

MPRA

Munich Personal RePEc Archive

Inertia in social learning from a summary statistic

Larson, Nathan

University of Virginia

2008

Online at <https://mpra.ub.uni-muenchen.de/32143/>
MPRA Paper No. 32143, posted 11 Jul 2011 19:56 UTC

Inertia in Social Learning from a Summary Statistic

Nathan Larson
Department of Economics
University of Virginia

July 8, 2011

Abstract

We model normal-quadratic social learning with agents who observe a summary statistic over past actions, rather than complete action histories. Because an agent with a summary statistic cannot correct for the fact that earlier actions influenced later ones, even a small presence of old actions in the statistic can introduce very persistent errors. Depending on how fast these old actions fade from view, social learning can either be as fast as if agents' private information were pooled (rate n) or it can slow to a crawl (rate $\ln n$). We also examine extensions to learning from samples of actions, learning about a moving target, heterogeneous preferences, and biases toward own information.

1 Introduction

We introduce a model of social learning in which people learn from statistics over other people's past actions. In the baseline model, this statistic is an average over a large pool of past actions; one could think of it as some kind of macroeconomic indicator. In a variation, the statistic is an average over a small, idiosyncratic sample of actions; one could think of this as word of mouth learning. The key feature in both cases is that the signal *summarizes* history: context about the sequence of individual past actions is lost. This lack of context is costly because it prevents people from unraveling interdependencies among past actions generated by social learning in the past. Our main aims in the paper are to show that this typically creates a bias toward old information, and to study how this bias affects the speed of learning.¹

The baseline model PA (for population average) builds on Vives [19]. A new cohort of agents arrives at each stage (1, 2, ..., n , ... and so on) and must choose an action once and for all in order to minimize a quadratic loss function. The common optimal action θ is unknown, but each agent receives a private signal, normally distributed around θ , and an observational signal based on prior agents' behavior. After acting the agent passes into an observation pool: this is a collection of agents

¹Herding on a sub-optimal action is not a possibility in our model; the continuous action space ensures that learning always continues and beliefs eventually converge to the truth. However there are natural comparisons between the mechanisms that cause *slow* learning in our model and the factors that stop learning in herding models – we touch on these later in this section.

whose actions remain visible or on display, in a sense that we clarify momentarily. Agents who are already in the pool exit from it at some rate; this is a shorthand way to say that their actions no longer contribute to observations by future agents. The observational signal that an agent sees before deciding is the *mean* action in this pool. The only difference in the sample average (SA) variation of the model is that each agent sees the mean of a random sample of actions drawn from the observation pool, rather than the average over the whole pool.

For the baseline model, the observation pool can be interpreted as the set of past actors who contribute to some aggregate measurement. To illustrate, imagine that new cohorts of youths arrive periodically and make decisions about how much education to acquire. Each of them can see a statistic on the average educational attainment of adults in the labor force. In this case, the observation pool – the labor force – reflects the past education choices of people who are still working. Because older people are more likely to have left the labor force, the observational signal is (to a first approximation) a recency-weighted average of past choices.² As a second example, suppose that most people begin to save seriously for retirement around age forty, make once-and-for-all decisions about a savings rate, and agree that aiming for a 70% ratio of retirement income to current income is desirable.³ The savings rate needed to hit this goal is uncertain, but people can see a government statistic on the average savings rate among 40-50 year olds. In this case, the observation pool consists largely of decisions that are one to ten years old, but there may be some older decisions as well. To illustrate a case where the “statistic” is less literal, consider a small farmer in a developing country whose village has adopted a new crop. In deciding how heavily to irrigate his field he may look at the level of the local reservoir to get an indication of the average water use by farmers who planted earlier than him. The observation pool model is well suited to other situations where people see a blend of recent and older actions; by varying the rate at which old actions exit out of the pool, one can give the decisionmaker an average of all past actions, or only the most recent ones, or a recency-weighted blend.

The sample average model describes settings where no comprehensive public statistic is available, and people collect a few examples to guide their decisions. For example, a new employee organizing her retirement plan may ask a few of her co-workers about the stock-bond allocations they chose, or a sprinter interested in optimizing his diet may browse a few web sites to see the protein-to-carbohydrate ratios that other athletes mention. In each case the observation pool of actions that are available to sample may be tilted toward recent ones – co-workers of longer tenure at the firm retire or quit, and older websites may eventually be updated or taken down. To apply model SA to examples like these, we need two additional criteria to hold. First, the decisionmaker has poor contextual information about exactly when her samples acted and what observations they had access to at the time. This would be true if the employee does not know exactly when her co-

²Of course, other factors besides age can affect the set of predecessors that is available to observe. Selection of the observational signal along non-age dimensions creates additional interesting issues for social learning, but we will not tackle those issues here.

³Of course, the savings rate is not really a once-and-for-all decision, but there is considerable evidence (e.g. Duflo and Saez [9]) that people do not revise their retirement plans very often.

workers were hired, or if the athlete’s web sites are undated and lack references.⁴ Loosely, we may say that the decisionmaker has no way to distinguish one sampled action from another.⁵ Second, the way the decisionmaker handles context-free actions is to treat them all equally.⁶ Under these assumptions, the best estimate of θ that the decisionmaker can form from her sample will be the sample average. By assuming that the decisionmaker simply observes this sample average directly, model SA stylizes the idea of observing individual but context-free actions.

Learning is measured by the precision of agents’ estimates of θ , or inversely by the squared error in their actions. As a benchmark, if agents’ private signals were simply published as soon as they arrived, then action precision would grow at rate n (where n is the number of cohorts that have arrived); we call this the *full-information* rate of learning, or fast learning. Our main results for model PA (Proposition 2) and SA (Proposition 6) show that social learning proceeds at the log of the rate at which old actions exit from the observation pool.⁷ If old actions exit at least exponentially – that is, if no more than e^{-k} of a cohort’s actions remain in the pool k periods after that cohort acted – then learning is fast. If old actions exit only polynomially (e.g., only k^{-a} actions remain k periods later) then learning slows to the log of the full-information rate. This is exceedingly slow: not before roughly e^n cohorts have acted will public information be as good as it would have been if the first n cohorts had published their signals.

To illustrate the bias that gives rise to slow learning, suppose that old actions never die in model PA, so each agent observes the equally weighted average of all past actions. Write x_1, x_2, x_3, \dots for the average action of agents in cohorts 1, 2, 3, and so on. Consider all of the channels through which private information held by the cohort 1 agents becomes impounded in the observational signal seen at stage $n + 1$. Action x_1 appears in its own right, but it also transmits influence indirectly through its effect on x_2 , the effect of both x_1 and x_2 on x_3 , and so on. In this way, the indirect influence of an old action can come to dwarf its direct contribution to the summary statistic. Fast social learning requires old actions to be purged fast enough to counter not just their direct influence, but also this compounding indirect influence. To measure the bias toward old information, we calculate the fraction of the squared error in the observational signal that can be attributed to idiosyncratic

⁴This issue has received considerable scrutiny in the medical community in the context of studies about the quality of online medical information. In one meta-study covering up to 1300 websites [12], researchers found that 58% to 83% of information sources failed to report even minimal information about how current their content was (such as when the site was created or last updated).

⁵In reality, the employee probably has some inkling about the order of her co-workers’ decisions. But the main reason that the order of decisions is important for inference is that she must be able to assess who her co-worker could have relied on for advice, and who might have relied on him. This information may be much murkier to her.

⁶Two notes. First, since actions will turn out to be normally distributed about θ , an optimal estimate based on past actions will always be a linear combination with some weights. By “treat equally,” we mean the weights must be equal, and this leads to the sample average.

Second, demanding that an agent treat unordered actions equally is not an innocuous constraint. A statistically clever agent should try to find clues to the order of the actions in the way they are grouped (that is, which ones are most similar or least similar to each other), and treating actions differently based on these clues should improve the estimate of θ . However, this type of analysis appears terrifically complicated – one must estimate not just when predecessors acted, but also what they believed about when their own predecessors acted, and so on. Assuming equal treatment amounts to an assumption that agents are not this clever.

⁷Propositions 2 and 6 are actually cast in terms of the rate at which the observation pool grows, which is negatively related to how fast old actions die out. The more intuitive explanation here follows Corollary 1.

error in each prior private signal (Proposition 3). If the observation pool aggregated information efficiently, then the share of this error due to any single signal would tend to zero over time. However, when old agents disappear from the observation pool slowly (sub-exponentially), the idiosyncratic error they introduce becomes locked in, as a fraction of total error. In one example, the error in the signals of the very first cohort to act is responsible for at least 82% of the error of every subsequent agent, no matter how far into the future! An agent in the model understands that she is facing an echo chamber – from her point of view, her recent predecessors have parroted the same old information too much and relied on their new private signals too little. But to correct this problem, she would need to be able to trace the chain of dependencies from one action to the next, and her information is not rich enough to let her do this. This also helps to account for the similar results in the population average and sample average models; because errors in actions quickly become highly correlated (Proposition 5), larger samples are not much more informative than small ones.

In many practical settings, people’s optimal actions will be similar but not necessarily identical. For example, the optimal level of schooling for a particular student may depend on her aptitude, finances, and so forth. Thus we extend model PA to allow for heterogeneity: each individual’s private signal is now centered on her *personal* best action, which differs from θ by a relative taste shock (her ‘type’). If agents know their relative types (such as a student who understands how smart she is relative to her classmates), the prior results still apply; each agent just adjusts her best estimate of θ by the taste shock. However, if an agent does not know her type, she must rely relatively heavily on her private signal (since this is her only personalized information). In this case, heterogeneity induces a bias toward recent information, and so the persistence of old actions in the observation pool, by creating a countervailing bias, can actually be helpful rather than pernicious (Proposition 7). The logic of the recency bias begins by noting that an agent now views *all* of her predecessors as having over-relied on their private signals, since their taste shocks are not relevant to her. However, it is her *recent* predecessors whose over-reliance is most egregious to her, since these agents saw relatively better observational signals of θ to which (in her view) they should have deferred much more than they did.

While we have focused on social learning about a fixed parameter θ , in many settings the target of learning might be changing over time. To study this case, we extend the model to allow θ to drift according to a random walk. Not surprisingly, this intensifies the harm done by old actions in the observation pool, as those actions are now not only too influential, but also out of date. In fact, we show that if old actions do not exit at least exponentially fast, then the observational signal becomes worthless: its precision tends to zero, and agents rely entirely on their private signals (Proposition 9).

All of our analysis relies on a *uniform exit* assumption. In effect, this says that while the contributions of two old cohorts to the observational signal may decline at any rate one likes, their weights *relative to each other* do not change over time. Uniform exit buys a great deal of tractability – it gives the learning dynamics a recursive characterization – but it has no deep justification. In general under non-uniform exit the state space of the learning dynamics explodes, making analysis

very difficult, but we provide examples illustrating that a small chance of observing an arbitrarily old action can suffice to slow down learning. One of these is rather striking: suppose that only actions in the very first cohort have any persistence, perhaps due to some first-mover prominence; otherwise the observational signal is focused on the most recent actions. Furthermore, let the weight on the first-movers in the observational signal tend to zero over time. Unless that weight tends to zero fast enough, the long run rate of social learning will still be slow.

Much of the existing literature on sequential observational learning, following Banerjee [3] and Bikchandani, Hirshleifer, and Welch [5], has focused on herding. In this literature, as in our paper, an agent uses two sources of information to make a decision: first, some private information and second, her observations about how other agents have acted in similar decision problems.⁸ It is well understood that this generates an externality – an agent does not account for the fact that future agents will tend to be more interested in her exclusive private information than in her observational information (which they may be able to duplicate), thus she places too little weight on her private information when she acts. Still, this would not cause a problem for future agents if they were able to perfectly invert her action, thus recovering her private signal. The central insight of the herding literature is that if actions are imperfectly invertible, then this externality can slow down learning or even drive it to a halt. In many of the classic herding papers, this non-invertibility arises because the action space has lower dimension (often the action is binary) than posterior beliefs. With rare exceptions, this literature treats the full sequence of actions as observable in order to focus on the challenge of recovering beliefs from actions.

In a similar vein, Vives [19] assumes that the full sequence of actions is observed, but that each action is observed with noise, so once again beliefs cannot be perfectly recovered from actions. He shows that the rate of learning slows to $n^{\frac{1}{3}}$. Our model is very similar, but the only source of noise is the conflation of early and late actions – we show that this can be an even more severe impediment to learning.

Our paper is not the first to dispense with the assumption that the sequence of past actions is perfectly observed. Ellison and Fudenberg [10], [11] and Banerjee and Fudenberg [4] introduce models in which agents learn from aggregate statistics about samples of past actions (or outcomes). Because these models all give agents a binary action choice, their focus is necessarily on whether, rather than how quickly, complete learning is achieved. Like us, Smith and Sorensen [17] show the desirability of a sampling scheme that favors recent actions, but their results are also on whether learning is achieved, not its speed. Çelen and Kariv [7] revisit the binary action, sequential model under the extreme assumption that only the most recent action is observed. They find that while herding is never permanent (as it would be with perfectly observed histories), longer and longer spells of herding occur as time goes on. This highlights an important difference between models with and without an invertible mapping from beliefs to actions – in our model, when everyone observes her immediate predecessor, that observation is sufficient for the predecessor’s belief, and learning is at

⁸ Actually, one could imagine many decision outcomes that could be observed instead of, or in addition to, actions. As the informational issues are similar regardless of which outcome variable is observed, we follow most of the literature in focusing on observable actions.

the full information rate. In an influential paper, Smith and Sorensen [16] showed among other results that complete learning is achieved with finite actions if at least some agents have private signals that are arbitrarily precise. They point out that a contrarian action is extremely influential for later agents precisely because the contrarian must have been quite sure of herself to have deviated from the herd. Callander and Hörner [6] demonstrate that this intuition can apply even when an agent observes only the total number of choices of each type, rather than the full action sequence. They develop conditions under which later agents optimally tilt toward the minority, rather than the majority action.

Another branch of the literature, launched by Bala and Goyal [2], assumes that agents observe each other according to a social network, typically represented by a graph. Bala and Goyal show that if the network structure makes particular agents too influential, lock-in on the wrong action can occur. In DeMarzo, Vayanos, and Zwiebel [8], the network represents a channel for agents to share beliefs (rather than observe actions or outcomes) with each other. With complete information about the network structure, each agent faces a very challenging inference problem: she must understand exactly who learned what from whom in order to correct for correlation and redundancy in neighbors' beliefs. If agents use simpler inference heuristics, then the information of well-connected individuals tends to accumulate too much weight. Our paper shares with both of these the intuition that “over-observation” of a particular set of agents impedes learning, but in our case it is early actors rather than well-connected ones who are observed disproportionately often. Finally, Acemoglu, Dahleh, Lobel, and Ozdaglar [1] develop very general conditions on the network topology under which complete learning is attained. Our objective is complementary to all of these papers; in our model, complete learning occurs in the limit, and our focus is on how long this takes. The question is a very practical one because rate n and rate $\ln n$ learning are very different animals – the latter may resemble incomplete learning over any reasonable time horizon.

Finally from a technical point of view, assuming the full sequence of past actions to be perfectly observed has the advantage of permitting a compact, recursive representation of the learning dynamics. (This is its main appeal, since it is a heroic assumption for most practical applications.) The sequence of observations then constitutes a filtration, the most recent observation suffices as a state variable (since it is sufficient for all public information), and powerful tools like the Martingale Convergence Theorem can be brought into service. Our uniform exit assumption allows us to retain the tractability of a recursive model without requiring agents to have unrealistically detailed information about the past. Extensions to the model that preserve this recursive structure and avoid adding more state variables are relatively easily handled; we have discussed a few, but one can imagine many others. Extensions that bring in additional state variables are more challenging; a comparison of the proofs for models PA (one state variable) and SA (two) will give a sense of the additional complications.

The rest of the paper is laid out as follows. Section 2 introduces our baseline model (PA), while Section 3 characterizes the speed of learning and the persistence of old information. Section 4 develops the results for model SA. Section 5 develops extensions, including heterogeneous tastes

and the moving target. Section 6 examines examples that relax the uniform exit assumption, while Section 7 concludes.

2 The Model

There is an infinite sequence of cohorts, each consisting of a unit measure of agents, that arrive at stages $\{1, 2, \dots, n, \dots\}$. The agents in each cohort are identical, in the sense that they face the same decision and share the same utility function, so we focus on the decision problem of a representative agent from cohort n . To minimize ambiguous pronouns, we usually keep the convention that an agent whose decision being examined is female while her predecessors are male.⁹ At stage 0, Nature chooses once-and-for-all a value for $\theta \in \mathfrak{R}$, the parameter that the agents will try to estimate. Agents' priors on θ are diffuse.¹⁰ At stage n , agent n_i arrives (for $i \in [0, 1]$), observes information \mathcal{I}_{n_i} , and then chooses an action x_{n_i} to minimize the quadratic loss function

$$E((x_{n_i} - \theta)^2 | \mathcal{I}_{n_i})$$

The information set $\mathcal{I}_{n_i} = \{s_{n_i}, o_{n_i}\}$ contains two elements: an individual signal s_{n_i} , and an observation o_{n_i} about past agents' actions. The individual signal is given by $s_{n_i} = \theta + \varepsilon_n + \varepsilon_{n_i}$, where $\varepsilon_n \sim N(0, v_c)$ is i.i.d. across cohorts and $\varepsilon_{n_i} \sim N(0, v_{ind})$ is i.i.d. across individuals and cohorts. The cohort error captures the possibility that a common factor affects the estimates of all agents who act at stage n , while the individual error captures idiosyncratic noise across agents within a cohort. We follow Vives [19] in making the convention that the Strong Law of Large Numbers applies exactly to the individual errors for a cohort, and also to any measurable subset of a cohort. That is, we assume that for each cohort,¹¹

$$\int_A \varepsilon_{n_i} di = 0 \text{ for any measurable subset } A \subseteq [0, 1] . \quad (1)$$

This implies that the mean of the individual signals in cohort n is $\int_0^1 s_{n_i} di = \theta + \varepsilon_n$; as a matter of shorthand, we will refer to $s_n = \theta + \varepsilon_n$ as the 'cohort signal.' Similarly, we refer to $x_n = \int_{[0,1]} x_{n_i} di$ as the average action in cohort n , or simply the cohort action. For the most part, the analysis will only require keeping track of s_n and x_n , not the individual signals and actions. Let $v = v_c + v_{ind}$ be the variance of s_{n_i} . We require $v > 0$, but we allow the possibility of individual signals that are uncorrelated ($v_c = 0$) or perfectly correlated ($v_{ind} = 0$) across agents in a cohort. The structure of

⁹Of course, this convention makes gender time-inconsistent, but that is immaterial to the results.

¹⁰None of the results depend materially on diffuse priors, but by using them we will reduce the notational burden substantially.

¹¹Although property (1) is often used in the economics literature, integrating over a continuum of i.i.d random variables presents certain technical complications, as discussed by Judd (1985) and Feldman and Gilles (1985). Green (1994) shows that (1) can be justified if, instead of the usual Borel sets, one endows the population of the cohort with an uncountably generated sigma-algebra. For the sake of expositional simplicity, we will just assume that (1) holds. Note that idiosyncratic shocks will not play any role in the main results. If they are absent (that is, if $v_{ind} = 0$), then (1) can be dropped, and the main results do not change.

the game, and the primitives v_c and v_{ind} are common knowledge; furthermore, each agent knows which cohort she belongs to. The observational signal o_{n_i} of agent n_i depends on an observation pool \mathcal{P}_{n-1} which contains a measure P_{n-1} of the agents who arrived in stages 1 through $n - 1$. Throughout most of the paper, we make the following assumption about this observational signal.

Population Average (PA) Agent n_i observes the mean action \bar{x}_{n-1} taken by all of the agents in \mathcal{P}_{n-1} .

This could represent a situation in which agents have access to a public summary statistic about past actions. In Section 4, we will consider the following alternative assumption.

Sample Average (SA) Agent n_i observes the mean action of S individual agents drawn randomly from \mathcal{P}_{n-1} , for an integer $S \geq 1$. (Agent n_i does not observe which cohorts these agents were drawn from.)

After the agents in cohort n choose their actions x_{n_i} , the observation pool is updated. The new observation pool \mathcal{P}_n is constructed from \mathcal{P}_{n-1} by the following steps. First, a measure $d_{n-1} \leq 1$ of (randomly chosen) agents in \mathcal{P}_{n-1} exit from the observation pool. Then all of the cohort n agents are added to get \mathcal{P}_n , and the game proceeds to stage $n + 1$. Thus the size of the observation pool grows according to $P_n = P_{n-1} + 1 - d_{n-1}$. It is assumed throughout that the sequence $\{d_n\}$ converges, and we also impose the following condition on the departure of old actions from the observation pool.¹²

Uniform Exit Every agent in \mathcal{P}_{n-1} has the same chance $(1 - d_{n-1}/P_{n-1})$ to survive and be present in \mathcal{P}_n .

Remark 1 *Under these assumptions, the average action in the observation pool at stage n will be a weighted average of the first $n - 1$ cohort actions x_1 through x_{n-1} , with (weakly) lower weight on older cohorts.*

The first part of the remark follows because for any cohort $m < n$, the average action among cohort m agents remaining in \mathcal{P}_{n-1} is the same as the average action among those who have departed – both are equal to x_m , by (1). Older cohorts have smaller weights because their ranks have been culled more times. Uniform exit implies that the observational signal preserves the relative weights on the actions of different past cohorts. That is, for any two past cohorts m and m' , the contributions of both x_m and $x_{m'}$ to the average action in \mathcal{P}_{n-1} decline with n , but the ratio those contributions does not change over time. The appeal of the uniform exit condition is mainly technical; it opens the door to a recursive treatment of the observational signal without which the analysis would be much less tractable.¹³

¹²Requiring $\{d_n\}$ to converge is mainly for technical and expositional convenience. This rules out persistent, abrupt changes in the death rate, including cycles such as $(d_1, d_2, d_3, \dots) = (0, 1, 0, 1, \dots)$. Extending the analysis to handle such cases would be fairly routine but tedious.

¹³This assumption plays a similar role to the ‘recursive sampling’ assumption used by Smith and Sørensen [17]. Among other implications, this assumption will impose a degree of smoothness on how the representation of cohorts in \mathcal{P}_{n-1} changes with their vintage. For example, the assumption rules out finite moving averages in which agent n ’s observation puts positive weight on all cohorts more recent than $n - k$ and zero weight on earlier cohorts.

We define several closely related measures of learning: $\tau_{n_i} = \left(E((x_{n_i} - \theta)^2)\right)^{-1}$ is the precision of agent n_i 's estimate of θ , while $\kappa_n = E((\bar{x}_n - \theta)^2)$ is the squared error in the population average action after cohort n . The events at stage $n = 1$ differ slightly from the general case. These agents have no one to observe, so they must act on the basis of their individual signals s_{1_i} alone. Furthermore, since there is no prior sample for the death rate to act upon, we have $P_1 = 1$, where \mathcal{P}_1 includes the full complement of cohort 1 actions and nothing else.

The size, composition, and evolution of the observation population depends entirely on our assumptions about the exit rate d_n . For example, if $d_n = 1$ for all n , we say that there is *immediate exit*. In this case, $P_n = 1$, \mathcal{P}_n always consists of the most recent cohort, and each agent observes the mean action of the most recent cohort. Alternatively, if $d_n = 0$ for all n , then $P_n = n$ and the observation pool puts equal weight on all past actions. In this case, agent $n + 1$ observes an equally weighted blend of older and more recent actions. More generally, the larger d_n is, the smaller the contribution of older actions will be in the average that an agent observes. To make this a bit more precise, we define a few terms. Let $\tilde{d} = \lim_{n \rightarrow \infty} \frac{d_n}{P_n}$ be the limiting percentage exit rate; the existence of this limit is straightforward to establish.¹⁴ Let $\#_m^n$ be the measure of cohort m actions that remain in \mathcal{P}_{n-1} , the pool observed at stage n , and let $\lambda_m^n = \#_m^n / P_{n-1}$ be their fraction in that pool.

Def. We say that cohort m actions exit *asymptotically exponentially* if $\lim_{n \rightarrow \infty} \frac{\ln \#_m^n}{n-m} < 0$. Similarly, we say their share of the observation pool declines asymptotically exponentially if $\lim_{n \rightarrow \infty} \frac{\ln \lambda_m^n}{n-m} < 0$. We say that exit (or decline of λ_m^n) is *asymptotically slower than exponential* if the respective limit is zero.

Def. The observation pool is *bounded* if $P_n \rightarrow N^*$, for some $N^* < \infty$, and *unbounded* if $P_n \rightarrow \infty$.

Except for one special case ($d_n = 1$ for all n), the existence of these limits follows directly from the existence of \tilde{d} . Notice that asymptotically exponential exit implies that the measure of cohort m agents remaining shrinks roughly as $e^{-r(n-m)}$, for some $r > 0$. We will say the observation pool has property EE_1 (SE_1) if cohort m actions exit asymptotically exponentially (slower than exponentially) for all $m \geq 1$. Define properties EE_2 and SE_2 analogously for the share of cohort m actions in the observation pool. Proposition 1 shows that these various measures of the persistence of old actions are essentially equivalent.

Proposition 1 (*Equivalent characterizations of exit*) *Suppose exit is not immediate. (There is at least one cohort n' with $d_{n'} < 1$.) Then,*

1. $\tilde{d} > 0 \Leftrightarrow P_n$ is bounded $\Leftrightarrow EE_1 \Leftrightarrow EE_2$
2. $\tilde{d} = 0 \Leftrightarrow P_n$ is unbounded $\Leftrightarrow SE_1 \Leftrightarrow SE_2$

Excluding immediate exit is mainly for brevity of presentation – one may think of old actions as exiting faster than exponentially in this case, since the limits defining EE_1 and EE_2 diverge. Given

¹⁴ $\{P_n\}_{n \geq 1}$, as a monotonically increasing sequence, must either converge or tend to infinity. In either case, $\lim_{n \rightarrow \infty} \frac{d_n}{P_n}$ exists.

Proposition 1, in the sequel we will often refer simply to exponential or slower than exponential exit of old actions, without specifying $\#_m^n$ or λ_m^n , and with the qualifier “asymptotically” always implied.

3 Equilibrium

The first cohort is a special case: without predecessors to observe, and given her diffuse priors, the optimal action for agent 1_i is simply her signal: $x_{1_i} = s_{1_i}$. For cohorts $n \geq 2$, we begin by reviewing some standard results on normal learning. Suppose that agent n_i observes \bar{x}_{n-1} and s_{n_i} and believes that $\bar{x}_{n-1} \sim N(\theta, \kappa_{n-1})$, $s_{n_i} \sim N(\theta, v)$, and that \bar{x}_{n-1} and s_{n_i} are independent conditional on θ . Then she has a unique optimal action given by the precision-weighted average of her two signals:

$$\begin{aligned} x_{n_i} &= E(\theta | \bar{x}_{n-1}, s_{n_i}) = (1 - \alpha_n)\bar{x}_{n-1} + \alpha_n s_{n_i} \\ \alpha_n &= \begin{cases} 1 & \text{if } n = 1 \\ \frac{\kappa_{n-1}}{\kappa_{n-1} + v} & \text{if } n > 1 \end{cases}, \quad \tau_{n_i} = \kappa_{n-1}^{-1} + \frac{1}{v} \end{aligned} \quad (2)$$

These conditions apply for cohort 2 agents: they can infer how cohort 1 agents must have acted, and this permits them to infer that $\bar{x}_1 \sim N(\theta, \kappa_1)$, with $\kappa_1 = v_c$. By acting optimally, they ensure that \bar{x}_2 is a mixture of normally distributed, mean θ signals. Thus, $\bar{x}_2 \sim N(\theta, \kappa_2)$, where κ_2 depends on α_2 . But then, because cohort 3 agents can deduce how agents in the first two cohorts must have acted, they can infer κ_2 , and so they too act according to (2). This logic applies inductively to all future cohorts.

In order to provide a benchmark for the main results, we briefly discuss the learning path under full information and under perfect observability of actions. Then we characterize learning in our model with bounded and unbounded observation pools. Throughout, we will say that the sequence κ_n converges to zero at rate $g(n)$, for some increasing function $g(n)$, if the sequence $\{g(n)\kappa_n\}$ converges to a finite, positive constant. Similarly, τ_{n_i} diverges at rate $g(n)$ if $\left\{\frac{\tau_{n_i}}{g(n)}\right\}$ converges to a finite, positive constant. Note that these rates are measured with respect to the arrival rate of new information with new cohorts, not calendar time, which plays no particular role in the model. Of course, if we were to specify that new cohorts arrive at a certain rate per day or per year, it would be a simple matter to express the convergence rate of κ_n or τ_{n_i} in terms of calendar time.

3.1 Benchmarks

We define two alternative versions of our game which will be used as benchmarks. In the *full information* version of the game, agent n_i 's observational signal o_{n_i} contains the full sequence of individual signals for all prior agents. In the *perfect observability* version, o_{n_i} contains the full sequence of actions by all prior agents. In the main model, κ_n is the variance of the deviation from θ of the observational signal observed by cohort $n + 1$. By analogy with this, let κ_n^{FI} be the variance of the best estimate of θ that can be made from the observational signal available to a cohort $n + 1$

agent in the full information game (and similarly for κ_n^{PO} in the perfect observability version).

In the full information case, agent $(n+1)_i$'s observational signal $o_{(n+1)_i}$ is essentially equivalent to (s_1, \dots, s_n) , since she can integrate out the individual-specific error terms. Furthermore, since the prior cohorts' signals are i.i.d. conditional on θ , a sufficient statistic for them is their mean $\bar{s}_n = \frac{1}{n} \sum_{i=1}^n s_i$, where $\bar{s}_n \sim N(\theta, v_c/n)$. The optimal action based on the observational signal alone is \bar{s}_n , so $\kappa_n^{FI} = \frac{v_c}{n}$. In other words, if all signals become public upon arrival, then the precision of the observational signal improves at rate n . Of course, this is just the standard, textbook result for aggregating normal random variables. Agent $(n+1)_i$'s optimal action is a precision-weighted average of \bar{s}_n with her individual signal s_{n_i} . This is $x_{(n+1)_i}^{FI} = \frac{v_c s_{(n+1)_i} + n v \bar{s}_n}{v_c + n v}$ which delivers precision $\tau_{(n+1)_i}^{FI} = \frac{n}{v_c} + \frac{1}{v}$.

Under perfect observability, it becomes quickly evident that even though agents do not observe the sequence of signals directly, they can deduce the signals from the sequence of actions, and so learning is as if we were in the full information case. Suppose that $x_{m_i}^{PO} = x_{m_i}^{FI}$ for all $m \leq n$. That is, all agents in cohorts 1 through n act as though they had access to the sequence of past signals. Then the average action for each such cohort m must be a convex combination of the signals $\vec{s}_m = (s_1, \dots, s_m)$, with strictly positive weight on s_m .^{15,16} It follows that the vector $\vec{x}_n^{PO} = (x_1^{PO}, \dots, x_n^{PO})$ can be written $\vec{x}_n^{PO} = M \vec{s}_n$, where M is a lower triangular matrix with a strictly positive main diagonal. M is invertible, so a cohort $n+1$ agent can recover $\vec{s}_n = M^{-1} \vec{x}_n^{PO}$ by observing past cohort average actions \vec{x}_n^{PO} . Consequently, if all earlier agents acted as if they had access to full information about prior signals, then cohort $n+1$ agents do as well. Because cohort 1 trivially acts on full information, we can conclude inductively that all agents do, and so we have $x_n^{PO} = x_n^{FI}$ and $\kappa_n^{PO} = \kappa_n^{FI} = \frac{n}{v_c}$ for all n . Once again, learning is at rate n .

3.2 The Speed of Learning in Case PA

Now we return to our baseline model PA: an agent observes the average action in a pool that mixes recent actions and older ones. Define the partial sums of P_n^{-1} as $R_n = \sum_{m=1}^n \frac{1}{P_m}$. From Proposition 1, one can see that more persistent old actions tend to make P_n larger and R_n smaller; this relationship will be sharpened in Section 3.3. Our main result on the speed of learning is the following.¹⁷

Proposition 2 *If P_n is bounded, learning is at rate n . If P_n is not bounded, learning is at rate R_n , which is slower than n . (That is, $\frac{R_n}{n} \rightarrow 0$.) In particular, if d_n is bounded away from 1, learning is at rate $\ln n$.*

¹⁵In fact, it will be $x_m^{PO} = \frac{v s_1 + \dots + v s_{m-1} + v_c s_m}{v_c + (m-1)v}$, which follows from $x_{n_i}^{FI}$ above.

¹⁶A simpler approach would be to note that if $x_n^{PO} = \bar{s}_n$, then agent $n+1$ can compute its full information action from x_n^{PO} alone: $x_{n+1}^{PO} = \frac{1}{n+1} s_{n+1} + \frac{n}{n+1} x_n^{PO} = \bar{s}_{n+1}$. This depends on the fact that with normal signals, \bar{s}_n is sufficient for \vec{s}_n . However, we want to emphasize that sufficiency is not critical here because the full vector of signals can be recovered, hence the more circuitous approach.

¹⁷If P_n is bounded, then R_n grows at rate n . Thus a more concise statement of the proposition is that learning is at rate R_n , without qualification on P_n . We separate the bounded and unbounded cases only to emphasize when learning will be slower than the full information rate.

The proof is in the appendix. In the rest of this section, we develop the equations of motion that govern the learning process and sketch some intuition for the result. To begin, suppose that the agents in some cohort n can infer that the signal quality of the average action that they observe is κ_{n-1} . Then how heavily will these agents weight their own signals, and how will their actions feed into the κ_n observed by the next cohort? From (2), we have $x_{n_i} = (1 - \alpha_n) \bar{x}_{n-1} + \alpha_n s_{n_i}$, with $\alpha_n = \frac{\kappa_{n-1}}{v + \kappa_{n-1}}$. Thus, using (1), the average action within cohort n is $x_n = (1 - \alpha_n) \bar{x}_{n-1} + \alpha_n s_n$. Meanwhile, a measure d_{n-1} of agents exit from the observation population. Because these agents are selected randomly, the average action among these departers is identical to the average action among those who remain – both are equal to \bar{x}_{n-1} . Then, the new observation population is composed of a measure $P_{n-1} - d_{n-1}$ of remaining agents and a measure 1 of cohort n agents. The updated average action is

$$\bar{x}_n = \left(1 - \frac{\alpha_n}{P_n}\right) \bar{x}_{n-1} + \frac{\alpha_n}{P_n} s_n. \quad (3)$$

Because the errors in \bar{x}_{n-1} and s_n are independent, we have $\bar{x}_n \sim N(\theta, \kappa_n)$, where:

$$\kappa_n = \left(1 - \frac{\alpha_n}{P_n}\right)^2 \kappa_{n-1} + \left(\frac{\alpha_n}{P_n}\right)^2 v_c \quad (4)$$

We can express this difference equation in κ_n as

$$\begin{aligned} \kappa_n - \kappa_{n-1} &= -2 \frac{\alpha_n}{P_n} \kappa_{n-1} + \left(\frac{\alpha_n}{P_n}\right)^2 (v_c + \kappa_{n-1}) \\ &= -M_n \frac{\kappa_{n-1}^2}{v + \kappa_{n-1}} \end{aligned} \quad (5)$$

where $M_n = \left(2 - \frac{1}{P_n} \frac{v_c + \kappa_{n-1}}{v + \kappa_{n-1}}\right) \frac{1}{P_n}$.

To illustrate how the results of Proposition 2 come about, we introduce $\bar{\tau}_n = \frac{1}{\kappa_n}$, the precision of the observational signal after cohort n . (This is linked to τ_{n_i} , the precision of a cohort n agent's action x_{n_i} , by $\tau_{n_i} = \bar{\tau}_{n-1} + v$, since cohort n observes \bar{x}_{n-1} and combines it optimally with s_{n_i} .) If we express (5) in terms of $\bar{\tau}_n$, the equation of motion becomes particularly simple.

$$\bar{\tau}_n - \bar{\tau}_{n-1} = M_n \frac{\bar{\tau}_{n-1}}{v \bar{\tau}_{n-1} + 1 - M_n} \approx \frac{M_n}{v}$$

where the last approximation applies when κ_{n-1} is sufficiently small (so $\bar{\tau}_{n-1}$ is sufficiently large). If the population size is bounded, then M_n tends toward $(2 - \frac{1}{N^*} \frac{v_c}{v}) \frac{1}{N^*}$, and so $\bar{\tau}_n$ grows roughly as $((2 - \frac{1}{N^*} \frac{v_c}{v}) \frac{1}{N^*}) n$. Alternatively, if the population size is unbounded, then for large n we have the approximation $\bar{\tau}_n - \bar{\tau}_{n-1} \approx \frac{2}{P_n}$, so $\bar{\tau}_n$ eventually grows (and κ_n shrinks) like $2R_n$. However, in this case, R_n always grows more slowly than n . The worst outcome (as the last point of the proposition notes) is when the exit rate of old agents never catches up to the entry rate of new ones – that is, when d_n is bounded below 1. In this case, the size of the observation pool grows linearly in n , so $\bar{\tau}_n$ is on the order of $\sum_{m=1}^n \frac{1}{m} \approx \ln n$.

A more economic intuition for the slowdown in learning can be gleaned from (5), where the dominant term can be written $(\kappa_n - \kappa_{n-1})/\kappa_{n-1} \approx -2\alpha_n/P_n$. That is, the incremental percentage improvement in public information depends on how aggressively arriving agents incorporate new information in their actions (α_n) and on how much that new information is diluted by being mixed into the population average action (P_n^{-1}). Free-riding on the public signal tends to slow learning because α_n declines as κ_n becomes more precise. This is exacerbated when old actions exit slowly, because P_n grows and so the dilution of new information worsens over time.

This is a useful time to remark on the tractability purchased by the uniform exit assumption. Without this assumption, characterizing κ_n would require keeping track of the action and the remaining size of each preceding cohort – that is, we would have to monitor a state variable that increases in dimensionality over time. Under uniform exit, κ_{n-1} is sufficient for κ_n , so the dynamics are recursive and relatively simple.

3.3 The Persistence of Old Information in Case PA

To provide more insight about the causes of slow learning, we develop three measures of a cohort’s influence: the share of its action in the observation pool, the weight of its *signal* in that pool, and the fraction of the error in the observational signal that it is responsible for. One of the keys to slow learning is that the second two types of influence can persist long after most of the actions in that cohort have departed.

The fraction of cohort n ’s observation \bar{x}_{n-1} composed of cohort m actions was introduced earlier; it is λ_m^n . Notice that \bar{x}_{n-1} is a weighted average of cohort actions $\{x_1, x_2, \dots, x_{n-1}\}$, so we can write $\bar{x}_{n-1} = \sum_{m=1}^{n-1} \lambda_m^n x_m$, with $\sum_{m=1}^{n-1} \lambda_m^n = 1$. Next, notice that because each x_n is a linear combination of the average idiosyncratic signal of cohort n and past actions, we can also write \bar{x}_{n-1} as a weighted average of the cohort signals $\{s_1, s_2, \dots, s_{n-1}\}$: $\bar{x}_{n-1} = \sum_{m=1}^{n-1} \pi_m^n s_m$, where the weights π_m^n satisfy $\sum_{m=1}^{n-1} \pi_m^n = 1$. Finally, because the cohort signals are i.i.d. with variance v_c , the expected squared error in the observational signal \bar{x}_{n-1} can be written $\kappa_{n-1} = v_c \sum_{m=1}^{n-1} (\pi_m^n)^2$. Each term in this sum represents the contribution of cohort m information $s_m = \theta + \varepsilon_m$ to the ‘public error’ observed by the agents acting at stage n . The fraction of that public error that can be attributed to cohort m can therefore be written $\chi_m^n \equiv \frac{v_c(\pi_m^n)^2}{\kappa_{n-1}} = (\pi_m^n)^2 / \sum_{m=1}^{n-1} (\pi_m^n)^2$.

Proposition 1 has already begun to characterize how long old actions persist (in terms of the primitive d_n). While this has to do with how fast λ_m^n tends to zero with n , it will be convenient to focus instead on how fast $\ln(1/\lambda_m^n)$ tends to infinity with n .¹⁸ We will say that old actions exit at *log rate* $g(n)$ if $\{\ln(1/\lambda_m^n)\}_{n>m} \rightarrow \infty$ at rate $g(n)$. (So for example, if old actions exit exponentially, then they exit at log rate n .) For the persistence of old *information*, we study the rates at which π_m^n , and χ_m^n tend to zero with n (if they do); call these the decay rates of cohort m information and cohort m relative influence, respectively. If the observational signal were to aggregate private

¹⁸The reason is technical. It will turn out that $\lambda_m^n \rightarrow 0$ at rate $\exp(rR_n + o(R_n))$, for some constant r . The non-leading $o(R_n)$ terms are unwieldy and of no particular interest, since later results depend only on R_n . Focusing on $\ln(\lambda_m^n)$ streamlines the analysis, as these non-leading terms vanish relative to R_n .

information efficiently, then each cohort's signal would be equally weighted, and both π_m^n , and χ_m^n would tend to zero at rate n . To see how action and information persistence actually behave under social learning, consider how π_m^n , and χ_m^n evolve from stage n to $n+1$. In the first stage after cohort m acts, we have $\lambda_m^{m+1} = 1/P_m$ and $\pi_m^{m+1} = \alpha_m/P_m$. Thereafter, the weights evolve recursively according to:

$$\lambda_m^{n+1} = \left(1 - \frac{1}{P_n}\right) \lambda_m^n \quad \text{and} \quad \pi_m^{n+1} = \left(1 - \frac{\alpha_n}{P_n}\right) \pi_m^n \quad \text{for } n > m$$

Immediately after cohort m has acted, we have $\pi_m^{m+1} < \lambda_m^{m+1}$: \mathcal{P}_m contains every cohort m action, but those actions put only weight $\alpha_m < 1$ on the new information s_m .¹⁹ At each subsequent stage though, the impact of m 's action shrinks by a percentage $\frac{1}{P_n}$, but the impact of its information shrinks by the smaller percentage $\frac{\alpha_n}{P_n}$. This reflects the fact that signal s_m becomes impounded in the actions of later cohorts to the extent that those cohorts free-ride on their observational signal. As public information improves, α_n declines (agents free-ride more on the past), and so old information vanishes more and more slowly.

Proposition 3 provides a precise characterization of the relationship between old actions and old information. The proof requires a regularity condition (Poly) for cases in which the observation pool diverges slowly.

(Poly) If $P_n \rightarrow \infty$ and $P_n/R_n \rightarrow 0$, then there is some $b \in (0, 1)$ such that P_n grows at rate n^b .

The main substantive effect of this condition is to exclude from analysis certain cases in which the observation pool grows very slowly (sub-polynomially); in these excluded cases, old actions exit very slightly slower than exponentially.²⁰ The other effect is to impose a type of smoothness on slowly growing pools that is useful in proving parts of the next result.

Proposition 3 *Suppose (Poly) holds and exit is not immediate. The action, information, and relative influence of cohort m decay at the following rates.*

1. (Actions) $\ln \lambda_m^n \rightarrow -\infty$ at rate R_n .
2. (Information)²¹ $\pi_m^n \rightarrow 0$ at rate $(R_n)^\eta$, where $\eta = \left(2 - \frac{v_c}{v} \lim_{n \rightarrow \infty} (1/P_n)\right)^{-1}$.
3. (Relative Influence)
 - (a) If P_n is bounded with $P_n \rightarrow N^*$, then $\chi_m^n \rightarrow 0$ at rate $(R_n)^{\frac{v_c}{2vN^* - v_c}}$.
 - (b) If P_n is unbounded, then χ_m^n converges to a strictly positive constant as $n \rightarrow \infty$.

Immediate exit is excluded, as earlier, for the sake of parsimony. If exit is immediate, then information is aggregated efficiently: $\pi_m^n = \chi_m^n = \frac{1}{n}$. The following corollary is just a simple restatement of part (1).

¹⁹The first cohort is an exception; we have $\lambda_1^2 = \pi_1^2$ since $\alpha_1 = 1$.

²⁰To be precise, exit in these cases satisfies $\lim_{n \rightarrow \infty} \frac{\ln \lambda_m^n}{n} = 0$ and $\lim_{n \rightarrow \infty} \frac{\ln \lambda_m^n}{n^{1-b}} = -\infty$ for any $b > 0$.

²¹Note that $1/P_n$ converges regardless of whether P_n is bounded or diverges, so the limit embedded in the definition of η is well defined.

Corollary 1 *Fix any arbitrary cohort m . If cohort m exits from the observation pool at log rate $g(n)$, then learning is at rate $g(n)$.*

The main results here deserve emphasis. First, learning is tied to the rate at which old actions exit from the observation pool. If old actions exit at least exponentially (that is, at log rate n), then learning is at the full-information rate. If old actions disappear more slowly than this, then learning will be slower than rate n .

Second, the information of a cohort decays more slowly than its action (polynomially rather than exponentially in R_n). To put this in high relief, notice that we can approximate: $\ln(1/\pi_m^n) \approx \eta \ln(\ln(1/\lambda_m^n))$, where higher values of $\ln(1/\pi_m^n)$ and $\ln(1/\lambda_m^n)$ reflect faster departure of cohort m information and actions respectively. This approximation also suggests two channels through which cohort m information becomes particularly persistent if P_n is unbounded (rather than bounded). First, cohort m 's action sticks around longer (that is, $\ln(1/\lambda_m^n)$ is smaller) if P_n is unbounded, and this has a direct effect on the persistence of its information. The second effect is indirect: the persistence of cohort m 's information is greater *relative* to the persistence of its action when P_n is unbounded, because η is smaller. (If P_n is bounded (unbounded), $\eta > (=) \frac{1}{2}$.) This reflects a knock-on effect: the longer that an old action persists, the more time it has to influence later cohorts, and that influence persists after the original action is gone.

Third, if old actions exit slower than exponentially, then their relative influence never disappears! An example makes the point rather forcefully. Suppose the observation pool includes all past actions, so $d_n = 0$ and $P_n = n$.²² Then simulations show that χ_1^n tends to approximately 0.82. That is, as far into the future as one likes, 82% of the squared error in the public signal \bar{x}_n comes from the signal error of the very first cohort.

Fourth, given the equivalences laid out in Proposition 1, one might wonder whether there are other simple barometers of whether learning will be fast or slow. A natural candidate has to do with what we will call the *cumulative persistence* of old actions. Define cohort m 's cumulative persistence to be $\sum_{n=m+1}^{\infty} \lambda_m^n$; this captures the cumulative direct impact of cohort m actions, or equivalently (with an additional normalization) the average duration of a cohort m action in the observation pool. An immediate conjecture is that the boundary between fast and slow learning could be related to whether cumulative persistence is finite or infinite. Proposition 3 makes it easy to test this conjecture. To illustrate, consider three different scenarios in which old actions exit fast ($\lambda_m^n \propto e^{-n}$), slightly slower ($\lambda_m^n \propto e^{-\sqrt{n}}$), or quite slowly ($\lambda_m^n \propto 1/n$). In the first two scenarios, the cumulative persistence of every cohort is finite, while in the third, each cohort's cumulative persistence is infinite. However, social learning is at rate n only in the first scenario, and slower in the second two, so the conjecture cannot hold.²³

Finally, a few technical points should be clarified; a reader who wishes to avoid these details could skip ahead. The fact that none of the decay rates for cohort m depends on m is a consequence

²²For this example, we set $v = v_c = 1$.

²³Section 6.1 provides further insight about why the exit rate of old actions, rather than their cumulative persistence, is the most useful barometer for assessing the speed of learning.

of the uniform exit assumption. The main substantive implication of condition (Poly) is to exclude cases in which P_n diverges at a very slow, sub-polynomial rate; for example, $P_n = \ln n$. In such cases, Proposition 2 still applies – old actions exit slower than exponentially, and learning is slower than rate n – but the other quantities in Proposition 3 are difficult to characterize cleanly.²⁴ Finally, we note that the motivation for the rather long proof of part (2) is to get to the clean, and somewhat startling result in (3b). Settling for the convergence rate of $\ln \pi_m^n$ (similarly to part (1)) would be much simpler but would not provide enough precision to show that (3b) is true.

3.4 The Importance of Free-riding

At this point, we have seen slow learning can arise when old information is persistent and self-interested agents put less weight on their own signals than later generations would like them to. We have also seen (in the bounded P_n case), that the free-riding problem is not sufficient to slow down learning by itself, if old actions are weeded out relatively quickly. However, we have not yet seen whether free-riding is necessary to get slow learning. To answer this, consider a contrived situation in which a social planner has the power to choose the weights $\alpha_n \in [0, 1]$ that cohorts place on their own signals. This planner knows the path of P_n but has no other information or instruments at her disposal; in particular, she cannot control d_n and P_n and she cannot help the agents to communicate about their signals. She chooses a sequence $\{\alpha_n\}_{n \geq 1}$ with the sole objective of reducing the error in the public signal as quickly as possible.²⁵

Proposition 4 *The social planner can attain rate n learning, regardless of the rate at which old actions exit.*

Proof. Set $\alpha_n = \frac{P_n}{n}$. This is feasible (since $P_n \leq n$). A brief inspection of (4) shows that this policy attains $\kappa_n = \frac{v_c}{n}$ and $\bar{\tau}_n = n/v_c$ for all $n \geq 1$. ■

Proposition 4 makes it clear that both elements – old information that is sufficiently persistent, and agents who ignore the informational spillover generated by their actions – are required to get slow social learning.

4 Observing the Average of a Finite Sample (SA)

In this section we develop the sample average (SA) version of the model. The most surprising results revolve around the fact that there are essentially no benefits from observing larger samples – learning is no faster, and agents’ utility is no higher. This runs counter to the usual intuition that larger samples help to average out idiosyncratic mistakes. The usual intuition is not wrong *per se*,

²⁴The main reason for interest in these intermediate, sub-polynomial P_n cases would be to test the conjecture that learning is slower than rate n if and only if the relative influence of older cohorts does not decay. This would be true if (3b) could be extended to all unbounded P_n . However, one can find examples ($P_n = \ln n$ is one) for which this conjecture is false: learning is slower than rate n , but χ_m^n tends to 0 with n . Further details are available on request.

²⁵Since the result below has to do with feasibility rather than optimality, this objective function appears only for motivational purposes. If we were interested in deriving optimal $\{\alpha_n\}_{n \geq 1}$ policies, this objective would need to be stated with substantially more precision.

but social learning leads people to make highly correlated mistakes, and it is this common portion of their errors that comes to dominate the speed of learning. This helps to explain the fact that we see identical rates of learning in models SA and PA, even though the ‘sample’ in the latter case is the entire observation pool.

Let the observation pool evolve just as before, and let $x_{\mathcal{P}_{n-1}}$ refer to an arbitrary individual action belonging to \mathcal{P}_{n-1} . A typical member n_i of cohort n receives her idiosyncratic signal s_{n_i} as before, plus an observational signal $x_{S(n_i)}$. This observational signal is the mean of $S \geq 1$ individual actions $x_{\mathcal{P}_{n-1}}$ drawn randomly from \mathcal{P}_{n-1} . This is all that agent n_i knows about her observational signal; she does not observe the individual actions that comprise $x_{S(n_i)}$, nor does she observe the ages of those individual actions.²⁶ Of course, she will be able to draw certain inferences about the likely composition of $x_{S(n_i)}$ based on her knowledge about how \mathcal{P}_{n-1} evolves.

Because the members of a cohort will now receive different observational signals, some new notation will be required. Let $\hat{\kappa}_{n-1} = E\left((x_{\mathcal{P}_{n-1}} - \theta)^2\right)$ be a cohort n agent’s expectation of the squared error in an individual action drawn randomly from \mathcal{P}_{n-1} , and let $\hat{\rho}_{n-1} = E\left((x_{\mathcal{P}_{n-1}} - \theta)(x'_{\mathcal{P}_{n-1}} - \theta)\right)$ be her expectation of the covariance in the errors of two randomly drawn actions. Together, $\hat{\kappa}_{n-1}$ and $\hat{\rho}_{n-1}$ determine the precision of a cohort n agent’s observational signal, and consequently, her action and expected utility.

Lemma 1 *The expected squared error of a cohort n observational signal is*

$$\bar{\kappa}_{n-1} \equiv E\left((x_{S(n_i)} - \theta)^2\right) = \frac{1}{S}\hat{\kappa}_{n-1} + \frac{S-1}{S}\hat{\rho}_{n-1}.$$

The optimal action of a cohort n agent n_i with observational signal $x_{S(n_i)}$ and idiosyncratic signal s_{n_i} is $x_{n_i} = \hat{\alpha}_n s_{n_i} + (1 - \hat{\alpha}_n) x_{S(n_i)}$, where $\hat{\alpha}_n = \frac{\bar{\kappa}_{n-1}}{\bar{\kappa}_{n-1} + v}$. This agent’s expected utility is $-\hat{\alpha}_n v$.

We will work with $\hat{\kappa}_n$ as our measure of learning, but we shall see that $\bar{\kappa}_n$ and $\hat{\alpha}_n$ behave similarly. Notice that $-\hat{\kappa}_n$ is also the expected utility of an agent drawn randomly from \mathcal{P}_n , so we will sometimes refer to it as the mean utility in the observation pool. The expression for $\bar{\kappa}_{n-1}$ makes it clear that an increase in her sample size S only benefits an agent to the extent that the error covariance $\hat{\rho}_{n-1}$ among the sampled actions is low.

Also notice that in order to characterize the quantities of interest in the model, it is not necessary to describe the sampling history – the list specifying the identities of the agents whose actions appear in each sample. This sampling history is not a part of agents’ information sets, and neither their expectations $\hat{\kappa}_{n-1}$ and $\hat{\rho}_{n-1}$, nor their choice variables $\hat{\alpha}_n$ can condition on it. In contrast with the common observational signal of model PA, now some agents will be luckier with their sampling than others. The luckiest will sample recent actors who themselves sampled recent actors, and so on, while the unluckiest will sample very old actions. But because an agent does not know that she has been lucky or unlucky, this has no effect on behavior. Luck in sampling does have distributional

²⁶The case of $K = 1$ is an exception: if agent n_i ’s sample includes a single action, then of course $x_{S(n_i)}$ trivially reveals that action.

implications: some people within a cohort will be (unwittingly) better informed than average, and thus better off than the average utility of $-\hat{\alpha}_n v$, while others will be worse informed and worse off.²⁷ While it would be interesting to pursue these implications, we will not do so here.²⁸

The evolution of $\hat{\kappa}_n$ differs from that of κ_n in model PA (see (4)) because now we are dealing with the mean squared error of individual actions in \mathcal{P}_n rather than the squared error of the mean action in that pool. As the observation pool adjusts from \mathcal{P}_{n-1} to \mathcal{P}_n , a measure d_n of agents depart and the full measure 1 of cohort n agents are added. Because the departing agents are determined randomly, the expected squared error among the remaining $P_{n-1} - d_n$ agents remains equal to $\hat{\kappa}_{n-1}$. The expected squared error for entering agents is $\hat{\alpha}_n v$, so $\hat{\kappa}_n$ is simply:

$$\hat{\kappa}_n = \left(1 - \frac{1}{P_n}\right) \hat{\kappa}_{n-1} + \frac{1}{P_n} (\hat{\alpha}_n v) \quad , \quad n > 1 \quad (6)$$

The error covariance among sampled actions affects $\hat{\kappa}_n$ through the error $\hat{\alpha}_n v$ of the new arrivals: $\hat{\alpha}_n$ depends on squared error in their observational signal $\bar{\kappa}_{n-1}$, which in turn depends on $\hat{\rho}_{n-1}$. In order to develop a recursive characterization of that covariance, let us call a sampled action $x_{\mathcal{P}_n}$ ‘new’ if it occurred at the most recent stage n , or ‘old’ if it occurred earlier. Let $\hat{\rho}_n|_{new,new}$ be the conditional error covariance of two sampled actions that happen to be new:

$$\hat{\rho}_n|_{new,new} = E((x_{\mathcal{P}_n} - \theta)(x'_{\mathcal{P}_n} - \theta) \mid x_{\mathcal{P}_n} \text{ and } x'_{\mathcal{P}_n} \text{ are new})$$

and define $\hat{\rho}_n|_{new,old}$ and $\hat{\rho}_n|_{old,old}$ analogously. An action drawn from \mathcal{P}_n is new with probability $1/P_n$, so the unconditional error covariance of two action drawn from \mathcal{P}_n may be written

$$\hat{\rho}_n = \left(1 - \frac{1}{P_n}\right)^2 \hat{\rho}_n|_{old,old} + \frac{2}{P_n} \left(1 - \frac{1}{P_n}\right) \hat{\rho}_n|_{new,old} + \frac{1}{P_n^2} \hat{\rho}_n|_{new,new}$$

Next we address these conditional covariance terms, one by one. The first one, $\hat{\rho}_n|_{old,old}$, is simply $\hat{\rho}_{n-1}$, the error covariance from one stage ago. For the second term, suppose that $x_{\mathcal{P}_n}$ is the new action drawn from stage n (while $x'_{\mathcal{P}_n}$ occurred prior to stage n). Then $x_{\mathcal{P}_n}$ is a weighted average of a new idiosyncratic signal (with weight $\hat{\alpha}_n$) that is uncorrelated with earlier errors, and S actions randomly drawn from \mathcal{P}_{n-1} . Call these old actions $y_{old_1}, \dots, y_{old_S}$. Then the covariance of $x_{\mathcal{P}_n}$ and $x'_{\mathcal{P}_n}$ is

$$E((x_{\mathcal{P}_n} - \theta)(x'_{\mathcal{P}_n} - \theta) \mid x_{\mathcal{P}_n} \text{ is new, } x'_{\mathcal{P}_n} \text{ is old}) = \frac{1}{S} (1 - \hat{\alpha}_n) \sum_{s=1}^S E((x'_{\mathcal{P}_n} - \theta)(y_{old_s} - \theta)) .$$

²⁷Of course, both here and in model PA, there are also differences in outcomes due simply to idiosyncratic signals.

²⁸In particular, one might like to know about divergence: do the luckiest (unluckiest) agents learn at a faster (slower) rate than average? On this question, the answer appears to be a qualified no. The luckiest samplers consistently rely on their private signals too much (since they do not realize how accurate their observations are). A casual analysis (details on request) indicates that this limits their precision to around double that of an average agent. The very worst off agents make no progress over time (since sampling only from cohort 1 is always a possibility). However, no positive fraction of a cohort can have errors too much worse than average, lest they drag up the average.

But each expectation in the summation is just $\hat{\rho}_n|_{old,old}$, so we have $\hat{\rho}_n|_{new,old} = (1 - \hat{\alpha}_n) \hat{\rho}_n|_{old,old}$, or $\hat{\rho}_n|_{new,old} = (1 - \hat{\alpha}_n) \hat{\rho}_{n-1}$. Finally, for the case of two new actions, note that $x_{\mathcal{P}_n}$ and $x'_{\mathcal{P}_n}$ will both put weight $\hat{\alpha}_n$ on the cohort n error ε_n , while their observational signals have error covariance $\hat{\rho}_n|_{old,old}$, so we arrive at $\hat{\rho}_n|_{new,new} = (1 - \hat{\alpha}_n)^2 \hat{\rho}_{n-1} + \hat{\alpha}_n^2 v_c$. Combining these pieces, we have

$$\hat{\rho}_n = \left(1 - \frac{\hat{\alpha}_n}{P_n}\right)^2 \hat{\rho}_{n-1} + \left(\frac{\hat{\alpha}_n}{P_n}\right)^2 v_c \quad , \quad n > 1 \quad (7)$$

For completeness, note also that $\hat{\kappa}_1 = v$ and $\hat{\rho}_1 = v_c$.²⁹ Just as $\hat{\kappa}_n$ depends on $\hat{\rho}_{n-1}$ through the weight $\hat{\alpha}_{n-1}$, so does $\hat{\rho}_n$ depend on $\hat{\kappa}_{n-1}$ through $\hat{\alpha}_n$. Equations (6) and (7), along with $\hat{\alpha}_n$, fully characterize agents' behavior, and determine the rate of social learning with finite samples. The special case in which agents sample single actions ($S = 1$) is a bit simpler, as the evolution of $\hat{\kappa}_n$ is characterized by (6) alone.

As a first step, we ask how correlated the errors in sampled actions are. This correlation coefficient is

$$\frac{\hat{\rho}_n}{\hat{\kappa}_n} = E \left((x_{\mathcal{P}_n} - \theta) (x'_{\mathcal{P}_n} - \theta) \right) / \sqrt{E \left((x_{\mathcal{P}_n} - \theta)^2 \right) E \left((x'_{\mathcal{P}_n} - \theta)^2 \right)}.$$

The benefit of observing a larger sample is greatest when $\frac{\hat{\rho}_n}{\hat{\kappa}_n}$ is relatively small. Unfortunately for agents, the errors in their predecessors' actions become highly correlated very quickly.

Proposition 5 *If $S > 1$, then the error correlation of two actions drawn from \mathcal{P}_{n-1} tends to one: $\frac{\hat{\rho}_n}{\hat{\kappa}_n} \rightarrow 1$.*

We postpone discussing this result until after presenting rates of learning. The proof leverages the fact that the size of decreases $\hat{\kappa}_n - \hat{\kappa}_{n-1}$ in average squared error can be shown to be roughly proportional to $\hat{\kappa}_{n-1} - \hat{\rho}_{n-1}$. Meanwhile, decreases $\hat{\rho}_n - \hat{\rho}_{n-1}$ in the covariance are roughly proportional to $\hat{\rho}_{n-1}^2$. The decline that is first order in the variables happens fast, bringing $\hat{\kappa}_n$ close to $\hat{\rho}_n$ before the latter gets close to zero.

As a consequence of Proposition 5, the speed of social learning is governed by the rate at which the covariance of action errors is driven to zero. Consulting (7), we see an expression that is almost identical to (4), but with κ_n replaced by $\hat{\rho}_n$ and α_n replaced by $\hat{\alpha}_n$. When we observe that $\hat{\alpha}_n$ tends toward $\frac{\hat{\rho}_{n-1}}{\hat{\rho}_{n-1} + v}$ (because of Proposition 5), the resemblance becomes exact, giving us the following result.

Proposition 6 *(Rates of social learning in the sample average model)*

i) *If $S = 1$, then $\hat{\kappa}_n R_n \rightarrow v$.*

ii) *If $S > 1$, then $\hat{\kappa}_n R_n$, $\hat{\rho}_n R_n$, and $v \hat{\alpha}_n R_n$ all tend to $v\eta$, where $\eta = \left(2 - \frac{v_c}{v} \lim_{n \rightarrow \infty} \frac{1}{P_n}\right)^{-1}$.*

Thus, social learning is always at rate R_n .

²⁹So in fact, if we were willing to ignore the fact that $\hat{\kappa}_0$ and $\hat{\rho}_0$ are not defined, equations (6) and (7) actually apply at $n = 1$ as well.

Let us briefly summarize the results for the case ($P_n = n$) in which old actions never exit from the pool that later agents sample. If agents sample one predecessor each, without knowing when he acted or whom he observed, learning is at rate $\ln n$. If instead, each agent sees the average of two samples, the rate at which (expected squared) errors disappear remains $\ln n$, but the *level* of those errors is cut in half. If the size of the sample increases further, neither the rate of learning nor the level of errors improves. Furthermore, returning to our baseline model with each agent observing the average action in the entire population yields no further changes in the rate of learning, or the level of errors.

Why does the size of an agent’s sample make so little difference? It may be helpful (albeit slightly imprecise) to think of splitting an agent’s expected error $\hat{\kappa}_n$ into mistakes $\hat{\kappa}_n - \hat{\rho}_n$ that are idiosyncratic to her, and mistakes $\hat{\rho}_n$ that are common to other agents. A larger sample *does* help to drive the idiosyncratic portion $\hat{\kappa}_n - \hat{\rho}_n$ down to zero quickly.³⁰ However, the errors that remain are highly correlated across individuals, so larger samples do little to eradicate them.³¹

5 Extensions

In this section we add several new features to the baseline population average learning model, including heterogeneous preferences, a self-reliance bias, and drift in the optimal action θ .³² The first two features will imply, counter to our earlier results, that having older actions linger in the observation pool can be beneficial for the speed of learning. In both cases, the reason is that an agent no longer wishes to rely solely on her most recent predecessors for guidance because (from her point of view) they have over-relied on their idiosyncratic information. In contrast, if θ is a moving target, the disadvantages of relying on older information become more severe.

5.1 Heterogeneity

We return to the framework of the PA model, but now we assume that different individuals have different optimal decisions. We model this by assuming that agent n_i chooses x_{n_i} to minimize $E((x_{n_i} - \theta_{n_i})^2 | \mathcal{I}_{n_i})$, where $\theta_{n_i} = \theta + \mu_{n_i}$. The new term μ_{n_i} represents a shifter that is specific to the individual agent. These taste shocks are distributed i.i.d. (across cohorts and individuals) with $\mu_{n_i} \sim N(\theta, v_{ta})$, and their variance v_{ta} is common knowledge among agents.

As before, an agent’s information set \mathcal{I}_{n_i} always contains an observational signal equal to the average action in \mathcal{P}_{n-1} . We consider two different specifications for the agent’s idiosyncratic information. In the *known private taste* version of the model, agent n_i observes the pair of signals $\{s_{n_i}, \mu_{n_i}\}$, where $s_{n_i} = \theta + \varepsilon_n + \varepsilon_{n_i}$ as earlier. In other words, she sees a signal about the action

³⁰We do not demonstrate this formally, but expressions for the rate of change of $\hat{\kappa}_n - \hat{\rho}_n$ suggest that it declines faster if $\frac{S-1}{S}$ is large.

³¹It seems plausible that larger samples might have benefits if there were additional sources of idiosyncratic error in the model. For example, heterogeneous tastes (as developed for model PA in Section 5.1) generate a persistent source of idiosyncratic variation in actions that an observer could minimize with a large sample.

³²These extensions could be applied to model SA as well, but analysis would be a bit more tedious because of the extra state variable.

θ that is best for a typical agent, and she also knows which direction (and how much) to adjust that action to better suit herself. In the *shrouded private taste* version of the model, agent n_i observes the idiosyncratic signal $\tilde{s}_{n_i} = \theta_{n_i} + \varepsilon_n + \varepsilon_{n_i}$. The terms ε_n and ε_{n_i} represent cohort and individual-specific errors as before, with the added proviso that these errors are independent of the taste shocks.³³ Thus, the agent has a noisy signal about her personal optimal action θ_{n_i} , but she does not observe μ_{n_i} directly, so she does not know how different she is from a ‘typical’ agent. We include both specifications because both are plausible, and they have quite different implications for learning. In our years of schooling example, the taste shifter could reflect relative ability – perhaps the optimal amount of schooling is higher for smarter students. In this case, it may be reasonable to think that a student knows her taste shock μ_{n_i} – say by comparing herself to her classmates – but still needs to learn about the average costs and returns to schooling (as summarized by θ) before making a decision. In this case, the known taste model applies. Alternatively, suppose in our irrigation example that a farmer has information about the past returns to irrigating his own land – call this \tilde{s}_{n_i} – but does not know whether his land is predisposed (for geological reasons, perhaps) to need more or less water than the average plot. While his idiosyncratic information may *suggest* that he should use more or less water than other farmers have done (i.e., if $\tilde{s}_{n_i} - \bar{x}_{n-1}$ is positive or negative), he cannot be certain how much of that difference $\tilde{s}_{n_i} - \bar{x}_{n-1}$ represents private value and how much is just noise. In this case, the shrouded taste model would be appropriate.³⁴

5.1.1 Known Private Tastes

If agents can observe their relative private value shifters directly, the analysis remains very close to the original PA model. It is straightforward to see that the optimal action for agent n_i is to form her best estimate of θ given signals s_{n_i} and \bar{x}_{n-1} , and then add μ_{n_i} to get her optimal action. Thus,

$$x_{n_i} = \alpha_n s_{n_i} + (1 - \alpha_n) \bar{x}_{n-1} + \mu_{n_i}, \quad \text{where } \alpha_n = \frac{\kappa_{n-1}}{\kappa_{n-1} + v}$$

As in the shrouded private values case, the mean action of cohort n is $x_n = \alpha_n s_n + (1 - \alpha_n) \bar{x}_{n-1} = \alpha_n (\theta + \varepsilon_n) + (1 - \alpha_n) \bar{x}_{n-1}$. The evolution of \bar{x}_n and κ_n follow the same equations as in model PA. Because α_n is also chosen just as in model PA, the following corollary is immediate.

Corollary 2 *In the model with known private values, social learning about θ takes place at the same rate that applies in the baseline model PA.*

5.1.2 Shrouded Private Tastes

Other than the stated changes, the model is identical to our baseline model PA. Notice that there is no prospect for complete learning about *individual* targets, in the sense of $E \left((x_{n_i} - \theta_{n_i})^2 \right)$ tending to zero over time, because an agent’s only information about her taste shock comes from her private

³³Also, we extend assumption (1) to the taste shocks.

³⁴One can easily imagine intermediate cases that we will not pursue here. For example, rather than observe μ_{n_i} separately and perfectly, an agent could observe it separately but with noise.

signal. Instead, we focus on how quickly the true value of θ , the common part of the optimal action, is revealed. As earlier, let \bar{x}_n be the average action in the observation pool \mathcal{P}_{n-1} . We retain $\kappa_n = E((\bar{x}_n - \theta)^2)$ as our measure of learning.

We begin by observing that the average action \bar{x}_{n-1} is a less precise signal about agent n_i 's personal optimal action when tastes are very heterogeneous; the larger v_{ta} , the more weight she will want to shift onto her private signal in forming her action. To derive her optimal action, consider the errors in her two signals: $\bar{x}_{n-1} - \theta_{n_i} = (\bar{x}_{n-1} - \theta) - \mu_{n_i}$ and $\tilde{s}_{n_i} - \theta_{n_i} = \varepsilon_n + \varepsilon_{n_i}$. These errors are independent of each other, and moreover, the two components of the observational signal error – the deviation of \bar{x}_{n-1} from θ , and the deviation of θ from θ_{n_i} – are also independent. Thus the errors in \bar{x}_{n-1} and \tilde{s}_{n_i} have variances of $\kappa_{n-1} + v_{ta}$ and v respectively. As before, agent n_i 's optimal action is a precision-weighted average of these two signals.³⁵

$$x_{n_i} = \alpha_n \tilde{s}_{n_i} + (1 - \alpha_n) \bar{x}_{n-1}, \quad \text{where } \alpha_n = \frac{\kappa_{n-1} + v_{ta}}{\kappa_{n-1} + v_{ta} + v}$$

Notice that the own signal weight is bounded away from zero, $\alpha_n > v_{ta}/(v_{ta} + v)$, regardless of how accurately observational information reveals θ . Of course, this simply reflects the wedge created by the taste shock.

Let \tilde{s}_n be the mean of the cohort n private signals. The taste shocks wash out of this mean by assumption (1), so we have $\tilde{s}_n = \theta + \varepsilon_n$, and the mean action over cohort n is still given by $x_n = \alpha_n \tilde{s}_n + (1 - \alpha_n) \bar{x}_{n-1}$. The recursive expression (3) still characterizes the updated average action \bar{x}_n , and so the evolution of κ_n is still governed by

$$\kappa_n = \left(1 - \frac{\alpha_n}{P_n}\right)^2 \kappa_{n-1} + \left(\frac{\alpha_n}{P_n}\right)^2 v_c \quad (8)$$

Thus any differences in the speed of learning with shrouded private tastes (relative to model PA) can be attributed the fact that agents' "self-reliance" is bounded away from zero. Define $\bar{\alpha} = \frac{v_{ta}}{v_{ta} + v}$, the limiting weight placed on one's own signal under the assumption that $\kappa_n \rightarrow 0$, and $\bar{d} = \lim_{n \rightarrow \infty} d_n$. It is useful to define a quantity $\beta = \frac{2\bar{\alpha}}{1-\bar{d}}$ which is small when old actions are very persistent and tastes are similar, and large when tastes are quite different and old actions exit quickly.

Proposition 7 *Rates of social learning in the model with shrouded private tastes are as summarized in Table 1.*

The awkward statement of the final case – roughly, learning is neither faster nor slower than rate n^β – could be tightened at the cost of additional technical assumptions.³⁶ To summarize informally, greater persistence of old actions in the observation pool has a non-monotonic effect on the rate of learning. Learning is fastest (rate n) if the mass of departing old actions is smaller than the mass

³⁵ As earlier, we write α_n rather than α_{n_i} because the optimal weight is the same for all cohort n agents.

³⁶ An examination of the proof suggests that sufficiently quick convergence of d_n to \bar{d} (rate n^a for any $a > 0$ should do) suffices for the convergence of $n^\beta \kappa_n$.

Growth of P_n	Learning
(1) Bounded	Incomplete: κ_n converges to a strictly positive constant
(2) Unbounded, with $\frac{P_n}{n} \rightarrow 0$	$\kappa_n \rightarrow 0$ at rate P_n
(3) Unbounded, with $\frac{P_n}{n} \rightarrow 1 - \bar{d} \in (0, 2\bar{\alpha})$	$\kappa_n \rightarrow 0$ at rate n (= rate P_n)
(4) Unbounded, with $\frac{P_n}{n} \rightarrow 1 - \bar{d} > 2\bar{\alpha}$	$\kappa_n \rightarrow 0$ faster than rate $n^{\beta-\epsilon}$ and slower than $n^{\beta+\epsilon}$ (for any $\epsilon > 0$)

Table 1: Rates of learning in the shrouded private taste model. Reading from top to bottom, each case represents progressively slower exit of old actions. Case (4) may be read as: $n^{\beta-\epsilon}\kappa_n \rightarrow 0$ and $n^{\beta+\epsilon}\kappa_n \rightarrow \infty$, for any $\epsilon > 0$.

of new actions ($\bar{d} < 1$), but not too much smaller ($\bar{d} > 1 - 2\bar{\alpha}$). If old actions are more or less persistent than this, then learning slows down.

The reason for the change in results (relative to Proposition 2) is that private values induce agents to put more weight on their own information. As we saw in Proposition 4, increasing α_n can counteract the tendency of old actions to retard learning, but if α_n rises too much, a new problem arises – the observation pool can tilt too far toward the most recent information. If β is small, then excessive weight on old information remains the bottleneck in learning, while if $\beta = \infty$ (because $d_n \rightarrow 1$), excess weight on recent information becomes the bottleneck. When the persistence of old actions is just balanced by the ‘self reliance’ induced by heterogeneity, learning is fast.

Thus the speed of learning with heterogeneity depends critically on whether agents are able to separate information that pertains to them specifically (μ_{n_i}) from information about θ that is relevant to all agents. If this separation is possible, then they will respond to the individual-specific information while continuing to underweight (from a social point of view) their information about θ . If they cannot separate individual-specific from general information about the decision, then in their desire to respond to the former, their actions will as a side-effect bring more new information about θ along for the ride. Because there is underlying bias toward old information (for the reasons discussed in earlier sections), either private value scenario could lead to faster learning than the other, depending on how severe that underlying bias toward old information is.³⁷

5.2 Self Reliance Bias

Informally, one could interpret the weight α_n that an agent places on her idiosyncratic signal as “self reliance.” In the standard version of the model, self-reliance tends to zero as the public signal grows more precise, whereas in the shrouded private taste model, self-reliance is (rationally) bounded away from zero no matter how precise public information becomes. Experimental studies of herding and social learning often find that subjects are more self reliant than theory would predict; we will refer to this as “self reliance bias.” The clearest evidence for this bias comes from models of

³⁷Munshi [15] tests for a very similar issue in agricultural data from India, and finds results that are roughly consistent with ours. He finds weaker social learning about a crop (rice) with yields that are sensitive to farmer characteristics (like the soil type in a field) that are hard for other farmers to observe. Conversely, social learning is stronger for wheat, where individual heterogeneity is less important. In his setting, learning is largely from recent planting decisions (as well as yield outcomes), so the confounding effect of old actions seems likely to be small.

herding with binary actions and signals; in a meta-study pooling data from 13 different experiments, Weizsäcker [20] finds that when subjects’ private and observational signals conflict, they act on their private signals unless the posterior odds are at least 2:1 against being correct.³⁸ In the setting of our model, one can imagine a number of plausible reasons that people might exhibit a self reliance bias, including the following: (1) Subjects might overestimate the precision of their own private information. (2) Subjects might have doubts about the competence of their predecessors. (3) A subject might suspect that her predecessors’ actions reflect a different objective than her own. (4) Properly evaluating observational data requires a chain of relatively sophisticated inferences, while evaluating one’s private information is straightforward. Subjects may have more confidence in their *own* ability to handle the latter task without making mistakes.

To illustrate the implications of a self reliance bias as transparently as possible, we introduce the following assumption in our baseline PA model.³⁹

SRB For all $n > 1$, every agent places the same weight $\alpha > 0$ on her own signal.

While the main attraction of SRB is simplicity, it could be motivated as something akin to a stationarity requirement on strategies. In some situations it may be realistic to assume that an agent does not have a clear sense of how long learning has been going on prior to her arrival; in this case, her strategy should not condition on her cohort number. Under assumption SRB, the squared error in the population average action κ_n still evolves according to (4), but individual actions do not converge to θ . Agent n_i ’s action never gets closer to θ than an expected squared error of $\alpha^2 v$, regardless of n , due to the persistent introduction of new idiosyncratic error. When we frame learning in terms of κ_n below, the interpretation should be that this reflects how accurately an agent *could* predict θ after stage n if she were willing to rely on \bar{x}_n exclusively. Define $\beta = \frac{2\alpha}{1-d}$ similarly to the previous section.

Proposition 8 *Under assumption SRB, the rate of learning is essentially the same as in the shrouded private taste model. That is, if P_n is bounded, κ_n tends to a constant. If P_n is unbounded, then $\kappa_n \rightarrow 0$ at the rate specified for β in Proposition 7.*

The proof is omitted. It follows exactly the same lines as the proof of Proposition 7 but is simpler because in this case, α_n reaches its limiting value immediately. The logic is essentially the same as well – because a self reliance bias tilts the observation pool toward recent information, it can partially correct (or overcorrect) the underlying bias toward old information. Thus, this bias can be socially useful, even though from an individual’s point of view it represents a mistake.

³⁸Because the likelihood ratio for a single private signal was also 2:1 (that is, $\Pr(\text{correct signal}) / \Pr(\text{incorrect signal}) = 2$) in most of these studies, one interpretation is that subjects acted as though they had a second independent private signal confirming the first one. One of the constituent studies, Goeree et al. [13], also finds that subjects roughly double-count their own signals in the context of a structural quantal response model of behavior.

³⁹A natural alternative would be to assume that an agent overestimates the precision of her private signal. If each agent acts as though her signal has precision kv^{-1} instead of v^{-1} , with $k > 1$, then self-reliance will be roughly k times as large as it should be, for large n , but it will still tend to zero over time. It is straightforward to see that this will not change the analysis in Section 3 appreciably. The *level* of error in the public signal will be lower than if agents behaved rationally, but the rate of learning will be the same.

5.3 Moving Target

In this section, we start from the baseline PA model, but now we assume that the optimal action drifts over time. Instead of a static target θ , suppose that a cohort n agent n_i tries to minimize $E((x_{n_i} - \theta_n)^2 | \mathcal{I}_{n_i})$, where θ_n follows a random walk: $\theta_n = \theta_{n-1} + v_n$, where the increments are i.i.d $v_n \sim N(0, v_{mt})$. (Let θ_n take initial value θ_1 .) If the environment changes quickly relative to the rate at which agents act, this model may be more apt than our standard model. As one might expect, a moving target makes the persistence of older actions in the observation pool even more disadvantageous for learning than in our baseline model.

Agent n_i 's idiosyncratic signal is now $s_{n_i} = \theta_n + \varepsilon_n + \varepsilon_{n_i}$ (with ε_n and ε_{n_i} distributed as earlier) – thus cohort n 's signals embed information (not replicated in any earlier actions) about the most recent change in the target v_n . As earlier, \bar{x}_{n-1} is the average action in the observation pool. It is somewhat arbitrary whether to define the error in \bar{x}_{n-1} with respect to θ_{n-1} , its contemporary optimal action, or with respect to θ_n , the target that \bar{x}_{n-1} will be used to predict. We choose the former, defining $\kappa_{n-1}^{mt} = E\left((\bar{x}_{n-1} - \theta_{n-1})^2\right)$. Because $\bar{x}_{n-1} - \theta_n = (\bar{x}_{n-1} - \theta_{n-1}) - v_n$, it follows that the squared error of \bar{x}_{n-1} with respect to θ_n is $E\left((\bar{x}_{n-1} - \theta_n)^2\right) = \kappa_{n-1}^{mt} + v_{mt}$. For an agent in cohort n , the errors $(s_{n_i} - \theta_n)$ and $(\bar{x}_{n-1} - \theta_n)$ are independent, so her optimal estimate of θ_n puts weight $\alpha_n = \frac{\kappa_{n-1}^{mt} + v_{mt}}{\kappa_{n-1}^{mt} + v_{mt} + v}$ on her idiosyncratic signal. As before, the mean action of cohort n agents is $x_n = \alpha_n s_n + (1 - \alpha_n) \bar{x}_{n-1}$ (with $s_n = \theta_n + \varepsilon_n$), and the mean action in the observation pool updates according to $\bar{x}_n = (1 - \alpha_n/P_n) \bar{x}_{n-1} + (\alpha_n/P_n) s_n$. Together these imply that the error in the observation pool updates according to

$$\begin{aligned} \kappa_n^{mt} &= \left(1 - \frac{\alpha_n}{P_n}\right)^2 (\kappa_{n-1} + v_{mt}) + \left(\frac{\alpha_n}{P_n}\right)^2 v_c \quad , \quad \text{or equivalently,} \\ \gamma_n &= \left(1 - \frac{\alpha_n}{P_n}\right)^2 \gamma_{n-1} + \left(\frac{\alpha_n}{P_n}\right)^2 v_c + v_{mt} \quad , \quad \text{with } \alpha_n = \frac{\gamma_n}{\gamma_n + v} \end{aligned}$$

where we define $\gamma_n = \kappa_n + v_{mt} = E\left((\bar{x}_{n-1} - \theta_n)^2\right)$.

Clearly, there is no hope of complete learning in this setting; given the persistent change in the target, observing past actions \bar{x}_{n-1} can never shed light on the most recent change v_n in the optimal action. At best, γ_n could converge to v_{mt} , in which case the value of observing past actions (to newly arriving agents) is eventually stationary. At worst, γ_n could diverge, in which case newly arriving agents will eventually give up on observational information entirely and rely exclusively on their private signals ($\alpha_n \rightarrow 1$).

Proposition 9 *If $P_n \rightarrow N^* > 1$, then $\gamma_n \rightarrow \gamma_{N^*}^* > v_{mt}$. If P_n is unbounded, then γ_n diverges.*

As expected, stale information in the observation pool makes it substantially harder to learn about a moving target. If old actions die off at least exponentially (that is, $\ln \lambda_m^n \propto -n$ as in the bounded population case), then paying attention to past actions is always worthwhile: in the long run, agents put a weight of $\frac{v}{\gamma_{N^*}^* + v}$ on ‘precedent,’ and their expected utility tends toward $-\frac{\gamma_{N^*}^*}{\gamma_{N^*}^* + v} v$.

If old actions die off more slowly than this (P_n unbounded) then the current target drifts away from the observation pool over time, and in the long run, agents ignore precedent, use only their idiosyncratic information, and earn payoffs tending toward $-v$. The rate at which the environment is changing, as summarized by v_{mt} , also affects learning, but it plays a secondary role. The limit $\gamma_{N^*}^*$ is increasing in v_{mt} , so if old actions are purged sufficiently fast, then slower drift in the target implies that agents can put more weight on precedent and earn higher long run utility than they would if v_{mt} were larger. However, if old actions are more persistent, *any* drift in the target drives their long run payoffs down to $-v$.⁴⁰

One could interpret the heterogeneity and moving target extensions in this section as representing, respectively, transient and persistent shocks to the optimal decision. One natural intermediate case, which we will not treat, arises when the optimal decision varies according to an autoregressive process. The challenge in this case (and the reason that we do not pursue it) is that there does not appear to be a way to avoid representing the equations of motion with a set of state variables that grows with n .

6 Non-uniform Exit of Old Actions

While our uniform exit condition has proven to be convenient, there is no compelling reason that the persistence of old actions could not take some other form. In this section, we drop this condition in two examples and show that the principle of slow learning is robust. Without uniform exit, a curse of dimensionality arises: one must keep track of the composition of all past observation pools, and in general this would require a state variable whose size grows with n .⁴¹ Our examples focus on the situation in which an agent observes mainly the preceding cohort’s action, tainted by some earlier actions. We avoid the curse of dimensionality by looking at cases in which almost all past cohorts are treated identically, which allows learning to be described by only two or three state variables of fixed size. In the first example, the first ‘founding’ cohort’s action is the only confounder; other old actions are not persistent. In the second example, almost all old actions disappear immediately, but a small fraction of them remain observable indefinitely. The examples make it clearer that it is not uniform exit that drives slow learning – a vanishing but non-zero chance that an action is observed indefinitely far into the future can suffice.

6.1 Example: Only One Action Persists

In some settings, the very first decisions made may enjoy a special prominence. For example, early adopters of a new medical technique may receive a burst of media attention simply for being first. Then an agent facing the same decision later on may be disproportionately exposed to these first

⁴⁰We also have $\gamma_{N^*}^* \rightarrow \infty$ as $N^* \rightarrow \infty$, so there is a smooth transition between the bounded and unbounded population cases. Both of these comparative static conclusions are derived implicitly from the cubic equation that defines $\gamma_{N^*}^*$. (See the proof of Proposition 9). A closed form for $\gamma_{N^*}^*$ is available, but unilluminating.

⁴¹At stage n , one would need the composition of \mathcal{P}_{n-1} (an $n-1$ element vector), of \mathcal{P}_{n-2} (an $n-2$ element vector), and so forth.

actions if (for example) she carries out an internet search. In this section we model this idea by assuming that only the action of cohort 1 persists, while other old actions disappear immediately. Specifically, we assume that cohort n 's observational signal blends the most recent action x_{n-1} and the first action x_1 :

$$\bar{x}_{n-1} = (1 - w_{n-1})x_{n-1} + w_{n-1}x_1$$

where $\{w_n\}$ is an exogenous sequence of weights (with $w_n \in [0, 1]$). As usual, the fact that agents observe an average of recent and 'original' actions may be interpreted either literally, or as shorthand for an assumption that agents are exposed to many (undated) individual actions but store only a mental summary statistic as a way to cope with data overload.

A second purpose of studying this extreme case is to understand whether it is possible for a single cohort's persistent influence to slow down learning. With this in mind, let us define the cumulative persistence of a cohort 1 actions to be $\sum_{n=1}^{\infty} w_n$. (This is the natural extension of the definition introduced in Section 3.3.) In the interest of brevity, the results here are informal, and we focus on two illustrative cases.

Case 1: After a finite time, the first actors' choices are never observed.

More precisely, suppose that w_n is equal to one for all $n \leq K - 1$ and zero for all $n \geq K$, for some finite $K > 1$. Thus, agents in cohorts 2 through K observe only x_1 , but not more recent actions. From cohort $K + 1$ on, an agent observes the most recent action, unmuddled by any confounding effect from x_1 . In this case, it is not hard to see that learning must be at rate n , regardless of how large K is. Suppose that we set aside the actions of cohorts 2 through $K - 1$, and relabel cohort K as $\tilde{2}$, cohort $K + 1$ as $\tilde{3}$, and so on. Then we have cohort $\tilde{2}$ observing action x_1 , cohort $\tilde{3}$ observing $x_{\tilde{2}}$, cohort $\tilde{4}$ observing $x_{\tilde{3}}$, and so on, with cohort $\tilde{n} + 1$ observing $x_{\tilde{n}}$, where $\tilde{n} = n - K + 1$. But this is just the case in which every cohort observes its immediate predecessor, and so we know that learning is at rate \tilde{n} . In this case, the only effect of cohort 1's persistence is to stall learning for K stages. As long as K is finite, this has no lasting effect on the speed of learning.

Case 2: The first actors' choices exit gradually.

For convenience, we consider weight sequences of the form $w_n = 1/n^a$, for $a > 0$. As usual, we define κ_{n-1} as the expected squared error in \bar{x}_{n-1} . Recalling that the expected squared error in x_1 is v_c , we can write κ_{n-1} recursively as:

$$\kappa_{n-1} = (1 - w_{n-1})^2 \text{var}(x_{n-1}) + 2w_{n-1}(1 - w_{n-1})\rho_{n-1} + w_{n-1}^2 v_c.$$

In this expression, the second term with $\rho_{n-1} \equiv E((x_1 - \theta)(x_{n-1} - \theta))$ represents the persistent influence of the first-movers' signal error on the most recent action. The third term represents the direct effect of first-mover error on learning. Although this direct effect does limit learning in principle (κ_n cannot shrink faster than w_n^2), as usual it is the indirect effect that really applies the brakes. The full equations of motion for this model (including ρ_n) are derived in the appendix.

Result 1 *If first-mover influence disappears fast enough ($a \geq \frac{3}{2}$), then learning is fast ($\kappa_n \rightarrow 0$ at*

rate n). If first-mover influence disappears more slowly ($a < \frac{3}{2}$), then learning is at rate $\left(\frac{1}{w_n}\right)^{\frac{2}{3}} = n^{\frac{2}{3}a}$.

We designate this a “result” on the basis of (i) a persuasive but semi-rigorous analysis of the equations of motion (also in the appendix) and (ii) compelling computational evidence.⁴² Figure 1(a) shows log plots of $\ln \kappa_n$ versus $\ln n$ for the first 100,000 cohorts, for a range of values of a (from 0.3 to 1.8, in increments of 0.3). The curves are clearly close to linear for n large enough. For each curve, if its slope tends to a constant, then that constant must be the rate of learning (as a power of n).⁴³ Figure 1(b) plots these slopes $r_n \equiv \left|\frac{\Delta \ln \kappa_n}{\Delta \ln n}\right|$ (where $\Delta \ln \kappa_n = \ln \kappa_n - \ln \kappa_{n-1}$ and $\Delta \ln n = \ln n - \ln(n-1)$). Consistent with Result 1, these slopes appear to converge to $\min\left(\frac{2}{3}a, 1\right)$.⁴⁴

To put this result in context, recall Corollary 1. In loose terms, Corollary 1 says that if cohort 1 actions exit out of the observation pool à la $\{w_n\}$ – and (this is the uniform exit condition) all subsequent cohorts also disappear at a similar rate – then learning slows to rate $\ln(1/w_n)$. Result 1 suggests that if only one cohort persists in this way, learning still slows down, but not as much (rate $(1/w_n)^{2/3}$).

Next consider the cumulative persistence of cohort 1.⁴⁵ For Case 1, this is finite and simply equal to K . For Case 2, the cumulative persistence is finite whenever $a > 1$, and infinite otherwise. Here as in Section 3.3, finite cumulative persistence does not guarantee fast learning: if $a \in (1, \frac{3}{2})$, the direct impact of the first actions is bounded but the speed of learning is still slow. Furthermore, since K was arbitrary, one can easily create scenarios in which cumulative persistence is higher in Case 1 than in Case 2, yet the speed of learning is slower in Case 2. The cumulative persistence measure’s main deficiency is failing to account for the fact that blending an old action into the observational signal can do more damage in later stages, when the precision of the old action is much worse than that of the rest of the observational signal. When n is large, an agent correctly expects her observational signal to be good, and rationally puts high weight on it relative to her idiosyncratic signal. A small amount of pollution from x_1 at this stage will be particularly over-influential, since it will be treated with a deference that it does not deserve.

⁴²The analysis in the appendix points toward a proof, but there are technicalities that would appear to require considerable spilling of ink. The plots below are exemplary; computations with other parameters are also consistent with Result 1.


⁴³To be a bit more precise,

$$\ln(\kappa_n n^r) = \ln \kappa_1 + \sum_{m=2}^n (r - r_m) \Delta \ln n .$$


So if $\kappa_n \rightarrow 0$ at rate n^r , for some r , then r_n must converge to r .

⁴⁴For smaller a (specifically $a = 0.3$ and $a = 0.6$ here), although the difference $\frac{2}{3}a - r_n$ is small from the start, it appears to shrink rather slowly. For $a = 1.8$, the fact that r_n rises briefly above 1 – suggesting that κ_n shrinks faster than rate n^1 for a time – may seem to violate the constraint that learning cannot be faster than the full information rate obtained by pooling signals, which is rate n . There is no contradiction – κ_n is always strictly greater than $\frac{v_c}{n}$. The transient phase when $r_n > 1$ arises because the gap $(\kappa_n - \frac{v_c}{n})$ between actual and full information errors ‘burns off’ quickly as w_n shrinks.

⁴⁵I thank a referee for suggesting this possible line of demarcation between fast and slow learning.



(a) $\ln \kappa_n$ vs. $\ln n$



(b) Instantaneous rate of learning $r_n = \left| \frac{\Delta \ln \kappa_n}{\Delta \ln n} \right|$

Figure 1: Learning when only cohort 1 actions are persistent. Cohort 1 actions exit as $w_n = n^{-a}$, for six different values of a , ranging from 0.3 (lightest curve) to 1.8 (darkest curve), in increments of 0.3. Parameters $v = 2$, $v_c = 1$, first 10^5 cohorts computed.

6.2 Example: Hyperbolic Exit

To motivate the example, suppose that before cohort n acts, a government statistical agency publishes what it intends to be x_{n-1} , the average action in the most recent cohort. However, in the agency’s database, a small fraction of past actions are entered with a missing “date” field, and due to a programming error, these actions are included in the published statistic.

Formally, define y_n to be the (equally weighted) average of actions x_1 through x_n . An agent in cohort n receives the observational signal $\bar{x}_{n-1} = \beta x_{n-1} + (1 - \beta) y_{n-2}$, where we will think of β as close to one and $1 - \beta$ as the error rate. Under these assumptions, the share of action x_{n-1} in the observational signal falls from β in cohort n , to $\frac{1-\beta}{n-1}$ in cohort $n+1$, then to $\frac{1-\beta}{n}$ in cohort $n+2$, and so on. Because an action’s share declines dramatically at first, then more gradually, we refer to this as hyperbolic exit (by way of analogy with hyperbolic discounting). Notice that if $\beta = 0$, so that $\bar{x}_{n-1} = y_{n-2}$, we are almost back to the version of our baseline model in which d_n is set to zero. In that case, each agent observes an unweighted average of her predecessors, and learning is at rate $\ln n$. One can think of this example as an extension of that case in which the weight on non-immediate predecessors is dialed down arbitrarily low.

In this case, an analysis of learning requires three state variables: we define $\nu_n = \text{var}(\bar{x}_n)$, $\omega_n = \text{var}(y_n)$, and $\rho_n = \text{cov}(\bar{x}_n, y_n)$, where ν_n takes on the role of our measure of social learning. Because the expressions involved become lengthy, for this part we make the simplifying assumption that $v = v_c = 1$. An agent in cohort n then puts weight $\alpha_n = \frac{\nu_{n-1}}{1+\nu_{n-1}}$ on her own signal, and weight $1 - \alpha_n$ on \bar{x}_{n-1} . Using this, we can write the following difference equations for the evolution of \bar{x}_n

and y_n :

$$y_n = \frac{\alpha_n}{n} s_n + \frac{n-1}{n} y_{n-1} + \frac{1-\alpha_n}{n} \bar{x}_{n-1} \quad (9)$$

$$\bar{x}_n = \alpha_n \beta s_n + (1-\beta) y_{n-1} + (1-\alpha_n) \beta \bar{x}_{n-1} \quad (10)$$

Learning depends on the evolution of the covariance matrix of \bar{x}_n and y_n (that is, ω_n , ν_n , and ρ_n), which can be derived from (9) and (10). These equations of motion are relegated to the appendix, but we will make a few remarks about them.

Remark 2 *If $\beta = 1$, then the model reduces to PA with immediate exit, and social learning is at rate n . If $\beta = 0$, the model is essentially the same as model PA with no exit ($d_n = 0$), and social learning is at rate $\ln n$.*⁴⁶

One might expect an intermediate rate of learning when $\beta \in (0, 1)$. On the one hand, mixing old actions into the public signal degrades its precision (relative to the immediate exit case). In particular, the precision of \bar{x}_n cannot grow faster than that of y_n .⁴⁷ However, the old actions contained in y_n should be more precise (compared to the no exit case) because they will have relied more heavily on their own recent predecessors. (That is, just as the influence of y_n drags down the rate at which ν_n improves, the influence of \bar{x}_n may pull up the rate at which ω_n improves.) Proposition 10 refutes this compromise intuition: if $\beta < 1$, then the influence of old actions dominates and learning is at rate $\ln n$. To simplify the proof, consider “test sequences” of the form $R_n(a) \equiv \sum_{m=1}^n r_n(a)$, where $r_n(a) = n^{a-1}$. For $a \in (0, 1]$, $R_n(a)$ grows as n^a , while $R_n(0)$ grows as $\ln n$.

Proposition 10 *In the model with hyperbolic exit, suppose that ω_n , ν_n , and ρ_n all converge to zero at common rate $R_n(a)$ for some $a \geq 0$. If $\beta = 1$, then $a = 1$. If $\beta < 1$, then $a = 0$. That is, learning is at rate n if agents can observe only their immediate predecessors but only logarithmic in n if there is any persistence from earlier actions.*

Proposition 10 is couched narrowly and could probably be weakened. In particular, careful study of the equations of motion suggests that if $\beta < 1$, then $(\ln n) \omega_n$, $(\ln n) \nu_n$, and $(\ln n) \rho_n$ all must converge to $\frac{1-\beta}{2}$, but we will not attempt a proof of this. The result provides further confirmation that the tail behavior of noise from prior actions is critical to the speed of learning. Here, even when $1 - \beta$ is small, these tails are fat. For example, at stage n , noise from actions more than $n/2$ stages old comprises roughly a fraction $\frac{1-\beta}{2}$ of the observation pool, regardless of n . No matter how small $1 - \beta$ is, eventually the cumulative influence from these fat tails slows down learning.

7 Concluding Remarks

At its heart, social learning asks how effective people can be at pooling their individual information. The question is most interesting – and most realistic – when there are constraints on learning from

⁴⁶In contrast with the $d_n = 0$ case of model PA, x_{n-1} is omitted from the average that cohort n sees, but this is inessential.

⁴⁷Specifically, we have $\nu_n \geq (1-\beta)^2 \omega_{n-1}$, which follows from the $(1-\beta) y_{n-1}$ term in (10).

others. One such constraint, which inspired the herding literature, is that sometimes people may observe others' choices but not the beliefs that led to them. Another constraint, which is the focus of this paper, is that people may see a summary of others' choices, such as aggregate sales data, or a statistic on the news. A summary such as this will often do a poor job of summarizing the *information content* of others' choices. The data that an observer would need to form a better summary is not just the individual choices, but also their *context*: when each choice was made, which prior choices it relied on, and so forth. We show that when people learn from summaries of past actions, they cannot correct for a type of echo chamber effect (recent actions rely on older ones) that tends to give the earliest actions too much influence. This inertia can slow learning down so much that, for practical purposes, it is effectively incomplete. The inertial effect of old actions grows more pernicious in a changing world (the moving target model), but it can also become beneficial (e.g., with heterogeneous tastes, or a self reliance bias) when a decisionmaker wishes her samplees had paid *more* heed to the past.

Our model can also be applied if people observe others' actions individually but treat them equally in forming a mental summary statistic. Equal treatment might arise because contextual information is missing (as we suggest in our sample average model). Another rationale, not pursued in this paper, is that people have clues about context but either underestimate the need to correct for correlation among the actions they sample or find it too difficult. In this case, forming a sample average could be a cognitively simple rule of thumb.

For savvy learners with partial context on the choices they see, the first order correction would be to try to discount older choices, since their information is likely to be redundant. In this sense, a higher (exogenous) frequency of sampling recent versus old choices in our model could be interpreted as a proxy for better information about context. As expected, social learning is faster when people are able to collect observations with less redundancy. However, this reduced-form interpretation is not a substitute for an explicit model of inference from partially ordered data. Further work along these lines appears challenging, but would be welcome.⁴⁸

Poor information about the dependencies among people's choices is a problem for applied microeconomists as well as for the people they study. Certain aspects of our model (sequences of cohorts of agents, normally distributed errors) bear a loose similarity to simple econometric models of repeated cross-section data. While our model is not ready to estimate, it suggests a tractable way to build modest assumptions about what people know into structural empirical models of social learning.

Our analysis relies heavily on the tractability of continuous actions paired with normal errors. Developing a similar model with discrete actions, such as the binary action models that have been a workhorse of the herding literature, would seem to require substantially different methods. That

⁴⁸One intermediate approach, developed in an earlier version of this paper, is to retain the observation of a sample average but to allow the agent to choose the composition of this sample (within certain constraints). This structure is motivated by a screening story: for example our sprinter from the introduction might be able to screen older blogs out of her sample based on hairstyles or dated pop culture references, just as she might be able to avoid blogs that are more focused on distance running. One can show that agents generally prefer to screen for recent actions, but if the screening technology is imperfect, the slow learning results in the paper still apply.

said, there is no obvious reason why the factors that contribute to fast or slow learning in our model would not play a roughly similar role if actions were discrete. Other adaptations are more direct. For example, in our model, an agent’s action is his (posterior) point estimate of the unknown parameter. Thus, it could be recast as a model of learning from others’ beliefs rather than from their actions without changing the mathematics.⁴⁹ Opportunities to observe summary statistics about other people’s beliefs – in the form of prediction markets, online product ratings, and so forth – have proliferated of late, and learning from the “wisdom of crowds” has seized the imagination of the popular press.⁵⁰ However, crowds and markets can also make mistakes and propagate stale information.⁵¹ Our model provides a framework for thinking about how the wisdom of crowds may sometimes be underwhelming due to the inertia of early mistakes.

References

- [1] Daron Acemoglu, Munther Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *NBER Working Paper W14040*, May 2008.
- [2] Venkatesh Bala and Sanjeev Goyal. Learning from neighbours. *Review of Economic Studies*, 65(3):595–621, July 1998.
- [3] Abhijit Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, August 1992.
- [4] Abhijit Banerjee and Drew Fudenberg. Word of mouth learning. *Games and Economic Behavior*, 46:1–22, 2004.
- [5] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change in informational cascades. *Journal of Political Economy*, 100(5):992–1026, October 1992.
- [6] Steven Callander and Johannes Horner. The wisdom of the minority. *mimeo*, 46, 2006.
- [7] Bogachan Celen and Shachar Kariv. Observational learning under incomplete information. *Games and Economic Behavior*, 47(1):72–86, 2004.
- [8] Peter DeMarzo, Dimitri Vayanos, and Jeffrey Zwiebel. Persuasion bias, social influence and uni-dimensional opinions. *Quarterly Journal of Economics*, 118(3):909–968, 2003.

⁴⁹We refer to agents’ beliefs after acting but before utility is realized, which is reasonable for decisions, like investment in education, for which the payoff is deferred. Adapting the model to reflect learning from agents’ beliefs after realizing utility would not be much more trouble. The main change is that we would need to specify what an agent can learn from her utility outcome, an issue that is irrelevant in the standard interpretation of our model.

⁵⁰For example, see [18].

⁵¹In one piquant example of how old and new information may be jumbled, on September 8, 2008, United Airline’s stock price tumbled over 75% upon reports of an impending bankruptcy filing [14]. Sheepish investors soon discovered that the impetus for the sell-off had been a six year old news story that had accidentally been labeled as current. The mistake generated a wealth transfer that was probably in the neighborhood of \$100 million.

- [9] Esther Duflo and Emmanuel Saez. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Quarterly Journal of Economics*, 118(3):815–842, August 2003.
- [10] Glenn Ellison and Drew Fudenberg. Rules of thumb for social learning. *Journal of Political Economy*, 101(4):612–643, 1993.
- [11] Glenn Ellison and Drew Fudenberg. Word-of-mouth communication and social learning. *Quarterly Journal of Economics*, 110(1):93–125, 1995.
- [12] Gunther Eysenbach, John Powell, Oliver Kuss, and Eun-Ryoung Sa. Empirical studies assessing the quality of health information for consumers on the world wide web. *JAMA: The Journal of the American Medical Association*, 287(20):2691–2700, May 22/29, 2002.
- [13] Jacob Goeree, Robert McKelvey, Thomas Palfrey, and Brian Rogers. Self-correcting information cascades. *Review of Economic Studies*, 74(3):733–62, July 2007.
- [14] Micheline Maynard. A mistaken news report hurts united. *New York Times*, September 8 2008.
- [15] Kaivan Munshi. Social learning in a heterogeneous population: technology diffusion in the indian green revolution. *Journal of Development Economics*, 73(1):185 – 213, 2004.
- [16] Lones Smith and Peter Sorensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, March 2000.
- [17] Lones Smith and Peter Sorensen. Rational social learning with random sampling. *mimeo*, 2008.
- [18] James Surowiecki. *The Wisdom of Crowds*. 2004.
- [19] Xavier Vives. How fast do rational agents learn? *Review of Economic Studies*, 60(2):329–47, April 1993.
- [20] Georg Weizsäcker. Do we follow others when we should? a simple test of rational expectations. *American Economic Review*, 100(5):2340–60, December 2010.

8 Proofs

Any omitted proofs appear in the Supplementary Appendix. Many of the proofs use the following standard result, due to Stolz and Cesàro, which can be thought of as a generalization of L’Hôpital’s rule to discrete summations.

Theorem 1 *Let $\{A_n\}$ and $\{B_n\}$ be two sequences of real numbers, with $\{B_n\}$ strictly increasing and unbounded. If the limit*

$$\lim_{n \rightarrow \infty} \frac{A_{n+1} - A_n}{B_{n+1} - B_n} = K$$

exists, then the limit $\lim_{n \rightarrow \infty} \frac{A_n}{B_n}$ also exists and is equal to K .

Corollary 3 Suppose that $\{A_n\}$ and $\{B_n\}$ satisfy Theorem 1 with limit K . Define $a_n = A_{n+1} - A_n$ and $b_n = B_{n+1} - B_n$. Let $\{C_n\}$ be strictly increasing and unbounded, with increments $c_n = C_{n+1} - C_n$. Suppose that

$$\lim_{n \rightarrow \infty} \frac{a_n - Kb_n}{c_n} = L$$

Then $\lim_{n \rightarrow \infty} \frac{A_n - KB_n}{C_n}$ exists and is equal to L .

Proof. This is a direct application of the theorem. ■

Proposition 1

Start with \tilde{d} and P_n . Since d_n is assumed to converge, let $d_n \rightarrow \bar{d}$. As the population P_n is weakly increasing, $\tilde{d} > 0$ implies $\bar{d} > 0$ and P_n converges to a finite limit. Conversely, $P_n \rightarrow N^* < \infty$ implies $d_n \rightarrow 0$, so $\tilde{d} = \frac{\bar{d}}{N^*} > 0$. For Part 2, $\tilde{d} = 0$ implies that either $P_n \rightarrow \infty$, or $\bar{d} = 0$ (which in turn implies $P_n \rightarrow \infty$). Conversely, since $d_n \in [0, 1]$, we have $P_n \rightarrow \infty$ implies $\frac{d_n}{P_n} \rightarrow 0$.

Next for arbitrary m , observe that all cohort m agents remain in P_m to be seen by cohort $m+1$ agents. A measure $\frac{d_m}{P_m}$ exit after cohort $m+1$ leaving $1 - \frac{d_m}{P_m}$ in P_{m+1} . Of these, a fraction $\frac{d_{m+1}}{P_{m+1}}$ leave after cohort $m+2$ acts, leaving a total of $\left(1 - \frac{d_m}{P_m}\right) \left(1 - \frac{d_{m+1}}{P_{m+1}}\right)$ in P_{m+2} . Thus we have $\#_m^{m+1} = 1$ and $\#_m^n = \prod_{k=m}^{n-2} \left(1 - \frac{d_k}{P_k}\right)$ for $n \geq m+2$. For properties EE₁ and SE₁ it suffices to look at the limit of $\frac{\ln \#_m^n}{n}$, since $\frac{n}{n-m} \rightarrow 1$ with n . We have $\frac{\ln \#_m^n}{n} = \frac{1}{n} \sum_{k=m}^{n-2} \ln \left(1 - \frac{d_k}{P_k}\right)$. Now apply Theorem 1 to get $\lim_{n \rightarrow \infty} \frac{\ln \#_m^n}{n} = \lim_{n \rightarrow \infty} \ln \left(1 - \frac{d_n}{P_n}\right)$, if the latter limit exists. We have $d_n \in [0, 1]$ and $P_n \geq 1$, so $\tilde{d} = \lim_{n \rightarrow \infty} \frac{d_n}{P_n} \in [0, 1]$. Furthermore, $\tilde{d} = 1$ is possible only if $d_n = 1$ for all n , a case that has been ruled out. Thus we have $\lim_{n \rightarrow \infty} \frac{\ln \#_m^n}{n} = \ln \left(1 - \tilde{d}\right) \leq 0$. Notice that this limit does not depend on m .

Because the fraction of cohort m actions in the observation pool observed by cohort n is $\lambda_m^n = \#_m^n / P_{n-1}$, we have $\lim_{n \rightarrow \infty} \frac{\ln \lambda_m^n}{n} = \ln \left(1 - \tilde{d}\right) - \lim_{n \rightarrow \infty} \frac{\ln P_n}{n}$. But P_n is bounded above by n , so $\lim_{n \rightarrow \infty} \frac{\ln P_n}{n} = 0$. Thus $\lim_{n \rightarrow \infty} \frac{\ln \lambda_m^n}{n} = \lim_{n \rightarrow \infty} \frac{\ln \#_m^n}{n} = \ln \left(1 - \tilde{d}\right)$ for all m , which implies EE₁ \Leftrightarrow EE₂ and SE₁ \Leftrightarrow SE₂. Finally observe that $\ln \left(1 - \tilde{d}\right) < 0 \Leftrightarrow \tilde{d} > 0$ and $\ln \left(1 - \tilde{d}\right) = 0 \Leftrightarrow \tilde{d} = 0$, which completes the proof.

Proposition 2

Part 1: $\kappa_n \rightarrow 0$ at rate R_n

Let $R_n = \sum_{m=1}^n \frac{1}{P_m}$ and $\bar{\tau}_n = \frac{1}{\kappa_n}$ as in the text and define $z_n = \frac{\bar{\tau}_n}{R_n} = \frac{1}{R_n \kappa_n}$. We will apply the Stolz-Cesàro Theorem to the increments of $\bar{\tau}_n$ and R_n . First note that R_n is strictly increasing and unbounded. (This follows from the fact that $1 \leq P_n \leq n$, so R_n grows at least as fast as the harmonic series $\frac{1}{1} + \frac{1}{2} + \dots + \frac{1}{n} + \dots$) We can write $\bar{\tau}_1 = 1$ and

$$\bar{\tau}_n - \bar{\tau}_{n-1} = \frac{M_n \bar{\tau}_{n-1}}{v \bar{\tau}_{n-1} + 1 - M_n} \quad (11)$$

Let

$$\begin{aligned} K_n &= \frac{\bar{\tau}_n - \bar{\tau}_{n-1}}{R_n - R_{n-1}} = \frac{P_n M_n \bar{\tau}_{n-1}}{v \bar{\tau}_{n-1} + 1 - M_n} \\ &= \left(2 - \frac{1}{P_n} \frac{v_c + \kappa_{n-1}}{v + \kappa_{n-1}}\right) \left(\frac{\bar{\tau}_{n-1}}{v \bar{\tau}_{n-1} + 1 - M_n}\right) \end{aligned}$$

To prove $z_n \rightarrow 0$ at rate R_n , it will suffice to show that $K_n \rightarrow K$ for some strictly positive K , and therefore (by the Stolz–Cesàro Theorem) that $z_n \rightarrow K$ as well. To show this, first observe that the first term in parentheses is strictly positive (because $P_n \geq 1$ and $v_c \leq v$), strictly increasing (because $\frac{1}{P_n}$ is weakly decreasing and κ_n strictly decreasing implies $\frac{v_c + \kappa_{n-1}}{v + \kappa_{n-1}}$ strictly decreasing), and bounded above; thus it converges to a strictly positive limit. This also implies that M_n converges to a weakly positive limit; define M by $\{M_n\} \rightarrow M$. For the second term in parentheses, note that $\bar{\tau}_n$ is strictly increasing (because κ_n is strictly decreasing) and either bounded (in which case $\bar{\tau}_n$ converges) or unbounded. In either case, ($\bar{\tau}_n$ converges or grows without bound), the second term in parentheses converges. Its limit is strictly positive because $\frac{\bar{\tau}_{n-1}}{v\bar{\tau}_{n-1}+1-M_n} \geq \frac{\bar{\tau}_{n-1}}{v\bar{\tau}_{n-1}+1} \geq \frac{1}{v+1}$ (where the last step uses $\bar{\tau}_n \geq 1$). Thus $K_n \rightarrow K$, with K strictly positive. Because this implies that $\bar{\tau}_n$ tends to infinity, the limiting constant is $K = \frac{2}{v} - \frac{v_c}{v^2} \lim_{n \rightarrow \infty} (1/P_n)$. (The limit of $1/P_n$ exists regardless of whether P_n is bounded.)

Part 2: Interpreting rate R_n learning

Let $f(n)$ be an arbitrary function on the natural numbers, and let $\Delta f_n = f(n) - f(n-1)$. Apply Theorem 1 to get $\lim_{n \rightarrow \infty} R_n/f(n) = \lim_{n \rightarrow \infty} 1/(P_n \Delta f_n)$, if the latter limit exists. Use $f(n) = n$ to get $\lim_{n \rightarrow \infty} R_n/n = \lim_{n \rightarrow \infty} 1/P_n$ which is either strictly positive and finite or zero, depending on whether the observation pool is bounded or unbounded. Next suppose that d_n is uniformly bounded away from 1. By assumption, d_n converges, so suppose $d_n \rightarrow 1 - \bar{g}$, with $\bar{g} > 0$. Set $f(n) = \ln n$ to get $\lim_{n \rightarrow \infty} \ln n/R_n = -\lim_{n \rightarrow \infty} P_n \ln(1 - 1/n)$. Write the expression in the last limit as $P_n \ln(1 - 1/n) = (P_n/n)(n \ln(1 - 1/n))$. But by assumption, $\lim_{n \rightarrow \infty} (P_n/n) = \bar{g}$, while $\lim_{n \rightarrow \infty} (n \ln(1 - 1/n)) = -1$, so $\lim_{n \rightarrow \infty} \ln n/R_n = \bar{g}$.

Lemmas 2, 3, 4, 5, and 6 are used in the proof of Proposition 3. Their proofs appear in the Supplementary Appendix.

Lemma 2 (Second order convergence rates) Suppose $\frac{\bar{\tau}_n}{R_n} \rightarrow K$, with $\eta = \frac{1}{Kv}$ as in the text. Let $H_n = \sum_{j=1}^n \frac{1}{P_j}$ if P_n is unbounded and $\frac{R_n}{P_n}$ diverges, or $H_n = \sum_{j=1}^n \frac{1}{R_j P_j}$ otherwise. Then $\frac{\bar{\tau}_n - KR_n}{H_n} \rightarrow L$, for some finite L . Furthermore, $\frac{R_n}{H_n} (\eta v - \kappa_n R_n) \rightarrow L (\eta v)^2$ and $\frac{R_n}{H_n} (\eta - \alpha_n R_n) \rightarrow L \eta^2 v$.

Lemma 3 $R_n P_n \geq n$ for all $n \geq 1$.

Lemma 4 (Discrete approximation of $\ln R_n$) Define $S_n = \sum_{k=1}^n \frac{1}{R_k P_k}$. The sequence $\{S_n - \ln R_n\}_{n \geq 1}$ converges.

Lemma 5 (Linear approximation of $\ln(1 - \alpha_n/P_n)$) Define c_k by $\ln(1 - \alpha_k/P_k) = -\alpha_k/P_k + c_k$. For arbitrary $m > 1$, the sum $\sum_{k=m+1}^{\infty} c_k$ converges.

Lemma 6 For arbitrary $m \geq 1$ and $q \in (0, 1)$, the sum $\sum_{k=m+1}^{\infty} \frac{1}{R_k^{2-q} P_k}$ converges.

Proposition 3

Part 1 Fix any arbitrary m and let $\tilde{\lambda}(n) = \ln \lambda_m^{n+1}$. Following the text, we can write

$$\tilde{\lambda}(n) = -\ln P_m + \sum_{k=m+1}^n \ln \left(1 - \frac{1}{P_k} \right)$$

We must show that $\left\{ \frac{\tilde{\lambda}(n)}{R_n} \right\}_{n > m} \rightarrow -L$, for some $L > 0$. By Theorem 1, it suffices to show that $\left\{ \frac{\tilde{\lambda}(n) - \tilde{\lambda}(n-1)}{R_n - R_{n-1}} \right\}_{n > m+1} =$

$\left\{P_n \ln \left(1 - \frac{1}{P_n}\right)\right\}_{n>m+1} \rightarrow -L$. If P_n is bounded then it converges to some $N^* > 1$, and the result follows directly. If P_n is unbounded, then the result follows from $\lim_{x \rightarrow 0} \frac{1}{x} \ln(1-x) = -1$.

Part 2 Define $\eta = \left(2 - \frac{v_c}{v} \lim_{n \rightarrow \infty} (1/P_n)\right)^{-1}$ as in the text. Observe from the proof of Proposition 2 that $\{\kappa_n R_n\}_{n \geq 1} \rightarrow \eta v$ and $\{\alpha_n R_n\}_{n \geq 1} \rightarrow \eta$. It suffices to show that the sequence $\{\ln(\pi_m^n(R_n)^\eta)\}_{n>m}$ converges, or equivalently that $\{\tilde{\pi}(n) + \eta \ln R_n\}_{n>m}$ converges, where $\tilde{\pi}(n) = \ln \pi_m^{n+1} = \ln(\alpha_m/P_m) + \sum_{k=m+1}^n \ln(1 - \alpha_k/P_k)$. The strategy is to approximate $\ln R_n$ by $S_n = \sum_{k=1}^n \frac{1}{R_k P_k}$ and approximate $\sum_{k=m+1}^n \ln(1 - \alpha_k/P_k)$ by $-\sum_{k=m+1}^n \alpha_k/P_k$. That is, we have

$$\ln(\pi_m^n(R_n)^\eta) = \sum_{k=m+1}^n \left(\frac{\eta}{R_k P_k} - \frac{\alpha_k}{P_k} \right) + Error_n^1 + Error_n^2 + C$$

where $Error_n^1 = \eta \left(\ln R_n - \sum_{k=m+1}^n \frac{1}{R_k P_k} \right)$, $Error_n^2 = \sum_{k=m+1}^n \left(\frac{\alpha_k}{P_k} + \ln \left(1 - \frac{\alpha_k}{P_k}\right) \right)$, and C is a (bounded) constant that does not depend on n . The sequences $\{Error_n^1\}_{n>m}$ and $\{Error_n^2\}_{n>m}$ converge by Lemmas 4 and 5, so to prove Part 2, it suffices to show that the sum $\sum_{k=m+1}^\infty a_k$ converges, where $a_k = \frac{1}{R_k P_k} (\eta - \alpha_k R_k)$. This will depend on second order effects in the speed of learning – namely, whether $\alpha_k R_k$ tends to η fast enough. Lemma 2 answers this question; we have $\frac{R_n}{H_n} (\eta - \alpha_n R_n) \rightarrow \tilde{L}$, for some finite constant \tilde{L} and a sequence H_n defined by $H_n = \sum_{j=1}^n \frac{1}{P_j^2}$ if P_n is unbounded and $\frac{R_n}{P_n}$ diverges, or $H_n = \sum_{j=1}^n \frac{1}{R_j P_j}$ otherwise. Let $a'_k = H_k / (R_k^2 P_k)$, and observe that $\lim_{k \rightarrow \infty} a_k / a'_k = \tilde{L}$. By the limit comparison test, to prove $\sum_{k=m+1}^\infty a_k$ convergent (and therefore to prove Part 2), it suffices to show that $\sum_{k=m+1}^\infty a'_k$ converges. There are two cases to consider, depending on which part of the piecewise definition of H_n applies.

Part 2, Case 1: P_n is either bounded, or R_n/P_n converges. Thus $H_n = \sum_{j=1}^n \frac{1}{R_j P_j}$.

By Lemma 4, $H_n / \ln R_n \rightarrow 1$, so by another application of the limit comparison test, it suffices to show that $\sum_{k=m+1}^\infty \frac{\ln R_k}{R_k^2 P_k}$ converges. Write this sum as $\sum_{k=m+1}^\infty \left(\frac{1}{R_k^{2-q} P_k} \right) \left(\frac{\ln R_k}{R_k^q} \right)$ for an arbitrary constant $q \in (0, 1)$. We have $\lim_{k \rightarrow \infty} \frac{\ln R_k}{R_k^q} = 0$ (since $R_k \rightarrow \infty$ with k), so it will suffice to show that $\sum_{k=m+1}^\infty \frac{1}{R_k^{2-q} P_k}$ converges. Lemma 6 shows this, completing this part of the proof. The spirit of Lemma 6 is to note that the summand may be written $(\Delta R_k) R_k^{q-2}$, which is something like the ‘derivative’ of R_k^{q-1} . Thus we may expect the summation to behave like the ‘definite integral’ $R_m^{q-1} - \lim_{k \rightarrow \infty} R_k^{q-1} = R_m^{q-1}$. Lemma 6 formalizes this intuition.

Part 2, Case 2: P_n is unbounded and R_n/P_n diverges. Thus $H_n = \sum_{j=1}^n \frac{1}{P_j^2}$.

This is the only case where we must rely on condition (Poly). By (Poly), $n^{-b} P_n$ converges to a strictly positive, finite limit. It is straightforward to use this fact, with applications of Theorem 1, to show that $n^{b-1} R_n$ and $n^{2b-1} H_n$ also converge to strictly positive finite limits. Together these imply that our summand a'_k tends to zero at rate k^{1+b} . (That is, $k^{1+b} a'_k$ converges to a strictly positive, finite limit.) But then, convergence of $\sum_{k=m+1}^\infty \frac{H_k}{R_k^2 P_k}$ is implied by the fact that the sum $\sum_{k=m+1}^\infty \frac{1}{k^{1+b}}$ converges.

Part 3 This follows more or less immediately from Part 2. By Part 2, for any cohort m , there exists some $L_m > 0$ such that $\{\pi_m^n(R_n)^\eta\}_{n>m} \rightarrow L_m$. By Proposition 2, $\frac{1}{\kappa_{n-1} R_{n-1}}$ converges to $K = \frac{1}{\eta v}$. Furthermore, $\frac{R_n}{R_{n-1}} = \frac{R_n}{R_n - \frac{1}{P_n}} \rightarrow 1$ since $R_n \rightarrow \infty$ and $\frac{1}{P_n} \leq 1$. Thus we can write

$$\chi_m^n(R_n)^{2\eta-1} = \left(\frac{v_c}{\kappa_{n-1} R_{n-1}} \right) \left(\frac{R_{n-1}}{R_n} \right) (\pi_m^n(R_n)^\eta)^2$$

Each term in parentheses on the righthand side converges to a positive limit, so $\left\{ \chi_m^n (R_n)^{2\eta-1} \right\}_{n>m}$ does as well. Parts 3a and 3b are merely a matter of emphasis. If $P_n \rightarrow N^* < \infty$, then $2\eta - 1 = \frac{2}{2 - \frac{v_c}{v} \frac{1}{N^*}} - 1 = \frac{v_c}{2vN^* - v_c} > 0$. However, if $P_n \rightarrow \infty$, then $2\eta - 1 = 0$, so in this case, χ_m^n tends toward a positive limit as $n \rightarrow \infty$.

Lemma 1

Label the individual actions that comprise $x_{S(n_i)}$ as $\{y_1, y_2, \dots, y_S\}$, so $x_{S(n_i)} = \frac{1}{S} \sum_{k=1}^K y_k$. The expected squared error in $x_{S(n_i)}$ can be written as

$$\begin{aligned} \bar{\kappa}_{n-1} &= \frac{1}{S^2} E \left(\left(\sum_{s=1}^S (y_s - \theta) \right)^2 \right) \\ &= \frac{1}{S^2} \left(\sum_{s=1}^S E \left((y_s - \theta)^2 \right) + 2 \sum_{s=1}^S \sum_{s'=s+1}^S E \left((y_s - \theta)(y_{s'} - \theta) \right) \right) \\ &= \frac{1}{S} \hat{\kappa}_{n-1} + \frac{S-1}{S} \hat{\rho}_{n-1} \end{aligned}$$

Agent n_i 's optimal weight $\hat{\alpha}_n$ solves

$$\hat{\alpha}_n = \arg \min_a E \left((as_{n_i} + (1-a)x_{S(n_i)} - \theta)^2 \right)$$

Because s_{n_i} and $x_{S(n_i)}$ are independent, conditional on θ , the minimand can be written as

$$a^2 E \left((s_{n_i} - \theta)^2 \right) + (1-a)^2 E \left((x_{S(n_i)} - \theta)^2 \right) = a^2 v + (1-a)^2 \bar{\kappa}_{n-1}$$

from which the optimal weight follows.

Proposition 5

As a preliminary step, we compute the increments of $\hat{\kappa}_n$ and $\hat{\rho}_n$:

$$\begin{aligned} \Delta \hat{\kappa}_n &= \hat{\kappa}_n - \hat{\kappa}_{n-1} = -\frac{1}{P_n} (\hat{\kappa}_{n-1} - \hat{\alpha}_n v) \\ \Delta \hat{\rho}_n &= \hat{\rho}_n - \hat{\rho}_{n-1} = -\frac{\hat{\alpha}_n \hat{\rho}_{n-1}}{P_n} \left(2 - \frac{1}{P_n} \frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} (\hat{\rho}_{n-1} + v_c) \right) \end{aligned}$$

Use the definitions of $\hat{\alpha}_n$ and $\bar{\kappa}_n$ to write $\Delta \hat{\kappa}_n$ as:

$$\begin{aligned} \Delta \hat{\kappa}_n &= -\frac{1}{P_n} \frac{1}{\bar{\kappa}_{n-1} + v} \left(\hat{\kappa}_{n-1}^2 + \frac{S-1}{S} (\hat{\kappa}_{n-1} - \hat{\rho}_{n-1}) (v - \hat{\kappa}_{n-1}) \right) \quad \text{or} \\ &= -\frac{1}{P_n} \frac{\hat{\kappa}_{n-1}}{\bar{\kappa}_{n-1} + v} \left(\frac{1}{S} \hat{\kappa}_{n-1} + \frac{S-1}{S} \left(v \left(1 - \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}} \right) + \hat{\rho}_{n-1} \right) \right) \end{aligned}$$

The following three supporting lemmas are proved in the Supplementary Appendix.

Lemma 7 $\frac{\hat{\kappa}_n}{\hat{\rho}_n}$ is bounded. (That is, $\hat{\rho}_n$ tends to zero no faster than $\hat{\kappa}_n$.)

Lemma 8 $\frac{1}{\hat{\rho}_n R_n}$ is bounded. (That is, $\hat{\rho}_n$ tends to zero no faster than rate R_n .)

Lemma 9 At least one of the following must hold: either $\frac{\hat{\kappa}_n - \hat{\rho}_n}{\hat{\rho}_n} \rightarrow 0$ or $(\hat{\kappa}_n - \hat{\rho}_n) R_n \rightarrow 0$.

Lemmas 7 and 8 put upper bounds on how fast the error variance vanishes, while Lemma 9 puts a lower bound on how fast the gap between $\hat{\kappa}_n$ and $\hat{\rho}_n$ must shrink. By Lemma 9, we either have $\frac{\hat{\kappa}_n - \hat{\rho}_n}{\hat{\rho}_n} \rightarrow 0$, in which case the proof is complete, or else $(\hat{\kappa}_n - \hat{\rho}_n) R_n \rightarrow 0$ holds. In the latter case we can write $\frac{\hat{\kappa}_n - \hat{\rho}_n}{\hat{\rho}_n} = ((\hat{\kappa}_n - \hat{\rho}_n) R_n) \left(\frac{1}{\hat{\rho}_n R_n} \right)$. Because the first term tends to zero and the second is bounded, we have $\frac{\hat{\kappa}_n - \hat{\rho}_n}{\hat{\rho}_n} \rightarrow 0$ for this case as well, which completing the proof.

Proposition 6

$S > 1$

To show that $\hat{\rho}_n R_n$ tends to a positive constant, we follow the proof of Lemma 8, adding the additional information (from Proposition 5) that $\frac{\hat{\kappa}_n - \hat{\rho}_n}{\hat{\rho}_n} \rightarrow 0$. We have

$$\lim_{n \rightarrow \infty} \frac{1}{\hat{\rho}_n R_n} = \lim_{n \rightarrow \infty} \frac{\Delta\left(\frac{1}{\hat{\rho}_n}\right)}{\Delta R_n} = \lim_{n \rightarrow \infty} \frac{\hat{\rho}_{n-1}}{\hat{\rho}_n} \frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} \left(2 - \frac{1}{P_n} \frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} (\hat{\rho}_{n-1} + v_c) \right)$$

On the righthand side, $\frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} = \frac{1}{\hat{\kappa}_{n-1} + v} \frac{\hat{\rho}_{n-1} + \frac{1}{S}(\hat{\kappa}_{n-1} - \hat{\rho}_{n-1})}{\hat{\rho}_{n-1}} \rightarrow \frac{1}{v}$. Furthermore, we have $\frac{\hat{\rho}_n}{\hat{\rho}_{n-1}} = \left(1 - \frac{\hat{\alpha}_n}{P_n} \right) + \hat{\alpha}_n \frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} \frac{v_c}{P_n^2} \rightarrow 1$ (because $\hat{\alpha}_n \rightarrow 0$). Thus we have

$$\lim_{n \rightarrow \infty} \frac{1}{\hat{\rho}_n R_n} = \lim_{n \rightarrow \infty} \frac{\Delta\left(\frac{1}{\hat{\rho}_n}\right)}{\Delta R_n} = \frac{1}{v} \left(2 - \frac{v_c}{v} \lim_{n \rightarrow \infty} \frac{1}{P_n} \right) = \frac{1}{\eta v}$$

Thus $\hat{\rho}_n \rightarrow 0$ at rate R_n . From this, it follows immediately that $\hat{\kappa}_n R_n$ tends to the same limit as $\hat{\rho}_n R_n$, while $\hat{\alpha}_n R_n \rightarrow \frac{1}{v} \lim_{n \rightarrow \infty} \hat{\rho}_n R_n$.

$S = 1$

If $S = 1$, then $\hat{\alpha}_n = \frac{\hat{\kappa}_{n-1}}{\hat{\kappa}_{n-1} + v}$ and so $\hat{\kappa}_n$ is characterized by

$$\Delta \hat{\kappa}_n = -\frac{1}{P_n} (\hat{\kappa}_{n-1} - \hat{\alpha}_n v) = -\frac{1}{P_n} \frac{\hat{\kappa}_{n-1}^2}{\hat{\kappa}_{n-1} + v}$$

Apply Theorem 1 to $\frac{1/\hat{\kappa}_n}{R_n}$, with $\Delta(1/\hat{\kappa}_n) = \frac{\Delta \hat{\kappa}_n}{\hat{\kappa}_n \hat{\kappa}_{n-1}}$ to get

$$\lim_{n \rightarrow \infty} \frac{1/\hat{\kappa}_n}{R_n} = \lim_{n \rightarrow \infty} \frac{\Delta\left(\frac{1}{\hat{\kappa}_n}\right)}{\Delta R_n} = \lim_{n \rightarrow \infty} \frac{\hat{\kappa}_{n-1}}{\hat{\kappa}_n} \frac{1}{\hat{\kappa}_{n-1} + v} = \frac{1}{v}$$

where the last step uses $\frac{\hat{\kappa}_{n-1}}{\hat{\kappa}_n} \rightarrow 1$ (which follows from $\frac{\Delta \hat{\kappa}_n}{\hat{\kappa}_{n-1}} = -\frac{1}{P_n} \frac{\hat{\kappa}_{n-1}}{\hat{\kappa}_{n-1} + v} \rightarrow 0$). Thus, $\hat{\kappa}_n R_n \rightarrow v$.

For subsequent results, we use the following lemma. While the result is surely not novel, we have not found this particular formulation in the literature.

Lemma 10 *Define a sequence Z_n by $Z_1 = Z_{init} > 0$ and $Z_n = (1 - b_{n-1}) Z_{n-1} + c_{n-1}$ if $n > 1$. Write the increments of this sequence as $\Delta Z_n \equiv Z_n - Z_{n-1} = -b_{n-1} \left(Z_{n-1} - \frac{c_{n-1}}{b_{n-1}} \right)$. Suppose that b_n and c_n are sequences with $0 < b_n < 1$ and $\sum_{n=1}^{\infty} b_n$ divergent, $c_n \rightarrow 0$, and $\frac{c_n}{b_n} \rightarrow Z \geq 0$. Then $Z_n \rightarrow Z$.*

Proposition 7

Case 1 from Table 1: P_n is bounded with $P_n \rightarrow N^*$

Let $F_P : [0, v_c] \rightarrow [0, v_c]$ be defined by $F_P(z) = \left(1 - \frac{\alpha(z)}{P}\right)^2 z + \left(\frac{\alpha(z)}{P}\right)^2 v_c$, where $\alpha(z) = \frac{z+v_{ta}}{z+v_{ta}+v}$. Notice that $\kappa_n = F_{P_n} \circ F_{P_{n-1}} \circ \dots \circ F_{P_2}(\kappa_1)$. We claim the following, deferring proofs:

Claim (1) F_P has a unique, strictly positive fixed point, denoted \bar{z}^P , satisfying

$$\left(2 - \frac{\alpha(z)}{P}\right) z - \frac{\alpha(z)}{P} v_c \Big|_{z=\bar{z}^P} = 0 \quad (12)$$

Claim (2) There is a function $b_P(z)$, satisfying $\frac{\alpha(0)}{P} < b_P(z) < 1$ for all $z \in [0, v_c]$, such that $F_P(z) - \bar{z}^P = (1 - b_P(z))(z - \bar{z}^P)$.

Claim (3) Define $\bar{z} \equiv \bar{z}^{N^*}$. We have $\{\bar{z}^{P_n}\} \rightarrow \bar{z}$.

A general intuition for the main result is that F_{P_n} tends to F_{N^*} , and iteration of F_{N^*} tends to \bar{z} , so $\kappa_n \rightarrow \bar{z}$. The proof is less direct. Define a variable $\omega_n \equiv \kappa_n - \bar{z}^{P_{n+1}}$, and note that $\omega_n = (F_{P_n}(\kappa_{n-1}) - \bar{z}^{P_n}) + (\bar{z}^{P_n} - \bar{z}^{P_{n+1}})$. Using Claim (2), we have

$$\begin{aligned} \omega_n &= (1 - b_{P_n}(\kappa_{n-1}))\omega_{n-1} + (\bar{z}^{P_n} - \bar{z}^{P_{n+1}}) \\ &= (1 - b_{n-1})\omega_{n-1} + c_{n-1}, \text{ setting } b_{n-1} = b_{P_n}(\kappa_{n-1}) \text{ and } c_{n-1} = \bar{z}^{P_n} - \bar{z}^{P_{n+1}} \end{aligned}$$

Now apply Lemma 10. We have $c_n \rightarrow 0$ by Claim (3), $\sum_{n=1}^{\infty} b_n$ divergent by Claim (2) and the fact that $\frac{\alpha(0)}{P_n} \rightarrow \frac{\alpha(0)}{N^*} > 0$, and $\frac{c_n}{b_n} \rightarrow 0$. Thus the lemma applies, and so $\omega_n \rightarrow 0$. Together with $\bar{z}^{P_n} \rightarrow \bar{z}$, this implies that $\kappa_n \rightarrow \bar{z}$.

Proof of Claim (1) The equation $F_P(z) - z = 0$ can easily be reduced to (12), which is quadratic in z . $F_P(z) - z = 0$ has one irrelevant negative solution (because $F_P(0) > 0$ and $\lim_{z \rightarrow -(v_{ta}+v)^+} F_P(x) = -\infty$). The remaining solution lies in $(0, v_c)$ (because $F_P(0) > 0$ and $F_P(z) - z|_{z=v_c} < 0$).

Proof of Claim (2) Using the definition of F_P and some tedious algebra, for arbitrary z and z' we can eventually write:

$$F_P(z) - F_P(z') = \left(1 - \frac{\alpha(z)}{P}\right)^2 (z - z') + \frac{\alpha(z) - \alpha(z')}{P} \left(\frac{\alpha(z)}{P}(z' + v_c) + \frac{\alpha(z')}{P}v_c - \left(2 - \frac{\alpha(z')}{P}\right)z'\right)$$

Next observe that $\alpha(z) - \alpha(z') = \frac{v}{(z'+v_{ta}+v)(z+v_{ta}+v)}(z - z') = \frac{1-\alpha(z)}{z'+v_{ta}+v}(z - z')$. Use this to write:

$$F_P(z) - F_P(z') = \left(\left(1 - \frac{\alpha(z)}{P}\right)^2 + \frac{1 - \alpha(z)}{P} \frac{1 - \alpha(z)}{z' + v_{ta} + v} \left(\frac{\alpha(z)}{P}(z' + v_c) + \frac{\alpha(z')}{P}v_c - \left(2 - \frac{\alpha(z')}{P}\right)z'\right)\right)(z - z')$$

Substitute in $z' = \bar{z}^P$, and use (12) to cancel righthand side terms, arriving at:

$$\begin{aligned} F_{P_n}(z) - \bar{z}^P &= \left(\left(1 - \frac{\alpha(z)}{P}\right)^2 + \frac{v + \bar{z}^n}{\bar{z}^n + v_{ta} + v} \frac{\alpha(z)(1 - \alpha(z))}{P^2}\right)(z - \bar{z}^P) \\ &= (1 - b_P(z))(z - \bar{z}^P) \end{aligned}$$

Expand and solve for $b_P(z)$ to get:

$$b_P(z) = \frac{\alpha(z)}{P} \left(2 - \frac{\alpha(z)}{P} - \frac{v + \bar{z}^n}{\bar{z}^n + v_{ta} + v} \frac{1 - \alpha(z)}{P}\right)$$

Observe that $b_P(z) < \frac{\alpha(z)}{P} \left(2 - \frac{\alpha(z)}{P}\right) \leq 1$ (where the strict inequality follows from $\alpha(z) < 1$). For the lower bound on $b_P(z)$, we have

$$b_P(z) > \frac{\alpha(z)}{P} \left(2 - \frac{1}{P}\right) > \frac{\alpha(0)}{P} > 0$$

Proof of Claim (3) The quadratic equation that defines \bar{z}^P is continuous in P .

Cases 2 and 3 from Table 1: P_n is unbounded and $\lim_{n \rightarrow \infty} \Delta P_n < 2\bar{\alpha}$

Recall that $\Delta P_n = 1 - d_n$ converges because (by assumption) d_n converges. Let $\bar{g} = \lim_{n \rightarrow \infty} \Delta P_n \geq 0$. Let $Z_n = P_n \kappa_n$. We have

$$\begin{aligned} Z_n &= \left(1 + \frac{\Delta P_n}{P_{n-1}}\right) \left(1 - \frac{\alpha_n}{P_n}\right)^2 Z_{n-1} + \frac{\alpha_n^2}{P_n} v_c \\ &= (1 - b_{n-1}) Z_{n-1} + c_{n-1} \end{aligned}$$

where $c_{n-1} = \frac{\alpha_n^2}{P_n} v_c$, and

$$b_{n-1} = \frac{1}{P_{n-1}} \left(2\alpha_n - \Delta P_n - \frac{\alpha_n^2}{P_n}\right)$$

It is immediate that $b_n \rightarrow 0$ and $c_n \rightarrow 0$ because P_n diverges. Furthermore, we have $\frac{c_n}{b_n} \rightarrow \frac{\bar{\alpha}^2 v_c}{2\bar{\alpha} - \bar{g}}$ which is strictly positive and finite. Because $\frac{\alpha_n^2}{P_n} \rightarrow 0$, b_{n-1} is eventually bounded below by (e.g.) $\frac{1}{2} \frac{2\bar{\alpha} - \bar{g}}{P_{n-1}}$. But $\sum_{n=1}^{\infty} \frac{1}{P_n}$ diverges, so $\sum_{n=1}^{\infty} b_n$ must diverge as well. Finally, there exists some \bar{n} such that $b_{n-1} \in (0, 1)$ for all $n \geq \bar{n}$. If we consider the sequence Z_n beginning with the \bar{n} th term $Z_{\bar{n}}$, all of the conditions of Lemma 10 are met, so we have $Z_n \rightarrow \frac{\bar{\alpha}^2 v_c}{2\bar{\alpha} - \bar{g}}$.

Case 4 from Table 1: P_n is unbounded and $\lim_{n \rightarrow \infty} \Delta P_n > 2\bar{\alpha}$

Observe that in this case, $\lim_{n \rightarrow \infty} \Delta P_n = \lim_{n \rightarrow \infty} \frac{P_n}{n} = \bar{g}$. Define $Z_n = P_n^\beta \kappa_n$. We have

$$\begin{aligned} Z_n &= P_n^\beta \kappa_{n-1} \left(1 - \frac{\alpha_n}{P_n}\right)^2 + \frac{\alpha_n^2}{P_n^{2-\beta}} v_c \\ &= \left(1 + \frac{\Delta P_n^\beta}{P_{n-1}^\beta} - \frac{2\alpha_n}{P_n^{1-\beta} P_{n-1}^\beta} + \frac{\alpha_n^2}{P_n^{2-\beta} P_{n-1}^\beta}\right) Z_{n-1} + \frac{\alpha_n^2}{P_n^{2-\beta}} v_c \end{aligned}$$

and so

$$\Delta Z_n = \left(\frac{\Delta(P_n^\beta)}{P_{n-1}^\beta} - \frac{2\alpha_n}{P_n^{1-\beta} P_{n-1}^\beta} + \frac{\alpha_n^2}{P_n^{2-\beta} P_{n-1}^\beta}\right) Z_{n-1} + \frac{\alpha_n^2}{P_n^{2-\beta}} v_c$$

where $\Delta(P_n^\beta) = P_n^\beta - P_{n-1}^\beta$. Expanding $P_n^\beta = (P_{n-1} + \Delta P_n)^\beta$, we have $\Delta(P_n^\beta) = \beta \frac{\Delta P_n}{P_{n-1}} + O\left(\left(\frac{\Delta P_n}{P_{n-1}}\right)^{2-\beta}\right)$.

Since ΔP_n and α_n converge, we have

$$\begin{aligned} \Delta Z_n &= \frac{1}{P_{n-1}} \left(\beta \Delta P_n - \left(\frac{P_{n-1}}{P_n}\right)^{1-\beta} 2\alpha_n + O\left(P_{n-1}^{\beta-1}\right)\right) Z_{n-1} + O\left(P_n^{\beta-2}\right) \\ &= \frac{1}{P_{n-1}} \left(\beta \Delta P_n - 2\alpha_n + O\left(\frac{1}{P_n}\right)\right) Z_{n-1} + O\left(P_n^{\beta-2}\right) \end{aligned}$$

(where the last step expands $\left(\frac{P_{n-1}}{P_n}\right)^{1-\beta} = \left(1 - \frac{\Delta P_n}{P_n}\right)^{1-\beta}$).

For $\beta \in (0, 1)$, the contribution of summed $O\left(P_n^{\beta-2}\right)$ terms to Z_n converges because $O\left(P_n^{\beta-2}\right) = O\left(n^{\beta-2}\right)$ with $\beta - 2 < -1$. Thus we can concentrate on the term in ΔZ_n that is proportional to Z_{n-1} . Suppose that

$\beta\bar{g} - 2\bar{\alpha} > \epsilon > 0$. Then for all n sufficiently large, we have $\beta\Delta P_n - 2\alpha_n + O\left(\frac{1}{P_n}\right) > \frac{\epsilon}{2}$ and $P_{n-1} < 2\bar{g}n$ and therefore

$$\Delta Z_n > \frac{\epsilon}{4\bar{g}} \frac{1}{n} Z_{n-1}$$

Thus if $\beta\bar{g} - 2\bar{\alpha} > 0$, then $Z_n \rightarrow \infty$.

Alternatively, suppose that $\beta\bar{g} - 2\bar{\alpha} < -\epsilon < 0$. Then for n sufficiently large, we have $\beta\Delta P_n - 2\alpha_n + O\left(\frac{1}{P_n}\right) < -\frac{\epsilon}{2}$. In this case, applying Lemma 10 implies $Z_n \rightarrow 0$.

Proposition 9

P_n unbounded

Note that $\left(1 - \frac{\alpha_n}{P_n}\right)^2 \rightarrow 1$. Then, for arbitrary $L > v_{mt}$, there exists some n_L such that $\left(1 - \frac{\alpha_n}{P_n}\right)^2 > 1 - \frac{v_{mt}}{2L}$ for all $n \geq n_L$. Define a new sequence y_n by

$$y_n = \begin{cases} \gamma_n & \text{if } n < n_L \\ \left(1 - \frac{v_{mt}}{2L}\right) y_{n-1} + v_{mt} & \text{if } n \geq n_L \end{cases}$$

By construction, $y_n \leq \gamma_n$ for all $n \geq 1$. Furthermore, $y_n \rightarrow 2L$, so there exists some n'_L such that $y_n > L$ for all $n \geq n'_L$. So *a fortiori*, $\gamma_n > L$ for all $n \geq n'_L$. Since L was arbitrary, this suffices to show that γ_n diverges.

P_n bounded

We use the following lemma adapted from Maxim Engers.

Lemma 11 *Suppose $\{F_n\}_{n \geq 1}$ is a sequence of functions (on a compact subset $Z \subseteq \mathfrak{R}$) that converges uniformly to a contraction mapping F . Let \bar{z} be the unique fixed point of F . For any initial $z_1 \in Z$, define a sequence $\{z_n\}_{n \geq 1}$ by $z_n = F_n(z_{n-1})$. Then $\{z_n\}_{n \geq 1} \rightarrow \bar{z}$.*

Suppose that $P_n \rightarrow N^*$. Let $Z = [v_{mt}, N^*(v_c + v_{mt})]$ and define $F_n : Z \rightarrow Z$ by $F_n(z) = \left(1 - \frac{1}{P_n} \frac{z}{z+v}\right)^2 z + \left(\frac{1}{P_n} \frac{z}{z+v}\right)^2 v_c + v_{mt}$. Let $F = \lim_{n \rightarrow \infty} F_n$, so we have $F(z) = \left(1 - \frac{1}{N^*} \frac{z}{z+v}\right)^2 z + \left(\frac{1}{N^*} \frac{z}{z+v}\right)^2 v_c + v_{mt}$. (The limits on Z are chosen so as to ensure that $F_n(Z) \subseteq Z$ for all $n \geq 1$, as well as for F .) It is straightforward to see that $F_n \rightarrow F$ uniformly, and furthermore, we have the following.

Lemma 12 *The function F , as defined above, is a contraction mapping.*

Thus F has a unique fixed point \bar{z} . We have $\gamma_1 = v_c + v_{mt} \in Z$, and $\gamma_n = F_n(\gamma_{n-1})$ for $n > 1$, so by Lemmas 11 and 12, $\{\gamma_n\}_{n \geq 1} \rightarrow \bar{z}$. Furthermore, because $F(v_{mt}) \neq v_{mt}$, we must have $\bar{z} > v_{mt}$.

9 Supplementary Appendix

9.1 Proofs Omitted from the Main Appendix

Lemma 2 *Suppose $\frac{\bar{\tau}_n}{R_n} \rightarrow K$, with $\eta = \frac{1}{Kv}$ as in the text. Let $H_n = \sum_{j=1}^n \frac{1}{P_j^2}$ if P_n is unbounded and $\frac{R_n}{P_n}$ diverges, or $H_n = \sum_{j=1}^n \frac{1}{R_j P_j}$ otherwise. Then $\frac{\bar{\tau}_n - KR_n}{H_n} \rightarrow L$, for some finite L . Furthermore, $\frac{R_n}{H_n} (\eta v - \kappa_n R_n) \rightarrow L(\eta v)^2$ and $\frac{R_n}{H_n} (\eta - \alpha_n R_n) \rightarrow L\eta^2 v$.*

Proof. Write $\Delta\bar{\tau}_n = \bar{\tau}_n - \bar{\tau}_{n-1}$ and $\Delta R_n = R_n - R_{n-1} = \frac{1}{P_n}$. With an eye toward applying Corollary 3, we recall from the proof of Proposition 2 that $\Delta\bar{\tau}_n$ and ΔR_n satisfy $\Delta\bar{\tau}_n = K_n \Delta R_n = \frac{K_n}{P_n}$, where

$$K_n = \frac{\bar{\tau}_n - \bar{\tau}_{n-1}}{R_n - R_{n-1}} = \left(2 - \frac{1}{P_n} \frac{v_c + \kappa_{n-1}}{v + \kappa_{n-1}}\right) \left(\frac{\bar{\tau}_{n-1}}{v\bar{\tau}_{n-1} + 1 - M_n}\right)$$

with $K_n \rightarrow K$. Then we can write $\Delta\bar{\tau}_n - K\Delta R_n = (K_n - K)\Delta R_n = \frac{1}{P_n}(K_n - K)$. For ease of exposition, separate K_n into $\tilde{K}_n = 2 - \frac{1}{P_n} \frac{v_c + \kappa_{n-1}}{v + \kappa_{n-1}}$ and $\hat{K}_n = \frac{\bar{\tau}_{n-1}}{v\bar{\tau}_{n-1} + 1 - M_n}$, with $\tilde{K}_n \rightarrow \tilde{K}$, $\hat{K}_n \rightarrow \hat{K}$, and $K = \tilde{K}\hat{K}$. We can write $K_n - K = \tilde{K}(D\hat{K}) + \hat{K}(D\tilde{K}) + (D\tilde{K})(D\hat{K})$, where $D\tilde{K} = \tilde{K}_n - \tilde{K}$ and $D\hat{K} = \hat{K}_n - \hat{K}$. Let $h_n = \frac{1}{P_n^2}$ if $\frac{R_n}{P_n}$ diverges, or $h_n = \frac{1}{R_n P_n}$ otherwise. Consider the limit:

$$\lim_{n \rightarrow \infty} \frac{\Delta\bar{\tau}_n - K\Delta R_n}{h_n} = \lim_{n \rightarrow \infty} \frac{\tilde{K}(D\hat{K}) + \hat{K}(D\tilde{K}) + (D\tilde{K})(D\hat{K})}{h_n P_n}$$

Take the numerator term by term. Regardless of whether P_n is bounded, $\hat{K} = \frac{1}{v}$. Thus, $D\hat{K} = \frac{\bar{\tau}_{n-1}}{v\bar{\tau}_{n-1} + 1 - M_n} - \frac{1}{v} = \frac{1}{v} \left(\frac{(M_n - 1)}{v\bar{\tau}_{n-1} + 1 - M_n}\right)$. Both \tilde{K} and M_n are bounded, so the limiting behavior of $\frac{\tilde{K}(D\hat{K})}{h_n P_n}$ is governed by $\frac{P_n}{\bar{\tau}_{n-1}}$ (if $P_n \rightarrow \infty$ and $\frac{R_n}{P_n}$ diverges) or $\frac{R_n}{\bar{\tau}_{n-1}}$ (otherwise). But $\bar{\tau}_{n-1}$ grows at rate R_n , so $\frac{\tilde{K}(D\hat{K})}{h_n P_n}$ tends to either zero (in the first case) or a positive constant (otherwise). Next, \tilde{K} equals $2 - \frac{1}{N^*} \frac{v_c}{v}$ or 2 depending on whether or not P_n is bounded. If P_n is bounded, then we have $\frac{\hat{K}(D\tilde{K})}{h_n P_n} = \hat{K} R_n (D\tilde{K})$, or, with some algebra, $-\frac{1}{P_n v^2} R_n \left(v_c \frac{N^* - P_n}{N^*} + \kappa_{n-1} \frac{v - v_c}{v + \kappa_{n-1}}\right)$. We have $R_n(N^* - P_n) \rightarrow 0$ ($N^* - P_n$ falls exponentially in n , while R_n grows linearly in n) and $R_n \kappa_{n-1}$ tends to a constant, so this term has a finite limit. If P_n is unbounded, then $D\tilde{K} = -\frac{1}{P_n} \frac{v_c + \kappa_{n-1}}{v + \kappa_{n-1}}$, and so $\frac{\hat{K}(D\tilde{K})}{h_n P_n}$ tends to a nonzero constant if $\frac{R_n}{P_n}$ diverges or tends to a constant. If $\frac{R_n}{P_n}$ tends to zero, then the dominant term of $\frac{\hat{K}(D\tilde{K})}{h_n P_n}$ is proportional to $\frac{R_n}{P_n}$, so $\frac{\hat{K}(D\tilde{K})}{h_n P_n}$ tends to zero as well. We omit the argument that $\lim_{n \rightarrow \infty} \frac{(D\tilde{K})(D\hat{K})}{h_n P_n}$ is finite, as it is very similar to the logic for the other two terms. Thus $\lim_{n \rightarrow \infty} \frac{\Delta\bar{\tau}_n - K\Delta R_n}{h_n}$ exists; call that limit L .

Next, apply Corollary 3. Define the strictly increasing sequence $H_n = \sum_{j=1}^n h_n$. To show that H_n is unbounded, first suppose that $\frac{R_n}{P_n}$ does not diverge. Then $H_n = \sum_{j=1}^n \frac{1}{R_j P_j}$ which diverges at rate $\ln R_n$ by Lemma 4. Alternatively, if P_n is unbounded and $\frac{R_n}{P_n}$ diverges, then $H_n = \sum_{j=1}^n \frac{1}{P_j^2} = \sum_{j=1}^n \frac{R_j}{P_j} \frac{1}{R_j P_j}$. In this case, the terms of $\sum_{j=1}^n \frac{1}{P_j^2}$ eventually dominate the terms of $\sum_{j=1}^n \frac{1}{R_j P_j}$ which we know to diverge. Thus H_n diverges in the second case as well. Thus, Corollary 3 can be applied to $\frac{\Delta\bar{\tau}_n - K\Delta R_n}{h_n}$ to arrive at $\lim_{n \rightarrow \infty} \frac{\bar{\tau}_n - K R_n}{H_n} \rightarrow L$.

The result that $\frac{R_n}{H_n}(\eta v - \kappa_n R_n) \rightarrow L(\eta v)^2$ follows by using $\kappa_n R_n \rightarrow K^{-1} = \eta v$. For the third result, use the definition of α_n to get $\frac{R_n}{H_n}(\eta - \alpha_n R_n) = \frac{1}{v} \frac{R_n}{H_n}(\eta v - \kappa_n R_n) + \frac{1}{v} \frac{(\kappa_{n-1} R_n)(\alpha_n R_n)}{H_n} - \frac{1}{v} \frac{\Delta \kappa_n R_n^2}{H_n}$. The second term tends to zero (since its numerator converges and its denominator diverges). For the third term, note that $\Delta \kappa_n R_n^2 = \Delta \tau_n \left(\frac{R_n}{\tau_{n-1}}\right) \left(\frac{R_n}{\tau_n}\right)$, with the terms in parentheses convergent and $\Delta \tau_n$ bounded. Thus, $\frac{1}{v} \frac{\Delta \kappa_n R_n^2}{H_n}$ tends to zero as well. The third result follows. ■

Lemma 3 $R_n P_n \geq n$ for all $n \geq 1$.

Proof. $R_1 P_1 = 1$, so this is satisfied for $n = 1$. For $n > 1$ we also have

$$\begin{aligned} R_n P_n - R_{n-1} P_{n-1} &= (R_n - R_{n-1}) P_n + R_{n-1} (P_n - P_{n-1}) \\ &= 1 + R_{n-1} (P_n - P_{n-1}) \geq 1 \quad , \end{aligned}$$

which suffices to prove the claim. ■

Lemma 4 (*Discrete approximation of $\ln R_n$*) Define $S_n = \sum_{k=1}^n \frac{1}{R_k P_k}$. The sequence $\{S_n - \ln R_n\}_{n \geq 1}$ converges.

Proof. Let $T_n = S_n - \ln R_n$, and $\Delta T_n = T_n - T_{n-1}$ for $n > 1$. We have

$$\begin{aligned} \Delta T_n &= \frac{1}{R_n P_n} + \ln \frac{R_{n-1}}{R_n} \\ &= \frac{1}{R_n P_n} + \ln \left(1 - \frac{1}{R_n P_n} \right) \end{aligned}$$

For $n > 1$, $\frac{1}{R_n P_n} \in (0, 1)$, and the function $x + \ln(1 - x) = -\left(\frac{x^2}{2} + \frac{x^3}{3} + \frac{x^4}{4} + \dots\right)$ is strictly negative for $x \in (0, 1)$, so we have $\Delta T_n < 0$.

Next, observe that we can equivalently write T_n as $T_n = \left(\frac{1}{R_n P_n} - \ln R_1\right) + \sum_{k=1}^{n-1} \left(\frac{1}{R_k P_k} - \ln \left(\frac{R_{k+1}}{R_k}\right)\right)$. For the summand, we have

$$\begin{aligned} \frac{1}{R_k P_k} - \ln \left(\frac{R_{k+1}}{R_k}\right) &= \frac{1}{R_k P_k} - \ln \left(1 + \frac{1}{R_k P_{k+1}}\right) \\ &\geq \frac{1}{R_k P_k} - \ln \left(1 + \frac{1}{R_k P_k}\right) \\ &> 0 \end{aligned}$$

where the second line uses $P_{k+1} \geq P_k$ and the third line uses the fact that $x - \ln(1 + x)$ is strictly positive for $x > 0$. Thus we have $T_n > \frac{1}{R_n P_n} - \ln R_1 = \frac{1}{R_n P_n} > 0$. But then, since T_n is decreasing and bounded below, it must converge. ■

Lemma 5 (*Linear approximation of $\ln(1 - \alpha_n/P_n)$*) Define c_k by $\ln(1 - \alpha_k/P_k) = -\alpha_k/P_k + c_k$. For arbitrary $m > 1$, the sum $\sum_{k=m+1}^{\infty} c_k$ converges.

Proof. It suffices to show that $\sum_{k=m+1}^{\infty} |c_k|$ converges. Write $z_k = \alpha_k/P_k$, so $c_k = z_k + \ln(1 - z_k)$. Observe that $z_k < 1$ (because $\alpha_k < 1$ and $P_k \geq 1$), that $z_k \rightarrow 0$ with k , and $z_k R_k P_k = \alpha_k R_k \rightarrow \eta$, where η is defined in the text. (This last limit follows from $\alpha_k R_k \rightarrow \frac{1}{v} \lim_{k \rightarrow \infty} \kappa_k R_k = \eta$.) For all k sufficiently large, we have $z_k < \frac{2\eta}{R_k P_k} \leq \frac{2\eta}{n}$, by the convergence of $z_k R_k P_k$ and Lemma 3 respectively. Furthermore, the function $|z + \ln(1 - z)|$ is increasing on $(0, 1)$, so we have $|c_k| < \left|\frac{2\eta}{k} + \ln\left(1 - \frac{2\eta}{k}\right)\right|$ for k sufficiently large. Thus, to prove the lemma, it suffices to show that $\sum_{k=m+1}^{\infty} \left|\frac{2\eta}{k} + \ln\left(1 - \frac{2\eta}{k}\right)\right|$ converges. This follows by applying the Integral Test using the test function $f(k) = \frac{2\eta}{k} + \ln\left(1 - \frac{2\eta}{k}\right)$. ■

Lemma 6 For arbitrary $m \geq 1$ and $q \in (0, 1)$, the sum $\sum_{k=m+1}^{\infty} \frac{1}{R_k^{2-q} P_k}$ converges.

Proof. Note that $R_k^{2-q} P_k \rightarrow \infty$ by Lemma 3. Let $c_k = \frac{1}{R_k^{2-q} P_k}$, so the sum of interest is $\sum_{k=m+1}^{\infty} c_k$. For the purpose of comparison, define $c'_k = R_{k-1}^{q-1} - R_k^{q-1}$. and observe that $\sum_{k=m+1}^{\infty} c'_k = R_m^{q-1} - \lim_{k \rightarrow \infty} R_k^{q-1} = R_m^{q-1}$ is convergent. Write $R_{k-1}^{q-1} = \left(R_k - \frac{1}{P_k}\right)^{q-1} = R_k^{q-1} \left(1 - \frac{1}{R_k P_k}\right)^{q-1}$ and take a binomial expansion to get $R_{k-1}^{q-1} = R_k^{q-1} \left(1 + \frac{1-q}{R_k P_k} + o\left(\frac{1}{R_k P_k}\right)\right)$. Use this to write c'_k as $c'_k = \frac{1-q}{R_k^{2-q} P_k} + o\left(\frac{1}{R_k^{2-q} P_k}\right)$. Comparing c_k to c'_k , we have $\lim_{k \rightarrow \infty} c'_k/c_k = 1 - q$, so the convergence of $\sum_{k=m+1}^{\infty} c_k$ is implied by the limit ratio test. ■

The next three lemmas support Proposition 5.

Lemma 7 $\frac{\hat{\kappa}_n}{\hat{\rho}_n}$ is bounded.

Proof. Write $\Delta\left(\frac{\hat{\kappa}_n}{\hat{\rho}_n}\right) = \frac{\hat{\kappa}_n}{\hat{\rho}_n} - \frac{\hat{\kappa}_{n-1}}{\hat{\rho}_{n-1}} = \frac{\hat{\kappa}_{n-1}}{\hat{\rho}_n} \left(\frac{\Delta\hat{\kappa}_n}{\hat{\kappa}_{n-1}} - \frac{\Delta\hat{\rho}_n}{\hat{\rho}_{n-1}} \right)$. We will show (i) that $\Delta\left(\frac{\hat{\kappa}_n}{\hat{\rho}_n}\right)$ is negative if $\frac{\hat{\kappa}_n}{\hat{\rho}_n}$ is sufficiently large, and (ii), $\Delta\left(\frac{\hat{\kappa}_n}{\hat{\rho}_n}\right)$ is bounded. This will suffice to show $\frac{\hat{\kappa}_n}{\hat{\rho}_n}$ has an upper bound. For (i), use the expressions above to write

$$\begin{aligned} \Delta\left(\frac{\hat{\kappa}_n}{\hat{\rho}_n}\right) &= -\frac{\hat{\kappa}_{n-1}}{\hat{\rho}_n} \frac{1}{P_n} \left(\frac{1}{\bar{\kappa}_{n-1}+v} \left(\frac{1}{S}\hat{\kappa}_{n-1} + \frac{S-1}{S} \left(v \left(1 - \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}} \right) + \hat{\rho}_{n-1} \right) \right) \right. \\ &\quad \left. - \hat{\alpha}_n \left(2 - \frac{1}{P_n} \frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} (v_c + \hat{\rho}_{n-1}) \right) \right) \\ &< -\frac{\hat{\kappa}_{n-1}}{\hat{\rho}_n} \frac{1}{P_n} \left(\frac{1}{\bar{\kappa}_{n-1}+v} \left(\frac{1}{S}\hat{\kappa}_{n-1} + \frac{S-1}{S} \left(v \left(1 - \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}} \right) + \hat{\rho}_{n-1} \right) \right) - 2\hat{\alpha}_n \right) \\ &= -\frac{\hat{\kappa}_{n-1}}{\hat{\rho}_n} \frac{1}{P_n} \frac{1}{\bar{\kappa}_{n-1}+v} \left(\frac{1}{S}\hat{\kappa}_{n-1} + \frac{S-1}{S} \left(v \left(1 - \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}} \right) + \hat{\rho}_{n-1} \right) - 2\bar{\kappa}_{n-1} \right) \\ &= -\frac{\hat{\kappa}_{n-1}}{\hat{\rho}_n} \frac{1}{P_n} \frac{1}{\bar{\kappa}_{n-1}+v} \left(\frac{S-1}{S} v \left(1 - \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}} \right) - \bar{\kappa}_{n-1} \right) \end{aligned}$$

But $\bar{\kappa}_{n-1}$ tends to zero with n (since $\hat{\kappa}_{n-1}$ and $\hat{\rho}_{n-1}$ do). Pick any arbitrary $L > 1$ and $\varepsilon_L = \frac{1}{2} \frac{S-1}{S} v \left(1 - \frac{1}{L} \right)$. There is some n^* such that $\bar{\kappa}_{n-1} < \varepsilon_L$ for all $n \geq n^*$. Then, for all $n \geq n^*$, if $\frac{\hat{\kappa}_{n-1}}{\hat{\rho}_{n-1}} > L$, then $\Delta\left(\frac{\hat{\kappa}_n}{\hat{\rho}_n}\right)$ is negative.

For (ii), note that $\frac{S-1}{S} v \left(1 - \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}} \right) - \varepsilon_L > -\frac{S-1}{S} v \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}}$, so for all $n \geq n^*$, we have $\left(\frac{S-1}{S} v \left(1 - \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}} \right) - \bar{\kappa}_{n-1} \right) > -\frac{S-1}{S} v \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}}$ and therefore

$$\begin{aligned} \Delta\left(\frac{\hat{\kappa}_n}{\hat{\rho}_n}\right) &< \frac{\hat{\kappa}_{n-1}}{\hat{\rho}_n} \frac{1}{P_n} \frac{1}{\bar{\kappa}_{n-1}+v} \left(\frac{S-1}{S} v \frac{\hat{\rho}_{n-1}}{\hat{\kappa}_{n-1}} \right) \\ &= \frac{\hat{\rho}_{n-1}}{\hat{\rho}_n} \frac{1}{P_n} \frac{v}{\bar{\kappa}_{n-1}+v} \frac{S-1}{S} \end{aligned}$$

Take $n^{**} \geq n^*$ large enough such that $\frac{\hat{\rho}_{n-1}}{\hat{\rho}_n} < 2$ for all $n \geq n^{**}$. (We can do this because $\frac{\hat{\rho}_{n-1}}{\hat{\rho}_n} = \frac{\hat{\rho}_{n-1}}{(1-\frac{\hat{\alpha}_n}{P_n})^2 \hat{\rho}_{n-1} + (\frac{\hat{\alpha}_n}{P_n})^2 v_c} < \frac{1}{(1-\frac{\hat{\alpha}_n}{P_n})^2}$ which tends to 1 since $\hat{\alpha}_n \rightarrow 0$ and $P_n \geq 1$.) Then we have $\Delta\left(\frac{\hat{\kappa}_n}{\hat{\rho}_n}\right) < 2 \frac{S-1}{S}$ for all $n \geq n^{**}$. Together, (i) and (ii) imply that for all $n \geq n^*$, $\frac{\hat{\kappa}_{n-1}}{\hat{\rho}_{n-1}}$ cannot exceed $L + 2 \frac{S-1}{S}$. Thus, $\frac{\hat{\kappa}_{n-1}}{\hat{\rho}_{n-1}}$ has a finite upper bound. ■

Lemma 8 $\frac{1}{\hat{\rho}_n R_n}$ is bounded. (That is, $\hat{\rho}_n$ tends to zero no faster than rate R_n .)

Proof. We proceed as in the proof of Proposition 2. Set up the ratio $\frac{1/\hat{\rho}_n}{R_n}$, where the increments of the numerator and denominator are $\Delta\left(\frac{1}{\hat{\rho}_n}\right) = \frac{1}{\hat{\rho}_n} - \frac{1}{\hat{\rho}_{n-1}} = -\frac{\Delta\hat{\rho}_n}{\hat{\rho}_n \hat{\rho}_{n-1}}$ and $\Delta R_n = \frac{1}{P_n}$. The proof of Theorem 1 can easily be extended to show that if the sequence $\frac{\Delta\left(\frac{1}{\hat{\rho}_n}\right)}{\Delta R_n}$ has an upper bound, then $\frac{1/\hat{\rho}_n}{R_n}$ has an upper bound as well. The sequence $\frac{\Delta\left(\frac{1}{\hat{\rho}_n}\right)}{\Delta R_n}$ is given by

$$\frac{\Delta\left(\frac{1}{\hat{\rho}_n}\right)}{\Delta R_n} = \frac{\hat{\rho}_{n-1}}{\hat{\rho}_n} \frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} \left(2 - \frac{1}{P_n} \frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} (\hat{\rho}_{n-1} + v_c) \right)$$

We have $\hat{\alpha}_n < \bar{\kappa}_{n-1} = \frac{1}{S}\hat{\kappa}_{n-1} + \frac{S-1}{S}\hat{\rho}_{n-1}$, so $\frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}} < \frac{1}{S} \frac{\hat{\kappa}_{n-1}}{\hat{\rho}_{n-1}} + \frac{S-1}{S}$. Since $\frac{\hat{\kappa}_{n-1}}{\hat{\rho}_{n-1}}$ is bounded by Lemma 7, $\frac{\hat{\alpha}_n}{\hat{\rho}_{n-1}}$ is as well. But then $\frac{\hat{\rho}_{n-1}}{\hat{\rho}_n}$ bounded, $P_n \geq 1$, and $\hat{\rho}_{n-1} \rightarrow 0$ imply that $\frac{\Delta\left(\frac{1}{\hat{\rho}_n}\right)}{\Delta R_n}$ is bounded. Thus, $\frac{1}{\hat{\rho}_n R_n}$ is bounded as claimed, and so $\hat{\rho}_n > \frac{C}{R_n}$ for some constant C . ■

Lemma 9 At least one of the following must hold: either $\frac{\hat{\kappa}_n - \hat{\rho}_n}{\hat{\rho}_n} \rightarrow 0$ or $(\hat{\kappa}_n - \hat{\rho}_n) R_n \rightarrow 0$.

Proof. We consider the limit of $(\hat{\kappa}_n - \hat{\rho}_n) R_n = \frac{R_n}{1/(\hat{\kappa}_n - \hat{\rho}_n)}$. By Theorem 1, this limit is equal to $\lim_{n \rightarrow \infty} \frac{\Delta R_n}{\Delta\left(\frac{1}{\hat{\kappa}_n - \hat{\rho}_n}\right)}$, where $\Delta\left(\frac{1}{\hat{\kappa}_n - \hat{\rho}_n}\right) = \frac{1}{\hat{\kappa}_n - \hat{\rho}_n} - \frac{1}{\hat{\kappa}_{n-1} - \hat{\rho}_{n-1}}$, if the latter limit exists. We have

$$\begin{aligned} \frac{\Delta R_n}{\Delta\left(\frac{1}{\hat{\kappa}_n - \hat{\rho}_n}\right)} &= -\frac{(\hat{\kappa}_{n-1} - \hat{\rho}_{n-1})(\hat{\kappa}_n - \hat{\rho}_n)}{P_n(\Delta\hat{\kappa}_n - \Delta\hat{\rho}_n)} \\ &< \frac{(\hat{\kappa}_{n-1} - \hat{\rho}_{n-1})(\hat{\kappa}_n - \hat{\rho}_n)}{\hat{\kappa}_{n-1} - \hat{\alpha}_n v - 2\hat{\alpha}_n \hat{\rho}_{n-1}} \end{aligned}$$

where the second line drops the term $(\hat{\alpha}_n/P_n)^2(\hat{\rho}_{n-1} + v_c)$ from $\Delta\hat{\rho}_n$. Rearranging the denominator, using the expressions for $\hat{\alpha}_n$ and $\hat{\kappa}_{n-1}$, we can write

$$\hat{\kappa}_{n-1} - \hat{\alpha}_n v - 2\hat{\alpha}_n \hat{\rho}_{n-1} = \left(\frac{S-1}{S} + \frac{\hat{\alpha}_n}{S}\right)(\hat{\kappa}_{n-1} - \hat{\rho}_{n-1}) - \hat{\alpha}_n \hat{\rho}_{n-1}$$

Using this in the inequality above, we have

$$\frac{\Delta R_n}{\Delta\left(\frac{1}{\hat{\kappa}_n - \hat{\rho}_n}\right)} < \frac{\hat{\kappa}_n - \hat{\rho}_n}{\left(\frac{S-1}{S} + \frac{\hat{\alpha}_n}{S}\right) - \frac{\hat{\alpha}_n \hat{\rho}_{n-1}}{\hat{\kappa}_{n-1} - \hat{\rho}_{n-1}}}$$

The numerator on the righthand side tends to zero (because $\hat{\kappa}_n$ and $\hat{\rho}_n$ do). In the denominator, $\left(\frac{S-1}{S} + \frac{\hat{\alpha}_n}{S}\right) \rightarrow \frac{S-1}{S}$. Thus, as long as $\frac{\hat{\alpha}_n \hat{\rho}_{n-1}}{\hat{\kappa}_{n-1} - \hat{\rho}_{n-1}} \not\rightarrow \frac{S-1}{S}$, we have $\frac{\Delta R_n}{\Delta\left(\frac{1}{\hat{\kappa}_n - \hat{\rho}_n}\right)} \rightarrow 0$, which implies $(\hat{\kappa}_n - \hat{\rho}_n) R_n \rightarrow 0$. Alternatively, consider the possibility that we do have $\frac{\hat{\alpha}_n \hat{\rho}_{n-1}}{\hat{\kappa}_{n-1} - \hat{\rho}_{n-1}} \rightarrow \frac{S-1}{S}$. But together with $\hat{\alpha}_n \rightarrow 0$, this implies that $\frac{\hat{\kappa}_n - \hat{\rho}_n}{\hat{\rho}_n} \rightarrow 0$. ■

Lemma 10 Define a sequence Z_n by

$$\begin{aligned} Z_1 &= Z_{init} > 0 \\ Z_n &= (1 - b_{n-1})Z_{n-1} + c_{n-1} \quad \text{if } n > 1 \end{aligned}$$

Write the increments of this sequence as $\Delta Z_n \equiv Z_n - Z_{n-1} = -b_{n-1}\left(Z_{n-1} - \frac{c_{n-1}}{b_{n-1}}\right)$. Suppose that b_n and c_n are sequences with $0 < b_n < 1$ and $\sum_{n=1}^{\infty} b_n$ divergent, $c_n \rightarrow 0$, and $\frac{c_n}{b_n} \rightarrow Z \geq 0$. Then $Z_n \rightarrow Z$.

Proof. Define $Y_n = Z_n - Z$, so that $Y_n = (1 - b_{n-1})Y_{n-1} + \gamma_{n-1}b_{n-1}$, where $\gamma_{n-1} = \frac{c_{n-1}}{b_{n-1}} - Z \rightarrow 0$. Iterate forward to get

$$Y_n = Y_1 \prod_{j=1}^{n-1} (1 - b_j) + \sum_{j=1}^{n-2} \left(\gamma_j b_j \prod_{k=j+1}^{n-1} (1 - b_k) \right) + \gamma_{n-1} b_{n-1}$$

Define $S_n = \prod_{j=1}^n \frac{1}{1-b_j}$ for $n \geq 1$, with $S_0 = 1$. We can write

$$Y_n = \frac{Y_1}{S_{n-1}} + \frac{1}{S_{n-1}} \sum_{j=1}^{n-1} \gamma_j b_j S_j$$

Define $\Delta S_n = S_n - S_{n-1}$, and observe that $\Delta S_n = b_n S_n$, so we have

$$Y_n = \frac{Y_1}{S_{n-1}} + \frac{1}{S_{n-1}} \sum_{j=1}^{n-1} \gamma_j \Delta S_j$$

We claim, deferring a proof for the moment, that $S_n \rightarrow \infty$, so the first term above converges to zero. Furthermore, we can apply Theorem 1 to the second term, yielding

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{\sum_{j=1}^{n-1} \gamma_j \Delta S_j}{S_{n-1}} &= \lim_{n \rightarrow \infty} \frac{\gamma_{n-1} \Delta S_{n-1}}{\Delta S_{n-1}} \\ &= \lim_{n \rightarrow \infty} \gamma_{n-1} = 0 \end{aligned}$$

Thus, $Y_n \rightarrow 0$ and so $Z_n \rightarrow Z$.

Finally, to show that $S_n \rightarrow \infty$, note that $\ln S_n = -\sum_{j=1}^n \ln(1 - b_j)$. Because $b_j < 1$, we have $\ln(1 - b_j) < -b_j$, so $\ln S_n > \sum_{j=1}^n b_j$. But the sum on the righthand side diverges by assumption, so we are done. ■

Lemma 11 Suppose $\{F_n\}_{n \geq 1}$ is a sequence of functions (on a compact subset $Z \subseteq \mathfrak{R}$) that converges uniformly to a contraction mapping F . Let \bar{z} be the unique fixed point of F . For any initial $z_1 \in Z$, define a sequence $\{z_n\}_{n \geq 1}$ by $z_n = F_n(z_{n-1})$. Then $\{z_n\}_{n \geq 1} \rightarrow \bar{z}$.

Proof. For $n \geq 1$, let $\xi_n = |z_n - \bar{z}|$. Then,

$$\xi_n = |F_n(z_{n-1}) - \bar{z}| \leq |F_n(z_{n-1}) - F(z_{n-1})| + |F(z_{n-1}) - \bar{z}| \leq \varepsilon_n + c\xi_{n-1}$$

where $c < 1$ is the modulus of contraction and $\varepsilon_n = \sup_{z \in Z} |F_n(z) - F(z)|$. Thus $\xi_n \leq \varepsilon_n + c\xi_{n-1}$, where $\varepsilon_n \rightarrow 0$ by the uniform convergence of $\{F_n\}_{n \geq 1}$. Thus, $\xi_n \rightarrow 0$. ■

Lemma 12 The function $F : Z \rightarrow Z$, with $Z = [v_{mt}, N(v_c + v_{mt})]$ and $F(z) = \left(1 - \frac{1}{N} \frac{z}{z+v}\right)^2 z + \left(\frac{1}{N} \frac{z}{z+v}\right)^2 v_c + v_{mt}$, is a contraction mapping.

Proof. Fix arbitrary $y, z \in Z$ with $y < z$. Write $\alpha_y = \frac{y}{y+v}$, $\alpha_z = \frac{z}{z+v}$, $\bar{\alpha} = \frac{1}{2}(\alpha_y + \alpha_z)$, and $\delta = N - 1 > 0$. Set $c_1 = \frac{\delta v_{mt}}{N^2(v_{mt}+v)} < 1$ and $c_2 = \frac{1}{2}$.

$$\frac{F(z) - F(y) < (1 - c_1)(z - y)}{\text{We can write}}$$

We can write

$$\begin{aligned} F(z) - F(y) &= \left(1 - \frac{\bar{\alpha}}{N}\right)^2 (z - y) + \frac{v_c}{N^2} (\alpha_z^2 - \alpha_y^2) \\ &\quad + \left(\left(1 - \frac{\alpha_z}{N}\right)^2 - \left(1 - \frac{\bar{\alpha}}{N}\right)^2\right) z + \left(\left(1 - \frac{\bar{\alpha}}{N}\right)^2 - \left(1 - \frac{\alpha_y}{N}\right)^2\right) y \end{aligned}$$

The last two terms are negative (because $\alpha_z > \bar{\alpha} > \alpha_y$). For the second term, use $\alpha_z - \alpha_y = \frac{v}{(y+v)(z+v)}(z - y)$ to get

$$\frac{v_c}{N^2} (\alpha_z^2 - \alpha_y^2) = \frac{2\bar{\alpha}}{N^2} \frac{v_c v}{(y+v)(z+v)} (z - y) \leq \frac{2\bar{\alpha}}{N^2} (1 - \alpha_y)(1 - \alpha_z)(z - y)$$

Consolidating, we have

$$F(z) - F(y) \leq \left(\left(1 - \frac{\bar{\alpha}}{N}\right)^2 + \frac{2\bar{\alpha}}{N^2} (1 - \alpha_y)(1 - \alpha_z)\right) (z - y)$$

Expand the factor multiplying $(z - y)$ on the righthand to get

$$\begin{aligned}
1 - 2\frac{\bar{\alpha}}{N} + \frac{\bar{\alpha}^2 + 2\bar{\alpha}(1 - \alpha_y)(1 - \alpha_z)}{N^2} &= 1 - \frac{\bar{\alpha}}{N^2} (2N - (\bar{\alpha} + 2(1 - \alpha_y)(1 - \alpha_z))) \\
&= 1 - \frac{\bar{\alpha}}{N^2} \left(2N - 2 + (3\bar{\alpha} - 2\bar{\alpha}^2) + \frac{1}{2}(\alpha_z - \alpha_y)^2 \right) \\
&\leq 1 - \frac{\bar{\alpha}}{N^2} (2\delta) \\
&\leq 1 - \frac{2\delta v_{mt}}{N^2(v_{mt} + v)} < 1 - c_1
\end{aligned}$$

$$\frac{F(z) - F(y)}{z - y} > -(1 - c_2)$$

If we define $\tilde{F}(\gamma) = F(\gamma) + \gamma$, it suffices to show that $\tilde{F}(z) - \tilde{F}(y) > c_2(z - y)$. For this, it suffices to have $\tilde{F}' > c_2$. Differentiate \tilde{F} to get

$$\begin{aligned}
\frac{d\tilde{F}(\gamma)}{d\gamma} &= 1 + \left(1 - \frac{\alpha}{N}\right)^2 + 2\left(\frac{v_c\alpha}{N^2} - \frac{\gamma}{N}\left(1 - \frac{\alpha}{N}\right)\right) \frac{d\alpha}{d\gamma} \\
&\geq 1 + \left(1 - \frac{\alpha}{N}\right)^2 - 2\frac{\gamma}{N}\left(1 - \frac{\alpha}{N}\right) \frac{d\alpha}{d\gamma} \\
&= 1 + \left(1 - \frac{\alpha}{N}\right)^2 - \frac{2}{N}\left(1 - \frac{\alpha}{N}\right) \frac{\gamma v}{(\gamma + v^2)} \\
&= 1 + \left(1 - \frac{\alpha}{N}\right)^2 - \frac{2}{N}\left(1 - \frac{\alpha}{N}\right) \alpha(1 - \alpha)
\end{aligned}$$

where $\alpha = \frac{\gamma}{\gamma + v}$. Using $\alpha(1 - \alpha) \leq \frac{1}{4}$ and $N > 1$, the negative term is bounded by $-\frac{1}{2}$. Thus $\frac{d\tilde{F}(\gamma)}{d\gamma} \geq \frac{1}{2} + \left(1 - \frac{\alpha}{N}\right)^2 > \frac{1}{2}$.

So to summarize, $|F(z) - F(y)| < (1 - \min(c_1, c_2))|z - y|$, so $F(\gamma)$ is a contraction mapping. ■

9.2 Results for the Non-uniform Exit Models in Section 6

9.2.1 Analysis for Section 6.1, Case 2: Gradual exit of the first cohort

The Equations of Motion

Denote the squared error in cohort n 's action by $\gamma_n \equiv E((x_n - \theta)^2)$. As in the text, the error in the observation pool after cohort n acts is $\kappa_n = E((\bar{x}_n - \theta)^2)$. We have $\gamma_1 = v_c$ and the covariance of the errors in x_1 and x_n is defined to be $\rho_n \equiv E((x_1 - \theta)(x_n - \theta))$. By assumption, the observation pool evolves according to

$$\bar{x}_{n-1} = (1 - w_{n-1})x_{n-1} + w_{n-1}x_1$$

so we have

$$\kappa_{n-1} = (1 - w_{n-1})^2 \gamma_{n-1} + 2w_{n-1}(1 - w_{n-1})\rho_{n-1} + w_{n-1}^2 v_c$$

An individual agent n_i who observes \bar{x}_{n-1} and s_{n_i} chooses the optimal action

$$x_{n_i} = \alpha_n s_{n_i} + (1 - \alpha_n) \bar{x}_{n-1} \quad , \quad \text{where } \alpha_n = \frac{\kappa_{n-1}}{\kappa_{n-1} + v}$$

Averaging over cohort n individuals, we have

$$x_n = \alpha_n s_n + (1 - \alpha_n) \bar{x}_{n-1} \quad \text{and} \quad \gamma_n = \alpha_n^2 v_c + (1 - \alpha_n)^2 \kappa_{n-1}$$

With algebra, this reduces to

$$\gamma_n = \alpha_n \bar{v}_n$$

where $\bar{v}_n = \alpha_n v_c + (1 - \alpha_n) v$ is a weighted average of the individual and cohort error variances. To characterize the covariance term, write x_n recursively as

$$x_n = \alpha_n s_n + (1 - \alpha_n) (1 - w_{n-1}) x_{n-1} + (1 - \alpha_n) w_{n-1} x_1$$

Then we can compute

$$\begin{aligned} \rho_n &= E((x_1 - \theta)(x_n - \theta)) \\ &= (1 - \alpha_n) (1 - w_{n-1}) \rho_{n-1} + (1 - \alpha_n) w_{n-1} v_c \end{aligned}$$

using the fact that $\gamma_1 = v_c$ and s_n is uncorrelated with the earlier errors. The equations for κ_n and ρ_n , along with the expressions for γ_n and α_n , and the exogenous sequence w_n , characterize the evolution of the model. We can substitute in to express the equations of motion as⁵²:

$$\kappa_n = \frac{\bar{v}_n}{\kappa_{n-1} + v} (1 - w_n)^2 \kappa_{n-1} + 2w_n (1 - w_n) \rho_n + w_n^2 v_c \quad (13)$$

$$\rho_n = \frac{v}{\kappa_{n-1} + v} ((1 - w_{n-1}) \rho_{n-1} + w_{n-1} v_c) \quad (14)$$

An informal derivation of the speed of learning

We focus on the case in which $a < \frac{3}{2}$, so learning is slow. Define sequences $Z_n = n^r \kappa_n$ and $\tilde{Z}_n = n^{\tilde{r}} \rho_n$. A full proof of Result 1 would demonstrate that (for the correct exponents r and \tilde{r}), Z_n and \tilde{Z}_n converge to strictly positive constants. In the analysis below, we *assume* that Z_n and \tilde{Z}_n converge for some (undetermined) values of r and $\tilde{r} \in (0, 1)$. Then we use logic similar to Lemma 10 to derive necessary conditions that r and \tilde{r} must satisfy. This suffices to pin down their values.

Suppose that r and \tilde{r} are such that $Z_n \rightarrow Z > 0$ and $\tilde{Z}_n \rightarrow \tilde{Z} > 0$. The first step is to express Z_n and \tilde{Z}_n in a form (using (13) and (14)) that allows us to appeal to Lemma 10. We have

$$Z_n = (1 - B_{n-1}) Z_{n-1} + C_{n-1} \quad \text{and} \quad \tilde{Z}_n = (1 - \tilde{B}_{n-1}) \tilde{Z}_{n-1} + \tilde{C}_{n-1}$$

where

$$\begin{aligned} B_{n-1} &= 1 - \left(1 - \frac{1}{n}\right)^{-r} \frac{\bar{v}_n}{v} (1 - \alpha_n) (1 - w_n)^2 \\ C_{n-1} &= 2n^{r-\tilde{r}} w_n (1 - w_n) \tilde{Z}_n + n^r w_n^2 v_c \\ \tilde{B}_{n-1} &= 1 - \left(1 - \frac{1}{n}\right)^{-\tilde{r}} (1 - \alpha_n) (1 - w_{n-1}) \\ \tilde{C}_{n-1} &= n^{\tilde{r}} w_{n-1} (1 - \alpha_n) v_c \end{aligned}$$

⁵²For clarity, \bar{v}_n (which is bounded between v_c and v) is left unsubstituted.

Next, we claim that B_{n-1} and \tilde{B}_{n-1} are bounded below by $\frac{1-r}{n}$ and $\frac{1-\tilde{r}}{n}$ respectively, so $\sum_{n=1}^{\infty} B_n$ and $\sum_{n=1}^{\infty} \tilde{B}_n$ diverge. To show this, note that $1 - \alpha_n \geq \frac{n-1}{n}$. (This follows because $\kappa_{n-1} \geq \frac{v_c}{n-1}$, since the righthand side is the full information variance after $n-1$ cohorts. Then use $1 - \alpha_n = \frac{v}{\kappa_{n-1} + v} \geq \frac{(n-1)v}{v_c + (n-1)v} \geq \frac{n-1}{n}$ (since $v_c \leq v$.) Then, because $\frac{\tilde{v}_n}{v} \leq 1$, we have

$$B_{n-1} \geq 1 - \left(1 - \frac{1}{n}\right)^{1-r} \quad \text{and} \quad \tilde{B}_{n-1} \geq 1 - \left(1 - \frac{1}{n}\right)^{1-\tilde{r}}$$

For $r \in (0, 1)$, the function $f(x) = 1 - (1-x)^{1-r}$ is convex, so $f(x) \geq f(0) - f'(0)x = (1-r)x$. Applied to the expressions above, we have $B_{n-1} \geq \frac{1-r}{n}$ and $\tilde{B}_{n-1} \geq \frac{1-\tilde{r}}{n}$.

Next, we argue that $Z_n \rightarrow Z$ implies that $C_{n-1}/B_{n-1} \rightarrow Z$. For this, write $\Delta Z_n = Z_n - Z_{n-1} = B_{n-1} \left(\frac{C_{n-1}}{B_{n-1}} - Z_{n-1} \right)$. If C_{n-1}/B_{n-1} were to diverge or to tend toward a limit different from Z , then the term in parentheses would be bounded away from zero. Given the lower bound on B_{n-1} , this would contradict the assumption that the partial sums of ΔZ_n converge. For the same reason, $\tilde{Z}_n \rightarrow \tilde{Z}$ implies that $\tilde{C}_{n-1}/\tilde{B}_{n-1} \rightarrow \tilde{Z}$.

Now consider the ratios B_{n-1}/C_{n-1} and $\tilde{B}_{n-1}/\tilde{C}_{n-1}$. For the latter, we have

$$\begin{aligned} \frac{\tilde{B}_{n-1}}{\tilde{C}_{n-1}} &= \frac{1 - \left(1 - \frac{1}{n}\right)^{-\tilde{r}} (1 - \alpha_n) (1 - w_{n-1})}{n^{\tilde{r}} w_{n-1} (1 - \alpha_n) v_c} \\ &= \frac{-\frac{\tilde{r}}{n} + \alpha_n + w_{n-1} + [2^{nd} \text{ order terms in } n^{-1}, \alpha_n, \text{ and } w_{n-1}]}{n^{\tilde{r}} w_{n-1} (1 - \alpha_n) v_c} \quad \text{so} \\ \lim_{n \rightarrow \infty} \frac{\tilde{B}_{n-1}}{\tilde{C}_{n-1}} &= \frac{1}{v_c} \lim_{n \rightarrow \infty} \left(-\tilde{r} n^{a-\tilde{r}-1} + \alpha_n n^{a-\tilde{r}} + n^{-\tilde{r}} \right) \end{aligned}$$

using $w_{n-1} = n^{-a} + o(n^{-a})$. Write the middle term as $\alpha_n n^{a-\tilde{r}} = \frac{\alpha_n}{\kappa_{n-1}} Z_{n-1} \left(\frac{n}{n-1} \right)^r n^{a-\tilde{r}-r}$ so we have

$$\lim_{n \rightarrow \infty} \frac{\tilde{B}_{n-1}}{\tilde{C}_{n-1}} = \tilde{Z}^{-1} = -\frac{\tilde{r}}{v_c} \lim_{n \rightarrow \infty} \left(n^{a-\tilde{r}-1} \right) + \frac{Z}{vv_c} \lim_{n \rightarrow \infty} n^{a-\tilde{r}-r}$$

This relationship can only hold with $r, \tilde{r} \in (0, 1)$ and \tilde{Z}^{-1} positive and finite if $a - \tilde{r} - r = 0$. In this case, the first term vanishes, and we have $\tilde{Z}^{-1} = \frac{Z}{vv_c}$. Note that the constraint that $r + \tilde{r} = a$ implies that $\tilde{r} < a$.

Next turn to B_{n-1}/C_{n-1} . Note that $\frac{\tilde{v}_n}{v} = 1 - \left(1 - \frac{v_c}{v}\right) \alpha_n$, so we have

$$\begin{aligned} \frac{B_{n-1}}{C_{n-1}} &= \frac{1 - \left(1 - \frac{1}{n}\right)^{-r} \frac{\tilde{v}_n}{v} (1 - \alpha_n) (1 - w_n)^2}{2n^{r-\tilde{r}} w_n (1 - w_n) \tilde{Z}_n + n^r w_n^2 v_c} \\ &= \frac{-\frac{r}{n} + \left(2 - \frac{v_c}{v}\right) \alpha_n + 2w_n + [2^{nd} \text{ order terms}]}{n^{r-\tilde{r}} w_n \left(2(1 - w_n) \tilde{Z}_n + w_n n^{\tilde{r}} v_c \right)} \end{aligned}$$

The constraint that $\tilde{r} < a$ gives us $w_n n^{\tilde{r}} \rightarrow 0$ (the last term in the denominator), so

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{B_{n-1}}{C_{n-1}} &= \frac{1}{2\tilde{Z}} \lim_{n \rightarrow \infty} \frac{-\frac{r}{n} + \left(2 - \frac{v_c}{v}\right) \alpha_n + 2w_n}{n^{r-\tilde{r}} w_n} \\ &= \frac{1}{2\tilde{Z}} \left(-r \lim_{n \rightarrow \infty} n^{a+\tilde{r}-r-1} + \left(2 - \frac{v_c}{v}\right) \lim_{n \rightarrow \infty} \alpha_n n^{a+\tilde{r}-r} + 2 \lim_{n \rightarrow \infty} n^{\tilde{r}-r} \right) \end{aligned}$$

For the middle term, substitute to get $\alpha_n n^{a+\tilde{r}-r} = \frac{\alpha_n}{\kappa_{n-1}} Z_{n-1} \left(\frac{n}{n-1}\right)^r n^{a+\tilde{r}-2r}$. Thus we have

$$\lim_{n \rightarrow \infty} \frac{B_{n-1}}{C_{n-1}} = Z^{-1} = \frac{1}{2\tilde{Z}} \left(-r \lim_{n \rightarrow \infty} n^{a+\tilde{r}-r-1} + \frac{(2 - \frac{v_c}{v})Z}{v} \lim_{n \rightarrow \infty} n^{a+\tilde{r}-2r} + 2 \lim_{n \rightarrow \infty} n^{\tilde{r}-r} \right)$$

Because $r < 1$, the exponents in the first two terms satisfy $a + \tilde{r} - r - 1 < a + \tilde{r} - 2r$. For this relationship to hold with $Z^{-1} > 0$, we must have either $[a + \tilde{r} - 2r = 0$ and $\tilde{r} - r \leq 0]$, or $[a + \tilde{r} - 2r \leq 0$ and $\tilde{r} - r = 0]$. The second pair of constraints will not work (as it would imply that $r = \tilde{r} > a$). If the first pair of constraints holds, then we will have $Z^{-1} = \frac{(2 - \frac{v_c}{v})Z}{2Zv}$, or $Z = \sqrt{\frac{2Zv}{2 - \frac{v_c}{v}}}$.

In summary, convergence of κ_n and ρ_n to zero at rates n^r and $n^{\tilde{r}}$ is consistent only if the following conditions hold

$$\begin{aligned} a - \tilde{r} - r &= 0 \\ a + \tilde{r} - 2r &= 0 \\ \tilde{r} - r &\leq 0 \end{aligned}$$

Solving these conditions yields the following rates:

$$r = \frac{2}{3}a \quad \text{and} \quad \tilde{r} = \frac{1}{3}a$$

Notice that $r \rightarrow 1$ as $a \rightarrow \frac{3}{2}$. This suggests that learning is no slower than rate n^1 for $a \geq \frac{3}{2}$, and therefore, the second part of Result 1 (because learning cannot be faster than rate n^1). The semi-rigorous approach above could be applied to the $a \geq \frac{3}{2}$ case with a bit more legwork.

9.2.2 Analysis for Section 6.2: Hyperbolic exit

Equations of motion

From the equations for y_n and \bar{x}_n , it is straightforward to derive the following.

$$\omega_n = \left(\frac{\alpha_n}{n}\right)^2 + \left(1 - \frac{1}{n}\right)^2 \omega_{n-1} + \left(\frac{1 - \alpha_n}{n}\right)^2 v_{n-1} + 2\frac{1}{n} \left(1 - \frac{1}{n}\right) (1 - \alpha_n) \rho_{n-1} \quad (15)$$

$$v_n = (\alpha_n \beta)^2 + (1 - \beta)^2 \omega_{n-1} + \beta^2 (1 - \alpha_n)^2 v_{n-1} + 2\beta (1 - \beta) (1 - \alpha_n) \rho_{n-1} \quad (16)$$

$$\rho_n = \frac{\alpha_n^2 \beta}{n} + (1 - \beta) \left(1 - \frac{1}{n}\right) \omega_{n-1} + \frac{\beta}{n} (1 - \alpha_n)^2 v_{n-1} \quad (17)$$

$$+ (1 - \alpha_n) \left(\left(1 - \frac{1}{n}\right) \beta + (1 - \beta) \frac{1}{n} \right) \rho_{n-1}$$

$$\alpha_n = \frac{v_{n-1}}{1 + v_{n-1}}$$

Proof of Proposition 10

If $\beta = 1$, so that each agent observes her immediate predecessor, then action x_n will be the efficient weighting of the first n signals, as discussed in Section 3.1, and the squared error in $y_n = x_n$ will be $\nu_n = \frac{1}{n}$. In this case, the population average action \bar{x}_n and the covariance ρ_n are not particularly germane to the

learning process, but for completeness, we discuss \bar{x}_n :

$$\bar{x}_n = \frac{1}{n} \sum_{j=1}^n \sum_{m=j}^n \frac{s_j}{m}, \text{ so } \omega_n = \frac{1}{n^2} \sum_{j=1}^n \left(\sum_{m=j}^n \frac{1}{m} \right)^2$$

By standard approximation methods, we can show that the summation grows as $2n$ (plus lower order terms), so ω_n also converges to zero at rate n .

If $\beta < 1$, then the system of difference equations can be written

$$\begin{aligned} \Delta\omega_n &= -\frac{2}{n}\omega_{n-1} + \frac{2}{n}\rho_{n-1} - \frac{2}{n}\rho_{n-1}\nu_{n-1} + O\left(\frac{\omega_{n-1}}{n^2}\right) + O\left(\frac{\omega_{n-1}^3}{n}\right) \\ \Delta\nu_n &= (1-\beta)^2\omega_{n-1} - (1-\beta^2)\nu_{n-1} + 2\beta(1-\beta)\rho_{n-1} - \beta^2\nu_{n-1}^2 - 2\beta(1-\beta)\rho_{n-1}\nu_{n-1} + O(\omega_{n-1}^3) \\ \Delta\rho_n &= (1-\beta)\left(1 - \frac{1}{n}\right)\omega_{n-1} + \beta\frac{1}{n}\nu_{n-1} - \left(\frac{\beta}{n} + (1-\beta)\left(1 - \frac{1}{n}\right)\right)\rho_{n-1} \\ &\quad - \beta\frac{\nu_{n-1}^2}{n} - \left(\beta + (1-2\beta)\frac{1}{n}\right)\rho_{n-1}\nu_{n-1} + \beta\rho_{n-1}\nu_{n-1}^2 + O\left(\frac{\omega_{n-1}^3}{n}\right) \end{aligned}$$

where we have used the facts that $\nu_n \leq \omega_n$ and $\rho_n \leq \omega_n$ in consolidating the ‘‘big O’’ terms. As an overview of the proof, notice that lowest order terms of $\Delta\omega_n$ can be written as $\Delta\omega_n \approx -\frac{2}{n}(\omega_{n-1} - \rho_{n-1} + \rho_{n-1}\nu_{n-1})$. By assumption, $\rho_{n-1}\nu_{n-1}$ shrinks at rate R_n^2 ; we will show that $\omega_{n-1} - \rho_{n-1}$ does as well. Then we have $\Delta\omega_n \propto -\frac{1}{nR_n^2} \propto -\frac{\omega_{n-1}^2}{n}$, which is similar to the standard version of model PA with no exit. As earlier, this generates a logarithmic rate of decrease in ω_n .

Now we proceed with the formal proof. Note that for any a , we have $\frac{r_n(a)}{R_n(a)} \rightarrow 0$. For brevity, we will drop the argument to $r_n(a)$ and $R_n(a)$ in most of what follows; R_n should always be interpreted to be the rate at which (by hypothesis) ω_n , ν_n , and ρ_n converge. We define $\tilde{\omega}_n = R_n\omega_n \rightarrow \tilde{\omega}$, $\tilde{\nu}_n = R_n\nu_n \rightarrow \tilde{\nu}$, and $\tilde{\rho}_n = R_n\rho_n \rightarrow \tilde{\rho}$ for some strictly positive $\tilde{\omega}$, $\tilde{\nu}$, and $\tilde{\rho}$. Next, define $z_n = \omega_n - \rho_n$, and observing that $\Delta\rho_n = (1-\beta)(\omega_{n-1} - \rho_{n-1}) - \beta\rho_{n-1}\nu_{n-1} + O\left(\frac{\omega_{n-1}}{n}\right) + O(\omega_{n-1}^3)$, we can write

$$\begin{aligned} z_n - z_{n-1} &= \Delta\omega_n - \Delta\rho_n \\ &= (1-\beta)(\rho_{n-1} - \omega_{n-1}) + \beta\rho_{n-1}\nu_{n-1} + O\left(\frac{\omega_{n-1}}{n}\right) + O(\omega_{n-1}^3) \\ &= -(1-\beta)z_{n-1} + \beta\rho_{n-1}\nu_{n-1} + O\left(\frac{\omega_{n-1}}{n}\right) + O(\omega_{n-1}^3) \end{aligned}$$

Rearranging and multiplying both sides by R_{n-1}^2 , we have

$$\begin{aligned} R_{n-1}^2 z_n - \beta R_{n-1}^2 z_{n-1} &= \beta R_{n-1}^2 \rho_{n-1} \nu_{n-1} + O\left(\frac{R_{n-1}^2 \omega_{n-1}}{n}\right) + O(R_{n-1}^2 \omega_{n-1}^3), \text{ or} \\ R_n^2 z_n \left(1 - 2\frac{R_{n-1}}{R_n} \frac{r_n}{R_n} - \left(\frac{r_n}{R_n}\right)^2\right) - \beta R_{n-1}^2 z_{n-1} &= \beta \tilde{\rho}_{n-1} \tilde{\nu}_{n-1} + O\left(\frac{R_{n-1}}{n} \tilde{\omega}_{n-1}\right) + O(\tilde{\omega}_{n-1}^2 \omega_{n-1}) \end{aligned}$$

Taking limits on both sides yields

$$(1-\beta) \lim_{n \rightarrow \infty} R_n^2 z_n = \beta \tilde{\rho} \tilde{\nu} + K$$

where K is a placeholder for the limit of the $O\left(\frac{R_{n-1}}{n} \tilde{\omega}_{n-1}\right)$ terms. If $\frac{R_n}{n} \rightarrow 0$, as we will show to be true shortly, then K must be zero. Notice that this implies that $\tilde{\omega}_n - \tilde{\rho}_n = R_n z_n \rightarrow 0$.

Next turn back to the equation of motion for $\tilde{\omega}_n$.

$$\begin{aligned}\Delta\tilde{\omega}_n &= R_n\omega_n - R_{n-1}\omega_{n-1} \\ &= R_{n-1}\Delta\omega_n + r_n\omega_{n-1} + r_n\Delta\omega_n\end{aligned}$$

Since the sum $\sum_{m=2}^n \frac{1}{m \ln m}$ diverges in n , the convergence of $\tilde{\omega}_n$ to $\tilde{\omega} > 0$ implies that if the expression $\hat{\omega}_n = (n \ln n) \Delta\tilde{\omega}_n$ converges, it must converge to 0. Using the expression above for $\Delta\tilde{\omega}_n$, we compute the constituent pieces of $\hat{\omega}_n = A_n + B_n + C_n$ as follows:

$$\begin{aligned}A_n &\equiv (n \ln n) R_{n-1} \Delta\omega_n = -2 \frac{\ln n}{R_{n-1}} (R_{n-1} (\tilde{\omega}_{n-1} - \tilde{\rho}_{n-1}) + \tilde{\rho}_{n-1} \tilde{\nu}_{n-1}) + O\left(\frac{\ln n}{n} \tilde{\omega}_{n-1}\right) + O\left(\tilde{\omega}_{n-1}^3 \frac{\ln n}{R_{n-1}^2}\right) \\ B_n &\equiv (n \ln n) r_n \omega_{n-1} = (n \ln n) \frac{r_n}{R_{n-1}} \tilde{\omega}_{n-1} \\ C_n &\equiv (n \ln n) r_n \Delta\omega_n = \frac{r_n}{R_{n-1}} A_n\end{aligned}$$

Consider the following possibilities for a :

$a > 0$

Then nr_n and $\frac{R_n}{\ln n}$ diverge. For term A_n , $\frac{\ln n}{R_n} \rightarrow 0$ and $\tilde{\omega}_{n-1} \rightarrow \tilde{\omega}$ imply that the two $O()$ terms converge to zero. For the first term, recall that $R_{n-1} (\tilde{\omega}_{n-1} - \tilde{\rho}_{n-1}) = R_{n-1}^2 z_{n-1}$, which was shown to converge, and $\tilde{\rho}_{n-1} \tilde{\nu}_{n-1} \rightarrow \tilde{\rho} \tilde{\nu}$. Thus, $\frac{\ln n}{R_n} \rightarrow 0$ implies that $\lim_{n \rightarrow \infty} A_n = 0$ (and therefore, $\lim_{n \rightarrow \infty} C_n = 0$). Next consider B_n . We have $\frac{nr_n}{R_{n-1}} \rightarrow a$, so B_n diverges, and therefore $\hat{\omega}_n \rightarrow \infty$, contradicting the convergence of $\tilde{\omega}$. Thus, ω_n , ν_n , and ρ_n do not converge to zero at any polynomial rate in n .

$a = 0$

In this case $nr_n \rightarrow 1$ and $\frac{R_n}{\ln n} \rightarrow 1$. For term A_n , as above, the two higher order terms vanish as $n \rightarrow \infty$, so we have

$$\begin{aligned}\lim_{n \rightarrow \infty} A_n &= -2 \left(\lim_{n \rightarrow \infty} R_{n-1} (\tilde{\omega}_{n-1} - \tilde{\rho}_{n-1}) + \tilde{\rho} \tilde{\nu} \right) = -\frac{2}{1-\beta} \tilde{\rho} \tilde{\nu}, \\ \lim_{n \rightarrow \infty} B_n &= \tilde{\omega}, \text{ and} \\ \lim_{n \rightarrow \infty} C_n &= 0\end{aligned}$$

This implies $(n \ln n) \Delta\tilde{\omega}_n \rightarrow \tilde{\omega} - \frac{2}{1-\beta} \tilde{\rho} \tilde{\nu}$. This is consistent with the convergence of $\tilde{\omega}_n$ if $\tilde{\omega} = \frac{2}{1-\beta} \tilde{\rho} \tilde{\nu}$, or equivalently, since we have already shown that $\tilde{\omega}_n - \tilde{\rho}_n \rightarrow \tilde{\omega} - \tilde{\rho} = 0$, if $\tilde{\nu} = \frac{1-\beta}{2}$. Thus, if ω_n , ν_n , and ρ_n all converge to zero at common rate R_n , it is possible that R_n could grow as $\ln n$.