DISCUSSION PAPER SERIES

No. 8514

ASSET PRICING UNDER RATIONAL LEARNING ABOUT RARE DISASTERS

Christos Koulovatianos and Volker Wieland

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP8514.asp**

ISSN 0265-8003

ASSET PRICING UNDER RATIONAL LEARNING ABOUT RARE DISASTERS

Christos Koulovatianos, Nottingham University
Volker Wieland, Goethe-Universität Frankfurt and CEPR

Discussion Paper No. 8514
August 2011

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 8514

August 2011

ABSTRACT

Asset Pricing under Rational Learning about Rare Disasters*

This paper proposes a new approach for modeling investor fear after rare disasters. The key element is to take into account that investors' information about fundamentals driving rare downward jumps in the dividend process is not perfect. Bayesian learning implies that beliefs about the likelihood of rare disasters drop to a much more pessimistic level once a disaster has occurred. Such a shift in beliefs can trigger massive declines in price-dividend ratios. Pessimistic beliefs persist for some time. Thus, belief dynamics are a source of apparent excess volatility relative to a rational expectations benchmark. Due to the low frequency of disasters, even an infinitely-lived investor will remain uncertain about the exact probability. Our analysis is conducted in continuous time and offers closed-form solutions for asset prices. We distinguish between rational and adaptive Bayesian learning. Rational learners account for the possibility of future changes in beliefs in determining their demand for risky assets, while adaptive learners take beliefs as given. Thus, risky assets tend to be lower-valued and price-dividend ratios vary less under adaptive versus rational learning for identical priors.

JEL Classification: C11, C61, D81, D83, D91, E21 and G11
Keywords: adaptive learning, asset pricing, Bayesian learning, beliefs, controlled diffusions and jump processes, learning about jumps and rational learning

Christos Koulovatianos
School of Economics
University of Nottingham
University Park, PF H32
Nottingham

Email: christos.koulovatianos
@nottingham.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=157640

Volker Wieland
Chair for Monetary Theory and Policy
Goethe-Universität Frankfurt
House of Finance, PF H 31
D-60323 Frankfurt
GERMANY

Email:
wieland@wiwi.uni-frankfurt.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=152242

* We thank Ravi Bansal, Luca Benzoni, John Campbell, Satyajit Chatterjee, Emmanuel Farhi, Sergey Slobodyan and Pietro Veronesi for useful discussions, comments, and suggestions. Of course, all remaining errors are our own.

Submitted 28 July 2011

1. Introduction

In the past twenty years, stock prices in the United States and other markets around the world experienced one boom and two busts. The boom took place in the second half of the 1990s, the first bust in the year 2000 and the second one at the end of 2008. **Figure 1** shows that, in the US market, fundamentals have played a role in both bust episodes. In both cases, dividends and earnings exhibited a drop of 20 to 30 percent within a short period of time. Strikingly, price-dividend (P-D) and price-earnings (P-E) ratios started to decline massively shortly after the drop in dividends. The fall in prices over and above the reduction in dividends or earnings was particularly pronounced and rapid after the 2008 episode. Clearly, achieving a better understanding of the factors that drive the movement of asset prices following a rare stock market crash is of great importance not only for researchers and investors, but also for policymakers keen to assess the extent of negative impact on overall economic activity. For example, they may wonder whether these asset price movements reflect information about the duration and frequency of such crashes, or whether they are driven by irrational fear and panic among investors.

This paper presents a new approach for modeling investor fear under uncertainty about the likelihood of rare disasters. It relates two different literatures on asset-pricing that have proceeded mostly separately from one another. First, there is the large literature on learning under parameter uncertainty. It recognizes that market participants lack knowledge of many key parameters characterizing financial markets. In a recent survey, Pastor and Veronesi (2009a), point out that "*many facts (in financial markets) that appear baffling at first sight seem less puzzling once we recognize that parameters are uncertain and subject to learning*".[1] The other literature aims to explain asset pricing puzzles as a consequence

[1] Recent studies include Pastor and Veronesi (2009b), who investigate the emergence of bubbles when average productivity of a new technology is uncertain and subject to learning, and Weitzman (2007), who argues

of disaster risk while maintaining the assumption of rational expectations. First proposed by Rietz (1988), this idea received renewed interest once Barro (2006, 2009) showed that empirically plausible disaster probabilities provide a powerful explanation of historical equity premia.[2] Weitzman (2007), however, questions the disaster risk literature by pointing to the "*inherent implausibility of being able to meaningfully calibrate rational-expectations-equilibria objective frequency distributions of rare disasters because the rarer the event the more uncertain is our estimate of its probability.*"

Our paper addresses Weitzman's criticism head-on and incorporates parameter uncertainty and Bayesian learning in a disaster-risk asset-pricing model. Similarly to Longstaff and Piazzesi (2004), Barro (2006, 2009) and Weitzman (2007), we use the Lucas (1978) exchange-economy asset-pricing model as a vehicle for conducting our analysis. As in Longstaff and Piazzesi (2004) and Barro (2006, 2009) we assume that dividends follow a jump-diffusion process in continuous time. This process includes a standard Brownian motion with drift that is interrupted by rare downward jumps. In the disaster risk literature, the probability that such a crash occurs within a given period of time may be fixed or time-varying but its stochastic properties are always assumed known to investors. We refer to this measure of the frequency of disasters as the hazard rate and consider a time-variable setup with the hazard rate switching between a high and low value at a given probability.[3] Under

that learning about the parameter that controls the spread of the distribution of future consumption growth helps explain equity premia and excess volatility relative to the rational expectations benchmark.

[2] Earlier on, Longstaff and Piazzesi (2004) had also calibrated a Rietz-type model with large downward dividend jumps using data from the Great Depression. More recently, Barro and Ursua (2008) put together a large international dataset on consumption and disasters and Barro et al. (2010) provided new estimates of the variability and persistence of disaster risk. Gabaix (2008, 2010), Wachter (2011) and Gourio (2008a,b) show that variable disaster risk serves to explain excessive volatility of price-dividend ratios. See also, LeRoy (2008) for a survey on asset-pricing excess volatility and existing approaches for diagnosing it in the data.

[3] The inverse of the hazard rate is the number of periods it takes for a jump to occur on average. So, if a hazard rate is high, then rare disasters are more frequent on average, meaning that any underlying sources of rare disasters create a more hazardous environment. Technically, nature's true disaster-shock process is a random mixture of two Poisson processes, one with a high (constant) hazard rate and one with a low one. The probability that the disaster is drawn from the high or low hazard rate process is also constant.

rational expectations—the starting point of our analysis—this probability is known, and consequently also the average hazard rate, that is the average frequency of disasters. Following Weitzman's critique, we proceed by treating this probability as unknown and model investors' beliefs and learning about the average hazard rate explicitly.[4] In other words, we consider Bayesian learning about the key parameter governing the frequency of disasters.

Bayes' rule implies that investors' beliefs abruptly drop to a more pessimistic level following a stock market crash. Investors suddenly fear that such disasters will occur much more frequently in the future than they had thought in the past. Here, pessimism or fear is not meant to suggest investor irrationality. Rather, increased pessimism simply means that investors' perceived value of the probability assigned to the high hazard rate case has risen. The exact definition of investor rationality under Bayesian learning will be laid out further below. Over time, investors revise their beliefs by repeatedly applying Bayes' rule. Bayesian learning makes efficient use of historical information and new data. In the absence of another crash, beliefs slowly turn more optimistic and learning implies a smooth reduction in the perceived probability of the high-hazard-rate case. Thus, investors' pessimistic beliefs exhibit a certain degree of persistence after a disaster has occurred.

In our model, asymptotic beliefs are unbiased. However, even infinitely-lived investors would never reach full confidence about the average frequency of disasters, as would be the case under rational expectations. This result is due to the slow arrival of information about the frequency of rare disasters. Despite using Bayes' rule for updating priors, posteriors never catch up with nature's parameters with infinite precision, even if an infinitely long history of actual data has been processed. These belief dynamics lend additional support to Weitzman's claim that analysis of rare-disaster-risk requires the modeling of subjective

[4] An alternative approach to modeling the impact of uncertainty about rare events on asset prices is offered by robust control in the presence of Knightian model uncertainty (see Liu, Pan and Wang (2005) for an implementation).

expectations and investor learning.

Asset prices depend on investors' beliefs about the unknown parameter governing the frequency of rare disasters. We solve analytically for the relationship between asset prices and beliefs regarding disaster probability. Using this asset pricing formula, we find that the abrupt increase in pessimism following a crash causes a sudden drop in price-dividend ratios. The extent of the decline in the asset price over and above the fall in dividends is entirely due to the shift in beliefs. Under rational expectations, jumps in prices and dividends would be proportional and the ratio would remain the same.[5] The subsequent persistence in pessimistic beliefs implies that the price-dividend ratio remains depressed for some time. It recovers slowly as long as no other crash occurs and investors' beliefs assign successively lower probabilities to the high-frequency-disaster case.

The link between asset prices and beliefs is derived by solving the dynamic optimization problem of the representative investor/household in our asset pricing model. In this context, we distinguish between a fully rational investor and one who learns in an adaptive fashion. Both, the rational and the adaptive learner base their decision on current beliefs regarding disaster probabilities that were obtained by applying Bayes rule to available data. The distinction between rational and adaptive Bayesian learning depends on whether or not the decision maker takes into account the dynamic transition equations of beliefs in her optimization problem, in addition to the other recursions governing laws of motion of state variables such as the dividend process. In other words, a rational learner knows that her beliefs will change in the future as new information arrives. In particular, as long as no

[5] Here our setup differs from other studies that introduce time-variable disaster risk in rational-expectations asset-pricing models (cf. Gabaix (2008, 2010), Gourio (2008b) and Wachter (2011)). While their specifications of time-variable disaster risk are useful for explaining excess volatility under rational expectations, we aim to show that such variations in P-D ratios could even be exclusively due to changes in investors' subjective perceptions of disaster risk. Our setup implies a constant P-D ratio under rational expectations, because the average hazard rate is known and independent of past developments.

crash occurs the perceived hazard rate will smoothly decline. The adaptive learner acts as if her beliefs will never change. Only as time advances, she re-calculates her estimate of the average hazard rate according to Bayes rule.[6] In this manner, adaptive Bayesian learning represents a well-defined deviation from fully rational behavior. Any difference between asset prices under adaptive versus rational Bayesian learning could then be characterized as being due to overly pessimistic or optimistic views.[7]

Under certain plausible conditions, we find that asset prices under rational learning are always higher than prices that follow from the behavior of adaptively learning investors for any given prior belief. This finding can be attributed to the fact that rational learners take into account that their estimates of disaster probabilities will change in the future. Specifically, in the absence of another crash they anticipate the gradual emergence of a more optimistic outlook. Thus, they demand more of the risky asset.

A recent paper that also investigates Bayesian learning about rare jumps is Benzoni, Colline-Dufresne and Goldstein (2011). These authors aim to explain the dramatic and lasting steepening of the implied volatility curve for equity index options after the 1987 stock market crash despite minimal changes in aggregate consumption. Similar to our approach, they consider learning about high versus low disaster intensity, but in their model jumps are in perceived dividends while the actual process is smooth. Benzoni et al (2011) provide a

[6] For other work distinguishing adaptive and rational Bayesian learning see Guidolin and Timmermann (2007), Cogley and Sargent (2008), and Koulovatianos, Mirman, and Santugini (2009). Adaptive learning reflects the anticipated utility concept studied, e.g., by Kreps (1998), Cogley and Sargent (2008), and Koulovatianos and Wieland (2011). Rational Bayesian learning may involve active experimentation, for example in the presence of multiplicative parameter uncertainty (see Mirman, Urbano and Samuelson (1993), Wieland (2000a,b) and Beck and Wieland (2002)). Recent work on asset pricing explores the role of Bayesian learning in booms and busts (see Benhabib and Dave (2011) for adaptive learning and Adam and Marcet (2010) for rational beliefs). Bansal and Shaliastovich (2011) present a model in which income and dividends are smooth but asset prices exhibit large moves. These jumps arise from rational learning by investors about an unobserved state.

[7] Our model with subjective beliefs about disaster risk and learning may also offer a more useful reference point for comparison with behavioral finance research on the consequences of investor sentiment and overreaction (cf. Barberis et al. (1998)) than the standard rational expectations benchmark.

numerical approximation to the solution for a given parameterization of their model, while we obtain an analytical solution to our model and derive price-dividend ratios explicitly as a function of investor beliefs. Furthermore, we distinguish between rational and adaptive learning and investigate the pricing implications.

Furthermore, we discovered an older yet unpublished study by Comon (2001) which also introduces rational Bayesian learning with extreme events.[8] Comon introduces parameter uncertainty regarding the hazard rate of rare dividend jumps in a variant of the Cox, Ingersoll and Ross (1985) exchange economy. Contrary to our approach, he assumes that prior subjective hazard rates of investors are Gamma distributed. One consequence of his framework is that learning only matters in influencing price-dividend ratios during the transition to rational expectations, which complicates empirical identification.

Finally, we calibrate our model and conduct dynamic simulations that illustrate the model's potential to capture key elements of the dynamic path of price-dividend ratios following the two crashes in the U.S. stock market in 2000 and 2008 shown in **Figure 1**. The calibration requires setting an initial prior belief on the average hazard rate of disasters. It turns out that it is possible to generate sudden drops followed by slowly improving P-D ratios under rational and adaptive learning. However, the adaptive learning simulation requires a prior belief that is roughly twice as optimistic as under rational learning. Since beliefs are the main driver of P-D dynamics in our model, it is of great interest to compare the behavior of model beliefs with survey data on investors' perception of the threat of a crash. Fortunately, such data is available in the form of Robert Shiller's Crash Confidence Index. The questionnaire underlying this data is explained in Shiller, Kon-Ya and Tsutsui (1996). Interestingly, our simulations of the boom and busts in the U.S. stock market are

[8] We are grateful to Pietro Veronesi for mentioning it in commenting on the first version of our paper and to Comon's adviser at Harvard, John Campbell, for scanning and sending us chapter 1 of his dissertation.

broadly consistent with dynamics of the beliefs indicated by the survey.

The remainder of the paper proceeds as follows. Section 2 presents the asset pricing model under rational expectations and demonstrates our solution approach. In Section 3, we then analyze and compare the decision making of households/investors that learn in an adaptive or rational Bayesian fashion. In the fourth section, closed-form solutions for asset prices are derived. To illustrate the power of the model to fit the behavior of price-dividend ratios and survey measures of beliefs following the last two crashes in the U.S. stock market, we present dynamic simulations conditional on given priors and the timing of these two busts. An extension to Epstein-Zin preferences, which do not restrict the intertemporal elasticity of substitution to be equal to the inverse of the coefficient of relative risk aversion, is discussed in Section 5. It is meant to address concerns regarding the special nature of the standard asset-pricing model with constant relative risk aversion (CRRA) preferences (see, for example, Barro (2009)and Wachter (2011)). Section 6 concludes.

2. Benchmark Model without Learning

The model is a simple representative-agent Lucas (1978) tree economy with rare jumps in the dividend process in the spirit of Barro (2006), analyzed in continuous time. The dividend process is given by,

$$\frac{dD(t)}{D(t)} = \mu dt + \sigma dz(t) + dq(t) \ , \tag{1}$$

in which $dz(t)$ is a standard Brownian motion, i.e., $dz(t) = \varepsilon(t)\sqrt{dt}$, with $\varepsilon(t) \sim N(0,1)$, for all $t \geq 0$. Moreover, $q(t)$ is a Poisson process driving random downward jumps in dividends of size $\zeta \cdot D(t)$, where $\zeta \in (0,1)$ is a random variable with given time-invariant distribution having compact support, $\mathcal{Z} \subset (0,1)$. In particular, the Poisson process $q(t)$ is characterized by,

$$dq(t) = \begin{cases} -\zeta & \text{with Probability } \lambda(t)\,dt \\ 0 & \text{with Probability } 1-\lambda(t)\,dt \end{cases} \tag{2}$$

The hazard rate of the Poisson process, $\lambda(t)$ in equation (2), is also random, taking two values only.[9] In particular, the density of λ is characterized by the binomial distribution,

$$\lambda(t) = \begin{cases} \lambda_h & \text{with Probability } \pi^* \\ \lambda_l & \text{with Probability } 1-\pi^* \end{cases}, \tag{3}$$

for all $t \geq 0$, and with $\lambda_h > \lambda_l > 0$, and $\pi^* \in (0,1)$. We assume that variables $z(t)$, $q(t)$, $\lambda(t)$, and ζ are all independent from each other at all times.

Consistently with the growing variable-disaster-risk literature (see, for example, Gabaix (2010) and Wachter (2011)) our variable-hazard-rate also assumes that *exogenous factors driving the hazard rate of rare disasters change over time.*[10] The extension that investors may be unaware of the frequency of such shifts in hazard rates, π^*, is a plausible environment of learning, which is the learning application we pursue in later sections.

Under rational expectations, at any instant, $t \geq 0$, two independent random events related to disasters are revealed. The first event is whether a rare disaster has occurred at t or not. The second event revealed is the regime of riskiness from which the rare-disaster event at t has been drawn (λ_h indicates rare-disaster risk coming from a more hazardous exogenous regime compared to the case of λ_l). Given these assumptions, Proposition 1 clarifies the stochastic law of motion governing the dynamics of the Poisson process $q(t)$.[11]

[9] By the term "hazard rate" we mean the average number of downward jumps in the dividend process per unit of time. The interpretation of the inverse of the hazard rate is the number of periods it takes for a jump to occur on average. So, if a hazard rate is high, then rare disasters are more frequent on average, meaning that any underlying sources of rare disasters create a more hazardous environment. Barro (2006) and Wachter (2011) both use the term "disaster probability" for expressing what is our hazard-rate concept throughout this paper.

[10] The stochastic process driving hazard-rate variability over time in Wachter (2011) and Gabaix (2010) is different from ours.

[11] For an introduction to Poisson processes and their application to continuous-time optimization modeling

Proposition 1 *Under the stochastic structure of hazard rates given by (3), the dynamics of* $q(t)$ *implied by nature, are governed by the stochastic equation,*

$$dq(t) = \begin{cases} -\zeta & \text{with Probability } \lambda^* dt \\ 0 & \text{with Probability } 1 - \lambda^* dt \end{cases} \tag{4}$$

where $\lambda^* \equiv E_{\pi^*}(\lambda) = \pi^* \lambda_h + (1 - \pi^*) \lambda_l$ *is the expected hazard rate.*

Proof See Appendix A. □

By construction, nature's stochastic process $\lambda(t)$ has independent increments. Equation (4) in Proposition 1 clarifies that, no matter which hazard rate applies at time $t \geq 0$ (i.e., λ_h vs. λ_l), at time instant $t + dt$, the hazard rate that applies for forward-looking decision making is equal to the hazard-rate expectation driven by parameter π^*, which is equal to $E_{\pi^*}(\lambda) = \pi^* \lambda_h + (1 - \pi^*) \lambda_l$. Given that $\lambda(t)$ and all other random variables comply with the independent-increments assumption, dividend-process innovations exhibit no persistence over time. In the representative-agent Lucas-tree-fruit economy consumption equals dividends in equilibrium. So, asset pricing under rational expectations fully reflects this feature of serially uncorrelated dividend innovations. By contrast, once we introduce rational learning, Bayes' rule implies that beliefs exhibit persistence, which is then reflected in pricing. So, our assumption that dividend innovations are serially uncorrelated by nature help in conveying the special role of belief-driven persistence introduced to asset pricing through Bayesian learning about disasters.

The Lucas-tree-fruit economy is an exchange economy of a large number of identical infinitely-lived agents of total mass equal to one. Agents trade only one risky asset, the market portfolio, which has returns given by equation (1). The representative agent maximizes

see, for example, Merton (1971, pp. 395-401), Kushner (1967, Ch. 9), Dreyfus (1965, pp. 470-472), Karlin and Taylor (1981, Ch. 16), Cox and Miller (1965, Ch. 9), Papoulis (1991, pp. 367-375) and Dixit and Pindyck (1994, pp. 85-87).

her expected lifetime utility given by,

$$E_0 \left[\int_0^\infty e^{-\rho t} \frac{c(t)^{1-\gamma} - 1}{1-\gamma} dt \right] \tag{5}$$

with $c(t)$ denoting an individual's consumption, with $\rho > 0$ being the rate of time preference, and with $\gamma \geq 0$ being the coefficient of relative risk aversion (the special case of $\gamma = 0$ corresponds to a risk-neutral investor, and is possible to be studied in the context of a representative-agent Lucas-tree economy). At any time $t \geq 0$, an individual holds $s(t) \geq 0$ shares of the risky asset. At time $t = 0$, the aggregate supply of the asset is $S(0) > 0$, and there is no new issuing of shares, so $S(t) = S(0)$ for all $t \geq 0$. Moreover, at time $t = 0$, the endowment of a representative individual is $s(0) = S(0)$. The budget constraint in continuous time is,[12]

$$ds(t) = \frac{1}{P(t)} [s(t) D(t) - c(t)] \, dt \ . \tag{6}$$

In what follows we construct the Hamilton-Jacobi-Bellman (HJB) equation in order to achieve two goals. First, the HJB equation introduces a recursive language that efficiently describes the distinction made between the adaptive learner vs. the rational learner regarding the way the two utilize new information in their forward-looking decision problem. Second, the HJB equation introduces a common solution technique through undetermined coefficients that tackles three problems: (a) the rational-expectations (RE) agent economy, (b) the adaptive-learner (AL) economy, and (c) the rational-learner (RL) economy.

[12] For the derivation of equation (6) from its discrete-time counterpart, which is

$$P_t s_t = s_{t-1} (P_t + D_t) - c_t \ ,$$

notice that the above equation can be re-written as

$$\Delta s_t \equiv s_t - s_{t-1} = \left(s_{t-1} \frac{D_t}{P_t} - \frac{c_t}{P_t} \right) \Delta t \ ,$$

where $\Delta t = 1$ under the convention that the discrete-time period length is unity. In continuous time, taking the limit $\Delta t \to 0$ results in (6) .

In order to solve the pricing problem an individual agent must determine her demand for the risky asset at any time $t \geq 0$. So, given any possible path $(P(t))_{t\geq 0}$ generated by a price function with $P(t) > 0$ for all $t \geq 0$, the agent must pick the paths $(s(t), c(t))_{t\geq 0}$ that maximize her utility given by (5), subject to (6) and (1). Yet, the determination of these demand functions is a stationary discounted dynamic programming problem that can be solved through a recursive functional-choice problem given by a HJB equation with no time indices.

In order to formulate the HJB equation of RE, we assume a pricing rule,

$$P = \Psi^{RE}(D) \ ,$$

and we denote a value function that is subject to the pricing rule $\Psi^{RE}(D)$ by $J^{RE}\left(s, D \mid \Psi^{RE}\right)$. The HJB equation is,

$$\rho J^{RE}\left(s, D \mid \Psi^{RE}\right) = \max_{c\geq 0}\left\{\frac{c^{1-\gamma}-1}{1-\gamma} + J_s^{RE}\left(s, D \mid \Psi^{RE}\right)\cdot\left[\frac{1}{\Psi^{RE}(D)}(sD-c)\right] + \right.$$
$$+ J_D^{RE}\left(s, D \mid \Psi^{RE}\right)\cdot \mu D + J_{DD}^{RE}\left(s, D \mid \Psi^{RE}\right)\frac{(\sigma D)^2}{2} +$$
$$\left. + \lambda^*\left\{E_\zeta\left[J^{RE}\left(s, (1-\zeta)D \mid \Psi^{RE}\right)\right] - J^{RE}\left(s, D \mid \Psi^{RE}\right)\right\}\right\} \ , \qquad (7)$$

with E_ζ denoting the expectations operator with respect to the random variable ζ only. In Appendix A we derive an analytical solution for the value function $J^{RE}\left(s, D \mid \Psi^{RE}\right)$, and show that, as reported by Barro (2006),[13]

$$P = \Psi^{RE}(D) = \frac{1}{\rho - \chi - \lambda^*\xi}\cdot D \ , \qquad (8)$$

where

$$\chi \equiv (1-\gamma)\left(\mu - \gamma\frac{\sigma^2}{2}\right) \ , \quad \text{and} \quad \xi \equiv E_\zeta\left[(1-\zeta)^{1-\gamma}\right] - 1 \ .$$

[13]Our parameter λ^* in the pricing formula matches up with parameter p in equation (17), page 839, in Barro (2006).

In the special case of $\gamma = 1$, $\chi = \xi = 0$, and the pricing function given by (8) implies $\Psi^{RE}(D) = D/\rho$. This means that the presence of risk does not affect pricing, no matter if this risk stems from the diffusion or from the jump process. In general,

$$\xi \gtreqless 0 \Leftrightarrow \gamma \gtreqless 1 \ . \tag{9}$$

The parametric relationship given by (9) implies that increasing the expected hazard rate, λ^*, decreases prices only if $\gamma < 1$. If $\gamma > 1$, which can be loosely interpreted as having higher risk aversion, increasing the expected hazard rate, λ^*, increases prices. This paradoxical result has been discussed by Bansal and Yaron (2004, p. 1487), and also by Barro (2009, p. 249). Both of these studies attribute the paradox to the fact that, with power utility, the coefficient of relative risk aversion and the elasticity of intertemporal substitution cannot be disentangled, as the one equals the reciprocal of the other. As a resolution to this rigid feature of CRRA preferences, the studies by Bansal and Yaron (2004) and Barro (2009) suggest the use of Epstein-Zin (1989) utility functions. In our learning application, expected hazard rates will be moving over time together with beliefs, so a parameter-value choice $\gamma < 1$ vs. $\gamma > 1$ becomes important for some additional mechanics related to learning. To tackle such calibration concerns we discuss possible extensions to Duffie-Epstein (1992a,b) preferences in Section 5.

3. Bayesian Learning about the Likelihood of Jumps

3.1 Characterization of Beliefs

Our formulation of beliefs aims at retaining analytical tractability when incorporating belief dynamics implied by Bayes rule. Thus, we take only a small but still very influential step away from rational expectations. Our investors cannot observe which hazard rate (λ_h vs. λ_l) is triggered by nature at any point in time. An investor observes stock market crashes but is unable to collect data on histories of hazard-rate realizations (dates and number of instances in which λ_h vs. λ_l have been triggered in the past). Given this ignorance, instead of considering π^* with infinite precision, at any time $t \geq 0$ the investor has subjective priors denoted by $\pi(t)$. We maintain that the investors' perceived hazard rate corresponds to the subjectively weighted average of λ_h or λ_l, with $\lambda_h > \lambda_l > 0$. In addition, she remains informed of the distribution of ζ, which is time-invariant. In sum, at any time $t \geq 0$, a learning investor's beliefs about random hazard rates are given by the Bernoulli distribution

$$\tilde{\lambda}(t) = \begin{cases} \lambda_h & \text{with Probability } \pi(t) \\ \lambda_l & \text{with Probability } 1 - \pi(t) \end{cases}, \tag{10}$$

where the tilde denotes random variables governed by distributions that depend on subjective beliefs about some parameters, as opposed to random variables governed by distributions that exclusively depend on nature's true parameters. The tilde refers to an investor's subjective perception about her beliefs, before new information has been revealed. Since there is a distinction between dynamics of ex-post beliefs (after information has arrived) and ex-ante beliefs (before information has arrived), a tilde denotes ex-ante beliefs. Following the same argument as the one in the proof of Proposition 1, we can directly state that, based on (10), the subjective perception of process $q(t)$ by the investor at time $t \geq 0$ is,

$$d\tilde{q}(t) = \begin{cases} -\zeta & \text{with Probability } \Lambda(\pi(t))\,dt \\ 0 & \text{with Probability } 1 - \Lambda(\pi(t))\,dt \end{cases}, \tag{11}$$

where

$$\Lambda(\pi(t)) \equiv \lambda_h \pi(t) + \lambda_l [1 - \pi(t)] \ , \tag{12}$$

i.e., perceived hazard rate equals the expected hazard rate according to priors $\pi(t)$, $E_\pi\left(\tilde{\lambda}(t)\right) = \Lambda(\pi(t))$. Notice the special case resulting from the definition given by (12), that $\lambda^* = \Lambda(\pi^*)$.

A similar characterization of beliefs has been used by Keller and Rady (2010). They study a game-theoretical application of learning with Poisson differential equations (Poisson bandits) analytically, while we analyze the optimal control problem of an investor in an asset pricing model with additional states due to dividends and shares.

3.2 Bayesian updating of prior beliefs on disaster risk

New information about disasters is restricted to observing whether a disaster has occurred or not. So, after new information has been incorporated, the only variable putting into motion beliefs about the perceived average hazard rate, $\Lambda(\pi(t)) = \lambda_h \pi(t) + \lambda_l [1 - \pi(t)]$, is probability $\pi(t)$. Proposition 2 characterizes the Bayesian-learning dynamics of $\pi(t)$.

Proposition 2 *After applying Bayes' rule, from a modeler's perspective, the dynamics of the posterior belief about π, based on the prior $\pi(t)$, are comprehensively characterized by the Poisson differential equation*

$$d\pi = -\delta\pi(1-\pi)\,dt + dq_\pi \ , \tag{13}$$

where $\delta \equiv \lambda_h - \lambda_l$, and

$$dq_\pi = \begin{cases} \frac{\lambda_h \pi}{\Lambda(\pi)} - \pi & \text{with Probability } \lambda^* dt \\ 0 & \text{with Probability } 1 - \lambda^* dt \end{cases} . \tag{14}$$

Proof See Appendix B. □

What equation (14) reveals is that beliefs about π jump from π to $\lambda_h \pi / \Lambda(\pi)$ whenever a jump in the dividend process occurs (with the average frequency of such events driven by nature's average hazard rate $\lambda^* = \Lambda(\pi^*)$). Whenever a jump does not occur, (13) reveals that average-hazard-rate beliefs, $\Lambda(\pi(t))$, decay towards λ_l, since the term $-\delta\pi(1-\pi)\,dt$ of

(13) considered alone, implies that $\pi(t) \to 0$ as $t \to \infty$. Yet, jumps in D affect the term dq_π in (13), and make π jump upwards, keeping $\pi(t)$ in the interior of the open interval $(0,1)$.[14]

Notice that equation (14) refers to learning from the modeler's (or nature's) perspective, which is the reason why nature's true parameter, λ^* has been used. The derivation of the law of motion for π when nature's true parameter, π^*, is unknown is the same as in the case of deriving equations (13) and (14). So, the only feature changing after this new derivation is that, from a learner's perspective who makes decisions inside a model, the perception of the average hazard rate for a rare disaster is $\Lambda(\pi)$ instead of $\lambda^* = \Lambda(\pi^*)$. Consequently, in (14) nature's true average hazard rate, λ^*, must be replaced by the average subjective hazard rate $\Lambda(\pi)$. Corollary 1 summarizes this observation, and we state it without proof, because the reasoning is exactly the same as in the proof of Proposition 2.

Corollary 1 *The dynamics of ex-ante perceived beliefs about π by Bayesian learning agents (denoted by $d\tilde{\pi}$), are given by,*

$$d\tilde{\pi} = -\delta\pi(1-\pi)\,dt + d\tilde{q}_\pi \ , \tag{15}$$

where

$$d\tilde{q}_\pi = \begin{cases} \frac{\lambda_h \pi}{\Lambda(\pi)} - \pi & \text{with Probability } \Lambda(\pi)\,dt \\ 0 & \text{with Probability } 1-\Lambda(\pi)\,dt \end{cases} . \tag{16}$$

Proposition 2 and Corollary 1 are equivalent to an application of optimal filtering from observations of point processes in Liptser and Shiryaev (2001). In fact, equations (15) and (16) correspond to equation (19.86) in Liptser and Shiryaev (2001, pp. 332-3). However, we provide a more intuitive proof that was derived independently, because we had only been

[14] In the next section we prove that that $\pi(t) \in (0,1)$ for all $t \geq 0$, and tends to move around π^* on average as $t \to \infty$ (see Proposition 3 about the asymptotic behavior of $E[\pi(t)]$).

made aware of Liptser and Shiryaev's result used by Benzoni et al. (2011) after having completed the first draft of our paper.

3.3 Long-Run Dynamics of Learning

A crucial question concerning belief dynamics is whether nature's average hazard rate, λ^*, is learned in the long run in the sense of Muth's (1961) concept of asymptotic convergence to rational expectations. Proposition 3 answers this question formally.

Proposition 3 *Let beliefs about the occurrence of a jump event be given by equation (13). Then for all* $\pi(0) = \pi_0 \in (0,1)$, $\pi(t) \in (0,1)$ *for all* $t \geq 0$, *and*

$$\lim_{t\to\infty} E\left[\pi(t)\right] = \pi^*, \tag{17}$$

while

$$\lim_{t\to\infty} Var\left[\pi(t)\right] = \frac{\left\{\lambda^* - 4\lambda_l + \left[(\lambda^*)^2 + 8\lambda^*\lambda_h\right]^{\frac{1}{2}}\right\}^2}{16\delta^2} - (\pi^*)^2 > 0 \ . \tag{18}$$

Proof See Appendix B. □

Proposition 3 states that, after collecting infinite rare-disaster data drawn from nature's realizations, the beliefs of learning investors about π are asymptotically unbiased, but learners do not reach infinite precision about this limiting average parameter.The variance of belief parameter $\pi(t)$ is bounded away from 0 as indicated by equation (18). Thus, rare disasters arrive at such low frequency that they leave even an infinitely-lived learner uncertain about disaster risk, despite that Bayes' rule is a statistically efficient way of analyzing historical data. In other words, asymptotic convergence to rational expectations is not achieved.

The result of non-asymptotic convergence to rational expectations supports the idea that modeling investors in an environment of learning is plausible. Since investors never reach infinite subjective precision in our model, we do not need to analyze whether investors are at a particular stage during their overall learning process. Empirically capturing priors on belief parameter $\pi(t)$ alone, is sufficient to describe the implied dynamics of learning in our model. In Section 4 we use survey data collected through a questionnaire described in Shiller et al. (1996) that approximate belief belief parameter $\pi(t)$ in order to calibrate our model.

4. Asset Pricing under Adaptive versus Rational Learning

The conceptual distinction between an adaptive learner (AL) and a rational learner (RL) is based on how the decision maker accounts for her own ignorance regarding π^*. AL is aware of her ignorance at the current time instant. However, she simply assumes that her beliefs will not change in the future. By contrast, RL, apart from being aware of her ignorance at the current time, is also aware of the future evolution of her beliefs according to Bayes' rule. So, RL approaches her lack of information in a fully rational manner, while AL is boundedly rational, because she neglects the knowledge that beliefs will be revised in the future with new data.

A direct way of distinguishing adaptive and rational Bayesian learning mathematically is by setting up the investor's decision-making problem by means of the Hamilton-Jacobi-Bellman (HJB) equations. HJB equations are optimization recursions incorporating and accommodating other recursions that reflect the stochastic structure of a problem. Corollary 1 in Section 3 identifies the recursions associated with Bayesian updating, i.e., equations (15) and (16). While equations (15) and (16) will be included in the HJB equation associated with RL, they will not figure in the HJB equation of AL. In the following, we start with AL's optimization problem and then proceed to RL's problem, which is more demanding from a technical perspective.

4.1 Asset Demand by an Adaptive Learner

At any time $t \geq 0$ prior beliefs are given by $\pi(t)$. The HJB equation of AL is,

$$\rho J^{AL}\left(s, D, \pi \mid \Psi^{AL}\right) = \max_{c \geq 0} \left\{ \frac{c^{1-\gamma} - 1}{1-\gamma} + J_s^{AL}\left(s, D, \pi \mid \Psi^{AL}\right) \cdot \left[\frac{1}{\Psi^{AL}(D, \pi)} (sD - c) \right] + \right.$$
$$+ J_D^{AL}\left(s, D, \pi \mid \Psi^{AL}\right) \cdot \mu D + J_{DD}^{AL}\left(s, D, \pi \mid \Psi^{AL}\right) \frac{(\sigma D)^2}{2} +$$
$$\left. + \Lambda(\pi) \left\{ E_\zeta \left[J^{AL}\left(s, (1-\zeta) D, \pi \mid \Psi^{AL}\right)\right] - J^{AL}\left(s, D, \pi \mid \Psi^{AL}\right) \right\} \right\}, \quad (19)$$

given a pricing rule $P = \Psi^{AL}(D, \pi)$, and with the dynamics of π being given by equation (13). First, notice that, unlike the case of RE's value function, AL's value function, $J^{AL}\left(s, D, \pi \mid \Psi^{AL}\right)$, now depends on prior beliefs, π. Second, notice that AL does not anticipate any learning in the future, despite that beliefs are continuously updated over time. This non-anticipation of learning by AL is captured by the fact that perceived dynamics of π (captured by equations (15) and (16)) are not incorporated in AL's HJB equation as it can be seen by (19). This absolute form of non-anticipation implies that, in the particular application with jumps and hazard rates that we are examining, at any instant $t \geq 0$, AL's perception of the average hazard rate is $\Lambda(\pi(t))$, and this perception is not anticipated to change in the future. This extreme concept of non-anticipation of learning is in accordance with the concept of anticipated utility implied by particular belief priors, analyzed, for example, in Kreps (1998) and Cogley and Sargent (2008).

4.2 Asset Demand by a Rational Learner

The HJB equation of RL is,

$$\rho J^{RL}\left(s, D, \pi \mid \Psi^{RL}\right) = \max_{c \geq 0} \left\{ \frac{c^{1-\gamma} - 1}{1-\gamma} + J_s^{RL}\left(s, D, \pi \mid \Psi^{RL}\right) \cdot \left[\frac{1}{\Psi^{RL}(D, \pi)} (sD - c) \right] + \right.$$
$$+ J_D^{RL}\left(s, D, \pi \mid \Psi^{RL}\right) \cdot \mu D + J_{DD}^{RL}\left(s, D, \pi \mid \Psi^{RL}\right) \frac{(\sigma D)^2}{2} +$$

$$-J_{\pi}^{RL}\left(s, D, \pi \mid \Psi^{RL}\right) \cdot \delta\pi\left(1-\pi\right) +$$

$$+\Lambda\left(\pi\right)\left\{E_{\zeta}\left[J^{RL}\left(s, \left(1-\zeta\right)D, \frac{\lambda_h \pi}{\Lambda\left(\pi\right)} \mid \Psi^{RL}\right)\right] - J^{RL}\left(s, D, \pi \mid \Psi^{RL}\right)\right\}\Bigg\}\ , \quad (20)$$

given a pricing rule $P = \Psi^{RL}(D, \pi)$, and while the dynamics of π are driven by equation (13). What distinguishes RL's problem from this of AL is that the subjective dynamic stochastic equations (15) and (16) that calculate the instantaneous change in beliefs, $d\tilde{\pi}(t)$, are incorporated in RL's problem (RL's anticipation of learning).

Two algebraic terms in RL's HJB equation capture RL's anticipation of learning. First, equation (15) implies that the anticipated change in beliefs in case no jump occurs at the immediately "adjacent" future instant is equal to $-\delta\pi(1-\pi)\,dt$. The impact of the instantaneous change in beliefs in case of no jumps on lifetime utility is captured by the partial derivative of J^{RL} with respect to π, J_{π}^{RL}, which multiplies the differential term $\delta\pi(1-\pi)$ in equation (20). Second, equation (15) indicates that, in the case a jump indeed occurs at the immediately "adjacent" future instant, the anticipated change in beliefs is equal to $d\tilde{q}_{\pi} = \lambda_h \pi / \Lambda(\pi) - \pi$, which appears in equation (16). The impact of a jump in D on lifetime utility can be seen by the whole last additive term of equation (20). In this last additive term of (20), RL acknowledges the impact of the jump on both the level of the dividend, D, and on the level of beliefs, π.[15]

4.3 Asset Prices under Adaptive Learning

Since AL's HJB equation differs from RE's HJB only in that λ^* in RE's problem has been replaced by $\Lambda(\pi)$, after following all steps in the proof of RE's asset pricing problem in

[15]In the case of AL's decision problem, only the impact of the jump on D is taken into account, as it can be seen by the last additive term of equation (19). Notice also that the last additive term of equation (20) is the partial outcome of multi-dimensional stochastic integration with respect to dimensions D and π during the process of applying the expectations operator on the right-hand side of RL's HJB equation.

Appendix A, we arrive at the pricing rule

$$\Psi^{AL}(D,\pi) = \frac{1}{\rho - \chi - \Lambda(\pi)\xi} \cdot D \ , \tag{21}$$

while J^{AL} is given by,

$$J^{AL}\left(s, D, \pi \mid \Psi^{AL}\right) = \begin{cases} \frac{\Psi^{AL}(D,\pi)}{D}\frac{(sD)^{1-\gamma}}{1-\gamma} - \frac{1}{\rho(1-\gamma)} & , \quad \text{if } \gamma \neq 1 \\ \frac{\Psi^{AL}(D,\pi)}{D}\ln(sD) + \phi^{AL}\left(\pi \mid \Psi^{AL}\right) & , \quad \text{if } \gamma = 1 \end{cases} \tag{22}$$

where

$$\begin{aligned} \phi^{AL}\left(\pi \mid \Psi^{AL}\right) &= \frac{1}{\rho}\frac{\Psi^{AL}(D,\pi)}{D}\left\{\left(\mu - \frac{\sigma^2}{2}\right) + \Lambda(\pi)\, E_\zeta\left[\ln(1-\zeta)\right]\right\} = \\ &= \frac{\mu - \frac{\sigma^2}{2} + \Lambda(\pi)\, E_\zeta\left[\ln(1-\zeta)\right]}{\rho^2} \ . \end{aligned} \tag{23}$$

In particular, RE's asset pricing problem is the special case of AL's problem when setting $\pi = \pi^*$. Yet, a crucial difference between the two pricing schemes is that in RE's problem the perceived average hazard rate is λ^* for all $t \geq 0$, and thus the P-D ratio implied by (8), $\Psi^{RE}(D)/D$, is constant over time. On the contrary, concerning AL's problem if for some $\hat{t} \geq 0$, $\pi\left(\hat{t}\right) = \pi^*$, then $\Psi^{AL}\left(D, \pi\left(\hat{t}\right)\right)/D = \Psi^{RE}(D)/D$, but this equality does not hold generally, as beliefs, π, driven by the stochastic equation (13) change continuously and permanently over time.

4.4 Asset Prices under Rational Learning

In Appendix C we prove that the pricing rule implied by RL's problem is

$$\Psi^{RL}(D,\pi) = \left[\pi\frac{1}{\rho - \chi - \lambda_h\xi} + (1-\pi)\frac{1}{\rho - \chi - \lambda_l\xi}\right] \cdot D \ , \tag{24}$$

while J^{RL} is given by,

$$J^{RL}\left(s, D, \pi \mid \Psi^{RL}\right) = \begin{cases} \frac{\Psi^{RL}(D,\pi)}{D}\frac{(sD)^{1-\gamma}}{1-\gamma} - \frac{1}{\rho(1-\gamma)} & , \quad \text{if } \gamma \neq 1 \\ \frac{\Psi^{RL}(D,\pi)}{D}\ln(sD) + \phi^{RL}\left(\pi \mid \Psi^{RL}\right) & , \quad \text{if } \gamma = 1 \end{cases} \tag{25}$$

where

$$\phi^{RL}\left(\pi \mid \Psi^{RL}\right) = \phi^{AL}\left(\pi \mid \Psi^{AL}\right) \ . \tag{26}$$

It is notable that in the case $\gamma = 1$, the asset pricing rule is common across RE, AL, and RL (the P-D ratio is equal to $1/\rho$ for all). Moreover, only in the case $\gamma = 1$, it is $J^{RL}\left(s, D, \pi \mid \Psi^{RL}\right) = J^{AL}\left(s, D, \pi \mid \Psi^{AL}\right)$.

Equations (24) and (21) lead to an immediate qualitative comparison between these two setups. Corollary 2 makes this comparison.

Corollary 2 *For all $D > 0$ and $\pi \in (0,1)$,*

$$\frac{\Psi^{AL}(D,\pi)}{D} \gtreqless \frac{\Psi^{RL}(D,\pi)}{D} \Leftrightarrow \gamma \gtreqless 1 \ , \tag{27}$$

and

$$\frac{\partial \frac{\Psi^{AL}(D,\pi)}{D}}{\partial \pi}, \frac{\partial \frac{\Psi^{RL}(D,\pi)}{D}}{\partial \pi} \gtreqless 0 \Leftrightarrow \gamma \gtreqless 1 \ . \tag{28}$$

Proof Let $g(\lambda) \equiv 1/(\rho - \chi - \lambda\xi)$, and assume that the model's parameters are such that $g(\lambda_h), g(\lambda_l) > 0$. Notice that $\Psi^{AL}(D,\pi) = g(\pi\lambda_h + (1-\pi)\lambda_l)$, and $\Psi^{RL}(D,\pi) = \pi g(\lambda_h) + (1-\pi) g(\lambda_l)$. Since $\xi \gtreqless 0 \Leftrightarrow \gamma \gtreqless 1$, it follows that $g''(\lambda) \lesseqgtr 0 \Leftrightarrow \gamma \gtreqless 1$, for all $\lambda \in [\lambda_l, \lambda_h]$. Given these facts above, equation (27) is proved using the definition of concavity. Equation (28) is proved directly from (24) and (21), after also taking into account that $\xi \gtreqless 0 \Leftrightarrow \gamma \gtreqless 1$. □

Equation (28) implies that pessimism (increasing π) leads to lower demand for assets and lower equilibrium price, only if $\gamma < 1$, i.e., only if the coefficient of relative risk aversion does not depart too much from the preferences of a risk-neutral investor. As discussed in Section 2 above, setting a value $\gamma > 1$ plays a counterintuitive role for asset pricing even in the rational-expectations setting. In the rational-expectations setting we have noticed that, if $\gamma > 1$, then assuming higher risk for the model implies higher asset prices. In our learning setup, through equation (28), a value $\gamma > 1$ implies that pessimism (increasing π) leads to higher demand for assets, which is another counterintuitive result. As Bansal and Yaron (2004, p. 1487), and Barro (2009, p. 249) comment, CRRA preferences imply a tight link between the coefficient of relative risk aversion and the elasticity of intertemporal substitution. They suggest using Epstein-Zin (1989) (Duffie-Epstein (1992a,b)) recursive preferences to resolve such counterintuitive mechanics. We return to this issue in section 5 in the context of learning.

The implication of equation (27) is intuitive once we notice the convexity of the function $g(\lambda) = 1/(\rho - \chi - \lambda\xi)$ when $\gamma < 1$, that was introduced in the proof of Corollary 2. As we have emphasized above, at any instant $t \geq 0$, AL has beliefs, π, and does not anticipate further learning. For the empirically plausible case in which the P-D ratio falls with increasing pessimism ($\gamma < 1$), the P-D ratio under AL-pricing is higher than the P-D ratio under RL-pricing, for the same priors $\pi \in (0, 1)$. This happens because RL sees the potential of learning something more optimistic in the future about π, which dominates anticipated pessimistic information that is possible to arrive and increases the demand for the asset. This dominance of the more optimistic information out of the whole range of information that is anticipated to arrive is demonstrated by the convexity of function $g(\lambda)$ when $\gamma < 1$.

4.5 Comparing dynamic model simulations to data on P-D ratios and survey expectations

This section serves two purposes. First, it uses dynamic simulations of a calibrated version of our asset pricing model to provide a simple visual illustration of the interaction of disaster risk, subjective beliefs of investors and price-dividend ratios. This illustration helps improve our understanding of the analytical results discussed in the preceding sections. Secondly, this section examines similarities between actual data on price-dividend ratios and survey expectations after a stock-market crash and such model simulations. In particular, we investigate whether the dynamics of subjective beliefs regarding disaster probabilities can cause P-D ratio drops and persistence similar to the U.S. stock market experience in the last two decades. We also check whether these belief dynamics remain roughly within the range of belief variations apparent in surveys of the perceived threat of such a crash. Such a comparison may help to motivate a thorough empirical investigation of the role of subjective beliefs relative to fundamentals in future research.

The U.S. data on P-D-ratios from **Figure 1**, is plotted again in the two top panels of **Figure 2** (dashed lines). The dashed lines in the two bottom panels of **Figure 2** represent survey-based beliefs about the likelihood of an imminent stock-market crash in the United States, produced using a survey method described in Shiller et al. (1996). Shiller's Crash Confidence Index (CCI) refers to the percentage of the respondents who stated that the probability of a stock-market crash occurring within the following semester is less than 10%.[16] So, the higher the CCI, the higher the optimism (more accurately, the higher the fraction of non-pessimistic survey respondents). At the time when the two incidents of the massive drops in dividends occurred, the Crash Confidence Index (CCI) was at its lowest level. Most interestingly, after the November 2008 crash, the CCI continued dropping for almost a year. Stock-market prices broadly seemed to follow the change in beliefs and declined substantially.[17]

The spirit of our exercise is to initiate a model simulation in the second semester of the

[16]Data are taken from the website
`http://icf.som.yale.edu/stock-market-confidence-indices-united-states`

[17]An alternative approach to measuring confidence using the cross-section of quarterly real GDP forecasts from the survey of professional forecasters is presented in Bansal and Shaliastovich (2010). They provide evidence that confidence and returns are negatively correlated and develop a model with jump-like confidence shocks and recency-biased learning.

year 1989 setting a level of initial beliefs for agent RL, π^{RL}_{1989}, which is close to values from CCI data. Then we impose two unforeseen jumps, one in the summer of year 2000, and one in the fall of 2008. To be able to simulate the model, we also need to calibrate a number of other parameters. In doing so, we choose values close to those used by Barro (2006) for explaining the equity premium puzzle. Investors' preference parameters are set to $\gamma = 0.2$ and $\rho = 2.28\%$. The parameters of the diffusion process with drift are set to $\mu = 2.5\%$ and $\sigma = 2\%$.[18] Regarding the magnitude of the impact of a disaster on dividends, ζ, we use a generic distribution, in which $\zeta = 20\%$ with probability 1, if a disaster occurs.

The hazard rates determining disaster risk are set to $\lambda_h = 1/5$, and $\lambda_l = 1/40$. $\lambda_h = 1/5$ implies an upper bound for pessimism, namely that sudden drops in the dividend process of magnitude 20% arrive once every five years on average. The most optimistic view, determined by $\lambda_l = 1/40$, is that such jumps arrive once every forty years on average. These values for λ_h and λ_l are not far from hazard rates motivated by rare-disaster data presented in Barro (2006, 2009). Moreover, the choice of upper bound provides a natural link to the Crash Confidence Index. Since the hazard rate reflects the average rate of disasters per year, the value $\lambda_h = 1/5$ implies the perceived probability that a disaster may occur with probability 10% each semester. Since $\pi(t)$ is the probability placed on such an event, the CCI index can be proxied by the value $1 - \pi(t)$, which can be viewed as the percentage of respondents who think that there is a lower than 10% probability of disaster (i.e., $1 - \pi(t)$ is interpreted as percentage of respondents who think that the hazard rate is lower than $\lambda_h = 1/5$)).

The two bottom panels of **Figure 2** plot the dynamics of $1 - \pi(t)$ (solid lines), under

[18]Barro (2006) uses these values in order to match historical data on consumption growth and volatility. An alternative calibration that would match growth and volatility data of dividends during the examined period would be: $\mu = 9\%$, $\sigma = 10\%$ and, $\rho = 7.4\%$. It would imply the same dynamics, because the magnitude of the expression $\rho - \chi$ is the same.

rational (bottom-left panel) and adaptive learning (bottom-right panel) that follow from equation (14) relative to the CCI index (dashed lines). The prior belief regarding the average hazard rate is to 85 percent for the RL and AL investors, that is $\pi_{1989}^{RL} = \pi_{1989}^{AL} = 85\%$. Thus, $1 - \pi_{1989} = 15\%$, which could be compared with a CCI index value indicating a 15% share of optimistic respondents. Since the dynamics of beliefs, $\pi(t)$, are driven by the same equation, and the disaster data is the same (namely, two crashes in 2000 and 2008 respectively), the evolution of beliefs is the same for both types of investors. These beliefs exhibit variations in the same range as the CCI index in the last two decades, namely between 20 to 50 percent. There is a gradual increase in optimism prior to the crashes in 2000 and 2008. A crash causes a drop to pessimistic levels that persists and is followed by a slow improvement. There are some differences and some similarities with the movements of the CCI index. This index did not rise so much before 2000. However, there is a local minimum around 2000, which is then followed by a slow improvement to optimistic heights prior to the global financial crisis. Then it rapidly declines reaching a minimum around the Lehman collapse, followed by another improvement.

The resulting dynamics of the price-dividend ratio are shown in the top panels of **Figure 2**. Though they share the same beliefs, asset demand by rational and adaptive learners is different and therefore also the evolution of price dividend ratios. As apparent in the top-left panel, RL investors anticipate more optimistic perceptions in the absence of another crisis and value the risky asset more highly. The simulation under rational learning thereby exhibits a continuing increase in the price-dividend ratio throughout the 1990s that almost reaches the observed U.S. stock market peak prior to the crash in 2000. The simulated P-D ratio then slowly rises from this depressed level to a lower peak followed by the rash in 2008. The comparison with the actual U.S. P-D data serves to illustrate that variations in subjective beliefs may well be capable of causing such dramatic movements.

Under adaptive learning (top right-hand panel) the movements in the P-D ratio are much smaller. This observation is fully consistent with Corollary 2 in Section 4.4, given that the preference parameter γ is set to a value below unity. AL investors act as if their beliefs will remain unchanged in the future. On balance they value the risky asset less than the RL investors. Thus, the P-D ratio remains substantially smaller under adaptive learning and its dynamics less pronounced. However, this simulation is not meant to propose that the assumption of adaptive learning is necessarily inconsistent with observed behavior the P-D ratio in the U.S. stock market. It is possible to change the calibration so as to achieve more pronounced movements in the P-D ratio under adaptive learning. For example, a more optimistic prior would result in higher valuations of the risky asset from the AL investors' perspective. As shown in the dynamic simulation reported in the right-hand-side panels of **Figure 3** (solid lines), an initial prior of $\pi^{AL}_{1989} = 57\%$, is sufficient to generate more dramatic rises and falls in the P-D ratio over time. This prior implies a level of optimism, $1 - \pi^{AL}_{1989} = 43\%$, that is more than double the value used in **Figure 2** and above the CCI data of that period.

Clearly, these simulations indicate that subjective belief dynamics can play an important role in understanding P-D ratios after stock market crashes. A thorough empirical investigation should be the subject of a future study. Before closing, however, we want to address a possibly important concern with regard to the theoretical specification of preferences we have used. Barro (2009) and others have suggested that asset-pricing models with CRRA preferences have difficulty matching certain empirical regularities because they restrict the intertemporal elasticity of substitution to be equal to the inverse of the coefficient of relative risk aversion. Instead, Epstein-Zin (1989) preferences allow to differentiate between risk aversion and the elasticity of substitution. Thus, in the next section, we provide an extension to Epstein-Zin preferences. In this case, it is only possible to obtain analytical solutions for adaptive-learner pricing.

5. Epstein-Zin Preferences

In the following we use the continuous-time formulation and parameterization of recursive "Epstein-Zin" preferences, suggested by Duffie and Epstein (1992a,b) to specify the AL investor;s utility, namely,

$$J\left(s\left(t\right), D\left(t\right), \pi\left(t\right)\ \mid \tilde{\Psi}^{AL}\right) = \int_t^\infty f\left(c\left(\tau\right)\ ,\ J\left(s\left(\tau\right), D\left(\tau\right), \pi\left(\tau\right)\ \mid \tilde{\Psi}^{AL}\right)\right) d\tau\ , \tag{29}$$

with $f\left(c, J\right)$ being a normalized aggregator of continuation utility, J, and current consumption, c, with

$$f\left(c, J\right) \equiv \rho\left(1-\gamma\right) \cdot J \cdot \frac{\left\{\frac{c}{\left[\left(1-\gamma\right)J\right]^{\frac{1}{1-\gamma}}}\right\}^{1-\frac{1}{\eta}} - 1}{1-\frac{1}{\eta}}\ , \tag{30}$$

where $\eta > 0$ denotes AL's elasticity of intertemporal substitution, while $\gamma > 0$ is the coefficient of relative risk aversion. Moreover, $\tilde{\Psi}^{AL}\left(D, \pi\right)$ denotes the pricing rule under Epstein-Zin preferences. Using our HJB solution approach, we show in Appendix D that,

$$\begin{aligned} \frac{\tilde{\Psi}^{AL}\left(D, \pi\right)}{D} &= \frac{1}{\rho - \left(1-\frac{1}{\eta}\right)\left(\mu - \gamma\frac{\sigma^2}{2}\right) - \Lambda\left(\pi\right)\frac{1-\frac{1}{\eta}}{1-\gamma}\left\{E_\zeta\left(1-\zeta\right)^{1-\gamma} - 1\right\}} = \\ &= \frac{1}{\rho - \frac{1-\frac{1}{\eta}}{1-\gamma}\left[\chi + \Lambda\left(\pi\right)\xi\right]}\ , \end{aligned} \tag{31}$$

which gives some more degrees of freedom for calibration. The special case of $\gamma = 1/\eta$, corresponds to pricing implied by (21) above, since this is the case in which Epstein-Zin preferences collapse to standard time-separable preferences with constant relative-risk aversion (we demonstrate this equivalence in Appendix D as well).

This adaptive learning specification with recursive preferences may serve as guidance for future research. Unfortunately, the case of rational learning would then require numerical approximation.

6. Conclusion

Recent research on rational-expectations asset-pricing models focuses on proposing variability in disaster risk as an explanation for several asset pricing puzzles and, in particular, for excessively volatile price-dividend (P-D) ratios (see, for example, Gabaix (2008, 2011), Wachter (2011), and Gourio (2008) and Barro et al. (2010)). Another line of research focuses on subjective beliefs and learning by investors and questions the assumption of knowledge of objective frequency distributions of disasters (for example, Weitzman (2007)). We have developed an asset pricing model with time variable disaster risk and Bayesian learning by imperfectly informed investors. We have also shown that this model helps understand episodes in which P-D ratios drop both rapidly and massively, at times intimately connected with jumps in the dividend process (e.g., see Figure 1). Such observations have also motivated research on bounded rationality and investor sentiment (see, for example, Barberis et al. (1998)). Instead of following such a research approach, here, we have proposed a theory that does not require relaxing rationality. Our analysis has only assumed limited information, i.e. we have relaxed that investors know everything about the structure of disaster-risk variability and we have introduced rational Bayesian learning. In addition, we have defined and analyzed a particular deviation from fully rational behavior in the form of adaptive Bayesian learning.

A key reason motivating our limited-information approach has been the particular nature of rare disasters. Given the slow rate at which rare disasters arrive, it is rather difficult to argue that investors confidently reach rational expectations about the average frequency of arrival of disasters (hazard rate). We have demonstrated that, in our setting, indeed, Bayes' rule does not lead to learning with perfect confidence even after infinite time has passed.

In our model, rational investors may be perfectly aware of their ignorance, and fully forward-looking, anticipating new information to arrive and future learning to take place. We show that in such an environment, Bayes' rule implies that beliefs jump to pessimistic levels after a rare disaster occurs. These jumps towards pessimism create massive jumps in demands for assets, and therefore imply massive downward jumps in P-D ratios. When disasters take long to occur, optimism gradually takes over, and it can lead to high P-D ratios. These dynamics imitate behavior that is often attributed to investor psychology, such as sudden investment freezing due to fear after a sudden event with dramatic short-run consequences occurs, and slow restoration of confidence after a long period of no stock-market crashes. So, our findings suggest that, under the assumption of not knowing the stochastic structure of rare events with dramatic short-run consequences, emotion and logic may meet each other, in the sense that what is perceived as emotion can be fully rationalized. An evolutionary psychology perspective might suggest that our results formalize an argument that the instinct of fear is an endowment by nature that complements rationality.

Asset-pricing dynamics in our illustrative simulations are qualitatively similar between adaptive and rational learners. However, there are substantial quantitative differences. The study of such quantitative differences between adaptive and rational learning could be an interesting topic for future research in asset-pricing models. Other important extensions would concern belief-heterogeneity among investors and second-order learning about rare disasters. Our setup and analysis could also be generalized in order to include learning about the possibility of upward jumps. For example, the emergence of a new general-purpose technology, may motivate optimistic expectations for a “new economy” with sudden bursts of investor enthusiasm, triggered by rare upward jumps in the dividend process (sometimes triggered by the sudden massive entry of new firms).

7. Appendix A

Proof of Proposition 1

Since the arrival times of jumps are exponentially distributed, at any time $t \geq 0$,

$$\Pr\{\text{a jump occurs within the time interval } [t, t+\Delta t]\} = F_t(t+\Delta t) \ , \tag{32}$$

where F_t is the distribution function of future times of disaster conditional upon being at the time instant t. If the hazard rate of disasters was a fixed parameter $\bar{\lambda}$, then $F_t(t+\Delta t)$ would be equal to $1 - e^{-\bar{\lambda}\Delta t}$. Yet, under the assumption that the hazard rate is stochastic, with the hazard-rate distribution being governed by equation (3),

$$F_t(t+\Delta t) = \pi^* \left(1 - e^{-\lambda_h \Delta t}\right) + (1-\pi^*)\left(1 - e^{-\lambda_l \Delta t}\right) \ . \tag{33}$$

Before we take the limit $\Delta t \to 0$, we must accommodate the idea that, in continuous time, jumps can occur only once at every instant t. For this reason we modify the definition given by (32) as,

$$\Pr\{\text{a jump occurs only once within the time interval } [t, t+\Delta t]\} = F_t(t+\Delta t) - O(\Delta t) \ ,$$

in which $O(\Delta t)$ is the asymptotic order symbol (recall that any function $\Theta(\Delta t) \geq 0$ is $O(\Delta t)$ if $\lim_{\Delta t \to 0}[\Theta(\Delta t)/\Delta t] = 0$). So, from the distribution function F_t we can construct its corresponding density function f_t through simple differentiation, i.e.,

$$\Pr\{\text{a jump occurs only once within the time interval } [t, t+dt)\} = dF_t(t) = F_t'(t) \cdot dt \ , \tag{34}$$

with

$$F_t'(t) = \lim_{\Delta t \to 0} \frac{F_t(t+\Delta t) - F_t(t) - O(\Delta t)}{\Delta t} \ .$$

Notice that for computing $F_t'(t)$, $F_t(t) = 0$, and $\lim_{\Delta t \to 0}[O(\Delta t)/\Delta t] = 0$ by definition, and given equation (33) L'Hôpital's rule must be used. Doing so leads to $F_t'(t) = \lambda^*$, and together with equation (34) the probabilities associated with the occurrence vs. nonoccurrence of disasters appearing in equation (4) are proved. □

Proof of the asset pricing equation under rational expectations (Equation (8))

For the derivation of equation (8), we first show that J^{RE} is given by,

$$J^{RE}\left(s, D \mid \Psi^{RE}\right) = \begin{cases} \frac{\Psi^{RE}(D)}{D}\frac{(sD)^{1-\gamma}}{1-\gamma} - \frac{1}{\rho(1-\gamma)} & , \text{ if } \gamma \neq 1 \\ \frac{\Psi^{RE}(D)}{D}\ln(sD) + \phi^{RE}\left(\Psi^{RE}\right) & , \text{ if } \gamma = 1 \end{cases} \tag{35}$$

in which

$$\phi^{RE}\left(\Psi^{RE}\right) = \frac{1}{\rho}\frac{\Psi^{RE}(D)}{D}\left\{\left(\mu - \frac{\sigma^2}{2}\right) + \lambda^* E_\zeta\left[\ln(1-\zeta)\right]\right\} = \frac{\mu - \frac{\sigma^2}{2} + \lambda^* E_\zeta\left[\ln(1-\zeta)\right]}{\rho^2} . \tag{36}$$

The first-order conditions of (7) are,

$$c^{-\gamma} = \frac{1}{\Psi^{RE}(D)} \cdot J_s^{RE}\left(s, D \mid \Psi^{RE}\right) . \tag{37}$$

In order to solve the differential equation given by (7) subject to (37), we take a guess on the general functional form of $J^{RE}\left(s, D \mid \Psi^{RE}\right)$ with undetermined coefficients. First, we examine the case $\gamma \neq 1$, taking the guess,

$$J^{RE}\left(s, D \mid \Psi^{RE}\right) = a + b\frac{(sD)^{1-\gamma}}{1-\gamma} , \tag{38}$$

in which the undetermined coefficients may depend on Ψ^{RE}, and thus be functionals of the form $a\left(\Psi^{RE}\right)$ and $b\left(\Psi^{RE}\right)$. We drop the dependence of α and b on Ψ^{RE} for notational simplicity. Equation (38) implies,

$$J_s^{RE}\left(s, D \mid \Psi^{RE}\right) = bs^{-\gamma}D^{1-\gamma} , \tag{39}$$

$$J_D^{RE}\left(s, D \mid \Psi^{RE}\right) = bs^{1-\gamma}D^{-\gamma} ,$$

and,

$$J_{DD}^{RE}\left(s, D \mid \Psi^{RE}\right) = -\gamma b s^{1-\gamma} D^{-\gamma-1} \ . \tag{40}$$

Combining equation (37) with (39) gives,

$$c = \left[b \frac{D}{\Psi^{RE}\left(D\right)}\right]^{-\frac{1}{\gamma}} sD \ . \tag{41}$$

Since all agents are identical, in equilibrium there is no trade among individuals, and the representative agent's demand for assets is $s\left(t\right) = S\left(t\right) = S\left(0\right)$ for all $t \geq 0$. This means that $ds\left(t\right) = 0$ for all $t \geq 0$. So, the budget constraint, equation (6), implies that, for all $t \geq 0$, each household consumes her dividend, i.e.

$$c = sD \ . \tag{42}$$

From (42) and (41) we obtain,

$$b = \frac{\Psi^{RE}\left(D\right)}{D} \tag{43}$$

Plugging equations (38) through (43) into the HJB equation given by (7), we arrive at,

$$\rho \left[a + \frac{1}{\rho\left(1-\gamma\right)}\right] + \left[1 - \left(\rho - \chi - \lambda^{*}\xi\right) b\right] \frac{\left(sD\right)^{1-\gamma}}{1-\gamma} = 0 \ . \tag{44}$$

Setting both the constant of equation (44) and the factor of $\left(sD\right)^{1-\gamma}$ equal to zero, we obtain both the pricing function given by (8) and the functional form of $J^{RE}\left(s, D \mid \Psi^{RE}\right)$ given by the branch of (35) corresponding to the case of $\gamma \neq 1$.

For the case in which $\gamma = 1$, the guess for $J^{RE}\left(s, D \mid \Psi^{RE}\right)$ is,

$$J^{RE}\left(s, D \mid \Psi^{RE}\right) = a_1 + b_1 \ln\left(sD\right) \ ,$$

and the same procedure as above leads to the expression given by the branch of (35) corresponding to the case of $\gamma \neq 1$.

8. Appendix B – Dynamics of Beliefs about Jumps

Proof of Proposition 2 We derive the law of motion given by (13) in two parts. The first part deals with the belief dynamics when a jump occurs at a particular instant. The second part derives belief dynamics when a jump does not occur at a particular instant.

Part 1. Belief evolution when a jump occurs at a particular instant

For the evolution of these beliefs over time, let's fix a time interval $\Delta t > 0$ of arbitrary length (Δt does not need to be arbitrarily short), and let's consider that within Δt a jump event occurs. Since the investor knows that the distribution governing the jump is exponential (it is only that the hazard rate is unclear to him), for any perceived hazard rate $\tilde{\lambda}$, the probability that the jump event occurs at least once within the interval $[t, t + \Delta t)$, is determined by the distribution function of an exponential random variable,

$$\Pr\{\text{jumps occur within the time interval } [t, t + \Delta t)\} = 1 - e^{-\tilde{\lambda}\Delta t} \ . \tag{45}$$

Based on (45), and applying Bayes' rule to the case where a jump has been observed to have occurred within the time interval $[t, t + \Delta t)$, the posterior belief for the probability that $\tilde{\lambda} = \lambda_h$ (i.e., the update to the prior $\pi(t)$) is,

$$\pi(t + \Delta t) = \frac{\pi(t)\left(1 - e^{-\lambda_h \Delta t}\right)}{\pi(t)\left(1 - e^{-\lambda_h \Delta t}\right) + [1 - \pi(t)]\left(1 - e^{-\lambda_l \Delta t}\right)} \ . \tag{46}$$

The next step is to take the limit $\Delta t \downarrow 0$ on both sides of equation (46). Yet, for treating the LHS of (46) we need somewhat different notation. In order to distinguish between prior and posterior π at instant $t \geq 0$, let's keep the notation $\pi(t)$ for the prior and use $\pi(t) + d\pi(t)|_{jump}$ in order to capture the posterior in this particular case where the jump has occurred within the time interval $[t, t + \Delta t)$ when $\Delta t \downarrow 0$. So,

$$\pi(t) + d\pi(t)|_{jump} = \lim_{\Delta t \downarrow 0} \frac{\pi(t)\left(1 - e^{-\lambda_h \Delta t}\right)}{\pi(t)\left(1 - e^{-\lambda_h \Delta t}\right) + [1 - \pi(t)]\left(1 - e^{-\lambda_l \Delta t}\right)} \ ,$$

and after applying L'Hôpital's rule it is,

$$\pi(t) + d\pi(t)\,|_{jump} = \frac{\lambda_h \pi(t)}{\Lambda(\pi(t))} \; . \tag{47}$$

A key observation about (47) is that whenever a jump occurs in the data, then beliefs jump as well: beliefs discontinuously move from $\pi(t)$ to $\lambda_h \pi(t) / \Lambda(\pi(t))$.

Part 2. Belief evolution when a jump does not occur at a particular instant

Let's fix a time interval $\Delta t > 0$ of arbitrary length (again, Δt does not need to be arbitrarily short), and let's consider that within Δt no jump event occurs at all. For any perceived hazard rate $\tilde{\lambda}$ from the perspective of the investor, the probability that the jump event occurs at least once within the interval $[t, t + \Delta t)$, is determined by the distribution function of an exponential distribution, namely,

$$\Pr\{\text{No jumps occur within the time interval } [t, t + \Delta t)\} = e^{-\tilde{\lambda}\Delta t} \; . \tag{48}$$

In light of (48), after applying Bayes' rule to the case where no jump has occurred within the time interval $[t, t + \Delta t)$, the posterior belief for the probability that $\tilde{\lambda} = \lambda_h$ (i.e., the update to the prior $\pi(t)$) is,

$$\pi(t + \Delta t) = \frac{\pi(t)\, e^{-\lambda_h \Delta t}}{\pi(t)\, e^{-\lambda_h \Delta t} + [1 - \pi(t)]\, e^{-\lambda_l \Delta t}} \; . \tag{49}$$

A problem with (49) can be seen immediately: this posterior belief is well-defined when Δt is bounded away from 0, yet, for the case of our interest, $\Delta t \downarrow 0$, the posterior belief seems to remain constant. However, this is not correct. Even without occurrence of jumps beliefs still change over time.[19] The key to calculating this posterior belief in continuous time is to observe that, whenever no jumps occur, the posterior belief will follow a *deterministic trend*

[19] Observe also that the expression given on the RHS of (47) is a strictly increasing and strictly concave function of $\pi(t)$ with two fixed points in the unitary square: 0 and 1. So, if beliefs remain the same while jumps are not occurring, then successive jumps will make $\pi(t)$ to converge to 1 as $t \to \infty$.

for that instant. So, in the absence of random events at that instant, it is possible to take a conventional derivative. This means that we can subtract $\pi(t)$ from both sides of (49), divide both sides of (49) by Δt, and obtain,

$$\frac{\pi(t+\Delta t)-\pi(t)}{\Delta t}=\pi(t)\left[1-\pi(t)\right]\frac{\frac{e^{-(\lambda_h-\lambda_l)\Delta t}-1}{\pi(t)e^{-(\lambda_h-\lambda_l)\Delta t}+1-\pi(t)}}{\Delta t}\ ,$$

which implies (after applying L'Hôpital's rule),

$$\lim_{\Delta t\downarrow 0}\frac{\pi(t+\Delta t)-\pi(t)}{\Delta t}=-\left(\lambda_h-\lambda_l\right)\pi(t)\left[1-\pi(t)\right]\ ,$$

or,

$$d\pi(t)\,|_{no\ jump}=-\delta\pi(t)\left[1-\pi(t)\right]dt\ . \tag{50}$$

Combining equations (47) and (50) leads to (13). $\square$

Proof of Proposition 3 Equation (13) has two parts, a deterministic part and a stochastic part. The deterministic part is the first term of the RHS of (13), which is equal to $-\delta\pi(1-\pi)$, and it defines a deterministic first-order differential equation,

$$\dot{\pi}=-\delta\pi(1-\pi)\ , \tag{51}$$

which can be re-written as,

$$\dot{\pi}=\delta\pi^2-\delta\pi\ . \tag{52}$$

Equation (52) is a Bernoulli differential equation. So, we can use the Bernoulli transformation

$$z_\pi(t)\equiv\pi(t)^{-1}\quad\text{for all } t\geq 0\ . \tag{53}$$

From (53) it is,

$$\dot{z}_\pi(t)=-\pi(t)^{-2}\cdot\dot{\pi}(t)\qquad\text{for all } t\geq 0\ . \tag{54}$$

So, after multiplying both sides of equation (52) by $-\pi^{-2}$ and also substituting (53) and (54) it is,

$$\dot{z}_\pi = \delta z_\pi - \delta \ . \tag{55}$$

The solution to equation (55) is,

$$z_\pi (t) - 1 = e^{\delta t} \left[z_\pi (0) - 1 \right] \ ,$$

and substituting (53) gives

$$\pi (t) = \frac{1}{1 + e^{\delta t} \frac{1-\pi_0}{\pi_0}} \ , \quad \text{for all } t \geq 0 \ . \tag{56}$$

Equation (56) shows that, no matter how much time has passed without any jumps occurring, probability π always stays within the open interval $(0, 1)$.

The second part of equation (13) is stochastic and given by a jump process, such that the probability jumps from its original level π to the level given by $\lambda_h \cdot \pi / \Lambda (\pi)$. The statement

$$\pi \in (0, 1) \Rightarrow \frac{\lambda_h \pi}{\Lambda (\pi)} \in (0, 1) \tag{57}$$

holds, because $\lambda_h \cdot \pi / \Lambda (\pi) > 0$ for all $\pi \in (0, 1)$, and $\lambda_h \cdot \pi / \Lambda (\pi) < 1 \Leftrightarrow \pi < 1$, which is also true for all $\pi \in (0, 1)$. Combining (57) with (56) proves the part of the proposition which states that for all $\pi (0) = \pi_0 \in (0, 1)$, $\pi (t) \in (0, 1)$ for all $t \geq 0$.

Applying analytical techniques that pertain to Poisson differential equations on (13), we obtain,[20]

$$\frac{E (d\pi)}{dt} = -\delta \pi (1 - \pi) + \lambda^* \left[\frac{\lambda_h \pi}{\Lambda (\pi)} - \pi \right] \ . \tag{58}$$

Using the fact that $\Lambda (\pi^*) = \lambda^*$ implies $\pi^* = (\lambda^* - \lambda_l) / \delta$, after some algebra, equation (58) gives,

$$\dot{\pi}^e = \delta^2 \pi (1 - \pi) \frac{\pi^* - \pi}{\Lambda (\pi)} \ . \tag{59}$$

[20]See, for example, Merton (1971, pp. 395-401) and Kushner (1967, pp. 16-22).

For calculating the limit $\lim_{t\to\infty} E[\pi(t)]$, notice that, according to Dynkin's formula (we denote the Dynkin operator by $\mathcal{D}$),

$$E[\pi(t)] = \pi(0) + E\left[\int_0^t \mathcal{D}\pi(\tau)\, d\tau\right] \ , \tag{60}$$

in which,

$$\mathcal{D}\pi(\tau) = \delta^2 \pi (1-\pi) \frac{\pi^* - \pi}{\Lambda(\pi)} \ , \tag{61}$$

which is a formula based on equation (59). The expectations operator in the expression $E\left[\int_0^t \mathcal{D}\pi(\tau)\, d\tau\right]$ of the RHS of (60) transforms the dynamics of (59) into

$$\dot{\pi}^e = \delta^2 \pi^e (1-\pi^e) \frac{\pi^* - \pi^e}{\Lambda(\pi^e)} \ , \tag{62}$$

where $\pi^e \equiv E(\pi)$. To see that the expectations operator in the expression $E\left[\int_0^t \mathcal{D}\pi(\tau)\, d\tau\right]$ of the RHS of (60) leads to (62), fix any $t \geq 0$ and any $\Delta t > 0$, and consider equation (60) expressed as,

$$\pi^e(t+\Delta t) = E[\pi(t+\Delta t)] = \pi(t) + E\left[\int_t^{t+\Delta t} \mathcal{D}\pi(\tau)\, d\tau\right] \ . \tag{63}$$

If we set $\Delta t > 0$ arbitrarily small, then (63) and (61) give rise to an approximate recursion with respect to π^e, given by the difference equation,

$$\pi^e(t+\Delta t) = \pi^e(t) + \delta^2 \pi^e(t)[1-\pi^e(t)] \frac{\pi^* - \pi^e(t)}{\Lambda(\pi^e(t))} \cdot \Delta t \ , \tag{64}$$

for all discrete periods with interval length $[t\ ,\ t+\Delta t]$ and any $t \geq 0$. Equation (64) is a deterministic equation, since its initial conditions are non-stochastic ($\pi^e(0) = \pi(0) = \pi_0$). Equation (64) is a construction by approximation that leads to differential equation (62) by subtracting $\pi^e(t)$ from both sides of (64), dividing by Δt, and taking the limit $\Delta t \to 0$. Equation (62) implies dynamics given by,

$$\text{for all } \pi^e \in (0,1)\,, \quad \dot{\pi}^e \gtreqless 0 \Leftrightarrow \pi^e \lesseqgtr \pi^* \ . \tag{65}$$

With the help of a one-dimensional phase diagram it can be verified that $\lim_{t\to\infty} E\left[\pi\left(t\right)\right] = \lim_{t\to\infty} \pi^{e}\left(t\right) = \pi^{*}$, which proves equation (17) of the proposition.

In order to prove equation (18), after applying analytical techniques that pertain to Poisson differential equations on (13), we obtain,

$$\frac{E\left(d\pi^{2}\right)}{dt} = -2\delta\pi^{2}\left(1-\pi\right) + \lambda^{*}\left\{\left[\frac{\lambda_{h}\pi}{\Lambda\left(\pi\right)}\right]^{2} - \pi^{2}\right\} \ ,$$

which simplifies to,

$$\frac{E\left(d\pi^{2}\right)}{dt} = \delta\pi^{2}\left(1-\pi\right)\left[\frac{\lambda^{*}\left(\lambda_{h}+\lambda_{l}+\delta\pi\right)}{\Lambda\left(\pi\right)^{2}} - 2\right] \ . \tag{66}$$

Using the same argument as above, we can show that (66) yields

$$\frac{E\left(d\pi^{2}\right)}{dt} = \delta E\left(\pi^{2}\right)\left\{1-\left[E\left(\pi^{2}\right)\right]^{\frac{1}{2}}\right\}\left[\frac{\lambda^{*}\left\{\lambda_{h}+\lambda_{l}+\delta\left[E\left(\pi^{2}\right)\right]^{\frac{1}{2}}\right\}}{\Lambda\left(\left[E\left(\pi^{2}\right)\right]^{\frac{1}{2}}\right)^{2}} - 2\right] \ . \tag{67}$$

For notational simplicity, we can use the transformation $z \equiv E\left(\pi^{2}\right)$, which makes (67) be expressed as,

$$\dot{z} = \delta z\left(1-z^{\frac{1}{2}}\right)\left[\frac{\lambda^{*}\left(\lambda_{h}+\lambda_{l}+\delta z^{\frac{1}{2}}\right)}{\Lambda\left(z^{\frac{1}{2}}\right)^{2}} - 2\right] \ , \tag{68}$$

since $\pi \in (0,1)$. Because the term $\delta z\left(1-z^{\frac{1}{2}}\right)$ in (68) is always positive for all $\pi \in (0,1)$, we can focus on the sign of the expression in the bracket of the right-hand side of equation (68), which is determined by the sign of the expression

$$f\left(z^{\frac{1}{2}}\right) \equiv z - \frac{\delta\pi^{*}-3\lambda_{l}}{2\delta}z^{\frac{1}{2}} - \frac{\lambda^{*}\lambda_{h}-\lambda_{l}+\left(\lambda^{*}-1\right)\lambda_{l}}{2\delta^{2}} \ , \tag{69}$$

and

$$f\left(z^{\frac{1}{2}}\right) \lesseqqgtr 0 \Leftrightarrow \dot{z} \gtreqqless 0 \ . \tag{70}$$

There exist two real roots for the quadratic form given by (69), namely,

$$z_{1,2}^{\frac{1}{2}} = \frac{\lambda^{*}-4\lambda_{l} \pm \left(\lambda^{*}\right)^{\frac{1}{2}}\left(\lambda^{*}+8\lambda_{h}\right)^{\frac{1}{2}}}{4\delta} \ ,$$

and it is easy to verify that one root is negative, while the other is positive. Since $\pi \in (0,1)$ for all $t \geq 0$, we discard the negative root and we keep the positive root which is,

$$\hat{z}^{\frac{1}{2}} = \left[\widehat{E\left(\pi^2\right)}\right]^{\frac{1}{2}} = \frac{\lambda^* - 4\lambda_l + \left(\lambda^*\right)^{\frac{1}{2}}\left(\lambda^* + 8\lambda_h\right)^{\frac{1}{2}}}{4\delta} \ ,$$

implying that,

$$\hat{z} = \widehat{E\left(\pi^2\right)} = \left[\frac{\lambda^* - 4\lambda_l + \left(\lambda^*\right)^{\frac{1}{2}}\left(\lambda^* + 8\lambda_h\right)^{\frac{1}{2}}}{4\delta}\right]^2 \ . \tag{71}$$

Most importantly, $f\left(z^{\frac{1}{2}}\right) \lesseqqgtr 0 \Leftrightarrow z \lesseqqgtr \hat{z}$ for all $z \in (0,1)$, so through the aid of a one-dimensional phase diagram, the relationship given by (70) confirms that $\hat{z}$ is globally stable, for all $z \in (0,1)$. This means that as $t \to \infty$, $E\left(\pi^2\right) \to \widehat{E\left(\pi^2\right)}$. Using the fact that, asymptotically, $Var\left(\pi\right) = \widehat{E\left(\pi^2\right)} - \left(\pi^*\right)^2$, proves (18). Given that $\pi^* = \left(\lambda^* - \lambda_l\right)/\delta$, after some algebra, it can be shown that $\widehat{E\left(\pi^2\right)} - \left(\pi^*\right)^2 > 0 \Leftrightarrow \lambda_h > \lambda^*$. The right-hand side of this equivalence is a true statement, completing the proof of the proposition. □

9. Appendix C – Rational Learner's Solution

This appendix is devoted to proving equations (24) and (25). The first-order conditions of (20) are given by,

$$c^{-\gamma} = \frac{1}{\Psi^{RL}\left(D,\ \pi\right)} \cdot J_s^{RL}\left(s, D,\ \pi \middle|\ \Psi^{RL}\right) \ , \tag{72}$$

and the guess we take for the undetermined-coefficients functional form of J^{RL} in the case $\gamma \neq 1$ is

$$J^{RL}\left(s, D,\ \pi \middle|\ \Psi^{RL}\right) = \kappa + \left(a + b\pi\right)\frac{\left(sD\right)^{1-\gamma}}{1-\gamma} \tag{73}$$

in which the undetermined coefficients, κ, a, and b, may depend on Ψ^{RL}, but we do not denote this dependence for notational simplicity. Moreover, undetermined coefficients a and b are different from these defined in other appendices. Equation (73) implies

$$J_s^{RL}\left(s, D,\ \pi \middle|\ \Psi^{RL}\right) = \left(a + b\pi\right)s^{-\gamma}D^{1-\gamma} \ ,$$

and after this is combined with (72), the implied formula for consumption is,

$$c = \left[\frac{D}{\Psi^{RL}(D,\ \pi)} \cdot (a + b\pi) \right]^{-\gamma} sD \ . \tag{74}$$

So, provided that the guess given by (73) proves to be correct, the market-clearing condition $c = sD$ combined with (74) implies that the P-D ratio is,

$$\frac{\Psi^{RL}(D,\ \pi)}{D} = a + b\pi \ . \tag{75}$$

Focusing on the case $\gamma \neq 1$, and using the guess given by (73) in order to calculate J_s^{RL}, J_D^{RL}, J_{DD}^{RL}, and J_π^{RL}, substitution of the resulting functions into the HJB equation (20), together with the market-clearing condition $c = sD$, after some algebra, results in the following expression,

$$\nu_1 + (\nu_2 + \nu_3 \cdot \pi) \frac{(sD)^{1-\gamma}}{1-\gamma} = 0 \ , \tag{76}$$

in which,

$$\nu_1 = \rho \left[\kappa + \frac{1}{\rho(1-\gamma)} \right] \ ,$$

$$\nu_2 = (\rho - \chi - \lambda_l \xi)\, a - 1 \ ,$$

and

$$\nu_3 = b(\rho - \chi - \lambda_h \xi) + a\delta\xi \ .$$

Ideally, it should be possible to make equation (76) hold for any levels of the model's variables, π, s, and D. The functional form on the left-hand side of equation (76) reveals that the only way to have equation (76) hold for any arbitrary levels of π, s, and D is to set $\nu_1 = \nu_2 = \nu_3 = 0$. Indeed, there exist unique values for the undetermined coefficients κ, a, and b, that make conditions $\nu_1 = \nu_2 = \nu_3 = 0$ to hold. These values are, $\kappa = -1/[\rho(1-\gamma)]$, $a = 1/(\rho - \chi - \lambda_l\xi)$, and $b = -a\delta\xi/(\rho - \chi - \lambda_h\xi)$, and after these values are combined with (75), and (73), equation (24) and the part of equation (25) that refers to the case in which $\gamma \neq 1$ are both validated.

For the case $\gamma = 1$, we take a guess of the form

$$J^{RL}\left(s, D,\ \pi \middle|\ \Psi^{RL}\right) = a_1 \cdot \ln\left(sD\right) + b_1 \pi + \kappa_1\ ,$$

and following the same procedure as above we arrive at the part of equation (25) that refers to the case in which $\gamma = 1$.

10. Appendix D

To see that preferences given by equations (29) and (30) lead to standard time-separable preferences with constant relative-risk aversion when $\gamma = 1/\eta$, notice that (30) implies,

$$f\left(c, J\right)|_{\gamma=\frac{1}{\eta}} = \rho \frac{c^{1-\gamma}}{1-\gamma} - \rho J\ . \tag{77}$$

To simplify notation, let's denote J at time $t \geq 0$ by $J\left(t\right)$, and let's use $J'\left(t\right)$ in order to denote the total derivative of J with respect to time evaluated at time t. Using this notation, equation (29) implies that $J'\left(t\right) = -f\left(c\left(t\right), J\left(t\right)\right)$, so (77) gives,

$$-J'\left(t\right) = \rho \frac{c\left(t\right)^{1-\gamma}}{1-\gamma} - \rho J\left(t\right)\ .$$

Multiplying the above equation by $\left(1/\rho\right) e^{-\rho t}$, and integrating with respect to time, gives,

$$-\frac{1}{\rho}\int_0^\infty e^{-\rho t} J'\left(t\right) dt = \int_0^\infty e^{-\rho t} \frac{c\left(t\right)^{1-\gamma}}{1-\gamma} dt - \int_0^\infty e^{-\rho t} J\left(t\right) dt\ .$$

After applying integration by parts in order to calculate the integral on the left-hand side of the above equation, we obtain,

$$\frac{1}{\rho}\left[J\left(0\right) - \lim_{t\to\infty} e^{-\rho t} J\left(t\right)\right] = \int_0^\infty e^{-\rho t} \frac{c\left(t\right)^{1-\gamma}}{1-\gamma} dt\ . \tag{78}$$

With lifetime utility $J(t)$ being always bounded (i.e., $\lim_{t\to\infty} J(t) < \infty$, too), $\lim_{t\to\infty} e^{-\rho t} J(t) = 0$, so (78) implies

$$J(0) = \rho \int_0^\infty e^{-\rho t} \frac{c(t)^{1-\gamma}}{1-\gamma} dt \ , \tag{79}$$

which reveals that $\gamma = 1/\eta$ implies that continuation utility is time separable of the form analyzed throughout Sections 2-4 of the paper.

Proof of the asset pricing equation under adaptive expectations and Epstein-Zin preferences (Equation (31))

For the derivation of equation (31), notice that the form of the HJB equation with Epstein-Zin preferences is,

$$\begin{aligned} 0 = \max_{c\geq 0} \Bigg\{ & f\left(c, J\left(s, D, \pi \mid \tilde{\Psi}^{AL}\right)\right) + J_s\left(s, D, \pi \mid \tilde{\Psi}^{AL}\right) \cdot \left[\frac{1}{\tilde{\Psi}^{AL}(D,\pi)}(sD - c)\right] + \\ & + J_D\left(s, D, \pi \mid \tilde{\Psi}^{AL}\right) \cdot \mu D + J_{DD}\left(s, D, \pi \mid \tilde{\Psi}^{AL}\right) \frac{(sD)^2}{2} + \\ & + \Lambda(\pi) \left\{ E_\zeta \left[J\left(s, (1-\zeta) D, \pi \mid \tilde{\Psi}^{AL}\right)\right] - J\left(s, D, \pi \mid \tilde{\Psi}^{AL}\right)\right\} \Bigg\} \ . \end{aligned} \tag{80}$$

First-order conditions are,

$$f_c\left(c, J\left(s, D, \pi \mid \tilde{\Psi}^{AL}\right)\right) = \frac{J_s\left(s, D, \pi \mid \tilde{\Psi}^{AL}\right)}{\tilde{\Psi}^{AL}(D,\pi)} \ . \tag{81}$$

Our guess for the functional form of J is,

$$J\left(s, D, \pi \mid \tilde{\Psi}^{AL}\right) = b(\pi) \frac{(sD)^{1-\gamma}}{1-\gamma} \ , \tag{82}$$

and we denote $b(\pi)$ by b for notational simplicity. Since $c = SD$ in equilibrium, (81) implies,

$$\frac{D}{\tilde{\Psi}^{AL}(D,\pi)} = \rho b^{\frac{\frac{1}{\eta}-1}{1-\gamma}} \ . \tag{83}$$

Substituting (82) and its implied derivatives into (80) gives, after some algebra,

$$\rho b^{\frac{\frac{1}{\eta}-1}{1-\gamma}} = \rho - \left(1 - \frac{1}{\eta}\right)\left(\mu - \gamma\frac{\sigma^2}{2}\right) + \frac{1-\frac{1}{\eta}}{1-\gamma}\Lambda\left(\pi\right)\left\{E_{\zeta}\left(1-\zeta\right)^{1-\gamma} - 1\right\} \ . \tag{84}$$

Combining (83) with (84) leads to (31).

REFERENCES

Adam, K. and A. Marcet (2010): Booms and busts in asset prices, working paper.

Aitchison, J. and Brown, J. A. C. (1957): The Log-normal Distribution, Cambridge (UK), Cambridge University Press.

Bakshi, G. and G. Skoulakis (2010): Do subjective expectations explain asset pricing puzzles?, Journal of Financial Economics, 98, 462-477.

Bansal, R. and A. Yaron (2004): Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, Journal of Finance, 59, 1481-1509.

Bansal, R. and I. Shaliastovich (2010): Confidence Risk and Asset Prices, American Economic Review: Papers and Proceedings 100, 537-541.

Bansal, R. and I. Shaliastovich (2011): Learning and Asset-Price Jumps, Review of Financial Studies, 24, 8, 2738-2780.

Barberis, N., A. Shleifer and R. Vishny (1998): A model of investor sentiment, Journal of Financial Economics, 49, 307-343.

Barro, R. J. (2006): Rare disasters and asset markets in the twentieth century, Quarterly Journal of Economics, 121, 823-866.

Barro, R. J. (2009): Rare Disasters, Asset Prices, and Welfare Costs, American Economic Review, 99, 243-264.

Barro, R. J., E. Nakamura, J. Steinsson and J. F. Ursua, (2010): Crises and Recoveries in an Empirical Model of Consumption Disasters. NBER Working Paper No. w15920.

Barro, R. J., and J. F. Ursua (2008): Macroeconomic crises since 1870, NBER Working Paper No. w13940.

Barro, R. J., and J. F. Ursua (2009): Stock-Market Crashes and Depressions, NBER Working Paper No. w14760.

Beck, G. and V. Wieland, (2002): Learning and control in a changing economic environment, Journal of Economics Dynamics and Control , 26, 1359-1377.

Benhabib, J., and C. Dave (2011): Learning, large deviations and rare events, NBER Working Paper No. w16816.

Benzoni, L., P. Collin-Dufresne and R. S.Goldstein (2011): Explaining asset pricing puzzles associated with the 1987 market crash, forthcoming, Journal of Financial Economics, doi:10.1016/j.jfineco.2011.01.008

Cogley, T. and Sargent, T. (2008): Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making, International Economic Review, 49, 185-221.

Comon, E. (2001): Extreme Events and the Role of Learning in Financial Markets, chapter 1, unpublished dissertation, Harvard University.

Cox, J. C., J. E. Ingersoll and S. A. Ross (1985): An Intertemporal General Equilibrium Model of Asset Prices, Econometrica, 53, 363-384.

Cox., D. R., and H. D. Miller (1965): The Theory of Stochastic Processes, London: Methuen & Co. Ltd.

Dixit, A. K., and R. S. Pindyck (1994), Investment under Uncertainty, New Jersey: Princeton University Press.

Dreyfous, S. E. (1965): Dynamic Programming and the Calculus of Variations, *R-441-PR, A Report Prepared for United States Airforce Project Rand*, The Rand Corporation, August 1965.

Duffie, D. and L. Epstein (1992a): Stochastic Differential Utility, Econometrica, 60, 353-394.

Duffie, D. and L. Epstein (1992b): Asset Pricing with Stochastic Differential Utility, Review of Financial Studies, 5, 411-436.

Epstein, L., and S. Zin (1989): Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, Econometrica, 57, 937-969.

Gabaix, X. (2008): Variable Rare Disasters: A Tractable Theory of Ten Puzzles in Macro-Finance. American Economic Review Papers and Proceedings 98, 64-67.

Gabaix, X. (2010): Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. Mimeo, NYU.

Gourio, F. (2008b): Disasters and recoveries. American Economic Review Papers and Proceedings 98, 68–73.

Gourio, F. (2008a): Time-series predictability in the disaster model, Finance Research Letters, 5, 191-203.

Guidolin, M. and Timmermann, A. (2007): Properties of equilibrium asset prices under alternative learning schemes, Journal of Economic Dynamics and Control, 31, 161-217.

Karlin, S. and Taylor, H.M. (1981): A Second Course in Stochastic Processes, New York: Academic Press.

Keller, G. and Rady, S. (2010): Strategic Experimentation with Poisson Bandits, Theoretical Economics, 5, 275-311.

Koulovatianos, C., Mirman, L. J. and Santugini, M. (2009): Optimal growth and uncertainty: Learning, Journal of Economic Theory, 144, 280-295.

Koulovatianos, C., and V. Wieland (2011): Investment Decisions under Adaptive Bayesian Learning about Diffusions and Jumps, Mimeo, University of Nottingham.

Kreps, D., (1998) "Anticipated Utility and Dynamic Choice," in D.P. Jacobs, E. Kalai, and M. Kamien, eds., Frontiers of Research in Economic Theory, Cambridge University Press, 242–74.

Kushner, H. J. (1967): Stochastic Stability and Control, New York: Academic Press.

LeRoy, S. F. (2008): Excess volatility tests. The New Palgrave Dictionary of Economics. Second Edition. Eds. Steven N. Durlauf and Lawrence E. Blume. Palgrave Macmillan.

Liptser, R. S. and A. N. Shiryayev (1977): Statistics of Random Processes I. New York: Springer-Verlag.

Liptser, R. S. and A. N. Shiryayev (2001): Statistics of Random Processes: I, II. Springer-Verlag, NewYork, NY.

Liu, J., J. Pan and T. Wang (2005): An Equilibrium Model of Rare-Event Premia and Its Implication for Option Smirks, The Review of Financial Studies, 18, 131-164.

Longstaff, F. A. and M. Piazzesi (2004) : Corporate earnings and the equity premium, Journal of Financial Economics, 74, 401-421.

Lucas, R. E. Jr. (1978): Asset prices in an exchange economy, Econometrica, 46, 1429-1445.

Merton, R.C. (1971): Optimum Consumption and Portfolio Rules in a Continuous-Time Model, Journal of Economic Theory, 3, 373-413.

Mirman, L. J., L. Samuelson and A. Urbano (1993): Monopoly experimentation, International Economic Review, 34, 549-564.

Muth, J. F. (1961): Rational expectations and the theory of price movements, Econometrica, 29, 315-335.

Papoulis, A. (1991), Probability, Random Variables, and Stochastic Processes, 3rd Edition, New York: McGaw-Hill, Inc.

Pastor, L. and P. Veronesi (2009a) : Learning in financial markets, Annual Review of Financial Economics, 1, 361-381.

Pastor, L. and P. Veronesi (2009b) : Technological revolutions and stock prices, American Economic Reivew, 99, 1451-1483.

Rietz, T. A. (1988): The Equity risk premium: A solution. Journal of Monetary Economics, 22, 117-131.

Ross, S. M. (2003): Introduction to Probability Models, 8^{th} Edition, New York: Academic Press.

Shiller, R. J., F. Kon-Ya and Y. Tsutsui (1996): Why Did the Nikkei Crash? Expanding the Scope of Expectations Data Collection. Review of Economics and Statistics, 78, 156-164.

Stirzaker, D. (2005): Stochastic Processes and Models, Oxford University Press.

Wachter, J. (2011): Can time-varying risk of rare disasters explain aggregate stock market volatility? Mimeo, Wharton School, University of Pennsylvania.

Weitzman, M. L. (2007): Subjective expectations and asset-return puzzles. American Economic Review, 97, 1102-1130.

Wieland, V. (2000a): Learning by doing and the value of optimal experimentation, Journal of Economics Dynamics and Control , 24 (4), 501-534.

Wieland, V. (2000b): Monetary policy, parameter uncertainty and optimal learning, Journal of Monetary Economics, 46, 199-228.

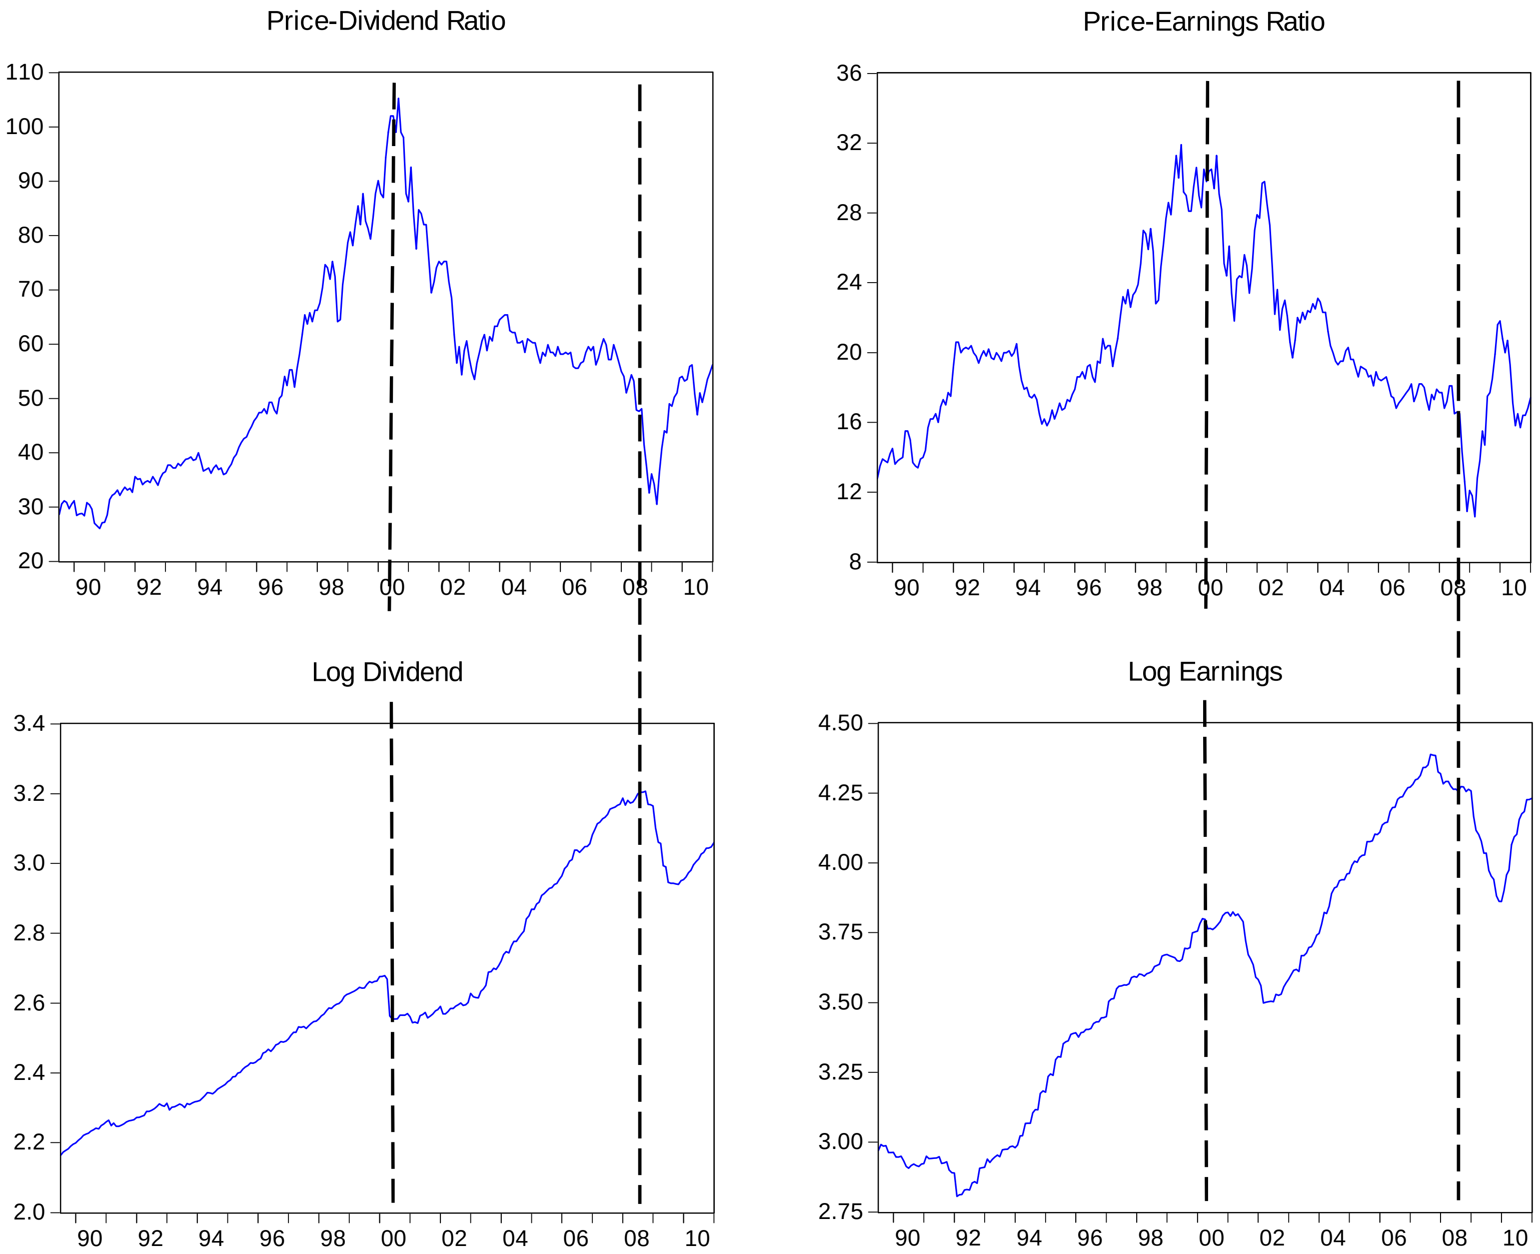

Figure 1 Monthly US stock-market data 1989-2010. Source: Datastream (TOTMKUS).

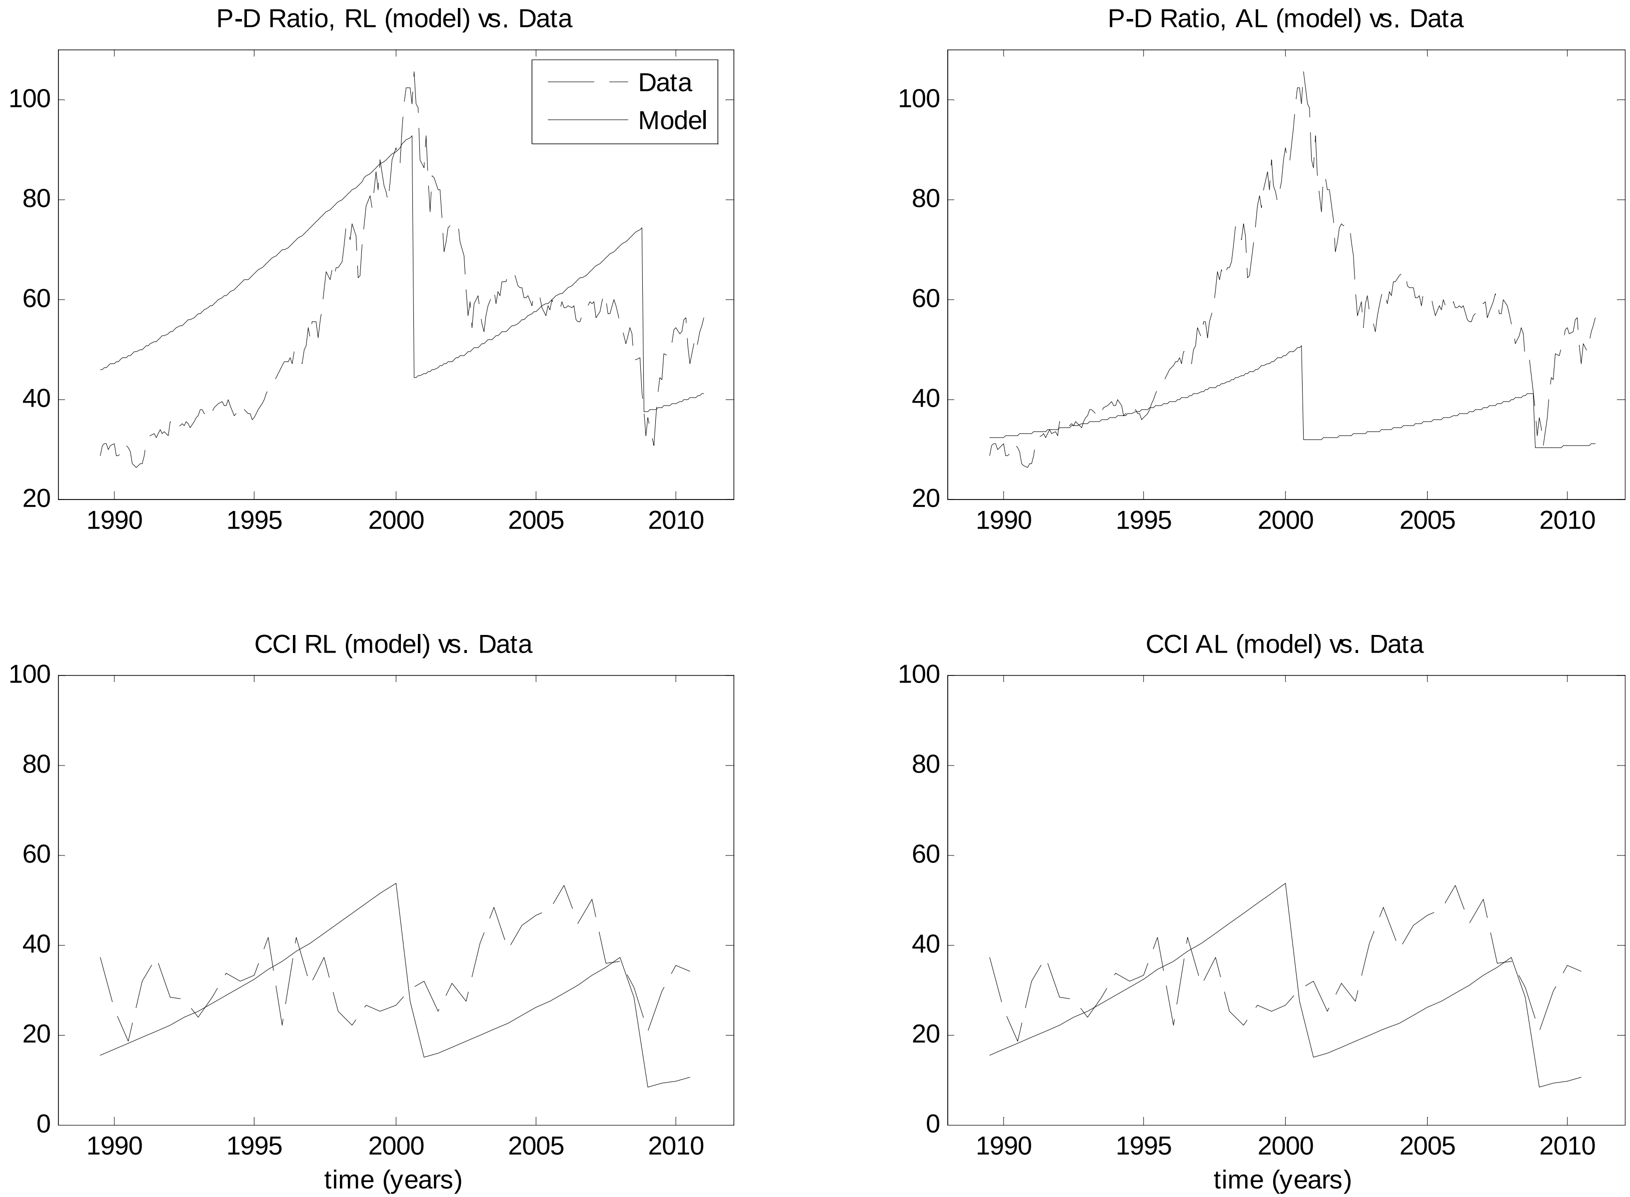

P-D Ratio, RL (model) vs. Data
Data
Model
P-D Ratio, AL (model) vs. Data
CCI RL (model) vs. Data
CCI AL (model) vs. Data
time (years)
time (years)

Figure 2

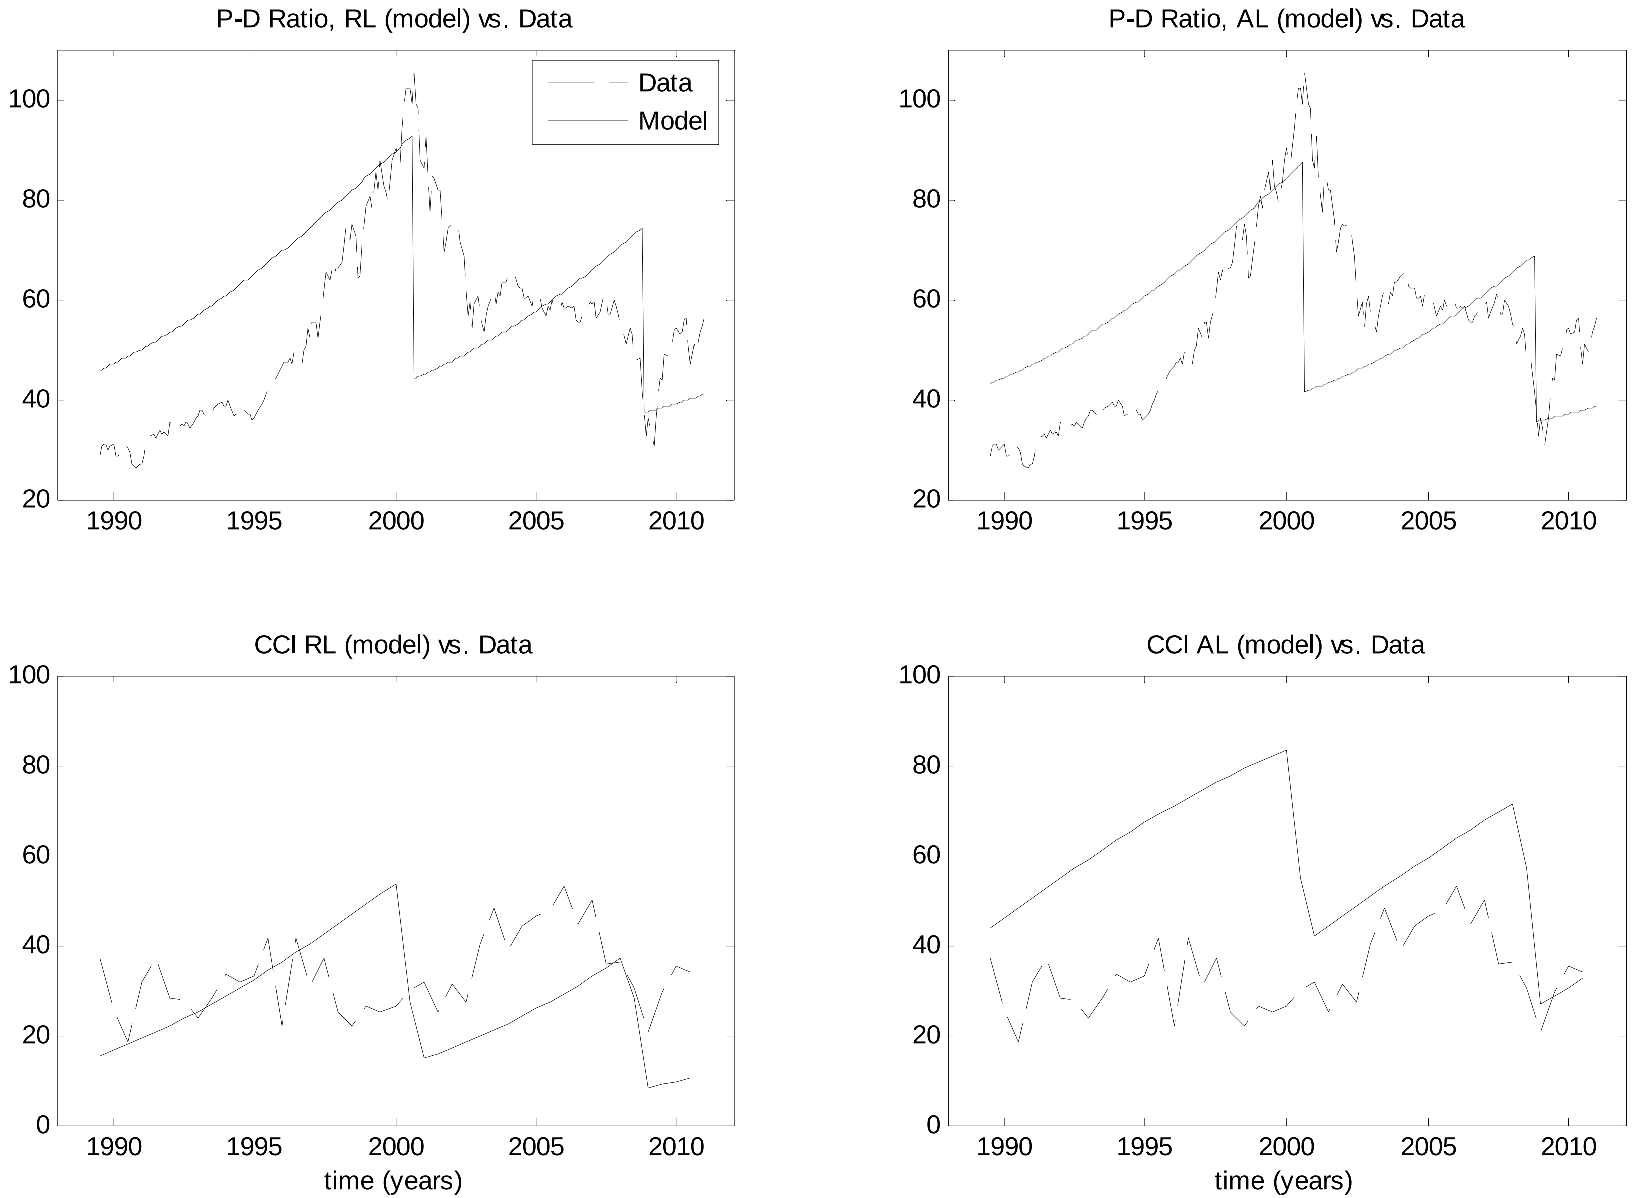
P-D Ratio, RL (model) vs. Data
Data
Model
100
80
60
40
20
1990
1995
2000
2005
2010
P-D Ratio, AL (model) vs. Data
100
80
60
40
20
1990
1995
2000
2005
2010
CCI RL (model) vs. Data
100
80
60
40
20
0
1990
1995
2000
2005
2010
time (years)
CCI AL (model) vs. Data
100
80
60
40
20
0
1990
1995
2000
2005
2010
time (years)

Figure 3