Learning by Matching *

PRELIMINARY AND INCOMPLETE, PLEASE DO NOT CIRCULATE

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

April 29, 2006

Abstract

We study how a continuum of agents learn about disseminated information in a dynamic beauty contest model when they do not observe aggregate variables, such as prices or quantities, but randomly observe each other's actions. We solve for the market equilibrium and find that the average learning curve is S-shaped: learning is slow initially, intensifies rapidly and finally converges slowly to the truth. We show that increasing public information always slows down learning in the long run. It also reduces welfare if agents are sufficiently patient, even when there is no coordination motive. Lastly, optimal diffusion of information requires that agents "strive to be different": agents need to be rewarded for choosing actions away from the population average.

*We'd like to thank, for fruitful discussions and suggestions, George-Marios Angeletos, Raquel Fernandez, Douglas Gale, Christian Hellwig, Boyan Jovanovic, Enrichetta Ravina, Stijn Van Nieuwerburgh, Laura Veldkamp, and Ivan Werning.

[†]Stanford University, e-mail: amador@stanford.edu

[‡]Department of Economics, University of California, Los Angeles, e-mail: poweill@gmail.com

1 Introduction

How does information diffuse in a population when there are no prices or quantities that aggregate the private information dispersed in the marketplace? What, if any, is the impact on diffusion dynamics and welfare of increasing public information? This paper addresses these questions in a dynamic beauty contest model based on two key assumptions. First, agents interact in a decentralized fashion and second, they cannot observe any endogenous aggregate.

We assume that, at the beginning of time, each of a continuum of agents receives both a private and a public signal about the state of the world. Every subsequent period, agents choose their actions to maximize a payoff which depends on i) the state of the world and ii) the average action in the population at the time. This dependence on population play directly creates a coordination motive. At the end of the period, every agent noisily observes the action of another randomly chosen agent. Because actions reflects current information, our continuum of agents progressively learn about the state of the world by randomly observing each others. This is the mechanism through which the initial private signals endogenously diffuse in the population.

We show that there exists an equilibrium where agents eventually learn all private information. The average belief in the population about the state of the world thus converges to the truth, but it does so along an S-shape curve as illustrated by the upper panel of Figure 1 and discussed further in Section 3.3. The learning curve is initially convex because of an information snowballing effect: agents learn from the learning of others. The learning curve is concave at the end because convergence to the truth implies that learning eventually slows down. In addition, because agents learn independently from one another, their learning histories are increasingly heterogeneous. This implies that the cross-sectional variance of beliefs increases at the beginning, as illustrated by the lower panel of Figure 1. This variance

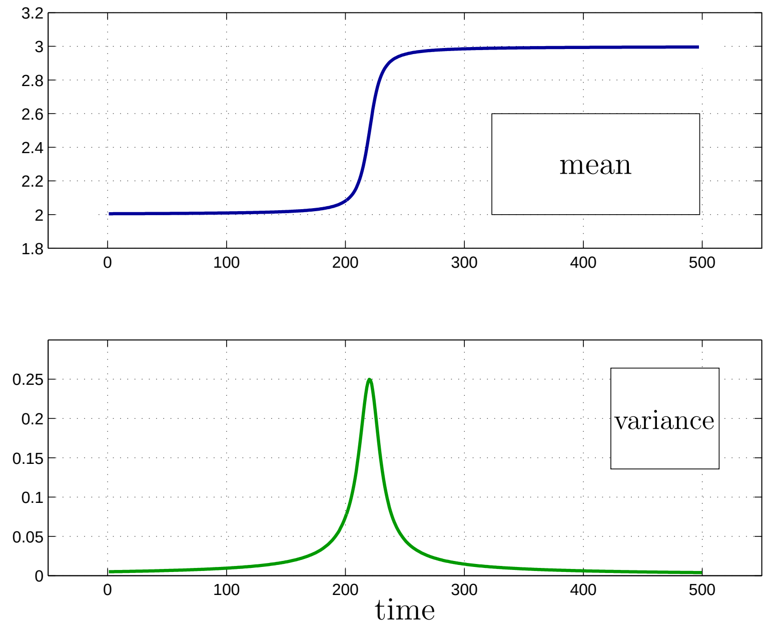

Figure 1: The distribution of beliefs over time.

eventually converges to zero as agents learn the truth.

Asymptotically, we show that the public information ends up crowding out private information: better public information at the beginning of time always slows down learning in the long run. This follows from an information externality. Indeed, with an increase in the precision of public information, an agent finds it optimal to load his action more heavily on the public signal than on his private information. The presence of observational noise implies that now it is harder for others to infer an agent's private information from his action. This effect slows down information diffusion. Note that, because of our continuum-of-players assumption, an agent has no incentive to take this effect into account when choosing his action.

Can it be possible then that better public information reduces welfare? By analyzing a continuous-time limit of the discrete-time model, we prove that a given marginal increase in the precision of the initial public signal is always welfare reducing if agents are sufficiently patient. In particular, the result encompasses the case when the agents have no coordination motive at all. Hence, differently from Morris and Shin (2002)), even in the absence of a payoff externality better public information can be welfare reducing.

In the last part of the paper we study the problem of a planner who seeks to maximize social welfare by telling agents what action to take as a time-varying affine function of their private beliefs. In the decentralized equilibrium, agents do not take into account the impact on aggregate learning of their actions and hence a learning externality appears. The planner internalizes this externality and would like agents to take actions more sensitive to their private beliefs. We show that this sensitivity changes non-monotonically over time. At the beginning, when information is very dispersed in the population, there is not much to learn from observing someone else's action and the planner prescribes a low sensitivity of actions to private information. After a while, learning has increased each agent's private information, and the planner finds it optimal to prescribe a high sensitivity. Eventually, agents know almost all the information. Then again, there is not much to learn from observing someone else's actions, and the planner prescribes a low sensitivity. Finally, we show that, in some cases, the planner's solution can be decentralized, in a beauty context spirit, by rewarding agents for taking actions away from the population play. In other words, agents should be rewarded for being different.

The results may apply to a broad range of economic interactions. For instance, in the macro economy, information is typically dispersed because households and firms know more about their local markets than about the economy as a whole. In addition, agencies collect and release macroeconomic information with long lags. In the meantime, firms and households learn about the state of the economy by interacting among each others.[1] One may also relate our setup to micro-level markets in which trade is typically bilateral and transaction prices are not released in real time. This is the case, for instance, for some over-the-counter

[1]This is the premise of Lucas (1972) and Phelps (1969). One may argue that asset prices efficiently aggregate private information. However, even asset prices appear to react to the release of macro information.

asset markets (see Edwards, Harris, and Piwowar (2004) for a study of the corporate-bond market).

Literature Review

Our work is related to the recent literature on the social value of public information (see Morris and Shin (2002), Hellwig (2005), and Angeletos and Pavan (2005)). In these models, public information may reduce welfare because a static payoff externality creates a coordination motive. Our contribution is to identify an alternative dynamic mechanism based on an information externality: in our model public information crowds out the diffusion of public information in the population.

Information externalities have been studied in the social learning literature (see, among many others, Vives (1993), Chamley and Gale (1994), and Vives (1997)). The maintained assumption of these models is that agents learn from public signals. The present paper makes the opposite assumption that, aside from the first period, agents do not observe any public signal. The two assumptions end up having strikingly different implications. Indeed, when agents learn from public signals, the learning speed is decreasing over time. This key implication is reversed in our model because agents learn from the learning of others, which creates an information snowballing effect: initially, learning speed increases over time. This implies that information diffuses along a S-shape, a pattern documented by a number of empirical studies of social learning (see Chapter 9 of Chamley (2004), and the reference therein). Recent work on social learning focused on learning in networks: Bala and Goyal (1998), Gale and Kariv (2003), Smith and Sorensen (2005) study deterministic networks with finite number of agents, Banerjee and Fudenberg (2004) provide a continuum-of-agents setup (see also DeMarzo, Vayanos, and Zwiebel (2003) for a network of boundedly rational agents). Because they lack tractability, these models end up focusing almost exclusively on the question of convergence to the truth. Our model of a random network with a continuum of agents can be solved in closed form, which allows us to take the learning-in-network literature a step further, with an analysis of transitional dynamics, welfare, and the impact of public information.

In Wolinsky (1990) seminal random-matching model of learning, information diffuses at the individual level but stays constant at the aggregate level: indeed, agents leave the economy after trading and uninformed agents continuously enter the economy. The issue of convergence when information diffuses on the aggregate has been subsequently addressed in Green (1991), Blouin and Serrano (2001), and in the independent work of Duffie and Manso (2006). Wallace (1997), Katzman, Kennan, and Wallace (2003), and Araujo and Shevchenko (2006) address learning about the money supply in Trejos and Wright (1995) random-matching model. For tractability, they assume that the money supply becomes public after either one or two periods. Araujo and Camargo (2006) relax this assumption in a Kiyotaki and Wright (1989) model, and study the government incentives to expand the money supply. Our setup is somewhat simpler than these models because agents do not learn from trading but from observing the action of others. The benefit of this simplification is that we can explicitly characterize the transitional dynamics of beliefs and study the welfare impact of public information.

The rest of the paper is organized as follows. Section 2 introduces the setup. Section 3 provides the transitional dynamics of the beauty contest equilibrium, and studies the impact on welfare and diffusion of a marginal increase in public information. Section 4 studies optimal information diffusion and section 5 concludes.

2 Setup: A Dynamic Beauty Contest

In this section we introduce the dynamic beauty contest. Our economy is populated by analysts who, every period, prepare a forecast of the state of the world. With some probability, at the end of every period, all analysts publicly announce their forecast, the state of the world is revealed, and an analyst's payoff is a function of how far his announced forecast is from i) the actual state of the world, and ii) the average forecast in the population. With the complementary probability, the state of the world is not revealed, each analyst gets to observe the forecast of a randomly chosen colleague up to some noise, and the economy moves to the next period.

The formal model is as follows. Time discrete and possibly runs to infinity. The economy is populated by a continuum of analysts indexed by $i \in [0,1]$. The state of the world is summarized by a parameter $\theta \in \mathbb{R}$, that all analysts take to be normally distributed with mean $\bar{\theta}$ and variance σ_0^2 (they share a common prior). For the rest of the paper, the common prior $\bar{\theta}$ is interpreted as a public signal: namely, at time $t = -1$, analysts have a completely diffuse prior and observe $\bar{\theta} = \theta + v_t$, for some v_t that is normally distributed with mean zero and variance σ_0^2.

However, analysts immediately become asymmetrically informed about θ: at the beginning of time, each analyst receives a signal $z_{i1} = \theta + w_{i1}$, where w_{i1} is normally distributed with mean zero and variance s_1^2. Signals are idiosyncratic in that the random variable w_{i1} is pairwise independent across analysts.

The timing of a period is as follows. At the start of each period $t \in \{1, 2, \ldots\}$, every analyst prepares his forecast $a_{it} \in \mathbb{R}$ for the period. At time $t + 1$, with probability $1 - \beta$, the game ends and, following Morris and Shin (2002), the analyst receives a payoff equal to:

$$U_{it} = -(a_{it} - \theta)^2 - \frac{b}{1-b}(L_{it} - \bar{L}_t), \tag{1}$$

where $b \in (-\infty, 1)$ and

$$L_{it} = \int_0^1 (a_{jt} - a_{it})^2 \, dj \tag{2}$$

$$\bar{L}_t = \int_0^1 L_{it} \, di. \tag{3}$$

An analyst trades off the distance of his announcement to the payoff-relevant parameter θ against the distance from the average announcement in the population. The parameter b captures the strength of the beauty contest: larger b means that an analyst worries more about staying close to average announcement. Lastly, the beauty contest is a zero-sum game. Indeed, the cross-sectional sum $\int_0^1 (L_{it} - \bar{L}) \, di$ of analysts' beauty-contest losses is equal to zero.

If θ is not revealed, then the game continues and every analyst observes the announcement of some other randomly chosen analyst, up to some noise. In particular, analyst i observes

$$a_{jt} + \varepsilon_{jt},$$

where j is drawn randomly according to a uniform distribution, independently across analysts and over time. Likewise, the noise is normally distributed with mean zero and variance σ_ε^2, and is idiosyncratic across analysts and over time.

Equilibria

The history for analyst i at time t is given by $h_{it} = \{z_{i1}, a_{j1} + \varepsilon_{j1}, \ldots, a_{jt-1} + \varepsilon_{jt-1}\}$. The strategies are mappings from the set of all possible histories at every time to possible announcements. The market solution is taken to be the Bayesian equilibrium of the dynamic game. At any point in time, given his beliefs, and taking as given the strategies of all other analysts, an analyst's strategy maximizes his expected payoff of the *current period*. This follows because a particular analyst's action is negligible by the continuum-of-agents assumption. In particular, let $a_i(h_{it})$ be the strategy of agent i with history h_i. Then it has to be the case that

$$a_i(h_{it}) = (1-b)E(\theta \,|\, h_{it}) + bE\left(\int_0^1 a_j(h_{jt})\,dj \,\middle|\, h_{it}\right), \tag{4}$$

which, together with the common prior assumption, implies that equilibrium strategies are symmetric.

3 Linear Equilibrium

In this section we characterize a Bayesian equilibrium of this beauty-contest game. We show that the learning curve is S-shaped and that a marginal increase in public information speeds up learning in the short run but slows it down in the long-run. In some case, when agents are sufficiently patient, it also reduces welfare.

3.1 Preliminary

We start by describing the analysts learning dynamics, under the three following hypotheses (which we verify hold in an equilibrium in the next subsection).

Hypothesis H1: at the beginning of each period $t \in \{1, 2, \ldots\}$, an analyst $i \in [0, 1]$ has observed a sequence $z_{i1}, \ldots, z_{it}$ of signals, where $z_{it} = \theta + w_{it}$, for some normal random variables w_{it} with mean zero and variance s_t^2.

Hypothesis H2: the sequence $w_{i1}, w_{i2}, \ldots$ is independent from θ for all $i \in [0, 1]$.

Hypothesis H3: the random variables w_{it} are almost surely independent across time and independent from $w_{i1}, w_{i2}, \ldots, w_{it-1}$. Moreover, for all $j \in [0, 1]$, w_{it} is almost surely independent from $w_{j1}, w_{j2}, \ldots, w_{jt}$.

We now provide a recursive characterization of the learning dynamics implied by Hypotheses (H1)-(H3). At the beginning of each period $t \in \{0, 1, 2, \ldots\}$, the prior of analyst $i \in [0, 1]$ is that θ is normally distributed with mean $\hat{\theta}_{it} \equiv E[\theta \mid z_{i1}, \ldots, z_{it}]$ and variance σ_t^2. Remember, the prior at the beginning of time (before receiving the first signal) is $\hat{\theta}_{i0} = \bar{\theta}$ and σ_0^2. We refer to the conditional expectation $\hat{\theta}_{it}$ as the "belief" of the analyst. We guess and verify in the proof of Proposition 1 that the cross-sectional distribution of these beliefs can be written

$$\hat{\theta}_{it} = (1 - x_t)\bar{\theta} + x_t\theta + u_{it}, \tag{5}$$

where $x_t \in [0, 1]$ is some constant, and u_{it} is a normal random variable with mean zero and variance τ_t^2, that is independent from θ. In words, equation (5) says that, conditional on the true value θ, the cross-sectional distribution of $\hat{\theta}_{it}$ is normally distributed with mean $(1-x_t)\bar{\theta}+x_t\theta$ and variance τ_t^2. The initial conditions are $x_0 = 0$ and $\tau_0 = 0$. The following Proposition applies standard linear-projection results (see for instance chapter 4 of Luenberger (1969)) in order to derive a recursive characterization of $\{x_t, \sigma_t^2, \tau_t^2\}_{t=0}^{\infty}$.

Proposition 1 (Learning dynamics). *For a given sequence $\{s_1^2, s_2^2, ...\}$ of variances and under hypothesis (H1)-(H3), at each time $t \in \{0, 1, \ldots, \}$, an analyst believes that θ is normally distributed with mean $\hat{\theta}_{it}$ and variance σ_t^2, where*

$$\hat{\theta}_{it+1} = \hat{\theta}_{it} + k_{t+1}(z_{it+1} - \hat{\theta}_{it}) \tag{6}$$

$$\sigma_{t+1}^2 = (1 - k_{t+1})\sigma_t^2 \tag{7}$$

where $k_{t+1} \equiv \sigma_t^2/(\sigma_t^2 + s_{t+1}^2)$, and σ_0^2 is given. In addition, the parameters governing cross-sectional distribution (5) of $\hat{\theta}_{it}$ are

$$x_{t+1} = 1 - \frac{\sigma_{t+1}^2}{\sigma_0^2} \tag{8}$$

$$u_{it+1} = (1 - k_{t+1})u_{it} + k_{t+1}w_{it+1} \tag{9}$$

$$\tau_{t+1}^2 = \sigma_0^2 x_{t+1}(1 - x_{t+1}). \tag{10}$$

Proof. In the appendix. □

Note that x_t is the reduction in the variance of an analyst's posterior relative to his initial belief. The following result will prove useful,

Lemma 1. *For a given sequence $\{s_1^2, s_2^2, \ldots\}$ of variances, and under hypotheses (H1)-(H3), the variance reduction x_t follows the recursion*

$$x_{t+1} = 1 - (1 - x_t)\frac{s_{t+1}^2/\sigma_0^2}{s_{t+1}^2/\sigma_0^2 + (1 - x_t)}, \tag{11}$$

for all $t \in \{1, 2, \ldots\}$ and with $x_1 = \sigma_0^2/(s_1^2 + \sigma_0^2)$.

Proof. To obtain the recursion for x_t, divide both sides of equation (7) by σ_0^2. The result follows by plugging the value of k_{t+1} and using equation (8). □

This previous result can be easily interpreted in terms of precisions. Note that $1/(\sigma_0(1 - x_t))$ is the precision of an agent beliefs about θ at time t. After observing the signal with precision $1/s_{t+1}$, the agent's new precision is the sum of his previous precision plus the precision of the new signal: $1/(\sigma_0(1 - x_{t+1})) = 1/(\sigma_0(1 - x_t)) + 1/s_{t+1}$. This is equivalent to equation (11).

So far in this section we have characterized the learning dynamics under hypotheses (H1)-(H3), and given a sequence $\{s_1^2, s_2^2, \ldots\}$ of variances. To obtain that sequence and to verify the validity of our hypotheses, we need to solve for the actions taken by the analysts in equilibrium. The next section proceeds to construct such an equilibrium.

3.2 A Linear Equilibrium

In this subsection we characterize an equilibrium in which an analyst's announcement is affine.

Suppose that the announcement of analyst $j \in [0, 1]$ at time $t \in \{1, 2, \ldots\}$ can be written as a linear combination of his time 0 prior and his current beliefs $F_t\bar{\theta} + G_t\hat{\theta}_{jt}$, for some $(F_t, G_t) \in \mathbb{R}_+^2$ to be determined. Note that, because all analysts are using the same linear coefficients at any time t, their different histories affect their actions only through their posterior beliefs. Substituting (5) we have that

$$a_{jt} = F_t\bar{\theta} + G_t\left((1-x_t)\bar{\theta} + x_t\theta + u_{jt}\right). \tag{12}$$

The recursion (9) implies that, for all $j \in [0,1]$, u_{jt} is a linear combination of $w_{j1}, w_{j2}, \dots, w_{jt}$. Then, by the induction hypothesis (H3), the u_{jt} are almost surely independent across analysts. Therefore, given that the u_{jt} have zero mean given θ, an informal application of the Law of Large Numbers shows that the average announcement is

$$\int_0^1 a_{jt}\,dj = F_t\bar{\theta} + G_t\left((1-x_t)\bar{\theta} + x_t\theta\right). \tag{13}$$

This implies that an analyst with history h_{it} expects the average announcement to be

$$E\left[\int_0^1 a_{jt}\,dj \,\middle|\, h_{it}\right] = F_t\bar{\theta} + G_t\left((1-x_t)\bar{\theta} + x_t\hat{\theta}_{it}\right).$$

Hence, equation (4) shows that analyst i's best reply is

$$a_{it} = (1-b)\hat{\theta}_{it} + b\left[F_t\bar{\theta} + G_t\left((1-x_t)\bar{\theta} + x_t\hat{\theta}_{it}\right)\right]. \tag{14}$$

Identifying with the coefficients of (12) and (14) shows that, in a linear equilibrium

$$F_t = \frac{b(1-x_t)}{1-bx_t} \tag{15}$$

$$G_t = \frac{1-b}{1-bx_t}. \tag{16}$$

In particular $F_t = 1 - G_t$, implying that an analyst's forecast can be written $\bar{\theta} + G_t\left(\theta_{it} - \bar{\theta}\right)$. If $b = 0$ there is no beauty contest and $G_t = 1$, meaning that the equilibrium analyst's strategy is simply to announce his belief $\hat{\theta}_{it}$. If $b \in (0,1)$ then $G_t \in (0,1)$ implying that analysts underweight their belief θ_{it} relative to the common prior $\bar{\theta}$: analysts strive to look alike. Conversely if $b < 0$, then $G_t > 1$ and analysts strive to look different.

The last thing to do in order to complete our characterization of an equilibrium, is to verify that our maintained hypotheses (H1)-(H3) hold, and to determine the sequence s^2_{t+1}.

Equation (12) implies that observing the announcement of a randomly chosen agent $j \in [0,1]$ up to some noise ε_{jt} amounts to observing

$$z_{it+1} = \theta + \frac{u_{jt}}{x_t} + \frac{\varepsilon_{jt}}{G_t x_t} \equiv \theta + w_{it+1}, \tag{17}$$

where

$$w_{it+1} \equiv u_{jt}/x_t + \varepsilon_{jt}/(G_t x_t), \tag{18}$$

which verifies hypothesis (H1). Hypothesis (H2) and (H3) are verified because of the following intuitive reason. Random matching with a continuum of analysts implies that any two analysts have almost surely observed different colleagues at any previous time. Also, any of those colleagues have almost surely observed different colleagues previously, and so on. This together with the normality and the linear strategies, imply that conditional on θ, the signals received by observing others analysts are normally distributed, independent through time and across analysts. The formal statement is as follows.

Proposition 2 (Existence). *There exists a linear equilibrium in which hypotheses (H1)-(H3) hold. The coefficients F_t and G_t of the linear strategy are given by (15) and (16). The variance of w_{it+1} is*

$$s_{t+1}^2 = \sigma_0^2 \frac{x_t(1-x_t) + \alpha(1-bx_t)^2/(1-b)^2}{x_t^2} \tag{19}$$

and the "variance reduction" x_t evolves according to

$$x_{t+1} = H(x_t, \alpha, b) \equiv x_t + \frac{x_t^2(1-x_t)^2}{x_t(1-x_t^2) + \alpha(1-bx_t)^2/(1-b)^2} \tag{20}$$

for all $t \in \{1, 2, \ldots\}$, where $\alpha \equiv \sigma_\varepsilon^2/\sigma_0^2$ and with initial condition $x_1 = \sigma_0^2/(s_1^2 + \sigma_0^2)$.

Proof. In the appendix. □

3.3 Information Aggregation Dynamics

We first study the dynamics of the "variance reduction" variable x_t.

Proposition 3 (Asymptotic revelation). *If $x_1 = 0$, then $x_t = 0$ for all $t \in \{1, 2, \ldots\}$. Otherwise, if $x_1 \neq 0$, then the variance reduction goes to 1 as t goes to infinity.*

Proof. The function $H(\,\cdot\,, \alpha, b)$ is continuous and such that i) $H(x, \alpha, b) > x$ for all $x \in (0, 1)$ and ii) $H(0, \alpha, b) = 0$ and $H(1, \alpha, b) = 1$. □

If $x_1 = 0$, there is nothing to learn and an analyst's belief stays the same forever. If $x_1 > 0$ then, asymptotically, an analyst's belief converges to the truth. The following Proposition shows that time path of the cross-sectional distribution (5) of beliefs has the two following qualitative features:

Proposition 4 (S-shaped average, Hump-shaped variance). *The average belief $\hat{\theta}_t \equiv \int_0^1 \hat{\theta}_{it}\, di$ converges to θ along a S-shaped curve. Namely, there is some time $t_s \geq 0$ such that $|\hat{\theta}_{t+1} - \hat{\theta}_t|$ is increasing if and only if $t \leq t_s$. The variance τ_t^2 of the cross-sectional belief distribution converges to zero following a hump-shaped curve. Namely, there is some time $t_h \geq 0$ such that $\tau_{t+1}^2 - \tau_t^2 \geq 0$ if and only if $t \leq t_h$.*

Proof. The first point requires some brute force in the appendix. The second point follows from (10). □

The results of the Proposition are illustrated by the numerical calculations of Figure 2. The associated parameter values, used in all the numerical examples of this paper, are summarized in Table 1. The upper panel shows the time path of the average belief $\hat{\theta}_t$, assuming that the state of the world is $\theta = 3$. The lower panel shows the time path of the variance τ_t^2 of the cross-sectional belief distribution.

The learning curve is convex at the beginning for the following reason. By observing a random colleague, an analyst effectively observes the average belief $\hat{\theta}_t \equiv (1 - x_t)\bar{\theta} + x_t\theta$ up

Table 1: Parameter Values.

Parameter		Value
Variance of the prior	σ_0^2	1
Variance of the observational noise	σ_ε^2	1
Implied noise to signal ratio	α	1
Variance of the private signal noise	s_1^2	199
Implied initial variance reduction	x_1	0.005
Initial belief	$\bar{\theta}$	2
Beauty contest intensity	b	0
Probability of continuing	β	0.7

to two noises: a "sampling" noise u_{jt} and an exogenous observational noise ε_{jt}. However, as analysts learn, the average belief $\hat{\theta}_t$ loads more and more on the state of the world θ. This mitigates the negative impact on learning of the two noises and initially accelerates learning. Note that, because of convergence to the true value, learning cannot accelerate forever. Hence, at the end, the learning curve must be concave.

Lastly, note that if x_1 is large enough then $t_s = 0$. In that case, learning immediately start in the upper branch of the S, and learning speed is decreasing over time. The condition that x_1 is large enough is met when s_1^2/σ_0^2 is small. In other words, when information is not too dispersed so that analysts learn a great deal from their initial private information, then there is no information snowballing effect.

The hump-shape of the cross-sectional variance follows because analysts have independent learning histories that lead them to learn the same thing. Namely, at time zero, analysts' beliefs are all the same and the initial signals $\theta + w_{i1}$ create heterogenous beliefs. This implies that the distribution of beliefs fans out ($\tau_1 > \tau_0 = 0$). By continuity, if the distribution of beliefs remains concentrated ($\tau_1 \simeq 0$), then then the distribution continues to fan out ($\tau_2 > \tau_1$). Asymptotically analysts agree again, implying that τ_t must converges to zero.

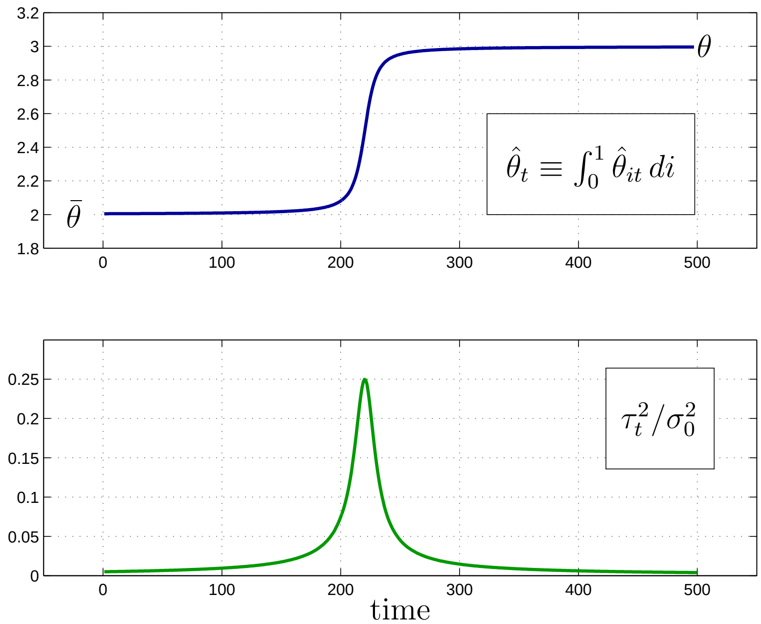

Figure 2: Aggregate Learning Dynamics.

The next proposition shows the intuitive results that reducing σ_ε^2, or reducing b leads to faster learning: indeed, analysts learn more from each others if the observational noise is smaller, or if other analysts make announcements that are more sensitive to their current beliefs.

Proposition 5 (A comparative static). *Consider* $(\sigma_\varepsilon^2(1), \sigma_\varepsilon^2(2)) \in \mathbb{R}_+^2$ *such that* $\sigma_\varepsilon^2(1) < \sigma_\varepsilon^2(2)$*, let* $x_1(k) = x_1$ *and* $x_{t+1}(k) = H(x_t(k), \sigma_\varepsilon^2(k)/\sigma_0^2, b)$ *for* $k \in \{1, 2\}$*. Then,* $x_t(1) > x_t(2)$ *for all* $t \in \{2, 3, \ldots\}$*. Similarly, consider* $(b(1), b(2)) \in (-\infty, 1)^2$ *such that* $b(1) < b(2)$*, let* $x_1(k) = x_1$ *and* $x_{t+1}(k) = H(x_t(k), \alpha, b(k))$ *for* $k \in \{1, 2\}$*. Then,* $x_t(1) > x_t(2)$ *for all* $t \in \{2, 3, \ldots\}$*.*

Proof. Follows from the fact that $H(x, \alpha, b)$ is strictly decreasing in both α and b. □

The following Proposition characterizes the asymptotic learning speed.

Proposition 6 (Asymptotic learning speed). *The sequence of* σ_t^2 *admits the following* *asymptotic expansion*

$$\sigma_t^2 = \frac{\sigma_\varepsilon^2}{t} + \sigma_\varepsilon^2 \left(1 + \frac{2\sigma_\varepsilon^2}{(1-b)\sigma_0^2}\right) \frac{\log(t)}{t^2} + O\left(\frac{1}{t^2}\right), \tag{21}$$

where $O(1/t^2)$ is a sequence bounded by M/t^2, for some $M \in \mathbb{R}_+$.

Proof. In the appendix. □

This tells us two things. First, to a first order approximation, learning occurs at speed σ_ε^2/t, as if each analyst were receiving a signal $\theta + \varepsilon_{it}$ every period. So, in our setup with independent learning histories, social learning resembles single-agent learning in the limit. This is in sharp contrast with the setup of Vives (1997) in which analysts share a common learning history. The second message of the Proposition is that public information has a negative impact on learning in the long run.

Corollary 1 (Short- and long-run impacts of public information). *Consider $(\sigma_0^2(1), \sigma_0^2(2))$ such that $\sigma_0^2(1) < \sigma_0^2(2)$. Let $x_1(k) = \sigma_0^2(k)/(s_1^2 + \sigma_0^2(k))$, $x_{t+1}(k) = H(x_t(k), \sigma_\varepsilon^2/\sigma_0^2(k), b)$, and $\sigma_t^2(k) = \sigma_0^2(k)(1 - x_t(k))$. Then, there exists some $1 \leq T_s < T_\ell$ such that, $\sigma_t^2(1) < \sigma_t^2(2)$ for all $t \in \{1, \ldots, T_s\}$ and $\sigma_t^2(1) > \sigma_t^2(2)$ for all $t \in \{T_\ell, T_\ell + 1, \ldots\}$.*

Proof. The result for $t \in \{1, \ldots, T_s\}$ follows from (7). For $t \in \{T_\ell, T_\ell + 1, \ldots\}$, it follows from Proposition (6). □

Imagine that, at time zero, analysts receive some public information regarding the state of the world. This decreases σ_0^2, and therefore speeds up learning in the short run. However, Corollary 1 shows that it slows learning down in the long run. Indeed, with better public information, an analyst's forecasts puts a higher weight on the common prior $\bar{\theta}$, and a lower weight on their belief θ_{it}. Together with the observational noise, this implies that analysts have less to learn from observing each others, and slows down learning.

3.4 The Welfare Cost of Public Information

This subsection shows that, in our setup, public information can reduce welfare. Our finding holds even when $b = 0$. Hence, in contrast with Morris and Shin (2002), the negative impact on welfare of public information does not rely on analysts having a coordination motive.

3.4.1 Utilitarian Welfare

We take our welfare criterion to be the equally weighted sum of analysts' expected utility. By the Law of Large Number, this criterion coincides with the ex-ante utility of a representative analyst

$$-E\left[\sum_{t=1}^{\infty}(1-\beta)\beta^{t-1}\left(a_t(h_{it})-\theta\right)^2\right] \tag{22}$$

where $(1-\beta)\beta^{t-1}$ is the probability that the game ends at time $t \in \{1, 2, \ldots\}$. Now, from equation (12), we know that $a_t(h_{it}) - \theta = (1 - G_t x_t)(\bar{\theta} - \theta) + u_{it}$, where u_{it} is independent from θ, has a mean zero and variance τ_t^2. This implies that

$$\begin{aligned} E\left[(a_t(h_{it})-\theta)^2\right] &= (1-x_tG_t)^2 E\left[(\theta-\bar{\theta})^2\right] + G_t^2 E\left[u_{it}^2\right] \\ &= \left[\sigma_0^2\left(1-x_t\right)\right]\frac{(1+(b-2)bx_t)}{(1-bx_t)^2} \equiv \sigma_t^2 w(x_t), \end{aligned} \tag{23}$$

where the second line follows from plugging (16) and $\tau_t^2 = \sigma_0^2 x_t(1-x_t)$ into the equation, and noting that $\sigma_t^2 = \sigma_0^2(1-x_t)$. So the utilitarian criterion can be written:

$$-\sum_{t=1}^{\infty}(1-\beta)\beta^{t-1}\sigma_t^2 w(x_t). \tag{24}$$

The flow welfare at time t is the variance of beliefs σ_t^2 multiplied by some function $w(x_t)$ that is greater than 1, and which is equal to 1 when $b = 0$.

3.4.2 A Continuous-time Approximation

This subsection proposes a continuous-time approximation of our setup which greatly facilitates welfare analysis. The approximation is obtained by letting the observational noise grow very large while, at the same time, letting analysts observe each others more and more frequently. Formally, let Δ be the amount of time between periods, and let us index the economy by Δ.

Assumption 1. The observational variance $\sigma_\varepsilon^2(\Delta)$ is such that, as Δ goes to zero, $\Delta\sigma_\varepsilon^2(\Delta)$ goes to some $\tilde{\sigma}_\varepsilon^2 \in \mathbb{R}_+$.

In order to obtain a proper continuous time limit, we also require that the probability $1-\beta(\Delta)$ of ending the game goes to zero in order Δ.

Assumption 2. The probability $\beta(\Delta)$ of continuing the game is such that, as Δ goes to zero, $(1-\beta(\Delta))/\Delta$ goes to some $r \in \mathbb{R}_+$.

In other words, in the limit as Δ goes to zero, the game ends at some Poisson arrival time with intensity r. Under these two assumptions, the evolution equation for the variance reduction is

$$\frac{x_{t+\Delta} - x_t}{\Delta} = \frac{x_t^2(1-x_t)^2}{x_t(1-x_t^2)\Delta + \Delta\sigma_\varepsilon^2(\Delta)/\sigma_0^2(1-bx_t)^2/(1-b)^2}$$

which, by taking the limit as Δ goes to zero, becomes

$$\dot{x}_t = \frac{\sigma_0^2}{\tilde{\sigma}_\varepsilon^2}(1-b)^2 \frac{x_t^2(1-x_t)^2}{(1-bx_t)^2} \tag{25}$$

The Welfare criterion (24) is

$$-\sum_{k=1}^{\infty}(1-\beta)\beta^k\sigma_{k\Delta}^2 w(x_{k\Delta}) = -\sum_{k=1}^{\infty}(r\Delta + o(\Delta))\, e^{-(r+o(1))k\Delta}\sigma_{k\Delta}^2 w(x_{k\Delta})$$

which, by taking the limit as Δ goes to zero, becomes

$$W(\sigma_0^2) = -r\int_1^{\infty}\sigma_t^2 w(x_t)\, e^{-r(t-1)}dt, \tag{26}$$

with initial condition $x_1 = \sigma_0^2/(\sigma_0^2 + s_1^2)$ and where $\sigma_t^2 = \sigma_0^2(1 - x_t)$.[2] We now derive a closed-form solution for the integral (26), which intuitively follows from guessing that $W(\sigma_0^2) = J(x_1)$, where the function $J(\cdot)$ solves the Hamilton-Jacobi-Bellman equation

$$rJ(x) = -r\sigma_0^2(1-x)w(x) + J'(x)\frac{\sigma_0^2}{\tilde{\sigma}_\varepsilon^2}(1-b)^2 \frac{x^2(1-x)^2}{(1-bx)^2}.$$

Direct integration of this ODE with the change of variable $y = (1-b)x/(1-x)$ and given the boundary condition $J(1) = 0$ provides a

Lemma 2 (closed-form Solution). *Let $G(y) \equiv y + 2\log(y) - 1/y$ and $f(y) \equiv (1 + (1 - b)y)/y^2/(1-b)$. Then, we have*

$$W(\sigma_0^2) = -r\tilde{\sigma}_\varepsilon^2 \int_{y_1}^{\infty} f(y) e^{-r\tilde{\sigma}_\varepsilon^2/s_1^2 \frac{G(y)-G(y_1)}{y_1}} dy, \tag{27}$$

where $y_1 = (1-b)\sigma_0^2/s_1^2$.

Proof. In the appendix. □

Based on formula (27), we show:

Theorem 1 (Welfare Cost of Public Information). *For all $\sigma_0^2 > 0$ the exists some $\eta > 0$ such that, for all $r\tilde{\sigma}_\varepsilon < \eta$, $W'(\sigma_0^2) > 0$.*

Proof. In the appendix. □

Theorem 1 implies that for any level σ_0^2 of information, a marginal increase in public information reduces welfare as long as the intensity r of finishing the game is low enough. Hence, in contrast with Morris and Shin (2002), an increase in public information can reduce welfare even if $b = 0$ and analysts' payoffs induce no coordination motives.

[2] We cannot provide an approximation theorem stating that the learning dynamics in the continuous time limit is indeed the limit of discrete-time learning dynamics as the time Δ between period goes to zero. We conjecture as much, and proceed. In appendix C, we provide numerical calculations suggesting that the result of Theorem 1 also holds in discrete time.

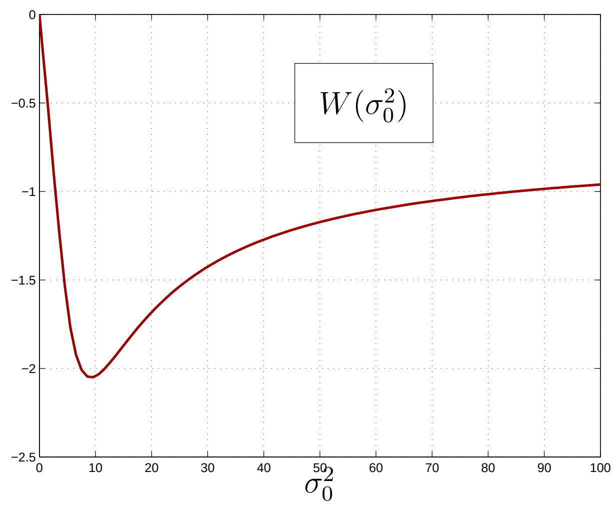

Figure 3: Welfare in the continuous-time model as a function of σ_0^2.

The Theorem does not imply, however, that $W(\sigma_0^2)$ can be monotonically increasing in σ_0^2. Indeed, revealing the state of the world would clearly improve welfare. By continuity, one might expect that a sufficiently large release of public information would also improve welfare. This intuition is confirmed by the numerical calculation of Figure 3: it shows that the function $W(\cdot)$ is non-monotonic. It first decreases but eventually increases if σ_0^2 is large enough.

4 Optimal Information Diffusion

Our dynamic beauty contest exhibits an information externality. Namely, in an equilibrium, an analyst does not internalize that his announcement constitutes valuable information for the analyst who is spying on him. This section addresses this externality by studying a problem of optimal information diffusion, subject to the learning technology. The tradeoff faced by the planner is as follows: in a static world (only one period) it would be efficient for the analysts to announce their beliefs. However, because of the dynamic nature of the problem, future analysts learn from the announcements made today, and hence a planner would like the analysts to make announcements that are even more sensitive to their beliefs. This generates a loss today in case that the game ends: indeed, analysts announcement will be far from the actual parameter value. On the other hand, this improves the dissemination of information tomorrow in case that the game continues.

We show that the planner requires that analysts strive to be different: they should make their forecast more sensitive to their private beliefs than in the static optimum. In addition, the optimal sensitivity varies non-monotonically over time. It is small at the beginning, large in the middle, and small again at the end.

4.1 The Planning Problem

A planner chooses functions $a_t(\cdot)$ mapping histories h_t into announcements, in order to maximize the ex-ante utility of a representative analyst:

$$-\sum_{t=1}^{\infty}(1-\beta)\beta^{t-1}\int (a_t(h_{it})-\theta)^2 \; P_t(dh_{it}, d\theta), \tag{28}$$

where P_t is the joint probability distribution over histories h_{it} and the state of the world θ. The planner is constrained by the learning technology which means that, at each time $t \in \{2, 3, \ldots, \}$, the probability distribution P_t is obtained by an application of Bayes' rule given P_{t-1} and given that analyst $i \in [0,1]$ observes the announcement $a(h_{jt}) + \varepsilon_{jt}$ of some randomly chosen analyst $j \neq i$.

In this section, we follow Vives (1997) and restrict attention to the class of time-varying affine announcements, whereby an analyst's announcement is restricted to be

$$a_{it} = F_t\bar{\theta} + G_t\hat{\theta}_{it}, \tag{29}$$

for some time-varying constants F_t and G_t and where, as before, $\hat{\theta}_{it}$ is an analyst expectation of θ conditional on his history h_{it-1}. In words, equation (29) means that an analyst announcement must be an affine function of his conditional expectations $\hat{\theta}_{it}$. Although the existence of a linear equilibrium makes it natural to study affine announcements, we could not prove that an unrestricted optimum is indeed affine. In section 4.3, we illustrate one virtue of an optimal affine announcements: as long as G_t is not too large, it can be implemented by letting analysts play a beauty contest game with an appropriately chosen weight $b \in (-\infty, 1)$.

Given our restriction (29), we can solve for the learning dynamics exactly as in the previous section. In particular, the results of Proposition 1 hold, with $s_{t+1}^2 = \sigma_0^2/x_t^2\left(x_t(1-x_t) + \alpha/G_t^2\right)$. Plugging this back into (11) and rearranging gives the transition function

$$x_{t+1} = x_t \left(1 + \frac{G_t^2 x_t (1-x_t)^2}{G_t^2 x_t (1-x_t^2) + \alpha}\right) \equiv g(x_t, G_t^2). \tag{30}$$

Let's now turn to the planner's objective. We first note that, by the Law of Iterated Expectations, $E(\hat{\theta}_{it}) = \bar{\theta}$. By definition, $V(\theta - \hat{\theta}_{it}) = \sigma_t^2$. Lastly, because $\hat{\theta}_{it}$ is a conditional expectation, if follows that $\hat{\theta}_{it}$ is orthogonal to $\theta - \hat{\theta}_{it}$. Therefore $V(\theta) = \sigma_0^2 = V(\hat{\theta}_{it}) + V(\theta - \hat{\theta}_{it})$, implying that $V(\hat{\theta}_{it}) = \sigma_0^2 - \sigma_t^2$. One can also verify these results by working directly on equation (5) for the cross-sectional distribution of $\hat{\theta}_{it}$. Taken together, these remarks imply that the planner's flow utility is

$$-E\left[(a_{it} - \theta)^2\right] = -\overline{\theta}^2 (F_t + G_t - 1)^2 - \sigma_0^2 (G_t - 1)^2 x_t - \sigma_0^2 (1 - x_t).$$

Note that the control F_t does not enter the transition function (30). Hence, maximizing the objective with respect to F_t reduces to a static quadratic optimization problem, whose solution is $F_t = 1 - G_t$. Replacing into the objective and ignoring the constants we can let the planner's flow utility be $x_t(G_t - G_t^2/2)$. Lastly, note that we can restrict attention to positive G_t. Indeed, if $G_t < 0$, then applying $-G_t$ yields a higher flow utility and leaves x_{t+1} unchanged. Hence, we can make the change of variable $\gamma_t \equiv G_t^2$ and let $G_t = \sqrt{\gamma_t}$.

4.2 Striving to be Different

An *admissible control* is a positive sequence $c = \{\gamma_t\}_{t=1}^{\infty}$. The set of admissible controls is denoted by $\mathcal{C}$. Given an admissible control c, the state x_t^c evolves according to the difference equation $x_{t+1}^c = g(x_t^c, \gamma_t)$, for all $t \geq 1$, where $x_1^c = x_1$ is given. The planner's inter-temporal utility is

$$U(x_1, c) = \sum_{t=1}^{\infty}(1-\beta)\beta^{t-1}u(x_t^c, \gamma_t) \ ,$$

where $u(x, \gamma) = x(\sqrt{\gamma} - \gamma/2)$. The *planner's problem* in sequence form is then

$$W(x_1) = \sup_{c \in \mathcal{C}} U(x_1, c). \tag{31}$$

Our first result in this section is to show that $W(\,\cdot\,)$ is strictly increasing: the higher the precision of analysts' beliefs, the higher the value to the planner:

Proposition 7 (Monotonicity). *The value function $W(\,\cdot\,)$ is strictly increasing. In addition, for every $x \in (0,1)$, there exists some $m \in \mathbb{R}_+$ and some $\varepsilon \in \mathbb{R}_+$ such that $0 < x' - x < \varepsilon$ implies that $W(x') - W(x) > m(x' - x)$.*

Proof. In the appendix. □

The second part of the proposition shows that the welfare gains from increasing x are (at least) of first-order. We study the following Bellman operator

$$T(f)(x) = \sup_{\gamma \in \mathbb{R}_+} \{(1-\beta)u(x,\gamma) + \beta \ f \circ g(x, \gamma)\} \,. \tag{32}$$

We apply standard dynamic programming arguments in the following Banach space. Given $(k, \eta) \in \mathbb{R}_+$, we let $\mathcal{X}(k, \eta)$ be the set of continuous functions $f : [0,1] \to \mathbb{R}_+$ such that $f(0) = 0$, bounded above by $1/2$ and such that, for all $0 < x < x' < \eta$,

$$f(x') - f(x) \leq k(x' - x). \tag{33}$$

Clearly, $\mathcal{X}(k, \eta)$ is a Banach space when equipped with the sup norm.[3]

Lemma 3. *The following holds,*

(i) For every $f \in \mathcal{X}(k, \eta)$, the supremum on the right-hand side of (32) is achieved.

(ii) For every $f \in \mathcal{X}(k, \eta)$, Tf is continuous and bounded.

(iii) There exists $(k, \eta) \in \mathbb{R}^2$ such that if $f \in \mathcal{X}(k, \eta)$ then $Tf \in \mathcal{X}(k, \eta)$.

Proof. In the appendix. □

This leads to the following Proposition

Proposition 8. *The operator T is a contraction mapping $\mathcal{X}(k, \eta)$ into itself. Therefore it has a unique fixed point $W(\cdot)$. Also, the function $W(\cdot)$ is the value of the planner's optimal control problem: it satisfies (31).*

Proof. The contraction follows by noticing that T satisfies Blackwell sufficient conditions for a contraction. Hence, the Contraction Mapping Theorem applies to (32). Lastly, given that the flow utility is bounded above, we can apply the Bellman Principle, implying that the function $W(\cdot)$ is indeed the value of the planner's. (See Theorems 3.2, 3.3, and 4.3 in Stokey and Lucas (1989)). □

The static planning problem ($\beta = 0$) is to maximize $u(x, \gamma)$ with respect to $\gamma \geq 0$. Its solution is to set $\gamma = 1$. In other words, the planner prescribes analysts to announce their belief $\hat{\theta}_{it}$. This result no longer holds in the dynamic problem under consideration when $\beta > 0$, because increasing γ_t speeds up learning. Indeed, (30) shows that, as long as $\alpha > 0$, x_{t+1} increases with γ_t. This is a symptom of the externality we seek to study and

[3] One might wonder what makes condition (33) useful: indeed, it is not needed for applying the Contraction Mapping Theorem, nor to show that the solution of the Bellman equation coincides with the planner's value function. The condition turns out to be useful for establishing properties of the planner's policy function. We come back to this remark when discussing the results of Proposition 9.

follows because analysts observe noisy observation of each others' announcements, of the form $(1-\sqrt{\gamma_t})\bar{\theta}+\sqrt{\gamma_t}\,\hat{\theta}_{jt}+\varepsilon_{jt}$. Therefore, increasing γ_t increases the signal-to-noise ratio and hence the informativeness of the announcement. Note however that if $\alpha=0$, then there is no noise and an analyst conditional expectation can be inferred perfectly from his action. In that case, γ has no impact on the dynamics of x_t and the planning solution is to let $\gamma_t=1$ at each time.

One might guess from this discussion that the planner finds it optimal to let $\gamma_t>1$ at each time because this allows the planner to increase x_{t+1}. This intuition is confirmed in the following Proposition:

Proposition 9. *Let* $\Gamma:[0,1]\to[0,\infty)$ *be the maximum correspondence of (32) when* $f=W$*. Then,*

(i) For all $x\in(0,1)$*,* $\Gamma(x)\subset(1,\infty)$

(ii) In an optimal solution, x_t^* *goes to 1 as time goes to infinity.*

(iii) Any optimal control is such that γ_t^* *goes to one as time goes to infinity.*

(iv) For any sequence $x_k\to 0$ *and any* $\gamma_k\in\Gamma(x_k)$*, we have that* $\gamma_k\to 1$ *as* $k\to\infty$*.*

Proof. Part (i) In the appendix. □

Part (i) of the Proposition tells that the planner finds it optimal to internalizes the information externality by prescribing $\gamma>1$. Indeed, because $u(x,\gamma)$ is maximized at $\gamma=1$, the welfare loss of increasing γ above one is of second order. On the other hand, Proposition 7 implies that the welfare gain of increasing x_{t+1} are (at least), of first order. Part (ii) of the Proposition tells that, asymptotically in the planner's solution, there is full revelation of the state of the world: the variance of beliefs goes to zero. Part (iii) says that an optimal control converges to 1 as time goes to infinity: there is nothing to learn in the limit, and the control approaches the static solution. This follows from point (*ii*), together with the fact that the maximum correspondence is upper hemi-continuous and satisfies $\Gamma(1) = \{1\}$. Part (*iv*) tells that any selection of the maximum correspondence $\Gamma(\cdot)$ is not monotonic. Specifically, the planner prescribes $\gamma \simeq 1$ for x close to zero, $\gamma \simeq 1$ for x close to 1, and $\gamma > 1$ for x bounded away from 0 and 1. That is, the social optimum (approximately) coincides with the private optimum at the boundaries $x \in \{0, 1\}$. In between, the planner speeds up information diffusion by prescribing $\gamma > 1$. Some intuition goes as follows. If no signal is revealed or if all information is revealed, everybody has the same posterior. This implies that analysts have nothing to learn from observing their colleagues' announcements. As a result, the dynamic optimum must coincide with the static optimum. By continuity, close to those extremes, the dynamic optimum almost coincides with the static optimum.

The main difficulty of Proposition 9 is to prove part (*iv*). One might think that it follows from upper hemi-continuity at zero: in fact, because $\Gamma(0) = \mathbb{R}$, upper hemi-continuity at zero imposes no restriction on the behavior of the maximum correspondence for x close to zero. Another approach is to take first-order conditions in equation (32) and write

$$(1-\beta)\frac{\partial u}{\partial \gamma} + \beta \frac{dW}{dx} \circ g(x,\gamma) \times \frac{\partial g}{\partial \gamma}(x,\gamma) = 0.$$

Now, as x goes to zero, we have that $\partial g / \partial \gamma$ also goes to zero, meaning that there is and less gain from increasing γ above one. Thus, one might expect that γ goes to one as x goes to zero. This argument runs into two difficulties. First, as x goes to zero, the loss $\partial u / \partial \gamma$ of increasing γ also goes to zero. And second, the value function needs not be differentiable, as our problem does not satisfy the convexity conditions required for an application of the Benveniste and Scheinkman (1979) Theorem. The second difficulty can be circumvented using property (33), which shows that the slope of the value function is bounded above for x close to zero.

4.3 Numerical Example

In this section we provide a numerical illustration of our results. We solve the planning problem on Matlab with a standard value-function-iteration algorithm (see, e.g., Chapter 12 of Judd (1999)). Figure 4 shows that the value function is indeed increasing in the variance reduction x and appears to be strictly quasi-concave but not concave. Figure 5 confirms that the maximum correspondence is not monotonic. Lastly, the upper panel of Figure 6 confirms that information diffuses faster under the planning solution.

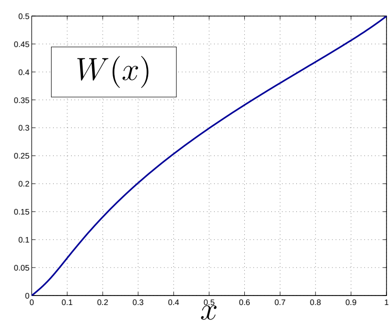

Figure 4: Value Function.

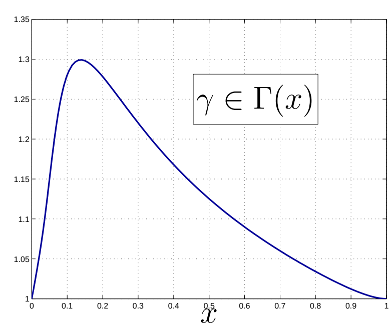

Figure 5: Policy Function.

The lower panel of Figure 6 provides a numerical answer to the following implementation question: is there a sequence of beauty-contest games implementing the planning solution? In other words, can we pick sequence $\{b_t\}_{t=1}^{\infty}$ of beauty-contest parameters such that, at each time, analysts announcements are socially optimal. One easily sees that this amounts to pick

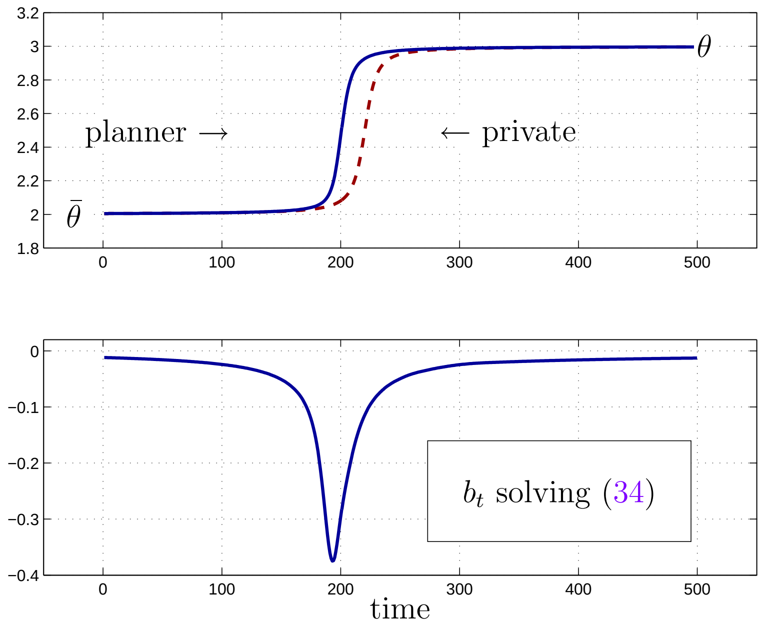

Figure 6: Aggregate Learning Dynamics.

some b_t such that

$$\sqrt{\gamma_t} = \frac{1 - b_t}{1 - b_t x_t}. \tag{34}$$

Since the planner prescribes $\gamma_t > 1$, it must be that $b_t < 0$, meaning that the planner gives monetary reward for making announcements *away* from the population average. What might prevent implementation is that, in a beauty-contest equilibrium, the weight $(1-b)/(1-bx_t)$ that an analyst puts on his own belief $\hat{\theta}_{it}$ is bounded above by $1/x_t$. Therefore, the planning solution can be implemented in a beauty-contest game if and only if, for every $x \in [0, 1]$, there is some $\gamma \in \Gamma(x)$ such that $\gamma \leq 1/x^2$. This property clearly holds when $x \simeq 0$ because, in that region, every $\gamma \in \Gamma(x)$ is close to 1. Unfortunately, we are not able to prove that this property holds for all $x \in [0, 1]$. The calculations shown in the lower panel of Figure 6 suggest however that, for some parameter values, the social optimum can be implemented in a sequence of beauty-contest games.

5 Conclusion

This papers studies how private information diffuses among a continuum of agents who interact at random. We show that agents learn the truth along a S-shape learning curve. In particular, at the beginning there is an information snowballing effect because agents learn from the learning and others. We show that larger public information at the beginning always slows down the the diffusion of private information in the economy, and sometimes reduce welfare. Further work may address the optimal timing of public information release in this economy.

A Appendix

A.1 Proof of Proposition 1

Standard projection formula (see, e.g., Luenberger (1969)) imply equations (6) and (7). Substituting (5) into (6) and identifying unknown coefficients, we obtain the recursions (8) and (9) of the Proposition. Note that, since w_{it+1} is independent from θ, our guess that u_{it+1} is independent from θ is verified. Because u_{it} is a linear combination of $w_{i0}, \ldots, w_{it}$, our hypothesis (H3) implies that it is independent from w_{it+1}. Hence, taking variance on both side of (9) implies equation (10) of the Proposition. Now, equation (8) can be written $1-x_{t+1} = (1-k_{t+1})(1-x_t)$, and equation (7) can be written $\sigma_{t+1}^2 = (1-k_{t+1})\sigma_t^2$. Hence, $1-x_{t+1}$ and σ_{t+1}^2 solve the same linear difference equation. This implies that the ratio $(1-x_t)/\sigma_t^2$ stays constant over time, that is $1-x_{t+1} = \sigma_{t+1}^2/\sigma_0^2$ which is recursion (8) of the Proposition. Because for $k_{t+1} = \sigma_t^2/(\sigma_t^2 + s_{t+1}^2)$, one easily verifies that equation (7) is equivalent to

$$\sigma_{t+1}^2 = (1-k_{t+1})^2\sigma_t^2 + k_{t+1}^2 s_{t+1}^2. \tag{35}$$

Subtracting equation (10) from equation (35) and dividing both sides by σ_0^2, we find

$$\frac{\sigma_{t+1}^2 - \tau_{t+1}^2}{\sigma_0^2} = (1-k_{t+1})^2 \frac{\sigma_t^2 - \tau_t^2}{\sigma_0^2}.$$

Therefore, the sequence $(1-x_t)^2$ and the sequence $(\tau_t^2 - \sigma_t^2)/\sigma_0^2$ solve the same linear difference equation. Because, $\tau_0 = 0$, they also have the same initial condition, implying that

$$\frac{\sigma_t^2 - \tau_t^2}{\sigma_0^2} = (1-x_t)^2 = \frac{\sigma_t^4}{\sigma_0^4}.$$

Rearranging, we obtain $\tau_t^2 = \sigma_0^2 x_t(1-x_t)$, which is equation (10) of the Proposition.

A.2 Proof of Proposition 2

Equation (9) implies that u_{jt} is a linear combination of $w_{j1}, w_{j2}, \ldots, w_{jt}$. By the induction hypothesis, all of the $w_{j1}, \ldots, w_{jt}$ are independent from θ. Since ε_{jt} is also independent from θ, our hypothesis (H2) that w_{it+1} is independent from θ is verified.

The last thing to verify is hypothesis (H3). First consider any $s \leq t$ and some $\ell \in [0,1]$. Since $\ell \neq j$ almost surely, our induction hypothesis implies that $w_{\ell s}$ is almost surely independent from $w_{j1}, w_{j2}, \ldots, w_{jt}$, and hence from u_{jt} (which is a linear combination of . Since $w_{\ell s}$ is also independent from ε_{jt}, our hypothesis that $w_{\ell s}$ is independent from w_{it+1} is verified. Now we have for any $\ell \neq i$:

$$w_{\ell t+1} = \frac{1}{x_t}\left(u_{nt} + \frac{1}{G_t}\varepsilon_{nt}\right),$$

for some $n \in [0,1]$ which is almost surely different from j. Our induction hypothesis implies that the sequences $w_{j1}, \ldots, w_{jt}$ and the sequences $w_{n1}, \ldots, w_{nt}$ are independent from one another, and therefore that u_{nt} is independent from u_{jt}. Since by assumption ε_{jt} is independent from ε_{nt}, our induction hypothesis that $w_{\ell t+1}$ is independent from w_{it+1} is verified.

Taking the variance of both sides of (18), substituting (16), we find

$$s_{t+1}^2 = \sigma_0^2 \frac{x_t(1-x_t) + \alpha(1-bx_t)^2/(1-b)^2}{x_t^2} \tag{36}$$

for all $t \in \{1, 2, \ldots\}$, with $\alpha \equiv \sigma_\varepsilon^2/\sigma_0^2$. Plugging (36) into (11) and rearranging gives the result. At time $t=0$ we have that

$$\frac{\sigma_1^2}{\sigma_0^2} = \left(1 + \frac{\sigma_0^2}{s_1^2}\right)^{-1}$$

meaning that $x_1 = s_1^2/(s_1^2 + \sigma_0^2)$, so we can also think of x_1 as a primitive.

A.3 Proof of Proposition 4

The average belief is $\hat{\theta}_t = (1-x_t)\bar{\theta} + x_t\theta$, implying that $|\hat{\theta}_{t+1} - \hat{\theta}_t| = (x_{t+1} - x_t)|\theta - \bar{\theta}|$. Now the recursion for x_t can be written $x_{t+1} = x_t + h(x_t)$. Some simple algebra shows that

$$\frac{\partial h}{\partial x} = \frac{x(1-x)^2}{(1-b)^2(x(1-x^2) + \alpha(1-bx)^2/(1-b)^2)^2} P(x),$$

where

$$\begin{aligned} P(x) &= (1-b)^2 x(1-2x-x^2) + 2\alpha(1-bx)\frac{1-2x+bx^2}{1-x} \\ &\equiv (1-b)^2 x R(x) + 2\alpha(1-bx)\frac{T(x)}{1-x}. \end{aligned}$$

Evidently, $P(0) = 2\alpha$ and $P(x) \to -\infty$ when $x \to 1$. So there exists some $x^* \in (0,1)$ such that $P(x^*) = 0$. Note that $R(\cdot)$ (respectively) $T(\cdot)$ has only one root x_R (respectively x_T) in the interval $[0,1]$. Because $R(x) < T(x)$, we must have $x_R < x_T$. Given that $R(x^*)$ and $T(x^*)$ must have opposite signs, it follows that $x^* \in (x_R, x_T)$. Now,

$$\begin{aligned} P'(x^*) &= (1-b)^2 R(x^*) + (1-b)^2 x^* R'(x^*) - 2\alpha b\frac{T(x^*)}{1-x^*} \\ &\quad + 2\alpha(1-bx^*)\frac{(1-x^*)T'(x^*) + T(x^*)}{(1-x^*)^2}. \end{aligned}$$

The first three terms are negative because $x_R > x^*$ implies that $R(x^*) < 0$, because $R'(x^*) < 0$, and because $x^* < x_T$ implies that $T(x^*) > 0$. As for the last term, we have

$$(1-x^*)T'(x^*) + T(x^*) = -1 + 2bx - bx^2 \leq -1 + 2bx - b^2x^2 \leq -(1-bx)^2 \leq 0,$$

because $b \in (0,1)$. Therefore, it follows that $P'(x^*) < 0$, establishing that x^* is the only zero of $P(\cdot)$ over the interval $[0,1]$. The above analysis shows that $x_{t+1} - x_t \equiv h(x_t)$ is increasing for $x_t \in (0, x^*)$, and decreasing for $x_t \in (x^*, 1)$. The time t_s of the Proposition is then largest time t such that $x_t \leq t_s$.

A.4 Proof of Proposition 6

We start with the change of variable $\pi_t \equiv 1/\sigma_t^2$. Then, it follows from equation (7) that

$$\pi_{t+1} = \pi_t + \frac{1}{s_{t+1}^2}.$$

In turns, equation (19) shows that

$$\begin{aligned} \frac{1}{s_{t+1}^2} &= \pi_0\left(1 - \frac{\pi_0}{\pi_t}\right)^2 \left\{\left(1 - \frac{\pi_0}{\pi_t}\right)\frac{\pi_0}{\pi_t} + \alpha\frac{\left[1 - b\left(1 - \frac{\pi_0}{\pi_t}\right)\right]^2}{[1-b]^2}\right\}^{-1} \\ &= \frac{\pi_0}{\alpha}(\pi_t - \pi_0)^2 \left\{\frac{(\pi_t - \pi_0)\pi_0}{\alpha} + \frac{[\pi_t - b(\pi_t - \pi_0)]^2}{[1-b]^2}\right\}^{-1} \\ &= \frac{\pi_0}{\alpha}(\pi_t - \pi_0)^2 \left\{(\pi_t - \pi_0)^2 + \pi_0\left[(\pi_t - \pi_0)\left(\frac{1}{\alpha} + \frac{2}{1-b}\right) + \frac{\pi_0}{(1-b)^2}\right]\right\}^{-1} \\ &= \frac{\pi_0}{\alpha}\left\{1 - \frac{\pi_0\left[(\pi_t - \pi_0)\left(\frac{2\alpha + (1-b)}{\alpha(1-b)} + \frac{\pi_0}{(1-b)^2}\right)\right]}{(\pi_t - \pi_0)^2 + \pi_0\left((\pi_t - \pi_0)\frac{2\alpha + 1 - b}{\alpha(1-b)} + \frac{\pi_0}{(1-b)^2}\right)}\right\} \\ &= \frac{\pi_0}{\alpha}\left\{1 - \frac{\pi_0(2\alpha + 1 - b)}{\alpha(1-b)}\left[\frac{1}{\pi_t} + O\left(\frac{1}{\pi_t^2}\right)\right]\right\}, \end{aligned}$$

Now, plugging back $\alpha = \sigma_\varepsilon^2/\sigma_0^2 = \pi_0/\pi_\varepsilon$ into the last equation, we obtain

$$\pi_{t+1} = \pi_t + \pi_\varepsilon \left[1 - \frac{1}{\pi_t}\left(\pi_\varepsilon + \frac{2\pi_0}{1-b}\right)\right] + O\left(\frac{1}{\pi_t^2}\right). \tag{37}$$

We already know that $\pi_t \to +\infty$ as $t \to +\infty$. Hence, there is a T such that, for all $t \geq T$,

$$\pi_\varepsilon \left[1 - \frac{1}{\pi_t}\left(\pi_\varepsilon + \frac{2\pi_0}{1-b}\right)\right] + O\left(\frac{1}{\pi_t^2}\right) \geq \frac{\pi_\varepsilon}{2},$$

implying that $\pi_{t+1} \geq \pi_t + \pi_\varepsilon/2$. Summing from T to $t \geq T$, we obtain that $\pi_t \geq \pi_T + \pi_\varepsilon/2\,(t-T)$ for all $t \geq T$, and therefore that $1/\pi_t = O(1/t)$ as t goes to infinity. Plugging this back into (37) gives $\pi_{t+1} = \pi_t + \pi_\varepsilon + O\,(1/t)$. Summing over times then implies that

$$\pi_t = \pi_\varepsilon t + O(\log(t)). \tag{38}$$

Taking the inverse of (38) gives

$$\frac{1}{\pi_t} = \frac{1}{\pi_\varepsilon t}\left[1 + O\left(\frac{\log(t)}{t}\right)\right]^{-1} = \frac{1}{\pi_\varepsilon t} + O\left(\frac{\log(t)}{t^2}\right),$$

where the last equality follows because $\log(t)/t$ goes to zero as t goes to infinity. Plugging back this last equation into (37) provides

$$\pi_{t+1} = \pi_t + \pi_\varepsilon \left\{1 - \left[\frac{1}{\pi_\varepsilon t} + O\left(\frac{\log(t)}{t^2}\right]\right)\left(\pi_\varepsilon + \frac{2\pi_0}{1-b}\right)\right\} + O\left(\frac{1}{t^2}\right). \tag{39}$$

Summing over times once again shows that

$$\pi_t = \pi_\varepsilon t - \left(\pi_\varepsilon + \frac{2\pi_0}{1-b}\right)\log(t) + O(1), \tag{40}$$

where the $O(1)$ term follows from the fact that the series $1/t^2$ and $\log(t)/t^2$ are absolutely convergent. The Proposition then follows directly from inverting (40).

A.5 Proof of Lemma 2

After making the change of variable $y_t = (1-b)x_t/(1-x_t)$, the ordinary differential equation (25) can be written

$$\dot{y}_t = \gamma\left(\frac{y_t}{1+y_t}\right)^2 = \frac{\gamma}{G'(y_t)} \tag{41}$$

where $\gamma \equiv (1-b)\sigma_0^2/\tilde{\sigma}_\varepsilon^2$ and $G(\,\cdot\,)$ is the strictly increasing function $G\,(y) \equiv y - 1/y + 2\log\,(y)$. Multiplying both sides of (41) by $G'(y_t)$ and integrating from $s=1$ to $s=t$ shows that the solution y_t of the ODE is defined implicitly by

$$\gamma(t-1) = \; G\,(y_t) - G\,(y_1)\,. \tag{42}$$

plugging $x_t = y_t/(1-b+y_t)$ and $y_t = G^{-1}\,(\gamma(t-1) + G(y_1))$ back into the integral (24), we make the change of variable $t - 1 = 1/\gamma(G(y) - G(y_1)$. We obtain, after some algebra, the formula of the Lemma.

A.6 Proof of Theorem 1

Let $\delta = \tilde{\sigma}_\varepsilon^2/s_1^2$. Because $y_1 = (1-b)\sigma_0^2/s_1^2$, welfare increases in σ_0^2 if and only if

$$V\,(y_1) = -\int_{y_1}^{\infty} \frac{(1+(1-b)y)}{(1-b)\,y^2} e^{-r\delta\frac{G(y)-G(y_1)}{y_1}} dy$$

is increasing in y_1. Taking the derivative with respect to y_1, we obtain

$$V'(y_1) = \left(\frac{1+(1-b)y_1}{y_1^2(1-b)}\right) - \int_{y_1}^{\infty}\left(\frac{1+(1-b)y}{y^2(1-b)}\right)\frac{\partial}{\partial y_1}\left(-r\delta\frac{[G(y)-G(y_1)]}{y_1}\right)e^{-r\delta\frac{[G(y)-G(y_1)]}{y_1}}dy.$$

Using the fact that $\int_{y_1}^{\infty} r\delta\frac{G'(y)}{y_1}e^{-r\delta\frac{[G(y)-G(y_1)]}{y_1}}dy = 1$ we get that

$$V'(y_1) =$$
$$\int_{y_1}^{\infty}\left[\left(\frac{1+(1-b)y_1}{y_1^2(1-b)}\right)\frac{G'(y)}{y_1} - \left(\frac{1+(1-b)y}{y^2(1-b)}\right)\frac{\partial}{\partial y_1}\left(-\frac{[G(y)-G(y_1)]}{y_1}\right)\right]r\delta e^{-r\delta\frac{[G(y)-G(y_1)]}{y_1}}dy,$$

which can be written more compactly as into:

$$V'(y_1) = \int_{y_1}^{\infty}\frac{G'(y)}{y_1}\Phi(y,y_1)r\delta e^{-r\delta\frac{[G(y)-G(y_1)]}{y_1}}dy,$$

where

$$\Phi(y,y_1) = \left(\frac{1+(1-b)y_1}{y_1^2(1-b)}\right) - \left(\frac{1+(1-b)y}{y^2(1-b)}\right)\frac{1}{G'(y)}\frac{[G(y)-G(y_1)]}{y_1} - \left(\frac{1+(1-b)y}{y^2(1-b)}\right)\frac{G'(y_1)}{G'(y)}$$

Because $G'(y) = \frac{(1+y)^2}{y^2}$, we have

$$\Phi(y,y_1)(1-b) = \left(\frac{1+(1-b)y_1}{y_1^2}\right) - \left(\frac{1+(1-b)y}{(1+y)^2}\right)\frac{G(y)-G(y_1)}{y_1} - \left(\frac{1+(1-b)y}{(1+y)^2}\right)G'(y_1)$$

So calculating the limit when $y \to \infty$, we find that

$$\lim_{y\to\infty}\Phi(y,y_1)(1-b) = \left(\frac{1+(1-b)y_1}{y_1^2}\right) - \frac{1-b}{y_1} = \frac{1}{y_1^2} > 0.$$

Therefore there exists some $\varepsilon > 0$ and some $y^* \in (y_1,\infty)$ such that $\Phi(y,y_1) > \varepsilon$ for all $y \geq y^*$. Then we can write

$$\begin{aligned}
V'(y_1) &= \int_{y_1}^{y^*}\Phi(y,y_1)\frac{G'(y)}{y_1}r\delta e^{-r\delta\frac{G(y)-G(y_1)}{y_1}}dy \\
&\quad + \int_{y^*}^{\infty}\Phi(y,y_1)\frac{G'(y)}{y_1}r\delta e^{-r\delta\frac{G(y)-G(y_1)}{y_1}}dy \\
&\geq \int_{y_1}^{y^*}\Phi(y,y_1)\frac{G'(y)}{y_1}r\delta e^{-r\delta\frac{G(y)-G(y_1)}{y_1}}dy \\
&\quad +\varepsilon\int_{y^*}^{\infty}\frac{G'(y)}{y_1}r\delta e^{-r\delta\frac{G(y)-G(y_1)}{y_1}}dy
\end{aligned}$$

Now remember that $\delta = \tilde{\sigma}_\varepsilon^2/s_1^2$ and take the limit as $r\tilde{\sigma}_\varepsilon^2$ goes to zero, to find that

$$\lim_{r\tilde{\sigma}_\varepsilon^2\to 0}\int_{y_1}^{y^*}\Phi(y,y_1)\frac{G'(y)}{y_1}r\delta e^{-r\delta\frac{G(y)-G(y_1)}{y_1}}dy = 0, \tag{43}$$

after noticing that $|\Phi(y,y_1)|$ is continuous and hence bounded for over $[y_1,y^*]$, with $y_1 > 0$. On the other hand, the second term is

$$\begin{aligned}
\varepsilon\int_{y^*}^{\infty}\frac{G'(y)}{y_1}r\delta e^{-r\delta\frac{G(y)-G(y_1)}{y_1}}dy &= \varepsilon e^{-r\delta\frac{G(y^*)-G(y_1)}{y_1}}\left[-e^{-r\delta\frac{G(y)-G(y^*)}{y_1}}\Big|_{y^*}^{\infty}\right] \\
&= \varepsilon e^{-r\delta\frac{G(y^*)-G(y_1)}{y_1}} \geq 0. \qquad (44)
\end{aligned}$$

Taken together, (43) and (44) imply that

$$\lim_{r\tilde{\sigma}_\varepsilon^2\to 0}V'(y_1) \geq \varepsilon > 0$$

and we are done.

A.7 Proof of Proposition 7

Taking derivatives shows that

$$\frac{\partial u}{\partial x} = \sqrt{\gamma} - \gamma/2 \tag{45}$$

$$\frac{\partial u}{\partial \gamma} = x\left(\frac{1}{2\sqrt{\gamma}} - \frac{1}{2}\right) \tag{46}$$

$$\frac{\partial g}{\partial x} = \frac{\alpha^2 + x(1-x)\left[2\alpha\gamma(2-x) + 2\gamma^2 x(1-x)\right]}{\left[\alpha + \gamma x(1-x^2)\right]^2} \tag{47}$$

$$\frac{\partial g}{\partial \gamma} = \alpha \frac{x^2(1-x)^2}{\left[\alpha + \gamma x(1-x^2)\right]^2} \tag{48}$$

$$\frac{\partial u}{\partial x}\frac{\partial g}{\partial \gamma} - \frac{\partial u}{\partial \gamma}\frac{\partial g}{\partial x} \tag{49}$$

$$= \frac{1}{D}\Bigg((\sqrt{\gamma} - \gamma/2)\alpha x^2(1-x)^2 \tag{50}$$

$$-x/2(1/\sqrt{\gamma} - 1)(\alpha^2 + x(1-x)(2\alpha\gamma(2-x) + 2\gamma^2 x(1-x))\Bigg),$$

where $D = \left[\alpha + \gamma x(1-x^2)\right]^2$. Given that $\sqrt{\gamma} - \gamma/2 \geq \sqrt{\gamma} - \gamma$, a sufficient condition for (49) to be strictly positive is

$$\left(\frac{1}{\sqrt{\gamma}} - 1\right)\left(\gamma\alpha x^2(1-x)^2 - x/2(\alpha^2 + 2\alpha\gamma x(1-x)(2-x) + 2\gamma^2 x(1-x))\right) > 0$$

$$\Leftrightarrow \left(\frac{1}{\sqrt{\gamma}} - 1\right)\left(\gamma\alpha x(1-x)\left(x - x^2 - 2x + x^2\right) - x/2(\alpha^2 + 2\gamma^2 x(1-x))\right) > 0$$

$$\Leftrightarrow \left(1 - \frac{1}{\sqrt{\gamma}}\right)\left(\gamma\alpha x^2(1-x) + x/2(\alpha^2 + 2\gamma^2 x(1-x))\right) > 0 \tag{51}$$

Now let's consider an initial condition $x_1 \in (0,1)$ together with some optimal control c.[4] The associated sequence of state is $\{x_t^c\}_{t\geq 1}$. Let's also consider some other initial condition $\hat{x}_1 > x_1$. We have

$$W(\hat{x}_1) - W(x_1) = W(\hat{x}_1) - U(x_1, c) \geq U(\hat{x}_1, \hat{c}) - U(x_1, c), \tag{52}$$

for any admissible control $\hat{c}$. We pick $\hat{c}$ such that $U(\hat{x}_1, \hat{c}) > U(x_1, c)$, as follows: we let $\hat{\gamma}_t = 1$ as long as $g(\hat{x}_t^c, 1) \geq g(x_t^c, \gamma_t)$. At the first time τ such that $g(\hat{x}_\tau^c, 1) < g(x_\tau^c, \gamma_\tau)$, we choose the $\hat{\gamma}_\tau$ solving $g(\hat{x}_\tau, \hat{\gamma}_\tau) = g(x_\tau^c, \gamma_\tau)$. Thereafter, for all $t > \tau$, we let $\hat{\gamma}_t = \gamma_t$. To summarize, for all $t < \tau$, $\hat{\gamma}_t = 1$ and $\hat{x}_t^c \geq x_t^c$. For $t = \tau$, $1 < \hat{\gamma}_\tau \leq \gamma_\tau$ and $\hat{x}_\tau^c \geq x_\tau^c$. For $t > \tau$, $\hat{\gamma}_t = \gamma_t$ and $\hat{x}_t^c = x_t^c$. Therefore, we have

$$\frac{U(\hat{x}_1, \hat{c}) - U(x_1, c)}{1 - \beta} = \sum_{t=1}^{\tau-1} \beta^t \left(u(\hat{x}_t^c, 1) - u(x_t^c, \gamma_t)\right) + \beta^\tau \left(u(\hat{x}_\tau^c, \hat{\gamma}_\tau) - u(x_t^c, \gamma_t)\right). \tag{53}$$

The first $\tau - 1$ terms are all strictly positive because $\max_{\gamma \geq 0} u(x, \gamma) = u(x, 1) = x/2$ which is increasing in x. Because $\hat{x}_{\tau+1}^c = x_{\tau+1}^c$, the time-$\tau$ term can be written $v(\hat{x}_\tau^c) - v(x_\tau^c)$ where $v(x) = u(x, \psi(x))$ and $\psi(x)$ solves $g(x, \psi(x)) = x_{\tau+1}^c$. An application of the Implicit Function Theorem (see, e.g. Taylor and Mann (1983)) shows that $\psi'(x) = -(\partial g/\partial x)/(\partial g/\partial \gamma) < 0$. Hence

$$v'(x) = \frac{\partial u}{\partial x} - \frac{\partial u}{\partial \gamma}\frac{\partial g/\partial x}{\partial g/\partial \gamma}.$$

[4]Existence of an optimal control follows from the dynamic programming argument of Proposition 8

Because $\partial g/\partial\gamma > 0$, this shows that $v'(x) > 0$ if and only if

$$\frac{\partial u}{\partial x}\frac{\partial g}{\partial \gamma} - \frac{\partial u}{\partial \gamma}\frac{\partial g}{\partial x} > 0. \tag{54}$$

Now we note that, by construction, $\psi(\hat{x}_t^c) > 1$. Because $\psi(\,\cdot\,)$ is decreasing, $\psi(x) > 1$ for all $x \in [x_\tau^c, \hat{x}_\tau^c]$. Using (54) and (51), one can see that $v'(x) > 0$ for all $x \in [x_\tau^c, \hat{x}_\tau^c]$. Therefore, the time-$\tau$ term of (53) is also strictly positive. This shows that the value function is strictly increasing.

For the second part of the proof, we note that if $\tau > 1$ in equation (53), then (52) shows that

$$\frac{W(\hat{x}_1) - W(x_1)}{1-\beta} \geq u(\hat{x}_1, 1) - u(x_1, \gamma_1) \geq u(\hat{x}_1, 1) - u(x_1, 1) = \frac{1}{2}\left(\hat{x}_1 - x_1\right), \tag{55}$$

where the second inequality follows because $u(x, \cdot\,)$ is maximized at $\gamma = 1$. If, on the other hand, $\tau = 1$ in equation (53), then (52) implies that

$$\frac{W(\hat{x}_1) - W(x_1)}{1-\beta} \geq v(\hat{x}_1) - v(x_1) \geq \frac{v'(x_1)}{2}(\hat{x}_1 - x_1), \tag{56}$$

for $\hat{x}_1$ close enough to x_1, and where $v(\,\cdot\,)$ is the function defined above. Letting $m \equiv (1-\beta)\min\{v'(x_1)/2, 1/2\}$ completes the proof.

A.8 Proof of Lemma 3

Part (i) We have

$$\begin{aligned}
&(1-\beta)u(x,\gamma) + \beta f \circ g(x,\gamma) \\
&\leq (1-\beta)u(x,\gamma) + \beta/2 = (1-\beta)x(\sqrt{\gamma} - \gamma/2) + \beta/2.
\end{aligned} \tag{57}$$

This means that, given some $x \in (0, 1]$, for $\gamma \geq \bar{\gamma}$, the left-hand side of (57) is negative. Since the right-hand side of (32) is positive at $\gamma = 0$, it follows that the suppremum over $\mathbb{R}_+$ is equal to the supremum over the compact $[0, \bar{\gamma}]$. Because of continuity, the supremum is achieved.

Part (ii)**.** Continuity of Tf follows from the Theorem of the Maximum (see Theorem 3.6 in Stokey and Lucas (1989)). Lastly, because $u(x,\gamma) \leq x/2 \leq 1/2$, we have $Tf \leq 1/2$.

Part (iii) In this paragraph we show how to pick some $(k, \eta) \in \mathbb{R}_+^2$ such that, if f satisfies the Lipschitz property (33), then Tf also satisfies it. Consider some $(k, \eta) \in \mathbb{R}_+^2$ and pick some $x \in (0, \eta)$. Let γ be a maximizer of (32) at x. For any $\gamma' < \gamma$, we have

$$\begin{aligned}
&(1-\beta)u(x,\gamma) - (1-\beta)u(x,\gamma') + \beta f \circ g(x,\gamma) - \beta f \circ g(x,\gamma') \geq 0 \\
\Rightarrow\quad &(1-\beta)u(x,\gamma) - (1-\beta)u(x,\gamma') + \beta k(g(x,\gamma) - g(x,\gamma')) \geq 0 \\
\Rightarrow\quad &(1-\beta)\frac{\partial u}{\partial \gamma}(x,\gamma) + \beta k \frac{\partial g}{\partial \gamma}(x,\gamma) \geq 0 \\
\Rightarrow\quad &(1-\beta)\frac{x}{2}\left(\frac{1}{\sqrt{\gamma}} - 1\right) + \beta k \frac{\alpha x^2(1-x)^2}{[\alpha + \gamma x(1-x^2)]^2} \\
\Rightarrow\quad &(1-\beta)\frac{x}{2}\left(\frac{1}{\sqrt{\gamma}} - 1\right) + \frac{\beta}{\alpha} k x^2 (1-x)^2 \geq 0 &(58)\\
\Rightarrow\quad &\gamma \leq \frac{1-\beta}{[1-\beta-2\beta k/\alpha x(1-x)^2]^2} \equiv \phi(k,x), &(59)
\end{aligned}$$

for η small enough. In the above, the first line follows because γ is a maximizer of (32), the second line follows from (33) together with the fact that $g(x,\gamma)$ is increasing in γ, the third line follows from dividing both side by $\gamma - \gamma' \geq 0$ and letting $\gamma' \to \gamma$, the fourth line follows by substituting in the expression (45) and (48) for the partial derivatives, the fifth line follows from $\alpha + \gamma x(1-x^2) \geq \alpha$, and the last line from rearranging, noting that $x > 0$ and that, if η is small enough, then $1 - \beta - 2\beta/\alpha k x(1-x)^2$ is positive for all $x \in (0, \eta)$. Now consider $0 < x < x' < \eta$. Let γ and γ' the respective maximizers of (32). We have

$$\begin{aligned}
& Tf(x') - Tf(x) \\
= \;& (1-\beta)u(x',\gamma') + \beta f \circ g(x',\gamma') - (1-\beta)u(x,\gamma) - \beta f \circ g(x,\gamma) \\
= \;& (1-\beta)u(x',\gamma') - (1-\beta)u(x,\gamma') + \beta f \circ g(x',\gamma') - \beta f \circ g(x,\gamma') \\
& + (1-\beta)u(x,\gamma') + \beta f \circ g(x,\gamma') - (1-\beta)u(x,\gamma) - \beta f \circ g(x,\gamma) \\
\leq \;& (1-\beta)u(x',\gamma') - (1-\beta)u(x,\gamma') + \beta f \circ g(x',\gamma') - \beta f \circ g(x,\gamma') \qquad (60) \\
\leq \;& (1-\beta)(x'-x)\left(\sqrt{\gamma'} - \gamma'/2\right) + \beta k \left(g(x',\gamma') - g(x,\gamma')\right) \qquad (61) \\
\leq \;& (1-\beta)(x'-x)/2 + \beta k \frac{\partial g}{\partial x}(x'',\gamma')\,(x'-x). \qquad (62)
\end{aligned}$$

where $x'' \in [x, x']$. In the above, inequality (60) follows because γ maximizes (32) at x, inequality (61) follows because of (33), and inequality (62) follows because $\sqrt{\gamma} - \gamma/2 < 1/2$ together with a first-order Taylor expansion of $g(\,\cdot\,,\gamma')$. Hence, it follows from (62) that a sufficient condition for the Lipschitz condition (33) to hold for Tf is

$$\begin{aligned}
& (1-\beta)/2 + \beta k \frac{\partial g}{\partial x}(x'',\gamma') \leq k \\
\Leftrightarrow \;& k\Big(1 - \beta \frac{\partial g}{\partial x}(x'',\gamma')\Big) \geq (1-\beta)/2. \qquad (63)
\end{aligned}$$

Now we also have

$$\begin{aligned}
\frac{\partial g}{\partial x}(x'',\gamma') &= \frac{\alpha^2 + 2x''(1-x'')\left(\alpha\gamma(2-x'') + \gamma^2 x''(1-x'')\right)}{[\alpha + \gamma' x''(1-x''^2)]^2} \\
&\leq 1 + 2x''(1-x'')\left[\gamma'/\alpha(2-x'') + \gamma'^2/\alpha^2 x''(1-x'')\right] \\
&\leq 1 + 2x''(1-x'')\left[\phi(x',k)/\alpha(2-x'') + \phi(x',k)^2/\alpha x'(1-x')\right] \\
&\equiv \psi(x',x'',k),
\end{aligned}$$

where $\phi(x,k)$ is the function defined in equation (59). Therefore, a sufficient condition for (63) to hold is that

$$k\left(1 - \beta\psi(x',x'',k)\right) \geq (1-\beta)/2 \qquad (64)$$

for all $(x',x'') \in [0,\eta]^2$. Now let $k = 1$. Then, because $\psi(0,0,k) = 1$, (64) is satisfied when $x' = x'' = 0$ with a strict inequality. By continuity, there exists $\eta > 0$ such that, (64) holds for all $(x',x'') \in [0,\eta]^2$. This completes the proof.

A.9 Proof of proposition 9

[i] $\Gamma(x) \subseteq (1, \infty)$ for $x \in (0,1)$. Take some $x \in [0,1]$ and consider the function

$$w(\gamma) \equiv (1-\beta)u(x,\gamma) + \beta W \circ g(x,\gamma). \qquad (65)$$

Note that $W(\,\cdot\,)$ and $g(x,\,\cdot\,)$ are both strictly increasing functions. Since $u(x,\,\cdot\,)$ is increasing for $\gamma \in [0,1)$, it follows that $w(\,\cdot\,)$ is strictly increasing in $\gamma \in [0,1)$, implying that $\Gamma(x) \subseteq [1,\infty)$.

Then, for $\gamma > 1$ and close enough to 1, we have

$$\begin{aligned}
w(\gamma) - w(1) &= (1-\beta)u(x,\gamma) - (1-\beta)u(x,1) + \beta W \circ g(x,\gamma) - \beta W \circ g(x,1) \\
&\geq (1-\beta)u(x,\gamma) - (1-\beta)u(x,1) + \beta m \left(g(x,\gamma) - g(x,1)\right),
\end{aligned}$$

for some $m > 0$ given by the second part of Proposition 7. Dividing both sides of the equation by $\gamma - 1$, and letting γ go to 1 shows that

$$\lim_{\gamma \to 1} \frac{w(\gamma) - w(1)}{\gamma - 1} = 0 + \beta m \frac{\partial g}{\partial \gamma}(x,1) > 0,$$

which shows that $\gamma = 1$ cannot maximize $w(\gamma)$. Therefore, $\Gamma(x) \subset (1,\infty)$.

[ii]. Convergence of the state We already know that, when $\gamma = 1$ for all t, the state x_t^1 converges towards one as t goes to infinity. Suppose that the planner chooses an optimal control $\gamma_t^* \in \Gamma(x_t^*)$. Because $\gamma_t^* \in [0,1]$ and because g is decreasing in γ, we have $x_{t+1}^* \geq g(x_t^*, 1)$. By induction, this implies that $1 \geq x_t^* \geq x_t^1$. Therefore, x_t^* goes to 1 as time goes to infinity.

[iii]. Convergence of the control Suppose $x_1 = 1$. Then, $x_{t+1} = x_t = 1$ for all time: at each time, the planner has to solve a static optimization problem whose unique solution is $\gamma_t = 1$. This shows in turns that $\Gamma(1) = \{1\}$. Now suppose that the planner chooses an optimal control $\gamma_t^* \in \Gamma(x_t^*)$. Because $\gamma_t^* \in [0,1]$ it has a converging subsequence $\gamma_{t_k}^*$. Since $x_{t_k}^*$ goes to one, and since (by the Theorem of the Maximum) the correspondence Γ is upper hemicontinuous, we know that $\lim \rho_{t_k}^* \in \Phi(1) = \{1\}$. Hence, the only accumulation point of the sequence γ_t^* is 1. Therefore, the sequence γ_t^* goes to one.

[iv]. Non-monotonicity Equation (59) shows that, for all $x \in (0, \eta)$, $\gamma \in \Gamma(x)$ implies that $\gamma \leq \phi(x,k)$. Since $\phi(0,k) = 1$ and $\gamma \geq 1$, this implies that $\Gamma(x)$ goes to 1 as x goes to 1. Namely, for all sequence $x_k \to 1$ and all $\gamma_k \in \Gamma(x_k)$, we have $\gamma_k \to 1$.

B Learning without Observational Noise

This appendix solves for information diffusion and welfare when agents can observe each others' action without the exogenous informational noise. Equation (20) shows that, when $\alpha = \sigma_\epsilon^2/\sigma_0^2 = 0$,

$$1 - x_{t+1} = \frac{1 - x_t}{1 + x_t},$$

for $t \in \{1, 2, \ldots\}$. Since $\sigma_t^2 = \sigma_0^2(1 - x_t)$, we obtain

$$\sigma_{t+1}^2 = \frac{\sigma_t^2}{2 - \sigma_t^2/\sigma_0^2}.$$

Now, with the change of variable $\pi_t = 1/\sigma_t^2$, we obtain that $\pi_{t+1} = 2\pi_t - \pi_0$. Therefore, $\pi_{t+1} = \pi_0 + 2^t(\pi_1 - \pi_0)$. Plugging back the initial condition that $\pi_1 = \pi_0 + 1/s_1^2$, we obtain that

$$\sigma_t^2 = \frac{\sigma_0^2 s_1^2}{s_1^2 + 2^{t-1}\sigma_0^2}, \tag{66}$$

for $t \in \{1, 2, \ldots\}$. Taken together, these calculations imply

Lemma 4. *If $\sigma_\varepsilon^2 = 0$, then, asymptotically $\sigma_t^2 \sim s_1^2/2^t$. Let's consider $\sigma_0(1)^2 < \sigma_0(2)^2$ and denote the subsequent sequence of variances be $\{\sigma_t^2(k)\}_{t=1}^\infty$. Then, for all $t \in \{1, 2, \ldots\}$, we find that $\sigma_t^2(1) \leq \sigma_t^2(2)$.*

One can easily verify that, without informational noise, the time path of the distribution of beliefs also has a S-shaped mean and a hump-shaped variance. However, two other features of the learning dynamics are sharply different than with informational noise. First, analysts learn much faster, at the geometric rate $1/2^t$ instead of the linear rate $1/t$. Second, the impact of public information is unambiguous: it never slows down the diffusion of private information and always improves welfare.

C Welfare Cost in Discrete Time

This appendix provides numerical calculations suggesting that the properties of the continuous-time model of section 3.4 also holds in the discrete-time model that we study in the rest of the paper. The parameters are chosen in the spirit of our continuous-time limit: namely, we choose a large observational noise $\sigma_\epsilon^2 = 100$ and a discount factor $\beta = 0.995$ that is close to one. We also set $s_1^2 = 199$ as in our previous numerical examples. Figure 7 shows the discrete-time social welfare as a function of σ_0^2. Its shape turns out to be similar to that of the continuous-time social welfare of Figure 3. This suggests in particular that the non-monotonicity property of Theorem 1 also holds in discrete time.

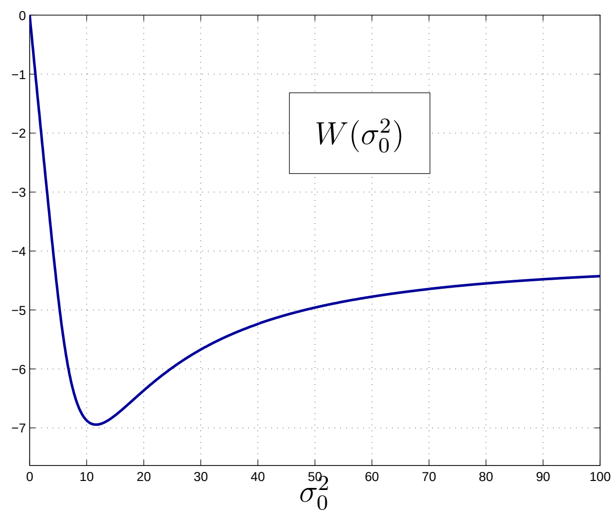

Figure 7: Welfare in the discrete-time model as a function of σ_0^2.

References

Angeletos, G.-M., and A. Pavan (2005): "Efficient Use of Information and Welfare Analysis in Economies with Complementarities and Asymmetric Information," . 5

Araujo, L., and B. Camargo (2006): "Information, Learning, and the Stability of Fiat Money," *Journal of Monetary Economics, Forthcoming*. 6

Araujo, L., and A. Shevchenko (2006): "Price Dispersion, Information and Learning," *Journal of Monetary Economics, Forthcoming*. 6

Bala, V., and S. Goyal (1998): "Learning from Neighbours," *Review of Economic Studies*, 65, 595–621. 5

Banerjee, A., and D. Fudenberg (2004): "Word-of-mouth learning," *Game and Economic Behavior*, 46, 1–22. 5

Benveniste, L. M., and J. A. Scheinkman (1979): "On the Differentiability of the Value Function in Dynamic Models of Economics," *Econometrica*, 47, 727–732. 27

Blouin, M. R., and R. Serrano (2001): "A Decentralized Market with Common Values Uncertainty: Non-Steady States," *Review of Economic Studies*, 68, 323–346. 6

Chamley, C. (2004): *Rational Herds*. Cambridge University Press, Cambridge. 5

Chamley, C., and D. Gale (1994): "Information Revelation and Strategic Delays in a Model of Investment," *Econometrica*, 62, 1065–1085. 5

DeMarzo, P., D. Vayanos, and J. H. Zwiebel (2003): "Persuasion Bias, Social Influence, and Uni-Dimensional Opinions," *Quaterly Journal of Economics*, 118, 909–968. 5

Duffie, D., and G. Manso (2006): "Information Diffusion in Large Population," . 6

Edwards, A. K., L. Harris, and M. S. Piwowar (2004): "Corporate Bond Market: Transparency and Transaction Costs," Working Paper, University of Southern California. 5

Gale, D., and S. Kariv (2003): "Bayesian Learning in Social Networks," Working Paper, New York University. 5

Green, E. J. (1991): "Eliciting Traders' Knowledge in "Frictionless" Asset Markets, Working Paper," . 6

Hellwig, C. (2005): "Heterogeneous Information and the Welfare Effects of Public Information Disclosures," . 5

Judd, K. L. (1999): *Numerical Methods in Economics*. MIT Press, Boston. 28

Katzman, B., J. Kennan, and N. Wallace (2003): "Output and Price Level Effects of Monetary Uncertainty in a Matching Model," *Journal of Economic Theory*, 108, 217–255. 6

Kiyotaki, N., and R. Wright (1989): "On Money as a Medium of Exchange," *Journal of Political Economy*, 97, 927–954. 6

Lucas, R. E. J. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4, 103–124. 4

Luenberger, D. G. (1969): *Optimization by Vector Space Methods*. Wiley, John and Sons, New York. 10, 31

Morris, S., and H. S. Shin (2002): "The Social Value of Public Information," *American Economic Review*, 92, 1521–1534. 4, 5, 7, 18, 20

Phelps, E. S. (1969): *Microeconomic Fundations of Employment and Inflation Theory*. Wiley, John and Sons, New York. 4

Smith, L., and P. N. Sorensen (2005): "Rational Social Learning by Random Sampling," Working Paper, University of Michigan. 5

Stokey, N. L., and R. E. Lucas (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge. 25, 36

Taylor, A. E., and R. W. Mann (1983): *Advanced Calculus*. Wiley, John and Sons, New-York. 35

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

Vives, X. (1993): "How Fast do Rational Agents Learn?," *Review of Economic Studies*, 60, 329–347. 5

——— (1997): "Learning from others: a Welfare Analysis," *Games and Economic Behavior*, 20, 177–200. 5, 17, 22

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

Wolinsky, A. (1990): "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, 58, 1–23. 6