

HOPE FOR THE BEST, PLAN FOR THE WORST

KENNETH KASA AND XIAOWEN LEI

ABSTRACT. This paper studies asset pricing when individuals struggle to strike a balance between doubt and hope. We argue this internal struggle is consistent with recent evidence from neuroscience. We operationalize it using the robust control and filtering approach of Hansen and Sargent (2008). Our key innovation is to assume that filtering is *optimistically* biased. Investment decisions, however, reflect doubts about model specification, and are pessimistically biased. We show that doubts about model specification, combined with optimistically distorted beliefs about dividend growth can not only explain low average price/dividend ratios and high average returns, but can also generate the sort of large procyclical swings in price/dividend ratios that are observed in the data. High and volatile returns occur despite the fact that investors have low (approximately logarithmic) risk aversion. The model's belief distortions are empirically plausible, with detection error probabilities between 10-20%.

JEL Classification Numbers: G12, D81

I am large, I contain multitudes. –

WALT WHITMAN (*Song of Myself*, 1855)

1. INTRODUCTION

The 1980s and 90s were a dismal period for asset pricing research. The empirical flaws of standard models highlighted by Shiller (1981), Hansen and Singleton (1983) and Mehra and Prescott (1985) were remarkably resilient to theoretical modifications. Fixing one problem invariably exacerbated other problems. Economists could not even explain the unconditional first moments of risky versus safe assets.


Since then the asset pricing literature has advanced significantly. The equity premium and risk free rate puzzles are no longer puzzles. If you want to know why mean returns on risky assets are so high while returns on safe assets are so low there are a menu of model options available. There are theories based on ‘exotic preferences’, e.g., the backward-looking habit persistence model of Campbell and Cochrane (1999) or the forward-looking long-run risk model of Bansal and Yaron (2004). There are models based on even more exotic preferences, e.g., Prospect Theory (Benartzi and Thaler (1995)). Alternatively, there are theories based on beliefs, which deviate from the Rational Expectations Hypothesis, e.g., Barillas, Hansen, and Sargent (2009) or Ju and Miao (2012). Other theories add frictions, leading to incomplete markets. Although initially this seemed like the obvious candidate, it proved to be surprisingly difficult to implement empirically. While Constantinides and Duffie (1996) showed that highly persistent idiosyncratic labor income shocks

Date: January, 2026.

We would like to thank Sev Hou for helpful comments.

could in principle explain the equity premium, data on labor income didn't seem to support this. However, recent work by Constantinides and Ghosh (2017) obtains promising results by allowing for 'rare disasters' in idiosyncratic labor income, as in Barro (2006).

Given these recent successes, the asset pricing literature has become more ambitious, and is now focused on a broader spectrum of data, e.g., the *conditional* moments of asset returns and their comovements with other observed data. Figure 1 in particular remains a challenge. It depicts (annualized) monthly data on the Price/Earnings ratio in the USA for the period 1890-2023 (obtained from Shiller's website).



The flat line is the unconditional mean, around 17.5. Its level reflects of the relatively high mean return on the stock market. Taking the reciprocal as an estimate of the dividend yield, and adding 2% as the mean growth rate of dividends yields a 7-8% (annualized, real) return. What is most striking about Figure 1, however, is the volatility. The PE rose above 30 during the 'roaring 20s' before it collapsed at the outset of the Great Depression. It rose even higher, to above 40, at the peak of the 'dotcom boom' of the 1990s. Conversely, it fell into the single digits during the bear markets of the early 1900s and 1970s. This volatility was the original inspiration for Shiller's (1981) classic paper, which showed that these fluctuations could not be explained by revisions of rational expectations of future dividend growth. Shiller argued that bubbles provide the most likely explanation. Other researchers focus on the possibility of a strongly countercyclical risk premium. No consensus has yet emerged on either possibility.

Our paper proposes a new possibility. It is not based on bubbles or risk aversion. Nor is it based on incomplete markets and idiosyncratic risk. Instead, it is based on beliefs. It is perhaps most closely related to Hansen and Sargent (2010). Hansen and Sargent were responding to criticisms of the Bansal-Yaron (2004) long-run risks model. For example, Beeler and Campbell (2012) noted that empirical evidence in support of the model's two key ingredients, i.e., stochastic growth and stochastic volatility, is rather weak. Hansen and Sargent (2010) introduce model uncertainty, and show that it is enough that agents *suspect* that growth rates *might* drift, even if in fact they do not. They are also able to generate apparently strong countercyclical fluctuations in risk prices by incorporating robust model averaging. Agents think that bad times will be persistent, whereas good times will be transitory. However, in Hansen and Sargent (2010) these fluctuations do not

represent changes in the willingness of agents to bear risk. Instead, they represent changes in the price of model uncertainty.

We make one small, but important, change to Hansen and Sargent’s Fragile Beliefs model. Instead of supposing that agents revise their beliefs in a defensively pessimistic manner, we assume beliefs are distorted *optimistically*. Operationally, this is accomplished by simply ‘flipping the sign’ of Hansen and Sargent’s robust filtering parameter. Now agents think good times will persist, while bad times are transitory. This generates the sort of prolonged booms and sharp sudden crashes that are observed in the data (see Figure 1). To generate procyclical movements in the PE ratio, we suppose relative risk aversion is slightly less than one.¹ In some respects, our model resembles the ‘intrinsic bubbles’ model of Froot and Obstfeld (1991). However, in our model transversality conditions are obeyed and we do not confront the uncomfortable question of how agents’ beliefs coordinate on a bubble. Even more important, we maintain Hansen and Sargent’s assumption that agents doubt the specification of their dividend growth models, and so introduce *pessimistic* drift distortions into their models. This allows us to generate a low mean PE ratio (and hence high mean equity returns), while at the same time generating high volatility and procyclical movements in the PE ratio.

Anytime a model is based on distorted, ‘nonrational’ beliefs, it is important to demonstrate that such distortions are plausible, and do not leave too many \$5 bills on the sidewalk. We respond to this challenge in two ways. First, following Hansen and Sargent (2008), we argue that a bit of defensive pessimism is *not* irrational. It is simply an acknowledgment of the limits of your own knowledge. As such, defensive pessimism could even be interpreted as a form of ‘higher-order’ rationality. The real question is why agents do not respond to model uncertainty in a Bayesian manner. For that, we refer readers to the detailed discussion in Hansen and Sargent (2008). Here we simply note that the founding father of Bayesian decision theory, Savage (1972), cautioned against the misapplication of ‘small worlds’ Bayesian methods to ‘large worlds’ problems (see, e.g., p. 16). We claim that the problem of investing in the stock market and placing value on new technologies is a large worlds problem. Even if one grants us our nonBayesian approach, we still confront the practical question of *quantifying* the degrees of optimism and pessimism. For that we again follow Hansen and Sargent (2008) by calibrating our belief parameters to ‘detection error probabilities’, so that agents do not base their forecasts and investment decisions on empirically implausible models of dividend growth.

Our second response is to argue that the distortions we employ are consistent with evidence from the psychology literature, and more recently, from the neuroscience literature. The novel aspect of our model is that agents are being *simultaneously* optimistic *and* pessimistic. They base dividend forecasts on optimistically biased estimates of the underlying unobserved growth state. At the same time, they base investment decisions on a pessimistically biased model of the process generating observed dividends. Isn’t this schizophrenic, or at least bipolar?² The key to reconciling this apparent contradiction is to note that the problems of forming and revising beliefs about the future on the one

¹If risk aversion exceeded one, then the higher discount rate associated with good news about future dividends would dominate the higher expected cash flows, and the PE ratio would be countercyclical.

²Kasa (2012) showed how in certain cases these two biases could exactly offset each other and produce what looks like, to an outside observer, Rational Expectations.

hand, and taking actions based on those beliefs on the other, involve different cognitive processes, and take place (primarily but not exclusively) in different parts of the brain. As discussed below, neuroscientific evidence suggests that learning and belief revision takes place in the left hemisphere of the brain, and is optimistically biased, in the sense that positive signals elicit over reactions while negative signals elicit under-reaction. This bias is tempered, however, by the right hemisphere, which is regarded as the locus of a person's 'alert system', which produces cautionary, pessimistically biased responses to the information being transmitted from the left hemisphere. Of course, the Pop Science notion that the world is divided into left-brain people, who are analytical, and right-brain people, who are emotional, is a vast oversimplification. The fact is, unless afflicted by a stroke or some other brain injury, everyone is both a left and right-brained person, and the two halves work together in intricate ways to optimize an individual's response to an ever changing environment. Still, brain lateralization *is* real, and should not be too surprising to economists, who are attuned to the benefits of specialization and the gains from trade.

By relating stock market dynamics to nonrational belief dynamics our paper is consistent with recent work by Bordalo, Gennaioli, Porta, and Shleifer (2024). They directly measure expectations using survey data, and argue that it is important to distinguish between short-term and long-term growth forecasts. Bordalo et. al. (2024) show that long-term earnings growth forecasts overreact to good news, and that this overreaction can explain many asset pricing puzzles, both in aggregate time series data and in firm level cross-sectional data. There are three key differences between our paper and theirs. First, and most importantly, we provide an explicit 'micro-foundation' for belief overreaction, based on recent advances in neuroscience. Second, we argue that a form of persistent pessimism is also important, which allows our model to also explain high average returns. Third, we validate the empirical magnitudes of our belief distortions using detection error probabilities, rather than using survey data. This allows us to consider a much longer sample period, which can be important when focusing on low frequency distortions. Despite these differences, we view our work as quite complementary to theirs.

The remainder of the paper is organized as follows. Section 2 reviews evidence from the psychology and neuroscience literatures on the co-existence of optimism and pessimism. As a benchmark, Section 3 computes asset prices when agents are Bayesians. Here we rediscover the equity premium and volatility puzzles. Asset prices are too high (and mean returns too low), and are nearly unresponsive to revisions in beliefs about future dividend growth. Section 4 presents our main results. We show how drift distortions in the dividend growth model and Kalman filtering equation can be used to operationalize and quantify the internal tensions between optimism and pessimism that agents struggle to balance. Now prices become much more responsive to belief revisions, and come close to matching the volatility of the observed Price/Earnings ratio displayed in Figure 1. Section 5 shows that the belief distortions we introduce are empirically plausible, in the sense that detection error probabilities are not too low. Finally, Section 6 offers concluding remarks, and a technical appendix contains proofs and mathematical details.

2. OPTIMISM & PESSIMISM: EVIDENCE FROM PSYCHOLOGY AND NEUROSCIENCE

The central premise of this paper is that individuals confront a tension between optimism and pessimism, and that this tension manifests itself differently in different decisions. We start by briefly reviewing evidence in support of this tension, coming first from the psychology literature, and then more recently from the neuroscience literature.

2.1. Psychology. We are certainly not the first persons to suggest that individuals balance optimism and pessimism. The homespun wisdom of our paper’s title is so old that no one knows who first said it. We also would bet that references to it could be found in Shakespeare, if one cared to look. However, the first formal, scientific study that we are aware of is Kahneman and Tversky (1979). They argue that the combination of ‘bold forecasts’ and ‘timid actions’ derives from a common underlying psychological bias, namely, ‘narrow framing’. Narrow framing refers to the tendency to view problems as more unique than they really are. Individuals become excessively risk averse because they do not adequately account for the risk reducing benefits of diversification and the law of large numbers. In the context of forecasting, Kahneman and Tversky refer to narrow framing as the ‘inside view’, which anchors predictions of exogenous future events on your own plans and scenarios, thus creating an illusion of excessive control and optimism.

More recently, Haselton and Nettle (2006) adopt an evolutionary perspective. They appeal to something called ‘Error Management Theory’, and argue that if the costs of false positives and false negatives have been asymmetric over the course of evolutionary history, then natural selection could produce a population of ‘paranoid optimists’. The provocative example they use to motivate their analysis is the case of ‘sexual inference’, i.e., the tendency of men to overestimate a woman’s sexual interest. If the cost of a missed opportunity to spread your genes exceeds the cost of a slap in the face, then men might develop an optimism bias when gauging their own chances. Although this theory does not directly predict that individuals display optimism when forecasting, but pessimism when investing, it could be made to generate such behavior under plausible conditions. In particular, if the cost of being too optimistic when deciding on a portfolio allocation exceeds the cost of being too pessimistic, then agents might appear to be excessively risk averse. This is in fact precisely the theory behind ‘loss aversion’ and Prospect Theory. It could even arise with standard preferences in the presence of borrowing constraints or subsistence levels. Conversely, conditional on participation, agents might be optimistically biased when forecasting if the costs of pessimism about future returns exceed the costs of being too optimistic. This combination might occur in a ‘Gold Rush’ or ‘Superstars’ economy, which produces small probabilities of huge fortunes.

At a more general level, our model of Hoping For the Best While Planning for the Worst is consistent with a common thread running through much of the psychology literature, namely, the view that individuals are composites of competing, or duelling selves. Kahneman’s (2011) well known book, *Thinking, Fast and Slow*, contains dozens of examples. This composite view of individuals has even become common within economics, examples being the hyperbolic discounting model of Laibson (1997) and Hansen and Sargent’s (2008) model of ‘evil agents’ and robust control. Interestingly, the notion that individuals are composites of competing interests is also consistent with evidence from the rapidly developing field of neuroscience. We turn to this evidence next.

2.2. Neuroscience. Evidence of brain lateralization emerged somewhat by chance. Epileptic seizures are abnormal electrical discharges that spread from one half of the brain to the other. Currently, there are effective medications that inhibit these discharges. This wasn't the case back in the 1960s, however. A neurologist named Roger Sperry proposed the idea of cutting the corpus callosum (a tract of nerve fibers separating the two halves of the brain) in an effort to contain and mitigate the electrical discharges causing epileptic seizures. The procedure turned out to be a success, and Sperry was awarded the Nobel Prize in 1981.

Over time, however, as these split-brain patients continued to be tested and monitored, some subtle changes began to be detected. For example, patients would often move their heads when processing visual information, or make symbolic hand movements. Using the fact that human vision is 'cross-wired', with the left hemisphere processing information from the right visual field (and vice versa), experiments were run where different kinds of information were shown to first one eye and then the other. It was discovered that when a picture of a face was shown to the right eye (and so processed by the left-side of the brain) it was unrecognized, but was easily recognized when shown to the left eye. Conversely, when a passage of text was shown to the left eye, the patient could not read it, but had no trouble when it was shown to the right eye. Although the idea that language and speech are processed in the left-side of the brain can be traced back to 1800s, these studies were the first solid evidence in support of it.³

For the purposes of our paper, the most interesting of these kinds of studies were performed by Wolford, Miller, and Gazzaniga (2000). They engaged subjects in a probability-guessing experiment. In these experiments subjects are shown random sequences of two outcomes, e.g., a red or green light appears on their computer screen. The probabilities of the two outcomes are different, e.g., red appears with probability $3/4$ while green appears with probability $1/4$. Subjects are rewarded based on the number of correct guesses. It has long been known that humans engage in 'frequency matching'. In the above example, they guess red about 75% of the time. This is clearly suboptimal. To maximize your payoff, you should *always* guess red. Interestingly, nonhumans (e.g., rats) usually optimize. Inspired by the earlier research on split-brain patients, Wolford et. al. added a new twist to this classic experiment. They played the game separately on the left- and right-hemispheres (by displaying the events to only the right and left eyes, respectively). Interestingly, they found that the left hemisphere engages in frequency matching, whereas the right hemisphere optimizes.

In his popular book *Human: The Science Behind What Makes Your Brain Unique*, Gazzaniga (2008), who is a student of Sperry, summarizes the left/right brain research as follows:

This is consistent with the hypothesis that the left-hemisphere interpreter constructs theories to assimilate perceived information into a comprehensible whole. By going beyond simply observing events to asking why they happened, a brain can cope with such events more effectively if they happen again...Accuracy remains high in the right hemisphere, however, because it

³The idea first arose from autopsies on people with speech impediments. People with speech impediments were observed to have damage in the left-side of their brains.

does not engage in these interpretive processes. The advantage of having such a dual system is obvious. The right hemisphere maintains an accurate record of events, leaving the left hemisphere free to elaborate and make inferences about the material presented. (Gazzaniga (2008), p. 296)

What, you might ask, does this have to do with stock market investment? Investing involves *two* conceptually different activities. One must first formulate some belief about the future consequences of any given investment. This involves developing theories and hypotheses about how the future may unfold. This is a creative process, which is akin to engaging in frequency matching. We claim this problem is primarily processed in the left-hemisphere, and is optimistically biased. Second, an individual must decide *how much* of his scarce capital he should devote to the left-hemisphere's theories. We claim this investment allocation decision is allocated to the alert system in the right hemisphere, and is pessimistically biased. How do we know the left-hemisphere is optimistically biased while the right hemisphere is pessimistic? Much of the rapid progress made in neuroscience during the past 20 years derives from the development of new technologies that allow scientists to *directly* measure and isolate brain activity. Scientists no longer have to rely on autopsies, illnesses, or accidents. These recent studies provide strong evidence that indeed the left hemisphere is optimistic while the right hemisphere is pessimistic.

The most common of these new technologies is Functional Magnetic Resonance Imaging, or fMRI. It is popular because it is relatively inexpensive and is the least invasive. It exploits the correlation that exists in the brain between neural activity and blood flow. However, because it does not directly measure neural activity, it is relatively imprecise, and must be coupled with statistical methods that extract signal from noise.⁴ More accurate technologies directly measure neural activity. Two examples are Transcranial Direct Current Stimulation (tDCS) and Transcranial Magnetic Stimulation (TMS). These methods temporarily stimulate (or inhibit) neural activity by applying electrical currents or magnetic pulses to selective regions of the brain. Clearly, this is more invasive. Sharot, Kanei, Marston, Korn, Rees, and Dolan (2012) employed TMS to study whether belief updating could be influenced by disrupting either the left or right hemispheres. Participants were asked to estimate the likelihood of various adverse life events (e.g., Alzheimers, crime, etc.). After each trial participants were then given new information stating the likelihood of these events for individuals matching their characteristics (e.g., age, race, gender, etc.). Finally, participants were given the opportunity to revise their estimates based on the new information. They found that participants with a disrupted left hemisphere became relatively more responsive to 'bad news' (population probability > prior), whereas participants with disrupted right hemispheres became relatively more responsive to 'good news' (population probability < prior). Fecteau, Knoch, Sultani, Boggio, and Pascual-Leone (2007) employed a tDCS study to examine risk-taking, the second stage of our investment problem. Participants were asked to play a gambling game, consisting of 100 trials of guessing the color of the box hiding a winning ticket. Each trial had six boxes of two colors, with the relative proportions differing from trial to trial. Payoffs were structured so that riskier bets yielded higher returns. Treated patients received tDCS

⁴fMRI studies are heavily used in the recent neuroeconomics literature. See Fehr and Rangel (2011) for a survey.

stimulation to the right hemisphere. They found that the average bets of treated patients were significantly less risky, which supports the notion that the right hemisphere is in charge of regulating risk taking behavior.⁵ These are just two examples of the many neuroscientific studies that support our maintained assumption that belief formation in the left hemisphere is optimistically biased, while risk taking behavior is regulated in the right hemisphere and is pessimistically biased. Hecht (2013) provides a comprehensive survey.

Finally, before moving on to the model, we should note that not everyone is convinced that these neuroscientific advances call for a rethinking of economic theory. In a widely cited critique of neuroeconomics, Gul and Pesendorfer (2008) state the following:

Since expected utility theory makes predictions only about choice behavior, its validity can be assessed only through choice evidence. If economic evidence leads us to the conclusion that expected utility theory is appropriate in a particular set of applications, then the inability to match this theory to the physiology of the brain might be considered puzzling. But this puzzle is a concern for neuroscientists, not economists.... (p. 25)

Our central argument is simple: neuroscience evidence cannot refute economic models because the latter make no assumptions or draw no conclusions about the physiology of the brain. Conversely, brain science cannot revolutionize economics because it has no vehicle for addressing the concerns of the latter. Economics and psychology differ in the questions they ask. Therefore, abstractions that are useful for one discipline will typically not be useful for the other. The concepts of a preference, a choice function, demand function, GDP, utility, etc. have proven to be useful abstractions in economics. The fact that they are less useful for the analysis of the brain does not mean that they are bad abstractions in economics. (p. 26)

We’ve included this extended quotation because we agree with it. Like the behaviorists before them, early advocates of neuroeconomics were arguably too devoted to tearing down economics, rather than building it up. Our goal is to show that insights from neuroscience can indeed improve economic models, even if you’re like Gul and Pesendorfer and only care about observed choices. If an alternative theory can explain asset prices equally well, we should prefer the theory that is more consistent with the ‘microfoundations’ of human behavior.

3. A BAYESIAN BENCHMARK

We start by studying an economy inhabited by individuals who face no tension between doubt and hope. They are Bayesians. This will provide a useful benchmark.

3.1. Environment. Consider a frictionless, complete markets economy, consisting of identical infinitely lived agents. Their only choice in life is to decide how much to invest in a ‘tree’ yielding a flow of nondurable ‘fruit’, which provides consumption utility.⁶

⁵It should be noted that the motivation for their study was to develop potential treatments for other types of risky behavior, e.g. taking drugs, not stock market investment!

⁶Zero net supply bonds are redundant in this economy. However, we can still compute a market-clearing (shadow) interest rate.

The optimal choice is characterized by an Euler equation. Exploiting the law-of-iterated-expectations and imposing a transversality condition, we can iterate forward the Euler equation and obtain the following present value asset pricing equation:

$$P_t U'(C_t) = E_t \int_t^\infty e^{-\delta(s-t)} U'(C_s) D_s ds \quad (3.1)$$

where δ is the rate of time preference. Notice that the marginal utility of consumption plays the role of a ‘stochastic discount factor’. When times are bad and consumption is low, marginal utility is high, which raises the value of the reduced flow of dividends. This trade-off becomes clear if we assume utility takes the constant relative risk aversion form, $U(C) = \frac{1}{1-\gamma} C^{1-\gamma}$, where γ is the coefficient of relative risk aversion. Imposing the equilibrium condition, $C_t = D_t$, we get

$$P_t D_t^{-\gamma} = E_t \int_t^\infty e^{-\delta(s-t)} D_s^{1-\gamma} ds \quad (3.2)$$

This expresses the equilibrium price, P_t , as an expectation of the exogenous flow of future dividends. Notice that the impact of future dividends on the current price is ambiguous. If $\gamma > 1$ then the discount rate effect dominates the cash flow effect, and asset prices decline in response to good news about future dividends. Conversely, if $\gamma < 1$ then prices rise. This tends to create problems when taking this model to the data. Returns are high on average, which calls for a high γ . However, prices are also procyclical. They are high when times are good, and low in bad times. In principle, this could be consistent with $\gamma > 1$ if there is sufficient mean reversion in consumption and dividends, so that current good times create expectations of future bad times. However, the data suggest that mean reversion is quite weak. In the next section we avoid this dilemma by following Barillas, Hansen, and Sargent (2009), and introduce a new parameter that captures model uncertainty. This allows us to keep $\gamma < 1$ in order to generate procyclicality, while at the same time generating high average returns via a model uncertainty premium. For now, however, in this Bayesian world, we are stuck. We shall see that when $\gamma < 1$ the equilibrium Price/Dividend ratio is far too high (and mean returns far too low).

To evaluate the expectation in eq. (3.2) agents must take a stand on the process generating dividends. For now we assume agents have no doubts about this process. Dividends are known to follow the process

$$\frac{dD_t}{D_t} = \mu_t dt + \sigma dB_t \quad (3.3)$$

where B_t is a Brownian motion. Notice that the mean growth rate, μ_t , is stochastic. Not only that, it is unobserved. However, it is known to be governed by the following mean-reverting Ornstein-Uhlenbeck process

$$d\mu_t = \rho(\bar{\mu} - \mu_t)dt + \sigma_\mu dB_t^\mu \quad (3.4)$$

where $\bar{\mu}$ is the long-run average growth rate and $\rho > 0$ is a mean reversion parameter. For simplicity, we assume the Brownian motion driving the growth rate, B_t^μ , is uncorrelated with the Brownian motion driving dividends. So far, this looks a lot like the Bansal-Yaron

(2004) long-run risks model. Remember, however, that preferences are time separable here, so that the long-run risk created by eqs. (3.3)-(3.4) is not priced in equilibrium.⁷

3.2. Learning. Since the growth rate is unobserved, agents must formulate an estimate, $\hat{\mu}_t$, and sequentially revise it as new information arrives. Given knowledge of the process (3.4), they can do this by employing the Kalman filter.


$$d\hat{\mu}_t = \rho(\bar{\mu} - \hat{\mu}_t)dt + \frac{Q_t}{\sigma}d\hat{B}_t \quad (3.5)$$

$$dQ_t = \left(\sigma_\mu^2 - 2\rho Q_t - \frac{Q_t^2}{\sigma^2} \right) dt \quad (3.6)$$

Eq. (3.5) is the sequence of conditional means and eq. (3.6) is the sequence of conditional variances, Q_t . Note that the conditional variance evolves deterministically. The innovation process, $d\hat{B}_t$ is defined as follows

$$d\hat{B}_t \equiv \frac{1}{\sigma} \left(\frac{dD_t}{D_t} - \hat{\mu}_t dt \right)$$

From the perspective of the agent, it is an i.i.d. process. It is simply the (standardized) error when forecasting the current dividend growth rate. Figure 2 plots $\hat{\mu}_t$ along with the sequence of actual earnings growth rates



The parameter values that were used are contained in Table 1. Note that the preference parameters, (γ, δ) , are not needed to implement the Kalman filter. The remaining four parameters $(\sigma, \sigma_\mu, \rho, \bar{\mu})$ were calibrated from observed earnings growth.⁸

⁷Epstein, Farhi, and Strzalecki (2014) criticize the Bansal-Yaron model on introspection grounds. They argue that the implied premium on the early resolution of uncertainty is implausible.

⁸Following Shiller, we use earnings rather than dividends, since changes in tax laws have reduced the incentive for firms to pay dividends, which produces a trend in the Price/Dividend ratio.

TABLE 1
BENCHMARK PARAMETER VALUES (ANNUALIZED)

γ	δ	σ	σ_μ	ρ	$\bar{\mu}$
0.955	.0116	.0163	.00087	.0096	.0185

The mean reversion parameter, ρ , was obtained by fitting an ARMA(1,1) process to the first difference in log earnings, which is the discretized version of eqs. (3.3)-(3.4). The (annualized) coefficient on lagged earnings growth is 0.99 (highly significant), which then suggests a value of $1 - .99 \approx .01$ for ρ . This turns out to be an important parameter in determining the dynamics of the PE ratio.

Figure 2 shows that the predicted growth rate (red line) tracks the actual growth rate well on average, but not surprisingly, is less volatile. The question now becomes – can this process of belief revision explain the observed fluctuations in the PE ratio seen in Figure 1?

3.3. Equilibrium. To evaluate the expected present value in (3.2) we employ the tools of dynamic programming. Define the right-hand side of (3.2) by the function $V(D, \hat{\mu})$. To simplify, we assume that the learning process has already converged to its steady state, so that Q_t is constant, and hence, not a state variable. It is given by the unique positive root of the following quadratic

$$\bar{Q}^2 + 2\rho\sigma^2\bar{Q} - \sigma_\mu^2\sigma^2 = 0$$

When expressed in terms of the agent's own information set, the state variable D is governed by the process

$$\frac{dD_t}{D_t} = \hat{\mu}_t dt + \sigma d\hat{B}_t \quad (3.7)$$

while the state variable $\hat{\mu}$ is governed by the Kalman filter in eq. (3.5). The Hamilton-Jacobi-Bellman (HJB) equation associated with this problem is then given by the following partial differential equation (PDE)

$$\delta V(D, \hat{\mu}) = D^{1-\gamma} + \hat{\mu}DV_D + \frac{1}{2}\sigma^2 D^2 V_{DD} + \rho(\bar{\mu} - \hat{\mu})V_{\hat{\mu}} + \frac{1}{2}\left(\frac{\bar{Q}}{\sigma}\right)^2 V_{\hat{\mu}\hat{\mu}} + \bar{Q}DV_{D\hat{\mu}} \quad (3.8)$$

where subscripts denote partial derivatives. To solve this PDE we can employ a standard guess-and-verify/separation of variables strategy. Conjecture that the solution takes the following form

$$V(D, \hat{\mu}) = F(\hat{\mu})D^{1-\gamma} \quad (3.9)$$

One can readily verify that a common $D^{1-\gamma}$ term appears in each term of eq. (3.8), and so can be cancelled out. Note that given the form of the solution in (3.9), and the fact that the left-side of (3.2) contains the term $D^{-\gamma}$, the equilibrium Price/Dividend ratio is then simply given by the function $F(\hat{\mu})$. It is determined by the solution of the following linear, nonhomogeneous, 2nd-order ODE:

$$\delta F(x) = 1 + (1 - \gamma) \left(x - \frac{1}{2} \sigma^2 \gamma \right) F(x) + \rho(\bar{\mu} - x) F'(x) + \frac{1}{2} (\bar{Q}/\sigma)^2 F''(x) + (1 - \gamma) \bar{Q} F'(x) \quad (3.10)$$

where for notational convenience in what follows, we substitute x in place of $\hat{\mu}$. Unfortunately, this equation does not have a convenient closed-form solution in terms of named functions. However, it can be solved easily using a series expansion. The details are provided in Appendix A. This gives us

Proposition 3.1. *With Bayesian beliefs, the equilibrium Price/Earnings ratio is given by*

$$F(x) = e^{\lambda x} \left\{ \sum_{n=0}^{\infty} a_n (A_0 - g\rho x)^n \right\} \quad (3.11)$$


where $\lambda \equiv (1 - \gamma)/\rho$, $g \equiv 2(\sigma/\bar{Q})^2$, and $A_0 \equiv [\rho\bar{\mu} + \bar{Q}(1 - \gamma)]$. The a_n coefficients are given by the following recursion

$$a_n = - \frac{[a_{n-2}(c - (n-2)b)(n-2)! + H\eta^{n-2}]}{n!} \quad n \geq 3 \quad (3.12)$$

with $a_0 = -H/c$ and $a_1 = a_2 = 0$. (The constants H and c are defined in Appendix A).

This series converges rapidly for all economically reasonable parameter values. The dominant factor driving fluctuations in the PE ratio is the exponential function, $\exp(\lambda x)$. If λ is large, the PE ratio will be volatile. It will be large when γ is close to 0 (risk-neutrality) and ρ is small (high persistence in dividend growth rates). Since empirically ρ is close to 0, relative risk aversion must be close to 1. Otherwise, the PE ratio will be *too* volatile. Of course, if $\gamma = 1$ and preferences are logarithmic, then the PE ratio becomes constant and equal to $1/\delta$. From the results in Appendix A, one can show that when $\gamma < 1$ growth uncertainty, $\bar{Q}$, effectively increases the expected long-run growth rate (ie, a higher A_0), and hence the mean PE ratio. This is the Jensen's inequality effect highlighted by Pastor and Veronesi (2003).

Figure 3 plots eq. (3.11) (magenta line) using the parameters in Table 1 and the sequence of state estimates, $\hat{\mu}_t$, given by the Kalman filter in eq. (3.5) (red line in Fig. 2).



Evidently, as many others have discovered before us, we can see that Bayesian learning in a Complete Markets/Representative Agent economy really doesn't help resolve the Equity Premium Puzzle. The mean PE ratio is far too high (corresponding to a low mean return). Of course, if we amplify γ enough we can reduce the mean PE ratio. Unfortunately, this would also produce countercyclical PE fluctuations, which is strongly counterfactual. This tension is what motivated Bansal and Yaron (2004) to use recursive preferences.

On the bright side, Figure 3 shows that when $\gamma < 1$ Bayesian learning can produce large procyclical fluctuations in PE ratios. This echoes the results of Weitzman (2007). Weitzman shows that in models featuring stochastic latent variables, like ours, posteriors about asset returns are *very* sensitive to priors about the parameters of the latent variable process. In the next section we exploit this sensitivity to explain both the Equity Premium Puzzle and the volatility and comovement puzzles. Our contribution relative to Weitzman (2007) is to provide a theory of why agents might want to perturb their prior in a way that better fits the data.

4. ASSET PRICING WHEN AGENTS HAVE A HUMAN BRAIN

Although this section contains the paper's main result, most of the heavy lifting was done in the previous section. A key advantage of the Hansen-Sargent approach to ambiguity and model uncertainty is that it requires relatively minor modifications of 'business as usual'. We just need to add drift distortions to the dividend process (3.7) and the Kalman filter in (3.5). Both distortions are constrained by a relative entropy cost. As in Barillas, Hansen, and Sargent (2009), the distortion to the dividend process is negative, and reflects doubts about model specification. As discussed in Section 2, this pessimism is being driven by the alert system of right hemisphere. The distortion to the Kalman filter is

motivated by Hansen and Sargent (2007). However, here there is an important difference. Hansen and Sargent (2007) introduce a pessimistic drift distortion to the Kalman filter, reflecting an agent's doubts about his own prior. As discussed in Section 2, here we do precisely the opposite. Our belief distortion is *optimistic*, reflecting the left hemisphere's role in formulating optimistically biased plans and expectations.

The interaction between the left and right hemispheres of the brain can be formalized via the following dynamic zero sum game version of the expected present value of dividends

$$V(D, \hat{\mu}) = \min_{h_1} \max_{h_2} \tilde{E}_t \int_t^\infty e^{-\delta(s-t)} \left\{ D_s^{1-\gamma} + \frac{1}{2}(1-\gamma) (\varepsilon_1^{-1} h_{1s}^2 - \varepsilon_2^{-1} h_{2s}^2) \right\} ds \quad (4.13)$$

where the quadratic cost functions in the h_{it} sequences represent relative entropy penalties. The free parameters $(\varepsilon_1, \varepsilon_2)$ determine the level of pessimism and optimism. They are scaled by $(1-\gamma)^{-1}$ to put them in units of utility. They are closely related to the 'decision costs' in Brunnermeier and Parker's (2005) model of Optimal Expectations. Notice that if $\varepsilon_i = 0$ then the entropy costs become infinite, and the optimal choices are $h_{1t} = h_{2t} = 0$. In this case, the model collapses to the Bayesian world of Section 3. In fact, as emphasized by Hansen and Sargent (2008), there is an ex-post Bayesian interpretation of this problem, which allows us to view the agent as a Bayesian, but with a particular (distorted) prior. Hence, there is no necessary inconsistency between our approach and Bayesian approaches to neuroscience (see, e.g., Doya, Ishii, Pouget, and Rao (2007) and Parr, Pezzulo, and Friston (2022)).⁹

Expectations are taken with respect to the following drift distorted dividend process

$$dD_t = (\hat{\mu}_t + \sigma h_{1t}) D_t dt + \sigma D_t d\tilde{B}_t \quad (4.14)$$

and the following drift distorted sequence of state estimates

$$d\hat{\mu}_t = [(\rho(\bar{\mu} - \hat{\mu}_t) + (\bar{Q}/\sigma) h_{2t}) dt + \frac{\bar{Q}}{\sigma} d\hat{B}_t] \quad (4.15)$$

The idea here is that the cautious right hemisphere controls the model distortion sequence, h_{1t} , while the bold and imaginative left hemisphere controls the belief distortion, h_{2t} .

Performing the optimization and subbing the optimal h_i functions into HJB equation yields the following modified HJB equation

$$\begin{aligned} \delta V(D, \hat{\mu}) = & D^{1-\gamma} + \hat{\mu} D V_D + \frac{1}{2} \sigma^2 D^2 V_{DD} + \rho(\bar{\mu} - \hat{\mu}) V_{\hat{\mu}} + \frac{1}{2} \left(\frac{\bar{Q}}{\sigma} \right)^2 V_{\hat{\mu}\hat{\mu}} + \bar{Q} D V_{D\hat{\mu}} \\ & - \frac{1}{2} \frac{\varepsilon_1}{1-\gamma} \sigma^2 D^2 (V_D)^2 + \frac{1}{2} \frac{\varepsilon_2}{1-\gamma} (\bar{Q}/\sigma)^2 (V_{\hat{\mu}})^2 \end{aligned} \quad (4.16)$$

The first few terms are identical to the Bayesian HJB equation in (3.8). Notice, however, that the final two terms make the equation *nonlinear*. Nonlinear PDEs are a no man's land. They can only be solved numerically, or approximately via perturbation expansions. Another problem is that this nonlinearity will in general make the solution nonhomothetic. Depending on whether $\gamma \leq 1$, optimism and pessimism will either wax or wane as the economy grows.

⁹Interestingly, Mangalam (2025) criticizes the Bayesian Brain literature along lines that are remarkably similar to Hansen and Sargent's (2008) criticisms of the Bayesian interpretation of robust control.

To deal with these problems we employ a trick first proposed by Maenhout (2004). This involves scaling the cost parameters $(\varepsilon_1, \varepsilon_2)$. Notice that if ε_1 is scaled by V , so that $\varepsilon_1 = \bar{\varepsilon}_1/V$, where $\bar{\varepsilon}_1$ is constant, then the term involving $(V_D)^2$ becomes homothetic, and a conjectured solution of the form $V \sim D^{1-\gamma}$ will be consistent. Behaviorally, homotheticity is being induced by making the cost of pessimism increase as the economy grows (assuming $\gamma < 1$). The filtering cost parameter, ε_2 , requires a slightly different scaling. Suppose as before we conjecture a separable solution of the form $V = D^{1-\gamma}F(\hat{\mu})$. Then notice that if we employ the scaling $\varepsilon_2 = (1 - \gamma)\bar{\varepsilon}_2/(D^{1-\gamma}F')$, then equation becomes linear and homothetic.

Given this scaling we can once again posit a solution of the form $V = D^{1-\gamma}F$, and after cancelling out $D^{1-\gamma}$, obtain the following *linear* ODE for the PE ratio

$$\begin{aligned} \delta F(x) = & 1 + (1 - \gamma) \left(x - \frac{1}{2}\sigma^2\gamma \right) F(x) + \rho(\bar{\mu} - x)F'(x) + \frac{1}{2}(\bar{Q}/\sigma)^2 F''(x) + (1 - \gamma)\bar{Q}F'(x) \\ & - \frac{1}{2}\bar{\varepsilon}_1\sigma^2(1 - \gamma)F(x) + \frac{1}{2}\bar{\varepsilon}_2(\bar{Q}/\sigma)^2 F'(x) \end{aligned} \quad (4.17)$$

Mathematically, this is the exact same equation as before. All that has changed are the coefficient functions. We can solve it using a series expansion just like before. This gives us


Proposition 4.1. *When agents hope for the best, but plan for the worst, the (scaled) equilibrium PE ratio is given by the same expression as the Bayesian PE ratio,*

$$F(x) = e^{\lambda x} \left\{ \sum_{n=0}^{\infty} a_n (A_0 - g\rho x)^n \right\}$$

All that changes are the definitions of the (A_0, c, H, η) parameters, which then changes the sequence of expansion coefficients, a_n , given by the recursion in eq. (3.12).

Because λ is independent of the $\bar{\varepsilon}_i$ parameters, and because $a_1 = a_2 = 0$, to a 3rd-order approximation the effects of doubt and hope operate via the constant term, a_0 . By inspection of eq. (4.17), doubts about the model just inject a constant downward adjustment to growth, $\frac{1}{2}\sigma^2\bar{\varepsilon}_1(1 - \gamma)$, into the F coefficient on the right-hand side of (4.17). This reduces a_0 and the mean PE ratio. At the same time, (homothetic) optimistic filtering just injects a constant upward adjustment to the long-run growth rate, $\frac{1}{2}(\bar{Q}/\sigma)^2$, into the F' coefficient. This increases a_0 and the mean PE ratio.

Figure 4 plots the PE ratio in an economy where agents hope for the best, but plan for the worst. The parameters are the same as before. The new parameters, $(\bar{\varepsilon}_1, \bar{\varepsilon}_2)$ were set to $\bar{\varepsilon}_1 = 1190$ and $\bar{\varepsilon}_2 = 12$. By themselves, these values are meaningless, since they are not unit free. What matters is the magnitude of the implied distortions. The pessimism parameter $\bar{\varepsilon}_1$ implies a 71 basis point reduction in the growth rate of dividends. Conversely, the optimism parameter $\bar{\varepsilon}_2$ implies a 55 basis point positive distortion to $\bar{\mu}$.




Evidently, the mean PE ratio is now closer to the data. However, a 71 basis point pessimistic distortion to the dividend process results in only a modest decrease in the mean PE ratio. Predicted mean returns only increase by about 130 basis points. Of course, we could increase this effect by increasing $\bar{\epsilon}_1$, but as we shall see below, this would produce an excessively low detection error probability. Also, notice that the volatility is substantially reduced. In fact, the human brain coefficient of variation is about the same as the Bayesian coefficient of variation. The challenge now is to increase volatility while at the same time decreasing the mean PE ratio.

4.1. Numerical Solution. The previous series solution is convenient for developing some intuition for how the parameters influence the moments of the predicted PE ratio. However, it suffers from a serious disadvantage. Scaling the cost parameters to deliver a linear PDE implies that the distortions are *constant*. This can still increase mean returns (as seen above), but it seems clear that allowing the distortions to be state-dependent will make it easier for the model to match the data. Moreover, state dependent optimism and pessimism are consistent with the survey evidence presented in Bordalo et. al. (2024). (See below for further discussion).

Figure 5 presents the Bayesian and Human Brain solutions for the PE ratio from the unscaled/nonlinear PDE in equation (4.16). We solve the nonlinear Hamilton-Jacobi-Bellman equation using a perturbation expansion in the optimism and pessimism parameters. The zeroth-order term captures the Bayesian case, and the two first-order terms account for the human brain. The zeroth-order solution is obtained by solving a linear elliptic PDE using finite differences on a uniform grid, with Neumann boundary conditions enforced through one-sided stencils. The first-order terms solve forced versions of

the same PDE, where the forcing terms depend nonlinearly on the zeroth-order solution and its derivatives. A full description of the numerical method is provided in Appendix B.¹⁰



Evidently, the fit is much improved. The predicted PE ratio is both lower and more volatile. The mean PE ratio is now 27.5, only modestly above the actual value of 17.5. The standard deviation is 13.4, which is actually higher than the observed value of 7.42. However, if one corrects for differences in the mean by computing the coefficient of variation, the volatilities match up pretty well. The predicted value is .489, while the actual is .424. Perhaps most impressive is the model's ability to match the rise and fall of the PE ratio during the Roaring 20s/Great Depression and Dotcom Boom periods. Table 2 summarizes the results


TABLE 2
ACTUAL VS. PREDICTED MOMENTS

	Mean	st. dev.	cvar	$\text{corr}(PE, \hat{PE})$
Data	17.5	7.42	.424	—
Bayes	95.1	11.6	.122	.402
Scaled	43.7	5.47	.125	.279
Numerical	27.5	13.4	.489	.345

Crucially, both the model and filtering distortions now fluctuate over time. The Left-Brain/Right-Brain free parameters are set to $(\bar{\varepsilon}_1, \bar{\varepsilon}_2) = (.0841, .302)$, which produces a mean pessimistic drift distortion of about 2.5% and a mean optimistic filtering distortion

¹⁰The Matlab code is available upon request and posted on the author's webpage.


of almost 1% (both annualized). Figure 6 depicts the time-series of the distortions. Note, the filtering distortion is scaled to be in units of long-run mean growth (by dividing by the mean reversion parameter ρ).



Interestingly, notice that the two distortions are negatively correlated. When optimism about growth increases, the agent becomes more pessimistic about his model. Comparing these distortions with actual and predicted earnings in Figure 2, notice that the optimistic left-hemisphere overreacts to good news when growth is strong, especially given the high persistence of growth. At the same time, the pessimistic right-hemisphere urges restraint, and when $\gamma < 1$ this restraint becomes stronger during good times. This internal tension produces heightened volatility, especially during economic expansions.

4.2. Comparison to Bordalo et. al. (2024). Below we report evidence that these distortions are empirically plausible. Following Hansen and Sargent (2008), we do this by computing detection error probabilities. However, as emphasized by Bordalo et. al. (2024), one could instead adopt a less model dependent strategy by simply *asking* people what they expect growth to be over various horizons. In this section we show that our results are quite consistent with this survey evidence.

Figure 7 plots the optimistic distortion to $\bar{\mu}$ from Figure 6, corresponding to the shorter sample reported in Bordalo et. al. (2024), i.e., 1980-2020. It should be compared to Figure 1 in their paper. Notice that our theory-implied long-run growth forecast distortion exhibits a very similar pattern. Long-run growth forecasts decline by about 15 basis points during the 1980s, then rise rapidly, by about 30 basis points during the Dotcom boom period of the 1990s. These optimistic long-run growth forecasts are then reversed during the Financial Crisis of 2008-09.



The fact that our model can endogenously reproduce the pattern of expectations in the survey data serves as a useful cross-check on our results.

5. DETECTION ERROR PROBABILITIES

The key mechanism in our model is belief distortions. Investors are defensively pessimistic when constructing models of dividend growth, but conditional on their models, they ‘hope for the best’ when formulating and revising forecasts about future dividend growth. We argued above that evidence from neuroscience supports this tension between optimism and pessimism. Unfortunately, the neuroscience literature does not provide much guidance on the *magnitude* of these distortions. We have introduced two new free parameters (the relative entropy costs on pessimism and optimism), so it seems rather obvious we could tune these parameters to explain both high and volatile returns. How do we know that the resulting parameter values are empirically plausible? Is there any external evidence we can consider that would discipline our calibration?

Here we follow Hansen and Sargent (2008) and quantify plausibility using ‘detection error probabilities’. Agents are only allowed to consider models that could have plausibly generated the observed data. Specifically, we view agents as statisticians, who attempt to discriminate among models using likelihood ratio statistics. When likelihood ratio statistics are large, detection error probabilities are small, and models are easy to distinguish. Detection error probabilities will be small when models are very different, or when there is a lot of data. Classical inference is based on adopting a null hypothesis, and fixing the probability of falsely rejecting this hypothesis. Detection error probabilities treat the null and alternative symmetrically, and average between Type I and Type II errors. In

particular,

$$DEP = \frac{1}{2}\text{Prob}(H_A|H_0) + \frac{1}{2}\text{Prob}(H_0|H_A)$$

Hence, a DEP is analogous to a p-value. Our results would therefore be implausible if the DEP is too small. Small DEP's imply agents are hedging against models that could easily be rejected using observed data.

In general, entropy constrained distortions are state dependent, and so the error probabilities must be computed using monte carlo simulation (Hansen and Sargent (2008, Chpt. 9)). However, in the scaled version of our model, the resulting distortions are *constant*. Hence, the null and alternative models for log dividend growth are both linear Gaussian ARMA(1,1) models, which only differ in their intercepts. This allows us to compute detection error probabilities using a simple z -score.¹¹

By quasi-differencing the Kalman filter for $\hat{\mu}_t$, (4.17), and then substituting into the discretized version of the dividend growth model in (4.16) we obtain

$$[1 - (1 - \rho)L]\Delta \log(D_{t+1}) = \rho\bar{\mu} + (\Delta_\mu - \rho\Delta_d) + \sigma \left[\frac{\bar{Q}}{\sigma^2} + 1 - (1 - \rho)L \right] d\hat{B}_{t+1} \quad (5.18)$$

where $\Delta_\mu > 0$ is the distortion to the Kalman filter and $\Delta_d > 0$ is the distortion to the dividend growth model. As noted in the previous section, our benchmark parameters imply $\Delta_\mu = .0055$ and $\Delta_d = .0071$ (both annualized). There are a couple of important points to notice here. First, notice the two distortions *offset* each other. This makes it harder to detect them, based solely on observed data. Second, notice that the dividend model distortion is scaled by ρ , which is close to zero, given the observed weak mean reversion in dividend growth. This means that detection error probabilities are driven primarily by the optimistic filtering distortion. Intuitively, when stochastic growth features weak mean reversion, it becomes difficult to isolate the mean growth rate.

Since the detection errors only depend on the intercept, the Type I and Type II error probabilities are the same, and we obtain

$$DEP = \Phi \left(\frac{-\sqrt{T}|\Delta_\mu - \rho\Delta_d|}{\hat{\sigma}} \right)$$

where Φ denotes the standard normal CDF and $\hat{\sigma}$ is the steady state standard deviation of the error term in (5.22). It is given by

$$\hat{\sigma} = \sigma \sqrt{(1 + \bar{Q}/\sigma)^2 + (1 - \rho)^2}$$

¹¹We take a couple short cuts. First, we assume Q_t has converged to its steady state, and so is constant. Second, we ignore the MA(1) structure of the model's error term. Since this tends to bias toward rejection, it should bias our detection error probabilities downwards, which reinforces our results. It would be straightforward to consider the ARMA(1,1) case using the frequency domain procedure outlined in Taniguchi and Kakizawa (2000, pgs. 490-500).

If we set $T = 130 \times 12$, which is roughly equal to the length of our data set, we find $DEP = .097$. Thus, even after 130 years our left-brain/right-brain investor would not be able to resolve their optimism/pessimism tension with much statistical confidence.¹²

In the nonlinear version of our model, in Section 4.1, the above strategy no longer works, since distortions become state-dependent. Anderson, Hansen, and Sargent (2003) provide a Chernoff-type DEP upper bound formula, which involves integrating the distortions over time (eq. (8.11) in their paper). However, in practice this bound is often not very tight. So instead we pursue an asymptotic (large deviations) approximation, which gives the rate at which detection errors decline as $T \rightarrow \infty$. This rate is state independent. Using results from Kutoyants (2004, chpt. 7) we have¹³

$$DEP \approx \frac{1}{2} \exp \{ -TE(h_2^2) \} = .197 \quad (5.19)$$

where, since $\rho \approx 0$, we ignore the model distortions, $h_1(D, \hat{\mu})$. Since they are offsetting, this reinforces our results. The expectation, $E(h_2^2)$, is computed using the distorted probability measure, by simply averaging the squared Kalman filter distortions. Since Type I and Type II error rates are equal in the large deviations limit, it does not matter which probability measure we use. Essentially, what we are computing is the probability that a positive (negative) drift diffusion ends up at a negative (positive) value. Since the magnitude of the drift and the diffusion coefficients are identical, the error rates are the same. Remember, this is really an *upper* bound to the actual DEP, since we are ignoring a term that goes to zero at a slower than exponential rate. Although .197 is reassuringly high, it would nonetheless be useful to check it by running a Monte Carlo simulation to compute the actual detection error probability.

6. CONCLUSION

This paper has shown that when agents are pessimistic about their models but optimistic about their beliefs, stock returns will be high on average and Price/Earnings ratios will exhibit large procyclical fluctuations. This is true even though markets are complete and agents have modest (nearly logarithmic) risk aversion. The key innovation in our paper is to link this apparently schizophrenic combination of optimism and pessimism to recent work in neuroscience, which shows that beliefs are processed primarily by the optimistically-biased left-hemisphere of the brain, while decisions are primarily processed by the pessimistically biased right-hemisphere.

Our belief-based explanation of asset return dynamics requires the assumption that relative risk aversion is low, in particular, $\gamma < 1$. This may seem empirically implausible in light of the past 40 years of asset pricing research. For example, although Long-Run Risk models are able to reduce γ by positing recursive preferences, they still require $\gamma > 1$ to match the equity premium. Following Barillas, Hansen, and Sargent (2009), we are able to keep γ low by substituting model uncertainty for risk aversion. We avoid implausibly low DEPs by offsetting this with a form of optimism about long-run growth. Also, keep

¹²In reality, agents are finite-lived and seem to place more weight on their own experiences (Malmendier and Nagel (2011)). In this case, detection error probabilities would of course be even larger, and would persist (within the economy) indefinitely.

¹³A more precise statement is $\lim_{T \rightarrow \infty} \frac{1}{T} \log(\text{Type I} + \text{Type II}) = -E(h_2^2)$.

in mind that values of γ in the neighborhood of 1.0 are consistent with most estimates of life-cycle models using microdata.

An important task for future work is to examine the robustness of these results to alternative model specifications and parameter values. Preliminary analysis suggests that model-implied price volatility is quite sensitive to mean reversion in earnings (as parameterized by ρ). In the Introduction we noted that our paper is similar to Bordalo, Gennaioli, Porta, and Shleifer (2024). One advantage of their work is to show that belief overreaction also explains cross-sectional asset pricing puzzles. Since earnings persistence differs across firms, it would be interesting to see whether our model can also shed light on firm-level stock return data.

APPENDIX A. PROOF OF PROPOSITION 3.1

Since the ODE we want to solve is linear, the general solution is the sum of a particular solution and two linearly independent solutions of the homogeneous equation. Unlike Froot and Obstfeld (1991), we are not interested in intrinsic bubbles, which correspond to solutions of the homogeneous equation. However, these solutions are still useful, since they lead to a particular solution of the nonhomogeneous equation via a ‘variation of parameters’ strategy. Unfortunately, the homogeneous solutions take the form of confluent hypergeometric functions (see Abramowitz and Stegun (1972, chpt. 13)), which then must be integrated when constructing the particular solution. This makes an ‘analytical’ solution complicated and difficult to interpret. So instead we resort to a simple series solution (see, e.g., Boyce and DiPrima (1977, chpt. 4)).

Begin by writing eq. (3.10) as follows:

$$F''(x) + (b_0 - b_1x)F'(x) + (b_2 + b_3x)F(x) = -g \quad (\text{A.20})$$

where $g = 2(\sigma/Q)^2$, $b_0 = g[\rho\bar{\mu} + (1 - \gamma)Q]$, $b_1 = g\rho$, $b_2 = -g[\delta + \frac{1}{2}\sigma^2\gamma(1 - \gamma)]$, and $b_3 = g(1 - \gamma)$. We can eliminate the x -dependence from the coefficient of F by writing the solution as follows:

$$F(x) = v(x)u(x)$$

and assuming that

$$v(x) = e^{\lambda x} \quad \lambda = (1 - \gamma)/\rho$$

This produces the following equation for $u(x)$:

$$u''(x) + (A_0 - b_1x)u'(x) + A_1u(x) = -ge^{-\lambda x} \quad (\text{A.21})$$

where $A_0 = b_0 + 2\lambda$ and $A_1 = b_2 + \lambda b_0 + \lambda^2$. Next, if we define the change-of-variables

$$z = A_0 - b_1x$$

we obtain the equation

$$u''(z) - bzu'(z) + cu(z) = -He^{\eta z} \quad (\text{A.22})$$

where $b = 1/b_1$, $c = A_1/b_1^2$, $H = (g/b_1^2)e^{-\lambda A_0/b_1}$, and $\eta = \lambda/b_1$. Equation (A.22) can be solved easily using a series expansion. Write the solution in the following form

$$u(z) = \sum_{n=0}^{\infty} a_n z^n \quad (\text{A.23})$$

where the a_n are undetermined coefficients. Using this conjectured solution to take derivatives, expanding the exponential function on the right-hand side, and then matching coefficients on z produces the following recursion for the a_n coefficients:

$$0 = 2a_2 + ca_0 + H \quad (\text{A.24})$$

$$a_n = -\frac{[a_{n-2}(c - (n-2)b)(n-2)! + H\eta^{n-2}]}{n!} \quad n \geq 3 \quad (\text{A.25})$$

We can pin down a unique solution by exploiting the fact that $u = 1/\delta$ when $\gamma = 1$. This tells us that

$$\begin{aligned} a_0 &= -\frac{H}{c} \\ a_1 &= 0 \end{aligned}$$

which then implies $a_2 = 0$. Finally, reversing the change-of-variables gives us the following solution to eq. (3.10)

$$F(x) = e^{\lambda x} \left\{ \sum_{n=0}^{\infty} a_n (A_0 - b_1x)^n \right\}$$

as stated in Proposition 3.1.

APPENDIX B. NUMERICAL SOLUTION OF THE HJB EQUATION

To solve the nonlinear HJB equation 4.16, we employ an asymptotic expansion in the parameters ε_1 and ε_2 . The value function $V(D, \hat{\mu})$ is approximated as:

$$V(D, \hat{\mu}) \approx V_0(D, \hat{\mu}) + \varepsilon_1 V_1(D, \hat{\mu}) + \varepsilon_2 V_2(D, \hat{\mu})$$

where V_0 is the zeroth-order solution (Bayesian case) and V_1 and V_2 are first-order corrections capturing pessimism and optimism distortions, respectively. This decomposition converts the original nonlinear PDE into a sequence of linear problems that can be solved easily. The expansion is valid when ε_1 and ε_2 are small, which aligns with our economic interpretation of moderate optimism and pessimism concerns.

B.1. Zeroth-Order Solution (V_0). We assume a separable form $V_0(D, \hat{\mu}) = F(\hat{\mu})D^{1-\gamma}$, which emerges when $\varepsilon_1 = \varepsilon_2 = 0$. The function $F(\hat{\mu})$ satisfies the second-order ODE:

$$\delta F = 1 + (1 - \gamma)(\hat{\mu} - \frac{1}{2}\gamma\sigma^2)F + [\rho(\bar{\mu} - \hat{\mu}) + (1 - \gamma)Q]F' + \frac{1}{2}(Q/\sigma)^2 F''$$

We discretize this ODE on a uniform grid $\{\hat{\mu}_i\}$ with N points and spacing $\Delta\hat{\mu}$. Derivatives are approximated as:

- Interior points ($i = 2, \dots, N - 1$):

$$F'_i \approx \frac{F_{i+1} - F_{i-1}}{2\Delta\hat{\mu}}, \quad F''_i \approx \frac{F_{i+1} - 2F_i + F_{i-1}}{\Delta\hat{\mu}^2}$$

- Boundary conditions ($i = 1$ and $i = N$), using Neumann stencils:

$$\text{Left: } \frac{-3F_1 + 4F_2 - F_3}{2\Delta\hat{\mu}} = 0, \quad \text{Right: } \frac{F_{N-2} - 4F_{N-1} + 3F_N}{2\Delta\hat{\mu}} = 0$$

– *Derivation:* From Taylor expansions of F_2 and F_3 around F_1 , solving for F'_1 .

The discretized system takes the form $\mathbf{A}\mathbf{F} = \mathbf{b}$, where A is an $N \times N$ sparse matrix that captures the finite difference scheme. It combines the diagonal terms from the F coefficient, the off-diagonal terms from the first and second derivatives, and the modified boundary rows from the Neumann conditions. $\mathbf{b}$ is a vector of ones (from the constant term in the ODE), and the solution F represents the PE ratio function in the Bayesian case. The elements of matrix $\mathbf{A}$ are as follows:

- Interior rows ($i = 2, \dots, N - 1$):

$$A_{i,i-1} = \frac{\bar{Q}^2}{2\sigma^2\Delta\hat{\mu}^2} - \frac{\rho(\bar{\mu} - \hat{\mu}_i) + (1 - \gamma)\bar{Q}}{2\Delta\hat{\mu}},$$

$$A_{i,i} = \delta - (1 - \gamma) \left(\hat{\mu}_i - \frac{\gamma\sigma^2}{2} \right) - \frac{\bar{Q}^2}{\sigma^2\Delta\hat{\mu}^2},$$

$$A_{i,i+1} = \frac{\bar{Q}^2}{2\sigma^2\Delta\hat{\mu}^2} + \frac{\rho(\bar{\mu} - \hat{\mu}_i) + (1 - \gamma)\bar{Q}}{2\Delta\hat{\mu}}.$$

- Boundary rows:

– Left ($i = 1$):

$$A_{1,1:3} = \left[-\frac{3}{2\Delta\hat{\mu}}, \frac{4}{2\Delta\hat{\mu}}, -\frac{1}{2\Delta\hat{\mu}} \right], \quad \mathbf{b}_1 = 0$$

– Right ($i = N$):

$$A_{N,N-2:N} = \left[\frac{1}{2\Delta\hat{\mu}}, -\frac{4}{2\Delta\hat{\mu}}, \frac{3}{2\Delta\hat{\mu}} \right], \quad \mathbf{b}_N = 0$$

B.2. First-Order Corrections (V_1 and V_2). The first-order corrections satisfy similar PDEs with additional source terms:

- For $V_1 = G(\hat{\mu})D^{2(1-\gamma)}$:

$$\frac{1}{2}\varepsilon_1(1-\gamma)^2F^2 \text{ appears as the forcing term}$$

- For $V_2 = H(\hat{\mu})D^{2(1-\gamma)}$:

$$-\frac{1}{2}(Q/\sigma)^2(F')^2 \text{ appears as the forcing term}$$

Both G and H satisfy equations of the form $AG = b^G$ and $AH = b^H$ with the same matrix A as before. Using the numerical solution of F and its derivative F' , we can construct

$$b^G = \frac{1}{2}\varepsilon_1(1-\gamma)^2F^2, \quad b^H = -\frac{1}{2}(Q/\sigma)^2(F')^2$$

Thus, we just need to solve

$$G = A^{-1}b^G, \quad H = A^{-1}b^H$$

Finally, the full first-order approximation is constructed by the following assembly.

$$V(D, \hat{\mu}) = F(\hat{\mu})D^{1-\gamma} + \varepsilon_1 G(\hat{\mu})D^{2(1-\gamma)} + \varepsilon_2 H(\hat{\mu})D^{2(1-\gamma)}$$

REFERENCES

- ABRAMOWITZ, M., AND I. A. STEGUN (1972): *Handbook of Mathematical Functions*. Dover Publications.
- ANDERSON, E. W., L. P. HANSEN, AND T. J. SARGENT (2003): “A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection,” *Journal of the European Economic Association*, 1, 68–123.
- BANSAL, R., AND A. YARON (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59, 1481–1509.
- BARILLAS, F., L. P. HANSEN, AND T. J. SARGENT (2009): “Doubts or Variability?,” *Journal of Economic Theory*, 144, 2388–418.
- BARRO, R. J. (2006): “Rare Disasters and Asset Markets in the 20th Century,” *Quarterly Journal of Economics*, 121, 823–66.
- BEELER, J., AND J. Y. CAMPBELL (2012): “The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment,” *Critical Finance Review*, 1, 141–82.
- BENARTZI, S., AND R. H. THALER (1995): “Myopic Loss Aversion and the Equity Premium Puzzle,” *Quarterly Journal of Economics*, 110, 73–92.
- BORDALO, P., N. GENNAIOLI, R. L. PORTA, AND A. SHLEIFER (2024): “Belief Overreaction and Stock Market Puzzles,” *Journal of Political Economy*, 132, 1450–84.
- BOYCE, W. E., AND R. C. DIPRIMA (1977): *Elementary Differential Equations and Boundary Value Problems*. John Wiley & Sons, third edn.
- BRUNNERMEIER, M. K., AND J. A. PARKER (2005): “Optimal Expectations,” *American Economic Review*, 95, 1092–1118.
- CAMPBELL, J. Y., AND J. H. COCHRANE (1999): “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107, 205–51.
- CONSTANTINIDES, G. M., AND D. DUFFIE (1996): “Asset Pricing With Heterogeneous Consumers,” *Journal of Political Economy*, 104, 219–40.
- CONSTANTINIDES, G. M., AND A. GHOSH (2017): “Asset Pricing With Countercyclical Household Consumption Risk,” *Journal of Finance*, 72, 415–59.
- DOYA, K., S. ISHII, A. POUGET, AND R. RAO (2007): *Bayesian Brain: Probabilistic Approaches to Neural Coding*. MIT Press.
- EPSTEIN, L. G., E. FARHI, AND T. STRZALECKI (2014): “How Much Would You Pay to Resolve Long-Run Risk?,” *American Economic Review*, 104, 2680–97.
- FECTEAU, S., D. KNOCH, N. SULTANI, P. BOGGIO, AND A. PASCUAL-LEONE (2007): “Diminishing Risk-Taking Behavior by Modulating Activity in the Prefrontal Cortex: A Direct Current Stimulation Study,” *Journal of Neuroscience*, 14, 12500–05.
- FEHR, E., AND A. RANGEL (2011): “Neuroeconomic Foundations of Economic Choice – Recent Advances,” *Journal of Economic Perspectives*, 25, 3–30.
- FROOT, K. A., AND M. OBSTFELD (1991): “Intrinsic Bubbles: The Case of Stock Prices,” *American Economic Review*, 81, 1189–1214.
- GAZZANIGA, M. S. (2008): *Human: The Science Behind What Makes Your Brain Unique*. HarperCollins.
- GUL, F., AND W. PESENDORFER (2008): “The Case for Mindless Economics,” in *The Foundations of Positive and Normative Economics: A Hand Book*, ed. by A. Caplin, and A. Schotter. Oxford University Press.
- HANSEN, L. P., AND T. J. SARGENT (2007): “Recursive Robust Estimation and Control without Commitment,” *Journal of Economic Theory*, 136, 1–27.
- (2008): *Robustness*. Princeton University Press.
- (2010): “Fragile Beliefs and the Price of Uncertainty,” *Quantitative Economics*, 1, 129–62.
- HANSEN, L. P., AND K. J. SINGLETON (1983): “Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns,” *Journal of Political Economy*, 91, 249–65.
- HASELTON, M. G., AND D. NETTLE (2006): “The Paranoid Optimist: An Integrative Evolutionary Model of Cognitive Biases,” *Personality and Social Psychology Review*, 10, 47–66.
- HECHT, D. (2013): “The Neural Basis of Optimism and Pessimism,” *Experimental Neurobiology*, 22, 173–99.

- JU, N., AND J. MIAO (2012): “Ambiguity, Learning, and Asset Returns,” *Econometrica*, 80, 559–91.
- KAHNEMAN, D. (2011): *Thinking, Fast and Slow*. Doubleday.
- KAHNEMAN, D., AND D. LOVALLO (1993): “Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking,” *Management Science*, 39, 17–31.
- KASA, K. (2012): “A Behavioral Defense of Rational Expectations,” Working Paper 12-05, Dept. of Economics, Simon Fraser University.
- KUTOYANTS, Y. A. (2004): *Statistical Inference for Ergodic Diffusion Processes*. Springer-Verlag.
- MAENHOUT, P. J. (2004): “Robust Portfolio Rules and Asset Pricing,” *Review of Financial Studies*, 17, 951–83.
- MALMENDIER, U., AND S. NAGEL (2011): “Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?,” *Quarterly Journal of Economics*, 126, 373–416.
- MANGALAM, M. (2025): “The Myth of the Bayesian Brain,” *European Journal of Applied Physiology*, pp. doi.org/10.1007/500421-025-05855-6.
- MEHRA, R., AND E. PRESCOTT (1985): “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 15, 145–61.
- PARR, T., G. PEZZULO, AND K. J. FRISTON (2022): *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press.
- PASTOR, L., AND P. VERONESI (2003): “Stock Valuation and Learning About Profitability,” *Journal of Finance*, 58, 1749–89.
- SAVAGE, L. J. (1972): *The Foundations of Statistics*. Dover Publications, second revised edn.
- SHAROT, T., R. KANEI, D. MARSTON, C. KORN, G. REES, AND R. DOLAN (2012): “Selectively Altering Belief Formation in the Human Brain,” *Proceedings of the National Academy of Sciences*, 109, 17058–62.
- SHILLER, R. J. (1981): “Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?,” *American Economic Review*, 71, 421–36.
- TANIGUCHI, M., AND Y. KAKIZAWA (2000): *Asymptotic Theory of Statistical Inference for Time Series*. Springer-Verlag.
- WEITZMAN, M. L. (2007): “Subjective Expectations and Asset-Return Puzzles,” *American Economic Review*, 97, 1102–30.
- WOLFORD, G., M. B. MILLER, AND M. GAZZANIGA (2000): “The Left Hemisphere’s Role in Hypothesis Formation,” *Journal of Neuroscience*, 20, 1–4.

DEPARTMENT OF ECONOMICS, SIMON FRASER UNIVERSITY, 8888 UNIVERSITY DR., BURNABY BC V5A 1S6

Email address: kkasa@sfu.ca

URL: <https://www.sfu.ca/~kkasa/>

DEPARTMENT OF ECONOMICS, UNIVERSITY OF GUELPH, 50 STONE RD E, GUELPH, ON CANADA. N1G 2W1

Email address: leix@uoguelph.ca

URL: <http://www.uoguelph.ca/lang/people/xiaowen-lei>