Learning from Private and Public Observations of Others' Actions*

Manuel Amador[†] and Pierre-Olivier Weill[‡]

June 2, 2009

Abstract

We study the diffusion of dispersed private information in a large economy. We assume that agents learn from the actions of others through two channels: a public channel, that represents learning from prices, and a private channel that represents learning from local interactions. We show that, when agents learn *only* from the public channel, an initial release of public information increases agents' total knowledge at all subsequent times and increases welfare. When a private learning channel is present, this result is reversed: more initial public information reduces agents asymptotic knowledge by an amount in order of $\log(t)$ units of precision. When agents are sufficiently patient, this reduces welfare.

*For fruitful discussions and suggestions, we'd like to thank George-Marios Angeletos, Bruno Biais, Ricardo Caballero, Amil Dasgupta, Chris Edmond, Raquel Fernandez, Doireann Fitzgerald, Alessandra Fogli, Boyan Jovanovic, Douglas Gale, Christian Hellwig, Thomas Mariotti, Andy Skrzypacz, Stijn Van Nieuwerburgh, Dimitri Vayanos, Laura Veldkamp, Xavier Vives, Ivan Werning, and seminar participants at ECB, UCSD, Princeton, NYU, Harvard Junior Lunch, Stanford Junior Lunch, University of Montreal, LSE, Toulouse, Columbia, MIT, Penn State, USC, Society of Economics Dynamics Meeting (2006), Stanford Institute for Theoretical Economics (2006), Southwestern Economic Theory Conference (2007) at USC, Minnesota Workshop in Macroeconomic Theory (2006), Federal Reserve Bank of Cleveland Summer Workshop in Money, Banking, and Payment (2006), IESE's Conference on Complementarities and Information (2007), for their comments. We are grateful to the Associate Editor and two anonymous referees for comments that greatly improved the paper. We also thank Romans Pancs for excellent research assistance.

†Department of Economics, Stanford University, Stanford, CA 94305-6072, e-mail: amador@stanford.edu

‡Department of Economics, University of California Los Angeles, Bunche 8283, Los Angeles, Box 951477, CA 90095, e-mail: poweill@econ.ucla.edu

1 Introduction

Households and firms have knowledge from their local markets about aggregate economic conditions. This dispersed private knowledge diffuses in the population over time as households and firms learn from each other valuable bits of information that are not yet known to everyone. Some channels of information diffusion are private: for example private information comes from observing the actions of others in local markets. Other channels are public, such as the information aggregated in prices, or in the macroeconomic figures published by agencies. In this paper, we study how private information dispersed in a large economy diffuses through such private and public channels. The main result of the paper is to show that these two channels of information diffusion generate different diffusion dynamics and more importantly, generate opposite social values of public information. When agents learn only from public channels, a release of public information would always increase welfare. By contrast, when agents also learn from private channels, a release of public information can reduce welfare.

Our baseline model is set in continuous time and builds on the discrete-time environments of Vives (1993, 1997). We consider a continuum of agents who, at time zero, receive both public and private signals about the state of the world. This is the only exogenous source of information in the model. After time zero, each agent takes an action at every moment until some random time when the state of the world is revealed and her payoff is realized. We assume that an agent's payoff decreases with the distance between the sequence of actions and the revealed state, and it is independent of the actions of any other agent. After receiving initial information, but before the state of the world is revealed, each agent learns continuously by observing two noisy signals about the actions of others. The first signal is private, only observed by the agent: this is the private learning channel and represents information gathered through private communication and local interactions. The second signal is public, shared with everyone: it is the public learning channel and represents an endogenous aggregate variable, such as a price or some macroeconomic aggregate.

We solve for an equilibrium in which agents eventually learn the truth. Like in the herding literature, there is a public information externality: better public information reduces the informational content of both the public and the private channels and slows down learning. Indeed, with an increase in the precision of public information, each agent's action becomes more sensitive to the public information that is already known by everyone, and hence useless for the purpose of learning. At the same time, each agent's action becomes less sensitive to private information. This makes it harder to glean private information from the noisy observation of agents' actions, and thus generates a negative learning externality: information will diffuse more slowly through all channels. The private learning channel amplifies the externality: since each agent accumulates less information from the private channel, her actions become less sensitive to private information, slowing down information diffusion even more, and a negative feedback loop ensues.

Now imagine that a benevolent agency holds some relevant, but partial, information and ponders whether to release it publicly at time zero. A release has the direct beneficial effect of making agents' current decisions better informed. However, it has the negative effect of slowing the diffusion of private information. We show that when agents only learn through the public channel, the negative externality is sufficiently weak that a release of public information increases the amount known at all times and is always socially beneficial. By contrast, when agents also learn through the private channel, the negative externality is amplified: it so slows information diffusion that all agents end up less informed in the long run. More precisely, we show that increasing the precision of the initial public signal by one unit reduces knowledge at time t by an amount in order of $\log(t)$ units of precision.

The social benefit of releasing public information depends then on the trade off between increasing the amount known currently by all agents, and reducing the amount known in the future. We show that if agents have enough time to learn from others' actions before the state of the world is revealed, then a given marginal increase in the precision of the initial public signal is always socially costly.

In the final section of the paper we explore the robustness of our results by analyzing the socially optimal diffusion of information: we study the problem of a planner that can choose the sensitivity of the agents' actions to their private and public forecasts, as in Vives (1997). We show that, after the planner corrects the information externality, public information always increases *ex-ante* social welfare. Surprisingly, welfare is reduced *ex-post*: the planner finds it optimal to make agents less informed in the long run, by the same $\log(t)$ amount as in the decentralized equilibrium.

Related Literature

Our work is related to the recent literature on the social value of public information. As initially shown by Morris and Shin (2002), public information can reduce welfare in the presence of a payoff externality. Angeletos and Pavan (2007) provided a complete characterization of the effect in general linear-quadratic models, and Hellwig (2005) studied the implications for a monetary economy. However, the models used so far have been essentially static, and abstracted from learning. Our contribution is to analyze an alternative mechanism based on a dynamic information externality: in our baseline model there are no payoff externalities, but public information slows down the diffusion of private information in the population.[1]

The social learning literature, started by Banerjee (1992) and Bikhchandani et al. (1992), has been concerned with information externalities. Its central result is the possibility of informational cascades and herds: agents may choose to disregard their private information, acting solely on the basis of the public information, and take the "wrong" action. Thus public information, by facilitating the emergence of herds, can reduce welfare. However, the standard herding models are sequential move games, and the appearance of cascades and herds requires bounded beliefs and a discrete set of actions. Our model is closest to Vives (1993) and Vives (1997) instead, where beliefs are unbounded and actions lie in a continuous space. The maintained assumption among these social-learning papers is that private information diffuses through public channels. Our paper allows information to diffuse through both public and private channels, and we show that this has different implications for dynamics and welfare.

Burguet and Vives (2000) present a model where agents exert effort to collect private information, and show that public information reduces the incentives of agents to gather new information and can reduce welfare. In our model, agents do not collect new information: we focus instead on the question of how information, that is already in the hands of agents, aggregates and diffuses.[2] In many situations it is reasonable to expect that

[1] Morris and Shin (2005) set up a model in which a central bank ends up learning less from the actions of private agents after disclosing public information. But since private agents do not accumulate information over time, they do not suffer from our dynamic learning externality. In fact, absent payoff externalities, public information improves social welfare in their model.

[2] Our results regarding the social value of public information complements theirs, and would apply in a costly information acquisition set up after the agents have stopped gathering information but could still observe signals about the aggregate market behavior.

agents can learn at no cost from each other, for example through their ongoing market interactions or through the observation of public prices.

Strategic experimentation models study the interplay between information acquisition and information diffusion. Recent examples include Bolton and Harris (1999), Bergemann and Välimäki (2000), Décamp and Mariotti (2004), Keller et al. (2005), and Rosenberg et al. (2007). The key inefficiency is a free-riding problem: given that other agents' actions generate observable new information, an individual agent has less incentive to learn on his own by taking the costly action. No such free riding is at a play in our model where information is generated at no cost, and where an individual agent's decision problem is not affected by other agents' contemporaneous and future decisions, but only by their past decisions.

Some recent work on social learning has focused on learning in networks: Bala and Goyal (1998), Gale and Kariv (2003), and Smith and Sørensen (2005) study deterministic networks with a finite number of agents, Banerjee and Fudenberg (2004) provide a continuum-of-agents setup, and DeMarzo et al. (2003) propose a network of boundedly rational agents. The private learning channel of the present paper is arguably a reduced form model of local interactions through networks. However, our model is tractable enough to address questions that would be more difficult in an explicit network model.

Another body of research uses search-and-matching models to study how agents learn from local interactions with others (see for example Wolinsky 1990 seminal work and the recent work of Kircher and Postlewaite 2008). The issue of convergence to the truth has also been addressed in Green (1991), Blouin and Serrano (2001), Duffie and Manso (2007), and Golosov et al. (2007). Wallace (1997), Katzman et al. (2003), Araujo and Shevchenko (2006), and Araujo and Camargo (2006) address learning about a money supply shock in a Trejos and Wright's (1995) random-matching model. Our setup, although somewhat simpler, is related to this literature as the private signals about aggregate actions can be interpreted as the result of random local interactions. The benefit of this simplification is that we can characterize transitional dynamics of beliefs and compute the social value of public information.

The rest of this paper is organized as follows. Section 2 introduces the setup. Section 3 solves for an equilibrium. Section 4 studies the impact on dynamics and welfare of changes in the quality of public information. Section 5 analyzes optimal information diffusion. Section 6 concludes. The appendix collects all the proofs not in the main text. An addendum to this paper (Amador and Weill, 2007a), presents additional results and extensions of our model.

2 Set up

Time is continuous and runs forever. We fix a probability space $\{\Omega, \mathcal{G}, Q\}$ together with an information filtration $\{\mathcal{G}_t, t \geq 0\}$ satisfying the usual conditions (Protter, 1990). The economy is populated by a $[0,1]$-continuum of agents whose payoff depend on some unknown state of the world $x \in \mathbb{R}$. At each time before some exponentially distributed "day of reckoning" $\tau > 0$, with parameter λ, each agent takes an action $a_{it} \in \mathbb{R}$. At time τ, the state of the world x is revealed and each agent receives the payoff

$$- \int_0^{\tau} (a_{it} - x)^2 \, dt.$$

Agents are endowed with a diffused common prior that x is normally distributed with mean zero and zero precision.[3] Over time, they observe a public signal Z_t and a private signal z_{it}. The initial realizations of these signals are

$$Z_0 = x + \frac{W_0}{\sqrt{P_0}} \quad \text{and} \quad z_{i0} = x + \frac{\omega_{i0}}{\sqrt{p_0}} \tag{1}$$

where W_0 and $(\omega_{i0})_{i \in [0,1]}$ are normally distributed with mean zero and variance one, pairwise independent, and independent of everything else. In equation (1), P_0 and p_0 represent the respective precisions of the public and the private signal.

The initial signal Z_0 represents information released by a public agency. By continuously varying its precision, P_0, we will obtain the impact on diffusion and welfare of varying the size of an information release. The continuum of initial private signals, $(z_{i0})_{i \in [0,1]}$, make agents asymmetrically informed about x, and represents dispersed information about aggregate economic conditions.

At all times after time zero the public and the private signal evolve according to the

[3]The assumption of a diffused prior is without loss of generality. In an online Addendum available from the authors website, we consider the general case of a prior with a strictly positive precision, and show that our results for the dynamics of beliefs and welfare do not change.

stochastic differential equations:

$$dZ_t = A_t dt + \frac{dW_t}{\sqrt{P_\varepsilon}} \quad \text{and} \quad dz_{it} = A_t dt + \frac{d\omega_{it}}{\sqrt{p_\varepsilon}}, \tag{2}$$

where $A_t \equiv \int_0^1 a_{it}\, di$ is the cross-sectional average action at time t, and where W and $(\omega_i)_{i \in [0,1]}$ are pairwise independent Wiener processes with initial conditions W_0, ω_{i0}, and are independent from everything else.

There is a key difference between the initial signal realization (1) and the subsequent signal realizations (2): the former is centered around the true state of the world, x, while the later is centered around the average action, A_t. This means that, after time zero, all the information about x that agents learn comes from others, and there is no additional arrival of "new" information. Indeed, the realizations dZ_t and dz_{it} are signals about the information that others accumulated in the past so, ultimately, they are signals about the initial realizations Z_0 and z_{i0}.

The public signal, dZ_t, could represent the information conveyed by some endogenous aggregate variable (such as prices). The private signal, dz_{it}, on the other hand, captures the decentralized gathering of information. One could think, for instance, of local interaction and of private communication, such as gossips.[4, 5]

Given a process A for the average action, we let $\mathcal{G}_{it}$ be the filtration generated by $\{(Z_s, z_{is}), 0 \leq 0 \leq t\})$, representing all the information available to agent $i \in [0,1]$ at any time $t > 0$. Because agents are infinitesimal, their actions do not affect the average action process A, and hence do not affect the information they receive. So, the agents' inter-temporal problems are essentially static, and together with their quadratic payoffs, this implies that optimal actions are the expectation

$$a_{it} = \mathbb{E}\left[x \mid \mathcal{G}_{it}\right] \tag{3}$$

of the random variable x, conditional on their information filtration $\{\mathcal{G}_{it}, \, t \geq 0\}$. Finally,

[4] In previous work, Amador and Weill (2006), we suggest that specification (2) may arise when each agent continuously observes, with idiosyncratic noise, the action of other randomly chosen agents. Intuitively, observing the action of a randomly chosen agent amounts to sampling from a distribution centered around the average action, A_t. When the time between periods and the precision of the noise go to zero at the same rate, we informally arrive at specification (2).

[5] In Amador and Weill (2007b), we provide a different interpretation of the endogenous private signal, dz_{it}: they are generated because agents receive exogenous private signals about the noise in public endogenous aggregates. We show how such signals about the noise naturally arise in a monetary economy.

in an equilibrium, these individual actions have to generate the average:

$$A_t = \int_0^1 a_{it}\, di. \tag{4}$$

We summarize all the above in the following:

Definition 1. *An* equilibrium *is a collection of processes a_i and A solving (2), (3), and (4).*

3 An Equilibrium

We now show that there exists an equilibrium in which agent i's action at any time is the convex combination of two forecasts of the state of the world: a public forecast shared with everyone in the economy, which we denote by $\hat{X}_t$, and a private forecast, denoted by $\hat{x}_{it}$, containing all the information observed by agent i and no one else.

Let us guess for now that these forecasts are normally distributed, independent given x, and that their precisions are common knowledge. Denoting by P_t and p_t the precision of the public and the private forecast, Bayesian updating implies that the action taken by agent i at time t is, then,

$$a_{it} = \mathbb{E}[x|\mathcal{G}_{it}] = \frac{P_t}{P_t + p_t}\hat{X}_t + \frac{p_t}{P_t + p_t}\hat{x}_{it}\ , \tag{5}$$

a "precision weighted" convex combination of the public and private forecasts. The public forecast is, of course, the same for every agent. The private forecasts, on the other hand, are unbiased and based on independent private information: thus, their cross-sectional average must equal x. These observations mean that the average action is

$$A_t = \int_0^1 a_{it} di = \frac{P_t}{P_t + p_t}\hat{X}_t + \frac{p_t}{P_t + p_t}x.$$

Recall that the public and private signals dZ_t and dz_{it} of equation (2) are centered around the average action A_t. But since the public forecast, $\hat{X}_t$, and the precisions, P_t and p_t, are all common knowledge, we obtain an informationnally equivalent set of signals

$$xdt + \frac{dW_t}{\sqrt{P_\varepsilon \left(\frac{p_t}{P_t+p_t}\right)^2}} \quad \text{and} \quad xdt + \frac{d\omega_{it}}{\sqrt{p_\varepsilon \left(\frac{p_t}{P_t+p_t}\right)^2}} \tag{6}$$

after first subtracting off $P_t/(P_t+p_t)\hat{X}_t\,dt$ from dZ_t and dz_{it}, and then dividing by $p_t/(P_t+p_t)$. Equation (6) defines two equivalent "transformed" public and private signals which have the convenient properties of being centered around x and independent given x.

It it then straightforward to verify our guess. We let the public forecast $\hat{X}_t$ be the expectation of x conditional on common prior and the history of the "transformed" public signal defined by the left side of equations (1) and (6) . Similarly, the private forecast $\hat{x}_{it}$ is the expectation of x conditional on the history of the "transformed" private signal, defined by the right side of equations (1) and (6). The precisions of the public and private forecasts, P_t and p_t, are readily characterized by a system of Ordinary Differential Equations (ODEs):

$$dP_t = P_\varepsilon \left(\frac{p_t}{P_t + p_t} \right)^2 dt \tag{7}$$

$$dp_t = p_\varepsilon \left(\frac{p_t}{P_t + p_t} \right)^2 dt\ , \tag{8}$$

with initial conditions P_0 and p_0, respectively. The intuition follows standard Bayesian updating formulas with independent signals: for instance, the change in the public precision, dP_t, is equal to the precision of the transformed public signal $P_\varepsilon \left(\frac{p_t}{p_t+P_t} \right)^2$.

The ODEs show that the informativeness of the public and the private signals at time t is a function of the precisions P_t and p_t of the public and private forecasts. This informativeness decreases with P_t and increases with p_t. That is, the more the agents know privately, the more informative the new signals become, and the faster agents learn. Improvements in public knowledge have the opposite effect: they slow down subsequent learning. This suggests that public and private learning affect the diffusion of information in the economy differently.

Closed form solution and learning asymptotic

Note that ODE (7) is equal to ODE (8) multiplied by $P_\varepsilon/p_\varepsilon$. So, as long as $p_\varepsilon > 0$, this implies that $P_t - (P_\varepsilon/p_\varepsilon)p_t$ stays constant over time, and $(P_t - P_0) = (p_t - p_0)P_\varepsilon/p_\varepsilon$. Plugging this into equation (8), we obtain

$$\dot{p}_t = p_\varepsilon \left(\frac{p_t}{\alpha + \beta p_t} \right)^2 \tag{9}$$

where $\alpha = P_0 - (P_\varepsilon/p_\varepsilon)p_0$ and $\beta = 1 + P_\varepsilon/p_\varepsilon$. Hence given an initial condition for the precision p_0 of the private forecast, and using equation (9), it is possible to solve analytically for p_t.

Theorem 1 (A closed form solution)**.** *There exists an equilibrium in which an agent's belief can be decomposed into two independent public and private forecasts with respective precisions P_t and p_t solving:*

- *for $p_e = 0$:*

$$p_t = p_0 \quad \text{and} \quad P_t + p_t = \left(3p_0^2 P_\varepsilon t + (p_0 + P_0)^3\right)^{1/3}, \tag{10}$$

- *for $p_e > 0$:*

$$H(p_t) = H(p_0) + p_\varepsilon t \quad \text{and} \quad P_t + p_t = \alpha + \beta p_t, \tag{11}$$

where $H(p) = 2\alpha\beta \log p + \beta^2 p - \alpha^2/p$, $\alpha = P_0 - (P_\varepsilon/p_\varepsilon)p_0$, and $\beta = 1 + P_\varepsilon/p_\varepsilon$.

This result allow us to characterize the asymptotic dynamics of the entire system:

Corollary 1 (Precisions asymptotic)**.** *(i) The precision of the private forecast monotonically converges to infinity as long as $p_\varepsilon > 0$, (ii) the ratio $p_t/(P_t + p_t)$ monotonically converges to $p_\varepsilon/(p_\varepsilon + P_\varepsilon)$, and (iii) as $t \to \infty$ the total precision, $P_t + p_t$, is such that*

$$P_t + p_t = \begin{cases} \left(3p_0^2 P_\varepsilon t\right)^{1/3} + Q_t, & \text{for } p_\varepsilon = 0 \\ \left(\frac{p_\varepsilon}{P_\varepsilon + p_\varepsilon}\right)^2 (p_\varepsilon + P_\varepsilon)\, t + 2\left(\frac{P_\varepsilon}{p_\varepsilon} p_0 - P_0\right) \log(t) + R_t, & \text{for } p_\varepsilon > 0. \end{cases} \tag{12}$$

where Q_t converges to zero and R_t is bounded.

To explain these asymptotic results, let us start with the case $p_\varepsilon > 0$. Note that part (ii) of the corollary shows that in the limit, as time goes to infinity, the precision of the signals generated by the average action converges to a strictly positive number as long as $p_\varepsilon > 0$. By equations (7) and (8), the sum of the precisions of the endogenously generated private and public signals converges in the limit to:

$$\lim_{t\to\infty} \left(\frac{p_t}{P_t + p_t}\right)^2 (P_\varepsilon + p_\varepsilon) = \left(\frac{p_\varepsilon}{P_\varepsilon + p_\varepsilon}\right)^2 (P_\varepsilon + p_\varepsilon), \tag{13}$$

and this is the coefficient on t in the asymptotic expansion of the precision of total beliefs. As long as $p_\varepsilon > 0$, when time goes to infinity, the social learning process converges to a situation that is as if agents repeatedly observe independent signals centered around x with precision given by (13).

When $p_\varepsilon = 0$, the asymptotic behavior is quite different. In this case, from part (ii) of the Corollary, the precision of the signals generated by the average action converges to zero. As a result, the speed of convergence of the learning process is greatly reduced asymptotically. In this case, the corollary is the continuous-time counterpart of the well known result of Vives (1993): when social learning is constrained to public observations of the average action, the precision of total beliefs goes to infinity at rate $t^{1/3}$. This rate is one order of magnitude slower than the rate at which beliefs would converge if agents were to observe, for example, noisy exogenous signals of x in every period with i.i.d. noise – in that case the precision would go to infinity at a linear rate t. However, the precision of beliefs is going to infinity at a linear rate when $p_\varepsilon > 0$, even though no new information is being exogenously provided to agents: that is, the social learning generated from endogenous private signals, no matter how noisy, is sufficient to restore the speed of convergence to its usual linear rate. When $p_\varepsilon > 0$, the accumulation of knowledge is taking place through both the public and the private social learning channels, and this is the key element that maintains the informativeness of the endogenous signals bounded away from zero, thus avoiding that the rate of convergence be dramatically reduced. The existence of a private social learning channel is important, and as will be shown below, it will also be crucial for the social value of public information.

4 The Impact of Public Information

4.1 Comparative dynamics

We fist study the impact of public information on precision dynamics.

4.1.1 When there is no endogenous private learning channel

We first show that when agents do not learn privately from others' action, then more public information at the beginning increases agents' knowledge at all times. We also argue that this result is robust: it only requires the social learning process to be smooth.

We consider a generalized version of our model where, in addition to observing a public signal of others' actions, we let agents accumulate exogenous information over time. Think for instance, of additional public and private exogenous signals centered around x. However, social learning is restricted to be *public* and "*smooth*": agents continuously observe public signals about the average action in population. This generalization leads us to formulate the following ODE:

$$dP_t = P_\varepsilon \left(\frac{\pi_t}{P_t + \pi_t} \right)^2 dt + d\Pi_t, \tag{14}$$

where Π_t and π_t denote the cumulative precision of exogenous public and private information. Note that we allow for Π_t and π_t to have jumps: that is, the exogenous information could arrive in a lumpy fashion.

Proposition 1. *Suppose exogenous public and private information increase over time according to piecewise continuously differentiable cumulative precisions Π_t and π_t. Consider two initial levels $P'_0 > P_0$ of public information and their associated paths P'_t and P_t. Then $P'_t > P_t$ for all t.*

The proposition follows from i) the standard mathematical result that two different solutions of a (possibly time dependent) ODE can never cross,[6] and ii) from our assumption that lumpy arrivals of public information can only arise from exogenous signals, i.e. $d\Pi_t$ does not depend on the current level of public information.

Namely, the differential equation implies that P'_t and P_t are continuous except when there is a lumpy arrival of public information, in which case they increase by the exact same amount. Thus, given $P'_0 > P_0$, if there were some time such that $P'_t \leq P_t$ then there must be a time $s < t$ such that $P_s = P'_s$. But since both precisions solve the same differential equation the paths of P_t and P'_t would be the same for all times, which is a contradiction.

In the Addendum,[7] we show that the result of Proposition 1 also holds in Vives (1993) original discrete time setup, under either one of the following three conditions: if the initial public precision is large, if the release in public information occurs after the first period, or if the observational noise per period is small. A large observational noise per

[6]Our ODE is slightly non-standard since we allow for positive jumps. However, it is straightforward to extend the standard existence and uniqueness result to this case. See the online Addendum.

[7]The Addendum is available from the authors website

period results in a small amount of endogenous learning per period, which is the natural discrete-time counterpart of our continuous time Brownian setup where information flows smoothly from the endogenous learning channel. In fact, the result of the Proposition holds for any "smooth" endogenous public learning channel: formally, the same proof would go through if the first term on the right-hand side of the ODE is replaced by some bounded locally Lipschitz function $F(t, P_t)$ of time and current public precision.

4.1.2 When there is an endogenous private learning channel

We now show that the effect of public information release changes dramatically when the accumulation of private information is endogenous.

From the asymptotic expansion (12) one sees that a unit increase in the initial public precision, P_0, decreases total precision by approximately $2\log(t)$ at time t. Thus, small differences in time-zero public information result in asymptotically unbounded differences in total information: after an initial release of public information, agents eventually know strictly less. This is in sharp contrast with the result obtained when there was no private social learning.

To understand how this $\log(t)$-impact comes about, let us go back to the ODE (9) for $\dot{p}_t$:

$$\dot{p}_t = p_\varepsilon \left(\frac{p_t}{\alpha + \beta p_t} \right)^2 = p_\varepsilon \left(\frac{\alpha}{p_t} + \beta \right)^{-2},$$

where the term inside the first square is the weight agents put on their private information at time t.

Note that, since $\alpha = P_0 - P_\varepsilon / p_\varepsilon p_0$, it follows that an increase in P_0 reduces the weight that agents put on their private information, and therefore reduces $\dot{p}_t$. But we know from Corollary 1 that this negative effect on the weight disappears in the limit, as the weight on private information converges to $p_\varepsilon / (p_\varepsilon + P_\varepsilon)$, a constant that is independent on initial conditions. However, as we will show below, the negative effect washes out at a slow speed of $1/t$, where t is given by the asymptotic learning speed. Thus, an increase in initial public precision reduces the rate of precision accumulation, $\dot{p}_t$, by an amount in order $1/t$. These vanishing reductions in $\dot{p}_t$ add up, however, to unbounded differences, as $\int dt/t = \log(t)$.

To see this, first note that $\dot{p}_t$ can be approximated by:

$$\dot{p}_t = \frac{p_\varepsilon}{\beta^2} - \frac{2p_\varepsilon}{\beta^3}\frac{\alpha}{p_t} + o\left(\frac{1}{p_t}\right), \tag{15}$$

as $p_t \to \infty$. The first term of the Taylor expansion, $p_\varepsilon(0+\beta)^{-2} = p_\varepsilon/\beta^2$, is the asymptotic rate of private precision accumulation. It reflects the fact that the weight on private information goes to a strictly positive constant, $1/\beta$, and implies that private precision goes to infinity at a linear speed:

$$p_t = \frac{p_\varepsilon}{\beta^2}t + o(t). \tag{16}$$

Next, we turn to the second term of the Taylor expansion. The coefficient $-2p_\varepsilon/\beta^3$ measures the effect of a change in α/p_t on $\dot{p}_t$, evaluated at the limit $\alpha/p_t = 0$. This coefficient is strictly negative because, even in the limit, the equilibrium weight that agents put on private precision remains sensitive to changes in α/p_t.[8] After plugging the approximation (16) in (15), the second term of the Taylor expansion becomes

$$-\frac{2p_\varepsilon}{\beta^3}\frac{\alpha}{p_t} = -\frac{2}{\beta}\frac{\alpha}{t}.$$

This term describes how fast the precision of the endogenous private signal is converging to its long run value of p_ε/β^2. Since $\alpha = P_0 - p_\varepsilon/P_\varepsilon p_0$, it also determines the rate at which the negative impact of increasing P_0 washes out in the long run. Namely, a unit increase in P_0 reduces the time t accumulation of private precision by $(2/\beta)/t$ units, so the negative impact of public information is indeed vanishing with time. But it is doing so at a slow speed: these small $(2/\beta)/t$ differences in the *change* of private precision accumulation add up through time to unbounded differences in the *level* of private precision, which will be in order $\int(2/\beta)(dt/t) = (2/\beta)\log(t)$. The corresponding differences in total precision, $\alpha + \beta p_t$, are thus obtained after multiplying by β, and are in order $2\log(t)$. We

[8]This is indeed the result of the equilibrium behavior. Contrast this with the following: suppose that agents were to follow an exogenous rule that forces them take an action with a weight on private information equal to some given smooth and decreasing function of P_t/p_t. Since this weight governs the accumulation of both public and private information, this dynamic system still has the property that $P_t = \alpha + (\beta - 1)p_t$ so we can write the weight as some function $f(\alpha/p_t) \in [0,1]$. Let us also impose the reasonable assumption that $f(0) > 0$: in words, as the agents accumulate lots of private information, the rule put some positive weight on it. The evolution of private information is then given by $\dot{p}_t = p_\varepsilon\left(f\left(\alpha/p_t\right)\right)^2$, and p_t will grow linearly in the limit. Whether the second term of a Taylor expansion in α/p_t of the ODE is strictly negative will depend on whether $f'(0)$ is strictly negative: that is, whether a change in α/p_t has a first order negative effect on the limiting weight.

summarize this discussion by:

Corollary 2. *Let $p_\varepsilon > 0$. Consider two initial levels $P'_0 > P_0$ of public information and their associated paths of private and public precisions (p'_t, P'_t) and (p_t, P_t). Then, for all $M > 0$ there exists $\bar{t} < \infty$ such that $p'_t + P'_t + M < p_t + P_t$ for all $t > \bar{t}$.*

Intuitively, a release of public information has a self-reinforcing negative feedback on the accumulation of private information. Better initial public information causes agents to put more weight on the public information and less weight on their accumulated private information. This reduces subsequent information accumulation from all channels, in particular from the private social learning channel. But this implies that agents accumulate less private information, so they will put even more weight on their public information, and less on their private information, reducing further private information accumulation, and a negative feedback ensues. The corollary makes explicit that the negative feedback is always strong enough to eventually overtake the initial gains in information.

4.2 The Social Value of Public Information

Whether an initial release of public information is socially beneficial depends on a trade-off between a short-term gain and a long-term loss. The short-term gain is that public information initially improves the precision of agents, which can now make better decisions. The long-term loss, is that public information eventually reduces the amount known by everyone as long as the private social learning channel is active. In the proposition that follows, we provide conditions that ensure that the long-term loss dominates: we show that if the state is revealed in a sufficiently long time, on average, then a marginal increase in public information always reduces utilitarian welfare. Hence, unlike Morris and Shin (2002), we conclude that even in the absence of a payoff externality, more public information can be welfare reducing.

Let the welfare criterion be the equally weighted sum of agents' expected utility. By the law of large numbers, this criterion coincides with the *ex-ante* utility of a representative agent,

$$W \equiv -\lambda \mathbb{E}\left[\int_0^\tau (a_{it} - x)^2 \, dt\right] = -\lambda \mathbb{E}\left[\int_0^\infty e^{-\lambda t} (a_{it} - x)^2 \, dt\right] = -\int_0^\infty \frac{\lambda e^{-\lambda t}}{\alpha + \beta p_t} \, dt,$$

where we have normalized welfare by the intensity λ. The first equality follows from the random end time τ being geometrically distributed: $\lambda e^{-\lambda t}$ is the probability density that the economy ends at time $t > 0$. The second equality follows from $a_{it} = \mathbb{E}[x \mid \mathcal{G}_{it}]$ so $\mathbb{E}\left[(a_{it} - x)^2 \mid \mathcal{G}_{it}\right] = 1/(P_t + p_t) = 1/(\alpha + \beta p_t)$, and an application of Fubini's Theorem.

Public information increases the total precision $\alpha + \beta p_t$ of agents' beliefs in the short run. But, as shown by Corollary 2, in the long run it results in unbounded losses of total precision, $\alpha + \beta p_t$. Because the welfare function is the present value of $-1/(\alpha + \beta p_t)$, it is natural to conjecture that as long as λ is close enough to 0 (i.e. the state is revealed in a long time, on average), public information reduces welfare.

Although intuitive, this result does not directly follow because, even when λ goes to zero, the trade-off between the short-term gain and the long-term loss remains non-trivial. Indeed, since $1/(\alpha + \beta p_t)$ converges to zero, the flow losses of increasing public information are vanishingly small. However, the next theorem shows that it is always possible to find a small enough λ to make a marginal releases of information welfare-reducing:

Theorem 2 (Social Value of Public Information)**.** *For all p_0 and P_0, there exists an $\eta > 0$ such that $0 < \lambda < \eta$ implies $\partial W / \partial P_0 < 0$.*

Theorem 2 tells us that a *marginal* increase in public information reduces welfare, as long as agents are sufficiently patient, which in our model means that the state is revealed in a sufficiently long time ($1/\lambda$ is high enough).

The theorem does not imply, however, that welfare is a monotonically decreasing function of P_0. Indeed, an infinite increase in precision would reveal the state of the world and would clearly improve welfare. By continuity, one might expect that a sufficiently large release of public information also would improve welfare. This intuition is confirmed by the numerical calculation of Figure 1: it shows that welfare is a non-monotonic function of P_0. It first decreases but eventually increases if P_0 is large enough.

In contrast, when $p_\varepsilon = 0$ (only the public social learning channel is active), an increase in public information always increases welfare, given that it increases the precision of agents' information at all times. Thus, the channel of information diffusion is crucial for the social value of public information.

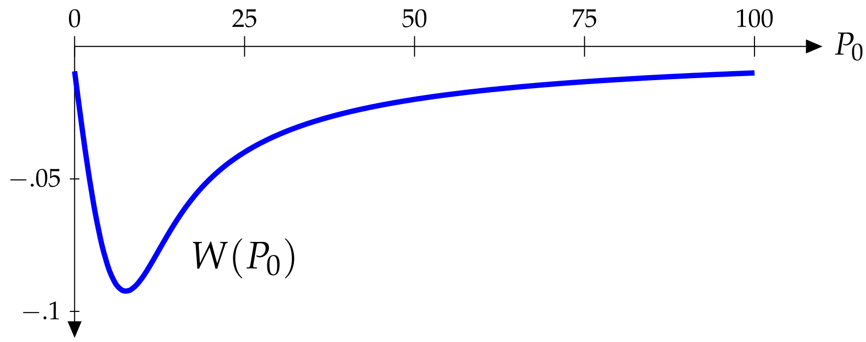

Figure 1: This figures plots welfare as function of the precision, P_0, of initial public information. We choose $p_0 = 0.05$, $p_\varepsilon = 1$, $P_\varepsilon = 0.01$ and $\lambda = 0.001$.

5 Optimal Information Diffusion

This section analyzes the dynamic learning externality by studying the socially optimal diffusion of information. As in Vives (1997), we let a planner choose the weight that agents put on their private information when taking their actions, in order to maximize utilitarian welfare. The planner can't alter the channels through which information diffuses: i.e. he is subject to the constraint that agents only learn from public and private signals of others' average actions. We show that the planner will induce agents to put more weight into their private information than the equilibrium allocation, as a way to correct the information externality. We show that releases of public information are now always beneficial. However, they still reduce the amount known by all agents after some finite time.

A planner chooses an adapted action process a_i in order to maximize the *ex-ante* utility of a randomly chosen agent, subject to the learning technology. In setting up our planning problem, we follow Vives (1997) and restrict attention to actions that are convex combinations of the public and private forecasts defined in Section 3:

$$a_{it} = (1 - \gamma_t)\hat{X}_t + \gamma_t \hat{x}_{it}. \tag{17}$$

That is, instead of using the individually optimal weight $p_t/(P_t + p_t)$ on private information, the planner prescribes agents to use some other deterministic weight $\gamma_t \in [0, 1]$.[9]

[9]In the online Addendum, we allow for actions that are general affine function of the public and the private forecast. However, we show that it is always optimal for the planner to use a convex combination.

We can then solve for the learning dynamics exactly as we did before in the equilibrium analysis. Namely, given a weight γ_t, the precisions P_t and p_t of the public and private forecasts solve:

$$dP_t = P_\varepsilon \gamma_t^2 dt \quad \text{and} \quad dp_t = p_\varepsilon \gamma_t^2 dt. \tag{18}$$

As before, the speed of learning is controlled by the weight, γ_t, that agents put on their private information. Also, these these laws of motion imply that P_t is an affine function of p_t: $P_t = \alpha + (\beta - 1)p_t$ where α and β take the same values as before, $\alpha = P_0 - p_0 P_\varepsilon / p_\varepsilon$ and $\beta = 1 + P_\varepsilon / p_\varepsilon$. This affine transformation simplifies the analysis because it allows to formulate the planner's problem in a one-dimensional state space. One should bear in mind, however, that the transformation imposes a restriction between the parameter α and the initial condition p_0: $p_0 = (\alpha - P_0)/(\beta - 1)$.

The utility flow of a representative agent is

$$\begin{aligned} -\mathbb{E}\left[(a_{it} - x)^2\right] &= -\mathbb{E}\left[\left(\gamma_t(x_{it} - x) + (1 - \gamma_t)(\hat{X}_t - x)\right)^2\right] \\ &= -\left(\frac{\gamma_t^2}{p_t} + \frac{(1-\gamma_t)^2}{\alpha + (\beta - 1)p_t}\right), \end{aligned}$$

where the first line follows from plugging in (17), and the second line from noting that the private and public forecast errors are independent with respective variance $1/p_t$ and $1/(\alpha + (\beta - 1)p_t)$. Let us, then, define the following welfare function for $p \geq p_0$:

$$v(p, \gamma | \alpha) \equiv -\lambda \int_0^\infty e^{-\lambda t} \left(\frac{\gamma_t^2}{p_t} + \frac{(1-\gamma_t)^2}{\alpha + (\beta - 1)p_t}\right) dt,$$

subject to $\dot{p}_t = p_\varepsilon \gamma_t^2$ with initial condition p. Let us denote by $V(p|\alpha)$ the suppremum of $v(p, \gamma|\alpha)$ in the set of admissible controls γ. This function gives the planner's value along any socially optimal precision path starting at p_0.

Our first result states that, once the planner has corrected for the information externality, the welfare effects of public information are always positive:

Proposition 2. *The value function $V(p_0|\alpha)$ is increasing in $\alpha = P_0 - P_\varepsilon / p_\varepsilon p_0$.*

This is intuitive: starting with a higher α, the planner can always choose the same sequence of weights on private information that he would have followed with a lower one. These same weights would then imply the same time path of private precision, p_t. But a higher α increases total precision, $\alpha + \beta p_t$, and thus welfare. The planner's ability to control the weight guarantees that the social value of public information is always positive.

We now proceed to analyze the impact of a public information release on precision dynamics, which requires a tighter characterization of the planner's optimal weight.

5.1 The Optimal Weight and Comparative Dynamics

In Appendix C, we show that the value function $V(p|\alpha)$ is differentiable with respect to p and solves the Hamilton Jacobi Bellman (HJB) equation:

$$\lambda V(p|\alpha) = \max_{\gamma \in [0,1]} \left\{ -\lambda \left(\frac{\gamma^2}{p} + \frac{(1-\gamma)^2}{\alpha + (\beta - 1)p} \right) + p_\varepsilon \gamma^2 V'(p|\alpha) \right\}. \tag{19}$$

Taking derivative with respect to the weight, γ, yields:

$$-2\lambda \left(\frac{\gamma}{p} + \frac{\gamma - 1}{\alpha + (\beta - 1)p} \right) + 2p_\varepsilon \gamma V'(p|\alpha). \tag{20}$$

The second term of (20), $2p_\varepsilon \gamma V'(p|\alpha)$, corresponds to the marginal welfare gain of speeding up information dissemination.[10] The first term, on the other hand, corresponds to the marginal change in agents' forecast error. This term is equal to zero if the planner chooses the individually optimal weight $p/(\alpha + \beta p)$ of the decentralized equilibrium we studied before, since it minimizes the forecast error. Clearly, the planner has incentive to increase its weight above the individual optimum $p/(\alpha + \beta p)$: at the margin, it has no impact on agents' forecast error, while the marginal welfare gain of speeding up information dissemination is strictly positive, given that $V'(p|\alpha) > 0$. Thus, the socially optimal weight satisfies:

$$\gamma^\star(p) > \frac{p}{\alpha + \beta p}. \tag{21}$$

Note that this implies that precision increases faster in the social optimum than in the decentralized equilibrium, and, in turns, that it also converges to infinity: therefore, full revelation is socially optimal.

Finally, after plugging $\gamma = 1$ in the right-hand side of (19), it follows that $0 > \lambda V(p|\alpha) \geq -\lambda/p + p_\varepsilon V'(p|\alpha)$, which implies that (20) is strictly negative when evaluated at $\gamma = 1$,

[10] Note that the value function is increasing: indeed if the planner applies the same optimal control but with a higher initial condition, his flow utility is higher at each time. This implies that $V'(p|\alpha) \geq 0$. In Appendix C we use the envelope condition to show that the derivative is, in fact, strictly positive.

and thus $\gamma \neq 1$.

Although we are not able to solve the HJB in closed-form, we are able to characterize the asymptotic behavior of the optimal control:

Theorem 3. *Suppose that $p_\varepsilon > 0$. In the planner's solution, as $p \to \infty$, the socially optimal weight admits the following asymptotic expansion:*

$$\gamma^\star(p) = \underbrace{\frac{1}{\beta} - \frac{\alpha}{\beta^2}\frac{1}{p}}_{\text{individual optimum}} + \underbrace{\frac{1}{\lambda}\frac{(\beta-1)p_\varepsilon}{\beta^2}\frac{1}{p}}_{\text{planner's correction}} + O\left(\frac{1}{p^2}\right) \tag{22}$$

The first two terms of the asymptotic expansion are the same as in the equilibrium behavior. The correction that the planner makes shows up in the third term. As can be seen, this third term is positive but vanishes as the precision converges to infinity, implying that the planner's weight converges to the equilibrium weight, $1/\beta$. As agents are close to learning the truth, the welfare gain from speeding up information diffusion becomes small, and the planner chooses to reduce the forecast error instead.

The asymptotic expansion of $\gamma^\star(p)$ also reveals another feature of the solution: impatience reduces the planner's asymptotic correction to the weight. This is intuitive, as a more impatient planner discounts more heavily the future benefits of correcting the information externality.

Importantly, the asymptotic expansion in (22) allows us to provide a characterization of the total beliefs, and the effects of public information on them. Surprisingly, our main equilibrium comparative statics still holds: an increase in initial public information eventually leads to less total knowledge,

Corollary 3. *Suppose that $p_\varepsilon > 0$. In the planner's solution, the total precision of agents beliefs at time t equals:*

$$P_t + p_t = (P_\varepsilon + p_\varepsilon)\left(\frac{p_\varepsilon}{P_\varepsilon + p_\varepsilon}\right)^2 t + 2\left(\frac{P_\varepsilon}{p_\varepsilon}p_0 - P_0 + \underbrace{\frac{1}{\lambda}\left(\frac{P_\varepsilon p_\varepsilon}{p_\varepsilon + P_\varepsilon}\right)}_{\text{planner's correction}}\right)\log(t) + S_t$$

where S_t is bounded.

The corollary follows from the asymptotic expansion (22) of $\gamma^\star(p)$, using a similar argument as in the equilibrium case. In particular, the expansion of $\gamma^\star(p)$ shows that the effects of both the initial public precision (as represented by α) and the planner's correction term are of order $1/p$, which implies that they will have a $\log(t)$-impact on the total precision at time t. The crucial observation is that the planner correction term is independent of α, and hence, differences in initial public information will continue to have the same negative and unbounded $\log(t)$-impact on total precision.

An immediate corollary is that any initial increase in public information eventually decreases total knowledge. So, even though public information always increases *ex-ante* welfare, we find that *ex-post* welfare will eventually decrease. This results arises even though the planner is appropriately correcting the externality by making agents put more weight into their private information than they would do in equilibrium.

Lastly, we observe that Corollary 3 is is a result of optimality, not feasibility: although, after a release of public information, the planner could choose a weight that increases knowledge at all times (as noted after Proposition 2), he chooses not to do so, and knowledge is reduced in the long run.

6 Conclusion

This papers analyzes how private information diffuses among a continuum of agents who learn from both public and private observations of each others' actions. We show that when agents learn from a private channel, a release of public information at the beginning always slows the diffusion of private information in the economy, eventually reduces the amount known by everyone, and sometimes reduces welfare. We studied the optimal diffusion of information and showed that, relative to the private optimum, the planner corrects the learning externality by recommending agents to put more weight on their private information. We show that the social value of public information after the planner corrects the externality is positive. However, a release of public information will eventually lead to agents knowing strictly less, as in the equilibrium.

A Proof of Theorem 1

We prove the theorem using a guess and verify approach. Namely, we start by guessing that agents' signals, in equation (2) are observationally equivalent to the transformed signals of equation (6). We then derive the dynamics of agents' beliefs and of the average action. We then verify our guess, given this stochastic process for the average action.

A.1 Beliefs dynamics

Assume, then, that agents observe the pair of public and private signals given by the Stochastic Differential Equations (SDE) (6), with initial conditions (1). Let the public forecast, $\hat{X}_t$ be the expectation of x conditional on the public signal on the left-hand side of (6), given the common prior. Similarly, let public forecast, $\hat{x}_{it}$ be the expectation of x conditional on the private signal on the left-hand side of (6), given a fully diffuse prior. Denote by P_t and p_t the associated public and private precision. Then, we have:

Lemma 1 (Dynamics of Private and Public Forecasts)**.** *The private and public forecasts* $(\hat{x}_{it}, \hat{X}_t)$ *and the precisions* (p_t, P_t) *solve the system of SDE*

$$d\hat{x}_{it} = \frac{p_\varepsilon}{p_t}\frac{p_t^2}{(P_t+p_t)^2}\left[\left(x-\hat{x}_{it}\right)\,dt + \frac{d\omega_{it}}{\sqrt{p_\varepsilon}\frac{p_t}{P_t+pi_t}}\right] \tag{23}$$

$$d\hat{X}_t = \frac{P_\varepsilon}{P_t}\frac{p_t^2}{(P_t+p_t)^2}\left[\left(x-\hat{X}_t\right)\,dt + \frac{dW_t}{\sqrt{P_\varepsilon}\frac{p_t}{P_t+p_t}}\right] \tag{24}$$

$$dp_t = p_\varepsilon\left(\frac{p_t}{P_t+p_t}\right)^2 dt \tag{25}$$

$$dP_t = P_\varepsilon\left(\frac{p_t}{P_t+p_t}\right)^2 dt\;, \tag{26}$$

with the initial conditions p_0*, and* P_0*. In addition, the above system can be integrated into*

$$\hat{x}_{it} = x + \frac{1}{p_t}\left[\sqrt{p_0}\omega_{i0} + \int_0^t \sqrt{p_\varepsilon}\frac{p_t}{P_t+p_t}d\omega_{iu}\right] \tag{27}$$

$$\hat{X}_t = x + \frac{1}{P_t}\left[\sqrt{P_0}W_0 + \int_0^t \sqrt{P_\varepsilon}\frac{p_t}{P_t+p_t}dW_u\right] \tag{28}$$

Equations (24)-(25) follows from a direct application of one-dimensional continuous-time Kalman filtering formula (see, for instance, Oksendal (1985), pages 85-105). In order to derive equation (28), we multiply both sides of (24) by P_t. We find

$$\begin{aligned}
& P_t d\hat{X}_t = P_\varepsilon \frac{p_t^2}{(P_t+p_t)^2}\left[\left(x-\hat{X}_t\right)\,dt + \frac{dW_t}{\sqrt{P_\varepsilon}\frac{p_t}{P_t+p_t}}\right] \\
\Rightarrow\quad & P_t d\hat{X}_t + dP_t\left(\hat{X}_t - x\right) = \sqrt{P_\varepsilon}\frac{p_t}{P_t+p_t}dW_t \\
\Rightarrow\quad & d\left[P_t\left(\hat{X}_t - x\right)\right] = \sqrt{P_\varepsilon}\frac{p_t}{P_t+p_t}dW_t \\
\Rightarrow\quad & P_t\left(\hat{X}_t - x\right) - P_0\left(\hat{X}_0 - x\right) = \int_0^t \sqrt{P_\varepsilon}\frac{p_u}{P_u+p_u}dW_u \\
\Rightarrow\quad & \hat{X}_t = \frac{P_0}{P_t}\hat{X}_0 + \left(1-\frac{P_0}{P_t}\right)x + \frac{1}{P_t}\int_0^t \sqrt{P_\varepsilon}\frac{p_u}{P_u+p_u}dW_u \qquad (29)
\end{aligned}$$

where the second line follows from the fact that $dP_t = P_\varepsilon p_t^2/(P_t+p_t)^2$. Because P_t is a deterministic function of time it follows that $d[(\hat{X}_t - x)P_t] = d\hat{X}_t P_t + (\hat{X}_t - x)dP_t$, which implies the third line. The fourth line follows from integrating the third line from $u = 0$ to $u = t$, and the fifth line follows from rearranging. Now note that $\hat{X}_0$ and P_0 are the posterior mean and precision at time zero, after observing the public signal $Z_0 = x + W_0/\sqrt{P_0}$ and starting from the fully diffused common prior. Therefore $\hat{X}_0 = Z_0$. Equation (28) then follows from plugging $\hat{X}_0 = Z_0 = x + W_0/\sqrt{P_0}$ back into (29). Equation (27) follows from exactly the same algebraic manipulations.

Under our guess, the signals that generate the public and the private forecast, $\hat{X}_t$ and $\hat{x}_{it}$, are independent conditional on x. Hence, an agent's forecast conditional on all his information will be a linear combination of the public and his private forecasts, with weights given by their respective precisions:

$$\frac{P_t}{P_t+p_t}\hat{X}_t + \frac{p_t}{P_t+p_t}\hat{x}_{it}. \qquad (30)$$

The precision of the agents belief conditional on all his information is $P_t + p_t$. QED

A.2 Verifying the guess

Let $\tilde{Z}_t$ and $\tilde{z}_{it}$ be the "transformed" signals of equation (6), with initial condition (1). We need to verify that the filtration generated by the transformed signals $(\tilde{Z}_t, \tilde{z}_{it})$ is the same as that generated by the signals of others' action, (Z_t, z_{it}). First, after plugging $A_t = P_t/(P_t+p_t)\hat{X}_t + p_t/(P_t+p_t)x$ into equation (2) we find that

$$\begin{aligned}
dZ_t &= \frac{P_t}{P_t+p_t}\hat{X}_t\,dt + \frac{p_t}{p_t+P_t}d\tilde{Z}_t \\
dz_{it} &= \frac{P_t}{P_t+p_t}\hat{X}_t\,dt + \frac{p_t}{p_t+P_t}d\tilde{z}_{it}
\end{aligned}$$

Keeping in mind that $\hat{X}_t$, by construction, is adapted to the filtration generated by $(\tilde{Z}_t, \tilde{z}_{it})$, the above equation shows that the filtration generated by (Z_t, z_{it}) is included in the filtration generated by $(\tilde{Z}_t, \tilde{z}_{it})$. To show the reverse inclusion, first rearrange the above equation into

$$d\tilde{Z}_t = \frac{P_t + p_t}{p_t}\left(dZ_t - \frac{P_t}{P_t + p_t}\hat{X}_t\,dt\right) \tag{31}$$

$$d\tilde{z}_{it} = \frac{P_t + p_t}{p_t}\left(dz_{it} - \frac{P_t}{P_t + p_t}\hat{X}_t\,dt\right) \tag{32}$$

Now, we also know from Lemma 1 that

$$d\hat{X}_t = \frac{P_\varepsilon}{P_t}\left(\frac{p_t}{P_t + p_t}\right)^2\left[\left(x - \hat{X}_t\right)\,dt + \frac{dW_t}{\sqrt{P_\varepsilon \frac{p_t}{P_t+p_t}}}\right] = \frac{P_\varepsilon}{P_t}\left(\frac{p_t}{P_t + p_t}\right)^2\left(-\hat{X}_t\,dt + d\tilde{Z}_t\right)$$

After plugging equations (31) in the equation above and rearranging, we find:

$$d\hat{X}_t = \frac{P_\varepsilon}{P_t}\left(\frac{p_t}{P_t + p_t}\right)^2\left\{-\left(1 + \frac{P_t}{p_t}\right)\hat{X}_t\,dt + \frac{P_t + p_t}{p_t}\,dZ_t\right\}$$

Therefore, $\hat{X}_t$ is adapted to the filtration generated by Z. Together with (31), this means that $\tilde{Z}_t$ is adapted to the filtration generated by Z_t. Together with (31) and (32) this implies that $(\tilde{Z}_t, \tilde{z}_{it})$ is adapted to the filtration generated by (Z_t, z_{it}). QED

A.3 Closed form solution

This can be verified directly. QED

B Proof of Corollary 1

Part (i): When $p_\varepsilon = 0$, this follows directly from the solution (10). When $p_\varepsilon > 0$, then from ODE (9) it is clear that p_t is strictly increasing, so it has a limit as $t \to \infty$. The limit must be infinite otherwise, as $t \to \infty$, the left-hand-side of equation (11) would remain bounded, while the right hand side would go to infinity.
Part (ii): When $p_\varepsilon = 0$ the ratio is $p_0/(p_0 + P_t)$ and is clearly decreasing towards zero. When $p_\varepsilon > 0$, we have

$$\frac{p_t}{P_t + p_t} = \frac{p_t}{\alpha + \beta p_t} = \frac{1}{\frac{\alpha}{p_t} + \beta}. \tag{33}$$

This implies that the ratio is converging towards $1/\beta = p_\varepsilon/(p_\varepsilon + P_\varepsilon)$. Since p_t is increasing, the ratio is strictly increasing if $\alpha > 0$, strictly decreasing if $\alpha < 0$, and constant if $\alpha = 0$.
Part (iii): When $p_\varepsilon = 0$, the result follows directly from the closed form solution (10). When $p_\varepsilon > 0$, we note that:

$$\begin{aligned} \dot{p}_t &= \left(\frac{p_t}{\alpha + \beta p_t}\right)^2 p_\varepsilon = \left(\frac{\alpha}{p_t} + \beta\right)^{-2} p_\varepsilon \\ &= \frac{{p_\varepsilon}^2}{\beta} - \frac{2\alpha p_\varepsilon}{\beta^3}\frac{1}{p_t} + O\left(\frac{1}{p_t^2}\right), \end{aligned}$$

where $O(1/p^2)$ is the standard Landau notation for a function $f(p)$ such that $p^2 f(p)$ is bounded. An application of Lemma 2 below show that:

$$p_t = \frac{p_\varepsilon{}^2}{\beta} t - \frac{2\alpha}{\beta} \log(t) + C_t,$$

for some bounded function C_t. The result follows after noting that $P_t + p_t = \alpha + \beta p_t$. QED

Lemma 2 (ODE Asymptotics). *Suppose*

$$\dot{x}_t = A + \frac{B}{x_t} + O\left(\frac{1}{x_t^2}\right) \tag{34}$$

and suppose that $x_t \to \infty$. *Then*

$$x_t = At + \frac{B}{A} \log(t) + C_t,$$

for some bounded function C_t.

Let T be such that $A - \epsilon \leq \dot{x}_t \leq A + \epsilon$ for all $t > T$, where $0 < \epsilon < A$. Then $(A - \epsilon)(t - T) \leq x_t - x_T \leq (A + \epsilon)(t - T)$, and thus $1/x_t = O(1/t)$. Plugging this back into the ODE (34) we have that $\dot{x}_t = A + O(1/t)$, which after integrating delivers that $x_t = At + O(\log(t))$. Taking the inverse of this last equation gives:

$$\frac{1}{x_t} = \frac{1}{At}\left[1 + O\left(\frac{\log(t)}{t}\right)\right]^{-1} = \frac{1}{At}\left[1 + O\left(\frac{\log(t)}{t}\right)\right] = \frac{1}{At} + O\left(\frac{\log(t)}{t^2}\right)$$

Plugging back into the ODE (34):

$$\dot{x}_t = A + \frac{B}{At} + O\left(\frac{\log(t)}{t^2}\right)$$

where it was used that $x_t = O(t)$ implies that $1/x_t^2 = O(1/t^2)$. Now, the function $\log(t)/t^2$ is absolutely integrable, so the final result follows by integrating both sides of this equation. QED

B.1 Proof of Theorem 2

Using equation (9) the welfare can be rewritten as,

$$W = -\int_0^\infty \left(\frac{\alpha + \beta p_t}{p_t^2}\right) \frac{1}{p_\varepsilon} e^{-\lambda t} \left(\dot{p}_t \, dt\right) = -\frac{1}{p_\varepsilon} \int_{p_0}^\infty \left(\frac{\alpha + \beta p}{p^2}\right) e^{-\lambda\left(\frac{H(p) - H(p_0)}{p_\varepsilon}\right)} dp. \tag{35}$$

where the last equality follows from $p_\varepsilon t = H(p_t) - H(p_0)$ for $H(p) \equiv 2\alpha\beta \log p + \beta^2 p - \alpha^2/p$ and, given that p_t monotonically approaches infinity through time, a change in the integrating variable from t to p. Thus, the welfare function W written this way depends on the initial precision P_0 of public information only through $\alpha = P_0 - (P_\varepsilon/p_\varepsilon)p_0$. A marginal increase in the precision of the initial public signal decreases social welfare if and only if $\partial W/\partial\alpha < 0$.

We thus take the derivative of W with respect to α. We find

$$\begin{aligned}\frac{\partial W}{\partial \alpha} &= -\frac{1}{p_\varepsilon}\int_{p_0}^{\infty}\left\{\frac{1}{p^2}-\frac{\alpha+\beta p}{p^2}\frac{\lambda}{p_\varepsilon}\frac{\partial}{\partial \alpha}\left(H(p)-H(p_0)\right)\right\}e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp \\ &= -\frac{1}{p_\varepsilon}\int_{p_0}^{\infty}\left\{\frac{1}{p^2}-\frac{\alpha+\beta p}{p^2}\frac{\lambda}{p_\varepsilon}\left[-2\alpha\left(\frac{1}{p}-\frac{1}{p_0}\right)+2\beta\log\left(\frac{p}{p_0}\right)\right]\right\}e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp. \qquad (36)\end{aligned}$$

We now integrate the first term $\int_{p_0}^{\infty}(-1/p^2)e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp$ of integral (36) by part, noting that $-1/p^2 = d/dp\,(1/p-1/p_0)$. This gives

$$\begin{aligned} & -\int_{p_0}^{\infty}\frac{1}{p^2}e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp \\ = & \left[\left(\frac{1}{p}-\frac{1}{p_0}\right)e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\right]_0^{\infty}+\int_{p_0}^{\infty}\left(\frac{1}{p}-\frac{1}{p_0}\right)H'(p)e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp \\ = & \int_{p_0}^{\infty}\left(\frac{1}{p}-\frac{1}{p_0}\right)H'(p)e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp, \qquad (37)\end{aligned}$$

because $H(p)\to\infty$ as $p\to\infty$. We manipulate the second term of the integral as follows:

$$\begin{aligned} & \int_{p_0}^{\infty}\frac{\alpha+\beta p}{p^2}\frac{\lambda}{p_\varepsilon}\left[-2\alpha\left(\frac{1}{p}-\frac{1}{p_0}\right)+2\beta\log\left(\frac{p}{p_0}\right)\right]e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp \\ = & \int_{p_0}^{\infty}\frac{H'(p)}{H'(p)}\frac{\alpha+\beta p}{p^2}\frac{\lambda}{p_\varepsilon}\left[\frac{2\alpha(p-p_0)}{pp_0}+2\beta\log\left(\frac{p}{p_0}\right)\right]e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp \\ = & \int_{p_0}^{\infty}\left[\frac{p}{\alpha+\beta p}\right]^2\frac{\alpha+\beta p}{p^2}\left[\frac{2\alpha(p-p_0)}{pp_0}+2\beta\log\left(\frac{p}{p_0}\right)\right]\frac{\lambda}{p_\varepsilon}H'(p)e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp \\ = & \int_{p_0}^{\infty}\frac{1}{\alpha+\beta p}\left[\frac{2\alpha(p-p_0)}{pp_0}+2\beta\log\left(\frac{p}{p_0}\right)\right]\frac{\lambda}{p_\varepsilon}H'(p)e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp, \qquad (38)\end{aligned}$$

where the third line follows from the fact that $H'(p)=[(\alpha+\beta p)/p]^2$. Plugging (37) and (38) into the above equation (36) gives:

$$\frac{\partial W}{\partial \alpha} = \frac{1}{p_\varepsilon}\int_{p_0}^{\infty}\Phi(p,p_0)\frac{\lambda}{p_\varepsilon}H'(p)e^{-\lambda/p_\varepsilon(H(p)-H(p_0))}\,dp, \qquad (39)$$

where

$$\Phi(p,p_0) = \left[\frac{1}{p}-\frac{1}{p_0}\right]+\frac{1}{\alpha+\beta p}\left[\frac{2\alpha(p-p_0)}{pp_0}+2\beta\log\left(\frac{p}{p_0}\right)\right].$$

Now since $\Phi(p,p_0)\to -1/p_0$ as $p\to\infty$, there exists some p^* such that $\Phi(p,p_0) < -1/(2p_0)$ for all $p>p^*$.

Letting $M^* = \sup_{p \in [p_0, p^*]} \Phi(p, p_0)$, equation (39) implies that

$$\begin{aligned}
\frac{\partial W}{\partial \alpha} &= \frac{1}{p_\varepsilon} \int_{p_0}^{p^*} \Phi(p, p_0) \frac{\lambda}{p_\varepsilon} H'(p) e^{-\lambda/p_\varepsilon (H(p) - H(p_0))} \, dp \\
&\qquad + \frac{\lambda}{p_\varepsilon} \int_{p^*}^{\infty} \Phi(p, p_0) \frac{\lambda}{p_\varepsilon} H'(p) e^{-\lambda/p_\varepsilon (H(p) - H(p_0))} \, dp \\
&\leq \frac{1}{p_\varepsilon} \left\{ M^* \int_{p_0}^{p^*} \frac{\lambda}{p_\varepsilon} H'(p) e^{-\lambda/p_\varepsilon (H(p) - H(p_0))} \, dp - \frac{1}{2p_0} \int_{p^*}^{\infty} \frac{\lambda}{p_\varepsilon} H'(p) e^{-\lambda/p_\varepsilon (H(p) - H(p_0))} \, dp \right\} \\
&= \frac{1}{p_\varepsilon} \left\{ M^* \left(1 - e^{-\lambda/p_\varepsilon (H(p) - H(p_0))} \right) - \frac{1}{2p_0} e^{-\lambda/p_\varepsilon (H(p) - H(p_0))} \right\}
\end{aligned}$$

The term inside the curly brackets is negative as long as λ is small enough. QED

C Proofs omitted in Section 5

In all what follows, we keep the parameters (α, β) the same, and we omit the dependence of functions on the parameters (α, β). In setting up and solving the planner's problem, we also impose that the control γ_t lies in a compact set, which we take to be the interval $[0, 1]$. However, as will become clear, this constraint is never binding.

C.1 Lipschitz continuity

We start by establishing that the value function is Lipschitz continuous:

Lemma 3. *For all $p' > p$:*

$$V(p') - V(p) \leq (p' - p) \left(\frac{1}{p^2} + \frac{\beta - 1}{(\alpha + (\beta - 1)p)^2} \right)$$

Indeed, consider two initial conditions $p' > p$. By definition of the value function, for any $\varepsilon > 0$ there is some control γ' such that $v(p', \gamma') \geq V(p') - \varepsilon$, and $V(p) \geq v(p, \gamma')$. Combining these two inequalities, we obtain:

$$\begin{aligned}
V(p') - V(p) &\leq v(p', \gamma') - v(p, \gamma') + \varepsilon \\
&= \int_0^\infty \lambda e^{-\lambda_t} \left[\gamma_t'^2 \left(\frac{1}{p_t} - \frac{1}{p_t'} \right) + \right. \\
&\qquad \left. + (1 - \gamma_t')^2 \left(\frac{1}{\alpha + (\beta - 1)p_t} - \frac{1}{\alpha + (\beta - 1)p_t'} \right) \right] dt + \varepsilon \\
&= \int_0^\infty \lambda e^{-\lambda_t} \left[\frac{\gamma_t'^2 (p_t' - p_t)}{p_t' p_t} + \frac{(1 - \gamma_t')^2 (\beta - 1)(p_t' - p_t)}{(\alpha + (\beta - 1)p_t)(\alpha + (\beta - 1)p_t')} \right] dt + \varepsilon \\
&\leq \int_0^\infty \lambda e^{-\lambda_t} \left[\frac{(p' - p)}{p' p} + \frac{(\beta - 1)(p' - p)}{(\alpha + (\beta - 1)p)(\alpha + (\beta - 1)p')} \right] dt \\
&\leq (p' - p) \left(\frac{1}{p^2} + \frac{\beta - 1}{(\alpha + (\beta - 1)p)^2} \right) + \varepsilon.
\end{aligned}$$

where the second to last line follows since $\gamma'_t \in [0,1]$, and because $\dot{p}'_t = \dot{p}_t = p_\varepsilon \gamma'_t$ implies that $p'_t - p_t = p' - p$ and $p'_t > p_t \geq p_0$ for all times. The result obtains by letting $\varepsilon \to 0$. QED

C.2 Dynamic programming

Let us define

$$V'(p^+) \equiv \limsup_{p' \to p^+} \frac{V(p') - V(p)}{p' - p}.$$

Since Lipschitz continuity implies absolute continuity, we know from Rudin (1974) Theorem 7.20 that the value function is differentiable almost everywhere: therefore, $V'(p^+)$ coincides is the classical derivative almost everywhere and

$$V(p') = V(p) + \int_p^{p'} V'(x^+)\, dx, \tag{40}$$

for all $p' > p$. Equipped with this definition of the derivative, the results of Chapter III in Bardi and Capuzzo-Dolcetta (1997) allow us to state (see the Addendum for a step-by-step explanation):

Lemma 4. *The value function solves the HJB equation (19) shown in the text for all p. Let $\gamma^*(p)$ achieve the maximum in (19)*

$$\gamma^*(p) \equiv \min\left\{1, \frac{p}{\alpha + \beta p - (\alpha + (\beta - 1)p) p p_\varepsilon V'(p^+)/\lambda}\right\}. \tag{41}$$

Then $\gamma_t^* = \gamma^*(p_t)$ where $p_t = p + \int_0^t p_\varepsilon (\gamma_t^*)^2 dt$ is an optimal control for the planner's problem with initial condition $p > p_0$.

As explained in the text, one easily shows that $\gamma^*(p) < 1$. Together with the finding that $\gamma^*(p) \geq p/(\alpha + \beta p)$, this implies that the constraint $\gamma^*(p) \in [0,1]$ is never binding. Next, plugging $\gamma^*(p)$ back into the HJB, one finds that

$$V(p) = -\frac{1 - \gamma^*(p)}{\alpha + (\beta - 1)p}. \tag{42}$$

But $V(p)$ is continuous: therefore, $\gamma^*(p)$ and, by implication, $V'(p^+)$, are also continuous. Then, it follows from (40) that $V(p)$ is, in fact, continuously differentiable.

Plugging the expression for $\gamma^*(p)$ as a function of $V'(p)$ into (42), one finds that $V(p)$ solves the ordinary differential equation (ODE):

$$V'(p) = \frac{\lambda}{p_\varepsilon} \frac{V(p)(\alpha + \beta p) + 1}{p\left[V(p)(\alpha + (\beta - 1)p) + 1\right]}. \tag{43}$$

Clearly, since the right-hand-side is continuous, so is the left hand side, and so on, implying that $V(p)$ admits continuous derivatives at all orders. This allows us to use the envelope condition in (19), and obtain:

$$V'(p) = \lambda \frac{\gamma^*(p)^2}{p^2} + \lambda \frac{(\beta - 1)(1 - \gamma^*(p))^2}{(\alpha + (\beta - 1)p)^2} + p_\varepsilon \gamma^*(p)^2 V''(p).$$

Integrating this up along the socially optimal path of precision starting at p, and using that $\lim_{t \to \infty} e^{-\lambda t} V'(p_t) = 0$ (which follows from the Lipschitz bound), gives:

Lemma 5. *The value function is continuously differentiable and its derivative is:*

$$V'(p) = \int_0^\infty \left(\frac{(\gamma_t^*)^2}{(p_t^*)^2} + \frac{(\beta-1)(1-\gamma_t^*)^2}{(\alpha+(\beta-1)p_t^*)^2} \right) \lambda e^{-\lambda t}\, dt, \tag{44}$$

where p_t^* and γ_t^* are, respectively, the socially optimal path of precision and the socially optimal weight starting at p. In particular, $V'(p) > 0$ for all p.

Plugging that $\gamma_t^* \geq p_t^*/(\alpha + \beta p_t^*) > 0$ in the right-hand side of (44) implies that the derivative is strictly positive.

C.3 Preliminary asymptotic results

Our first preliminary result is

Lemma 6. *As $p \to \infty$, $\gamma^*(p) = 1/\beta + O(1/p)$.*

To prove this result first note that, since $\gamma(p) \in [0,1]$, it follows from (44) that:

$$V'(p) \leq \frac{1}{p^2} + \frac{\beta-1}{(\alpha+(\beta-1)p)^2}.$$

Now plugging this into the (41), this gives

$$\begin{aligned}
&\frac{p}{\alpha+\beta p} \leq \gamma^*(p) \leq \frac{p}{\alpha+\beta p - \frac{p_\varepsilon}{\lambda}\left(\frac{\alpha+(\beta-1)p}{p} + \frac{(\beta-1)p}{\alpha+(\beta-1)p}\right)} \\
\Leftrightarrow &\left[\frac{\alpha}{p}+\beta\right]^{-1} \leq \gamma^*(p) \leq \left[\frac{\alpha}{p}+\beta - \frac{p_\varepsilon}{\lambda p}\left(\frac{\alpha+(\beta-1)p}{p} + \frac{(\beta-1)p}{\alpha+(\beta-1)p}\right)\right]^{-1} \\
\Leftrightarrow &\left[\beta + O\left(\frac{1}{p}\right)\right]^{-1} \leq \gamma^*(p) \leq \left[\beta + O\left(\frac{1}{p}\right)\right]^{-1},
\end{aligned}$$

and the result follows. QED

Next, we prove:

Lemma 7. *As $p \to \infty$, $p^2 V'(p) = 1/\beta + O(1/p)$.*

To see this, multiply both sides of (44) by p^2 and obtain

$$p^2 V'(p) = \int_0^\infty \frac{p^2 \gamma(p_t^*)^2}{(p_t^*)^2} \lambda e^{-\lambda t}\, dt + \int_0^\infty \frac{p^2(\beta-1)(1-\gamma(p_t^*))^2}{(\alpha+(\beta-1)p_t^*)^2} \lambda e^{-\lambda t}\, dt. \tag{45}$$

We start by showing that the first integral is $1/\beta^2 + O(1/p)$. Indeed, since we know from Lemma 6 that $\gamma(p) = 1/\beta + O(1/p)$, it follows that

$$\gamma(p_t^*)^2 = \frac{1}{\beta^2} + O\left(\frac{1}{p_t^*}\right) = \frac{1}{\beta^2} + O\left(\frac{1}{p}\right)$$

where the second equality follows because $p \leq p_t$ for all t. Substituting this in equation (45) and subtracting

$1/\beta^2$, we find:

$$\begin{aligned}
\left|\int_0^\infty \frac{p\gamma(p_t^*)^2}{(p_t^*)^2}\lambda e^{-\lambda t}\,dt - \frac{1}{\beta^2}\right| &= \left|\int_0^\infty \left(\frac{p^2\left(\frac{1}{\beta^2}+O\left(\frac{1}{p}\right)\right)}{(p_t^*)^2} - \frac{1}{\beta^2}\right)\lambda e^{-\lambda t}\,dt\right| \\
&= \left|\int_0^\infty \frac{p^2 + O(p) - (p+(p_t^*-p))^2}{\beta^2\,(p+(p_t^*-p))^2}\lambda e^{-\lambda t}\,dt\right| \\
&= \left|\int_0^\infty \frac{O(p) - 2p(p_t^*-p) + (p_t^*-p)^2}{\beta^2\,(p+(p_t^*-p))^2}\lambda e^{-\lambda t}\,dt\right| \\
&\leq \int_0^\infty \frac{O(p) + 2p(p_t^*-p) + (p_t^*-p)^2}{\beta^2 p^2}\lambda e^{-\lambda t}\,dt \\
&\leq \int_0^\infty \left(O\left(\frac{1}{p}\right) + \frac{2p_\varepsilon t}{\beta^2 p} + \frac{(p_\varepsilon t)^2}{\beta^2 p^2}\right)\lambda e^{-\lambda t}\,dt = O\left(\frac{1}{p}\right)
\end{aligned}$$

where the last inequality follows because $p_t^* - p = \int_0^t p_\varepsilon(\gamma_s^*)^2\,ds$ and $\gamma_s^* \in [0,1]$. Following the same steps one shows that the second integral in (45) is $(\beta-1)/\beta^2 + O(1/p)$, and the result follows. QED

C.4 Proof of Theorem 3

Combining equation (41) and Lemma 7, we obtain:

$$\begin{aligned}
\gamma(p) &= \frac{p}{(\alpha+\beta p-(\alpha+(\beta-1)p)p\frac{p_\varepsilon}{\lambda}V'(p^+|\alpha))} = \left[\frac{\alpha}{p}+\beta-\frac{1}{p}\left((\beta-1)+\frac{\alpha}{p}\right)\frac{p_\varepsilon}{\lambda}p^2V'(p|\alpha)\right]^{-1} \\
&= \left[\frac{\alpha}{p}+\beta-\frac{1}{p}\left((\beta-1)+\frac{\alpha}{p}\right)\frac{p_\varepsilon}{\lambda}\left(\frac{1}{\beta}+O\left(\frac{1}{p}\right)\right)\right]^{-1} = \left[\beta+\frac{1}{p}\left(\alpha-\frac{p_\varepsilon}{\lambda}\frac{\beta-1}{\beta}\right)+O\left(\frac{1}{p^2}\right)\right]^{-1} \\
&= \frac{1}{\beta}\left[1-\frac{1}{p}\left(\frac{\alpha}{\beta}-\frac{1}{\lambda}\frac{(\beta-1)p_\varepsilon}{\beta^2}\right)\right]+O\left(\frac{1}{p^2}\right),
\end{aligned}$$

as claimed. QED

C.5 Proof Corollary 3

Theorem 3 implies that

$$\gamma^*(p)^2 = \frac{1}{\beta^2}\left[1-\frac{2}{p}\left(\frac{\alpha}{\beta}-\frac{1}{\lambda}\frac{(\beta-1)p_\varepsilon}{\beta^2}\right)\right]+O\left(\frac{1}{p^2}\right).$$

Given that $\dot{p}_t^* = p_\varepsilon\gamma(p_t^*)^2$, this implies:

$$\dot{p}_t^* = \frac{p_\varepsilon}{\beta^2} - \frac{1}{p_t^*}\left(\frac{2p_\varepsilon\alpha}{\beta^3} - \frac{2p_\varepsilon}{\lambda}\frac{\beta-1}{\beta^4}\right) + O\left(\frac{1}{(p_t^*)^2}\right),$$

and the result follows from Lemma 2. QED

References

Amador, Manuel and Pierre-Olivier Weill, "Learning by Matching," 2006. 8

_ **and** _ , "Addendum to "Learning from private and public observations of others' actions"," 2007. Working Paper, Stanford and UCLA. 6

_ **and** _ , "Learning from Prices: Public Communication and Welfare," 2007. Working Paper. 8

Angeletos, George-Marios and Alessandro Pavan, "Efficient Use of Information and Social Value of Information," *Econometrica*, 2007, *75*, 1103–1142. 4

Araujo, Luis and Andrei Shevchenko, "Price Dispersion, Information and Learning," *Journal of Monetary Economics, Forthcoming*, 2006. 6

_ **and Braz Camargo**, "Information, Learning, and the Stability of Fiat Money," *Journal of Monetary Economics, Forthcoming*, 2006. 6

Bala, Venkatesh and Sanjeev Goyal, "Learning from Neighbours," *Review of Economic Studies*, 1998, *65*, 595–621. 5

Banerjee, Abhijit and Drew Fudenberg, "Word-of-mouth learning," *Game and Economic Behavior*, 2004, *46*, 1–22. 5

Banerjee, Abhijit V, "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, August 1992, *107* (3), 797–817. 4

Bardi, Marco and Italo Capuzzo-Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, Boston-Basel-Berlin: Birkhauser, 1997. 30

Bergemann, Dirk and Jusso Välimäki, "Experimentation in Markets," *Review of Economic Studies*, 2000, *67*, 213–234. 5

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades," *Journal of Political Economy*, October 1992, *100* (5), 992–1026. 4

Blouin, Max R. and Roberto Serrano, "A Decentralized Market with Common Values Uncertainty: Non-Steady States," *Review of Economic Studies*, 2001, *68*, 323–346. 5

Bolton, Partick and Christopher Harris, "Strategic Experimentation," *Econometrica*, 1999, *67*, 349–374. 5

Burguet, Roberto and Xavier Vives, "Social learning and costly information acquisition," *Economic Theory*, 2000, *15* (1), 185–205. 4

Chamley, Christophe, *Rational Herds*, Cambridge: Cambridge University Press, 2004.

Décamp, Jean-Paul and Thomas Mariotti, "Investment timing and learning externality," *Journal of Economic Theory*, 2004, *18*, 80–102. 5

DeMarzo, Peter, Dimitri Vayanos, and Jeffrey H. Zwiebel, "Persuasion Bias, Social Influence, and Uni-Dimensional Opinions," *Quaterly Journal of Economics*, 2003, *118*, 909–968. 5

Duffie, Darrell and Gustavo Manso, "Information Diffusion in Large Population," *American Economic Review, Papers and Proceedings*, 2007, *97*. 6

Fernandez, Raquel, "Culture as Learning: The Evolution of Female Labor Force Participation over a Century," 2007. NYU mimeo.

Fogli, Alessandra and Laura Veldkamp, "Nature or Nurture? Learning and the Geography of Female Labor Force Participation," 2007. Federal Reserve Bank of Minneapolis Staff Report 386.

Gale, Douglas and Shachar Kariv, "Bayesian Learning in Social Networks," 2003. Working Paper, New York University. 5

Golosov, Mikhail, Aleh Tsyvinski, and Guido Lorenzoni, "Decenralized Trading with Private Information," 2007. Working Paper, MIT and Yale. 6

Green, Edward J., "Eliciting Traders' Knowledge in "Frictionless" Asset Markets, Working Paper," 1991. 5

Hellwig, Christian, "Heterogeneous Information and the Welfare Effects of Public Information Disclosures," 2005. 4

Jovanovic, Bojan and Yaw Nyarko, "A Bayesian Learning Model Fitted to a Variety of Empirical Learning Curves," *Brooking Papers: Microeconomics*, 1995, pp. 247–304.

Katzman, Brett, John Kennan, and Neil Wallace, "Output and Price Level Effects of Monetary Uncertainty in a Matching Model," *Journal of Economic Theory*, 2003, *108*, 217–255. 6

Keller, Godfrey, Sven Rady, and Martin Cripps, "Strategic Experimentation with exponential bandits," *Econometrica*, 2005, *73*, 39–68. 5

Kircher, Philipp and Andrew Postlewaite, "Strategic Firms and Endogenous Consumer Emulation," *Quarterly Journal of Economics*, 2008, *123*, 621–661. 5

Morris, Stephen and Hyun Song Shin, "The Social Value of Public Information," *American Economic Review*, 2002, *92*, 1521–1534. 4, 17

_ **and** _ , "Central Bank Transparency and the Signal Value of Prices," *Brookings Papers on Economic Activity*, 2005, *2*, 1–66. 4

Oksendal, Bernt, *Stochastic Differential Equations*, Berlin: Springer, 1985. 24

Protter, Philip, *Stochastic Integration and Differential Equations*, New York: Springer-Verlag, 1990. 6

Rosenberg, Dinah, Eilon Solan, and Nicolas Vieille, "Social learning in one-arm bandit problems," *Econometrica*, 2007, *75*, 1591–1611. 5

Rudin, Walter, *Real and Complex Analysis*, New York: McGraw-Hill, 1974. 30

Smith, Lones and Peter Norman Sørensen, "Rational Social Learning by Random Sampling," 2005. Working Paper, University of Michigan. 5

Trejos, Alberto and Randall Wright, "Search, Bargaining, Money, and Prices," *Journal of Political Economy*, 1995, *103* (1), 118–141. 6

Vives, Xavier, "How Fast do Rational Agents Learn?," *Review of Economic Studies*, 1993, *60*, 329–347. 2, 4, 12, 13

__ , "Learning from others: a Welfare Analysis," *Games and Economic Behavior*, 1997, *20*, 177–200. 2, 4, 18, 19

Wallace, Neil, "Short-Run and Long-Run Effects of Changes in Money in a Random-Matching Model," *Journal of Political Economy*, 1997, *105*, 1293–1307. 6

Wolinsky, Asher, "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, 1990, *58*, 1–23. 5