Rationally Misplaced Confidence*

Derek Lemoine

University of Arizona, NBER, and CEPR
McClelland Hall 401, 1130 E Helen St, Tucson, AZ, 85721-0108, USA
dlemoine@email.arizona.edu

University of Arizona Working Paper 18-05

October 2021

First version: March 2018

I show that persistent underconfidence and overconfidence arise from rational Bayesian learning when effort and ability are complementary and payoffs are affected by a general class of shocks. Which arises depends on the decision-making environment. If greater effort either reduces or does not affect the variance of outcomes, then agents become underconfident on average, but if greater effort increases that variance by enough, then agents become overconfident on average. The mechanism is that agents learn away overconfidence and underconfidence at asymmetric rates because (i) Bayesian updating requires that their sensitivity to new information depend on their effort choices and (ii) their optimized effort choices depend on beliefs about their own ability. As one implication, management can induce overconfidence in employees by providing feedback that is conditionally vague.

JEL: D83, G41, M54

Keywords: learning, effort, confidence, ability, talent, feedback, appraisals

---

*I gratefully acknowledge an allocation of computer time from the University of Arizona High Performance Computing facility.

Shallow men believe in luck, believe in circumstances: It was somebody's name, or he happened to be there at the time, or, it was so then, and another day it would have been otherwise. Strong men believe in cause and effect.

– Ralph Waldo Emerson, *The Conduct of Life* (1860)

1 Introduction

How much should we read into our successes and failures? Can we reduce our exposure to luck by trying harder? Emerson's "strong men" believe that their efforts have predictable consequences, but Emerson's "shallow men" attribute the outcomes of their efforts to chance. The former learn a lot about themselves by observing the fruits of their efforts, whereas the latter do not infer as much from these outcomes. I here show a surprising result: rational Bayesian agents on average misjudge their own ability, and whether they become overconfident or underconfident on average depends on the predictability of their efforts' consequences.

Overconfidence is now generally recognized as an important factor in many markets. For instance, overconfidence can explain financial market anomalies (Daniel and Hirshleifer, 2015), the persistence of entrepreneurs (Astebro et al., 2014), and corporate investment and merger decisions (Malmendier and Tate, 2005, 2008, 2015). Experimental evidence suggests that underconfidence is also prevalent (e.g., Kirchler and Maciejovsky, 2002; Blavatskyy, 2009; Clark and Friesen, 2009; Urbig et al., 2009; Larkin and Leider, 2012; Murad et al., 2016), and through its link to depression, underconfidence may be especially important for wellbeing (Beck, 2002; Layard and Clark, 2015). Economists have sought to understand how over- and underconfidence can persist in the face of contrary data.

I propose a unified model in which persistent overconfidence and persistent underconfidence endogenously emerge from Bayesian updating by rational agents who have neoclassical utility functions, do not exhibit behavioral biases, never stop learning, and might even initially hold well-calibrated beliefs about their own ability.[1] Agents' rewards depend on effort choices, unknown ability, and unobserved shocks. Effort and ability are complementary, so

[1]Previous literature deviates in one or more of these dimensions, as described in Section 7. I focus on overconfidence in the sense of what Moore and Healy (2008) call "overestimation", reflecting a misjudgment of absolute ability. A distinct literature considers what Moore and Healy (2008) call "overprecision", in which agents underestimate the variance of outcomes (e.g., Daniel et al., 1998; Burnside et al., 2011). And yet another distinct literature studies overconfidence in the sense of what Moore and Healy (2008) call "overplacement" and the psychology literature calls the "better-than-average effect", which refers to the tendency for a majority of the population to judge their own abilities as being better than a majority of the population. I connect the results to overplacement in Appendix C.

agents apply more effort when they think they are of higher ability. Agents learn about their ability from the rewards they observe. Their rewards provide signals of their ability that are drawn from a member of an exponential family of distributions, which encompasses several important named families of distributions (Barndorff-Nielsen, 2006).[2] Agents' effort choices affect how much they learn from each reward because (i) effort and ability are complementary and (ii) effort potentially affects the variance of the unobserved shock.

I show that agents become underconfident on average if effort improves the signal of ability contained in each reward and become overconfident on average if effort confuses the signal of ability contained in each reward. The signal of ability is reward per unit effort, and the variance of the signal is equal to the variance of the unobserved shocks divided by effort squared. Additional effort improves the signal of ability if the shocks' variance decreases in effort or is independent of effort,[3] and it confuses that signal if the shocks' variance increases sufficiently strongly in effort. Agents therefore become underconfident on average if the shocks' variance decreases in effort or is independent of effort, and they become overconfident on average if it increases sufficiently rapidly in effort. We should expect agents to display neither average overconfidence nor average underconfidence only in the knife-edge case where the shocks' variance increases in effort at just the right rate. Empirical and experimental researchers should not be too quick to attribute detected overconfidence or underconfidence to behavioral biases.

As an example, let the variance of rewards be due purely to mean-zero external shocks whose distribution is independent of agents' effort choices. Like Emerson's "strong man", agents understand that their efforts have a consistent effect on outcomes, with luck playing only a supporting role that is independent of effort. For instance, running harder improves their times by a consistent amount that depends on their ability. When agents choose high effort, the observed reward contains a stronger signal of their true ability: because effort and ability are complementary, high effort increases the contribution of ability to outcomes. Imagine that agents' priors are well-calibrated at time 0 (i.e., centered around their true abilities). For some agents, the unobserved shock happens to take on a high value at time

[2]The critical feature of exponential families is that the posterior mean is a weighted average of the prior mean and the signal. For example, this is a well-known feature in normal-normal updating models, and normal distributions are members of an exponential family. The analysis will require that the variance function be quadratic (Morris, 1982; Morris and Lock, 2009), which permits the most prominent exponential families such as the normal, Poisson, binomial, negative binomial, and gamma distributions, with the latter nesting the exponential and chi-square distributions as special cases.

[3]When the variance of rewards is independent of effort, complementarity between effort and ability means that an agent obtains a completely uninformative signal as effort approaches zero and a perfectly informative signal as effort approach infinity.

0, so that they perceive a surprisingly high reward at time 0. As a result, they raise their central estimates of their ability and choose greater effort at time 1. Because they are now overconfident, their time 1 rewards will, on average, be surprisingly small, leading them to reduce their time 2 ability estimates towards the true values. Following the average time 1 reward, these agents will still be overconfident at time 2 but less so than at time 1. Indeed, because their high time 1 effort made their beliefs especially sensitive to the observed time 1 reward, they will tend to be only slightly overconfident by time 2.

In contrast, some agents receive an unobserved shock that happens to take on a low value at time 0. These agents reduce their central estimates of their ability and choose lower effort at time 1. Because they are now underconfident, their time 1 rewards will, on average, be surprisingly large, leading them to raise their time 2 ability estimates towards the true value. Following the average time 1 reward, these agents will still be underconfident at time 2 but less so than at time 1. Because their low time 1 effort made their beliefs especially insensitive to the observed time 1 reward, their underconfidence may still be nearly as severe at time 2 as it was at time 1.

On average, agents still have well-calibrated beliefs at time 1 because they adjust their beliefs symmetrically in response to high or low time 0 shocks. However, agents tend to be underconfident at time 2: on average, their central estimates are below their true abilities because their posterior beliefs are more sensitive to the observed reward when their effort is high. Agents learn away time 0 shocks especially quickly when these shocks lead them to raise their central estimates of their own ability, and they learn away time 0 shocks especially slowly when these shocks lead them to lower their central estimates of their own ability. I show that this average underconfidence persists arbitrarily far into the future, vanishing only in the limit as infinite data accumulate.

In many contexts, the role of luck will diminish as agents apply more effort. For instance, perhaps running harder smooths out variations in tempo that arise due to distractions or topography. When effort reduces the variance of rewards, agents extract even more information from observed outcomes under high effort. The asymmetry in learning speeds described above becomes even more pronounced. By the foregoing logic, average underconfidence will again endogenously emerge and persist.

Rational updating can also endogenously generate overconfidence. Now let greater effort increase exposure to luck by enough to increase the variance of the signal that agents extract from observed outcomes. Like Emerson's "shallow man", agents understand that their efforts are largely modulated by circumstance. For example, running harder here reduces the consistency of their times by increasing the consequences of each day's minor variations in weather, fitness, or diet. Agents' beliefs are especially sensitive to news following low effort choices. Because they choose low effort when they lack confidence in their own ability, they learn away overly low ability estimates especially rapidly. And they learn away overly high ability estimates only slowly because their high efforts lead to especially noisy outcomes. When these agents hold overly high ability estimates, they will tend to receive bad news but attribute any news more to chance than to their own ability. These agents become overconfident on average, an effect that persists into future periods and vanishes only in the limit as infinite data accumulate.

The predicted dynamics are consistent with evidence in Hoffman and Burks (2020). They document that long-haul truckers demonstrate both overconfidence and underconfidence when predicting the miles they will drive in the coming week. On average, the truckers are overconfident, and that average overconfidence declines only slowly as truckers gain more experience on the job. The present model may also explain an apparent irrationality in their structural model of truckers' beliefs: Hoffman and Burks (2020) estimate that truckers perceive the variance of their productivity shocks to be greater than the true variance, leading them to update beliefs about their own productivity only slowly. For truckers to have ended up overconfident on average, the present analysis requires that the variance of their productivity shocks be high when their effort and confidence are high. Effort is an omitted variable in Hoffman and Burks (2020), as they recognize. Because the perceived variance in Hoffman and Burks (2020) is identified by the observed speed of learning, most of the identifying variation is likely to come from truckers who are especially overconfident and thus have more to learn. If unobserved effort choices endogenously increase variance for those truckers, then a one-size-fits-all estimate of perceived variance will primarily reflect their high variance and may be greater than the observed variance of productivity shocks across all truckers. Truckers' belief updating may yet be rational once we account for effort.[4]

The next section presents an analytically transparent two-period model with normal distributions. Section 3 generalizes the distributional assumptions and extends the horizon to infinitely many periods. Section 4 considers normally distributed signals with infinitely many periods. Section 5 describes how management can design feedback to induce overcon-

[4]The present analysis is consistent with Hoffman and Burks (2020) even if truckers truly do misperceive shocks' variance: an earlier working paper version showed that the critical determinant of under- or overconfidence is the perception of the relationship between variance and effort, not the true relationship between variance and effort. When truckers are overconfident on average, an empirical analysis should detect truckers as perceiving high variance on average, whether or not the variance is in fact high for any of them.

fidence in employees. Section 6 relates the theoretical setting to prominent empirical work on overconfidence. Section 7 reviews related literature. Section 8 proposes opportunities for further work. The final section contains proofs. The appendix provides additional results and numerical examples.

2 Closed-Form Derivation with Normal Distributions

Begin by considering an example that allows for closed-form solutions and accessible exposition.[5] There is a continuum of agents, indexed by i and of measure 1. In every period $t \geq 0$, agent i chooses how much effort e_{it} to apply to an activity. The agent's cost of applying effort is $c_i(e_{it})$, with $c_i(\cdot) \in C^2$, $c_i'(\cdot) > 0$, and $c_i''(\cdot) > 0$. The activity provides reward π_{it}, which depends on the chosen level of effort, on the agent's fixed ability z_i, and on a random shock ϵ_{it}:

$$\pi_{it} = e_{it} z_i + \sqrt{f(e_{it})}\, \epsilon_{it}, \tag{1}$$

with $f(\cdot) \in C^1$ and strictly positive. Effort and ability are complementary.[6] The shock is independent over agents and time, observed only via its effect on payoffs, and, in this section, normally distributed with mean zero and variance normalized to 1.[7] Throughout, bars denote population-average variables.

Agents' true abilities $z_i > 0$ are unknown to them. In this section, agent i's beliefs about her ability z_i are summarized by a normal distribution with mean μ_{it} and variance $\Sigma_{it} < \infty$. I describe agent i as overconfident at time t if $\mu_{it} > z_i$ and as underconfident if $\mu_{it} < z_i$. In order to avoid stacking the desk for misplaced confidence, assume for now that each agent's beliefs are initially well-calibrated: $\mu_{i0} = z_i$.[8] Finally, to allow for simpler exposition, assume in this section that z_i and Σ_{i0} are constant over agents.

[5]See Appendix B for a numerical implementation.

[6]Bénabou and Tirole (2002) extensively motivate complementarity between effort and ability. In laboratory experiments, Chen and Schildberg-Hörisch (2019) show that higher estimates of one's own ability induce additional effort, as implied by the present setting. Effort and ability are also typically presented as complementary in psychology literature that describes how agents choose effort and infer ability (e.g., Nicholls and Miller, 1984).

[7]Altering the shock's variance is equivalent to rescaling $f(e_{it})$.

[8]I will relax this assumption in Section 3. One might argue that fully rational agents would understand that initial beliefs are correct and thus immediately infer their true ability from these initial beliefs. However, we will soon see that agents' mean beliefs can depart from true ability, so this argument is not robust to agents having already experienced histories of nonzero length. More fundamentally, rationality requires Bayesian updating but does not pin down agents' priors, which are always model primitives. We can always, if we wish, constrain these primitives by a further assumption of rational expectations. As discussed below, the assumption that $\mu_{i0} = z_i$ can be interpreted as a form of rational expectations assumption.

The function $f(\cdot)$ determines the role of luck (i.e., the variance of rewards), and its derivative determines how effort choices affect the the role of luck. When $f'(\cdot) = 0$, the role of luck is independent of effort. When $f'(\cdot) > 0$, trying harder amplifies the role of luck. But when $f'(\cdot) < 0$, trying harder gives an agent more control over outcomes.[9]

Each agent chooses e_{it} to maximize expected period payoffs:[10]

$$\max_{e_{it}} E_{it}\left[\pi_{it} - c_i(e_{it})\right],$$

where $E_{it}[\cdot]$ denotes expectations conditioned on μ_{it} and Σ_{it}, without knowledge of z_i or the realizations of ϵ_{it}. Each agent's optimal choice of effort e^*_{it} satisfies the first-order necessary condition:

$$c'_i(e^*_{it}) = E_{it}\left[z_i + \frac{1}{2}\frac{f'(e^*_{it})}{\sqrt{f(e^*_{it})}}\epsilon_{it}\right],$$

which implies that $c'_i(e^*_{it}) = \mu_{it}$. Optimal effort e^*_{it} is an increasing function of μ_{it}. In order to focus on the issue at hand, we want to rule out zero-effort traps, where agents stop receiving new signals of ability.[11] Therefore assume in this section that $f(\cdot)$ is sufficiently small relative to z_i and μ_{i0} that the probability that $\mu_{it} \leq 0$ is vanishingly small. (Sections 3 and 4 use more elegant approaches to rendering zero-effort traps irrelevant.)

Each agent updates their beliefs about their ability z_i upon observing realized payoffs π_{it}, with $s_{it} \triangleq \pi_{it}/e_{it}$ constituting the signal of ability. Agents are Bayesian learners. The combination of normally distributed beliefs and normally distributed shocks generates a conjugate Bayesian updating rule:

$$\mu_{i(t+1)} = \left(\Sigma_{it}^{-1}\mu_{it} + \frac{e_{it}^2}{f(e_{it})}s_{it}\right)\left(\Sigma_{it}^{-1} + \frac{e_{it}^2}{f(e_{it})}\right)^{-1}, \tag{2}$$

$$\Sigma_{i(t+1)} = \left(\Sigma_{it}^{-1} + \frac{e_{it}^2}{f(e_{it})}\right)^{-1}. \tag{3}$$

[9] The case with $f'(\cdot) < 0$ can be interpreted as an "internal locus of control", with $f'(\cdot) \geq 0$ being an "external locus of control" (see Lybbert and Wydick, 2018). See Hestermann and Yaouanq (2021) for an alternate formulation of locus of control.

[10] Appendix A shows that the results survive when agents are not myopic.

[11] Although interesting for its connection to depression (de Quidt and Haushofer, 2017) and to poverty traps (Lybbert and Wydick, 2018), allowing for this possibility would serve to obscure the relationships of interest here.

Define

$$w(e_{it}, \Sigma_{it}) \triangleq \frac{\frac{e_{it}^2}{f(e_{it})}}{\Sigma_{it}^{-1} + \frac{e_{it}^2}{f(e_{it})}} \in [0, 1] \tag{4}$$

as the weight that time t agents place on the time t signal. Writing w_{it} for short, equation (2) becomes:

$$\mu_{i(t+1)} = (1 - w_{it})\mu_{it} + w_{it} s_{it}. \tag{5}$$

Consider the population-average central estimate, $\bar{\mu}_1$. From equation (5),

$$\bar{\mu}_1 = \int_0^1 [(1 - w_{i0})\mu_{i0} + w_{i0} z_i] \,\mathrm{d}i = \bar{z}.$$

On average, agents' beliefs remain well-calibrated at time 1. The population-average central estimate at time 2 is:

$$\bar{\mu}_2 = \int_0^1 \left[[1 - w(e_{i1}, \Sigma_{i1})]\mu_{i1} + w(e_{i1}, \Sigma_{i1}) z_i + w(e_{i1}, \Sigma_{i1}) \frac{\sqrt{f(e_{i1})}}{e_{i1}} \epsilon_{i1} \right] \mathrm{d}i.$$

e_{i1} differs across agents because it depends on μ_{i1} and thus on ϵ_{i0}. Using $\bar{\mu}_1 = \bar{z}$, we have:

$$\bar{\mu}_2 = \bar{z} - \int_0^1 w(e_{i1}, \Sigma_{i1}) \; [\mu_{i1} - \bar{\mu}_1] \,\mathrm{d}i.$$

As a good Bayesian, each individual agent does not expect her beliefs to change over time (see Appendix C): an agent does not know her true z_i, so she takes expectations over both z_i and the ϵ. Nonetheless, average beliefs do change in a predictable fashion once we condition on the true z_i, as might a researcher with an objective measure of agents' abilities. In particular, the average central estimate tends to drift away from $\bar{z}$ unless w_{i1} and μ_{i1} are uncorrelated over agents. Recognizing that the μ_{i1} are normally distributed and using that z_i and Σ_{i1} are constant over agents,[12] apply Stein's Lemma:

$$\int_0^1 w(e_{i1}, \Sigma_{i1}) \; [\mu_{i1} - \bar{\mu}_1] \,\mathrm{d}i = w_{i0}^2 \frac{f(e_{i0})}{e_{i0}^2} \int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}} \frac{\mathrm{d}e_{i1}}{\mathrm{d}\mu_{i1}} \,\mathrm{d}i. \tag{6}$$

Recall that $\mathrm{d}e_{i1}/\mathrm{d}\mu_{i1} > 0$. Therefore, the average belief drifts away from $\bar{z}$ unless $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} = 0$. Agents become underconfident on average if $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} > 0$, and agents become overconfident on average if $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} < 0$. From equation (4), only in a knife-edge

[12]From equation (3), Σ_{i1} depends only on Σ_{i0} and e_{i0}.

case does $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} = 0$.

Observe from (4) that e_{i1} affects w_{i1} through

$$Var[s_{it}|e_{it}, z_i] = \frac{f(e_{it})}{e_{it}^2}, \tag{7}$$

which is the variance of the likelihood. It is easy to see that $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} > 0$ if and only if that variance decreases in e_{i1}.[13] Therefore agents become underconfident on average if effort decreases the variance of the likelihood and become overconfident on average if effort increases the variance of the likelihood. In the special case that $f(\cdot)$ is independent of e_{it}, agents become underconfident on average because larger e_{i1} then unambiguously increases the precision of agent i's time t signal $s_{i1} = \pi_{i1}/e_{i1}$ and thus decreases the variance of the likelihood.

Consider a population of agents for whom additional effort decreases the variance of the likelihood and thus makes $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} > 0$. Some agents happen to receive a positive shock in period 0. They become overconfident and choose greater effort in period 1. At that time, they tend to receive shocks that correct their overconfidence (pushing their beliefs back towards z_i). Because $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} > 0$, they learn especially rapidly from these period 1 shocks and so on average enter period 2 only mildly overconfident, with an estimate close to z_i. In contrast, some agents happen to receive negative shocks in period 0. These agents become underconfident and choose a low effort level in period 1. The period 1 rewards tend to correct their underconfidence (pushing their estimates back towards z_i), but because $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} > 0$, they do not learn much from those rewards. These agents therefore tend to remain nearly as underconfident entering period 2 as they were entering period 1. On average, period 2 agents are underconfident, as the link between effort and information processing leads them to learn away overconfident period 1 beliefs faster than they learn away underconfident period 1 beliefs.

If, instead, $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} < 0$, then period 2 agents are overconfident on average, as the link between effort and information processing leads them to learn away underconfident period 1 beliefs faster than they learn away overconfident period 1 beliefs. An experimentalist who observed agents in period 2 and compared their self-estimates to measured true ability would conclude that they tend to be overconfident, but this overconfidence is in fact justified by the combination of the data that the agents generated and how their choices generated the data.

[13] Section 4 relates the sign of $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1}$ to the elasticity of $f(e_{it})$.

Section 6 relates the setting to prominent empirical work on overconfidence. As a brief example, consider a student taking a test. The student chooses how much to focus on each question. Greater focus matters more for students with high ability than for students with low ability. Upon seeing the results of the test, students update their beliefs about their own ability, adjusting for how hard they tried on the test. This story is consistent with evidence from a recent field experiment: Gneezy et al. (2019) show that incentivizing students to exert more effort on a standardized test does improve test scores (effort matters for outcomes and responds to incentives) and improves test scores most strongly for higher-ability students (effort is complementary to ability). The authors highlight that cross-sectional comparisons of test scores across countries can mislead policymakers when students' unobserved effort differs across cultures. Here, we recognize that the students themselves are likely to account for their own effort choices when interpreting their own test scores. We see that their average beliefs will not accurately reflect their own abilities. And the type of inaccuracy depends on empirically testable phenomena. They underestimate ability on average if trying harder makes tests better signals of ability (as when effort reduces random mistakes arising from inattention), and they overestimate ability on average if trying harder increases the chance of random mistakes (as when effort means forsaking sleep to study longer).

Rationality dictates how agents update beliefs in response to new information but not what prior beliefs agents should hold.[14] Rational expectations is an additional assumption sometimes imposed to pin down prior beliefs. In order to raise the hurdle for obtaining misplaced confidence, I have modeled each agent as initially holding well-calibrated beliefs (which I will relax in Section 3). This can be considered a form of rational expectations. When z_i is heterogeneous (as it will be in Section 3), agents' observable initial effort choices track their abilities. It is as if each agent had observed signals of their own ability prior to time 0. An alternate form of rational expectations would fix each agent's μ_{i0} and Σ_{i0} to match the population distribution of the z_i, as if agents had observed signals of other agents' abilities prior to time 0.[15] In that case, $\bar{\mu}_0 = \bar{z}$ but agents' observable effort choices are initially independent of their true abilities. The evolution of the average belief ($\bar{\mu}_t$) is now the same as each agent's expectation of her mean posterior belief ($E_{i0}[\mu_{it}]$). Unsurprisingly, Appendix C shows that Bayesian agents' expectations of their own posterior beliefs are martingales, so $\bar{\mu}_t = \bar{z}$ at all $t \geq 0$ and we no longer find average misplaced confidence emerging. However, in many contexts it is more plausible that agents have some imperfect idea of their own

[14] See footnote 1 in van den Steen (2004) and citations therein.

[15] Of course, agents' beliefs no longer match the population distribution after time 0 because the distribution of the z_i is fixed yet, from (3), $\Sigma_{it} < \Sigma_{i0}$ for $t > 0$.

ability (with $\mu_{i0} \approx z_i$ and $\Sigma_{i0} > 0$) than that agents perfectly understand the population-wide distribution of ability but know nothing about their own standing within it. We should expect average over- or underconfidence to emerge in environments consistent with the former assumption but not in environments consistent with the latter assumption. Further, as we shall verify formally in subsequent sections, the appearance of average misplaced confidence is not a knife-edge result, so we should expect average misplaced confidence in environments even roughly consistent with the former assumption. In contrast, there is no reason to predict a lack of average misplaced confidence in environments that are only roughly consistent with the latter assumption.[16]

3 Infinitely Many Periods and Non-Normal Shocks

The model in Section 2 allows for a simple transparent derivation driven merely by first-order conditions, Bayes' Rule, and Stein's Lemma, but one may wonder whether misplaced confidence persists beyond two periods, whether it arises under non-normal distributions, and whether it arises for priors that are not initially well-calibrated. Here I expand the analysis to infinitely many periods, broaden the distributional assumptions, and relax the restriction on the prior at the cost of rendering closed-form solutions such as (6) impossible.

Let rewards be $\pi_{it} = e_{it} z_i + \nu_{it}$, with the ν_{it} indicating random shocks that are not directly observed by agents. In contrast to Section 2, z_i is here potentially heterogeneous and the ν_{it} are potentially non-normal. Conditional on e_{it}, the ν_{it} are identically and independently distributed over time and agents, with mean zero.[17] The rest of the setting is familiar from Section 2. In particular, each z_i is strictly positive and fixed over time, I define $\mu_{it} \triangleq E_{it}[z_i]$, optimal effort satisfies $c_i'(e_{it}^*) = \mu_{it}$ for $\mu_{it} > 0$, and agent i's time t signal of her ability is $s_{it} \triangleq \pi_{it}/e_{it}$.[18]

[16] The lack of average misplaced confidence when agents' prior beliefs match the population distribution of the z_i is in fact likely to be a knife-edge case for the following reason. As time passes, each μ_{it} converges to z_i, in which case we approach a setting with well-calibrated beliefs. The reason we would not predict average misplaced confidence in some sufficiently late period is that different agents approach their true abilities from different directions, and this heterogeneity is such that it exactly offsets the tendency towards average over- or underconfidence that manifests as agents' beliefs become well-calibrated. Failing to find average misplaced confidence is therefore unlikely to be robust to small perturbations in agents' priors around the population distribution of the z_i.

[17] We could permit the ν_{it} to have nonzero means by subtracting $E[\nu_{it}]$ from π_{it} in the definition of s_{it} below.

[18] Following others (e.g., Heidhues et al., 2018), I model agents choosing effort myopically because the exposition is clearer. The restriction is inessential: forward-looking agents' effort choices depend on μ_{it} and also on higher moments of the prior, but these effort choices converge to the myopic effort choices as time

Letting p_x be the density function of random variable x, I assumed in Section 2 that the likelihood $p_{s_{it}}(s_{it}|e_{it}, z_i)$ was a normal density. Here I instead assume that each z_i can be mapped to a θ_{it} such that $p_{s_{it}}(s_{it}|e_{it}, \theta_{it})$ is a member of a regular natural exponential family (NEF) of distributions (Morris, 2006).[19] (Note that this mapping can depend on e_{it}, which is the source of the t subscript on θ_{it}.) Formally,

$$p_{s_{it}}(s_{it}|e_{it}, \theta_{it}) = \exp[\theta_{it} s_{it} - M_i(\theta_{it}; e_{it})], \tag{8}$$

where M_i is continuous in e_{it} and the θ_{it} for which $\exp[M_i(\theta_{it}; e_{it})] < \infty$ constitute a nonempty open set in $\mathbb{R}$ for all $e_{it} > 0$. The families of distributions that satisfy (8) (and hence are NEFs) include the normal, Poisson, binomial, negative binomial, and gamma distributions, with the latter nesting the exponential and chi-square distributions as special cases. M_i is known as the "cumulant function" (distinct from the cumulant-generating function) because its kth derivative is the kth cumulant. Its form identifies a specific exponential family of distributions (such as the family of normal distributions). And within that specific family, the "natural parameter" θ_{it} indexes a specific distribution (as the mean does for a normal distribution with known variance).

Because $E_{it}[s_{it}|\theta_{it}] = M_i'(\theta_{it}; e_{it})$ and $z_i = E_{it}[s_{it}|\theta_{it}]$, we have $z_i = M_i'(\theta_{it}; e_{it})$.[20] Therefore $\mu_{it} = E_{it}[M_i'(\theta_{it}; e_{it})]$. As is well known (e.g., Barndorff-Nielsen, 1978; Consonni and Veronese, 1992), there is a bijection between θ_{it} and z_i.[21] Plugging $\theta(z_i; e_{it})$ into (8) would yield the likelihood in terms of z_i (i.e., $p_{s_{it}}(s_{it}|e_{it}, z_i)$), which is known as the mean parameterization.

Following Morris (1982) and subsequent literature (e.g., Consonni and Veronese, 1992; Morris and Lock, 2009), the "variance function" for the time t likelihood is

$$V_i(z_i; e_{it}) \triangleq \frac{\partial^2 M_i(\theta(z_i; e_{it}); e_{it})}{\partial \theta^2}.$$

The variance function gives the variance of the signal as a function of its conditional expectation, z_i. I restrict attention to distributions with variance functions in an especially

passes and also under the conditions of Proposition 2 below. Also see Appendices A and B.1.

[19] Any exponential family can be reparameterized as a natural exponential family (e.g., Gutiérrez-Peña and Smith, 1997). The ν_{it} might not come from the same family as s_{it}.

[20] The effect of effort on the mapping from z_i to θ_i follows from the implicit function theorem.

[21] $M_i'' > 0$ because it is the variance (i.e., the second cumulant), which implies that z_i increases monotonically in θ_{it}.

prominent and well-studied class:

$$V_i(z_i; e_{it}) = \zeta_{i2}(e_{it})\, z_i^2 + \zeta_{i1}(e_{it})\, z_i + \zeta_{i0}(e_{it}), \tag{9}$$

with V_i finite and with the coefficients $\zeta_{i2}(e_{it})$, $\zeta_{i1}(e_{it})$, and $\zeta_{i0}(e_{it})$ each a differentiable function of e_{it}. By satisfying (8) and (9), the signal is conditionally distributed according to a regular natural exponential family with quadratic variance function (NEF-QVF). Morris (1982) shows that there are six types of NEF-QVFs. These include the five most important NEFs: the normal, Poisson, gamma, binomial, and negative binomial families of distributions. As an example, equation (9) was satisfied in Section 2, where the variance of s_{it} conditional on e_{it} and z_i had, from equation (7), $\zeta_{i0}(e_{it}) = \frac{f(e_{it})}{e_{it}^2}$, $\zeta_{i1}(\cdot) = 0$, and $\zeta_{i2}(\cdot) = 0$ (the latter two relationships are always true for normal distributions).

The agent's time 0 prior over θ_{i0} is the standard conjugate prior, which is also a member of an exponential family (Diaconis and Ylvisaker, 1979):

$$p_{\theta_{i0}}(\theta_{i0}|\mathcal{I}_{i0}) = K_{i0} \exp\left[n_{i0} x_{i0} \theta_{i0} - n_{i0} M_i(\theta_{i0}; e_{i0})\right], \tag{10}$$

with $n_{i0} > \zeta_{i2}(e_{i0})$, x_{i0} in the convex hull of the support of $p_{s_{i0}}(\cdot|e_{i0}, \theta_{i0})$, $K_{i0} > 0$ a normalizing constant, and $\mathcal{I}_{i0}$ indicating agent i's time 0 information set. Theorem 1 of Diaconis and Ylvisaker (1979) implies $K_{i0} < \infty$. Below Assumption 2, I will relate the parameter x_{i0} to the mean of the prior and will interpret the parameter n_{i0}. Plugging $\theta(z_i; e_{it})$ into (10) would generate what Consonni and Veronese (1992) label the D-Y conjugate family of priors over the mean parameter (i.e., the D-Y conjugate prior over z_i).[22]

Because the variance function satisfies (9), the prior over z_i in (10) is a member of one of the Pearson families of distributions (Morris, 1983). Further, the variance function for the time 0 prior is (Morris, 1983; Morris and Lock, 2009):

$$\tilde{V}_{i0}(\mu_{i0}) = \frac{V_i(\mu_{i0}; e_{i0})}{n_{i0} - \zeta_{i2}(e_{i0})}. \tag{11}$$

Requiring $n_{i0} > \zeta_{i2}(e_{i0})$ ensures that the variance of the prior exists and is finite.[23] If one were to exogenously increase the variance of the signal, then the endogenous parameter n_{i0} must increase in order to hold the variance of the prior fixed. We will soon see that an

[22] Consonni and Veronese (1992) show that the D-Y conjugate prior over z_i is conjugate to (8) expressed in terms of z_i when the variance function is quadratic as in (9) but not necessarily otherwise.

[23] In Section 2, $\Sigma_{i0} < \infty$ implied $n_{i0} > 0$. Because $\zeta_{i2}(\cdot) = 0$ for normal distributions, $n_{i0} > \zeta_{i2}(e_{i0})$.

increase in n_{i0} means that the prior is weighted more heavily in the posterior mean, so increasing the variance of the signal has the intuitive effect of reducing the posterior mean's sensitivity to the signal.

I have thus far presented a fairly general framework with conventional groupings of distributional families. I now impose two more particular assumptions. The first assumption restricts the distribution of signals (which is determined by the distribution of the ν_{it}). If the prior and the set of possible signals permit agents to believe they have no ability, then all agents eventually fall into a low-confidence trap in which they choose zero effort and never revise their beliefs further. These traps would obscure the mechanism of interest here. The next assumption will ensure that agents avoid low-confidence, zero-effort traps:[24]

Assumption 1. $p_{s_{it}}$ *has support only in the weakly positive numbers.*

Most families of distributions that are NEF-QVF satisfy this restriction. The prominent exception is the family of normal distributions, which I analyze separately in Section 4.

The second assumption restricts agents' mean beliefs to initially have error no larger than δ:

Assumption 2. *For all i and for given $\delta \geq 0$, $\mu_{i0} > 0$ and $\mu_{i0} \in [z_i - \delta, z_i + \delta]$.*

Importantly, this assumption permits the possibility that priors are well-calibrated, as in Section 2.

From Theorem 2 of Diaconis and Ylvisaker (1979), $x_{i0} = E_{i0}[M_i'(\theta_{i0}; e_{i0})]$. Therefore $x_{i0} = \mu_{i0}$. And from equation (2.10) in Diaconis and Ylvisaker (1979), Bayesian updating implies

$$\begin{aligned} \mu_{i1} =& \frac{n_{i0}}{n_{i0}+1}\mu_{i0} + \frac{1}{n_{i0}+1}s_{i0} \\ =& (1 - w_{i0})\mu_{i0} + w_{i0}s_{i0}. \end{aligned} \tag{12}$$

The weight $w_{i0} \in (0, 1)$ that agent i places on the time 0 signal decreases in n_{i0}, and the weight $1 - w_{i0}$ that agent i places on time 0 prior beliefs increases in n_{i0}. This is why n_{i0} is commonly thought of as the sample size of the prior (Diaconis and Ylvisaker, 1979). Using conjugacy of the prior and repeating the steps, we find that the prior at any time t has the

[24] More precisely, Assumption 1 ensures that agents avoid these traps once combined with Assumption 2 below, which ensures that agents do not begin in a trap.

form of (10) with $K_{it} < \infty$, that its variance function has the form of (11), and that

$$\mu_{i(t+1)} = (1 - w_{it})\mu_{it} + w_{it}s_{it} \tag{13}$$

at any time $t > 0$, with $w_{it} \in (0, 1)$. Assumptions 1 and 2 imply $\mu_{i(t+1)} > 0$. The linearity of the posterior mean in the prior mean and the signal seen in both (12) and (13) is the critical feature of exponential families for the present analysis. In Section 2, this same linearity appeared in equation (5) because normal distributions are members of an exponential family.

Working backwards through time,

$$\begin{aligned}\mu_{it} - z_i =& \frac{w_{i(t-1)}}{e_{i(t-1)}}\nu_{i(t-1)} + (1 - w_{i(t-1)})(\mu_{i(t-1)} - z_i) \\ =& \sum_{j=1}^{t}\left(\prod_{k=1}^{j-1}(1 - w_{i(t-k)})\right)\frac{w_{i(t-j)}}{e_{i(t-j)}}\nu_{i(t-j)} + \left(\prod_{k=1}^{t-1}(1 - w_{i(t-k)})\right)(1 - w_{i0})[\mu_{i0} - z_i]\end{aligned} \tag{14}$$

when effort choices are nonzero. From the law of iterated expectations,

$$\begin{aligned}\int_0^1 [\mu_{it} - z_i]\,\mathrm{d}i =& \sum_{j=2}^{t}\int_0^1 \frac{w_{i(t-j)}}{e_{i(t-j)}}\, Cov_{i(t-j)|z_i}\left[\prod_{k=1}^{j-1}(1 - w_{i(t-k)}), \nu_{i(t-j)}\right]\mathrm{d}i \\ &+ \int_0^1 \left(\prod_{k=1}^{t-1}(1 - w_{i(t-k)})\right)(1 - w_{i0})[\mu_{i0} - z_i]\,\mathrm{d}i,\end{aligned} \tag{15}$$

where $Cov_{i(t-j)|z_i}$ is the covariance conditional on z_i and on $\nu_{i(t-j-h)}$, for $h \in \{1, ..., t-j\}$. Consider a case in which agents' beliefs are initially well-calibrated, meaning that $\mu_{i0} = z_i$. If the covariance is negative (positive) for all i and j, then $\bar{\mu}_t < (>)\, \bar{z}$, indicating average underconfidence (overconfidence). A negative covariance means that observing large values of π_{it}/e_{it} induces agent i to place additional weight on later periods' news shocks and less weight on later periods' priors. Thus, observing a signal of high ability undercuts itself: by altering later effort choices, it induces agent i to downweight this same signal when forming later posteriors. Agents' posteriors end up driven by the more pessimistic signals of ability. Conversely, a positive covariance means that observing a large value of π_{it}/e_{it} makes agent i's posterior less sensitive to later signals. Agents' posteriors end up driven by the more optimistic signals of ability.

It is again true that agents' beliefs are not biased on average in period 1 if agents' beliefs are initially well-calibrated:

$$\bar{\mu}_1 = \int_0^1 [(1 - w_{i0})\mu_{i0} + w_{i0} z_i] \, \mathrm{d}i$$
$$\in \left[\bar{z} - \delta \int_0^1 (1 - w_{i0}) \, \mathrm{d}i, \; \bar{z} + \delta \int_0^1 (1 - w_{i0}) \, \mathrm{d}i \right],$$

so that $\bar{\mu}_1 = \bar{z}$ if $\mu_{i0} = z_i$ for all i (i.e., if Assumption 2 holds with $\delta = 0$). Now consider subsequent periods. The following lemma relates effort to the weight placed on the signal in agents' updating equations:

Lemma 1. *Let Assumptions 1 and 2 hold. Then:*

i

$$\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0 \quad \textit{if} \quad \frac{\partial V_i(\mu_{i0}; e_{i0})}{\partial e_{i0}} < 0 \;\textit{and}\; \frac{\mathrm{d}\zeta_{i2}}{\mathrm{d}e_{i0}} \leq 0,$$

ii

$$\mathrm{d}w_{i0}/\mathrm{d}e_{i0} < 0 \quad \textit{if} \quad \frac{\partial V_i(\mu_{i0}; e_{i0})}{\partial e_{i0}} > 0 \;\textit{and}\; \frac{\mathrm{d}\zeta_{i2}}{\mathrm{d}e_{i0}} \geq 0$$

Proof. Agents set their priors independently of e_{i0}, so the variance of the prior must be independent of e_{i0}. Differentiating $\tilde{V}_{i0}$ with respect to e_{i0}, setting the derivative to zero, and rearranging, we have, for all $e_{i0} > 0$,

$$\frac{\mathrm{d}n_{i0}}{\mathrm{d}e_{i0}} = \frac{\partial V_i(\mu_{i0}; e_{i0})}{\partial e_{i0}} \frac{n_{i0} - \zeta_{i2}(e_{i0})}{V_i(\mu_{i0}; e_{i0})} + \frac{\mathrm{d}\zeta_{i2}}{\mathrm{d}e_{i0}}.$$

Under the conditions of part (i), this implies that $\mathrm{d}n_{i0}/\mathrm{d}e_{i0} < 0$ and, because $w_{i0} \triangleq 1/(n_{i0}+1)$ from (12), that $\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0$. Under the conditions of part (ii), this implies that $\mathrm{d}n_{i0}/\mathrm{d}e_{i0} > 0$ and that $\mathrm{d}w_{i0}/\mathrm{d}e_{i0} < 0$.

□

As is intuitive, agents' sensitivity to new information increases in effort (and thus in confidence) if additional effort reduces the variance of the signal and decreases in effort if additional effort increases the variance of the signal. We should expect the covariance in (15) to be negative (and underconfidence to emerge) if effort reduces the variance of the signal, and we should expect the covariance in (15) to be positive (and overconfidence to emerge) if effort increases the variance of the signal.

The following proposition shows that average bias indeed can again emerge after period 1, with the direction of average misplaced confidence depending on how effort affects the variance of the signal:

Proposition 1. *Let Assumptions 1 and 2 hold. For all* $t > 1$*, there exists* $\hat{\delta} > 0$ *such that if* $\delta < \hat{\delta}$ *and the moments of the distribution over each* ν_{it} *are sufficiently small, then*

i

$$\bar{\mu}_t < \bar{z} \quad \textit{if} \quad \frac{\partial V_i(\mu_{i0}; e_{i0})}{\partial e_{i0}} < 0 \textit{ and } \frac{\mathrm{d}\zeta_{i2}}{\mathrm{d}e_{i0}} \leq 0,$$

ii

$$\bar{\mu}_t > \bar{z} \quad \textit{if} \quad \frac{\partial V_i(\mu_{i0}; e_{i0})}{\partial e_{i0}} > 0 \textit{ and } \frac{\mathrm{d}\zeta_{i2}}{\mathrm{d}e_{i0}} \geq 0$$

Proof. The proof is by induction. See Section 9.

□

Average misplaced confidence emerges after period 1 and persists in later periods, vanishing only in the limit as agents accumulate infinite data. Importantly, average misplaced confidence emerges even when the deck is most stacked against it, with agents' priors all initially well-calibrated (so $\delta = 0$). As in Section 2, the direction of average misplaced confidence depends on the effects of effort on the variance of the likelihood. From Lemma 1, additional effort increases (decreases) the weight that agents place on the signal when it decreases (increases) the variance of the likelihood. Because effort increases in mean beliefs and new signals tend to correct mistaken beliefs, agents learn away mistaken beliefs faster when they are overconfident (underconfident), as already described following equation (6). We therefore predict average underconfidence (overconfidence) when additional effort decreases (increases) the variance of the likelihood.[25]

The variance of the time 0 likelihood is

$$Var_{i0}[s_{i0}|e_{i0}, \theta_{it}] = \frac{1}{e_{i0}^2} Var_{i0}[\nu_{i0}].$$

The variance of the likelihood decreases in effort if either the variance of ν_{i0} is independent of effort or the variance of ν_{i0} decreases in effort. The variance of the likelihood increases in effort only if the variance of ν_{i0} increases sufficiently strongly in effort. The variance of the likelihood is independent of effort only in a knife-edge case in which effort increases the

[25] The requirement that the moments of ν_{it} be small ensures that second-order terms do not dominate the effect on average confidence.

variance of ν_{i0} at just the right rate.[26] Proposition 1 therefore formalizes the intuition from the introduction about how the effect of effort on the variance of the unobserved shocks generates average misplaced confidence by inducing asymmetric learning speeds.

4 Infinitely Many Periods and Normally Distributed Shocks

Assumption 1 ruled out normal distributions in order to avoid a zero-effort trap ever being optimal, but the most closely related empirical work assumes normally distributed shocks (Hoffman and Burks, 2020). I here examine normally distributed shocks and therefore directly generalize Section 2 to infinitely many periods.

Let the reward π_{it} be as in equation (1), with normally distributed ϵ_{it} having mean zero and variance $\sigma^2 > 0$. Now let agent i's time 0 prior for z_i be truncated-normal with support in $[a_i, b_i]$, for $a_i \in [0, z_i)$ and $b_i \in (z_i, \infty]$. The lower truncation point rules out the possibility of beliefs justifying the choice of zero effort. A Bayesian's posterior is also truncated-normal, with a_i and b_i still the truncation points. Use μ_{it} and Σ_{it} to denote the mean and variance of the corresponding untruncated normal distribution. The updating rules are as in equations (2) and (3) from Section 2.[27] The rest of the setting is familiar from previous sections.

Define the elasticity of $f(e_{it})$ as $\chi(e_{it}) \triangleq e_{it}\, f'(e_{it})/f(e_{it})$.

Lemma 2. *$Var[s_{it}|e_{it}, z_i]$ decreases in e_{it} if and only if $\chi(e_{it}) < 2$, and w_{it} increases in e_{it} if and only if $\chi(e_{it}) < 2$.*

Proof. The first part follows from (7). The second part follows from the first part and (4). □

Additional effort increases the marginal effect of ability on the reward (i.e., the variance of ϵ_{it}/e_{it} declines in e_{it}), but when $f'(\cdot) > 0$, additional effort also increases the variance of the reward. The second effect dominates if and only if $\chi(e_{it}) > 2$.

Define agent i's maximum likelihood estimate of z_i as $\phi_{it} \triangleq \max\{a_i, \min\{b_i, \mu_{it}\}\}$. This maximum likelihood estimate is the mean of the corresponding untruncated distribution and converges to the mean of the truncated distribution as the mass beyond the truncation points

[26] In addition, if we eliminated complementarity between effort and ability (making effort choices independent of ability) or eliminated agents' freedom to choose effort (making effort exogenous), then $\mathrm{d}e_{it}/\mathrm{d}\mu_{it} = 0$ and the proof shows that we would predict neither underconfidence nor overconfidence.

[27] Truncation changes the posterior within its support only through the normalization factor.

vanishes. The following corollary describes the evolution of overconfidence, here measured as $\phi_{it} - z_i$:[28]

Proposition 2. *Let Assumption 2 hold. For all $t > 1$, there exists $\hat{\delta} > 0$ such that if $\delta < \hat{\delta}$ and σ^2 and each Σ_{i0} are sufficiently small, then*

i

$$\bar{\phi}_t < \bar{z} \quad \textit{if} \quad \chi(e_{i0}) < 2,$$

ii

$$\bar{\phi}_t > \bar{z} \quad \textit{if} \quad \chi(e_{i0}) > 2.$$

Proof. Equation (3) and the nonoptimality of zero effort choices imply $\Sigma_{i0} > \Sigma_{it}$ for $t > 0$. As Σ_{it} becomes small, the time t posterior and prior both become approximately normal, with μ_{it} and ϕ_{it} converging to $E_{it}[z_i]$. The proposition then follows from Lemma 2 and the proof of Proposition 1 since the family of normal distributions is NEF-QVF. □

In Section 2, average misplaced confidence emerged in period 2. We now learn that this average misplaced confidence persists in all later periods, vanishing only as agents accumulate infinite data. Researchers studying this population would detect average under- or overconfidence almost regardless of which period they happen to sample from. Proposition 2 relates average overconfidence and underconfidence to model primitives in a transparent fashion: the sign of $\chi(e_{i0})-2$ determines whether average overconfidence or underconfidence emerges, and that sign depends on observable characteristics of the decision-making environment captured by $f(e_{i0})$.

Lemma 2 establishes that the variance of the likelihood decreases in effort if and only if $\chi(e_{i0}) < 2$, so Proposition 2 provides results analogous to those in Proposition 1. As before, average underconfidence emerges if additional effort reduces the variance of the likelihood and average overconfidence emerges if additional effort increases the variance of the likelihood. Therefore we have again formalized the intuition about effort-dependent variance, asymmetric learning speeds, and average misplaced confidence given in the introduction.

[28] ϕ_{it} approaches $E_{it}[z_i]$ under the conditions of the corollary, so the present measure of overconfidence is in practice similar to the measure used in previous sections.

5 Application: Conditionally Vague Feedback from Management

Organizations may want to induce overconfidence in their employees (Gibbs, 1991; Gervais and Goldstein, 2007; Hoffman and Burks, 2020). The present analysis suggests a novel way that a principal can manage an agent's confidence without deceit. The principal need only commit to giving the agent more information after bad performances than after good performances. This conditionally vague feedback rule helps the agent to learn away mistakenly low beliefs about her own ability faster than she learns away mistakenly high beliefs.

For example, an employer could require one-on-one performance reviews with underperforming employees but not with overperforming employees. In these reviews, employees would gain insight into other factors that could have affected performance, such as broader market conditions. Underperforming employees would quickly learn away shocks due to bad luck, but overperforming employees would only slowly learn away shocks due to good luck.

As a second example, managers who rate their employees' performance could finely divide ratings among underperforming employees—highlighting the degree to which their performance was merely bad luck—while compressing ratings for highly performing employees. Indeed, Cappelli and Conyon (2018), among others, find evidence of just such a skew in the distribution of employee ratings.[29]

These types of interventions require the employer to abide only by simple rules governing the detail of feedback given. Because that feedback is honest, employees have no incentive to alter their behavior based on knowledge of this feedback rule. Employees who understand the feedback rule may understand that it tends to make employees overconfident on average, but any individual employee can do no better than to update as a Bayesian based on the information received and the effort choices made.

6 Relation to Empirical Work on Overconfidence

The critical ingredients for rational overconfidence are that ability and effort be complementary and that additional effort complicates learning about ability from observed payoffs. I now consider how these conditions fit prominent empirical work on overconfidence.

First, some of the most prominent field evidence for the importance of overconfidence

[29] Cappelli and Conyon (2018) also report that ratings vary over time for a given employee, apparently responding to performance (as required to match the present setting).

comes from chief executive officers' (CEOs') investment decisions. Malmendier and Tate (2005, 2015) measure overconfidence from how CEOs exercise the stock options granted to them. They show that CEOs who are overconfident by this measure tend to invest more when cash flow is abundant, in accord with predictions. The primary theoretical explanation for compensating executives (or other employees) through stock options is that firm owners seek to resolve a principal-agent problem by aligning owners' and executives' incentives but may be constrained from providing such incentives through salary adjustments or bonuses (Hall and Liebman, 1998). In the benchmark principal-agent framework, the executive's action space is effort (Murphy, 1999). Stock options are more commonly granted to executives than to salaried workers and least commonly granted to hourly workers (Hall and Murphy, 2002). If compensation tracks ability, then we can conclude that firm owners perceive effort and ability to be complementary. Further, increasing investment may increase the marginal effect of ability on firm value even as it increases exposure to diverse stochastic factors. If the latter effect is sufficiently strong, executives may have a hard time learning away overconfidence. The critical ingredients of the present model are in place, making CEOs' overconfidence plausibly rational. Future work should examine how overconfidence evolves following the types of high cash flow events that encourage high investment.

Second, much work considers overconfidence in investors (Daniel and Hirshleifer, 2015). In particular, overconfidence is linked to the volume of trade (e.g., Barber and Odean, 2001). Barber and Odean (2002) show that investors begin trading online just after they experience large returns, which the authors interpret as increasing investors' estimates of their own ability. Online trading enables greater effort by reducing frictions and providing more data to analyze. Deciding to move online following high returns suggests that investors view effort and ability as complementary. Graham et al. (2009) show that investors with higher regard for their own ability tend to invest in more asset classes (in particular, foreign markets). Investors will perceive unconventional asset classes to be especially noisy if they do not understand them as well. The critical ingredients of the present model are again in place. Rational investors might learn away overconfidence relatively slowly, leading researchers to detect average overconfidence among the population of investors.

Finally, consider a setting whose fit to the present model is less clear. Niederle and Vesterlund (2007) explore overconfidence in a laboratory environment. They show that subjects' performance on an addition task tends to improve when participating in a tournament instead of being paid piece-rate. This result suggests that effort matters. To learn whether effort and ability are complementary, we would like to know whether applying more effort improves higher-ability subjects' performance by more. Unfortunately, this is difficult to assess from the reported results. Future laboratory experiments should explore whether their subjects and schemes are consistent with the present paper's mechanism.

7 Related Literature

The proposed mechanism for generating persistent over- or underconfidence appears to be novel. A first set of papers describes agents' motivations to become overconfident, whether because optimism increases utility (Brunnermeier and Parker, 2005) or because confidence helps to overcome the tendency to procrastinate (Bénabou and Tirole, 2002). In contrast, the present setting is neoclassical: agents' expected payoffs are maximized when they correctly estimate their own ability.

A second set of papers generates overconfidence by assuming that individuals use a biased updating process (see Hirshleifer, 2001, 2015). For instance, individuals overly attribute successes to their own ability and failures to chance (e.g., Daniel et al., 1998; Gervais and Odean, 2001), or individuals forget failures more often than successes (Compte and Postlewaite, 2004). The present setting generates persistent overconfidence as a result of rational Bayesian learning.[30]

A third set of papers studies selection mechanisms that can make the majority of a population of Bayesian updaters believe that each of their abilities are better than average (e.g., Zábojník, 2004; Köszegi, 2006; Krähmer, 2007; Jehiel, 2018). When agents choose to stop collecting information once they receive a sufficiently positive signal about themselves or about the payoffs to some activity, high confidence is an absorbing state that attracts an ever greater share of the population. These settings have the flavor of bandit models, as there are actions that do not generate information about the outcomes of other actions.[31] In contrast, agents here never stop updating beliefs about the payoffs from all possible actions. van den Steen (2004) also considers a selection mechanism, with "overoptimism" emerging as a type of winner's curse when agents choose among a set of actions. Here average overconfidence emerges only over time, as agents' updated beliefs lead them to choose actions that have the side effect of making the speed of learning asymmetric around their true ability.

[30] Moore and Healy (2008) show how Bayesian updating can generate overestimation if agents' beliefs are not initially well-calibrated. I show overestimation emerging even when prior beliefs do not already demonstrate overestimation. Benoît and Dubra (2011) show how Bayesian updating can generate the appearance of overplacement. Here the overconfidence is not apparent: it affects observable effort choices.

[31] A similar mechanism underpins the model of self-control in Ali (2011).

The proposed mechanism is more closely related to two recent papers.[32] First, Silva (2017) demonstrates how the asymmetric speed at which agents learn following good and bad shocks can generate systematic overconfidence in a two-period model with normal distributions. The agent receives outside help following an early signal that he is of high quality but not after an early signal that he is of low quality. Because he is aware that he begins receiving help but does not know how good that help is, an agent who saw a good outcome in the first period will weight that outcome especially strongly when forming a posterior in period 2. The present paper is similar in generating overconfidence when asymmetric rates of learning make the time t posterior more sensitive to high rewards. However, here the mechanism is that agents themselves affect their ability to learn from signals as a byproduct of their optimal effort choices. I avoid postulating an additional source of noise that arises only after certain types of rewards.[33]

Second, Hestermann and Yaouanq (2021) study an agent who is uncertain about his own fixed ability and also about some feature of the environment. The agent learns about both from a sequence of binary outcomes. If the agent is initially overconfident, then he rationally believes that good outcomes reflect his own ability whereas bad outcomes reflect a harsh environment. In this manner, overconfidence can persist for some time. We here see how overconfidence and underconfidence can emerge and persist even when agents' initial beliefs are well-calibrated and even when agents correctly understand their environments.[34]

Finally, some recent work considers how learning may confirm an agent's overconfidence. Heidhues et al. (2018) study when the actions chosen under an agent's permanently misspecified model of his own ability generate signals that do not lead the agent to question his incorrect beliefs about his own ability.[35] Fudenberg et al. (2017) consider the interaction between learning and a form of misspecification that places probability zero on the truth. In contrast, here agents have correctly specified models of the world, long-run limit beliefs

[32]The proposed mechanism also shares features with a macro model of uncertainty over the business cycle (Van Nieuwerburgh and Veldkamp, 2006). There, asymmetric learning speeds arise because aggregate production and the (uncertain) level of technology are complementary and chosen production levels increase in beliefs about technology.

[33]On a technical level, the current paper generalizes beyond two periods and beyond normally distributed shocks. Either generalization prevents closed-form solutions.

[34]Hestermann and Yaouanq (2021) also analyze asymptotic beliefs when agents can experiment. As in the third set of papers discussed above, overconfidence can persist due to decisions to stop collecting information. That overconfidence does not impose costs on agents. In contrast, agents here never stop experimenting and misplaced confidence is costly.

[35]Heidhues et al. (2018) allow for learning about one's own ability in an extension to their primary setting. There, learning generates overconfidence because they assume that the agent sees a biased signal of his ability. This extension is closely related to the papers discussed in the second paragraph of this section.

converge to the truth, and misplaced confidence emerges endogenously rather than being imposed ex ante.[36]

8 Further Work

We have seen that rational Bayesian agents become, on average, persistently overconfident when additional effort makes it harder to learn from observed payoffs and become, on average, persistently underconfident otherwise. The critical element is that agents who believe themselves to be of high ability choose to exert more effort. These results call for three further types of investigations. First, empirical research has detected average overconfidence in several settings (see Sections 1 and 6). Future work should test how average misplaced confidence varies across environments based on the relationships between effort and variance. Second, psychologists have connected "explanatory styles" to a range of outcomes (Seligman, 1991). Future work should test the connection between explanatory styles and beliefs about how effort affects the variance of rewards. Third, job search models with learning about ability (e.g., Papageorgiou, 2014; Groes et al., 2015) should consider the consequences of endogenously misplaced confidence, especially when the relationship between effort and variance differs by occupation. Finally, future theoretical work should investigate the dynamics of actual and estimated ability when ability is itself improved by the accumulation of effort over time. Overconfidence may then be self-fulfilling: environments that lead agents to become overconfident on average may also lead agents to attain higher ability on average.

9 Proof of Proposition 1

Because an NEF is characterized by its variance function (e.g., Morris, 1982),[37] we can write $w_{it} \triangleq w(\mu_{it}, \tilde{V}_{it}(\mu_{it}))$. Fix e_{it} and consider signals received at times $t-1$ and t. Label them s^H and s^L, with $s^H \neq s^L$. From equation (2.10) of Diaconis and Ylvisaker (1979), $\mu_{i(t+1)}$

[36] An earlier working paper version generalized the setting to allow agents to have misspecified models of the relationship between variance and effort. It showed that what determines whether underconfidence or overconfidence emerges is not the true data generating process but agents' beliefs about that process (i.e., the stories agents tell themselves about the relationship between effort and variance). The reason is that beliefs drive agents' asymmetric rates of learning from good and bad shocks.

[37] Technically, an NEF is characterized by the combination of its variance function and the domain of its variance function, but we will not be varying the latter.

does not depend on the order in which the signals were received. Therefore

$$
\begin{aligned}
&(1-w_{it}^{L})(1-w_{i(t-1)})\mu_{i(t-1)} + (1-w_{it}^{L})w_{i(t-1)}s^{L} + w_{it}^{L}s^{H} \\
=&(1-w_{it}^{H})(1-w_{i(t-1)})\mu_{i(t-1)} + (1-w_{it}^{H})w_{i(t-1)}s^{H} + w_{it}^{H}s^{L}, \qquad (16)
\end{aligned}
$$

with w_{it}^{k} indicating w_{it} following $s_{i(t-1)} = s^{k}$. This equation defines w_{it}^{L} as a function of w_{it}^{H}, s^{H}, and s^{L}. But w_{it}^{L} cannot depend on the value of s^{H} because w_{it}^{L} is determined before knowing that $s_{it} = s^{H}$. Therefore $w_{it}^{L} = (1-w_{it}^{H})w_{i(t-1)}$. Substituting in, w_{it}^{H} is independent of s^{L} if and only if

$$
w_{it}^{H} = \frac{w_{i(t-1)} - (w_{i(t-1)})^2}{1-(w_{i(t-1)})^2}.
$$

In that case, $w_{it}^{H} = w_{it}^{L}$ and (16) holds. Therefore w_{it} does not depend on μ_{it} if e_{it} is fixed.

Taylor-expanding w_{it} around μ_{i0} and using the foregoing result, we have

$$
w_{it} = w_{i0} + \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}}\frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}}(\mu_{it} - \mu_{i0}) + R_{it}, \qquad (17)
$$

where R_{it} is a polynomial in terms of order $(\mu_{it} - \mu_{i0})^2$ and higher. Using (14), Assumption 2, and $(\mu_{it} - \mu_{i0})^2 = [\mu_{it} - z_i - (\mu_{i0} - z_i)]^2$, observe that terms of order $(\mu_{it} - \mu_{i0})^2$ are of order δ and $\nu_{i(t-j)}$, for $j \in \{1, ..., t\}$.

I now prove (i) by induction. The proof of (ii) is directly analogous.

Induction step for part (i):

The induction hypothesis is that $\int_0^1 [\mu_{iN} - z_i]\,\mathrm{d}i < 0$ for some $N > 1$. Assumptions 1 and 2 ensure that $\mu_{iN} > 0$ and $e_{iN} > 0$. From equation (15),

$$
\begin{aligned}
\int_0^1 [\mu_{i(N+1)} - z_i]\,\mathrm{d}i =& \sum_{j=2}^{N+1} \int_0^1 \left[\frac{w_{i(N+1-j)}}{e_{i(N+1-j)}} Cov_{i(N+1-j)|z_i} \left[\prod_{k=1}^{j-1} (1 - w_{i(N+1-k)}), \nu_{i(N+1-j)} \right] \right] \mathrm{d}i \\
&+ \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i]\,\mathrm{d}i.
\end{aligned}
$$

Pull $1 - w_{iN}$ out of the product inside the covariance operator and linearly distribute the operator:

$$
\begin{aligned}
\int_0^1 [\mu_{i(N+1)} - z_i] \,\mathrm{d}i = & \sum_{j=2}^{N+1} \int_0^1 \left[\frac{w_{i(N+1-j)}}{e_{i(N+1-j)}} Cov_{i(N+1-j)|z_i} \left[\prod_{k=2}^{j-1} (1 - w_{i(N+1-k)}), \nu_{i(N+1-j)} \right] \right] \mathrm{d}i \\
& - \sum_{j=2}^{N+1} \int_0^1 \left[\frac{w_{i(N+1-j)}}{e_{i(N+1-j)}} Cov_{i(N+1-j)|z_i} \left[w_{iN} \prod_{k=2}^{j-1} (1 - w_{i(N+1-k)}), \nu_{i(N+1-j)} \right] \right] \mathrm{d}i \\
& + \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0}) [\mu_{i0} - z_i] \,\mathrm{d}i.
\end{aligned}
$$

The summand in the top line is zero for $j = 2$. Relabel the indices in the top line to obtain

$$
\begin{aligned}
\int_0^1 [\mu_{i(N+1)} - z_i] \,\mathrm{d}i = & \sum_{j=2}^{N} \int_0^1 \left[\frac{w_{i(N-j)}}{e_{i(N-j)}} Cov_{i(N-j)|z_i} \left[\prod_{k=1}^{j-1} (1 - w_{i(N-k)}), \nu_{i(N-j)} \right] \right] \mathrm{d}i \\
& - \sum_{j=2}^{N+1} \int_0^1 \left[\frac{w_{i(N+1-j)}}{e_{i(N+1-j)}} Cov_{i(N+1-j)|z_i} \left[w_{iN} \prod_{k=2}^{j-1} (1 - w_{i(N+1-k)}), \nu_{i(N+1-j)} \right] \right] \mathrm{d}i \\
& + \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0}) [\mu_{i0} - z_i] \,\mathrm{d}i.
\end{aligned}
$$

Substitute for the top line from (15):

$$
\begin{aligned}
\int_0^1 [\mu_{i(N+1)} - z_i] \,\mathrm{d}i = & \int_0^1 [\mu_{iN} - z_i] \,\mathrm{d}i \\
& - \sum_{j=2}^{N+1} \int_0^1 \frac{w_{i(N+1-j)}}{e_{i(N+1-j)}} Cov_{i(N+1-j)|z_i} \left[w_{iN} \prod_{k=2}^{j-1} (1 - w_{i(N+1-k)}), \nu_{i(N+1-j)} \right] \mathrm{d}i \\
& + \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0}) [\mu_{i0} - z_i] \,\mathrm{d}i.
\end{aligned}
$$

The induction hypothesis then implies

$$
\begin{aligned}
\int_0^1 [\mu_{i(N+1)} - z_i] \,\mathrm{d}i < & - \sum_{j=2}^{N+1} \int_0^1 \frac{w_{i(N+1-j)}}{e_{i(N+1-j)}} Cov_{i(N+1-j)|z_i} \left[w_{iN} \prod_{k=2}^{j-1} (1 - w_{i(N+1-k)}), \nu_{i(N+1-j)} \right] \mathrm{d}i \\
& + \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0}) [\mu_{i0} - z_i] \,\mathrm{d}i.
\end{aligned}
$$

Substitute for w_{iN} from (17) to obtain:

$$\int_0^1 [\mu_{i(N+1)} - z_i]\,\mathrm{d}i < -\sum_{j=1}^{N} \int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} \frac{w_{i(N-j)}}{e_{i(N-j)}} Cov_{i(N-j)|z_i} \left[(\mu_{iN} - z_i) \prod_{k=1}^{j-1} (1 - w_{i(N-k)}),\, \nu_{i(N-j)} \right] \mathrm{d}i$$
$$+ \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i]\,\mathrm{d}i$$
$$- \sum_{j=1}^{N} \int_0^1 \frac{w_{i(N-j)}}{e_{i(N-j)}} Cov_{i(N-j)|z_i} \left[R_{iN} \prod_{k=1}^{j-1} (1 - w_{i(N-k)}),\, \nu_{i(N-j)} \right] \mathrm{d}i.$$

The final line becomes arbitrarily small as δ and the moments of the ν_{it} become small. Thus, under the conditions of the proposition,

$$\int_0^1 [\mu_{i(N+1)} - z_i]\,\mathrm{d}i < -\sum_{j=1}^{N} \int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} \frac{w_{i(N-j)}}{e_{i(N-j)}} Cov_{i(N-j)|z_i} \left[(\mu_{iN} - z_i) \prod_{k=1}^{j-1} (1 - w_{i(N-k)}),\, \nu_{i(N-j)} \right] \mathrm{d}i$$
$$+ \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i]\,\mathrm{d}i + \int_0^1 X_{i(N+1)}\,\mathrm{d}i,$$

where $X_{i(N+1)}$ is arbitrarily small. Use the relationship between covariances and expectations, the fact that the ν_{it} are mean-zero, and the law of iterated expectations:

$$\int_0^1 [\mu_{i(N+1)} - z_i]\,\mathrm{d}i < -\int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} (\mu_{iN} - z_i) \sum_{j=1}^{N} \frac{w_{i(N-j)}}{e_{i(N-j)}} \nu_{i(N-j)} \prod_{k=1}^{j-1} (1 - w_{i(N-k)})\,\mathrm{d}i$$
$$+ \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i]\,\mathrm{d}i + \int_0^1 X_{i(N+1)}\,\mathrm{d}i.$$

Substitute for the summation on the top line from (14):

$$\int_0^1 [\mu_{i(N+1)} - z_i]\,\mathrm{d}i$$
$$< -\int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} (\mu_{iN} - z_i)^2\,\mathrm{d}i + \int_0^1 \left(\prod_{k=1}^{N} (1 - w_{i(N+1-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i]\,\mathrm{d}i$$
$$+ \int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} (\mu_{iN} - z_i) \left(\prod_{k=1}^{N-1} (1 - w_{i(N-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i]\,\mathrm{d}i + \int_0^1 X_{i(N+1)}\,\mathrm{d}i.$$

If $\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0$, then Assumption 2 implies that

$$\begin{aligned}
&\int_0^1 [\mu_{i(N+1)} - z_i]\,\mathrm{d}i - \int_0^1 X_{i(N+1)}\,\mathrm{d}i \\
<& -\int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}}\frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}}(\mu_{iN} - z_i)^2\,\mathrm{d}i \\
&+ \delta \int_0^1 \left(\prod_{k=1}^{N-1}(1 - w_{i(N-k)})\right)(1 - w_{i0})\left[(1 - w_{iN}) + \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}}\frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}}|\mu_{iN} - z_i|\right]\,\mathrm{d}i.
\end{aligned}$$

The right-hand side is strictly negative for all $\delta < \hat{\delta}$ when $\hat{\delta}$ is defined as

$$\hat{\delta} \triangleq \frac{\int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}}\frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}}(\mu_{iN} - z_i)^2\,\mathrm{d}i}{\int_0^1 \left(\prod_{k=1}^{N-1}(1 - w_{i(N-k)})\right)(1 - w_{i0})\left[(1 - w_{iN}) + \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}}\frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}}|\mu_{iN} - z_i|\right]\,\mathrm{d}i}.$$

$\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0$ implies that $\hat{\delta} > 0$, and Lemma 1 shows that the conditions of part (i) imply that $\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0$. Because $X_{i(N+1)}$ becomes arbitrarily small as δ and the moments of the ν_{it} become small, there exists $\hat{\hat{\delta}} \in (0, \hat{\delta})$ such that, when the conditions of the proposition and $\delta \leq \hat{\hat{\delta}}$ hold,

$$\int_0^1 [\mu_{i(N+1)} - z_i]\,\mathrm{d}i < 0.$$

We have proved the induction step.

Basis step for part (i):

From (13) and properties of ν_{i1},

$$\begin{aligned}
\int_0^1 [\mu_{i2} - z_i]\,\mathrm{d}i &= \int_0^1 \left[(1 - w_{i1})(\mu_{i1} - z_i) + \frac{w_{i1}}{e_{i1}}\nu_{i1}\right]\,\mathrm{d}i \\
&= \int_0^1 \left[(1 - w_{i1})(1 - w_{i0})(\mu_{i0} - z_i) + (1 - w_{i1})\frac{w_{i0}}{e_{i0}}\nu_{i0} + \frac{w_{i1}}{e_{i1}}\nu_{i1}\right]\,\mathrm{d}i \\
&= \int_0^1 \left[(1 - w_{i1})(1 - w_{i0})(\mu_{i0} - z_i) + (1 - w_{i1})\frac{w_{i0}}{e_{i0}}\nu_{i0}\right]\,\mathrm{d}i.
\end{aligned}$$

Substitute for w_{i1} from (17) and use properties of ν_{i0}:

$$\int_0^1 [\mu_{i2} - z_i] \,\mathrm{d}i = \int_0^1 (1 - w_{i1})(1 - w_{i0})(\mu_{i0} - z_i) \,\mathrm{d}i - \int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} (\mu_{i1} - z_i) \frac{w_{i0}}{e_{i0}} \nu_{i0} \,\mathrm{d}i \\ - \int_0^1 R_{i1} \frac{w_{i0}}{e_{i0}} \nu_{i0} \,\mathrm{d}i.$$

Use (12) and properties of ν_{i0}:

$$\int_0^1 [\mu_{i2} - z_i] \,\mathrm{d}i + \int_0^1 R_{i1} \frac{w_{i0}}{e_{i0}} \nu_{i0} \,\mathrm{d}i = - \int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} \left(\frac{w_{i0}}{e_{i0}} \right)^2 \nu_{i0}^2 \,\mathrm{d}i + \int_0^1 (1 - w_{i1})(1 - w_{i0})(\mu_{i0} - z_i) \,\mathrm{d}i.$$

Assumption 2 implies

$$\int_0^1 [\mu_{i2} - z_i] \,\mathrm{d}i + \int_0^1 R_{i1} \frac{w_{i0}}{e_{i0}} \nu_{i0} \,\mathrm{d}i < - \int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} \left(\frac{w_{i0}}{e_{i0}} \right)^2 \nu_{i0}^2 \,\mathrm{d}i + \delta \int_0^1 (1 - w_{i1})(1 - w_{i0}) \,\mathrm{d}i.$$

The right-hand side is strictly negative for all $\delta < \hat{\delta}$ when $\hat{\delta}$ is defined as

$$\hat{\delta} \triangleq \frac{\int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} \left(\frac{w_{i0}}{e_{i0}} \right)^2 \nu_{i0}^2 \,\mathrm{d}i}{\int_0^1 (1 - w_{i1})(1 - w_{i0}) \,\mathrm{d}i}.$$

$\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0$ implies that $\hat{\delta} > 0$, and Lemma 1 shows that the conditions in part (i) imply that $\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0$. The term with R_{i1} becomes arbitrarily small as δ and the moments of the ν_{it} become small. So there exists $\hat{\hat{\delta}} \in (0, \hat{\delta})$ such that, when the conditions of the proposition and $\delta \leq \hat{\hat{\delta}}$ hold,

$$\int_0^1 [\mu_{i2} - z_i] \,\mathrm{d}i < 0.$$

We have proved the basis step.

References

Ali, S. Nageeb (2011) "Learning self-control," *The Quarterly Journal of Economics*, Vol. 126, No. 2, pp. 857–893.

Astebro, Thomas, Holger Herz, Ramana Nanda, and Roberto A. Weber (2014) "Seeking the roots of entrepreneurship: Insights from behavioral economics," *Journal of Economic Perspectives*, Vol. 28, No. 3, pp. 49–70.

Barber, Brad M. and Terrance Odean (2001) "Boys will be boys: Gender, overconfidence, and common stock investment," *The Quarterly Journal of Economics*, Vol. 116, No. 1, pp. 261–292.

——— (2002) "Online investors: Do the slow die first?" *The Review of Financial Studies*, Vol. 15, No. 2, pp. 455–488.

Barndorff-Nielsen, O. (2006) "Exponential families," in Samuel Kotz, Campbell B. Read, N. Balakrishnan, Brani Vidakovic, and Norman L. Johnson eds. *Encyclopedia of Statistical Sciences*: American Cancer Society.

Barndorff-Nielsen, Ole E. (1978) *Information and Exponential Families in Statistical Theory*: John Wiley & Sons, Ltd.

Beck, Aaron T. (2002) "Cognitive models of depression," in Robert L. Leahy and E. Thomas Dowd eds. *Clinical Advances in Cognitive Psychotherapy: Theory and Application*, New York: Springer Publishing Company, pp. 29–61.

Bénabou, Roland and Jean Tirole (2002) "Self-confidence and personal motivation," *The Quarterly Journal of Economics*, Vol. 117, No. 3, pp. 871–915.

Benoît, Jean-Pierre and Juan Dubra (2011) "Apparent overconfidence," *Econometrica*, Vol. 79, No. 5, pp. 1591–1625.

Blavatskyy, Pavlo R. (2009) "Betting on own knowledge: Experimental test of overconfidence," *Journal of Risk and Uncertainty*, Vol. 38, No. 1, pp. 39–49.

Brunnermeier, Markus K. and Jonathan A. Parker (2005) "Optimal expectations," *American Economic Review*, Vol. 95, No. 4, pp. 1092–1118.

Burnside, Craig, Bing Han, David Hirshleifer, and Tracy Yue Wang (2011) "Investor overconfidence and the forward premium puzzle," *The Review of Economic Studies*, Vol. 78, No. 2, pp. 523–558.

Cappelli, Peter and Martin J. Conyon (2018) "What do performance appraisals do?" *ILR Review*, Vol. 71, No. 1, pp. 88–116.

Chen, Si and Hannah Schildberg-Hörisch (2019) "Looking at the bright side: The motivational value of confidence," *European Economic Review*, Vol. 120, p. 103302.

Clark, Jeremy and Lana Friesen (2009) "Overconfidence in forecasts of own performance: An experimental study," *The Economic Journal*, Vol. 119, No. 534, pp. 229–251.

Compte, Olivier and Andrew Postlewaite (2004) "Confidence-enhanced performance," *American Economic Review*, Vol. 94, No. 5, pp. 1536–1557.

Consonni, Guido and Piero Veronese (1992) "Conjugate priors for exponential families having quadratic variance functions," *Journal of the American Statistical Association*, Vol. 87, No. 420, pp. 1123–1127.

Daniel, Kent and David Hirshleifer (2015) "Overconfident investors, predictable returns, and excessive trading," *Journal of Economic Perspectives*, Vol. 29, No. 4, pp. 61–88.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam (1998) "Investor psychology and security market under- and overreactions," *The Journal of Finance*, Vol. 53, No. 6, pp. 1839–1885.

de Quidt, Jonathan and Johannes Haushofer (2017) "Depression through the lens of economics: A research agenda," in Christopher B. Barrett, Michael R. Carter, and Jean-Paul Chavas eds. *The Economics of Poverty Traps*.

Diaconis, Persi and Donald Ylvisaker (1979) "Conjugate priors for exponential families," *The Annals of Statistics*, Vol. 7, No. 2, pp. 269–281.

Fudenberg, Drew, Gleb Romanyuk, and Philipp Strack (2017) "Active learning with a misspecified prior," *Theoretical Economics*, Vol. 12, No. 3, pp. 1155–1189.

Gervais, Simon and Itay Goldstein (2007) "The positive effects of biased self-perceptions in firms," *Review of Finance*, Vol. 11, No. 3, pp. 453–496.

Gervais, Simon and Terrance Odean (2001) "Learning to be overconfident," *The Review of Financial Studies*, Vol. 14, No. 1, pp. 1–27.

Gibbs, Michael J. (1991) "An economic approach to process in pay and performance appraisals," working paper.

Gneezy, Uri, John A. List, Jeffrey A. Livingston, Xiangdong Qin, Sally Sadoff, and Yang Xu (2019) "Measuring success in education: The role of effort on the test itself," *American Economic Review: Insights*, Vol. 1, No. 3, pp. 291–308.

Graham, John R., Campbell R. Harvey, and Hai Huang (2009) "Investor competence, trading frequency, and home bias," *Management Science*, Vol. 55, No. 7, pp. 1094–1106.

Groes, Fane, Philipp Kircher, and Iourii Manovskii (2015) "The U-shapes of occupational mobility," *The Review of Economic Studies*, Vol. 82, No. 2, pp. 659–692.

Gutiérrez-Peña, E. and A. F. M. Smith (1997) "Exponential and Bayesian conjugate families: Review and extensions," *Test*, Vol. 6, No. 1, pp. 1–90.

Hall, Brian J. and Jeffrey B. Liebman (1998) "Are CEOs really paid like bureaucrats?" *The Quarterly Journal of Economics*, Vol. 113, No. 3, pp. 653–691.

Hall, Brian J. and Kevin J. Murphy (2002) "Stock options for undiversified executives," *Journal of Accounting and Economics*, Vol. 33, No. 1, pp. 3–42.

Heidhues, Paul, Botond Köszegi, and Philipp Strack (2018) "Unrealistic expectations and misguided learning," *Econometrica*, Vol. 86, No. 4, pp. 1159–1214.

Hestermann, Nina and Yves Le Yaouanq (2021) "Experimentation with self-serving attribution biases," *American Economic Journal: Microeconomics*, Vol. 13, No. 3, pp. 198–237.

Hirshleifer, David (2001) "Investor psychology and asset pricing," *The Journal of Finance*, Vol. 56, No. 4, pp. 1533–1597.

——— (2015) "Behavioral finance," *Annual Review of Financial Economics*, Vol. 7, No. 1, pp. 133–159.

Hoffman, Mitchell and Stephen V. Burks (2020) "Worker overconfidence: Field evidence and implications for employee turnover and firm profits," *Quantitative Economics*, Vol. 11, No. 1, pp. 315–348.

Jehiel, Philippe (2018) "Investment strategy and selection bias: An equilibrium perspective on overoptimism," *American Economic Review*, Vol. 108, No. 6, pp. 1582–97.

Kirchler, Erich and Boris Maciejovsky (2002) "Simultaneous over- and underconfidence: Evidence from experimental asset markets," *Journal of Risk and Uncertainty*, Vol. 25, No. 1, pp. 65–85.

Köszegi, Botond (2006) "Ego utility, overconfidence, and task choice," *Journal of the European Economic Association*, Vol. 4, No. 4, pp. 673–707.

Krähmer, Daniel (2007) "Equilibrium learning in simple contests," *Games and Economic Behavior*, Vol. 59, No. 1, pp. 105–131.

Larkin, Ian and Stephen Leider (2012) "Incentive schemes, sorting, and behavioral biases of employees: Experimental evidence," *American Economic Journal: Microeconomics*, Vol. 4, No. 2, pp. 184–214.

Layard, Richard and David M. Clark (2015) *Thrive: How Better Mental Health Care Transforms Lives and Saves Money*, Princeton, New Jersey: Princeton University Press.

Lybbert, Travis J. and Bruce Wydick (2018) "Poverty, aspirations, and the economics of hope," *Economic Development and Cultural Change*, Vol. 66, No. 4, pp. 709–753.

Malmendier, Ulrike and Geoffrey Tate (2005) "CEO overconfidence and corporate investment," *The Journal of Finance*, Vol. 60, No. 6, pp. 2661–2700.

——— (2008) "Who makes acquisitions? CEO overconfidence and the market's reaction," *Journal of Financial Economics*, Vol. 89, No. 1, pp. 20–43.

——— (2015) "Behavioral CEOs: The role of managerial overconfidence," *Journal of Economic Perspectives*, Vol. 29, No. 4, pp. 37–60.

Moore, Don A. and Paul J. Healy (2008) "The trouble with overconfidence," *Psychological Review*, Vol. 115, No. 2, pp. 502–517.

Morris, Carl N. (1982) "Natural exponential families with quadratic variance functions," *The Annals of Statistics*, Vol. 10, No. 1, pp. 65–80.

——— (1983) "Natural exponential families with quadratic variance functions: Statistical theory," *The Annals of Statistics*, Vol. 11, No. 2, pp. 515–529.

——— (2006) "Natural exponential families," in Samuel Kotz, Campbell B. Read, N. Balakrishnan, Brani Vidakovic, and Norman L. Johnson eds. *Encyclopedia of Statistical Sciences*: American Cancer Society.

Morris, Carl N. and Kari F. Lock (2009) "Unifying the named natural exponential families and their relatives," *The American Statistician*, Vol. 63, No. 3, pp. 247–253.

Murad, Zahra, Martin Sefton, and Chris Starmer (2016) "How do risk attitudes affect measured confidence?" *Journal of Risk and Uncertainty*, Vol. 52, No. 1, pp. 21–46.

Murphy, Kevin J. (1999) "Executive compensation," in Orley C. Ashenfelter and David Card eds. *Handbook of Labor Economics*, Vol. 3: Elsevier, pp. 2485–2563.

Nicholls, J. G. and A. T. Miller (1984) "Development and its discontents: The differentiation of the concept of utility," in J. G. Nicholls ed. *The Development of Achievement Motivation*, Greenwich, CT: JAI Press, pp. 185–218.

Niederle, Muriel and Lise Vesterlund (2007) "Do women shy away from competition? Do men compete too much?" *The Quarterly Journal of Economics*, Vol. 122, No. 3, pp. 1067–1101.

Papageorgiou, Theodore (2014) "Learning your comparative advantages," *The Review of Economic Studies*, Vol. 81, No. 3, pp. 1263–1295.

Seligman, Martin E. P. (1991) *Learned Optimism*, New York: Alfred A. Knopf.

Silva, Francisco (2017) "Inducing overconfidence," *Economic Inquiry*, Vol. 55, No. 1, pp. 451–460.

van den Steen, Eric (2004) "Rational overoptimism (and other biases)," *American Economic Review*, Vol. 94, No. 4, pp. 1141–1151.

Urbig, Diemo, Julia Stauf, and Utz Weitzel (2009) "What is your level of overconfidence? A strictly incentive compatible measurement of absolute and relative overconfidence," Working Paper 09-20, Utrecht School of Economics.

Van Nieuwerburgh, Stijn and Laura Veldkamp (2006) "Learning asymmetries in real business cycles," *Journal of Monetary Economics*, Vol. 53, No. 4, pp. 753–772.

Zábojník, Ján (2004) "A model of rational bias in self-assessments," *Economic Theory*, Vol. 23, No. 2, pp. 259–282 (2004).

Appendix

Appendix A extends the analysis of Section 2 to the case of forward-looking agents. Appendix B provides numerical examples. Appendix C formally demonstrates why an individual agent does not expect to become overconfident.

A Forward-Looking Effort Choices

In Section 2, I assumed that agents choose effort myopically. I now extend the analysis of Section 2 to the case of forward-looking agents, who account for the informational value of their effort choices.[38]

Agents know the world will end after period 2. Agents now choose effort to maximize expected present value with per-period discount factor $\beta \in (0,1)$. For exposition, follow Section 2 in assuming that z_i and Σ_{i0} are the same across agents and also assume that $c_i(\cdot)$ is the same for all agents. Following the previous derivation,

$$\bar{\mu}_2 = \bar{z} - \int_0^1 w(e_{i1}, \Sigma_{i1}) \, [\mu_{i1} - \bar{\mu}_1] \, \mathrm{d}i, \tag{A-1}$$

with

$$\int_0^1 w(e_{i1}, \Sigma_{i1}) \, [\mu_{i1} - \bar{\mu}_1] \, \mathrm{d}i = w_{i0}^2 \, \frac{f(e_{i0})}{e_{i0}^2} \int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}} \frac{\mathrm{d}e_{i1}}{\mathrm{d}\mu_{i1}} \, \mathrm{d}i. \tag{A-2}$$

In the final period, agents choose effort such that

$$c_i'(e_{i2}^*) = \mu_{i2}.$$

Because e_{i2}^* is independent of Σ_{i2}, agent i's period 2 value function is also independent of Σ_{i2}. We have

$$V_{i2}(\mu_{i2}) = \max_{e_{i2}} E_{i2} \left[e_{i2} z_i - c(e_{i2}) \right]$$

and, by the envelope theorem,

$$V_{i2}'(\mu_{i2}) = e_{i2}^*(\mu_{i2}).$$

In period 1, agent i solves:

$$V_{i1}(\mu_{i1}, \Sigma_{i1}) = \max_{e_{i1}} E_{i1} \Bigg[e_{i1} z_i - c_i(e_{i1}) + \beta V_{i2}(\mu_{i2}) \Bigg].$$

The first-order condition is:

$$c_i'(e_{i1}) = \mu_{i1} + \beta E_{i1} \left[e_{i2}^*(\mu_{i2}) \, \frac{\mathrm{d}\mu_{i2}}{\mathrm{d}e_{i2}} \right].$$

[38] Heidhues et al. (2018) also extend their primary analysis to forward-looking agents. I cannot use their proof technique because they are interested in the long-run limit distribution of beliefs. Here beliefs always do eventually converge to the true value of ability. I am instead interested in average beliefs after a finite number of observations. The main text describes further differences between the two papers.

The new term on the right-hand side captures agents' recognition of the informational value of effort. Observe that:

$$E_{i1}\left[\frac{\mathrm{d}\mu_{i2}}{\mathrm{d}e_{i1}}\right] = E_{i1}\left[\frac{\partial w(e_{i1},\Sigma_{i1})}{\partial e_{i1}}\left[z_i-\mu_{i1}+\frac{\sqrt{f(e_{i1})}}{e_{i1}}\epsilon_{i1}\right]+\frac{1}{2}w(e_{i1},\Sigma_{i1})\frac{f'(e_{i1})}{e_{i1}\sqrt{f(e_{i1})}}\epsilon_{i1}\right.$$
$$\left.-\,w(e_{i1},\Sigma_{i1})\frac{\sqrt{f(e_{i1})}}{e_{i1}^2}\epsilon_{i1}\right]$$
$$=0.$$

A Bayesian agent does not expect his actions to change his central estimate in one direction or the other. The first-order condition becomes

$$0 = \mu_{i1} - c_i'(e_{i1}^*) + \beta Cov_{i1}\left(e_{i2}^*(\mu_{i2}),\ \frac{\mathrm{d}\mu_{i2}}{\mathrm{d}e_{i1}}\right),$$

which defines e_{i1}^* as a function of μ_{i1} and Σ_{i1}. For given e_{i1}^*, μ_{i1} does not affect the covariance in the second line. The right-hand side therefore depends directly on μ_{i1} only through the first term, so e_{i1}^* increases in μ_{i1} if the second-order condition holds at time 1.[39] In that case, the right-hand side of equation (A-2) is strictly positive if $\partial w_{i1}/\partial e_{i1} > 0$ and is strictly negative if $\partial w_{i1}/\partial e_{i1} < 0$. From equation (A-1), forward-looking agents with $\partial w_{i1}/\partial e_{i1} > 0$ become underconfident on average by period 2 and forward-looking agents with $\partial w_{i1}/\partial e_{i1} < 0$ become overconfident on average by period 2. This is the same result as in the analysis of myopic agents, even though the chosen levels of effort will be different.

B Numerical Examples

I now consider numerical examples, using an infinite-horizon version of the setting from Section 2. Under the parameterization, the results are nearly identical to instead using the truncated-normal prior from Section 4 with $a_i = 0$ and $b_i = \infty$. Let $z_i = \mu_{i0} = 20$, $\Sigma_{i0} = 16$, $f(e_{it}) = 16\,[e_{it}]^\alpha$, and $c_i(e_{it}) = 10\,e_{it}^2$. Note that $\bar{z} = z_i$ and, defining χ as in Section 4, $\chi(e_{it}) = \alpha$. From (4), observe that $\partial w_{i0}/\partial e_{i0} > 0$ if and only if $\chi(e_{i0}) < 2$. I simulate 1 million agents.

The top left panel of Figure A-1 shows that, as demonstrated analytically, agents become overconfident on average for $\alpha > 2$ and become underconfident on average for $\alpha < 2$. The degree of over- or underconfidence is larger when α is farther from 2. As time passes, agents' average beliefs converge towards their true ability, but average biases remain even after 100 periods.

[39] It is well-known that the second-order condition may not hold in models with active learning because the value function is convex in the priors (Nyarko, 1994). Following Easley and Kiefer (1988), the second-order condition holds if $c_i''(\cdot)$ is sufficiently large or β is sufficiently small. Failure of the second-order condition can lead to very different types of policy programs, beyond the scope of the present analysis (see Balvers and Cosimano, 1993).

The remaining panels of Figure A-1 fix $\alpha = 0$, so that the variance of rewards is independent of effort in these three panels. The top right panel plots the distribution of μ_{it} for $t \in \{1, 2, 3, 4, 5, 10\}$. μ_{i1} is normally distributed but the other distributions are skewed.[40] The distribution of μ_{it} becomes progressively narrower as data accumulate.

The lower left panel of Figure A-1 confirms the results of Proposition 1: $\bar{\mu}_1$ equals $\bar{z}$ but $\bar{\mu}_t$ drops below $\bar{z}$ for $t > 1$. $\bar{\mu}_t$ does approach $\bar{z}$ again as t goes to infinity, but this approach is slow. The maximum average bias arises in period 4. The average bias is still 78% of this maximum in period 10 and 12% of this maximum in period 100. The circles show that agents' uncertainty about their ability does decline quickly as they observe additional data, but their beliefs nonetheless remain biased on average.

The lower right panel plots $Cov[\mu_{it}, w_{it}]$ (crosses) as well as the correlation (circles) between μ_{it} and w_{it}. The covariance and correlation are positive because agents with large μ_{it} choose high effort e_{it} and because w_{it} increases in e_{it}. The covariance is especially positive in early periods when agents are most uncertain about their own ability. The covariance approaches zero after the first few periods not because μ_{it} and w_{it} become uncorrelated over long horizons (the correlation in fact remains clearly positive even at long horizons) but because the variance of each variable declines strongly as agents become more certain of their ability.

B.1 Forward-Looking Agents

Now consider the implications of dynamically optimizing effort. Let $J_i(\mu_{it}, \Sigma_{it})$ denote the present value of agent i's optimal policy program from any time t with estimate μ_{it} and variance Σ_{it}. The agent's effort choices and value function solve the following Bellman equation:

$$J_i(\mu_{it}, \Sigma_{it}) = \max_{e_{it}} E_{it}\left[\pi_{it} - c_i(e_{it}) + \beta J_i(\mu_{i(t+1)}, \Sigma_{i(t+1)})\right],$$

where the agent's per-period discount factor is $\beta \in (0, 1)$. I solve the Bellman equation via value function iteration.[41] The left panel of Figure A-2 considers the role of α, again for a population of 1 million agents. The solid lines give $\bar{\mu}_t$ for myopic agents and the dotted lines give $\bar{\mu}_t$ for forward-looking agents with $\beta = 0.5$.[42] We see that mean beliefs follow remarkably similar paths in the two populations, even though the forward-looking agents on average choose greater effort for all α. The right panel sets $\alpha = 0$ and plots the percentage difference between myopic agents and forward-looking agents with $\beta = 0.9$. Active learning

[40]This skew arises because, first, $\mu_{i(t+1)}$ depends on w_{it} (which is a nonlinear function of μ_{it}) and, second, because the realized signal π_{it}/e_{it} is a nonlinear function of e_{it}. Both sources of skewness vanish when $f(e_{it}) = A\, e_{it}^2$ for some $A > 0$.

[41]I use the collocation method, with 10^2 Chebyshev nodes and 10^2 Chebyshev basis functions. I integrate via Gauss-Legendre quadrature. The domain of approximation for μ_{it} extends from 1 to 50, and the domain of approximation for Σ_{it} ranges from 0 to $(11/10)\Sigma_{i0}$. To simulate, I parallelize over 1 million agents using the University of Arizona High Performance Computing facility.

[42]I use $\beta = 0.5$ in this first analysis because the model becomes more difficult to solve with $\beta = 0.9$ and α large.

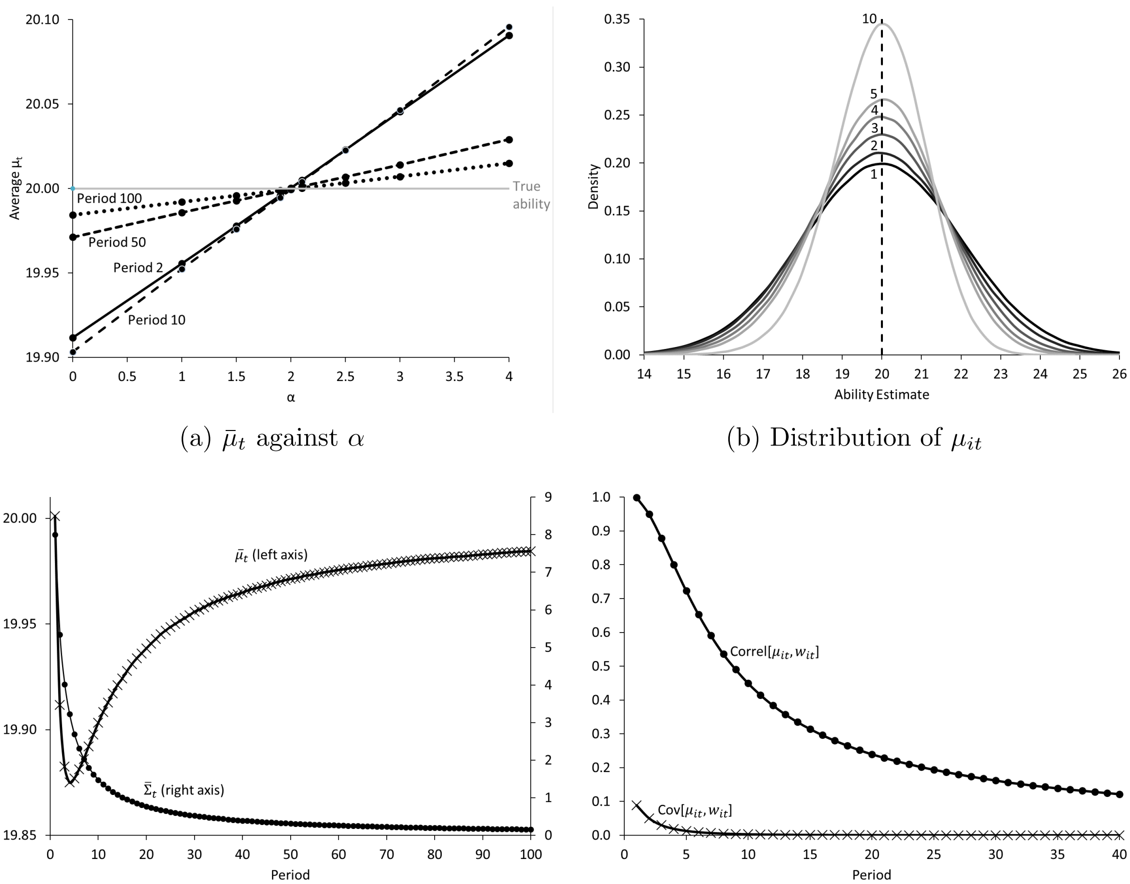

(a) $\bar{\mu}_t$ against α

(b) Distribution of μ_{it}

(c) $\bar{\mu}_t$ and $\bar{\Sigma}_t$

(d) $Cov(\mu_{it}, w_{it})$ and $Correl(\mu_{it}, w_{it})$

Figure A-1: The top left panel varies $\chi(\cdot) = \alpha$. In the other panels, $\chi(\cdot) = \alpha = 0$. The dashed horizontal (vertical) line in the top left (right) panel indicates the true ability $z_i = \bar{z} = 20$. All plots sample one million trajectories for ϵ_{it}.

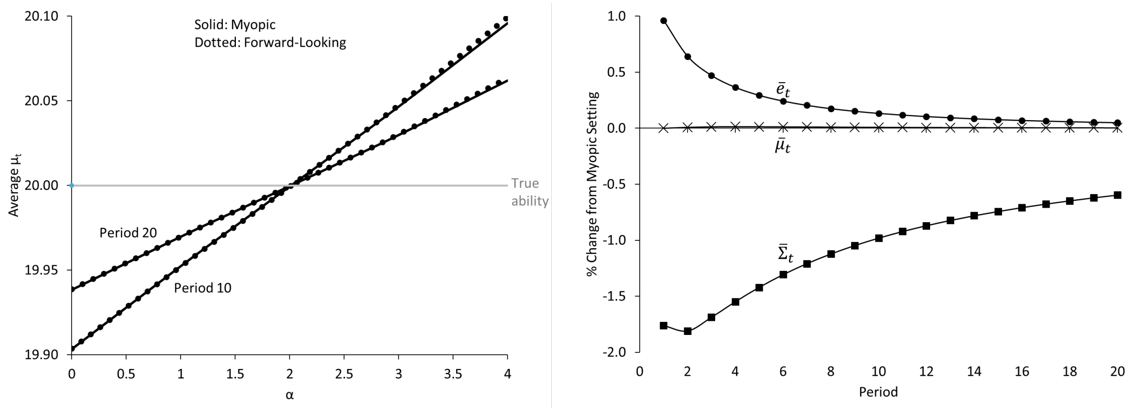

(a) $\bar{\mu}_t$ against α (b) $\bar{\mu}_t$, $\bar{e}_t$, and $\bar{\Sigma}_t$

Figure A-2: Left: Average estimates for myopic (solid) and forward-looking agents ($\beta = 0.5$, dotted) against $\chi(\cdot) = \alpha$. Right: The difference in average effort, estimates, and uncertainty between a myopic agent and a forward-looking agent ($\beta = 0.9$), with $\chi(\cdot) = \alpha = 0$. All plots sample 1 million trajectories for ϵ_{it}.

(i.e., experimentation) motivations matter: forward-looking agents choose higher effort on average (circles) in order to learn their own ability faster (squares). Remarkably, however, $\bar{\mu}_t$ (crosses) differs by less than 0.015% between the forward-looking and myopic agents. Forward-looking agents become underconfident almost exactly as shown in the lower left panel of Figure A-1, despite reducing their uncertainty faster.

C Agent i's beliefs are martingales

How can a rational, Bayesian agent expect to become either overconfident or underconfident? It is important to recognize that I have described the evolution of average beliefs under the true data generating process: I measure over- and underconfidence with respect to agents' true abilities. An individual agent does not know her true ability. In the setting of Section 2, let $\tilde{\mu}_{it}$ denote agent i's time 0 expectation of her own time t ability estimate μ_{it}:

$$\tilde{\mu}_{i1} = \int\int \left[[1 - w(e_{i0}, \Sigma_{i0})]\mu_{i0} + w(e_{i0}, \Sigma_{i0}) z_i + w(e_{i0}, \Sigma_{i0}) \frac{\sqrt{f(e_{i0})}}{e_{i0}} \epsilon_{i0} \right] h_{i0}(z_i, \epsilon_{i0})\, \mathrm{d}\epsilon_{i0}\, \mathrm{d}z_i$$
$$= \mu_0,$$

where $h_{i0}(\cdot, \cdot)$ is the joint density under agent i's time 0 prior. We earlier saw that period 1 own-ability estimates are μ_0 on average, and we now see that each agent indeed expects this outcome. Now consider agent i's expectations of her own period 2 beliefs:

$$\tilde{\mu}_{i2} = \int\int\int \left[[1 - w(e_{i1}, \Sigma_{i1})]\mu_{i1} + w(e_{i1}, \Sigma_{i1}) z_i + w(e_{i1}, \Sigma_{i1}) \frac{\sqrt{f(e_{i1})}}{e_{i1}} \epsilon_{i1} \right] h_{i0}(z_i, \epsilon_{i0}, \epsilon_{i1})\, \mathrm{d}\epsilon_{i0}\, \mathrm{d}\epsilon_{i1}\, \mathrm{d}z_i,$$

where μ_{i1} and e_{i1} depend on ϵ_{i0} and z_i. We have:

$$\tilde{\mu}_{i2} = \mu_{i0} - w(e_{i0}, \Sigma_{i0}) \frac{\sqrt{f(e_{i0})}}{e_{i0}} Cov_{i0}\left[w(e_{i1}, \Sigma_{i1}), \epsilon_{i0}\right] + [1 - w(e_{i0}, \Sigma_{i0})] Cov_{i0}\left[w(e_{i1}, \Sigma_{i1}), z_i\right],$$

where Cov_{i0} indicates a covariance based on the time 0 information set of agent i. Applying Stein's Lemma to each covariance yields:

$$\begin{aligned}\tilde{\mu}_{i2} =& \mu_0 - w_{i0}^2 \frac{f(e_{i0})}{e_{i0}^2} E_{i0}\left[\frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}} \frac{\mathrm{d}e_{i1}}{\mathrm{d}\mu_{i1}}\right] + w_{i0}[1 - w_{i0}]\Sigma_{i0} E_{i0}\left[\frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}} \frac{\mathrm{d}e_{i1}}{\mathrm{d}\mu_{i1}}\right] \\ =& \mu_{i0},\end{aligned}$$

where the middle term on the right-hand side of the first line is the right-hand side of equation (6) and the second equality uses, from equation (4), $w_{i0} f(e_{i0})/e_{i0}^2 = [1 - w_{i0}]\Sigma_{i0}$. Agents anticipate the covariance between w_{i1} and ϵ_{i0} captured in equation (6). Because agents do not know z_i, they also account for the covariance between w_{i1} and z_i. These two covariances exactly cancel, leaving only the μ_{i0}.[43] As is true in general, the Bayesian agent's posterior belief is a martingale. She does not expect her central estimate of her own ability to move in any particular way from her initial best estimate.[44] Each agent knows that the population of agents will tend to develop misplaced confidence on average, but each agent can do no better than to update as a Bayesian using whatever particular signals she happens to receive.

References from the Appendix

Balvers, Ronald J. and Thomas F. Cosimano (1993) "Periodic learning about a hidden state variable," *Journal of Economic Dynamics and Control*, Vol. 17, No. 5, pp. 805–827.

Easley, David and Nicholas M. Kiefer (1988) "Controlling a stochastic process with unknown parameters," *Econometrica*, Vol. 56, No. 5, pp. 1045–1064.

[43]Imagine that $\partial w_{i1}/\partial e_{i1} > 0$, so that agents become underconfident on average. High (low) values of z_i mean that the agent is underconfident (overconfident) in time 0. The resulting high (low) values of π_{i0} tend to correct the agent's initially mistaken beliefs, leading her to increase (decrease) her period 1 effort. In period 1, the agent learns the true z_i especially rapidly when effort is high (because z_i is high) and especially slowly when effort is low (because z_i is low). The agent's uncertainty about z_i drags the agent's expected central estimate up over time because she approaches the truth faster when z_i is high. For the Bayesian agent, the tendency of ϵ_{i0} to drag the expected central estimate down over time exactly cancels the tendency of z_i to drag the expected central estimate up over time.

[44]If we assume that agent i does not perceive other agents' rewards, z_i can differ by agent, and agent i applies the same prior over z_i to other agents as to herself, then we can interpret $\tilde{\mu}_{it}$ as agent i's belief at time t about other agents' average abilities. Agent i therefore always assesses the average agent to have ability μ_{i0}. If agents in fact become over- (under-) confident on average, then the population of agents demonstrates over- (under-) placement in the taxonomy of Moore and Healy (2008), allowing researchers to detect a better-than-average (worse-than-average) effect.

Heidhues, Paul, Botond Köszegi, and Philipp Strack (2018) "Unrealistic expectations and misguided learning," *Econometrica*, Vol. 86, No. 4, pp. 1159–1214.

Moore, Don A. and Paul J. Healy (2008) "The trouble with overconfidence," *Psychological Review*, Vol. 115, No. 2, pp. 502–517.

Nyarko, Yaw (1994) "On the convexity of the value function in Bayesian optimal control problems," *Economic Theory*, Vol. 4, No. 2, pp. 303–309.