Rationally Misplaced Confidence*

Derek Lemoine

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

University of Arizona Working Paper 18-05

March 2018

I show that persistent underconfidence and overconfidence each arise from rational learning about one's own abilities. If an agent chooses to exert more effort when more confident and believes that greater effort reliably improves outcomes, then the agent learns away overconfidence faster than he learns away underconfidence. The agent ends up underconfident on average. In contrast, the agent ends up overconfident on average if he believes that greater effort increases his exposure to chance. The results are consistent with evidence that underconfidence and overconfidence are both widespread in the population, and the mechanism is consistent with the modern understanding of depression.

JEL: D83, D90, G41

Keywords: learning, effort, confidence, ability, talent, explanatory style, depression, therapy

*I thank David Hirshleifer for helpful comments.

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? Many people, like Emerson's "strong men", believe that their efforts will have predictable consequences, but many others, like Emerson's "shallow men", are quick to attribute the outcomes of their efforts to chance. People in the first group learn a lot about themselves by observing the fruits of their efforts, whereas people in the second group do not infer as much from these outcomes. We will here see a surprising result: rational Bayesian agents systematically misjudge their own ability, and whether they end up overconfident or underconfident depends on their beliefs about the predictability of their efforts' consequences.

Overconfidence is now generally recognized as an important factor in many markets. For instance, overconfidence among traders can explain financial market anomalies (Daniel and Hirshleifer, 2015), overconfidence can explain why some people persist as entrepreneurs (Astebro et al., 2014), and overconfidence among business executives can affect corporate investment decisions (Malmendier and Taylor, 2015). Economists have therefore sought to model the psychological biases and motivations that can enable overconfidence to persist even in the face of contrary data. However, recent experimental evidence suggests that underconfidence is more prevalent than overconfidence (e.g., Kirchler and Maciejovsky, 2002; Blavatskyy, 2009; Clark and Friesen, 2009; Urbig et al., 2009; Murad et al., 2016). And underconfidence may be more important for wellbeing. The welfare cost of depression, often considered a product of underconfidence, is enormous. Depression afflicts around one in six adults in developed countries and contributes more to developed countries' disease burden than does any single class of physical illness (Layard and Clark, 2015). In the U.S. alone, major depressive disorders affect 15.4 million adults and cost \$210 billion per year (Greenberg et al., 2015). Given the pervasiveness and costs of underconfidence, it is important that we also understand its origins and the prospects for treatment.

I here propose a unified model in which persistent overconfidence and persistent underconfidence both emerge from Bayesian updating by rational agents who have neoclassical utility functions, do not exhibit behavioral biases, and initially hold correct expectations of their own ability.[1] An agent chooses his effort level in each period to maximize his expected reward. The reward depends on both his effort level and his unknown ability and

[1] As we will see, previous literature deviates in one or more of these three dimensions.

is subject to unobserved shocks. Because ability and effort are complementary, the agent applies more effort when he thinks he is of high ability. Upon observing the reward, the agent updates his beliefs about his ability. The agent's effort choices affect the strength of the signal provided by the observed reward. I show that whether the agent ends up overconfident or underconfident on average depends on his understanding about how his effort level affects the variance of his reward. He ends up overconfident on average if he thinks that the variance increases sufficiently strongly in his effort, and he ends up underconfident on average otherwise. Only in a knife-edge case will we expect an agent to display neither overconfidence nor underconfidence.

Begin by considering an agent who believes that the variance of his rewards is completely due to temporary shocks to his ability.[2] Greater effort amplifies the marginal effect of his durable ability but also amplifies the marginal effect of temporary shocks to his ability. Greater effort therefore does not affect the signal of his true ability provided by the observed reward. We are interested in the distribution of the agent's posterior beliefs at some future time t. I show that if the agent initially estimates his own ability correctly (i.e., if his prior is centered around the true value), then his posterior at any future time t is, on average, centered around his true ability as well. For the agent to be underconfident or overconfident on average, he would have to start with an incorrect estimate of his own ability. This special case is consistent with standard intuition.

Now consider an agent who believes that the variance of his rewards is completely due to external shocks that are independent of his effort choices. Like Emerson's "strong man", he believes that his efforts have a consistent effect on outcomes. For instance, he believes that running harder should improve his time by a consistent amount. The signal contained in his observed reward now depends on his chosen effort level. When he chooses high effort, he believes that he increases the marginal effect of his true ability on the observed reward without changing the variance of the reward. As a result, the observed reward contains a stronger signal of his true ability and he strongly adjusts his posterior beliefs upon observing the reward.

Imagine that the agent has a prior centered around his true ability at time 0. Also imagine that the unobserved shock happens to take on a high value at time 0, so that the agent perceives a surprisingly high reward at time 0. As a result, he raises his central estimate of his ability and chooses greater effort at time 1. Because he is now overconfident, his time 1 reward will, on average, be surprisingly small and will lead him to reduce his time 2 ability estimate towards the true value. Following the average time 1 reward, the agent will still be overconfident at time 2 but less so than at time 1. Indeed, because his high time 1 effort made his beliefs especially sensitive to the observed time 1 reward, he will tend to be only slightly overconfident by time 2.

Now imagine that the unobserved shock happens to take on a low value at time 0. In

[2] We will see that what matters is not the actual variance of the rewards but the agent's beliefs about this variance. The results do not depend on whether the agent has correct beliefs about the variance.

this case, the agent reduces his central estimate of his ability and chooses lower effort at time 1. Because he is now underconfident, his time 1 reward will, on average, be surprisingly large and will lead him to raise his time 2 ability estimate towards the true value. Following the average time 1 reward, the agent will still be underconfident at time 2 but less so than at time 1. But because his low time 1 effort made his beliefs especially insensitive to the observed time 1 reward, his underconfidence may still be nearly as bad at time 2 as it was at time 1.

If we average across these two possibilities, the agent does not display any bias at time 1 because he adjusts his beliefs symmetrically in response to high or low time 0 shocks. However, the agent is underconfident on average at time 2: his average central estimate is below his true ability. The critical mechanism is that the agent's posterior beliefs are more sensitive to the observed reward when his effort is high. The agent learns away time 0 shocks especially quickly when they lead him to raise his central estimate of his own ability, and he learns away time 0 shocks especially slowly when they lead him to lower his central estimate of his own ability. The Bayesian agent becomes underconfident on average despite his initially correct beliefs about his own ability. In fact, I show that he will remain underconfident on average in any future period, approaching correct beliefs only asymptotically as he accumulates infinite data.[3]

Rational updating can also endogenously generate overconfidence. Now let the agent believe that greater effort increases the variance of his reward. Like Emerson's "shallow man", he believes that his efforts are largely modulated by circumstance. For example, he believes that running harder reduces the consistency of his time by increasing the consequences of each day's minor variations in weather, fitness, or diet. This agent's beliefs will be especially sensitive to news following low effort choices. Because he chooses low effort when he lacks confidence in his own ability, he learns away overly low ability estimates especially rapidly. And he learns away overly high ability estimates only slowly because he believes that his high efforts lead to especially noisy outcomes. When this agent has an overly high ability estimate, he will tend to receive bad news, but he attributes his failures more to chance than to his own ability. This agent becomes overconfident on average and will remain so in all future periods.

The proposed mechanism for generating underconfidence is consistent with modern theories of depression. As described in Section 6, many psychologists have concluded that depression is about processing information in a way that promotes a negative self-image. As a result, the agent withdraws from the world and becomes trapped in a state of inactivity.

[3]The agent does not expect to become biased. The bias arises in the eyes of an outside observer who knows the agent's true ability. The outside observer could be a manager, a teacher, an experimenter, or a social planner, depending on the setting. The agent himself forms expectations by accounting not only for the possible sequences of unobserved shocks but also for the possible values of his true ability. As a good Bayesian, the agent does not expect his central estimate to drift over time. We will formally demonstrate this difference between the agent's expectations and the outside observer's expectations in Section 3.

We here see precisely such a mechanism leading to persistent underconfidence. When agents interact with the world (i.e., apply high effort) and believe that this interaction has regular consequences, they are quick to ascribe negative outcomes to their own inadequate ability rather than to chance. As their ability estimates fall, they choose lower effort levels and thus have a hard time learning away these incorrect beliefs. Being trapped in a state of persistent self-doubt and low effort is due not merely to having received bad news but to how the agent processes such news.

Cognitive behavioral therapy approaches to treating depression aim to help agents to re-engage with the world (change their behavior) and learn to process information differently (change cognition). These approaches have been shown to not only treat an instance of depression but to prevent a relapse into depression.[4] We here see the benefit of cognitive behavioral therapy's twin goals. An agent can be trapped in an overly pessimistic state when extreme negative shocks reduce the agent's estimate of his own ability and thereby reduce his effort. If the agent believes that he receives a worse signal of his own ability when he chooses low effort, then this reduction in effort will delay him from learning that his ability is greater than he thinks it is. Making the agent re-engage with the world through higher effort choices will tend to produce payoffs that raise his estimate of his own ability. However, in order to prevent the agent's confidence from trending down again, it is critical to also change the agent's information processing: he needs to stop believing that greater effort should increase his control over outcomes. Interventions that help an agent to re-engage with the world can alleviate a case of depression, but only techniques that also change his style of information processing can make a relapse less likely.

The proposed mechanism is, to my knowledge, novel in the literature on overconfidence.[5]

[4] This is not to suggest that there is no neurological component to depression. The recommended approaches combine cognitive behavioral therapy with pharmacological treatments; however, randomized controlled trials have found cognitive behavioral therapy to be more effective than pharmacological treatments at preventing relapse (e.g., Seligman, 1991; Scott, 1996; Hollon et al., 2005; Butler et al., 2006; Dobson et al., 2008; Layard and Clark, 2015). Further, there is evidence that cognitive behavioral therapy reduces the hyperactivity in depressives' amygdala while also stimulating their underactive prefrontal cortexes, whereas pharmacological treatments primarily help only the amygdala (DeRubeis et al., 2008; Layard and Clark, 2015).

[5] We here focus on overconfidence in the sense of what Moore and Healy (2008) call "overestimation." A literature in finance has focused on 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). Much other 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. Many authors have described selection mechanisms that can make a population of Bayesian updaters demonstrate overplacement (van den Steen, 2004; Zábojník, 2004; Köszegi, 2006; Jehiel, 2016): the common ingredient is that actors choose to stop collecting information once they receive a sufficiently positive signal about themselves or about the payoffs to some activity, so that high confidence is an absorbing state that attracts an ever greater share of the population. Benoît and Dubra (2011) show how Bayesian updating can lead too many agents to think that they are above average when likely events reinforce beliefs in one's own quality, and Santos-Pinto and

A first set of papers describes agents' motivations to choose to become overconfident, whether because optimism increases utility (Brunnermeier and Parker, 2005), because confidence directly improves outcomes (Compte and Postlewaite, 2004), or because confidence helps to overcome the tendency to procrastinate (Bénabou and Tirole, 2002). In contrast, the present setting is neoclassical: the agent's payoffs are maximized when the agent has a correct estimate of his own ability. A second set of papers generates overconfidence by assuming that individuals use a biased updating process (called "biased self-attribution") that overly attributes successes to their own ability and failures to chance (e.g., Daniel et al., 1998; Gervais and Odean, 2001).[6] Many believe (e.g., Hirshleifer, 2001, 2015) that such biases are necessary for a neoclassical model to generate persistent overconfidence, but the present setting generates persistent overconfidence (and also persistent underconfidence) as a result of rational Bayesian learning.

The proposed mechanism is more closely related to two recent papers. First, Silva (2017) also demonstrates how the asymmetric speed at which agents learn following good and bad shocks can generate systematic overconfidence. However, there the critical asymmetry is exogenously imposed: the agent is assumed to receive outside help following an early signal that he is of high quality but not after an early signal that he is of low quality, and this unobserved outside help subsequently prevents him from learning about his own ability. In contrast, the present paper's asymmetric speeds of learning emerge endogenously from the interaction between agents' effort choices and their statistical models of the world. Second, Hestermann and Yaouanq (2016) study an agent who is uncertain about his own fixed ability and also about some feature of the environment. 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 quite a while. We here see how persistent overconfidence and underconfidence can emerge even when the agent's initial beliefs are well-calibrated and even when the agent correctly understands his environment.[7]

The proposed model is also consistent with recent research in both psychology and man-

Sobel (2005) model heterogeneity in the mapping from a set of skills to an ability index.

[6]Though not explicitly about overconfidence, the model of confirmatory bias in Rabin and Schrag (1999) has a similar flavor.

[7]The present setting allows for the possibility of a misspecified statistical model in order to clarify that overconfidence and underconfidence depend not on the actual data generating process but on the agent's beliefs about that process. Misspecification does not drive the results. Heidhues et al. (2017) assume both that an agent is overconfident and that he does not realize that he might be overconfident. They study when the actions chosen under this permanently misspecified model of his own ability can generate signals that would not lead the agent to question his incorrect beliefs about his own ability. The present work is similar in allowing for both overconfidence and misspecification, but here overconfidence emerges endogenously, the misspecification affects only beliefs about the source of variance rather than ruling out correct beliefs about ability, and misspecification is not necessary for the results. Fudenberg et al. (2017) also consider the interaction between learning and a form of misspecification that places probability zero on the truth. We here study a less extreme form of misspecification that is not critical to the setting.

agement that has emphasized the role of beliefs about one's own self-efficacy in determining performance. Stronger beliefs about self-efficacy have been shown to increase an agent's chosen level of effort and the persistence of an agent's effort in the face of bad shocks (Bandura, 1982; Wood and Bandura, 1989; Stajkovic and Luthans, 1998; Bandura, 2001; Tenney et al., 2015).[8] We will see both effects here. Further, this same literature's discussion of learning also matches our setting: beliefs about one's own ability are thought to adjust not to absolute outcomes but to outcomes as filtered through previous beliefs about ability. We will formally model this updating process and highlight the previously overlooked importance of beliefs about how effort choices interact with randomness in the world.

The next section describes the setting. Section 3 analyzes a two-period example. Section 4 shows that the agent's posterior estimates drift away from the correct estimate on average. Section 5 provides a numerical example. Section 6 connects the analysis to the modern understanding of depression. Section 7 extends the analysis to the case of forward-looking agents. The final section concludes. The appendix contains proofs.

2 Setting

In every period t, an agent chooses how much effort e_t to apply to an activity. His cost of applying effort is $c(e_t)$, with $c(\cdot) \in C^2$, $c'(\cdot) > 0$, and $c''(\cdot) > 0$. The activity provides reward π_t, which depends on the chosen level of effort, on the agent's ability z, and on a random shock ϵ_t:

$$\pi_t = e_t z + \sqrt{f(e_t)}\, \epsilon_t. \tag{1}$$

Effort and ability are complementary.[9] The shock is mean-zero, normally distributed, and serially uncorrelated, with variance σ^2. $f(\cdot) > 0$ determines the degree to which the noise depends on the agent's effort. Conditional on z, the variance of the agent's time t payoff is $f(e_t)\, \sigma^2$.

The agent does not know his own ability z. His beliefs about his ability are summarized by a normal distribution with mean μ_t and variance Σ_t. The agent believes that rewards π_t are generated as

$$\pi_t = e_t z + \sqrt{g(e_t)}\, \epsilon_t, \tag{2}$$

with $g(\cdot) \in C^1$. The function $g(\cdot) > 0$ determines the degree to which the agent believes that the variance of his rewards is driven by shocks to his ability. When $g(\cdot) = f(\cdot)$, the agent correctly understands the data generating process, but when $g(\cdot) \neq f(\cdot)$, the agent's statistical model is misspecified. When $g'(\cdot) > 0$, the agent believes that trying harder amplifies risk. For example, running harder can increase the chance of an especially fast

[8] Tenney et al. (2015) find that beliefs about self-efficacy do not improve performance directly but instead improve performance through an increase in effort. This mechanism is consistent with the present setting but not with models such as Compte and Postlewaite (2004).

[9] See Bénabou and Tirole (2002) for extensive motivation of complementarity between effort and ability.

time and also, via premature exhaustion, the chance of an especially slow time; studying for a test through the night can increase the variance of one's score through the chance of decreased alertness; and not trying on a test may generate a low grade regardless of ability whereas trying hard exposes one to the variance of the test's quality. In contrast, when $g'(\cdot) < 0$, the agent believes that trying harder gives him more control over outcomes. For example, running harder may produce a more even pace; studying longer for a test may reduce the chance of mistakes; and trying harder on a multiple-choice test may ensure a grade that faithfully reflects ability whereas not trying may produce pure randomness.

The agent chooses e_t to maximize his expected per-period payoffs:[10]

$$\max_{e_t} \hat{E}_t \left[\pi_t - c(e_t)\right],$$

where $\hat{E}_t$ indicates the agent's expectations at his time t information set. The agent's optimal choice of effort e_t^* satisfies the first-order necessary condition:

$$c'(e_t^*) = \hat{E}_t \left[z + \frac{1}{2} \frac{g'(e_t^*)}{\sqrt{g(e_t^*)}} \epsilon_t \right],$$

which implies that

$$c'(e_t^*) = \mu_t. \tag{3}$$

Optimal effort e_t^* is an increasing function of μ_t. Throughout, we assume that z and μ_0 are much greater than zero, so that we need only consider strictly positive effort choices, and we omit the asterisk on e_t when clear.

Upon observing the payoffs π_t, the agent updates his beliefs about his ability z. The agent is a Bayesian learner,[11] so that the mean and variance of his beliefs evolve as

$$\mu_{t+1} = \left(\Sigma_t^{-1}\mu_t + \frac{e_t}{g(e_t)}\sigma^{-2}\pi_t\right)\left(\Sigma_t^{-1} + \frac{e_t^2}{g(e_t)}\sigma^{-2}\right)^{-1}, \tag{4}$$

$$\Sigma_{t+1} = \left(\Sigma_t^{-1} + \frac{e_t^2}{g(e_t)}\sigma^{-2}\right)^{-1}. \tag{5}$$

[10] In Section 7, I analyze a forward-looking agent who accounts for the informational value of his effort choices.

[11] The assumption that the agent chooses effort myopically but learns optimally is consistent with two prominent frameworks. First, "planner-doer" models of the self posit an internal conflict between a far-sighted planner and a myopic doer who actually makes day-to-day decisions (e.g., Thaler and Shefrin, 1981; Fudenberg and Levine, 2006). Second, much work in macroeconomics has used "anticipated utility" frameworks, in which the decision-maker updates beliefs from period to period as a Bayesian but formulates policies as if current beliefs will never change (see Kreps, 1998). The decision-maker learns only passively, without considering the informational value of his potential actions. In Section 7, we extend the analysis to an agent who learns actively.

Define

$$w(e_t, \Sigma_t) \triangleq \frac{\frac{e_t^2}{g(e_t)}\sigma^{-2}}{\Sigma_t^{-1} + \frac{e_t^2}{g(e_t)}\sigma^{-2}} \in [0, 1] \tag{6}$$

as the weight that the time t agent places on the signal extracted from the observed reward π_t when updating his beliefs, with $1 - w(e_t, \Sigma_t)$ the weight placed on the prior μ_t. Writing w_t for short, equation (4) becomes:

$$\mu_{t+1} = (1 - w_t)\mu_t + w_t \frac{\pi_t}{e_t}. \tag{7}$$

The signal is π_t/e_t, not π_t: the agent knows his choice of e_t and adjusts the observed reward π_t for this choice when constructing the signal of z.

We will be interested in the evolution of the agent's estimate of his own ability under the true data generating process (1) when the agent has unbiased beliefs at time 0. Formally, we assume that

$$\mu_0 = z,$$

and we study $E_0[\mu_t]$, where E_0 indicates expectations under the true data generating process at the time 0 information set. $E_0[\mu_t]$ averages over the possible sequences $\{\epsilon_s\}_{s=0}^{t-1}$ with z given.[12]

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. (2017) show that incentivizing students to exert more effort on a standardized test does improve test scores (effort matters for outcomes and responds to incentives) and improves test scores most strongly for higher-ability students (effort is complementary to ability).[13] 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, and we consider whether their beliefs will, on average, accurately reflect their abilities.

3 Two-Period Analysis

We begin with an analysis of the evolution of beliefs over the first two periods before turning to the full analysis.

[12] The agent's expectation operator $\hat{E}$ averages over both the sequences of ϵ and the possible values of z. Therefore $E_0[\mu_t] = \hat{E}_0[\mu_t | z = \mu_0]$.

[13] Levitt et al. (2016) also show that incentivizing effort improves scores on standardized tests.

From equation (7), we have:

$$\begin{aligned}E_0[\mu_1] =&[1 - w(e_0, \Sigma_0)]\mu_0 + w(e_0, \Sigma_0)z + w(e_0, \Sigma_0)E_0\left[\frac{\sqrt{f(e_0)}}{e_0}\epsilon_0\right]\\ =&[1 - w(e_0, \Sigma_0)]\mu_0 + w(e_0, \Sigma_0)z.\end{aligned}$$

Using $\mu_0 = z$, we have:

$$E_0[\mu_1] =z.$$

On average, the agent's beliefs remain properly calibrated at time 1. Further, μ_1 is normally distributed, with variance $\sigma^2\, w_0^2\, g(e_0)/e_0^2$.

Now consider the agent's central estimate at time 2:

$$E_0[\mu_2] =E_0\left[[1 - w(e_1, \Sigma_1)]\mu_1 + w(e_1, \Sigma_1)z + w(e_1, \Sigma_1)\frac{\sqrt{f(e_1)}}{e_1}\epsilon_1\right].$$

e_1 is a random variable because it depends on μ_1.[14] Using $E_0[\mu_1] = z$, we have:

$$\begin{aligned}E_0[\mu_2] =&E_0\left[1 - w(e_1, \Sigma_1)\right] E_0[\mu_1] + E_0[w(e_1, \Sigma_1)]z - Cov_0\left[w(e_1, \Sigma_1), \mu_1\right]\\ =&z - Cov_0\left[w(e_1, \Sigma_1), \mu_1\right].\end{aligned}$$

The agent's central estimate tends to drift away from z unless the covariance is zero. From Stein's Lemma,

$$Cov_0\left[w(e_1, \Sigma_1), \mu_1\right] =\sigma^2\, w_0^2\, \frac{g(e_0)}{e_0^2} E_0\left[\frac{\partial w(e_1, \Sigma_1)}{\partial e_1}\frac{\mathrm{d}e_1}{\mathrm{d}\mu_1}\right].$$

Recall that $\mathrm{d}e_1/\mathrm{d}\mu_1 > 0$. Therefore, the agent's beliefs tend to drift away from z unless $\partial w(e_1, \Sigma_1)/\partial e_1 = 0$, so that additional effort does not affect the agent's ability to learn from the observed reward π_1. If $\partial w(e_1, \Sigma_1)/\partial e_1 > 0$, then the covariance is strictly positive and the period 2 agent will, on average, underestimate his own ability. If the agent happens to receive a positive shock in period 0, then he becomes overconfident and chooses greater effort in period 1. At that time, he will tends to receive shocks that correct his overconfidence (pushing his beliefs back towards z). If $\partial w(e_1, \Sigma_1)/\partial e_1 > 0$, then he learns especially rapidly from these period 1 shocks and so on average enters period 2 only mildly overconfident, with an estimate close to z. However, if the agent happens to receive a negative shock in period 0, then the agent becomes underconfident and chooses a low effort level in period 1. If $\partial w(e_1, \Sigma_1)/\partial e_1 > 0$, then the agent does not learn much from the period 1 reward π_1.

[14] Σ_1 is not random because, from equation (5), it depends only on Σ_0 and e_0. The analysis in Section 4 will account for Σ_t being random for $t > 1$.

That reward will tend to correct his underconfidence (pushing his estimate back towards z), but the agent's beliefs are not sensitive to this reward. The agent will tend to enter period 2 with an estimate that is only a bit improved from his period 1 estimate, remaining nearly as underconfident as he was in period 1. Averaging across these two cases, the period 2 agent is underconfident on average because the link between effort and information processing led him to learn away overconfident period 1 beliefs faster than he learned away underconfident period 1 beliefs.

But does the possibility that $E_0[\mu_2] \neq \mu_0$ violate the Bayesian precept that the agent cannot expect to revise his beliefs in a particular direction? The analysis thus far has been from the perspective of an outside observer who knows the agent's true ability z. We therefore took expectations only over sequences of ϵ_t. However, the agent does not know z. The agent therefore takes expectations over both z and ϵ_t. Using $\hat{E}_0[z] = \mu_0$, we have:

$$\begin{aligned}\hat{E}_0[\mu_1] =&[1 - w(e_0, \Sigma_0)]\mu_0 + w(e_0, \Sigma_0)\hat{E}_0[z] + w(e_0, \Sigma_0)\hat{E}_0\left[\frac{\sqrt{g(e_0)}}{e_0}\epsilon_0\right] \\ =&\mu_0.\end{aligned}$$

The agent does not expect his central estimate μ to change from period 0 to period 1. Now consider the agent's expectation of his central estimate at time 2:

$$\hat{E}_0[\mu_2] =\hat{E}_0\left[[1 - w(e_1, \Sigma_1)]\mu_1 + w(e_1, \Sigma_1)z + w(e_1, \Sigma_1)\frac{\sqrt{g(e_1)}}{e_1}\epsilon_1\right].$$

Using $\hat{E}_0[\epsilon_1] = 0$ and substituting for μ_1, we have:

$$\begin{aligned}\hat{E}_0[\mu_2] =&\hat{E}_0\left[1 - w(e_1, \Sigma_1)\right]\hat{E}_0\left[\mu_1\right] - \widehat{Cov}_0\left[w(e_1, \Sigma_1), w(e_0, \Sigma_0)z\right] - \widehat{Cov}_0\left[w(e_1, \Sigma_1), w(e_0, \Sigma_0)\frac{\sqrt{g(e_0)}}{e_0}\epsilon_0\right] \\ &+ \hat{E}_0\left[w(e_1, \Sigma_1)\right]\hat{E}_0[z] + \widehat{Cov}_0\left[w(e_1, \Sigma_1), z\right].\end{aligned}$$

Using $\hat{E}_0[\mu_1] = \hat{E}_0[z] = \mu_0$, this becomes:

$$\hat{E}_0[\mu_2] =\mu_0 + [1 - w(e_0, \Sigma_0)]\widehat{Cov}_0\left[w(e_1, \Sigma_1), z\right] - w(e_0, \Sigma_0)\frac{\sqrt{g(e_0)}}{e_0}\widehat{Cov}_0\left[w(e_1, \Sigma_1), \epsilon_0\right].$$

Using Stein's Lemma, we find:

$$\begin{aligned}\widehat{Cov}_0\left[w(e_1, \Sigma_1), z\right] =&\Sigma_0\hat{E}_0\left[\frac{\partial w(e_1, \Sigma_1)}{\partial e_1}\frac{\mathrm{d}e_1}{\mathrm{d}\mu_1}\frac{\mathrm{d}\mu_1}{\mathrm{d}z}\right], \\ \widehat{Cov}_0\left[w(e_1, \Sigma_1), \epsilon_0\right] =&\sigma^2\hat{E}_0\left[\frac{\partial w(e_1, \Sigma_1)}{\partial e_1}\frac{\mathrm{d}e_1}{\mathrm{d}\mu_1}\frac{\mathrm{d}\mu_1}{\mathrm{d}\epsilon_0}\right].\end{aligned}$$

Substituting yields:

$$\hat{E}_0[\mu_2] = \mu_0 + w(e_0,\Sigma_0)[1-w(e_0,\Sigma_0)]\Sigma_0\hat{E}_0\left[\frac{\partial w(e_1,\Sigma_1)}{\partial e_1}\frac{\mathrm{d}e_1}{\mathrm{d}\mu_1}\right] - w(e_0,\Sigma_0)^2\frac{g(e_0)}{e_0^2}\sigma^2\hat{E}_0\left[\frac{\partial w(e_1,\Sigma_1)}{\partial e_1}\frac{\mathrm{d}e_1}{\mathrm{d}\mu_1}\right]$$
$$= \mu_0,$$

where the second equality uses, from equation (6), $w(e_0,\Sigma_0)\frac{g(e_0)}{e_0^2}\sigma^2 = [1-w(e_0,\Sigma_0)]\Sigma_0$. The agent does not expect his central estimate to drift over time.

Why the difference between $E_0[\mu_t]$ and $\hat{E}_0[\mu_t]$? Assume that $\partial w(e_1,\Sigma_1)/\partial e_1 > 0$. We previously saw that $Cov_0[w_1,\mu_1] > 0$ implies that $E_0[\mu_2] < z$, despite $\mu_0 = z$. From the agent's perspective, however, there are conflicting effects that cancel. First, uncertainty about ϵ_0 generates $\widehat{Cov}_0[w_1,\epsilon_0] > 0$ and tends to push $\hat{E}[\mu_2]$ below μ_0. High (low) values of ϵ_0 tend to generate high (low) values of π_0, which lead the agent to become overconfident (underconfident) in period 1. As in our previous story, the agent learns away mistaken beliefs especially rapidly when overconfident and especially slowly when underconfident, leading him to become underconfident on average. Second, uncertainty about the true value z generates $\widehat{Cov}_0[w_1,z] > 0$ and tends to push $\hat{E}[\mu_2]$ above μ_0. High (low) values of z mean that the agent is underconfident (overconfident) in time 0. The resulting high (low) values of π_0 tend to correct the agent's initially mistaken beliefs, leading him to increase (decrease) his period 1 effort. In period 1, the agent learns the true z especially rapidly when effort is high (because z is high) and especially slowly when effort is low (because z is low). The agent's uncertainty about z drags the agent's average central estimate upwards over time because he approaches the truth faster when z is high. For the Bayesian agent, the tendency of ϵ_0 to generate underconfidence exactly cancels the tendency of z to generate overconfidence. The agent does not expect his beliefs to drift one way or the other, even though on average they will.

4 General Analysis

We now establish that the drift in the agent's beliefs persists beyond period 2 and connect this drift to the agent's beliefs about the variance of π_t.

The elasticity of $g(e)$ with respect to e will play a critical role. Define this elasticity as $\chi(e) \triangleq e\,g'(e)/g(e)$. The following lemma shows that $\chi(e)$ determines how effort choices affect posterior beliefs:

Lemma 1. *For given Σ_t, w_t increases in e_t if $\chi(e_t) < 2$, decreases in e_t if $\chi(e_t) > 2$, and is independent of e_t if $\chi(e_t) = 2$. Σ_{t+1} decreases in e_t if and only if w_t increases in e_t.*

Proof. Differentiating, we have:

$$\frac{\partial w(\Sigma_t,e_t)}{\partial e_t} = w_t(1-w_t)\left[\frac{2}{e_t} - \frac{g'(e_t)}{g(e_t)}\right].$$

This is strictly positive if $\frac{g'(e_t)}{g(e_t)}e_t < 2$, is strictly negative if $\frac{g'(e_t)}{g(e_t)}e_t > 2$, and is zero if $\frac{g'(e_t)}{g(e_t)}e_t = 2$. The result for Σ_{t+1} follows straightforwardly. □

To gain intuition for this result, observe that greater effort leads the agent to place more weight on the time t signal if and only if this effort reduces the perceived variance of the time t signal. The agent believes that the variance of his time t signal π_t/e_t is $\sigma^2 g(e_t)/e_t^2$. The $g(e_t)$ reflects the agent's beliefs about the relationship between effort and the variance of the observed outcomes π_t. When $g'(\cdot) > 0$, the agent believes that additional effort increases the variance of the reward π_t. The e_t^2 captures a second effect of effort. Because effort and ability are complementary, additional effort increases the marginal effect of ability on the reward. This effect helps the agent to learn faster: it works to reduce the perceived variance of the signal π_t/e_t. If $\chi(e_t) > 2$, then the two effects conflict and the first effect dominates; if $\chi(e_t) \in [0, 2)$, then the second effect dominates; and if $\chi(e_t) < 0$, then the two effects go in the same direction. The two effects cancel when $\chi(e_t) = 2$: we then have $g(e_t) = A\,e_t^2$ for $A > 0$ and the variance of the perceived signal becomes $A\,\sigma^2$, which is independent of effort.

We assumed that the agent is unbiased at time 0. We are interested in whether the agent remains unbiased in expectation. We have the following lemma and proposition, which rely on approximations that are good as long as σ^2 and Σ_0 are not too large.

Lemma 2.

1. *If $\chi(e^*(z)) < 2$, then $Cov_0[w_t, \mu_t] \geq 0$ for all $t \geq 1$.*
2. *If $\chi(e^*(z)) > 2$, then $Cov_0[w_t, \mu_t] \leq 0$ for all $t \geq 1$.*
3. *If $\chi(e^*(z)) = 2$, then $Cov_0[w_t, \mu_t] \approx 0$ for all $t \geq 1$.*
4. *If effort is fixed exogenously, then $Cov_0[w_t, \mu_t] = 0$ for all $t \geq 1$.*

Proof. See appendix. □

Proposition 1.

1. *$E_0[\mu_1] = z$.*
2. *If $\chi(e^*(z)) < 2$, then $E_0[\mu_t] < z$ for all $t > 1$.*
3. *If $\chi(e^*(z)) > 2$, then $E_0[\mu_t] > z$ for all $t > 1$.*
4. *If $\chi(e^*(z)) = 2$, then $E_0[\mu_t] \approx z$ for all $t > 1$.*
5. *If the agent commits to future effort levels at time 0, then $E_0[\mu_t] = z$ for all $t \geq 1$.*

Proof. See appendix. □

The proposition says that all agents continue to have unbiased beliefs (in expectation) at time 1, but whether their later beliefs become biased depends on $\chi(e^*(z))$. From the tower property, we have

$$E_0[\mu_{t+1}] = E_0[w_t]z + (1 - E_0[w_t])E_0[\mu_t] - Cov_0[w_t, \mu_t]. \tag{8}$$

When $\chi(e^*(z)) = 2$, the variance of the perceived signal is independent of effort choices around $e^*(z)$. From Lemma 2, we have $Cov_0[w_t, \mu_t] = 0$. In this case, the amount the agent learns from a period's observation does not depend on the effort level he chooses. From equation (8), we always expect new data to move the agent's beliefs towards z. Because the agent's central estimates tend to move towards z at a rate that is independent of the agent's bias, the agent's beliefs will tend to remain centered around the truth if they begin at the truth. This result is consistent with intuition from standard models of unbiased learning.

Matters are different when $\chi(e^*(z)) \neq 2$. Now, from Lemma 2, $Cov_0[w_t, \mu_t] \neq 0$. The informativeness of the reward π_t now varies with mean beliefs μ_t and thus affects the expected drift in beliefs. Consider the case with $\chi(e^*(z)) < 2$. Here, additional effort enhances the signal-to-noise ratio in the reward π_t. As a result, the agent's posterior becomes especially sensitive to the observed π_t when effort is high: $Cov_0[w_t, \mu_t] > 0$. High effort levels correspond to cases in which the agent's central estimate has drifted above z. The next shock is likely to pull the agent's mean back towards z, and it will do so especially rapidly because of how the agent updates beliefs under high effort levels. In contrast, when negative shocks happen to push the agent's central estimate below z, the agent will be slower to revise his beliefs back towards z because his chosen effort will be low. The agent will thus tend to get stuck with overly pessimistic self-evaluations for longer than he is stuck with overly optimistic self-evaluations. From the perspective of time 0, the agent's future central estimate will, on average, be too pessimistic. In the opposite case of $\chi(e^*(z)) > 2$, high effort levels hinder updating. The foregoing logic then implies that the agent's future central estimates will, on average, be too optimistic.[15]

A special case is of particular interest for its intuitive interpretation of $g(e_t)$. Let the agent have two possible models for the variance of the random shocks: a first model in which shocks are to his ability z, and a second model in which shocks are purely external. In the first model, the marginal payoff from greater effort is stochastic, but in the second model,

[15] As infinite data accumulates, the agent's beliefs are consistent in the sense that they converge, in expectation, to z. To see this, note that the sequence of data to be explained can be written as $(\pi_1/e_1, \pi_2/e_2, ..., \pi_n/e_n)$. Each observation is drawn from a normal distribution, with a variance that changes with e_t and a mean that is constant over time. The observations are independently distributed. Choi and Ramamoorthi (2008) review consistency when observations are independently, non-identically distributed. We have a further wrinkle in that we allow the agent's model to be potentially misspecified. However, it is known that posteriors converge to the value that minimizes the Kullback-Leibler divergence from the agent's model to the true data generating model (e.g., Gelman et al., 2004), and it is easy to show that $\mu_t = z$ is the minimizer in our setting.

the marginal payoff from greater effort is deterministic. The agent places weight $\gamma \in [0,1]$ on the first model and weight $1-\gamma$ on the second model. Formally, we have:

Assumption 1. $g(e_t) = \gamma e_t^2 + (1-\gamma)$

When $\gamma = 1$, the agent believes that $\pi_t = e_t(z + \epsilon_t)$, and when $\gamma = 0$, the agent believes that $\pi_t = e_t z + \epsilon_t$. The following corollary describes how the drift in the agent's beliefs depends on γ:

Corollary 2. *Let Assumption 1 hold. Then* $E_0[\mu_t] < z$ *for all* $t > 1$ *if* $\gamma < 1$ *and* $E_0[\mu_t] \approx z$ *for all* $t \geq 1$ *if* $\gamma = 1$.

Proof. Follows from Proposition 1 once we recognize that $\gamma < 1$ implies $\chi(\cdot) < 2$ and that $\gamma = 1$ implies $\chi(\cdot) = 2$. □

Return to the example from Section 2 of students taking tests and learning about their ability from their scores. Consider two groups of students, both satisfying Assumption 1. The first group of students believes that applying greater effort has stochastic consequences ($\gamma = 1$), and the second group believes that applying greater effort increases their reward deterministically ($\gamma = 0$). When a student in the first group tries hard on a test and gets a bad score, he does not adjust his beliefs about his own ability very strongly because the student writes off the result to bad luck. The student's self-assessments remain, in expectation, centered around his true ability. However, we expect students in the second group to have a different fate. These students believe that trying hard should allow their talent to shine through. When these students see a bad test score after trying especially hard, they infer a lack of ability. This more negative assessment of their own ability leads them to apply less effort on the next test. On average, this next test will tell them that they are more talented than they believe themselves to be, but because they know that their effort level was low, they do not pay as much attention to this score. They therefore tend to retain an overly negative self-assessment. The key determinant of students' average beliefs and effort choices is not the actual data generating process but rather their beliefs about whether additional effort provides a better or worse signal of their ability.

5 A Numerical Example

In order to make these ideas concrete, we now consider a numerical example. Impose $c(e_t) = \frac{1}{2}e_t^2$. Let $z = \mu_0 = 20$, $\sigma^2 = \Sigma_0 = 16$, $f(e_t) = e_t^2$, and $g(e_t) = 1$. We therefore have $\chi(\cdot) = 0$.

The top panels of Figure 1 plot the distribution of μ_t (left) and w_t (right) for $t \in \{1, 2, 3, 4, 5, 10\}$. μ_1 is normally distributed but the other distributions are skewed.[16] The

[16] This skew arises because, first, μ_{t+1} depends on w_t (which is a nonlinear function of μ_t) and, second, because the realized signal π_t/e_t is a nonlinear function of e_t. Both sources of skewness vanish when Assumption 1 holds with $\gamma = 1$ and $f(\cdot) = g(\cdot)$.

distribution of μ_t becomes progressively narrower as data accumulates. The negative skew in μ_t translates into a skew in w_t in the first few periods, before the lower bound of 0 starts to affect the distribution of w_t. The skew in w_t becomes less pronounced as Σ_t falls, for the agent does not pay much attention to π_t (regardless of e_t) as he becomes certain of his true ability.

The lower left panel of Figure 1 confirms the results of Proposition 1: $E_0[\mu_t]$ equals z at $t = 1$, but $E_0[\mu_t]$ drops below z at $t = 2$ and remains below z for all greater t. $E_0[\mu_t]$ does approach z again as t goes to infinity, but this approach is slow. The maximum average bias arises in period 2. The average bias is still 40% of this maximum in period 10, 4% of this maximum in period 100, and 0.8% of this maximum in period 500. The circles show that the agent's uncertainty about his ability does decline quickly as he observes additional data, but his beliefs nonetheless remain biased on average.

The lower right panel plots $Cov_0[\mu_t, w_t]$ (crosses) as well as the correlation (circles) between μ_t and w_t. The covariance and correlation are positive because states of the world with large μ_t are states in which the agent chooses high effort e_t and because w_t increases in e_t. We see that the covariance is especially positive in early periods when the agent is most uncertain about his own ability. This strongly positive covariance explains why $E_0[\mu_t]$ declines over those early periods. The covariance approaches zero after the first few periods not because μ_t and w_t become uncorrelated over long horizons (the correlation in fact remains clearly positive even at long horizons) but because the variance of each variable declines strongly as the agent becomes more certain of his ability.

6 Application to Cognitive Behavioral Therapy

The foregoing analysis sheds new light on the nature of depression and on possible reasons why current treatments are successful. In the early and middle decades of the twentieth century, psychologists sought the reasons for the negative thoughts that characterize depression. In the 1960s and 1970s, a revolution in psychology viewed negative thinking as the content of depression (Seligman, 1991). Aaron Beck, one of the forefathers of modern clinical psychology, writes,

> The cognitive theory of depression is based essentially on an information-processing model. A pronounced and prolonged negative biasing of this process is manifest in the characteristic thinking disorder in depression (selective abstraction, overgeneralization, negative self-attributions). (Beck, 2002, 29)

Beck refers to the negative filter as a form of "automatic thinking". Seligman (1991) refers to it as a "pessimistic explanatory style" and emphasizes how agents with this explanatory style attribute bad events to their own pervasive self rather than to transient, chance outcomes. This negative filter causes the agent to doubt his own ability and leaves the agent especially

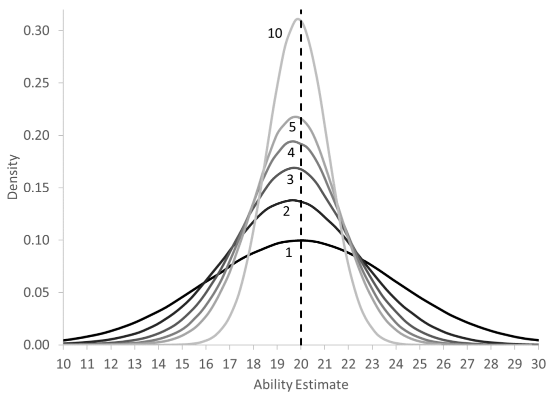

(a) Evolution of the estimate μ_t

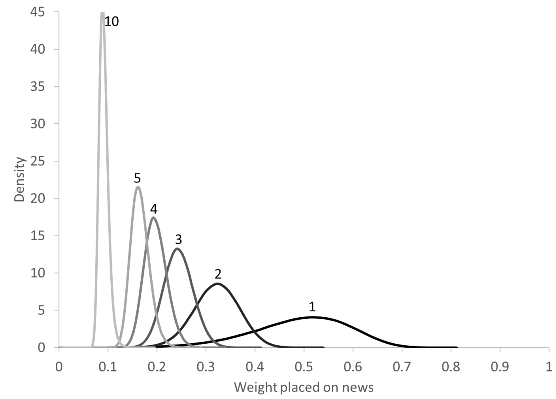

(b) Evolution of the news weight w_t

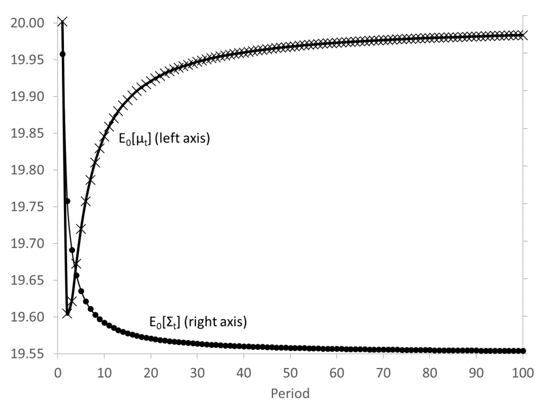

(c) $E_0(\mu_t)$ and $E_0(\Sigma_t)$

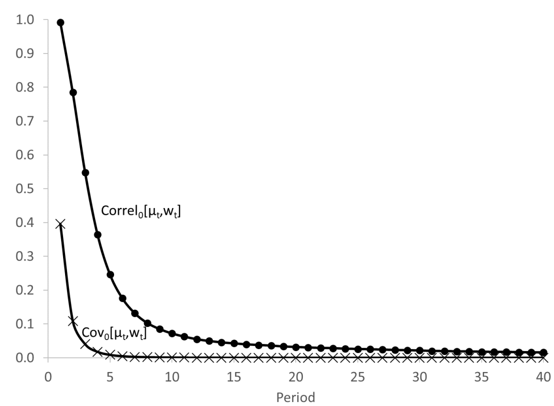

(d) $Cov_0(\mu_t, w_t)$ and $Correl_0(\mu_t, w_t)$

Figure 1: Top: The distribution of μ_t (left) and w_t (right) for $t \in \{1, 2, 3, 4, 5, 10\}$. The dashed vertical line in the top left panel indicates the true ability z, which is also μ_0. Bottom: The evolution of $E_0[\mu_t]$ and $E_0[\Sigma_t]$ (left), and of $Cov_0[\mu_t, w_t]$ and the correlation coefficient between μ_t and w_t (right). All plots sample one million trajectories for ϵ_t.

vulnerable to negative shocks. The agent withdraws from the world, displaying the passivity and inertia that characterize depressives' behavior. The agent's self-doubt reinforces his inactivity, and his inactivity closes off the possibility of receiving new signals that could ameliorate his self-doubt.[17]

Cognitive behavioral therapy emerged from this theory of depression. Cognitive behavioral therapy emphasizes two goals in treating depression: first, the therapist aims to change behavior by convincing the agent to re-engage with the world, and second, the therapist aims to change the agent's negative thinking by teaching new information processing techniques (Layard and Clark, 2015). In empirical tests, cognitive behavioral therapy has been found to effectively mitigate depression. Further, patients relapse at a lower rate than after other treatments, including pharmacological treatments (Seligman, 1991; Scott, 1996; Hollon et al., 2005; Butler et al., 2006; Dobson et al., 2008; Layard and Clark, 2015). Cognitive behavioral therapy seems to not just treat an instance of depression but also to "inoculate" patients against future depression (Seligman, 1991).

The present setting captures both the modern understanding of depression and the short- and long-run benefits of cognitive behavioral therapy. We see that when $\chi(\cdot) < 2$, an agent tends to end up underconfident because he learns away positive shocks quickly but does not learn away negative shocks quickly. When positive shocks lead him to a high estimate of his own ability, he increases his effort level and attributes future shocks to his own ability. But when his self-estimate is overly inflated, any news is likely to be negative. He therefore revises his self-estimate sharply downward and reduces his effort accordingly. This agent demonstrates the "pessimistic explanatory style" described by Seligman (1991), attributing this negative news to his own shortcomings rather than to chance. If the agent's ability estimate falls far enough, then the news is now, on average, going to be positive, but this overly pessimistic agent discounts that news because he knows that he did not apply much effort. The agent becomes underconfident on average because he discounts these positive shocks as largely due to chance even though he treated the initial negative shocks as especially reflective of his own ability. In contrast, agents with $\chi(\cdot) > 2$ do not believe that they should learn more from applying high effort levels. These agents tend to maintain an overly optimistic self-assessment because they write off the negative shocks that tend to accompany high effort choices as largely due to chance while emphasizing their own role in the positive shocks that tend to accompany low effort levels.

Especially acute negative shocks are known to trigger depression (Seligman, 1991; Beck, 2002). Following such a shock, an agent's estimate of his own ability may fall to an especially low level. His effort will also fall, so that he displays the passivity and inertia characteristic of depression. This low effort will allow only very slow learning. The agent thus persists in his negative beliefs despite receiving signals that his ability is not actually so low. Treatments

[17]Beck (2002, 59) writes, "Finally, the attribution to the self of deficiency removes any expectation of succor and reinforces the sense of *futility*. Thus, the cognitive content would tend to promote prolonged inactivity."

that lead the agent to re-engage with the world (either directly or by raising his estimate of his own ability) can be successful at curing his depression: the agent can return to a state with better self-confidence and greater effort once he starts receiving more informative signals as a result of applying more effort. However, as long as $\chi(\cdot) < 2$, his average estimates will once again tend to fall below his true ability. Only a treatment that changes his form of information processing can have longer-run effects in preventing a relapse of depression. If a treatment can teach the agent to stop expecting high effort to provide more control over the world, then the agent will be able to maintain higher confidence on average and will recover faster from negative shocks. This analysis is consistent with the goals of cognitive behavioral therapy.

We have developed a theory of rational depression that comports with leading explanations from psychology. Importantly, we do not need to postulate an ex ante behavioral bias in information processing, such as in Gervais and Odean (2001). Instead, we merely postulate that some agents believe (correctly or not) that greater effort provides a better signal of their own ability. Agents who believe that increasing effort should make the world respond in a more regular way quickly learn away high confidence states but learn away low confidence states only slowly. As a result, the "pessimistic explanatory style" described by Seligman (1991) emerges endogenously from rational updating: an agent with $\chi(\cdot) < 2$ tends to attribute bad events to his own pervasive ability, and such agents end up underconfident on average. Rather than being a behavioral bias that causes depression, a pessimistic explanatory style may simply reflect the combination of rational effort choices and the stories that Bayesian updaters tell themselves about the origin of chance events. A treatment for depression can have long-term success only if it changes the style of information processing, which may explain why cognitive behavioral therapy has proven more successful than other treatments.

7 Forward-Looking Effort Choices

We have thus far made only a single departure from full rationality: we assumed that agents choose effort myopically, even though they learn optimally.[18] However, forward-looking agents should account for the informational value of their effort choices. We now extend the analysis to the case of forward-looking agents. We will see that foresight does not break the link between effort choices and ability estimates that drives our results, we will establish that all previous results hold in a set of cases of particular interest, and we will see that forward-looking behavior only minimally affects the numerical example from Section 5.

The agent now chooses his effort to maximize his expected present value over an infinite planning horizon. Let the agent's per-period discount factor be $\beta \in (0, 1)$, and let $V(\mu_t, \Sigma_t)$

[18]Of course, this assumption is actually consistent with full rationality if agents' preferences do not place much weight on the future.

denote the present value of the agent's optimal policy program from any time t with estimate μ_t and variance Σ_t. The agent's effort choices and value function solve the following Bellman equation:

$$V(\mu_t, \Sigma_t) = \max_{e_t} \hat{E}_t \left[\pi_t - c(e_t) + \beta V(\mu_{t+1}, \Sigma_{t+1})\right],$$

subject to the transition equations (7) and (5). Effort will now be a function of the central estimate μ_t and also of the agent's uncertainty Σ_t about that estimate. Optimal effort $e^*(\mu_t, \Sigma_t)$ satisfies the following first-order condition:

$$\begin{aligned} c'(e_t) =& \mu_t + \beta \hat{E}_t \left[\frac{\partial V(\mu_{t+1}, \Sigma_{t+1})}{\partial \mu_{t+1}} \frac{\partial w_t}{\partial e_t} \left(z + \frac{\sqrt{g(e_t)}}{e_t}\epsilon_t - \mu_t\right)\right] \\ &+ \beta \hat{E}_t \left[\frac{\partial V(\mu_{t+1}, \Sigma_{t+1})}{\partial \mu_{t+1}} w_t \frac{\mathrm{d}[\sqrt{g(e_t)}/e_t]}{\mathrm{d}e_t}\epsilon_t\right] + \beta \hat{E}_t \left[\frac{\partial V(\mu_{t+1}, \Sigma_{t+1})}{\partial \Sigma_{t+1}} \frac{\partial \Sigma_{t+1}}{\partial e_t}\right]. \end{aligned}$$

Note that

$$\hat{E}_t \left[z + \frac{\sqrt{g(e_t)}}{e_t}\epsilon_t - \mu_t\right] = 0 \text{ and } \hat{E}_t \left[w_t \frac{\mathrm{d}[\sqrt{g(e_t)}/e_t]}{\mathrm{d}e_t}\epsilon_t\right] = 0.$$

Also note that $\partial \Sigma_{t+1}/\partial e_t$ is not random from the perspective of time t. The first-order condition becomes:

$$\begin{aligned} c'(e_t) =& \mu_t + \beta \hat{Cov}_t \left[\frac{\partial V(\mu_{t+1}, \Sigma_{t+1})}{\partial \mu_{t+1}}, z + \frac{\sqrt{g(e_t)}}{e_t}\epsilon_t\right] \frac{\partial w_t}{\partial e_t} \\ &+ \beta \hat{Cov}_t \left[\frac{\partial V(\mu_{t+1}, \Sigma_{t+1})}{\partial \mu_{t+1}}, \epsilon_t\right] w_t \frac{\mathrm{d}[\sqrt{g(e_t)}/e_t]}{\mathrm{d}e_t} + \beta \hat{E}_t \left[\frac{\partial V(\mu_{t+1}, \Sigma_{t+1})}{\partial \Sigma_{t+1}}\right] \frac{\partial \Sigma_{t+1}}{\partial e_t}. \end{aligned} \quad (9)$$

We see three new terms relative to the myopic agent's first-order condition (3). The new term on the first line accounts for the riskiness of the new information to be received following time t. The agent knows that his beliefs are likely to change, but he does not know in advance how they will change. The second argument of the covariance operator is the random component of the agent's time t signal π_t/e_t. When the covariance is positive, uncertainty about the rewards π_t increases the agent's expected payoff, so the agent moves his effort choice in a direction that increases the sensitivity of his posterior beliefs μ_{t+1} to π_t (i.e., if $\partial w_t/\partial e_t > 0$, the agent chooses greater effort). The first term on the second line is similar, except accounting for how the agent's effort choice affects the variance of the agent's posterior beliefs by affecting the variance of π_t. The final term on the second line accounts for the effect of effort choices on the agent's uncertainty about his own ability. This channel increases the agent's optimal effort choice when the agent prefers to be certain of his own ability ($\partial V(\mu_{t+1}, \Sigma_{t+1})/\partial \Sigma_{t+1} < 0$) and additional effort reduces his uncertainty about his ability ($\partial \Sigma_{t+1}/\partial e_t < 0$).

The results in the myopic setting were driven by the sensitivity of effort choices e_t to estimated ability μ_t. If these new terms break the link between effort choices and estimated ability, then we would expect that the main results would vanish and the agent's beliefs would remain unbiased on average. However, the next proposition shows that making agents forward-looking does not break the link between effort and estimated ability:

Proposition 3. *e_t^* cannot be independent of μ_t at all times t.*

Proof. Repeatedly applying the envelope theorem, we have:

$$\begin{aligned}\frac{\partial V(\mu_{t+1},\Sigma_{t+1})}{\partial \mu_{t+1}} =& e_{t+1}^* + \beta \hat{E}_{t+1}\left[\frac{\partial V(\mu_{t+2},\Sigma_{t+2})}{\partial \mu_{t+2}}(1-w_{t+1})\right]\\ =& e_{t+1}^* + \beta(1-w_{t+1})\sum_{i=1}^{\infty}\hat{E}_{t+1}\left[\beta^{i-1}e_{t+1+i}^*\prod_{j=1}^{i-1}(1-w_{t+1+j})\right].\end{aligned}$$

If e_t^* is independent of μ_t at all t, then so is Σ_t, and if both Σ_t and e_t^* are independent of μ_t, then so is w_t. And being independent of μ_t implies being independent of z and ϵ_{t-1}. Therefore if e_t^* is independent of μ_t at all t, then $\hat{Cov}_t\left[\frac{\partial V(\mu_{t+1},\Sigma_{t+1})}{\partial \mu_{t+1}}, z + \frac{\sqrt{g(e_t^*)}}{e_t^*}\epsilon_t\right] = 0$ for all t.

The envelope theorem also yields:

$$\begin{aligned}\frac{\partial V(\mu_{t+1},\Sigma_{t+1})}{\partial \Sigma_{t+1}} =& \beta \hat{E}_{t+1}\left[\frac{\partial V(\mu_{t+2},\Sigma_{t+2})}{\partial \mu_{t+2}}\frac{\partial \mu_{t+2}}{\partial w_{t+1}}\frac{\partial w_{t+1}}{\partial \Sigma_{t+1}} + \frac{\partial V(\mu_{t+2},\Sigma_{t+2})}{\partial \Sigma_{t+2}}\frac{\partial \Sigma_{t+2}}{\partial \Sigma_{t+1}}\right]\\ =& \beta \hat{Cov}_{t+1}\left[\frac{\partial V(\mu_{t+2},\Sigma_{t+2})}{\partial \mu_{t+2}}, z + \frac{\sqrt{g(e_{t+1}^*)}}{e_{t+1}^*}\epsilon_{t+1}\right]\frac{\partial w_{t+1}}{\partial \Sigma_{t+1}} + \beta \hat{E}_{t+1}\left[\frac{\partial V(\mu_{t+2},\Sigma_{t+2})}{\partial \Sigma_{t+2}}\right]\frac{\partial \Sigma_{t+2}}{\partial \Sigma_{t+1}}.\end{aligned}$$

Continuing in this fashion, we see that if $\hat{Cov}_{t+1}\left[\frac{\partial V(\mu_s,\Sigma_s)}{\partial \mu_s}, z + \frac{\sqrt{g(e_{s-1}^*)}}{e_{s-1}^*}\epsilon_{s-1}\right] = 0$ for all $s \geq t+2$, then $\partial V(\mu_{t+1},\Sigma_{t+1})/\Sigma_{t+1} = 0$. And we previously saw that if e_t^* is independent of μ_t at all t, then $\hat{Cov}_t\left[\frac{\partial V(\mu_{t+1},\Sigma_{t+1})}{\partial \mu_{t+1}}, z + \frac{\sqrt{g(e_t^*)}}{e_t^*}\epsilon_t\right] = 0$ for all t. Therefore, if e_t^* is independent of μ_t at all t, then $\partial V(\mu_{t+1},\Sigma_{t+1})/\Sigma_{t+1} = 0$.

Assume that e_t^* is independent of μ_t at all t. Then using the preceding results in equation (9), we find that optimal effort must satisfy $c'(e_t^*) = \mu_t$. But this contradicts the assumption that e_t^* is independent of μ_t. Therefore e_t^* must depend on μ_t at some t.

□

Because at least some effort choices must depend on estimated ability, endogenous effort choices may still create the asymmetrical rates of learning that generated the drift in the myopic agent's ability estimate.

We now establish a sufficient condition for the average drift in the forward-looking agent's ability estimate to be qualitatively similar to the myopic setting.

Proposition 4. *The results from Section 4 apply if $\chi(\cdot)$ is close to 2.*

Proof. Note that

$$\frac{\mathrm{d}}{\mathrm{d}e_t}\frac{\sqrt{g(e_t)}}{e_t} = \frac{1}{2}\frac{\sqrt{g(e_t)}}{e_t^2}\left[\chi(e_t) - 2\right] \quad \text{and} \quad \frac{\partial \Sigma_{t+1}}{\partial e_t} = \Sigma_{t+1}^2 \sigma^{-2}\frac{e_t}{g(e_t)}\left[\chi(e_t) - 2\right].$$

Lemma 1 established that $\partial w_t/\partial e_t \to 0$ as $\chi(e_t) \to 2$. We now see that the above terms also go to 0 as $\chi(e_t) \to 2$. Substituting into equation (9), we have that $c'(e_t) \to \mu_t$ as $\chi(e_t) \to 2$. Therefore, as $\chi(\cdot) \to 2$, forward-looking agents' optimal effort choices approach the effort choices made by myopic agents. Because forward-looking agents' updating rules differ from myopic agents' updating rules only through their realized effort choices, the evolution of forward-looking agents' beliefs will be similar to the evolution of myopic agents' beliefs when their effort choices are similar. Therefore, the results from the myopic setting hold for $\chi(\cdot)$ near 2.

□

Corollary 5. *Let Assumption 1 hold. Then there exists $\delta \in (0, 1]$ such that $E_0[\mu_t] < z$ for all $t > 1$ if $\gamma \in [1 - \delta, 1)$ and such that $E_0[\mu_t] \approx z$ for all $t \geq 1$ if $\gamma = 1$.*

Proof. Follows from Corollary 2 and Proposition 4. □

In our analysis of equation (9), we saw that the forward-looking agent's effort choices differ from the myopic agent's effort choices because the forward-looking agent accounts for how additional effort affects his rate of learning. When $\chi(\cdot)$ is close to 2, effort has only a small effect on the agent's rate of learning, so the forward-looking agent's effort choices are similar to the myopic agent's effort choices. Because the two settings differ only via the agents' effort choices, all of the results from the myopic setting apply to the forward-looking setting when $\chi(\cdot)$ is sufficiently close to 2.

We now return to the numerical example from Section 5. Recall that this example used $\chi(\cdot) = 0$. Figure 2 plots the percentage difference in the average ability estimate (crosses), the average effort (circles), and the average variance of the agent's beliefs (squares) between a myopic agent and a forward-looking agent.[19] We study a rather patient agent, with $\beta = 0.99$. The forward-looking agent chooses slightly less effort on average and as a result maintains slightly greater uncertainty. The forward-looking agent's effort choices do slightly increase most periods' correlation (not shown) between μ_t and w_t. However, these differences do not affect the drift in the agent's ability estimate to any appreciable degree: $E_0[\mu_t]$ differs by less than 0.01% between the two settings. On average, the forward-looking agent becomes persistently underconfident to about the same degree as the myopic agent becomes persistently underconfident.

[19] I solve the forward-looking model via value function iteration. I use the collocation method, with 5^2 Chebyshev nodes and a 5^2 Chebyshev basis. I integrate via Gauss-Legendre quadrature. The domain of approximation for μ covers $z \pm 3\sqrt{\Sigma_0}$, and the domain of approximation for Σ ranges from 0 to $\frac{3}{2}\Sigma_0$.

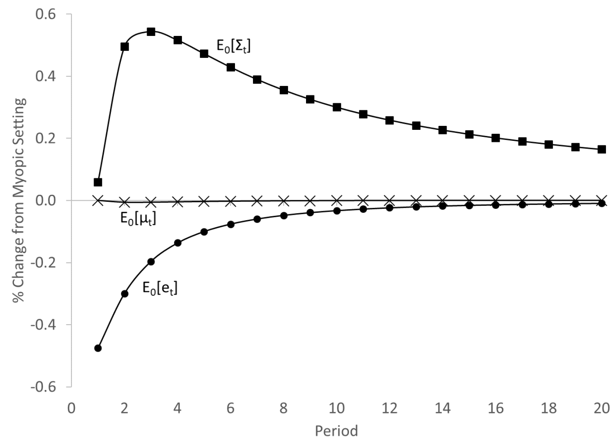

Figure 2: The difference in average effort, estimates, and uncertainty between a myopic agent and a forward-looking agent with $\beta = 0.99$. All paths use 10,000 draws.

8 Conclusion

We have seen that rational agents become, on average, persistently underconfident when they believe that additional effort makes outcomes more predictable and become persistently overconfident when they believe that additional effort makes outcomes less predictable. The critical element is that returns to effort vary with agents' beliefs about their own ability. As a result, agents' effort choices vary with their beliefs about their own ability. When agents believe that high effort provides a particularly informative signal of ability, agents will learn away mistaken overconfidence quickly (due to high effort) but will learn away mistaken underconfidence only slowly (due to low effort). Such agents end up underconfident on average. But if agents believe that high effort provides a particulaly poor signal of their ability, they end up overconfident on average. Only in a knife-edge case do they tend to end up with correctly calibrated beliefs.

These results call for three further types of investigations. First, psychologists have successfully connected "explanatory styles" to a range of outcomes. The present model provides a structural interpretation of explanatory styles, identifying the types of beliefs that can generate observed patterns. Future research in psychology should assess agents' views on the informativeness of effort for ability. Second, the results call for experiments that test the implications for confidence of manipulating the data generating process and also the implications for confidence of manipulating effort. In particular, recent work (e.g., Gneezy et al., 2017) has experimentally manipulated effort in order to show the importance of effort for outcomes. The present model suggests also exploring the importance of effort for agents' willingness to extract information from outcomes. Third, managers who do not internalize employees' cost of effort have an incentive to increase employees' effort, and employees may learn about their own ability from feedback such as compensation or evaluations. Future work should investigate how managers may design compensation or other feedback mechanisms so as to lead employees to become systematically overconfident and thereby choose high effort levels.

Appendix

Proof of Lemma 2

Part 1

Assume that $\chi(e^*(z)) < 2$. We proceed by induction.

Induction step:
The induction hypothesis is that $Cov_0[w_{t-1}, \mu_{t-1}] \geq 0$ for some $t \geq 2$. Fix $\bar{\Sigma} \in (0, \Sigma_0]$. Use a first-order approximation to $w(e^*(\mu_t), \Sigma_t)$ around $\mu_t = z$ and $\Sigma_t = \bar{\Sigma}$ to obtain

$$\begin{aligned} Cov_0[w_t, \mu_t] \approx & Cov_0\left[\mu_t, \left.\frac{\partial w_t}{\partial e_t}\right|_{(e^*(z),\bar{\Sigma})} \left.\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\right|_z \mu_t + \left.\frac{\partial w_t}{\partial \Sigma_t}\right|_{(e^*(z),\bar{\Sigma})} \Sigma_t\right] \\ = & \left.\frac{\partial w_t}{\partial e_t}\right|_{(e^*(z),\bar{\Sigma})} \left.\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\right|_z Var_0[\mu_t] + \left.\frac{\partial w_t}{\partial \Sigma_t}\right|_{(e^*(z),\bar{\Sigma})} Cov_0[\mu_t, \Sigma_t]. \end{aligned}$$

This approximation is good as long as σ^2 and Σ_0 are not too large. Note that:

$$\frac{\partial w_t}{\partial e_t} = (1 - w_t) w_t \left[2 - \chi(e(\mu_t))\right] \frac{1}{e(\mu_t)}, \quad \frac{\partial w_t}{\partial \Sigma_t} = \frac{w_t(1 - w_t)}{\Sigma_t}.$$

Substituting, we have:

$$Cov_0[w_t, \mu_t] \approx (1 - w(e^*(z), \bar{\Sigma}))\, w(e^*(z), \bar{\Sigma}) \left\{ \left[2 - \chi(e^*(z))\right] \frac{1}{e^*(z)} \left.\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\right|_z Var_0[\mu_t] + \frac{1}{\bar{\Sigma}} Cov_0[\mu_t, \Sigma_t] \right\}. \tag{10}$$

Approximate $Var_0[\mu_t]$ around $\mu_{t-1} = z$, $\epsilon_{t-1} = 0$, and $\Sigma_{t-1} = \bar{\Sigma}$:[20]

$$\begin{aligned} Var_0[\mu_t] \approx & Var_0[(1 - w(e^*(z), \bar{\Sigma}))\mu_{t-1} + w(e^*(z), \bar{\Sigma})\sqrt{f(e^*(z))}\epsilon_{t-1}/e^*(z)] \\ = & (1 - w(e^*(z), \bar{\Sigma}))^2 Var_0[\mu_{t-1}] + w(e^*(z), \bar{\Sigma})^2 f(e^*(z))\sigma^2/e^*(z)^2. \end{aligned}$$

Approximate $Cov_0[\mu_t, \Sigma_t]$ around $\mu_{t-1} = z$ and $\Sigma_{t-1} = \bar{\Sigma}$:

$$\begin{aligned} Cov_0[\mu_t, \Sigma_t] \approx & Cov_0\Bigg[(1 - w(e^*(z), \bar{\Sigma}))\mu_{t-1}, \\ & \frac{\partial \Sigma_t}{\partial e_{t-1}}\bigg|_{(e^*(z),\bar{\Sigma})} \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z \mu_{t-1} + \frac{\partial \Sigma_t}{\partial \Sigma_{t-1}}\bigg|_{(e^*(z),\bar{\Sigma})} \Sigma_{t-1}\Bigg] \\ = & Cov_0\Bigg[(1 - w(e^*(z), \bar{\Sigma}))\mu_{t-1}, \\ & - \Sigma(e^*(z), \bar{\Sigma})^2 \frac{e^*(z)}{\sigma^2 g(e^*(z))}\left[2 - \chi(e^*(z))\right] \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z \mu_{t-1} + \left(\frac{\Sigma(e^*(z), \bar{\Sigma})}{\bar{\Sigma}}\right)^2 \Sigma_{t-1}\Bigg] \\ = & Cov_0\Bigg[(1 - w(e^*(z), \bar{\Sigma}))\mu_{t-1}, \\ & - \Sigma(e^*(z), \bar{\Sigma})\frac{w(e^*(z), \bar{\Sigma})}{e^*(z)}\left[2 - \chi(e^*(z))\right] \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z \mu_{t-1} + \left(\frac{\Sigma(e^*(z), \bar{\Sigma})}{\bar{\Sigma}}\right)^2 \Sigma_{t-1}\Bigg] \\ = & - (1 - w(e^*(z), \bar{\Sigma}))\Sigma(e^*(z), \bar{\Sigma})\, w(e^*(z), \bar{\Sigma})\left[2 - \chi(e^*(z))\right] \frac{1}{e^*(z)} \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z Var_0[\mu_{t-1}] \\ & + (1 - w(e^*(z), \bar{\Sigma}))\left(\frac{\Sigma(e^*(z), \bar{\Sigma})}{\bar{\Sigma}}\right)^2 Cov_0[\mu_{t-1}, \Sigma_{t-1}], \end{aligned}$$

where the second-to-last equality uses $\Sigma(e^*(z), \bar{\Sigma})\frac{e^*(z)^2}{\sigma^2 g(e^*(z))} = w(e^*(z), \bar{\Sigma})$. These last two approximations are good as long as σ^2 and Σ_0 are not too large. Substitute all back into

[20] A first-order Taylor expansion of μ_t around some $\bar{\mu}$, $\bar{\Sigma}$, and $\bar{\epsilon}$ yields:

$$\begin{aligned} Var_0[\mu_t] \approx & Var_0\Bigg[(1 - \bar{w})\bar{\mu} + \bar{w}z + \bar{w}\frac{\sqrt{f(\bar{e})}}{\bar{e}}\bar{\epsilon} + (1 - \bar{w})[\mu_{t-1} - \bar{\mu}] + \bar{w}\frac{\sqrt{f(\bar{e})}}{\bar{e}}[\epsilon_{t-1} - \bar{\epsilon}] \\ & + \frac{\partial w_{t-1}}{\partial e_{t-1}}\bigg|_{(e^*(\bar{\mu}),\bar{\Sigma})} \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_{\bar{\mu}} [\bar{\mu} - z][\mu_{t-1} - \bar{\mu}] + \frac{\partial \left[w_{t-1}\frac{\sqrt{f(e_{t-1})}}{e_{t-1}}\right]}{\partial e_{t-1}}\Bigg|_{(e^*(\bar{\mu}),\bar{\Sigma})} \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_{\bar{\mu}} \bar{\epsilon}\,[\mu_{t-1} - \bar{\mu}] \\ & + \frac{\partial w_{t-1}}{\partial \Sigma_{t-1}}\bigg|_{(e^*(\bar{\mu}),\bar{\Sigma})} \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_{\bar{\mu}} [\bar{\mu} - z][\Sigma_{t-1} - \bar{\Sigma}] + \frac{\sqrt{f(e_{t-1})}}{e_{t-1}} \frac{\partial w_{t-1}}{\partial \Sigma_{t-1}}\bigg|_{(e^*(\bar{\mu}),\bar{\Sigma})} \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_{\bar{\mu}} \bar{\epsilon}\,[\Sigma_{t-1} - \bar{\Sigma}]\Bigg], \end{aligned}$$

where $\bar{w}$ is short for $w(e^*(\bar{\mu}), \bar{\Sigma})$ and $\bar{e}$ is short for $e^*(\bar{\mu})$. Substituting for $\bar{\mu} = z$ and $\bar{\epsilon} = 0$ yields the expression used in the proof.

$Cov_0[\mu_t, w_t]$ and use $1 - w(e^*(z), \bar{\Sigma}) = \Sigma(e^*(z), \bar{\Sigma})/\bar{\Sigma}$:

$$\begin{aligned}
Cov_0[w_t, \mu_t] =&(1 - w(e^*(z), \bar{\Sigma}))\, w(e^*(z), \bar{\Sigma}) \\
&\Bigg\{ [2 - \chi(e^*(z))]\, \frac{1}{e^*(z)}\, \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z \\
&\quad \left[(1 - w(e^*(z), \bar{\Sigma}))^2 Var_0[\mu_{t-1}] + w(e^*(z), \bar{\Sigma})^2 f(e^*(z))\sigma^2/e^*(z)^2\right] \\
&- (1 - w(e^*(z), \bar{\Sigma}))^2 w(e^*(z), \bar{\Sigma})\, [2 - \chi(e^*(z))]\, \frac{1}{e^*(z)}\, \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z Var_0[\mu_{t-1}] \\
&+ (1 - w(e^*(z), \bar{\Sigma}))^4 \frac{1}{\bar{\Sigma}} Cov_0[\mu_{t-1}, \Sigma_{t-1}] \Bigg\}.
\end{aligned} \tag{11}$$

The induction hypothesis that $Cov_0[\mu_{t-1}, w_{t-1}] \geq 0$ implies, from the time $t-1$ analogue of equation (10), that

$$\frac{1}{\bar{\Sigma}} Cov_0[\mu_{t-1}, \Sigma_{t-1}] \geq -\,[2 - \chi(e^*(z))]\, \frac{1}{e^*(z)}\, \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z Var_0[\mu_{t-1}].$$

We therefore have:

$$\begin{aligned}
\frac{Cov_0[w_t, \mu_t]}{(1 - w(e^*(z), \bar{\Sigma}))\, w(e^*(z), \bar{\Sigma})} \geq& [2 - \chi(e^*(z))]\, \frac{1}{e^*(z)}\, \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z \\
&\left[(1 - w(e^*(z), \bar{\Sigma}))^2 Var_0[\mu_{t-1}] + w(e^*(z), \bar{\Sigma})^2 f(e(e^*(z), \bar{\Sigma}))\sigma^2/e^*(z)^2\right] \\
&- (1 - w(e^*(z), \bar{\Sigma}))^2 w(e^*(z), \bar{\Sigma})\, [2 - \chi(e^*(z))]\, \frac{1}{e^*(z)}\, \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z Var_0[\mu_{t-1}] \\
&- (1 - w(e^*(z), \bar{\Sigma}))^4\, [2 - \chi(e^*(z))]\, \frac{1}{e^*(z)}\, \frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\bigg|_z Var_0[\mu_{t-1}].
\end{aligned} \tag{12}$$

Using $\chi(e^*(z)) < 2$, we have that $Cov_0[w_t, \mu_t] \geq 0$ if

$$\begin{aligned}
0 \leq& (1 - w(e^*(z), \bar{\Sigma}))^2 Var_0[\mu_{t-1}] + w(e^*(z), \bar{\Sigma})^2 f(e^*(z))\sigma^2/e^*(z)^2 \\
&- (1 - w(e^*(z), \bar{\Sigma}))^2 w(e^*(z), \bar{\Sigma}) Var_0[\mu_{t-1}] \\
&- (1 - w(e^*(z), \bar{\Sigma}))^4 Var_0[\mu_{t-1}] \\
\Leftrightarrow 0 \leq& w(e^*(z), \bar{\Sigma}) f(e^*(z))\sigma^2/e^*(z)^2 - (1 - w(e^*(z), \bar{\Sigma}))^3 Var_0[\mu_{t-1}].
\end{aligned} \tag{13}$$

Repeatedly approximating around $\mu_s = z$, $\epsilon_s = 0$, and w_s evaluated at $e_s = e^*(z)$ and $\Sigma_s = \bar{\Sigma}$, we have:

$$Var_0[\mu_{t-1}] \approx \frac{f(e^*(z))\sigma^2}{e^*(z)^2} \sum_{i=1}^{t-2} w(e^*(z), \bar{\Sigma})^2 \prod_{i+1}^{t-2} (1 - w(e^*(z), \bar{\Sigma}))^2.$$

Substituting into inequality (13), we have that $Cov_0[w_t, \mu_t] \geq 0$ if:

$$
\begin{aligned}
0 \leq & w(e^*(z), \bar{\Sigma}) \frac{f(e^*(z))\sigma^2}{e^*(z)^2} - (1 - w(e^*(z), \bar{\Sigma}))^3 \frac{f(e^*(Z))\sigma^2}{e^*(z)^2} \sum_{i=1}^{t-2} \left[w(e^*(z), \bar{\Sigma})^2 \prod_{i+1}^{t-2} (1 - w(e^*(z), \bar{\Sigma}))^2 \right] \\
\Leftrightarrow 0 \leq & w(e^*(z), \bar{\Sigma}) - (1 - w(e^*(z), \bar{\Sigma}))^3 \sum_{i=1}^{t-2} \left[w(e^*(z), \bar{\Sigma})^2 \prod_{i+1}^{t-2} (1 - w(e^*(z), \bar{\Sigma}))^2 \right].
\end{aligned}
$$

The final term is a geometric series with common ratio $(1 - w(e^*(z), \bar{\Sigma}))^2 < 1$. Using the fact that the value to which this series converges is maximized as $t \to \infty$, we have:

$$
\begin{aligned}
(1 - w(e^*(z), \bar{\Sigma}))^3 \sum_{i=1}^{t-2} \left[w(e^*(z), \bar{\Sigma})^2 \prod_{i+1}^{t-2} (1 - w(e^*(z), \bar{\Sigma}))^2 \right] \leq & \frac{(1 - w(e^*(z), \bar{\Sigma}))^3}{1 - (1 - w(e^*(z), \bar{\Sigma}))^2} w(e^*(z), \bar{\Sigma})^2 \\
= & \frac{(1 - w(e^*(z), \bar{\Sigma}))^3}{2 - w(e^*(z), \bar{\Sigma})} w(e^*(z), \bar{\Sigma}).
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
& w(e^*(z), \bar{\Sigma}) - (1 - w(e^*(z), \bar{\Sigma}))^3 \sum_{i=1}^{t-2} \left[w(e^*(z), \bar{\Sigma})^2 \prod_{i+1}^{t-2} (1 - w(e^*(z), \bar{\Sigma}))^2 \right] \\
\geq & w(e^*(z), \bar{\Sigma}) \left(1 - \frac{(1 - w(e^*(z), \bar{\Sigma}))^3}{2 - w(e^*(z), \bar{\Sigma})} \right) \\
\geq & 0.
\end{aligned}
$$

Inequality (13) therefore holds. Therefore, $Cov_0[w_t, \mu_t] \geq 0$ under the induction hypothesis that $Cov_0[w_{t-1}, \mu_{t-1}] \geq 0$.

Basis step:

We have

$$
\mu_1 = (1 - w_0)\mu_0 + w_0 \left[z + \frac{\sqrt{f(e_0)}}{e_0} \epsilon_0 \right], \text{ with } w_0 = \frac{\frac{e_0^2}{g(e_0)\,\sigma^2}}{\Sigma_0^{-1} + \frac{e_0^2}{g(e_0)\,\sigma^2}}.
$$

Note that w_0 is not random. μ_1 depends on only one random variable (ϵ_0), which is itself normal and enters linearly. Therefore μ_1 is normally distributed. By Stein's Lemma, we have:

$$
Cov_0[w_1, \mu_1] = Var_0(\mu_1) E_0 \left[\frac{\partial w_1}{\partial e_1} \frac{\mathrm{d} e_1}{\mathrm{d} \mu_1} \right]. \tag{14}
$$

Differentiating w_1, we have:

$$\frac{\partial w_1}{\partial e_1} = (1-w_1)w_1\left[2-\chi(e_1)\right]\frac{1}{e_1}, \tag{15}$$

which is strictly positive if and only if $\chi(e_1) < 2$. For σ^2 not too large,

$$E_0\left[\frac{\partial w_1}{\partial e_1}\frac{\mathrm{d}e_1}{\mathrm{d}\mu_1}\right] \approx \left.\frac{\partial w_1}{\partial e_1}\right|_{(e^*(z),\Sigma_1)} \left.\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\right|_z .$$

The assumption that $\chi(e^*(z)) < 2$ then implies $Cov_0[w_1, \mu_1] \geq 0$.

Part 2

Assume that $\chi(e^*(z)) > 2$. We again proceed by induction.

Induction step:
The induction hypothesis is that $Cov_0[w_{t-1}, \mu_{t-1}] \leq 0$ for some $t \geq 2$. Following the proof of the first part of the lemma and using the induction hypothesis that $Cov_0[w_{t-1}, \mu_{t-1}] \leq 0$, we find that inequality (12) becomes

$$\begin{aligned}\frac{Cov_0[w_t, \mu_t]}{(1-w(e^*(z),\bar{\Sigma}))\,w(e^*(z),\bar{\Sigma})} \leq & \left[2-\chi(e^*(z))\right]\frac{1}{e^*(z)}\left.\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\right|_z \\ & \left[(1-w(e^*(z),\bar{\Sigma}))^2 Var_0[\mu_{t-1}] + w(e^*(z),\bar{\Sigma})^2 f(e(e^*(z),\bar{\Sigma}))\sigma^2/e^*(z)^2\right] \\ & -(1-w(e^*(z),\bar{\Sigma}))^2 w(e^*(z),\bar{\Sigma})\left[2-\chi(e^*(z))\right]\frac{1}{e^*(z)}\left.\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\right|_z Var_0[\mu_{t-1}] \\ & -(1-w(e^*(z),\bar{\Sigma}))^4\left[2-\chi(e^*(z))\right]\frac{1}{e^*(z)}\left.\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\right|_z Var_0[\mu_{t-1}].\end{aligned}$$

Using $\chi(e^*(z)) > 2$, we have that $Cov_0[w_t, \mu_t] \leq 0$ if

$$0 \leq w(e^*(z),\bar{\Sigma}) f(e^*(z))\sigma^2/e^*(z)^2 - (1-w(e^*(z),\bar{\Sigma}))^3 Var_0[\mu_{t-1}].$$

The proof that this inequality holds is the same as the proof that inequality (13) holds in the first part of the lemma. Therefore $Cov_0[w_t, \mu_t] \leq 0$ under the induction hypothesis that $Cov_0[w_{t-1}, \mu_{t-1}] \leq 0$.

Basis step:
Following the basis step in the proof of the first part of the lemma, it is easy to show that $Cov_0[w_1, \mu_1] \leq 0$ for σ^2 not too large.

Part 3

Assume that $\chi(e^*(z)) = 2$. We again proceed by induction.

Induction step:
The induction hypothesis is that $Cov_0[w_{t-1}, \mu_{t-1}] \approx 0$ for some $t \geq 2$. Following the proof of the first part of the lemma, equation (11) becomes:

$$Cov_0[w_t, \mu_t] = (1 - w(e^*(z), \bar{\Sigma}))w(e^*(z), \bar{\Sigma})(1 - w(e^*(z), \bar{\Sigma}))^4 \frac{1}{\bar{\Sigma}} Cov_0[\mu_{t-1}, \Sigma_{t-1}].$$

The induction hypothesis that $Cov_0[\mu_{t-1}, w_{t-1}] \approx 0$ implies, from the time $t-1$ analogue of equation (10), that

$$\frac{1}{\bar{\Sigma}} Cov_0[\mu_{t-1}, \Sigma_{t-1}] \approx -[2 - \chi(e^*(z))] \frac{1}{e^*(z)} \left.\frac{\mathrm{d}e^*(\mu)}{\mathrm{d}\mu}\right|_z Var_0[\mu_{t-1}],$$

which equals zero under the assumption that $\chi(e^*(z)) = 2$. Therefore $Cov_0[w_t, \mu_t] \approx 0$ under the induction hypothesis that $Cov_0[w_{t-1}, \mu_{t-1}] \approx 0$.

Basis step:
Following the basis step in the proof of the first part of the lemma, it is easy to show that $Cov_0[w_1, \mu_1] \approx 0$ for σ^2 not too large.

Part 4

Let time t effort be fixed exogenously at $\bar{e}_t$. From equation (5), Σ evolves deterministically. Then, from equation (6), w_{t-1} is independent of random variables. Thus $Cov_0[w_t, \mu_t] = 0$.

Proof of Proposition 1

Part 1

Note that

$$E_0[\mu_1] = (1 - w_0)\mu_0 + w_0 z = \mu_0 + w_0(z - \mu_0) = \mu_0 = z, \tag{16}$$

where we recognize that $\mu_0 = z$ by assumption. This establishes the first part of the proposition.

Part 2

Assume that $\chi(e^*(z)) < 2$.

Induction step:
The induction hypothesis is that $E_0[\mu_t] < z$ for some $t > 1$. We have the following by the tower property:

$$\begin{aligned}E_0[\mu_{t+1}] =& E_0[E_0[\mu_{t+1}|\mu_t]]\\ =& E_0[\mu_t + w_t(z - \mu_t)]\\ =& E_0[w_t]z + (1 - E_0[w_t])E_0[\mu_t] - Cov_0[w_t, \mu_t]. \end{aligned} \tag{17}$$

From the induction hypothesis and the definition of w_t, we have $E_0[\mu_{t+1}] < z$ if $Cov_0[w_t, \mu_t] \geq 0$. The first part of Lemma 2 established that, in fact, $Cov_0[w_t, \mu_t] \geq 0$ for $t \geq 1$. Therefore $E_0[\mu_{t+1}] < z$ under the induction hypothesis that $E_0[\mu_t] < z$.

Basis step:
From equation (16), we have $E_0[\mu_1] = z$. From equations (14) and (15), we have $Cov_0[w_1, \mu_1] > 0$. Therefore, from equation (17), $E_0[\mu_2] < z$.

Part 3

Assume that $\chi(e^*(z)) > 2$.

Induction step:
The induction hypothesis is that $E_0[\mu_t] > z$ for some $t > 1$. From equation (17), the induction hypothesis, and the definition of w_t, we have $E_0[\mu_{t+1}] > z$ if $Cov_0[w_t, \mu_t] \leq 0$. The second part of Lemma 2 established that, in fact, $Cov_0[w_t, \mu_t] \leq 0$ for $t \geq 1$. Therefore $E_0[\mu_{t+1}] > z$ under the induction hypothesis that $E_0[\mu_t] > z$.

Basis step:
From equation (16), we have $E_0[\mu_1] = z$. From equations (14) and (15), we have $Cov_0[w_1, \mu_1] < 0$. Therefore, from equation (17), $E_0[\mu_2] > z$.

Part 4

Assume that $\chi(e^*(z)) = 2$.

Induction step:
The induction hypothesis is that $E_0[\mu_t] \approx z$ for some $t > 1$. From equation (17), the induction hypothesis, and the definition of w_t, we have $E_0[\mu_{t+1}] \approx z$ if $Cov_0[w_t, \mu_t] \approx 0$. The third part of Lemma 2 established that, in fact, $Cov_0[w_t, \mu_t] \approx 0$ for $t \geq 1$. Therefore $E_0[\mu_{t+1}] \approx z$

under the induction hypothesis that $E_0[\mu_t] \approx z$.

Basis step:
From equation (16), we have $E_0[\mu_1] = z$. From equations (14) and (15), we have $Cov_0[w_1, \mu_1] \approx 0$. Therefore, from equation (17), $E_0[\mu_2] \approx z$.

Part 5

Let time t effort be fixed exogenously at $\bar{e}_t$.

Induction step:
The induction hypothesis is that $E_0[\mu_t] = z$ for some $t > 1$. From equation (17), the induction hypothesis, and the definition of w_t, we have $E_0[\mu_{t+1}] = z$ if $Cov_0[w_t, \mu_t] = 0$. The fourth part of Lemma 2 established that, in fact, $Cov_0[w_t, \mu_t] = 0$ for $t \geq 1$. Therefore $E_0[\mu_{t+1}] = z$ under the induction hypothesis that $E_0[\mu_t] = z$.

Basis step:
From equation (16), we have $E_0[\mu_1] = z$. From equation (14), we have $Cov_0[w_1, \mu_1] = 0$. Therefore, from equation (17), $E_0[\mu_2] = z$.

References

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

Bandura, Albert (1982) "Self-efficacy mechanism in human agency," *American Psychologist*, Vol. 37, No. 2, pp. 122–147.

——— (2001) "Social cognitive theory: An agentic perspective," *Annual Review of Psychology*, Vol. 52, No. 1, pp. 1–26.

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

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

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

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

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

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

Butler, Andrew C., Jason E. Chapman, Evan M. Forman, and Aaron T. Beck (2006) "The empirical status of cognitive-behavioral therapy: A review of meta-analyses," *Clinical Psychology Review*, Vol. 26, No. 1, pp. 17–31.

Choi, Taeryon and R. V. Ramamoorthi (2008) "Remarks on consistency of posterior distributions," in *Pushing the Limits of Contemporary Statistics: Contributions in Honor of Jayanta K. Ghosh*, Vol. 3: Institute of Mathematical Statistics, pp. 170–186.

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

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

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

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

DeRubeis, Robert J., Greg J. Siegle, and Steven D. Hollon (2008) "Cognitive therapy vs. medications for depression: Treatment outcomes and neural mechanisms," *Nature Reviews Neuroscience*, Vol. 9, No. 10, pp. 788–796.

Dobson, Keith S., Steven D. Hollon, Sona Dimidjian, Karen B. Schmaling, Robert J. Kohlenberg, Robert Gallop, Shireen L. Rizvi, Jackie K. Gollan, David L. Dunner, and Neil S. Jacobson (2008) "Randomized trial of behavioral activation, cognitive therapy, and antidepressant medication in the prevention of relapse and recurrence in major depression," *Journal of Consulting and Clinical Psychology*, Vol. 76, No. 3, pp. 468–477.

Fudenberg, Drew and David K. Levine (2006) "A dual-self model of impulse control," *The American Economic Review*, Vol. 96, No. 5, pp. 1449–1476.

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

Gelman, Andrew, John B. Carlin, Hal S. Stern, and Donald B. Rubin (2004) *Bayesian Data Analysis*, Boca Raton, Fla: Chapman & Hall/CRC, 2nd edition.

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

Gneezy, Uri, John A. List, Jeffrey A. Livingston, Sally Sadoff, Xiangdong Qin, and Yang Xu (2017) "Measuring success in education: The role of effort on the test itself," Working Paper 24004, National Bureau of Economic Research.

Greenberg, Paul E., Andree-Anne Fournier, Tammy Sisitsky, Crystal T. Pike, and Ronald C. Kessler (2015) "The economic burden of adults with major depressive disorder in the United States (2005 and 2010)," *The Journal of Clinical Psychiatry*, Vol. 76, No. 2, pp. 155–162.

Heidhues, Paul, Botond Koszegi, and Philipp Strack (2017) "Unrealistic expectations and misguided learning," working paper.

Hestermann, Nina and Yves Le Yaouanq (2016) "It's not my fault! Ego, excuses and perseverance," working paper.

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

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

Hollon, Steven D., Robert J. DeRubeis, Richard C. Shelton, Jay D. Amsterdam, Ronald M. Salomon, John P. O'Reardon, Margaret L. Lovett, Paula R. Young, Kirsten L. Haman, Brent B. Freeman, and Robert Gallop (2005) "Prevention of relapse following cognitive therapy vs medications in moderate to severe depression," *Archives of General Psychiatry*, Vol. 62, No. 4, pp. 417–422.

Jehiel, Philippe (2016) "Investment strategy and selection bias: An equilibrium perspective on overoptimism," working paper.

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

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

Kreps, David (1998) "Anticipated utility and dynamic choice," in Donald P. Jacobs, Ehud Kalai, and Morton Kamien eds. *Frontiers of Research in Economic Theory*, Cambridge, England: Cambridge University Press.

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

Levitt, Steven D., John A. List, Susanne Neckermann, and Sally Sadoff (2016) "The behavioralist goes to school: Leveraging behavioral economics to improve educational performance," *American Economic Journal: Economic Policy*, Vol. 8, No. 4, pp. 183–219.

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

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

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

Rabin, Matthew and Joel L. Schrag (1999) "First impressions matter: A model of confirmatory bias," *The Quarterly Journal of Economics*, Vol. 114, No. 1, pp. 37–82.

Santos-Pinto, Luís and Joel Sobel (2005) "A model of positive self-image in subjective assessments," *American Economic Review*, Vol. 95, No. 5, pp. 1386–1402.

Scott, Jan (1996) "Cognitive therapy of affective disorders: a review," *Journal of Affective Disorders*, Vol. 37, No. 1, pp. 1–11.

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

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

Stajkovic, Alexander D. and Fred Luthans (1998) "Social cognitive theory and self-efficacy: Going beyond traditional motivational and behavioral approaches," *Organizational Dynamics*, Vol. 26, No. 4, pp. 62–74.

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

Tenney, Elizabeth R., Jennifer M. Logg, and Don A. Moore (2015) "(Too) optimistic about optimism: The belief that optimism improves performance," *Journal of Personality and Social Psychology*, Vol. 108, No. 3, pp. 377–399.

Thaler, Richard H. and H. M. Shefrin (1981) "An economic theory of self-control," *Journal of Political Economy*, Vol. 89, No. 2, pp. 392–406.

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

Wood, Robert and Albert Bandura (1989) "Social cognitive theory of organizational management," *Academy of Management Review*, Vol. 14, No. 3, pp. 361–384.

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