $\mathcal{X}''(\cdot)>0$ when $\kappa_1(\mu^{LF}-\overline{\theta})<\eta^m$, as $\mathcal{X}''(\cdot)<0$ when $\kappa_1(\mu^{LF}-\overline{\theta}) \in (\eta^m,\hat{\eta})$ and as $\mathcal{X}'(\cdot)<0$ when $\kappa_1(\mu^{LF}-\overline{\theta})>\hat{\eta}$, it follows that $\forall\overline{\theta}$, there exists a unique switching equilibrium if and only if

$$\left.\frac{\partial RHS}{\partial \mu}\right|_{\mu=\mu^{LF}=\eta^m/\kappa_1+\overline{\theta}} \le 1 \Leftrightarrow \mathcal{X}'(\eta^m) \le \frac{1+\alpha}{1-\alpha} \Leftrightarrow \frac{\sigma_\theta^2}{\sigma_\epsilon^2} \ge \frac{1}{2}[\mathcal{X}'(\eta^m)-1],$$

where the last equivalence follows from the fact that $\alpha = \sigma_\theta^2/(\sigma_\theta^2+\sigma_\epsilon^2)$.

Recall that

$$g(\mu) = \mu - \kappa_2\mathcal{X}(\kappa_1(\mu-\overline{\theta})) \tag{37}$$

and observe that Eq. (5) is equivalent to $g(\mu^{LF})=0$. If $\mu<0$, $g(\mu)<0$. Thus, $\mu^{LF}>0$. Hence, if $\overline{\theta}\le 0$, $\kappa_1(\mu^{LF}-\overline{\theta})>0$. It then follows from Lemma 3 that $\mathcal{X}'(\kappa_1(\mu^{LF}-\overline{\theta}))<0$. Hence, if $\overline{\theta}\le 0$, there exists a unique switching equilibrium.

Suppose that if Player i waits, she perfectly learns the state of the world, which gives an upper bound on the value of learning. Player i's gain of waiting then equals $\delta\Pr(\theta>0|\mu_i)E(\theta|\mu_i,\theta>0)$. Observe that for high enough a μ_i, $E(\theta|\mu_i,\theta>0)\approx E(\theta|\mu_i)=\mu_i$. As $\delta<1$, there exists a $\overline{\mu}<\infty$ such that $\overline{\mu}=\delta\Pr(\theta>0|\overline{\mu})E(\theta|\overline{\mu},\theta>0)$. If $\mu>\overline{\mu}$ Player i strictly prefers to invest at time one. Hence, $\mu^{LF}<\overline{\mu}<\infty$. As $\mu^{LF}\in(0,\overline{\mu})$, $\kappa_1(\mu^{LF}-\overline{\theta})\to-\infty$, as $\overline{\theta}\to\infty$. It then follows from Lemma 3 that $\lim_{\overline{\theta}\to\infty}((1-\alpha)/(1+\alpha))\mathcal{X}'(\kappa_1(\mu^{LF}-\overline{\theta}))=0$. By continuity, there exists a $\overline{\theta}_u$ such that if $\overline{\theta}\ge\overline{\theta}_u$, $((1-\alpha)/(1+\alpha))\mathcal{X}'(\kappa_1(\mu^{LF}-\overline{\theta}))\le 1$. Thus if $\overline{\theta}>\overline{\theta}_u$ there exists a unique switching equilibrium.

Recall from Lemma 3 that $\lim_{\delta\to 1}\hat{\eta}=-\infty$ and that $\mathcal{X}'(\eta)<0$ when $\eta>\hat{\eta}$. Thus, if δ is close to one, μ cuts $\kappa_2\mathcal{X}(\kappa_1(\mu-\overline{\theta}))$ when $\mathcal{X}'(\cdot)<0$, in which case equilibrium is unique. By continuity, there exists a $\overline{\delta}<1$ such that $\mathcal{X}'(\kappa_1(\mu^{LF}-\overline{\theta}))\le 0$ for all $\delta\ge\overline{\delta}$. $\square$

A.4. Proof of the equivalence between Eqs. (7) and (8)

In this proof, k denotes the p.d.f. of some random variable. For example, $k(\theta)=f((\theta-\overline{\theta})/\sigma_\theta)$. k obviously depends on the studied random variable. For example, it follows from our section "Definitions and Preliminaries" that $\theta\sim N(\overline{\theta},\sigma_\theta^2)$ and that $\mu_i|\theta\sim N(\alpha\theta+(1-\alpha)\overline{\theta},\sigma_\mu^2)$. Hence, $k(\theta)\neq k(\mu_i|\theta)$. In that sense, it would be more precise to use the notation k^θ and $k^{\mu_i|\theta}$ to respectively denote the p.d.f.'s of θ and $\mu_i|\theta$. In this proof, however, we avoid this cumbersome notation. This should not cause confusion. Observe that Eq. (7) can be rewritten as

$$\begin{aligned}\frac{1}{2}W = &\int \Pr(\mu_i>\mu^c|\theta)\theta k(\theta)\,d\theta + \delta\int \Pr(\mu_j>\mu^c,\mu_i\in[\min\{\underline{\mu},\mu^c\},\mu^c]|\theta)\theta k(\theta)\,d\theta\\ &+\delta\int \Pr(\mu_j<\mu^c,\mu_i\in[\min\{\underline{\mu}^0,\mu^c\},\mu^c]|\theta)\theta k(\theta)\,d\theta.\end{aligned} \tag{38}$$

Observe also that

$$\int \Pr(\mu_i>\mu^c|\theta)\theta k(\theta)\,d\theta = \int\int_{\mu^c}^{\infty}\frac{k(\mu_i,\theta)}{k(\theta)}\,d\mu_i\,\theta k(\theta)\,d\theta = \int_{\mu^c}^{\infty}\int \theta k(\theta|\mu_i)\,d\theta\,k(\mu_i)\,d\mu_i.$$

Trivially, $\mu_i=\int\theta k(\theta|\mu_i)d\theta$. Hence, the first integral of (38) is equal to $\int_{\mu^c}^{\infty}\mu_i k(\mu_i)\,d\mu_i$.

The second integral of (38) can be rewritten as

$$\iint_{\mu^c}^{\infty}\int_{\min\{\underline{\mu},\mu^c\}}^{\mu^c}\frac{k(\mu_j,\mu_i,\theta)}{k(\theta)}\,d\mu_i\,d\mu_j\,\theta k(\theta)\,d\theta.$$

Changing the order of integration, the above integral can be rewritten as

$$\int_{\min\{\underline{\mu},\mu^c\}}^{\mu^c}\underbrace{\int_{\mu^c}^{\infty}\underbrace{\int\theta k(\theta|\mu_i,\mu_j)\,d\theta}_{E(\theta|\mu_i,\mu_j)}\,k(\mu_j|\mu_i)\,d\mu_j}_{\Pr(\mu_j>\mu^c|\mu_i)E(\theta|\mu_i,\mu_j>\mu^c)}\,k(\mu_i)\,d\mu_i.$$

Hence, the second integral of (38) is equal to $\int_{\min\{\underline{\mu},\mu^c\}}^{\mu^c}\Pr(\mu_j>\mu^c|\mu_i)E(\theta|\mu_i,\mu_j>\mu^c)k(\mu_i)d\mu_i$.

Using an identical procedure, the third integral of (38) can be rewritten as $\int_{\min\{\underline{\mu}^0,\mu^c\}}^{\mu^c}\Pr(\mu_j<\mu^c|\mu_i)E(\theta|\mu_i,\mu_j<\mu^c)k(\mu_i)d\mu_i$. $\square$

---

[31] As a unit increase in $\overline{\theta}$ leads to a translation of $\mathcal{X}(\cdot)$ to the right by one unit (as shown in Fig. 1), it follows that there exists a unique $\overline{\theta}$ such that $\kappa_1(\mu^{LF}-\overline{\theta})=\eta^m$.

A.5. Proof of Proposition 2

The proof of this proposition is almost entirely explained in the body of the text. We are left to prove Inequality (12). It follows from (22) and from Lemma 2 that

$$\Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c) = [1-F(x(\mu^c,\mu_i))][\mu_i + \kappa_2 h(x(\mu^c,\mu_i))].$$

Recall that $(1-F(z))h(z) = f(z)$, that $f'(z) = -f(z)z$ and that $x(\mu^c,\mu_i) = (\mu^c - \alpha\mu_i - (1-\alpha)\overline{\theta})/\sigma_o$. Hence, the derivative of the right-hand side with respect to μ^c equals

$$\frac{1}{\sigma_o} f(x(\mu^c,\mu_i)) \left[-\mu_i + \kappa_2 \frac{(1-\alpha)\overline{\theta} + \alpha\mu_i - \mu^c}{\sigma_o} \right],$$

which is positive if and only if the term between square brackets is. It is straightforward to show that $\kappa_2/\sigma_o = 1/(1+\alpha)$. This insight permits us to conclude that the term between square brackets is positive if and only if $(1-\alpha)\overline{\theta} - \mu_i > \mu^c$.

A.6. Proof of Proposition 3

Let $\tilde{\mu} \equiv (1-\alpha)\overline{\theta} - \mu^{LF}$ and recall that $x(\mu^{LF},\mu_i) = (\mu^{LF} - \alpha\mu_i - (1-\alpha)\overline{\theta})/\sigma_o$.

Lemma 4. *Suppose $\overline{\theta} \le 0$. If, additionally, $\mu_i \in [\underline{\mu}^1, \tilde{\mu}]$, Player i wants the social planner to implement a higher cutoff. If $\mu_i \in [\tilde{\mu}, \mu^{LF}]$, Player i wants the social planner to implement a lower cutoff.*

Proof. Suppose $\mu_i = \tilde{\mu}$. As $\overline{\theta} \le 0$ and as $\mu^{LF} > 0$, Player *i* waits at time one. Furthermore, it follows from Inequality (13) that her gain of waiting is maximal under a laissez-faire policy. One has

$$\begin{aligned} E(\theta | \mu_i = \tilde{\mu}, \mu_j = \mu^{LF}) &= E\left(\theta \middle| s_i = -\frac{\mu^{LF}}{\alpha}, s_j = \frac{\mu^{LF} - (1-\alpha)\overline{\theta}}{\alpha}\right) \\ &= 0 \\ &< E(\theta | \mu_i = \tilde{\mu}, \mu_j > \mu^{LF}), \end{aligned}$$

where the second equality follows from Eq. (23), from the fact that $\frac{\beta}{2\alpha} = 1/(1+\alpha)$ and that $1-\beta = (1-\alpha)/(1+\alpha)$. It then follows from Point 1 of Lemma 2 that $\underline{\mu}^1 < \tilde{\mu}$. The lemma then follows from Inequality (12). □

Lemma 5. *One has* $\lim_{\overline{\theta} \to -\infty} \mu^{LF} = \lim_{\overline{\theta} \to -\infty} (1-\alpha)\overline{\theta} - \underline{\mu}^1$.

Proof. Rewriting $x(\mu^{LF},\tilde{\mu})$ using $\kappa_2 = \frac{1}{2}\beta\sigma_2$ and $\frac{\beta}{2\alpha} = 1/(1+\alpha)$ verifies that $\tilde{\mu} = -\kappa_2 x(\mu^{LF},\tilde{\mu})$. Furthermore, using the definition of ϕ, and that $\underline{\mu}^1$ is implicitly defined through $\phi(\underline{\mu}^1) = 0$, one has $\underline{\mu}^1 = -\kappa_2 h(x(\mu^{LF},\underline{\mu}^1))$. Therefore

$$(1-\alpha)\overline{\theta} - \mu^{LF} - \underline{\mu}^1 = \tilde{\mu} - \underline{\mu}^1 = \kappa_2 (h(x(\mu^{LF},\underline{\mu}^1)) - x(\mu^{LF},\tilde{\mu})). \tag{39}$$

Furthermore

$$x(\mu^{LF},\tilde{\mu}) = x(\mu^{LF},\underline{\mu}^1) + \frac{\underline{\mu}^1 - \tilde{\mu}}{\sigma_2}. \tag{40}$$

Inserting (40) into (39), and rearranging, yields

$$(\tilde{\mu} - \underline{\mu}^1)(1 + \tfrac{1}{2}\beta) = \kappa_2 (h(x(\mu^{LF},\underline{\mu}^1)) - x(\mu^{LF},\underline{\mu}^1)).$$

Recall that $h'(\eta) = h(\eta)[h(\eta) - \eta]$, that $h'(\eta) \in (0,1)$, and that $\lim_{\eta\to\infty} h(\eta) = \infty$. Hence, $h(\eta) > \eta$ and $\lim_{\eta\to\infty}(h(\eta) - \eta) = 0$. Since $\underline{\mu}^1 < 0$, $\lim_{\overline{\theta}\to-\infty} x(\mu^{LF},\underline{\mu}^1) = \infty$, which implies that

$$\lim_{\overline{\theta}\to-\infty} \kappa_2 (h(x(\mu^{LF},\underline{\mu}^1)) - x(\mu^{LF},\underline{\mu}^1)) = 0.$$

As $1 + \frac{1}{2}\beta > 0$, this implies that $\lim_{\overline{\theta}\to-\infty}(\tilde{\mu} - \underline{\mu}^1) = 0$. Rewriting this last equality yields the lemma. □

In the body of the text, we argued that $\Delta(\underline{\mu},\underline{\mu}) < 0$ and that $\Delta(\mu^{LF},\mu^{LF}) = 0$. Suppose that $\mu_i = \overline{\mu} > 0$. By definition, this means that if Player *i* gets bad news, she is indifferent between investing and not investing. Hence, one can think of her as someone who will always invest at time two—even if the other player did not do so at time one. Hence, $\Delta(\overline{\mu},\overline{\mu}) = (1-\delta)\overline{\mu} > 0$. If equilibrium is unique, those results imply that $\Delta(\mu^c,\mu^c) > 0$ whenever $\mu^c > \mu^{LF}$. In turn, this implies that the first term of Eq. (10) is negative if equilibrium is unique and if $\mu^c > (\mu^{LF},\overline{\mu}]$. Recall from Proposition 1 that equilibrium is unique whenever $\overline{\theta} < 0$. It follows from Lemma 5 that if $\overline{\theta} \to -\infty$ the third term of (10) is non-positive. Hence, $dU/d\mu^c < 0$ if $\mu^c \in (\mu^{LF},\overline{\mu}]$ and if $\overline{\theta}$ is sufficiently negative.

We are left to show that the social planner does not want to implement a $\mu^c > \overline{\mu}$ when the prior mean is sufficiently negative. From the explanations provided in the body of this paper, we know that if $\mu^c > \overline{\mu}$ all types between μ^c and $\underline{\mu}^0$

always invest at time two and that all types between $\underline{\mu}^1$ and $\underline{\mu}^0$ invest at time two only if they received good news. Hence

$$U = \int_{\mu^c}^{\infty} \mu_i f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i + \delta \int_{\underline{\mu}^0}^{\mu^c} \mu_i f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i + \delta \int_{\underline{\mu}^1}^{\underline{\mu}^0} \Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c) f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i.$$

$$< \int_{\underline{\mu}^0}^{\infty} \mu_i f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i + \delta \int_{\underline{\mu}^1}^{\underline{\mu}^0} \Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c) f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i \equiv \overline{U}.$$

Observe that $\overline{U} = U$ when $\mu^c = \overline{\mu}$. Hence, it suffices to show that $d\overline{U}/d\mu^c < 0$ when $\mu^c > \overline{\mu}$ and when $\overline{\theta}$ is sufficiently negative. Taking into account the fact that $E(\theta | \underline{\mu}^0, \mu_j < \mu^c) = E(\theta | \underline{\mu}^1, \mu_j > \mu^c) = 0$, it follows that $\forall \mu^c > \overline{\mu}$

$$\frac{d\overline{U}}{d\mu^c} = \frac{d\underline{\mu}^0}{d\mu^c}(1-\delta)\Pr(\mu_j > \mu^c | \underline{\mu}^0) E(\theta | \underline{\mu}^0, \mu_j > \mu^c) f\left(\frac{\underline{\mu}^0 - \overline{\theta}}{\sigma_\mu}\right) + \delta \int_{\underline{\mu}^1}^{\underline{\mu}^0} \frac{\partial}{\partial \mu^c} (\Pr(\mu_j > \mu^c | \mu_i) E(\theta | \mu_i, \mu_j > \mu^c)) f\left(\frac{\mu_i - \overline{\theta}}{\sigma_\mu}\right) d\mu_i. \quad (41)$$

It follows from Lemma 2 that $\underline{\mu}^0$ is implicitly defined by

$$\underline{\mu}^0 - \kappa_2 r\left(\frac{\mu^c - \alpha \underline{\mu}^0 - (1-\alpha)\overline{\theta}}{\sigma_o}\right) = 0.$$

It then follows from the implicit function theorem that

$$\frac{d\underline{\mu}^0}{d\mu^c} = -\frac{-\frac{\kappa_2}{\sigma_o} r'(\cdot)}{1 + \frac{\alpha \kappa_2}{\sigma_o} r'(\cdot)} = \frac{\frac{1}{1+\alpha} r'(\cdot)}{1 + \frac{1}{2}\beta r'(\cdot)} < 0,$$

where the inequality follows from the fact that $r'(\cdot) \in (-1, 0)$ and that $\beta \in [0, 1]$. Hence, the first term of Eq. (41) is negative. It follows from Lemma 5 that the gain of waiting of the inframarginal types (with the exception of $\underline{\mu}^1$) is decreasing in μ^c when $\overline{\theta} \to -\infty$. Hence, The second term of (41) is non-positive. Hence, $d\overline{U}/d\mu^c < 0$ when $\mu^c > \overline{\mu}$ and when $\overline{\theta}$ is sufficiently negative. □

References

Angeletos, G.-M., Hellwig, C., Pavan, A., 2007. Dynamic global games of regime change: learning, multiplicity and the timing of attacks. Econometrica 75, 711–756.

Angeletos, G.-M., Werning, I., 2006. Crises and prices: information aggregation, multiplicity, and volatility. American Economic Review 96, 1720–1736.

Aoyagi, M., 1998. Equilibrium delay in a simple dynamic model of investment. Economic Theory 12, 123–146.

Banerjee, A., 1992. A simple model of herd behavior. Quarterly Journal of Economics 107, 797–817.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informational cascades. Journal of Political Economy 100, 992–1026.

Caplin, A., Leahy, J., 1998. Miracle on sixth avenue: information externalities and search. Economic Journal 108, 60–74.

Chamley, C., 2004a. Delays and equilibria with large and small information in social learning. European Economic Review 48, 477–501.

Chamley, C., 2004b. Rational Herds, Economic Models of Social Learning. Cambridge University Press, UK.

Chamley, C., Gale, D., 1994. Information revelation and strategic delay in a model of investment. Econometrica 62, 1065–1085.

Choi, J.P., 1997. Herd behavior, the "penguin" effect and the suppression of informational diffusion: an analysis of informational externalities and payoff interdependencies. Rand Journal of Economics 28, 407–425.

Dasgupta, A., 2007. Coordination and delay in global games. Journal of Economic Theory 134, 195–225.

DeGroot, M.H., 1970. Optimal Statistical Decisions. McGraw-Hill Publishing Co., Irwin.

Doyle, M., 2010. Informational externalities, strategic delay, and optimal investment subsidies. Canadian Journal of Economics 43, 941–966.

Eyster, E., Rabin, M., 2010. Naive herding in rich-information settings. American Economic Journal: Microeconomics 2, 221–243.

Frisell, L., 2003. On the interplay of informational spillovers and payoff externalities. Rand Journal of Economics 34, 582–592.

Gossner, O., Melissas, N., 2006. Informational cascades elicit private information. International Economic Review 47, 297–325.

Greene, W.H., 1993. Econometric Analysis, Second ed. Macmillan, New York.

Greenspan, A., 2008. The Age of Turbulence, Adventures in a New World. Penguin Books, New York.

Gul, F., Lundholm, R., 1995. Timing and the clustering of agents' decisions. Journal of Political Economy 103, 1039–1066.

Gunay, H., 2008. The role of externalities and information aggregation in market collapse. Economic Theory 35, 367–379.

Heidhues, P., Melissas, N., 2010. Technology Adoption, Social Learning, and Economic Policy. GESY Discussion Paper 306.

Hendricks, K., Kovenock, D., 1989. Asymmetric information, information externalities, and efficiency: the case of oil exploration. Rand Journal of Economics 20, 164–182.

Johnson, N.L., Kotz, S., Balakrishnan, N., 1995. Continuous Univariate Distributions. John Wiley and Sons, NJ.

Levin, D., Peck, J., 2008. Investment dynamics with common and private values. Journal of Economic Theory 143, 114–139.

Morris, S., Shin, H.S., 1999. A theory of the onset of currency attacks. In: Agénor, P.-R., Miller, M., Vines, D., Weber, A. (Eds.), Asian Financial Crises: Onset, Contagion and Consequences, Cambridge University Press, UK.

Morris, S., Shin, H.S., 2000. Rethinking multiple equilibria in macroeconomic modelling. NBER Macroeconomics Annual 15, 139–161.

Morris, S., Shin, H.S., 2002. Social value of public information. American Economic Review 92, 1521–1534.

Morris, S., Shin, H.S., 2003. Global games: theory and applications. In: Dewatripont, M., Hansen, L., Turnovsky, S. (Eds.), Advances in Economics and Econometrics (Procedures of the Eighth World Congress of the Econometric Society), Cambridge University Press, UK.

Morris, S., Shin, H.S., 2004. Coordination risk and the price of debt. European Economic Review 48, 133–153.

Morris, S., Shin, H.S., 2005. Central bank transparency and the signal value of prices. Brookings Papers on Economic activity 2, 1–66.

Park, A., Sabourian, H., 2011. Herding and contrarian behavior in efficient financial markets. Econometrica 79, 973–1026.

Zhang, J., 1997. Strategic delay and the onset of investment cascades. Rand Journal of Economics 28, 188–205.