Addenda to
"Learning from Private and Public Observation of Others' Action"

Manuel Amador and Pierre-Olivier Weill

August 24, 2011

The following addenda contain extensions and extra results.

In Addendum I we show why it is without loss of generality to normalize some parameters of the model.

In Addendum II we show that ODE (15) has a unique solution.

In Addendum III we study the impact of releasing public information in Vives (1993) original discrete time model.

In Addendum IV we show that an optimal affine control is always a convex combination of the public and private forecasts.

In Addendum V we detail the dynamic programming result leading to the Hamilton Jacobi Bellman equation (20).

In Addendum VI we show that the results in the paper extend to the case of a non-degenerate common prior.

In Addendum VII we provide a the dynamic of the belief distribution in the population.

In Addendum VIII we provide some natural applications of our abstract learning model.

I Normalization

The main ODE of the paper is:

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

where $\alpha = P_0 - \frac{P_\varepsilon}{p_\varepsilon} p_0$ and $\beta = 1 + P_\varepsilon / p_\varepsilon$. Let us denote the solution of this ODE by $\phi(t; p_0, P_0, p_\varepsilon, P_\varepsilon)$, to make the dependence on initial conditions and exogenous parameter explicit. We now propose two normalizations. First, we have:

Lemma 13. *For all* $(p_0, P_0, p_\varepsilon, P_\varepsilon)$ *and all* $k > 0$*:*

$$\phi(t; p_0, P_0, p_\varepsilon, P_\varepsilon) = \phi\left(kt; p_0, P_0, \tfrac{p_\varepsilon}{k}, \tfrac{P_\varepsilon}{k}\right).$$

This means that scaling the level of noise precision up and down by k is equivalent to scaling time by the same factor. To see this, let $\tilde{p}_t \equiv \phi\left(kt; p_0, P_0, \frac{p_\varepsilon}{k}, \frac{P_\varepsilon}{k}\right)$ and note that scaling of p_ε and P_ε leaves α and β the same. So:

$$\begin{aligned} \dot{\tilde{p}}_t &= \frac{d}{dt}\phi\left(kt; p_0, P_0, \tfrac{p_\varepsilon}{k}, \tfrac{P_\varepsilon}{k}\right) = k\frac{p_\varepsilon}{k}\left(\frac{\tilde{p}_t}{\alpha + \beta\tilde{p}_t}\right)^2 \\ &= p_\varepsilon \left(\frac{\tilde{p}_t}{\alpha + \beta\tilde{p}_t}\right)^2. \end{aligned}$$

So p_t and $\tilde{p}_t$ solve the same ODE and have the same initial condition: they must coincide at all times. QED

Another normalization is given by:

Lemma 14. *For all* $(p_0, P_0, p_\varepsilon, P_\varepsilon)$ *and all* $k > 0$*:*

$$\phi\left(t; kp_0, kP_0, kp_\varepsilon, kP_\varepsilon\right) = k\phi(t; p_0, P_0, p_\varepsilon, P_\varepsilon). \tag{I.1}$$

This says that scaling all precisions, $(p_0, P_0, p_\varepsilon, P_\varepsilon)$, up and down by the same factor is equivalent to scaling up the path of private information. To see this, let

$\tilde{p}_t \equiv \frac{1}{k}\phi\left(t; 1, kp_0, kP_0, kp_\varepsilon, kP_\varepsilon\right)$, $\tilde{\alpha} = k\alpha$ and $\tilde{\beta} = \beta$. We have

$$\dot{\tilde{p}}_t = \frac{1}{k}\frac{d}{dt}\phi\left(t; kp_0, kP_0, kp_\varepsilon, kP_\varepsilon\right) = \frac{1}{k}kp_\varepsilon\left(\frac{k\tilde{p}_t}{\tilde{\alpha} + \beta k\tilde{p}_t}\right)^2 = p_\varepsilon\left(\frac{k\tilde{p}_t}{k\alpha + \beta k\tilde{p}_t}\right)^2$$
$$= p_\varepsilon\left(\frac{\tilde{p}_t}{\alpha + \beta\tilde{p}_t}\right)^2$$

and we are done. QED

Taken together, the two lemmas of this addendum imply that:

$$\phi(t; p_0, P_0, p_\varepsilon, P_\varepsilon) = p_0\phi\left(\frac{P_\varepsilon}{p_0}t; 1, \frac{P_0}{p_0}, \frac{p_\varepsilon}{P_\varepsilon}, 1\right).$$

II A note on the ODE of Section 4.1.1

Lemma 15. *Suppose that Π_t and π_t are increasing and piecewise continuously differentiable. Then, starting from any initial condition P_0, ODE (15) has a unique solution.*

Because of piecewise continuous differentiability, there exists a strictly increasing sequence $t_0 < t_2 < \ldots < t_K$ where $K \leq \infty$ and $t_K = \infty$, such that the functions Π_t and π_t are continuously differentiable on each interval $[t_k, t_{k+1})$ for all $k < K$. Starting from any initial condition P_{t_k}, we can construct a path of public precision by solving the ODE:

$$\dot{P}_t = P_\varepsilon \left(\frac{\pi_t}{P_t + \pi_t} \right)^2 + \dot{\Pi}_t$$

on every interval $[t_k, t_{k+1})$. Moreover, it is easy to see that the maximal solution of the ODE is defined on the entire interval. Otherwise, suppose the maximal solution was defined over $[t_k, T)$, for $T < t_{k+1}$. Since P_t is increasing, it has a limit as $t \to T$. Moreover,

$$\dot{P}_t \leq P_\varepsilon + \dot{\Pi}_t \Rightarrow P_t \leq P_\varepsilon (T - t_k) + \Pi_T - \Pi_{t_k},$$

so the limit is finite. But then, one can extend the solution of the ODE further starting at time T from $\lim_{t \to T} P_t$. This is a contradiction since we assumed that P_t were a maximal condition. Having constructed a solution over $[t_k, t_{k+1})$, we let

$$P_{t_{k+1}} = \lim_{t \to t_{k+1}^-} P_t + \Pi_{t_{k+1}^+} - \Pi_{t_k}, \tag{II.1}$$

and proceeding by induction we obtain a path of public precision which is well defined over $[0, \infty)$. Moreover, the standard uniqueness result applies: i.e. two solutions starting from the same initial condition must coincide at all times. Indeed, consider two solutions starting from the same initial condition and suppose that they start differing after some time $s > 0$. Then if $s \neq t_k$, by continuity we find that the two solutions coincide at s. If $s = t_k$, then by continuity the two solutions coincide at t_k^- and, because of (II.1) at t_k as well. Thus, the standard local uniqueness result imply that the two solution coincides just after s, which is a contradiction.

III Public information releases in Vives' discrete time model

In Vives (1993), public precision evolves according to:

$$P_{t+1} = P_t + P_\varepsilon \left(\frac{p_0}{p_0 + P_t} \right)^2 = F(P_t). \tag{III.1}$$

The graph of $F(P)$ is shown in Figure 2. Since it lies strictly above the 45 degree line, it follows that P_t is increasing and goes to infinity as t goes to infinity. Another property of that is key for explaining the effect of public information is that $F(P)$ is decreasing over the interval $[0, P_m]$, where

$$P_m = \max\{p_0 \left(\sqrt{2P_\varepsilon p_0} - 1 \right), 0\},$$

and increasing for $P > P_m$.

If $P_0 \geq P_m$, then more public information at the beginning increases knowledge in times that follow. Suppose it did not: that is, consider the first time t such that a precision sequence with more initial public information, P'_t, moves below a sequence with less initial public info, P_t. By construction at time $t-1$,

$$P'_{t-1} > P_{t-1} > P_0 \geq P_m.$$

Since $F(P)$ is increasing for $P > P_m$, it follows from these inequalities that $P'_t = F(P'_{t-1}) > P_t = F(P_{t-1})$, which is a contradiction.

If $P_0 < P_m$, then since $F(P)$ is decreasing at P_0 it follows that more initial public information (i.e. a slightly higher P_0) reduces P_1. Since $P_1 = F(P_0) \geq F(P_m) > P_m$, then the reasoning of the previous paragraph shows that the precision sequence is also lower at all subsequent times.

Lastly, note that, since $P_1 > P_m$ for any P_0, it follows that if the planner has one piece of public information to release either at time $t = 0$ or at time $t = 1$, he will always want to release it, at least at time $t = 1$.

Proposition 3. *Suppose public precision evolve according to the difference equation*

(III.1). Then, a marginal increase in time t public information, P_t, increases precision at all times $s \geq t$ if and only if (i) $t = 0$ and $P_0 \geq P_m$, or (ii) $t \geq 1$.

Note that if P_ε is small enough, then $P_m = 0$ and increasing public information at time zero always increase precision.

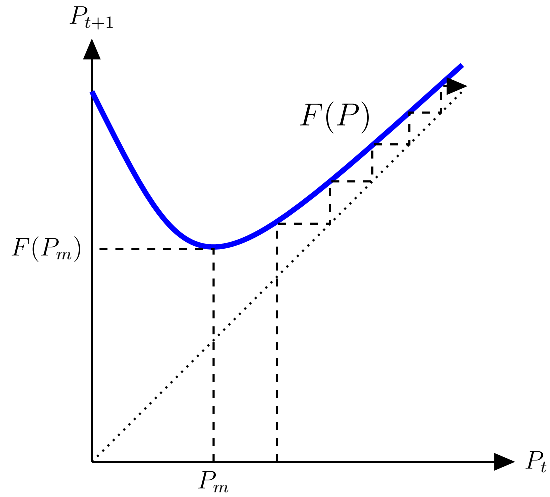

Figure 2: Precision dynamics in Vives.

IV General affine controls

Lemma 16. *In the planner's problem, an optimal affine control is of the form:*

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

Indeed, suppose that, at time t, the planner uses a control of the form:

$$a_{it} = \kappa_t + \Gamma_t \hat{X}_t + \gamma_t \hat{x}_{it} \tag{IV.1}$$

This implies that:

$$a_{it} - x = \kappa_t + (\Gamma_t + \gamma_t - 1)\hat{X}_t + (1-\gamma_t)\left(\hat{X}_t - x\right) + \gamma_t(\hat{x}_{it} - 1).$$

But all terms on the right-hand side are uncorrelated with one another each others. First, the forecast errors $\hat{X}_t - x$ and $\hat{x}_{it} - x$ are uncorrelated because they are based on independent information. The public forecast is uncorrelated from the public forecast error by construction. Lastly, the public forecast is independent from the private forecast error because private forecast errors are idiosyncratic. Thus:

$$\mathbb{E}\left[(a_{it} - x)^2\right] = \mathbb{E}\left[\left(\kappa_t + (\Gamma_t + \gamma_t - 1)\hat{X}_t\right)^2\right] + \frac{\gamma_t^2}{p_t} + \frac{(1-\gamma_t)^2}{\alpha + (\beta - 1)p_t}.$$

Since κ_t and Γ_t do not affect the law of motion of precision, it is clearly optimal for the planner to minimize the first term with respect to κ_t and Γ_t, i.e. to set $\kappa_t = 0$ and $\Gamma_t = 1 - \gamma_t$. QED

V The results leading to Lemma 9

Proposition 2.8, page 104 in Bardi and Capuzzo-Dolcetta (1997) states that the value function is a viscosity solution of an appropriate HJB equation. Note that Lemma 8 implies that $V(p)$ is locally Lipschitz. Thus, we can apply Theorem 2.40, in page 128, together with Remark 2.43, in page 131, to show that value function solves the following HJB equation,[1, 2]

$$\lambda V(p) = \sup_{(u,q)\in\Phi(p)} \left\{ u + qV'(p^+) \right\}, \tag{V.1}$$

where for every p

$$\Phi(p) \equiv \Big\{ (u,q) \Big| u = \theta u(p,\gamma) + (1-\theta)u(p,\gamma');\ q = \theta p_\varepsilon \gamma^2 + (1-\theta)p_\varepsilon \gamma'^2; \\ \text{for some } (\gamma,\gamma',\theta) \in [0,1]^3 \Big\},$$

and where

$$u(p,\gamma) = -\lambda \left(\frac{\gamma^2}{p} + \frac{(1-\gamma)^2}{\alpha + (\beta-1)p} \right).$$

Note that, for all $(u,q) \in \Phi(p)$, there is some $(\gamma,\gamma',\theta) \in [0,1]^3$ such that:

$$\begin{aligned} u + qV'(p^+) &= \theta u(p\gamma) + (1-\theta)u(p\gamma') + (\theta\gamma^2 + (1-\theta)\gamma'^2)p_\varepsilon V'(p^+) \\ &\leq u\left(p, \sqrt{\theta\gamma^2 + (1-\theta)\gamma'^2}\right) + (\theta\gamma^2 + (1-\theta)\gamma'^2)p_\varepsilon V'(p^+) \\ &= u\left(p, \gamma''\right) + \gamma''^2 p_\varepsilon V'(p^+). \end{aligned} \tag{V.2}$$

where inequality (V.2) follows from the concavity of $g \mapsto u(p, \sqrt{g})$ and

$$\gamma'' \equiv \sqrt{\theta\gamma^2 + (1-\theta)\gamma'^2} \in [0,1].$$

[1]The infinite-horizon optimal control problem of Bardi and Capuzzo-Dolcetta is formulated with the state space $\mathbb{R}^N$ and requires that the utility flow function to be bounded. This is different from our problem in which the state space is $\mathbb{R}^2_+$ and the utility flow unbounded as either p or P go to zero. We can nevertheless apply their results by using the change of variable $p = \max\{x, p_0\}$, where $x \in \mathbb{R}^2$.

[2]The derivative that appears in Bardi and Capuzzo-Dolcetta is a one-sided directional derivative, but in our one dimensional state space it coincides with the right-derivative shown in the appendix.

And thus we can reduce the search for an optimal control to a value of $\gamma \in [0,1]$, which delivers the HJB equation of the Lemma.

For the existence of the optimal control we use Theorem 2.61 part (ii) in page 142 of Bardi and Capuzzo-Dolcetta (1997) together with Remark 2.62 in page 142, which imply that

$$\gamma_t^* = \gamma^*(p_t) \equiv \arg\max_{\gamma \in [0,1]} \left\{ u(p_t^*, \gamma) + \gamma^2 V'((p_t^*)^+) \right\} \tag{V.3}$$

where $p_t^* = p + \int_0^t (\gamma_t^*)^2 dt$, is an optimal control.

VI Non-degenerate Common Prior

In this Addendum we provide the formulas for agents' belief when the prior is non degenerate, i.e. with a strictly positive precision $\bar{P}$. We assume that the precision of the initial public signal is $P_0 - \bar{P}$: this implies that, given the common prior and after observing the initial public signal, agents' posterior precision is P_0, just as in the paper. Then, one easily verify that all the results on precision and welfare go through identically. There is one small difference, in Lemma 1 concerning the dynamic of the public forecast:

Lemma 17. *With a non-degenerate prior precision $\bar{P}$, all formulas characterizing the dynamic of public precision, P_t, private precision, p_t, and private forecast, $\hat{x}_{it}$, stay the same. The only difference is that the stochastic integral for the public forecast becomes:*

$$\hat{X}_t = \left(1 - \frac{\bar{P}}{P_t}\right) x + \frac{1}{P_t}\left[\sqrt{P_0 - \bar{P}} W_0 + \int_0^t \sqrt{P_\varepsilon} \tfrac{p_t}{P_t + p_t} dW_u\right] \tag{VI.1}$$

Indeed, going through the exact same step as in the proof of Lemma 1, we arrive at (31). But then we have to plug in a different expression for $\hat{X}_0$, the posterior public forecast given the public signal and the common prior. Indeed, when $\bar{P} > 0$, we have:

$$\hat{X}_0 = \left(1 - \frac{\bar{P}}{P_0}\right)\left(x + \frac{W_0}{\sqrt{P_0 - \bar{P}}}\right).$$

The result then follows. Note that, when $\bar{P} = 0$, we obtain the same expression as in Lemma 1.

VII Cross-sectional Beliefs

VII.1 A Preview

In what follows, we characterize the dynamic behavior of the distribution of agents' beliefs in closed form. We show that, as long as there is a sufficiently active private information channel, the average belief converges along an S-shaped curve, as long as the initial information is sufficiently dispersed and the private channel is active.[3] Agents can become more and more confident about the information they gathered privately, taking actions that are increasingly sensitive to this private information. This makes the private and public learning channels more informative, generating more precise signals, faster learning and in some cases, a convex learning curve at the beginning. On the other hand, convergence to the truth implies that, in the limit, the average belief is concave. Although the average belief follows an S-shaped curve, we show that the cross-sectional dispersion of agents' beliefs will converge to zero along a hump-shaped curve. Indeed, when agents learn privately, they learn independently, and their learning histories are increasingly heterogeneous: hence, the dispersion of beliefs might increase early on. However, this dispersion eventually must converge to zero as agents learn the truth. Importantly, we show that when agents learn only from a public channel and the private channel is shut down, then the S and the hump disappear: the average belief converges to the truth along a concave curve, and the dispersion of beliefs converges to zero along a decreasing curve.

VII.2 The Results

Consider the model where the prior is non degenerate, $\bar{P} > 0$. In the present normal-quadratic framework, the distribution of beliefs in the population is also normal and

[3] The S-shaped pattern has been documented in a number of empirical studies of social learning. See Chamley (2004, chap. 9) and also Jovanovic and Nyarko (1995)). See also the subsequent learning models of Fernandez (2007) and Fogli and Veldkamp (forthcoming) to explain the S shape in women labor force participation.

can be characterized in closed form. To see this, recall that the action of agent i is

$$a_{it} = \frac{p_t}{p_t + P_t}\hat{x}_{it} + \frac{P_t}{p_t + P_t}\hat{X}_t \ , \tag{VII.1}$$

and that an agent's action is the same as her belief. Then, equations (VI.1) and (29) imply that cross-sectional beliefs are normally distributed, and therefore entirely characterized by their average and dispersion. Taking expectations conditional on x, it follows that

$$E\left[a_{it} \,|\, x\right] = E\left[A_t | x\right] = x\left(1 - \frac{\bar{P}}{p_t + P_t}\right) \ , \tag{VII.2}$$

and so the rate at which the expected average belief converges to the truth is the same as the rate at which the variance of an agent's belief, $1/(p_t + P_t)$, approaches zero.

The cross-sectional dispersion of beliefs, or actions, also can be computed. Given that the second term in the right hand side of equation (VII.1) is common across all agents, the cross-sectional dispersion is simply the variance of the first term,

$$\theta_t \equiv E[(a_{it} - A_t)^2] = \frac{p_t}{(p_t + P_t)^2} \ ,$$

where the cross-sectional variance of $\hat{x}_{it}$ is $1/p_t$. Note that the dynamics of the average belief, or action, are given by

$$\begin{aligned} \frac{dE\left[a_{it}|x\right]}{dt} &= x\bar{P}\frac{\dot{p}_t + \dot{P}_t}{(p_t + P_t)^2} = x\bar{P}(p_\varepsilon + P_\varepsilon)\left(\frac{p_t}{(p_t + P_t)^2}\right)^2 \\ &= x\bar{P}(p_\varepsilon + P_\varepsilon)\theta_t^2. \end{aligned} \tag{VII.3}$$

Hence, the expected action, $E[a_{it}|x]$, monotonically approaches x and its changes are proportional to the square of the cross-sectional dispersion, θ_t. Similarly, the dynamics of the cross-sectional dispersion are given by

$$\frac{d\theta_t}{dt} = \frac{\dot{p}_t}{(p_t + P_t)^2} - \frac{2p_t\left(\dot{p}_t + \dot{P}_t\right)}{\left(p_t + P_t\right)^2} = \frac{\dot{p}_t\,(p_\varepsilon + P_\varepsilon)}{(p_t + P_t)^2}\left(\frac{p_\varepsilon}{p_\varepsilon + P_\varepsilon} - 2\frac{p_t}{p_t + P_t}\right), \tag{VII.4}$$

where we used that $\dot{P}_t = P_\varepsilon/p_\varepsilon \dot{p}_t$. Recall that $p_t/(p_t + P_t)$ converges to $p_\varepsilon/(p_\varepsilon + P_\varepsilon)$

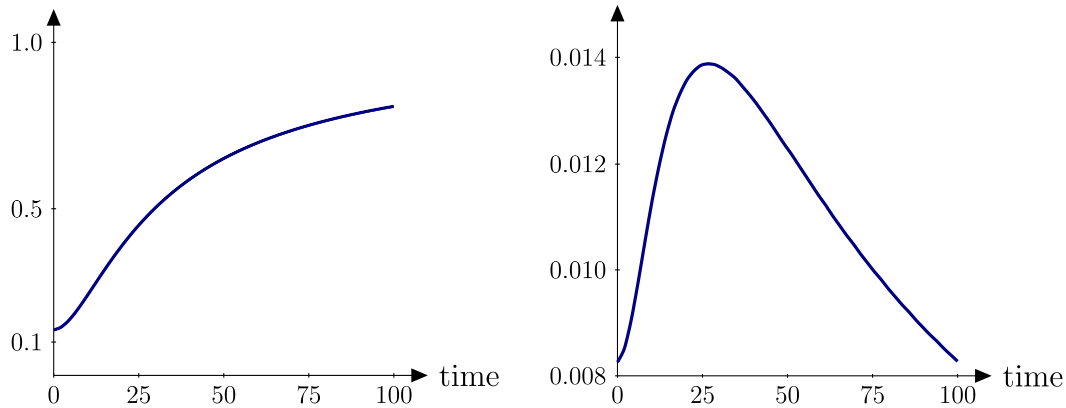

(a) Mean Belief when $x = 1$.

(b) Dispersion.

Figure 3: This is a particular example of the dynamics of cross-sectional beliefs when the path of the mean is S-shaped and the dispersion of belief is hump-shaped. We choose $\bar{P} = 9.5$, $P_0 = 10$, $p_0 = p_\varepsilon = P_\varepsilon = 1$.

monotonically. Together with the fact that $\dot{p}_t$ is positive at each time, this implies that the cross-sectional dispersion is eventually decreasing. It follows also that the cross-sectional dispersion will have a hump shape as long as it is increasing at time zero, that is as long as $p_\varepsilon/(p_\varepsilon + P_\varepsilon) > 2p_0/(p_0 + P_0)$.

Finally, note that if the cross-sectional dispersion, θ_t, is hump-shaped, then it follows from equation (VII.3) that the path of $E[a_{it}|x]$ is S-shaped. Moreover, the inflexion time of the S corresponds to the point of highest dispersion of actions. We have just proved the following:

Proposition 4 (S-shaped diffusion)**.** *The path of $E[a_{it}|x]$ monotonically approaches x as time goes to infinity. If $2p_0/(p_0 + P_0) < p_\varepsilon/(p_\varepsilon + P_\varepsilon)$ then there exists a $t_0 > 0$ such that $|dE[a_{it}|x]/dt|$ is increasing for all $t < t_0$ and decreasing for all $t > t_0$. Otherwise, $|dE[a_{it}|x]/dt|$ is decreasing for all t.*

As shown in the proposition and illustrated in Figure 3, the model is able to generate an S-shaped learning curve. The convex part of the S arises because of an information-snowballing effect as long as the initial private information is sufficiently dispersed, i.e. its precision is sufficiently small. In that case, as agents learn privately, they become increasingly confident about their private information. As a result, the weight they assign to their private information becomes larger and larger and both the private and the public signal become more and more informative: learning is faster and faster at the beginning. Eventually, as agents learn the truth, learning must slow down and the learning curve becomes concave.

The hump shape of the dispersion of belief is more standard, and arises because of the private learning channel. Initially, agents beliefs are concentrated close to their common prior. As time goes by, agents learn privately, independently of each other. Some become optimistic about the state, and others pessimistic, which means that the dispersion of beliefs increases at the beginning. However, this dispersion must converge to zero as agents asymptotically learn the truth.

VIII Implications and Extensions

In this Addendum we apply and generalizes our results. In a first extension, we relax our baseline-model assumption that agents do not learn from their own experience, in that they cannot observe their payoffs until the state is revealed. We show that relaxing this assumption does not significantly alter our results. Next, in a second extension, we provide an example of public signals being generated by the observation of market prices. We assume that firms face quadratic adjustment costs in investment and buy capital goods in a competitive market. After some random time to build a capital stock, these goods become inputs in a production technology with a common but initially unknown productivity. As anticipated by our baseline model, the price of capital generates a public signal centered around the average action, which in this case is the average investment. Different from our baseline model, though, this example features a payoff externality: the price of the capital affects firms' profits, and hence depends on the aggregate investment. We show that even in the presence of this externality, the learning dynamics and the welfare results remain the same as in the baseline model. The extensions demonstrate the tractability of our continuous-time set up, and suggest that the results derived are robust to more explicit economic environments.

VIII.1 Technology Diffusion

In many interesting economic situations, it is reasonable to assume that firms and households learn not only from observing others' actions but also from their own experience. Such cases arise naturally, for example, if agents observe their realized payoffs which are noisy signals of the underlying state. One might wonder whether the results obtained in the baseline model generalize to this type of environment. In this subsection, we modify our baseline model to allow for the possibility of learning from experience: a new technology of uncertain productivity has arrived and agents learn about it from their own use, and as before, from observing how others use it.

Suppose that agent i owns a technology that converts labor input, denoted by a_{it}, into units of a final good of price 1. The agent incurs a quadratic loss, $a_{it}^2/2$, from supplying labor. The productivity of this technology is assumed to be the sum of a time-varying and random idiosyncratic component, plus an unknown but constant component. We assume that this constant component, denoted by x, is common across the technologies of all agents in the economy. The accumulated output up to time τ is given by,

$$\int_0^\tau \left(x dt + \frac{1}{\tilde{p}_\varepsilon} d\tilde{\omega}_{it} \right) a_{it} \ ,$$

where the term in brackets is the sum of x and the idiosyncratic productivity process. We assume that $\tilde{\omega}_{it}$ is a standard Wiener process.

For simplicity, we remove public learning from the model: $P_\varepsilon = 0$. However, we still assume that agents privately can observe a signal centered around the average labor supply, with precision p_ε. Unlike before, we now let the agents also observe their current output processes: agents are learning from others, but also from their own use of this technology.

Agents optimally choose their labor decision, $a_{it} = E[x|\mathcal{G}_{it}]$. A linear equilibrium is obtained by decomposing the beliefs at any time into a public and a private forecast. Given that there is no public learning, the precision of the public forecast is a constant, P_0. The precision of the private forecast, p_t, follows the ODE

$$dp_t = \left(\frac{p_t}{p_t + P_0} \right)^2 p_\varepsilon dt + \tilde{p}_\varepsilon dt \ , \tag{VIII.1}$$

where the first term is the same as before, and the second term captures the knowledge that arises from observing your own output. The following proposition characterizes the dynamics of precision:

Proposition 5. *Define*

$$\tilde{H}(p) \equiv \frac{P_0 \sqrt{p_\varepsilon} (p_\varepsilon - \tilde{p}_\varepsilon)}{\sqrt{\tilde{p}_\varepsilon} (\tilde{p}_\varepsilon + p_\varepsilon)^2} \arctan \left(\frac{P_0 \tilde{p}_\varepsilon + (\tilde{p}_\varepsilon + p_\varepsilon) p}{P_0 \sqrt{p_\varepsilon}} \right) + \frac{p_\varepsilon \sqrt{\tilde{p}_\varepsilon} P_0}{(\tilde{p}_\varepsilon + p_\varepsilon)^2} \log \left(P_0^2 \tilde{p}_\varepsilon + 2 P_0 \tilde{p}_\varepsilon p + (\tilde{p}_\varepsilon + p_\varepsilon) p^2 \right) + \frac{p}{p_\varepsilon + \tilde{p}_\varepsilon} \ ,$$

then, the solution to equation (VIII.1) is implicitly given by $\tilde{H}(p_t) - \tilde{H}(p_0) = t$, and *the asymptotic expansion of p_t is*

$$p_t = (\tilde{p}_\varepsilon + p_\varepsilon)t - 2\frac{P_0 p_\varepsilon}{p_\varepsilon + \tilde{p}_\varepsilon}\log(t) + R_t,$$

where R_t is some bounded function.

Note that if $\tilde{p}_\varepsilon = 0$, this collapses to our baseline asymptotic expansion for $P_\varepsilon = 0$. As can be seen from the asymptotic expansion, an increase in P_0, will eventually reduce the precision p_t by any finite amount, and a version of Proposition 2 holds.

VIII.2 Investment

In this subsection we introduce a market setting into our baseline mode and show that the price that arises in equilibrium generates a public signal centered around the average action in the population. In doing this, we also extend the baseline model by introducing a particular payoff externality: the payoffs to any individual will depend on the equilibrium price, which in turn is affected directly by the actions of others.

Consider an economy with two goods: a capital good and a final good. Capital goods are produced at a marginal cost that is increasing in the aggregate capital stock. In particular, if K_t is the aggregate stock of capital good at time t then the cost of producing an extra unit of capital at time t is assumed to be cK_t, for some $c > 0$. The capital good sector is competitive, implying that the price of the capital good at time t is cK_t.

There are a continuum of final-goods producers who accumulate capital until some random time τ. At that point, the final-goods market opens, and the final-goods producers transform their capital stock into final goods using a linear technology with marginal product x. We assume the same information structure as in our baseline model: at time zero, final-goods producers initially are asymmetrically informed about the state x of technology, then learn about it over time through public and private channels to be described below, until the state is realized at time τ.

Let a_{it} be the amount invested at time t by firm i. Agents' profits are realized at the same time as x:

$$\int_0^\tau \left[x a_{it} - cK_t a_{it} - \frac{a_{it}^2}{2} \right] dt,$$

where the second term, $cK_t a_{it}$ is the cost of purchasing the investment and the final term, $a_{it}^2/2$, is a quadratic adjustment cost.

Let $A_t = \int_0^1 a_{it} di$ be the aggregate level of investment at time t. The total capital stock evolves according to

$$dK_t = A_t \, dt + \frac{1}{P_\varepsilon} dW_t \tag{VIII.2}$$

where the term $\frac{1}{P_\varepsilon} dW_t$ captures random external demand shocks for capital.

Profit maximization by final-goods producers implies that $a_{it} = E_{it}(x) - cK_t$. Plugging this back into the previous equation, the law of motion for the capital stock becomes

$$dK_t = \left(\int_0^1 E_{it}(x) di - cK_t \right) dt + \frac{dW_t}{\sqrt{P_\varepsilon}} \ .$$

The change in the price of the capital good, cdK_t, is an endogenous public signal about the state of technology, x. Given that the price cK_t of the capital good is known by everyone, observing cdK_t is equivalent to observing

$$dZ_t = \left(\int_0^1 E_{it}(x) di \right) dt + \frac{dW_t}{\sqrt{P_\varepsilon}} \ . \tag{VIII.3}$$

Similarly to our baseline model, here we assume that agents also observe a private signal centered around the average action, i.e. A_t up to some idiosyncratic noise,

$$\left(\int_0^1 E_{jt}(x) di - cK_t \right) dt + \frac{d\omega_{it}}{\sqrt{p_\varepsilon}} \ .$$

As before, given that the price cK_t of the capital good is known by everyone, this signal is observationally equivalent to

$$dz_{it} = \left(\int_0^1 E_{it}(x) di \right) dt + \frac{d\omega_{it}}{\sqrt{p_\varepsilon}} \ . \tag{VIII.4}$$

The information structure is characterized by equations (VIII.3) and (VIII.4), which are the equations in our baseline model: therefore, the evolution of precisions in our baseline model also characterizes the evolution of precisions in this market setting, and our previous results concerning the dynamics and comparative statics remain unchanged.

However, in principle, the welfare function is different because of the presence of the cost of capital in the profits of the final-goods producers. Specifically, let us define the welfare function as the *ex ante* profit of a final-goods producer,

$$W = E_0\left[\int_0^\infty e^{-\lambda t}\left(xa_{it} - \frac{a_{it}^2}{2} - cK_t a_{it}\right)dt\right] = \int_0^\infty E_0\left[e^{-\lambda t}\left(xa_{it} - \frac{a_{it}^2}{2} - cK_t a_{it}\right)\right]dt\ .$$

Optimality implies $a_{it} = E_{it}(x) - cK_t$. Substituting into the welfare function, we find

$$\tilde{W} = \int_0^\infty e^{-\lambda t}E_0\left[-(x - E_{it}(x))^2 + \frac{(x - cK_t)^2}{2}\right]\ .$$

The first term inside the brackets is the same welfare flow as in our baseline model. We then need to compute $\int_0^\infty e^{-\lambda t}E_0(x - cK_t)^2 dt$,

Proposition 6. *The welfare function, $\tilde{W}$, is*

$$\tilde{W} = -\frac{\lambda}{\lambda + 2c}\int_0^\infty \frac{e^{-\lambda t}}{p_t + P_t}\,dt + \frac{1/\bar{P} + c^2(K_0^2 + P_\varepsilon/\lambda)}{\lambda + 2c}\,dt.$$

The last two terms reflect initial conditions that do not matter for our welfare exercise, and the first term is the same as before, except for the multiplicative constant. Thus, the welfare result in Proposition 2, which obtained for the baseline model, also holds in this market setting.

VIII.3 Omitted Proofs

VIII.3.1 Proof of Proposition 5

We can rewrite the solution as,

$$a_0 + (\tilde{p}_\varepsilon + p_\varepsilon)t = p_t + P_0\frac{\sqrt{p_\varepsilon}(p_\varepsilon - \tilde{p}_\varepsilon)}{\sqrt{\tilde{p}_\varepsilon}(p_\varepsilon + \tilde{p}_\varepsilon)}\arctan\left[\frac{P_0\tilde{p}_\varepsilon + (p_\varepsilon + \tilde{p}_\varepsilon)p_t}{P_0\sqrt{p_\varepsilon\tilde{p}_\varepsilon}}\right] + \\ + \frac{P_0 p_\varepsilon}{p_\varepsilon + \tilde{p}_\varepsilon}\log\left[P_0^2\tilde{p}_\varepsilon + 2P_0\tilde{p}_\varepsilon p_t + (p_\varepsilon + \tilde{p}_\varepsilon)p_t^2\right]\ .$$

Given that p_t goes to ∞, it follows that $p_t/t \to (p_\varepsilon + \tilde{p}_\varepsilon)$ as t goes to ∞. Now,

$$p_t - (p_\varepsilon + \tilde{p}_\varepsilon)t + \frac{P_0 p_\varepsilon}{p_\varepsilon + \tilde{p}_\varepsilon}\log(t^2) = a_0 - P_0\frac{\sqrt{p_\varepsilon}(p_\varepsilon - \tilde{p}_\varepsilon)}{\sqrt{\tilde{p}_\varepsilon}(p_\varepsilon + \tilde{p}_\varepsilon)}\arctan\left[\frac{P_0\tilde{p}_\varepsilon + (p_\varepsilon + \tilde{p}_\varepsilon)p_t}{P_0\sqrt{p_\varepsilon\tilde{p}_\varepsilon}}\right] - \frac{P_0 p_\varepsilon}{p_\varepsilon + \tilde{p}_\varepsilon}\log\left[P_0^2\tilde{p}_\varepsilon/t^2 + 2P_0\tilde{p}_\varepsilon p_t/t^2 + (p_\varepsilon + \tilde{p}_\varepsilon)(p_t/t)^2\right] .$$

Note that, since $\arctan(\,\cdot\,)$ and p_t/t are bounded, and since p_t/t is bounded away from zero, the right-hand side is indeed a bounded function.

VIII.3.2 Proof of Proposition 6

Let $\mu_t = x - cK_t$, and $\gamma_t = p_t/(p_t + P_t)$, then

$$d\mu_t \;=\; c\left[(1-\gamma_t)\left(x - \hat{X}_t\right) - \mu_t\right]dt - c\frac{dW_t}{\sqrt{P_\varepsilon}}$$

and, by Ito's lemma

$$\begin{aligned}
& d\left(\mu_t^2\right) = 2\mu_t d\mu_t + d\mu_t d\mu_t = 2\mu_t d\mu_t + \frac{c^2}{P_\varepsilon}dt \\
\Rightarrow\;& d\left(\mu_t^2\right) = 2\mu_t c\left(1-\gamma_t\right)\left(x - \hat{X}_t\right)dt - 2c\mu_t^2 dt - \frac{2\mu_t c}{\sqrt{P_\varepsilon}}dW_t + \frac{c^2}{P_\varepsilon}dt \\
\Rightarrow\;& \mu_t^2 = \mu_0^2 + \int_0^t 2\mu_s c\left(1-\gamma_s\right)\left(x - \hat{X}_s\right)ds - 2c\int_0^t \mu_s^2 ds - \int_0^t \frac{2\mu_s c}{\sqrt{P_\varepsilon}}dW_s + \frac{c^2}{P_\varepsilon}t.
\end{aligned}$$

Taking expectations on both sides, we obtain

$$E_0\mu_t^2 = E_0\mu_0^2 + \frac{c^2}{P_\varepsilon}t + 2c\int_0^t \left(1-\gamma_s\right)E_0\mu_s\left(x - \hat{X}_s\right)ds - 2c\int_0^t E_0\mu_s^2 ds \qquad \text{(VIII.5)}$$

To compute $E_0\left[\mu_t\left(x - \hat{X}_t\right)\right]$, note the following:

$$\begin{aligned}
E_0\left[\mu_t\left(x - \hat{X}_t\right)\right] &= E_0\left[(x - cK_t)\left(x - \hat{X}_t\right)\right] \\
&= E_0\left[\left(x - \hat{X}_t\right)\left(x - \hat{X}_t\right)\right] + E_0\left[\left(\hat{X}_t - cK_t\right)\left(x - \hat{X}_t\right)\right],
\end{aligned}$$

but, by definition of the conditional expectation $\hat{X}_t$, the residual $x - \hat{X}_t$ is orthogonal to $\hat{X}_t$ and K_t. Therefore, the second term is zero and we find that

$$E_0\left[\mu_t\left(x-\hat{X}_t\right)\right] \;=\; E_0\left[\left(x-\hat{X}_t\right)\left(x-\hat{X}_t\right)\right] = \frac{1}{P_t}.$$

Let $\bar{\mu}_t \equiv E_0\mu_t^2$. Then plugging the above finding back into (VIII.5), and keeping in mind that $\gamma_t = p_t/(p_t+P_t)$, we obtain

$$\bar{\mu}_t = \bar{\mu}_0 + 2c\int_0^t \frac{1}{p_s+P_s}ds - 2c\int_0^t \bar{\mu}_s ds + \frac{c^2}{P_\varepsilon}t\ .$$

Differentiating,

$$\dot{\bar{\mu}}_t = \frac{2c}{p_t+P_t} + \frac{c^2}{P_\varepsilon} - 2c\bar{\mu}_t\ .$$

The above differential equation can now be solved for $\bar{\mu}_t$:

$$\begin{aligned}\bar{\mu}_t &= e^{-2ct}\bar{\mu}_0 + \int_0^t e^{-2c(t-s)}\left(\frac{2c}{p_s+P_s}+\frac{c^2}{P_\varepsilon}\right)ds\\ &= e^{-2ct}\bar{\mu}_0 + \frac{c}{2P_\varepsilon}\left(1-e^{-2ct}\right) + \int_0^t e^{-2c(t-s)}\frac{2c}{p_s+P_s}ds.\end{aligned}$$

Note that, $\lim_{c\to 0} A_t = 0$. The welfare function is, then,

$$\tilde{W} = \int_0^\infty e^{-\lambda t}\left(-\frac{1}{p_t+P_t} + e^{-2ct}\bar{\mu}_0 + \int_0^t e^{2c(\tau-t)}\left(\frac{2c}{p_\tau+P_\tau}+\frac{c^2}{P_\varepsilon}\right)d\tau\right)dt \quad \text{(VIII.6)}$$

Now switching the order of the two integrals, we find

$$\int_0^\infty e^{-\lambda t}\int_0^t e^{2c(\tau-t)}\frac{2c}{p_\tau+P_\tau}d\tau\, dt = \int_0^\infty e^{-\lambda\tau}\frac{2c}{r+2c}\frac{1}{p_\tau+P_\tau}\,d\tau.$$

Plugging back and integrating the other terms, we find:

$$\tilde{W} = -\frac{\lambda}{\lambda+2c}\int_0^\infty \frac{e^{-\lambda t}}{p_t+P_t}\,dt + \frac{A_0}{\lambda+2c} + \frac{c}{2P_\varepsilon}\frac{2c}{\lambda(\lambda+2c)}.$$

Now recall that $\bar{\mu}_0 = E_0[(x-cK_0)^2] = E_0[x^2 - 2cK_0x + c^2K_0^2] = 1/\bar{P} + c^2K_0^2$, and we are done.