Economic Theory
https://doi.org/10.1007/s00199-024-01618-0

RESEARCH ARTICLE

Rationally misplaced confidence

Derek Lemoine[1]

Received: 24 March 2023 / Accepted: 12 October 2024

Abstract
I show that persistent underconfidence and overconfidence can each arise from rational Bayesian learning when effort and ability are complementary. Which arises depends on the decision-making environment, and in particular on the effect that greater effort has on the variance of outcomes. 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 effort choices in turn depend on beliefs about their own ability. As one implication, I show that management can credibly induce additional effort from employees by designing feedback that generates average overconfidence through being conditionally vague.

Keywords Learning · Effort · Confidence · Ability · Talent · Feedback · Appraisals

JEL Classification D83 · G41 · M54

> *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 abil-

Conditionally Accepted Submission to Economic Theory.

✉ Derek Lemoine
dlemoine@email.arizona.edu

[1] Department of Economics, University of Arizona, NBER, and CEPR, McClelland Hall 401, 1130 E Helen St, Tucson, AZ 85721-0108, USA

Published online: 07 November 2024

ity, 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 may or may not 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 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 affects the variance of the signals that agents extract from the rewards they observe.[3]

I define average under- and overconfidence as emerging when the average of agents' posterior estimates differs from the average of agents' true abilities. I show that which emerges depends on how agents' effort choices affect the sensitivity of their posterior beliefs to new signals of their ability and on the tension within those signals between correcting old misperceptions and introducing new noise. A Bayesian's poste-

[1] Previous literature deviates in one or more of these dimensions, as described in Sect. 8. 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.

[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] For tractability, I model agents as choosing effort myopically in my primary analysis (e.g., Heidhues et al. 2018). The critical feature of effort choices is merely that they increase in agents' expectations of their own ability. Complementarity will tend to make this true for forward-looking agents as well. Moreover, forward-looking agents' effort choices converge to the myopically optimal effort choices as time passes. The appendix analyzes forward-looking agents.

rior beliefs are especially sensitive to new signals when the signals have low variance. 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. Here the variance of the unobserved shocks potentially depends on effort. Additional effort makes the signal of ability more precise if the shocks' variance decreases in effort or is independent of effort,[4] and it makes that signal less precise if the shocks' variance increases sufficiently strongly in effort. Complementarity between effort and ability means that an agent's chosen effort increases in her mean belief about her own ability. The precision of the signal an agent receives about her own ability thus generally depends on her mean belief, which means that the rate at which she adjusts her posterior towards a new signal depends on her mean belief.

In the short run, whether under- or overconfidence emerges on average depends on whether the difference between an agent's signal and prior mean tends to reflect random shocks or tends to reflect a correction to mistaken mean beliefs. First consider a case in which all agents initially happen to have mean beliefs that correspond to their own true ability: each agent's beliefs are initially well-calibrated to their own circumstances, which is one form of rational expectations. In that case, the noisy element to the signals that agents initially receive tends to make their beliefs miscalibrated. Agents whose effort choices lead them to learn this noise away faster tend to converge back to their true beliefs faster. Average misplaced confidence becomes determined by the agents who have not learned away this noise as quickly. So agents come to display average underconfidence if high effort makes the signal of their ability more precise (because they learn away overconfidence relatively quickly), agents come to display average overconfidence if high effort makes the signal of their ability less precise (because they learn away underconfidence relatively quickly), and agents display neither average overconfidence nor average underconfidence in the knife-edge case that the precision of their signal is independent of their effort choices.

Next consider a case in which all agents initially have the same mean belief but their true abilities are in fact symmetrically distributed around that mean belief. This case is an alternate form of rational expectations, in which each agent's mean belief is well-calibrated to the population average. Now the signals that agents receive tend to close the gap between their true abilities and their own expectations of their abilities. If that effect is sufficiently strong, then average misplaced confidence becomes determined by the agents who move their prior towards their true ability relatively slowly. So agents come to display average underconfidence if high effort makes the signal of their ability less precise (because initially overconfident agents reduce their estimates towards their true ability relatively quickly), agents come to display average overconfidence if high effort makes the signal of their ability more precise (because initially underconfident agents increase their estimates towards their true ability relatively quickly), and agents again display neither average overconfidence nor average underconfidence in the knife-edge case that the precision of their signal is independent of their effort choices.

As an example, let the variance of rewards be due to mean-zero external shocks whose symmetric distribution is independent of agents' effort choices. Like Emerson's

[4] When the shocks' variance 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 approaches infinity.

“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 each well-calibrated at time 0 (i.e., centered around each of their true abilities, as was described two paragraphs back). For some agents, the unobserved shock happens to take on a high value at time 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.

I show that a manager can take advantage of these dynamics to obtain more effort from employees who are learning about their own ability from feedback. The manager chooses the precision of the feedback that employees see. By making feedback sufficiently less informative when the manager observes employees choosing high effort, the manager induces employees to learn away bad shocks faster than good shocks. Importantly, this feedback rule is credible: it achieves its goal only if employees know the manager has designed it this way (in contrast to a rule that biased the feedback sent to employees), as the manager wants them to account for precision when updating their beliefs. Even though employees know the manager has designed feedback to induce them to become overconfident on average, they still can do no better than to update as Bayesians in response to any particular sequence of feedback.

The predicted dynamics of misplaced confidence 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 initially well-calibrated beliefs yet end 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.[5]

The next section presents an analytically transparent model with normal distributions. Sections 3 and 4 generalize the distributional assumptions and extend the predictions to infinitely many periods. Section 5 contains a numerical example. Sec-

[5] 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.

tion 6 models a manager designing feedback to obtain more effort from employees. Section 7 relates the theoretical setting to prominent empirical work on overconfidence. Section 8 reviews related literature. Section 9 proposes opportunities for further work. The final section contains additional analysis. The appendix analyzes forward-looking agents.

2 Closed-form derivation with normal distributions

Begin by considering an example that allows for closed-form solutions and accessible exposition.

2.1 Preliminaries

There is a continuum of agents, indexed by i and of measure 1. In each 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 $c_i(e_{it}) : \mathbb{R}^+ \to \mathbb{R}^+$, with $c_i(\cdot)$ twice continuously differentiable, $c_i'(0) = 0$, 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]

Agents' true abilities z_i are unknown to them. Average ability is strictly positive: $\int_0^1 z_i \,\mathrm{d}i > 0$. Agents update as Bayesians from the observed π_{it}. In this section, agent i's prior beliefs about her ability z_i are summarized by a normal distribution with mean $\mu_{i0} > 0$ and variance $\Sigma_{i0} \in (0, \infty)$, agents' prior beliefs are well-calibrated on average ($\mu_{i0} = \int_0^1 z_i \,\mathrm{d}i$, which implies μ_{i0} is independent of i), and all agents have the same prior (i.e., Σ_{i0} is also independent of i).[8] These assumptions permit a transparent derivation that highlights the main mechanisms. At the end of this section, I will analyze a case with heterogeneous priors. In subsequent sections, I will permit the prior to be non-normal, permit the prior to be heterogeneous, and permit $\mu_{i0} \neq \int_0^1 z_i \,\mathrm{d}i$.

[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). However, other work does explore substitutability between effort and ability, both empirically (Mueller and Dweck 1998) and theoretically (Deimen and Wirtz 2022).

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

[8] Rationality requires Bayesian updating but does not pin down agents' priors, which are always model primitives. See footnote 1 in van den Steen (2004) and citations therein. The assumption that $\mu_{i0} = \int_0^1 z_i \,\mathrm{d}i$ is a form of rational expectations assumption that restricts these primitives. Below equation (10), I discuss two stronger forms of rational expectations: one requires, in addition, that $\Sigma_{i0} = \mathrm{Var}[z_i]$, and the other requires that $\mu_{i0} = z_i$.

The function $f(\cdot)$ determines the role of luck (i.e., the variance of rewards), and its derivative determines how effort choices affect 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]:

$$\max_{e_{it} \geq 0} E_{it} \left[\pi_{it} - c_i(e_{it})\right],$$

where $E_{it}[\cdot]$ denotes expectations conditioned on the mean μ_{it} and variance Σ_{it} of agent i's beliefs at the beginning of time t, without knowledge of z_i or the realizations of ϵ_{it}. At an interior optimum, agent i'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 $c'_i(e^*_{it}) = \mu_{it}$.[11] Optimal effort e^*_{it} is an increasing function of μ_{it}. When $\mu_{it} < c'_i(0)$, the optimal choice of effort is $e^*_{it} = 0$.

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. The combination of normally distributed beliefs and normally distributed shocks generates a conjugate Bayesian updating rule:

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

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

Define

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

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 Le Yaouanq (2021) for an alternate formulation of locus of control.

10 The Appendix shows that the results of the two-period version of this analysis survive when agents are not myopic. The only difference with forward-looking agents is that they may adjust effort choices to actively experiment. Intuitively, there is no reason why active experimentation motives should oppose and overwhelm the mechanisms described here.

11 Convexity of the cost function guarantees that, when it exists, an interior solution is a local maximum.

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

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

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

Lemma 1 (Signal Quality and Effort) *Let $e_{it} > 0$. Then $\partial w(e_{it}, \Sigma_{it})/\partial e_{it} > 0$ if $\chi(e_{it}) < 2$, $\partial w(e_{it}, \Sigma_{it})/\partial e_{it} < 0$ if $\chi(e_{it}) > 2$, and $\partial w(e_{it}, \Sigma_{it})/\partial e_{it} = 0$ if $\chi(e_{it}) = 2$.*

Proof From (4),

$$\frac{\partial w(e_{it}, \Sigma_{it})}{\partial e_{it}} = (1 - w_{it})\frac{w_{it}}{e_{it}}[2 - \chi(e_{it})].$$

□

Additional effort increases the marginal effect of ability on the signal (i.e., the variance of ϵ_{it}/e_{it} declines in e_{it}), which works to increase w_{it}. But when $f'(\cdot) > 0$, additional effort also increases the variance of the signal, which works to reduce w_{it}. The second effect dominates if and only if $\chi(e_{it}) > 2$.[12]

Define the population-average central estimate $\bar{\mu}_t \triangleq \int_0^1 \mu_{it}\, \mathrm{d}i$ and the population-average ability $\bar{z} \triangleq \int_0^1 z_i\, \mathrm{d}i$. In the present example, $\bar{\mu}_0 = \bar{z}$. Imagine that a researcher can measure agents' average beliefs and average ability, as may be true in an experimental environment or in a large-sample econometric study. That researcher evaluates agents' confidence as follows:

Definition 1 Time t agents are *overconfident on average* if $\bar{\mu}_t > \bar{z}$ and are *underconfident on average* if $\bar{\mu}_t < \bar{z}$.

Observe that over- and underconfidence are defined with respect to true ability. Because individual agents do not know their own or others' true abilities, an individual agent cannot tell whether she displays over- or underconfidence. Although other definitions are possible, the above definitions are natural ones that fit many applications.[13]

[12] If $e_{it} = 0$, then $w_{it} = 0$ and the agent is in a zero-effort trap in which they stop learning and thus maintain $e_{it} = 0$ forever. This possibility is similar to the bandit-like models reviewed in Sect. 8. Although interesting for its connection to depression (de Quidt and Haushofer 2017) and to poverty traps (Lybbert and Wydick 2018), this possibility is not the mechanism of interest here. This possibility does, however, mean that the present section emphasizes predictions two periods ahead. In subsequent sections, the prior will make zero-effort traps impossible and will enable predictions arbitrarily many periods ahead.

[13] I express the setting in terms of a population of agents and results in terms of population averages, but one can equivalently interpret the setting as describing an individual j and the results as describing the expectation of μ_{jt} conditional on z_j. Definition 1 then says that we label the agent as over- or underconfident on average if a researcher who knows the agent's true ability expects the agent's posterior mean to be different from her true ability.

2.2 Analysis

Consider the population-average central estimate in period 1. From equation (5),

$$\begin{aligned}\bar{\mu}_1 &= \int_0^1 [(1-w_{i0})\mu_{i0} + w_{i0}s_{i0}]\,\mathrm{d}i \\ &= \int_0^1 \Big[\mu_{i0} + w(e_{i0},\Sigma_{i0})[z_i - \mu_{i0}]\Big]\,\mathrm{d}i + \int_0^1 w(e_{i0},\Sigma_{i0})\frac{\sqrt{f(e_{i0})}}{e_{i0}}\epsilon_{i0}\,\mathrm{d}i.\end{aligned}$$

The second integral is zero because ϵ_{i0} is mean zero and uncorrelated with either e_{i0} or Σ_{i0} across agents i. We extract w_{i0} from the first integral because it is independent of i by homogeneity of the time 0 prior. Therefore,

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

Because agents' beliefs are initially well-calibrated on average (i.e., because $\bar{\mu}_0 = \bar{z}$),

$$\bar{\mu}_1 = \bar{z}. \tag{6}$$

Agents' beliefs remain well-calibrated on average in period 1.

The population-average central estimate in period 2 is:

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

The second integral is zero because ϵ_{i1} is mean zero and uncorrelated with either e_{i1} or Σ_{i1} across agents i. Using that and (6),

$$\bar{\mu}_2 = \bar{z} - \int_0^1 w(e_{i1},\Sigma_{i1})\,(\mu_{i1} - z_i)\,\mathrm{d}i. \tag{7}$$

The covariance between w_{i1} and period 1 misplaced confidence determines whether beliefs remain well-calibrated on average in period 2. If w_{i1} tends to be large when agents are overconfident but small when agents are underconfident, then agents tend to learn away overconfidence relatively quickly but not to learn away underconfidence as quickly. Agents then become underconfident on average in period 2. If, instead, w_{i1} tends to be small when agents are overconfident but large when agents are underconfident, then agents tend to learn away underconfidence relatively quickly but not to learn away overconfidence as quickly. Agents then become overconfident on average in period 2. And if w_{i1} is uncorrelated with agents' period 1 confidence, then agents' beliefs remain well-calibrated on average in period 2.

The population-average central estimate in period $t+1$ is:

$$\bar{\mu}_{t+1} = \bar{z} - \int_0^1 w_{it}(\mu_{it} - z_i)\,\mathrm{d}i + (\bar{\mu}_t - \bar{z}), \tag{8}$$

using that ϵ_{it} is mean-zero and uncorrelated with either e_{it} or Σ_{it} across agents i. The first two terms are familiar from our analysis of $\bar{\mu}_2$. The third term appears only if average misplaced confidence has emerged in period t. It reflects the tendency of average misplaced confidence to persist. For average misplaced confidence to begin to disappear, the covariance between w_{it} and period t misplaced confidence must switch sign as t increases.[14]

We now analyze the sign of that covariance. Rewrite the integral from (7):

$$\int_0^1 w(e_{i1}, \Sigma_{i1})\,(\mu_{i1} - z_i)\,\mathrm{d}i = \int_0^1 w(e_{i1}, \Sigma_{i1})(\mu_{i1} - \bar{z})\,\mathrm{d}i - \int_0^1 w(e_{i1}, \Sigma_{i1})(z_i - \bar{z})\,\mathrm{d}i. \tag{9}$$

To progress towards an insightful analytic expression, assume either that z_i is normally distributed across agents or that z_i is constant across agents. In the former case, each integral on the right-hand side is the covariance between w_{i1} and a zero-mean normal variable, and in the latter case, the first integral on the right-hand side is such a covariance while the second integral on the right-hand side vanishes. In either case, use Stein's Lemma[15]:

$$\begin{aligned}\int_0^1 & w(e_{i1}, \Sigma_{i1})\,[\mu_{i1} - z_i]\,\mathrm{d}i \\ &= w_{i0}^2\left[\frac{f(e_{i0})}{e_{i0}^2} + \mathrm{Var}[z_i]\right]\int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}}\frac{\mathrm{d}e_{i1}(\mu_{i1})}{\mathrm{d}\mu_{i1}}\,\mathrm{d}i \\ &\quad - \mathrm{Var}[z_i]\int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}}\frac{\mathrm{d}e_{i1}(\mu_{i1})}{\mathrm{d}\mu_{i1}}\frac{\mathrm{d}\mu_{i1}}{\mathrm{d}z_i}\,\mathrm{d}i.\end{aligned}$$

Use $\mathrm{d}\mu_{i1}/\mathrm{d}z_i = w_{i0}$ from (5):

$$\int_0^1 w(e_{i1}, \Sigma_{i1})\,[\mu_{i1} - z_i]\,\mathrm{d}i = w_{i0}\left[w_{i0}\frac{f(e_{i0})}{e_{i0}^2} - (1 - w_{i0})\,\mathrm{Var}[z_i]\right]\int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}}\frac{\mathrm{d}e_{i1}(\mu_{i1})}{\mathrm{d}\mu_{i1}}\,\mathrm{d}i.$$

Substituting for w_{i0} from (4), equation (7) becomes:

$$\bar{\mu}_2 = \bar{z} - w_{i0}(1 - w_{i0})\Big[\Sigma_{i0} - \mathrm{Var}[z_i]\Big]\int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}}\frac{\mathrm{d}e_{i1}(\mu_{i1})}{\mathrm{d}\mu_{i1}}\,\mathrm{d}i. \tag{10}$$

[14] With normally distributed shocks, agent i has (with probability 1) $\mu_{it} < c_i'(0)$ for some t. So with probability 1, $e_{it} = 0$ and $\bar{\mu}_{t+1} = \bar{\mu}_t$ for all t sufficiently large. As a result, the present section does not contain firm predictions for $t > 2$. See footnote 12 for more on zero-effort traps, and see subsequent sections for settings that lack such traps and therefore permit additional analysis of misplaced confidence for $t > 2$.

[15] From equation (3), Σ_{i1} depends only on Σ_{i0} and e_{i0}. When μ_{i0} and Σ_{i0} are homogeneous, Σ_{i1} is constant over i and not random. In this case, the only random argument of w_{i1} is e_{i1}.

Misplaced confidence emerges on average in period 2 when the term in square brackets on the right-hand side is nonzero and the integral is nonzero.[16] From Lemma 1, the sign of $\partial w_{i1}/\partial e_{i1}$ depends on whether the elasticity $\chi(e_{i1})$ is greater or less than 2.

The two forces in square brackets determine the direction of average misplaced confidence. Begin by considering how potential heterogeneity in z_i affects misplaced confidence. This mechanism reflects the role of $z_i - \mu_{i0}$ in determining the direction in which agents update beliefs in response to the initial signal. This mechanism pushes agents to become underconfident on average in period 2 if and only if

$$\int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}} \frac{\mathrm{d}e_{i1}(\mu_{i1})}{\mathrm{d}\mu_{i1}} \,\mathrm{d}i < 0.$$

It otherwise pushes agents to become overconfident on average in period 2. Agents who happen to have high ability tend to receive surprisingly positive signals in period 0. A positive signal in period 0 leads them to increase their posterior mean and thus choose higher effort in period 1. If $\partial w_{i1}/\partial e_{i1} < 0$, then such agents learn about their high ability from the period 1 signal only slowly. In contrast, agents with low ability tend to receive surprisingly negative signals in period 0, which leads them to reduce their posterior mean and thus choose less effort in period 1. If $\partial w_{i1}/\partial e_{i1} < 0$, then such agents learn relatively quickly about their low ability from the period 1 signal. Heterogeneity in z_i therefore leads agents to become underconfident on average when $\partial w_{i1}/\partial e_{i1} < 0$ for all i. The opposite story holds if $\partial w_{i1}/\partial e_{i1} > 0$. In either case, the mechanism is the heterogeneous rates at which agents converge to their true abilities from initially mistaken priors.

To understand the second force in square brackets, consider a special case in which $\mathrm{Var}[z_i] = 0$, so that ability is homogeneous and every agent has well-calibrated beliefs in the initial period. In this case, the only force determining the direction in which agents update beliefs in period 1 is whether their news shocks ϵ_{i0} are positive or negative. Agents become underconfident on average in period 2 if and only if

$$\int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}} \frac{\mathrm{d}e_{i1}(\mu_{i1})}{\mathrm{d}\mu_{i1}} \,\mathrm{d}i > 0,$$

and agents become overconfident on average in period 2 if and only if this inequality is reversed. A positive ϵ_{i0} leads agents to become overconfident in period 1 and thereby to choose high effort in period 1. If $\partial w_{i1}/\partial e_{i1} > 0$, then such agents learn away the initial positive shock relatively quickly. In contrast, agents who receive a negative ϵ_{i0} choose less effort in period 1 and thereby learn away the negative shock relatively slowly. Heterogeneity in shocks ϵ_{i0} therefore leads agents to become underconfident on average when $\partial w_{i1}/\partial e_{i1} > 0$ for all i. The opposite story holds if $\partial w_{i1}/\partial e_{i1} < 0$. In either case, the mechanism is the heterogeneous rates at which agents learn away shocks.

The two mechanisms oppose each other. When $\mathrm{Var}[z_i]$ is large, there is substantial heterogeneity in ability and thus large scope for the prior mistakes that drive the

[16] From agents' first-order conditions, $\mathrm{d}e_{i1}/\mathrm{d}\mu_{i1} > 0$ if $\mu_{i1} \geq c_i'(0)$ and $\mathrm{d}e_{i1}/\mathrm{d}\mu_{i1} = 0$ otherwise.

first mechanism. When Σ_{i0} is large, prior mistakes are corrected relatively rapidly, so that most of the variation in beliefs is soon due to the effects of shocks and the second mechanism tends to be relatively strong. The first mechanism dominates when agents' mistakes about their true ability are more heterogeneous than their prior beliefs are diffuse (i.e., when $\mathrm{Var}[z_i] > \Sigma_{i0}$), and the second mechanism dominates when agents' mistakes about their true ability are less heterogeneous than their prior beliefs are diffuse (i.e., when $\mathrm{Var}[z_i] < \Sigma_{i0}$).[17]

Agents tend to maintain well-calibrated beliefs in period 2 in two cases. First, it could be that $\mathrm{Var}[z_i] = \Sigma_{i0}$, meaning that agents' prior beliefs match the population distribution of ability. This is a knife-edge case that could be imposed as a form of rational expectations. In this special case, the evolution of the population-average mean belief is equivalent to a Bayesian individual's expectation of how their own posterior mean would evolve over time (reinterpreting the i as indexing samples from an individual's potential sequences of news shocks rather than indexing individuals). By the martingale property of beliefs, the individual must not expect their beliefs to drift in any particular direction. And that is indeed what we see: $\bar{\mu}_2 = \bar{z}$, and because $\bar{z} = \mu_{i0}$ to the best of the agent's beliefs, $\bar{\mu}_2 = \mu_{i0}$.

Second, it could be that $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1} = 0$. However, from equation (4) and Lemma 1, this condition is truly a knife-edge case. We would need $f(\cdot)$ to have just the constant elasticity form in order to have $\partial w(e_{i1}, \Sigma_{i1})/\partial e_{i1}$ constant, and that form must have an elasticity of exactly 2 for that constant to be zero.

Finally, consider a stronger form of rational expectations in which each agent's prior happens to be well-calibrated around their own ability, so that $\mu_{i0} = z_i$ with z_i normally distributed. In that case, it is easy to see that equations (6) and (7) still hold, equation (9) becomes

$$\int_0^1 w(e_{i1}, \Sigma_{i1})\,(\mu_{i1} - z_i)\,\mathrm{d}i = \int_0^1 w(e_{i1}, \Sigma_{i1})(\mu_{i1} - \mu_{i0})\,\mathrm{d}i,$$

and equation (10) becomes

$$\bar{\mu}_2 = \bar{z} - w_{i0}(1 - w_{i0})\Sigma_{i0} \int_0^1 \frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}} \frac{\mathrm{d}e_{i1}(\mu_{i1})}{\mathrm{d}\mu_{i1}}\,\mathrm{d}i. \tag{11}$$

The prior mistake channel from (10) vanishes, leaving us with the shock channel whose strength grows with the variance of the time 0 prior.

Section 7 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

[17] When the first mechanism dominates in (10), the covariance in (11) can flip as t increases because prior mistakes are eventually narrowed. When the second mechanism dominates in (10), the covariance in (11) can flip as t increases because the posterior variance falls even as heterogeneity in agents' mistakes can increase (owing to heterogeneity in the sequences of shocks that agents receive).

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. When they are each especially uncertain about their own ability but they each have roughly similar abilities, 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). These predictions reverse when students have homogeneous, precise beliefs about their own ability but ability is in fact very heterogeneously distributed across them.

3 Non-normal shocks and priors from the natural exponential family

The model in Sect. 2 allows for a transparent derivation driven merely by first-order conditions, Bayes' Rule, and Stein's Lemma, but one may wonder whether misplaced confidence maintains the same sign beyond two periods, whether it arises under non-normal distributions, and whether it arises for priors that are not initially well-calibrated on average. Here I obtain concrete predictions for infinitely many periods by eliminating zero-effort traps, relax the distributional assumptions on priors and shocks, and relax the restriction that the prior mean is well-calibrated on average.

3.1 Preliminaries

Let rewards be $\pi_{it} = e_{it}z_i + \lambda\nu_{it}$, with the z_i strictly positive, the ν_{it} indicating potentially non-normal shocks that are not directly observed by agents, and $\lambda > 0$ a scaling parameter. Conditional on e_{it}, the ν_{it} are identically and independently distributed over time and agents, with mean zero.[18] As in Sect. 2, 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 $s_{it} \triangleq \pi_{it}/e_{it}$.

Letting p_x be the density function of random variable x, I assumed in Sect. 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 $\theta_{it} \triangleq \theta(z_i; e_{it})$ such that $p_{s_{it}}(s_{it}|e_{it}, \theta_{it})$ is a member of a regular

[18] The case with nonzero means is nested: 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.

natural exponential family (NEF) of distributions (Morris 2006)[19]:

$$p_{s_{it}}(s_{it}|e_{it}, \theta_{it}) = \exp[\theta_{it}s_{it} - M_i(\theta_{it}; e_{it})], \tag{12}$$

where M_i is continuous in e_{it} and where 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 (12) (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 used in Sect. 2). 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 $E_{it}[s_{it}|e_{it}, \theta_{it}] = M_i'(\theta_{it}; e_{it})$ and $z_i = E_{it}[s_{it}|e_{it}, \theta_{it}]$, 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 conditional on e_{it}.[21] Plugging $\theta(z_i; e_{it})$ into (12) 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 prominent and well-studied class:

$$V_i(z_i; e_{it}) = \zeta_{i2}(e_{it})\, z_i^2 + \zeta_{i1}(e_{it})\, z_i + \zeta_{i0}(e_{it}), \tag{13}$$

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 (12) and (13) for given e_{it}, 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 (13) was satisfied in Sect. 2, where the variance of s_{it} conditional on e_{it} and z_i had $\zeta_{i0}(e_{it}) = \frac{f(e_{it})}{e_{it}^2}$, $\zeta_{i1}(e_{it}) = 0$, and $\zeta_{i2}(e_{it}) = 0$ (the latter two relationships must be true for normal distributions).

[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_{it} 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} for given e_{it}.

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):

$$p_{\theta_{i0}}(\theta_{i0}|\mathcal{I}_{i0}) = K_{i0} \exp\left[n_{i0}x_{i0}\theta_{i0} - n_{i0}M_i(\theta_{i0}; e_{i0})\right], \tag{14}$$

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 $\theta(z_i; e_{i0})$ into (14) 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] I assume that the true z_i is within the support of the D-Y conjugate prior over z_i.

Because the variance function satisfies (13), the prior over z_i in (14) 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):

$$\tilde{V}_{i0}(\mu_{i0}) = \frac{V_i(\mu_{i0}; e_{i0})}{n_{i0} - \zeta_{i2}(e_{i0})}. \tag{15}$$

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 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 will also sometimes 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 by introducing an additional mechanism closer to that in bandit models (see Sect. 8). The first assumption will ensure that agents avoid low-confidence, zero-effort traps[24]:

Assumption 1 (Positive Signals) $p_{s_{it}}$ 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 Sect. 4.

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

---

22 Consonni and Veronese (1992) show that the D-Y conjugate prior over z_i is conjugate to (12) expressed in terms of z_i when the variance function is quadratic as in (13) but not necessarily otherwise.

23 In Sect. 2, $\Sigma_{i0} < \infty$ implied $n_{i0} > 0$ and, because $\zeta_{i2}(\cdot) = 0$ for normal distributions, $n_{i0} > \zeta_{i2}(e_{i0})$.

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. See footnote 12 on zero-effort traps.

Assumption 2 (Priors Not Too Miscalibrated) For all i and for given $\delta \geq 0$, $\mu_{i0} > 0$ and $\mu_{i0} \in [z_i - \delta, z_i + \delta]$.

This assumption permits the possibility that each prior is calibrated to the true population distribution of ability (the form of rational expectations analyzed in Sect. 2). This assumption also permits the possibility that each prior is well-calibrated at the individual level (the stronger form of rational expectations that led to equation (11)). But it also permits the possibility that agents' mean beliefs do not match their abilities. It is important to permit the possibility of well-calibrated priors because any Bayesian model can generate misplaced confidence if priors are miscalibrated but would not typically do so when priors are well-calibrated. And it is important to permit the possibility of miscalibrated priors to highlight that current results are not knife-edge.

Assumption 2 requires that priors not be "too" miscalibrated. We will also be interested in the implications of relaxing Assumption 2. In that case, we replace it with the following assumption:

Assumption 2' (Priors of Positive Ability) For all i, $\mu_{i0} > 0$.

Assumption 2' requires only that each agent initially believes they have strictly positive ability. We will see that bounding the degree of initial miscalibration (as in Assumption 2) is essential for predictions about the sign of misplaced confidence in the short run but not essential for predictions about misplaced confidence in the long run.

3.2 Analysis

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

$$\begin{aligned} \mu_{i1} &= \frac{n_{i0}}{n_{i0}+1}\mu_{i0} + \frac{1}{n_{i0}+1}s_{i0} \\ &= (1-w_{i0})\mu_{i0} + w_{i0}s_{i0}. \end{aligned} \tag{16}$$

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 form of (14) with $K_{it} < \infty$, that its variance function has the form of (15), and that

$$\mu_{i(t+1)} = (1-w_{it})\mu_{it} + w_{it}s_{it} \tag{17}$$

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 (16) and (17) is the critical feature of exponential families for the present analysis. In Sect. 2, this same linearity appeared in equation (5) because normal distributions are members of an exponential family.

The following proposition shows that average bias 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 (Average Misplaced Confidence in All Periods) *Let Assumptions 1 and 2 hold. For all $t > 1$, there exists $\psi > 0$ such that if $\delta + \lambda \leq \psi$, then*

i

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

ii

$$\bar{\mu}_t > \bar{z} \text{ if } \frac{\partial V_i(\mu_{i0}; e_{i0})}{\partial e_{i0}} > 0 \text{ and } \frac{\mathrm{d}\zeta_{i2}}{\mathrm{d}e_{i0}} \geq 0 \text{ for all } i.$$

Proof The proof is by induction. See Sect. 10.2. □

Average misplaced confidence emerges after period 1 and persists in later periods, vanishing only in the limit as agents accumulate infinite data.[25] 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 Sect. 2, the direction of average misplaced confidence depends on the effects of effort on the variance of the likelihood. From Lemma 2, 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 (under-confident), as already described following equation (10). We therefore predict average underconfidence (overconfidence) when additional effort decreases (increases) the variance of the likelihood.[26]

The variance of the time 0 likelihood is $\frac{1}{e_{i0}^2} \mathrm{Var}_{i0}[\nu_{i0}]$. This variance decreases in effort if either the variance of ν_{i0} is independent of effort or the variance of ν_{i0} decreases in effort. It 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 variance of ν_{i0} at just the right rate.[27] 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.

The following corollary relaxes Assumption 2.

[25] The true parameter z_i is within the support of the prior, so the agent's posterior will converge to a point mass on z_i. See, for instance, Diaconis and Freedman (1986).

[26] The requirement that μ_{i0} be within δ of z_i and that δ not be too large restricts the strength of the heterogeneous z_i term in (10).

[27] 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.

Corollary 2 (Average Misplaced Confidence in the Long Run) *Let Assumptions 1 and 2' hold. There exist $s \geq 0$ and $\psi > 0$ such that, for all $t \geq s$ and $\lambda \leq \psi$,*

i

$$\bar{\mu}_t < \bar{z} \text{ if } \frac{\partial V_i(\mu_{is}; e_{is})}{\partial e_{is}} < 0 \text{ and } \frac{\mathrm{d}\zeta_{i2}}{\mathrm{d}e_{is}} \leq 0 \text{ for all } i,$$

ii

$$\bar{\mu}_t > \bar{z} \text{ if } \frac{\partial V_i(\mu_{is}; e_{is})}{\partial e_{is}} > 0 \text{ and } \frac{\mathrm{d}\zeta_{i2}}{\mathrm{d}e_{is}} \geq 0 \text{ for all } i.$$

Proof For s sufficiently large, Assumption 2 holds with δ arbitrarily small. The result then follows from Proposition 1.

□

If we allow initial beliefs to be arbitrarily poorly calibrated, then we can no longer predict the direction of average misplaced confidence within two periods, but we can still predict that direction in the long run. It depends on how effort affects the variance function once beliefs become better calibrated. Intuitively, if initial beliefs display, for example, arbitrarily severe initial underconfidence on average, then this average underconfidence may persist for some time, but eventually agents learn away these initial mistakes and the learning dynamics described previously take over.[28]

4 Normally distributed shocks with truncated-normal priors

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 extend the model of Sect. 2 to obtain concrete predictions over infinitely many periods.

Let the reward π_{it} be as in equation (1), with $z_i > 0$ and 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. These statistics may be heterogeneous across agents. The updating rules are as in equations (2) and (3) from Sect. 2.[29] The rest of the setting is familiar from previous sections.

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

[28] As beliefs become better calibrated in the long run, the setting starts to look like that which led to equation (11).

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

distribution and converges to the mean of the truncated distribution as the mass beyond the truncation points vanishes. The following proposition describes the evolution of average overconfidence, here measured as $\bar{\phi}_t - \bar{z}$, under the assumption that the mean of each agent's initial nontruncated beliefs is approximately well-calibrated[30]:

Proposition 3 (Average Misplaced Confidence in All Periods (Normal Distributions)) *Let Assumption 2 hold. For all $t > 1$, there exists $\psi > 0$ such that if $\delta + \sigma \leq \psi$ and each Σ_{i0} is sufficiently small, then*

i

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

ii

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

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 1 and the proof of Proposition 1, since the family of normal distributions is NEF-QVF with $\zeta_{i2}(\cdot) = 0$. □

In Sect. 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 3 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 1 implies that the variance of the likelihood decreases in effort if and only if $\chi(e_{i0}) < 2$, so Proposition 3 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.

Finally, we can again relax Assumption 2.

Corollary 4 (Average Misplaced Confidence in the Long Run (Normal Distributions)) *Let Assumption 2' hold. There exist $s \geq 0$ and $\psi > 0$ such that, for all $t \geq s$ and $\sigma \leq \psi$,*

i

$$\bar{\phi}_t < \bar{z} \quad \textit{if} \quad \chi(e_{is}) < 2 \quad \textit{for all } i,$$

[30] ϕ_{it} approaches $E_{it}[z_i]$ under the conditions of the proposition, so $\bar{\phi}_t$ approaches $\bar{\mu}_t$ and the present measure of overconfidence is in practice similar to the measure used in previous sections.

ii

$$\bar{\phi}_t > \bar{z} \quad \textit{if} \quad \chi(e_{is}) > 2 \quad \textit{for all } i.$$

Proof For s sufficiently large, Assumption 2 holds and Σ_{is} is small. The result then follows from Proposition 3.

□

As in Sect. 3, allowing initial beliefs to be arbitrarily poorly calibrated costs us the ability to predict the direction of average misplaced confidence in the short run but does not stop us from predicting that direction in the long run.

5 Numerical example

I now consider a numerical example, adapting the setting of Sect. 2. Under the parameterization, the results are nearly identical to instead using the truncated-normal prior from Sect. 4 with $a_i = 0$ and $b_i = \infty$. For all i, 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 $\chi(e_{it}) = \alpha$. From (4), observe that $\partial w_{i0}/\partial e_{i0} > 0$ if and only if $\alpha < 2$. And in (10), $\mathrm{Var}[z_i] = 0$. I simulate one million agents.

The top left panel of Fig. 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.

The remaining panels of Fig. 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.[31] The distribution of μ_{it} becomes progressively narrower as data accumulate.

The lower left panel of Fig. 1 shows that $\bar{\mu}_1$ 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 $\mathrm{Cov}[\mu_{it}, w_{it}]$ (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

[31] 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 $A > 0$.

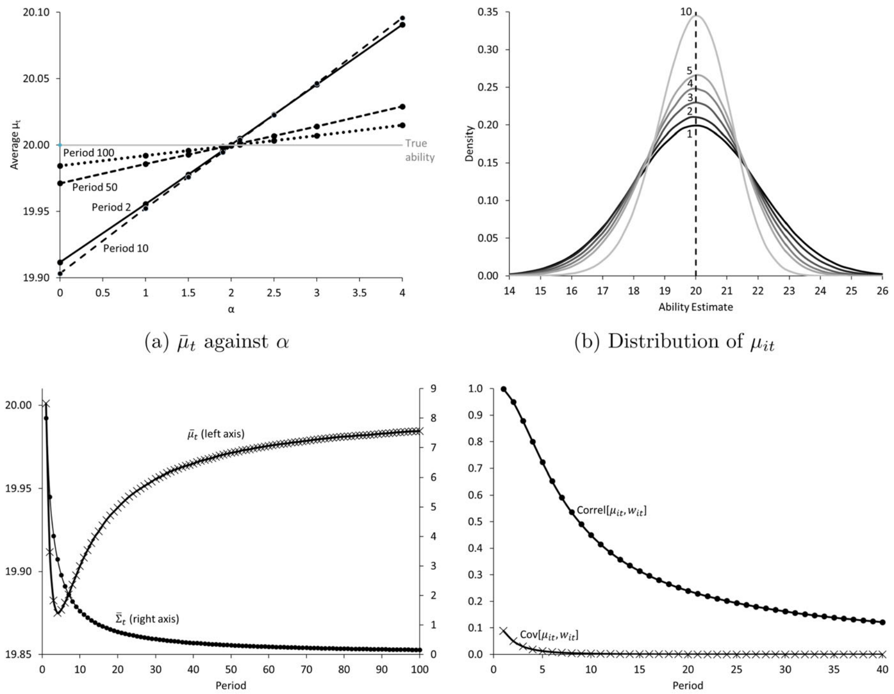

(a) $\bar{\mu}_t$ against α

(b) Distribution of μ_{it}

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

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

Fig. 1 The top left panel varies $\chi(\cdot) = \alpha$. In the top right panel, labels indicate the period t. 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}

positive even at long horizons) but because the variance of each variable declines strongly as agents become more certain of their ability.

6 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), especially when effort and ability are complementary. 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 low-effort performances than after high-effort 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.[32]

These types of interventions do not require the employer to know workers' true abilities. Instead, these interventions require the employer to have additional information about shocks that affect employees' outcomes and to commit to 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.

To formalize such stories, adopt the setting of Sect. 4, with Σ_{i0} small relative to $(b_i - a_i)^2$ and with $\mu_{i0} = z_i$ (i.e., $\delta = 0$ in Assumption 2). Agent i is an employee with cost function $c_i(e_{it}) = 0.5e_{it}^2$, so that the first-order condition implies $e_{it}^* = \mu_{it}$. The manager knows the employee's true ability z_i and observes the employee's effort. The manager chooses how the variance of the signal (i.e., of the feedback) that the employee receives changes with the employee's effort. In particular, parameterize $f(e_{it}) = e_{it}^{\alpha}$ (so that $\chi(e_{it}) \triangleq e_{it} f'(e_{it})/f(e_{it}) = \alpha$) and consider a manager choosing the parameter α. The manager pays a cost $C(\alpha)$ to deviate from some baseline feedback level $\bar{\alpha}$, where $C(\alpha) = 0.5(\alpha - \bar{\alpha})^2$.

The manager chooses α before period 0. After the manager chooses α, she announces it to the employee. The employee proceeds to choose time 0 effort, update as a Bayesian based on the time 0 signal and the known α, and so on through subsequent periods. The manager seeks to maximize expected period 2 effort, net of the costs of choosing α[33]:

$$\max_{\alpha} \left\{ E_0[e_{i2}^*(\mu_{i2})|z_i] - C(\alpha) \right\}.$$

where μ_{i2} follows (18) and expectations are over the sequences of shocks that the employee may receive. Taking a first-order approximation to the manager's objective around $\mu_{i2} = z_i$ (which is a good approximation for either Σ_{i0} or σ^2 small) and substituting $e_{it}^{*\prime} = 1$, the manager's problem becomes:

$$\max_{\alpha} \left\{ E_0[\mu_{i2} - z_i|z_i] - C(\alpha) \right\}.$$

[32] 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).

[33] Observe that the manager has no ability to affect expected effort in period 1 because the linear updating rule and symmetric shocks imply that $E_0[\mu_{i1}] = \mu_{i0}$, as in equation (6).

For Σ_{i0} small, equation (4) approximately holds and so does (11). The manager's problem becomes

$$\max_{\alpha} \left\{ -w_{i0}(1-w_{i0})\Sigma_{i0} E_0 \left[\frac{\partial w(e_{i1}, \Sigma_{i1})}{\partial e_{i1}} \frac{\mathrm{d} e_{i1}(\mu_{i1})}{\mathrm{d}\mu_{i1}} \middle| z_i \right] - C(\alpha) \right\},$$

which is equivalent to

$$\max_{\alpha} \left\{ -(2-\alpha) w_{i0}(1-w_{i0})\Sigma_{i0} E_0 \left[\frac{\Sigma_{i1}^{-1}}{\Sigma_{i1}^{-1} + e_{i1}^{2-\alpha}} \frac{e_{i1}^{1-\alpha}}{\Sigma_{i1}^{-1} + e_{i1}^{2-\alpha}} \middle| z_i \right] - C(\alpha) \right\}.$$

The first-order condition is:

$$\begin{aligned} \alpha =& \bar{\alpha} + w_{i0}(1-w_{i0})\Sigma_{i0} E \left[(1-w_{i1}) \frac{w_{i1}}{e_{i1}} \middle| z_i \right] \\ & - (2-\alpha) w_{i0}(1-w_{i0})\Sigma_{i0} E \left[(1-w_{i1}) \frac{w_{i1}}{e_{i1}} \ln e_{i1} [2w_{i1} - 1] \middle| z_i \right]. \end{aligned}$$

Let $\bar{\alpha} = 2$, so that a manager who chose not to spend any time or resources to modify feedback would leave the employee with $\partial w_{i1}/\partial e_{i1} = 0$ and thus retaining, in expectation, well-calibrated beliefs in period 2. Using the first-order condition, for $\alpha = 2$ to be the manager's optimal choice, it would have to be the case that

$$w_{i0}(1-w_{i0})\Sigma_{i0} E \left[(1-w_{i1}) \frac{w_{i1}}{e_{i1}} \middle| z_i \right] = 0.$$

But this condition clearly does not hold when $\Sigma_{i0}, \sigma^2 > 0$. So the manager does spend resources to choose α^* that will lead to expected misplaced confidence in period 2.

Will the manager choose to make the employee over- or underconfident in expectation? With $\bar{\alpha} = 2$, the manager's optimal α^* satisfies:

$$\alpha^* = 2 + \frac{w_{i0}(1-w_{i0})\Sigma_{i0} E \left[(1-w_{i1}) \frac{w_{i1}}{e_{i1}} \middle| z_i \right]}{1 - w_{i0}(1-w_{i0})\Sigma_{i0} E \left[(1-w_{i1}) \frac{w_{i1}}{e_{i1}} \ln e_{i1} [2w_{i1} - 1] \middle| z_i \right]}.$$

For μ_{i0} and z_i large and Σ_{i0} small, $\ln e_{i1}$ is strictly positive along most trajectories. And for Σ_{i0} small, $w_{i1} < 0.5$. In such a case, the denominator on the right-hand side is strictly positive and the manager chooses $\alpha^* > 2$. And per Proposition 3, the manager's choice $\alpha^* > 2$ induces the employee to become overconfident in expectation.[34]

This example shows that a manager who understands how employees update beliefs may choose to manipulate the variance of the feedback that an employee receives in order to increase the employee's expected effort. The employee understands that the manager is choosing α^* to induce overconfidence on average but can do no better

[34] In general, it is ambiguous whether $\alpha^* > 2$ or < 2. Future work should analyze the various factors determining how a manager should condition the variance of feedback on effort choices.

than to update as a Bayesian given that choice of α^*. Future work could explore more sophisticated strategies on the part of the manager. For instance, the manager may choose not to disclose some shocks, may use a different functional form for disclosure, or may vary the feedback rule over time. In addition, the present example considers a manager's ability to induce average overconfidence in a single employee, where averaging is over states of the world. If we instead consider a manager's ability to induce overconfidence among a population of employees, then we would need to consider the possibility that shocks are correlated across agents. Future work should explore implications of that correlation.

7 Relation to empirical work on overconfidence

The critical ingredients for rational overconfidence to emerge from well-calibrated beliefs are that ability and effort be complementary and that additional effort complicate 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 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), implying stock options may seek to induce additional effort. Stock options are more commonly granted to executives than to salaried workers and are least commonly granted to hourly workers (Hall and Murphy 2002). If we (loosely) take executives as being of higher ability, then one interpretation of remunerating through stock options is that principals believe it is especially important to induce effort from high-ability employees, as when effort and ability are 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 empirical 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. So one interpretation of investors moving online after experiencing high returns is that they want to apply more effort, suggesting 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, some previous empirical work may not fit the present model. In particular, Niederle and Vesterlund (2007) explore overconfidence in a laboratory environment. They show that subjects of both genders perform significantly better on an addition task when participating in a tournament instead of being paid piece-rate. One of their explanations is that this improved performance is a response to incentives. If this is the case, it seems likely that effort would be the mechanism, with effort responding to incentives and mattering for performance. 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.

8 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.[35]

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

[35] 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: if an event correlates with ability, then people who observe it conclude that they are better than average, and if that event is also the most likely outcome for most people, then most people rationally conclude that they are better than average. In that model, agents do not act on their beliefs about their abilities, which limits the observable implications of their beliefs, and signals arrive exogenously, which restricts their opportunities to learn their incorrect beliefs away. In the present model, overconfidence affects observable effort choices and thereby affects agents' payoffs and signals.

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.[36] 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.

The proposed mechanism is more closely related to two recent papers.[37] 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.[38]

Second, Hestermann and Le 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.[39]

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

36 A similar mechanism underpins the model of self-control in Ali (2011). Also, Deimen and Wirtz (2022) study a two-armed bandit model in which students who do not observe successes may stop choosing high effort and thereby become persistently underconfident.

37 The proposed mechanism also shares features with the macro model of uncertainty over the business cycle in 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. That mechanism is a special case of the mechanism studied here. Whereas asymmetric learning speeds there explain asymmetries over the course of a business cycle, I here consider a stationary setting and show how asymmetric learning speeds can generate and sustain biased mean beliefs.

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

39 Hestermann and Le 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.

permanently misspecified model of his own ability generate signals that do not lead the agent to question his incorrect beliefs about his own ability.[40] 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 converge to the truth, and misplaced confidence emerges endogenously rather than being imposed ex ante.[41,42]

9 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 several further types of investigations. First, empirical research has detected average overconfidence in several settings (see Sects. 1 and 7). 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.

40 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.

41 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 about the data generating process drive agents' asymmetric rates of learning from good and bad shocks.

42 Nielsen (2018) considers welfare evaluations in the presence of "rational overconfidence", which they define as a case in which agents' misspecified models of the world each match empirical averages but are, on average, overly precise. My version of rational overconfidence emerges dynamically from Bayesian updating and optimized effort choices, as opposed to being postulated directly.

10 Further analysis

10.1 Useful results for Sect. 3

Working backwards through time,

$$\begin{aligned}\mu_{it} - z_i =& \frac{w_{i(t-1)}}{e_{i(t-1)}} \nu_{i(t-1)} + (1 - w_{i(t-1)})(\mu_{i(t-1)} - z_i) \\ =& \lambda \sum_{j=1}^{t} \left(\prod_{k=1}^{j-1} (1 - w_{i(t-k)}) \right) \frac{w_{i(t-j)}}{e_{i(t-j)}} \nu_{i(t-j)} \\ &+ \left(\prod_{k=1}^{t-1} (1 - w_{i(t-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i] \end{aligned} \tag{18}$$

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

$$\begin{aligned}\int_0^1 [\mu_{it} - z_i] \, \mathrm{d}i =& \lambda \sum_{j=2}^{t} \int_0^1 \frac{w_{i(t-j)}}{e_{i(t-j)}} \operatorname{Cov}_{i(t-j)|z_i} \left[\prod_{k=1}^{j-1} (1 - w_{i(t-k)}), \nu_{i(t-j)} \right] \mathrm{d}i \\ &+ \int_0^1 \left(\prod_{k=1}^{t-1} (1 - w_{i(t-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i] \, \mathrm{d}i, \end{aligned} \tag{19}$$

where $\operatorname{Cov}_{i(t-j)|z_i}$ 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:

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

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 2 (Updating and Effort) *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 (16), 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.

10.2 Proof of Proposition 1

Because an NEF is characterized by its variance function (e.g., Morris 1982),[43] we can write $w_{it} \triangleq w(\mu_{it}, \tilde{V}_{it}(\mu_{it}))$. Fix e_{it} and consider two possible signals s^H and s^L received at times $t-1$ and t. From equation (2.10) of Diaconis and Ylvisaker (1979), $\mu_{i(t+1)}$ does not depend on the order in which the signals were received. Therefore

$$\begin{aligned}(1 - w_{it}^L)&(1 - w_{i(t-1)})\mu_{i(t-1)} + (1 - w_{it}^L)w_{i(t-1)}s^L + w_{it}^L s^H \\ &= (1 - w_{it}^H)(1 - w_{i(t-1)})\mu_{i(t-1)} + (1 - w_{it}^H)w_{i(t-1)}s^H + w_{it}^H s^L, \end{aligned} \tag{20}$$

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

[43] 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.

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 (20) 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

$$w_{it} = w_{i0} + \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} (\mu_{it} - \mu_{i0}) + R_{it}, \tag{21}$$

where R_{it} is a polynomial that is $\mathcal{O}(\mu_{it} - \mu_{i0})^2$ as $(\mu_{it} - \mu_{i0}) \to 0$ and thus, using Assumption 2, as $(\delta + \lambda) \to 0$. Observe that:

$$(\mu_{it} - \mu_{i0})^2 = [\mu_{it} - z_i]^2 + [\mu_{i0} - z_i]^2 - 2[\mu_{it} - z_i][\mu_{i0} - z_i].$$

Use (18):

$$\begin{aligned}
(\mu_{it} - \mu_{i0})^2 = {} & [\mu_{i0} - z_i]^2 \Bigg\{ 1 - 2 \left(\prod_{k=1}^{t-1} (1 - w_{i(t-k)}) \right) (1 - w_{i0}) \\
& + \left(\prod_{k=1}^{t-1} (1 - w_{i(t-k)}) \right)^2 (1 - w_{i0})^2 \Bigg\} \\
& + 2\lambda [\mu_{i0} - z_i] \Bigg\{ \left(\prod_{k=1}^{t-1} (1 - w_{i(t-k)}) \right) (1 - w_{i0}) - 1 \Bigg\} \\
& \sum_{j=1}^{t} \left(\prod_{k=1}^{j-1} (1 - w_{i(t-k)}) \right) \frac{w_{i(t-j)}}{e_{i(t-j)}} \nu_{i(t-j)} \\
& + \lambda^2 \Bigg\{ \sum_{j=1}^{t} \left(\prod_{k=1}^{j-1} (1 - w_{i(t-k)}) \right) \frac{w_{i(t-j)}}{e_{i(t-j)}} \nu_{i(t-j)} \Bigg\}^2 .
\end{aligned} \tag{22}$$

The third line is of order λ^2. Under Assumption 2, the first line is of order δ^2 and the second line is of order $\delta\lambda$. Then $(\mu_{it} - \mu_{i0})^2$ is $\mathcal{O}\big((\delta + \lambda)^2\big)$, and thus R_{it} is $\mathcal{O}\big((\delta + \lambda)^2\big)$ as $(\delta + \lambda) \to 0$.

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 (19),

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

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

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

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

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

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

Substitute for the top line from (19):

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

$$
\left[w_{iN} \prod_{k=2}^{j-1} (1 - w_{i(N+1-k)}),\, \nu_{i(N+1-j)} \right] \mathrm{d}i
$$

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

The induction hypothesis then implies

$$
\int_0^1 [\mu_{i(N+1)} - z_i]\, \mathrm{d}i < -\lambda \sum_{j=2}^{N+1} \int_0^1 \frac{w_{i(N+1-j)}}{e_{i(N+1-j)}} \mathrm{Cov}_{i(N+1-j)|z_i}
$$

$$
\left[w_{iN} \prod_{k=2}^{j-1} (1 - w_{i(N+1-k)}),\, \nu_{i(N+1-j)} \right] \mathrm{d}i
$$

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

Substitute for w_{iN} from (21) to obtain:

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

$$
< -\lambda \sum_{j=1}^{N} \int_0^1 \frac{\mathrm{d}w_{i0}}{\mathrm{d}e_{i0}} \frac{\mathrm{d}e_{i0}}{\mathrm{d}\mu_{i0}} \frac{w_{i(N-j)}}{e_{i(N-j)}} \mathrm{Cov}_{i(N-j)|z_i}
$$

$$
\left[(\mu_{iN} - z_i) \prod_{k=1}^{j-1} (1 - w_{i(N-k)}),\ \nu_{i(N-j)} \right] \mathrm{d}i
$$

$$
- \int_0^1 w_{iN} \left(\prod_{k=1}^{N-1} (1 - w_{i(N-k)}) \right) (1 - w_{i0})[\mu_{i0} - z_i]\, \mathrm{d}i
$$

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

Define

$$
X_{i(N+1)} \triangleq \sum_{j=1}^{N} \int_0^1 \frac{w_{i(N-j)}}{e_{i(N-j)}} \mathrm{Cov}_{i(N-j)|z_i} \left[R_{iN} \prod_{k=1}^{j-1} (1 - w_{i(N-k)}),\ \lambda \nu_{i(N-j)} \right] \mathrm{d}i.
$$

From (22) and Assumption 2, $X_{i(N+1)}$ is $\mathcal{O}\big((\delta + \lambda)^2\big)$ as $(\delta + \lambda) \to 0$.

Use the relationship between covariances and expectations, the fact that the ν_{it} are mean-zero, and the law of iterated expectations:

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

Substitute for the summation on the top line from (18):

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

Substitute for w_{iN} from (21):

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

The right-hand side is weakly negative. The second and third integrals on the left-hand side are each $\mathcal{O}\big((\delta+\lambda)^2\big)$ as $(\delta+\lambda) \to 0$. So for $(\delta+\lambda)$ sufficiently small,

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

We have proved the induction step.

Basis step for part (i):

From (17) and properties of ν_{i1},

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

Substitute for w_{i1} from (21) and use properties of ν_{i0}:

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

Use (16) and properties of ν_{i0}:

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

Assumption 2 implies

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

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

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

$\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0$ implies that $\hat{\delta} > 0$, and Lemma 2 shows that the conditions in part (i) imply that $\mathrm{d}w_{i0}/\mathrm{d}e_{i0} > 0$. The second integral on the left-hand side is $\mathcal{O}\big(\lambda(\delta+\lambda)^2\big)$ as $(\delta+\lambda) \to 0$. So for $(\delta+\lambda)$ sufficiently 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,

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

We have proved the basis step.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s00199-024-01618-0.

References

Ali, S.N.: Learning self-control. Q. J. Econ. **126**(2), 857–893 (2011)

Astebro, T., Herz, H., Nanda, R., Weber, R.A.: Seeking the roots of entrepreneurship: insights from behavioral economics. J. Econ. Perspect. **28**(3), 49–70 (2014)

Barber, B.M., Odean, T.: Boys will be boys: gender, overconfidence, and common stock investment. Q. J. Econ. **116**(1), 261–292 (2001)

Barber, Brad M., Odean, Terrance: Online investors: do the slow die first? Rev. Financ. Stud. **15**(2), 455–488 (2002)

Barndorff-Nielsen, O.: Exponential families. In: Kotz, S., Read, C.B., Balakrishnan, N., Vidakovic, B., Johnson, N.L. (eds.) Encyclopedia of Statistical Sciences. American Cancer Society, Atlanta (2006)

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

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

Bénabou, R., Tirole, J.: Self-confidence and personal motivation. Q. J. Econ. **117**(3), 871–915 (2002)

Benoît, J.-P., Dubra, J.: Apparent overconfidence. Econometrica **79**(5), 1591–1625 (2011)

Blavatskyy, P.R.: Betting on own knowledge: experimental test of overconfidence. J. Risk Uncertain. **38**(1), 39–49 (2009)

Brunnermeier, M.K., Parker, J.A.: Optimal expectations. Am. Econ. Rev. **95**(4), 1092–1118 (2005)

Burnside, C., Han, B., Hirshleifer, D., Wang, T.Y.: Investor overconfidence and the forward premium puzzle. Rev. Econ. Stud. **78**(2), 523–558 (2011)

Cappelli, P., Conyon, M.J.: What do performance appraisals do? ILR Rev. **71**(1), 88–116 (2018)

Chen, S., Schildberg-Hörisch, H.: Looking at the bright side: the motivational value of confidence. Eur. Econ. Rev. **120**, 103302 (2019)

Clark, J., Friesen, L.: Overconfidence in forecasts of own performance: an experimental study. Econ. J. **119**(534), 229–251 (2009)

Compte, O., Postlewaite, A.: Confidence-enhanced performance. Am. Econ. Rev. **94**(5), 1536–1557 (2004)

Consonni, G., Veronese, P.: Conjugate priors for exponential families having quadratic variance functions. J. Am. Stat. Assoc. **87**(420), 1123–1127 (1992)

Daniel, K., Hirshleifer, D.: Overconfident investors, predictable returns, and excessive trading. J. Econ. Perspect. **29**(4), 61–88 (2015)

Daniel, K., Hirshleifer, D., Subrahmanyam, A.: Investor psychology and security market under- and overreactions. J. Financ. **53**(6), 1839–1885 (1998)

de Quidt, J., Haushofer, J.: Depression through the lens of economics: a research agenda. In: Christopher, B. B., Michael, R. C., and Chavas, J.-P. (Eds.), The Economics of Poverty Traps (2017)

Deimen, I., Wirtz, J.: Control, cost, and confidence: perseverance and procrastination in the face of failure. Games Econom. Behav. **134**, 52–74 (2022)

Diaconis, P., Freedman, D.: On the consistency of Bayes estimates, Annals of Statistics, Vol. 14, No. 1, pp 1–26
Diaconis, P., Ylvisaker, D.: Conjugate priors for exponential families. Ann. Stat. **7**(2), 269–281 (1979)
Fudenberg, D., Romanyuk, G., Strack, P.: Active learning with a misspecified prior. Theor. Econ. **12**(3), 1155–1189 (2017)
Gervais, S., Goldstein, I.: The positive effects of biased self-perceptions in firms. Rev. Financ. **11**(3), 453–496 (2007)
Gervais, S., Odean, T.: Learning to be overconfident. Rev. Financ. Stud. **14**(1), 1–27 (2001)
Gibbs, M. J.: An economic approach to process in pay and performance appraisals, working paper (1991)
Gneezy, U., List, J.A., Livingston, J.A., Qin, X., Sadoff, S., Xu, Y.: Measuring success in education: the role of effort on the test itself. Am. Econ. Rev. Insights **1**(3), 291–308 (2019)
Graham, J.R., Harvey, C.R., Huang, H.: Investor competence, trading frequency, and home bias. Manag. Sci. **55**(7), 1094–1106 (2009)
Groes, F., Kircher, P., Manovskii, I.: The U-shapes of occupational mobility. Rev. Econ. Stud. **82**(2), 659–692 (2015)
Gutiérrez-Peña, E., Smith, A.F.M.: Exponential and Bayesian conjugate families: review and extensions. TEST **6**(1), 1–90 (1997)
Hall, B.J., Liebman, J.B.: Are CEOs really paid like bureaucrats? Q. J. Econ. **113**(3), 653–691 (1998)
Hall, B.J., Murphy, K.J.: Stock options for undiversified executives. J. Account. Econ. **33**(1), 3–42 (2002)
Heidhues, P., Köszegi, B., Strack, P.: Unrealistic expectations and misguided learning. Econometrica **86**(4), 1159–1214 (2018)
Hestermann, N., Le Yaouanq, Y.: Experimentation with self-serving attribution biases. Am. Econ. J. Microecon. **13**(3), 198–237 (2021)
Hirshleifer, D.: Investor psychology and asset pricing. J. Financ. **56**(4), 1533–1597 (2001)
Hirshleifer, D.: Behavioral finance. Annu. Rev. Financ. Econ. **7**(1), 133–159 (2015)
Hoffman, M., Burks, S.V.: Worker overconfidence: field evidence and implications for employee turnover and firm profits. Quant. Econ. **11**(1), 315–348 (2020)
Jehiel, P.: Investment strategy and selection bias: an equilibrium perspective on overoptimism. Am. Econ. Rev. **108**(6), 1582–97 (2018)
Kirchler, E., Maciejovsky, B.: Simultaneous over- and underconfidence: evidence from experimental asset markets. J. Risk Uncertain. **25**(1), 65–85 (2002)
Köszegi, B.: Ego utility, overconfidence, and task choice. J. Eur. Econ. Assoc. **4**(4), 673–707 (2006)
Krähmer, D.: Equilibrium learning in simple contests. Games Econom. Behav. **59**(1), 105–131 (2007)
Larkin, I., Leider, S.: Incentive schemes, sorting, and behavioral biases of employees: experimental evidence. Am. Econ. J. Microecon. **4**(2), 184–214 (2012)
Layard, R., Clark, D.M.: Thrive: How Better Mental Health Care Transforms Lives and Saves Money. Princeton University Press, Princeton, New Jersey (2015)
Lybbert, T.J., Wydick, B.: Poverty, aspirations, and the economics of hope. Econ. Dev. Cult. Change **66**(4), 709–753 (2018)
Malmendier, U., Tate, G.: CEO overconfidence and corporate investment. J. Financ. **60**(6), 2661–2700 (2005)
Malmendier, U., Tate, G.: Who makes acquisitions? CEO overconfidence and the market's reaction. J. Financ. Econ. **89**(1), 20–43 (2008)
Malmendier, U., Tate, G.: Behavioral CEOs: the role of managerial overconfidence. J. Econ. Perspect. **29**(4), 37–60 (2015)
Moore, D.A., Healy, P.J.: The trouble with overconfidence. Psychol. Rev. **115**(2), 502–517 (2008)
Morris, C.N.: Natural exponential families with quadratic variance functions. Ann. Stat. **10**(1), 65–80 (1982)
Morris, C.N.: Natural exponential families with quadratic variance functions: statistical theory. Ann. Stat. **11**(2), 515–529 (1983)
Morris, C.N.: Natural exponential families. In: Kotz, S., Read, C.B., Balakrishnan, N., Vidakovic, B., Johnson, N.L. (eds.) Encyclopedia of Statistical Sciences. American Cancer Society, Atlanta (2006)
Morris, C.N., Lock, K.F.: Unifying the named natural exponential families and their relatives. Am. Stat. **63**(3), 247–253 (2009)
Mueller, C.M., Dweck, C.S.: Praise for intelligence can undermine children's motivation and performance. J. Pers. Soc. Psychol. **75**, 33–52 (1998)
Murad, Z., Sefton, M., Starmer, C.: How do risk attitudes affect measured confidence? J. Risk Uncertain. **52**(1), 21–46 (2016)

Murphy, K.J.: Executive compensation. In: Ashenfelter, O.C., Card, D. (eds.) Handbook of Labor Economics, vol. 3, pp. 2485–2563. Elsevier, Berlin (1999)

Nicholls, J.G., Miller, A.T.: Development and its discontents: the differentiation of the concept of utility. In: Nicholls, J.G. (ed.) Develop. Achiev. Motiv., pp. 185–218. JAI Press, Greenwich, CT (1984)

Niederle, M., Vesterlund, L.: Do women shy away from competition? Do men compete too much? Q. J. Econ. **122**(3), 1067–1101 (2007)

Nielsen, C.K.: Rational overconfidence and social security: subjective beliefs, objective welfare. Econ. Theor. **65**(2), 179–229 (2018)

Papageorgiou, T.: Learning your comparative advantages. Rev. Econ. Stud. **81**(3), 1263–1295 (2014)

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

Silva, F.: Inducing overconfidence. Econ. Inq. **55**(1), 451–460 (2017)

van den Steen, E.: Rational overoptimism (and other biases). Am. Econ. Rev. **94**(4), 1141–1151 (2004)

Urbig, D., Stauf, J., Weitzel, U.: What is your level of overconfidence? A strictly incentive compatible measurement of absolute and relative overconfidence, Working Paper 09-20, Utrecht School of Economics (2009)

Van Nieuwerburgh, S., Veldkamp, L.: Learning asymmetries in real business cycles. J. Monet. Econ. **53**(4), 753–772 (2006)

Zábojník, J.: A model of rational bias in self-assessments. Econ. Theor. **23**(2), 259–282 (2004)