

Confidence Risk and Asset Prices

Ravi Bansal
Ivan Shaliastovich *

Current Version: November 2010

Comments Welcome

*Bansal (email: ravi.bansal@duke.edu) is affiliated with the Fuqua School of Business, Duke University, and NBER, and Shaliastovich (email: ishal@wharton.upenn.edu) is at The Wharton School, University of Pennsylvania. We would like to thank Alexander David, Cade Massey, Adriano Rampini, Robert Stambaugh, Pietro Veronesi, Amir Yaron and seminar participants at Duke, MIT (Sloan), Stockholm School of Economics, University of Amsterdam, University of Maryland, 2009 WFA meeting, and 2010 AEA meeting.

Confidence Risk and Asset Prices

Abstract

In the data, asset prices exhibit large negative moves at frequencies of about 18 months. These large moves are puzzling as they do not coincide, nor are they followed by any significant moves in the real side of the economy. On the other hand, we find that measures of investor's uncertainty about their estimate of future growth have significant information about large moves in returns. We set-up a recursive-utility based model in which investors learn about the latent expected growth using the cross-section of signals. To model the learning of the agents about the unobserved expected growth, we specify a belief-updating model (Kalman Filter is a special case) which incorporates the recency bias of investors in forecast formation. The uncertainty (confidence measure) about investor's growth expectations, as in the data, is time-varying and subject to large moves. In the confidence risks model, recency bias in conjunction with confidence risk fluctuations lead to large moves in asset prices. In calibrations we show that the model can account for the large return move evidence in the data, distribution of asset prices, predictability of excess returns and other key asset market facts.

1 Introduction

We highlight two puzzling data features which help isolate the economic risks that drive asset prices. First, in the data there is a strong evidence for large and predominantly negative moves in returns, which occur on average once every 18 months.¹ Consistent with these large negative moves, the distribution of asset prices in the data is non-Gaussian and exhibits negative skewness and heavy tails. Second, averaging across all the large asset-price moves, there is no persuasive evidence in the data for large contemporaneous or subsequent moves in the real side of the economy (consumption). That is, the once-in-18-month jumps in asset prices on average do not coincide, and nor are they followed by any significant large moves in the real side of the economy². The apparent lack of close connection between the asset prices and the real side of the economy indicates that the reasons for large asset-price moves can not be hardwired or built-in as large moves in the real side of the economy. This evidence begs the question, what risks can explain these large moves in asset prices? What is the compensation, in equilibrium, for these risks? Can we account for these large asset-price moves alongside the equity premium and the risk-free rate puzzles?

In this paper we present an equilibrium model that provides insights regarding these questions. Our model set-up follows a standard long-run risks specification of Bansal and Yaron (2004), which features Gaussian dynamics for consumption growth with time-varying expected growth and volatility; there are no large moves or jumps in the underlying consumption and dividend dynamics. A key feature of our model is that unlike the standard long-run risks specification, the expected growth is not directly observable, and investors learn about it using past history of the data and recent information in the cross-section of signals about expected growth. The cross-sectional standard deviation (the quality of the signals), referred to as the confidence measure, determines the uncertainty that investors face about future growth. In our model, fluctuations in confidence risk in conjunction with behaviorally motivated recency-biased expectations are responsible for triggering asset market jumps. In essence, our paper highlights non-Bayesian mechanisms for large moves in the financial markets.

To model the learning of agents about the unobserved expected growth, we specify a belief-updating model which incorporates the recency bias in forecast formation by investors — the optimal Kalman Filter is a special case of this belief-updating model. The recency bias specification is motivated by the evidence in Hogarth and Einhorn (1992), De Bondt and Thaler (1990), Kahneman and Tversky (1973) which suggests

¹Recent work featuring jump risks in asset prices includes Pan (2002), Andersen, Benzoni, and Lund (2002), Eraker, Johannes, and Polson (2003), Broadie, Chernov, and Johannes (2007).

² Both the frequency (18 months) and the magnitude of the jumps that we focus on are quite distinct from once in 600 months disaster states discussed in Gabaix (2007), Barro (2006) and Rietz (1988).

that the agents have a tendency to overweigh recent information and under-weight less salient data such as long-term averages. Such overweighing patterns are precluded in the standard optimal learning settings as agents optimally put less weight on recent information when its quality is low. Our specification of recency bias specification captures the intuition that as the information quality worsens, agents do not decrease the weight attached to the new signals, as in the Bayesian approach, and therefore overweigh the impact of recent information on their forecasts. Therefore, in equilibrium, asset prices decline more sharply (i.e. overreact) when confidence risk rises relative to the case of Bayesian learning. Our calibration and specification of the recency-biased expectation formation is consistent with the empirical findings in De Bondt and Thaler (1990) in that investors overweigh recent information more in states of high uncertainty.

In our model, investors have preference for the timing of resolution of uncertainty and demand compensation for short-run, long-run, consumption volatility and confidence risks. Our confidence risks specification includes large positive jumps. This is motivated by the empirical data and the theoretical model of Veldkamp (2006). We show that the confidence risks shocks receive quantitatively large risk compensation when investors' expectations are recency-biased. Hence, our approach to learning and confidence jump-risk channel can explain key features of returns data. Large positive shocks in the confidence measure endogenously translate into large negative jumps in returns. This can account for negatively skewed and heavy-tailed distribution of returns, even though consumption growth is Gaussian. Further, in the model both the expected excess returns and the price-dividend ratios can be predicted by the confidence measure, which can account for the predictability of excess returns by the price-dividend ratio in the data.

Our learning model features non-Bayesian, recency-biased expectation formation for predicting future consumption growth. A different approach to learning is presented in Hansen and Sargent (2006), who specify model-selection rules which capture investors' concerns about robustness and potential model misspecification in their expectation formations. A more commonly used approach is Bayesian Kalman-Filter learning which is featured in David (1997), Veronesi (1999), Veronesi (2000), Ai (2007), Bansal and Shaliastovich (2009)³. Our behaviorally-based approach for expectation formation, we show, is quantitatively important to match key asset market data, particularly regarding asset-price jumps.

The main target in this paper is to quantitatively explain the key non-Gaussian features of asset prices and at the same time account for the key dimensions of consumption and a confidence measure in the data. To give quantitative content to the model, we directly measure confidence using survey data. More specifically, we rely

³Rational learning is also featured in Detemple (1986), Gennotte (1986), Brennan (1998), Brennan and Xia (2001), David and Veronesi (2008), Croce, Lettau, and Ludvigson (2006).

on the cross-section of forecasts from the Survey of Professional Forecasters to construct the confidence measure as a cross-sectional variance of the average forecast in the data. These calculations are consistent with the theoretical setup of the model. We show that the average (consensus) forecast based on real GDP and industrial production data is persistent and very informative about future economy, even controlling for the history of the data. Next, we find that confidence measure in the data is quite small and on average equals about 1/15 of the volatility of the underlying macroeconomic series. Further, confidence measure exhibits substantial variation in time with occasional large positive spikes (high uncertainty). We find that the half-life of confidence shocks is about 6 months, so that the confidence shocks in the data are fairly short-lived. Our confidence risk measure highly correlates with the investor sentiment measure of ?. In general, our model provides a well-articulated equilibrium framework to analyze and interpret behaviorally-motivated investor sentiment.

We show that confidence measure contains important information about financial markets data. We provide strong statistical support for the large moves in the confidence measure in the data. In particular, large moves in the confidence measure are significantly related to contemporaneous large moves in returns, while there is no persuasive link in the data between large moves in real consumption and large moves in returns at the considered frequencies. Indeed, we show that the correlations of large move indicators in returns and in consumption are close to zero, both contemporaneously and in the future. On the other hand, the correlations of large move indicators in returns and confidence measure are about 35%. Our confidence measure also has significant information about asset valuations, even controlling for the aggregate volatility in the economy. We also consider the predictability of implied variance of returns by the confidence and macroeconomic volatility measures. Using the conditional quantile regressions, we find that large upward moves in the variance of returns are related to the fluctuations in confidence measure, rather than to the conditional volatility of the fundamentals in the economy. This indicates that large moves in asset prices are closely tied to the confidence risk that we feature in our model.

We calibrate the model to evaluate its quantitative implications for the equity markets. The calibration of consumption dynamics is standard and is designed to match the key features of the historical data. Parameters of confidence dynamics are calibrated to match unconditional moments of the series in the data, as well as its conditional distribution. In simulations we verify that the model can match well the key features of the confidence measure in the data, and in particular, along non-Gaussian dimensions. Based on the calibration of the model, we show that the model with fluctuating confidence risks can explain the negatively skewed and heavy-tailed distribution of returns, even though the consumption dynamics is Gaussian. In the model, as in the data, large moves in the confidence measure lead to large moves in asset prices and returns, though they are not mirrored in consumption growth rate.

For this reason, the model is also able to capture near zero correlation between the large moves in the real side of the economy and in returns, and significant correlation between large moves in the confidence measure and in returns, as found in the data.

Our model of confidence risks, which features recency bias and time-varying confidence of investors, provides a new channel for the variation in equity risk premium. In principle, in our model the variations in expected excess returns are driven by fluctuations in consumption volatility risks and confidence risks, however, quantitatively, almost all of the variation is due to fluctuations in confidence risks. Exploiting the fluctuations in confidence risks, we show that the model is capable of capturing short and long horizon predictability of future excess returns by price-dividend ratios. At the same time, the model is consistent with the general lack of predictability of future consumption growth rate by price-dividend ratios.

With learning, the full model generates an unconditional equity premium of about 5.3%. The key risk channels in the model come from confidence risks and long-run risks, which contribute about 1.7% each to the total premium. The compensations for short-run consumption shocks and volatility shocks are 1.2% and 0.8%, respectively. Using the empirical confidence measure and aggregate volatility in the data, we also construct the estimates of the conditional equity premium and show its decomposition into the sources of risk in the economy. The magnitudes of the risk compensation in the data are similar to those in the model. As the confidence risks are mostly jump risks, the model implies that the compensation for the jump risks in market returns is about 1/3 of the overall equity premium, which highlights the importance of the confidence risk channel for the asset markets.

The rest of the paper is organized as follows. In the next section we set up the model and discuss preferences of the representative agent and dynamics of the economy given the information set of investors. In Section 3 we solve for the asset prices and discount factor in the economy. Section 4 contains the empirical description of the confidence measure in the data, while the calibrations and asset-pricing implications of the economy are discussed in Section 5. The conclusion follows.

2 Model Setup

2.1 Preferences

We consider a discrete-time real endowment economy. The investors preferences over the uncertain consumption stream C_t can be described by the Kreps and Porteus (1978) recursive utility function of Epstein and Zin (1989):

$$U_t = \left\{ (1 - \delta)C_t^{\frac{1-\gamma}{\theta}} + \delta(E_t[U_{t+1}^{1-\gamma}])^{1/\theta} \right\}^{\frac{\theta}{1-\gamma}}, \quad (2.1)$$

where γ is a measure of a local risk aversion of the agent, ψ is the intertemporal elasticity of substitution and $\delta \in (0, 1)$ is the subjective discount factor. The conditional expectation is taken with respect to the date- t information set of the agent. For notational simplicity, we define

$$\theta = \frac{1 - \gamma}{1 - \frac{1}{\psi}}. \quad (2.2)$$

Note that when $\theta = 1$, that is, $\gamma = 1/\psi$, the above recursive preferences collapse to the standard case of expected utility. As is pointed out in Epstein and Zin (1989), in this case the agent is indifferent to the timing of the resolution of uncertainty of the consumption path. When the risk-aversion exceeds the reciprocal of IES, the agent prefers early resolution of uncertainty of consumption path, otherwise the agent has preference for late resolution of uncertainty. In the long-run risk model agents prefer early resolution of uncertainty.

As shown in Epstein and Zin (1989), the logarithm of the Intertemporal Marginal Rate of Substitution (IMRS) for these preferences is given by

$$m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1)r_{c,t+1}, \quad (2.3)$$

where $\Delta c_{t+1} = \log(C_{t+1}/C_t)$ is the log growth rate of aggregate consumption and $r_{c,t+1}$ is the log of the return (i.e., continuous return) on the asset which delivers aggregate consumption as its dividends each time period. This return is not observable in the data. It is different from the observed return on the market portfolio as the levels of market dividends and consumption are not equal: aggregate consumption is much larger than aggregate dividends. Therefore, we assume exogenous process for consumption growth and use a standard asset-pricing restriction,

$$E_t[\exp(m_{t+1} + r_{t+1})] = 1, \quad (2.4)$$

which holds for any continuous return $r_{t+1} = \log(R_{t+1})$ including the one on the wealth portfolio, to solve for the unobserved wealth-to-consumption ratio in the model. This enables us to express the discount factor in terms of the underlying state variables and shocks in the economy. We can then use the solution to the discount factor and the Euler equation (2.4) to price any asset in the economy.

We describe the dynamics of the real economy and the information set of the agent in the next section.

2.2 Consumption Process

The log consumption growth Δc_{t+1} process incorporates a time-varying mean x_t and stochastic volatility σ_t^2 :

$$\Delta c_{t+1} = \mu + x_t + \sigma_t \eta_{t+1}, \quad (2.5)$$

$$x_{t+1} = \rho x_t + \varphi_e \sigma_t \epsilon_{t+1}, \quad (2.6)$$

$$\sigma_{t+1}^2 = \sigma^2 + \nu_c (\sigma_t^2 - \sigma^2) + \varphi_w \sigma_t w_{c,t+1}. \quad (2.7)$$

where η_t , ϵ_t and $w_{c,t}$ are independent standard Normal shocks which capture short-run, long-run and volatility risks in consumption, respectively. Parameters ρ and ν_c determine the persistence of the conditional mean and variance of the consumption growth rate, while φ_e and φ_w govern their scale. Notably, short-run, long-run and consumption volatility shocks are Gaussian – there are no large moves (jumps) hard-wired into the underlying consumption process. The empirical motivation for the time-variation in the conditional moments of the consumption process comes from the long-run risks literature, see e.g. Bansal and Yaron (2004), Hansen, Heaton, and Li (2008) and Bansal, Kiku, and Yaron (2007b).

The agent knows the structure and parameters of the model and observes consumption volatility σ_t^2 , however, the true expected growth factor x_t is not directly observable and has to be inferred from the data. Investors form an estimate of the current state based on the past history of the data and recent information from the signals. Investors' estimate of the unobserved expected growth is subject to a learning error and is necessarily imprecise, and the amount of imprecision reflects the confidence of the investors about future economic growth. The time-variation in the quality of information about the future economy gives rise to fluctuating confidence and confidence risks, which we show have important asset-pricing implications in the economy.

2.3 Confidence About Growth

We consider a specification where in addition to observing current consumption, agents also receive n signals about the expected growth of the economy $x_{i,t}$, for $i = 1, 2, \dots, n$. These signals together with consumption data provide all the information about the expected growth state.

We assume that each signal deviates from the true state x_t by a random noise $\xi_{i,t}$,

$$x_{i,t} = x_t + \xi_{i,t}, \quad (2.8)$$

where the errors $\xi_{i,t}$ are randomly drawn from Normal distribution and are uncorrelated with fundamental shocks in the economy.

The date- t imprecision in signal i is captured by $V_{i,t}$:

$$\xi_{i,t} \sim N(0, V_{i,t}). \quad (2.9)$$

In general, the imprecision in the signal can be different across signals and can vary across time, hence subscripts i and t . However, we further assume that all signals i are ex-ante identical and come from the same distribution at each date t . Then, the precision of each signal is the same, and we denote $V_{0,t} \equiv V_{i,t}$ for all i .

As all the signals come from the same distribution and are ex-ante equally informative, the investor assigns same weight to each of them. That is, in the end the average signal is a sufficient statistic for the cross-section of all the individual ones. Define the average signal \bar{x}_t , which corresponds to the sample average of the individual signals. Then, using (2.8),

$$\bar{x}_t \equiv \frac{1}{n} \sum x_{i,t} = x_t + \xi_t, \quad (2.10)$$

where the uncertainty in average signal V_t and the average signal error are given by

$$V_t = \frac{1}{n} V_{0,t}, \quad \xi_t = \frac{1}{n} \sum \xi_{i,t}, \quad (2.11)$$

so that

$$\xi_t \sim N(0, V_t). \quad (2.12)$$

The uncertainty V_t determines the confidence of investors about their estimate of expected growth and is referred to it as confidence measure. In the model, confidence measure is assumed to be observable to investors. It can be estimated in the data

from the cross-section of the individual signals $x_{i,t}$. Indeed, the signal equation (2.8) implies that

$$E \left(\frac{1}{n-1} \sum_{i=1}^n (x_{i,t} - \bar{x}_t)^2 \right) = E \left(\frac{1}{n-1} \sum_{i=1}^n (\xi_{i,t} - \xi_t)^2 \right) = V_{0,t}, \quad (2.13)$$

so that the cross-sectional variance of the signals adjusted by the number of signals n can provide an estimate of the confidence measure V_t in the data.

The confidence measure in the model captures the uncertainty that the agents have about their estimate of future growth. The variation in the confidence measure across time reflects the fluctuations in the quality of information in the economy, so that at times when information is poor, signals are less precise and the uncertainty is high (V_t increases). The time-variation in the confidence measure and ensuing confidence risks, which become quantitatively important due to recency bias in investors' forecast formation, are the novel contribution of the model. Standard learning models, see for example David (1997), Veronesi (2000) and Brennan and Xia (2001), feature constant imprecision in observed signals, while Hansen and Sargent (2006) specify alternative learning rules robust to model misspecification. We discuss the specification of the confidence measure dynamics in the next section.

2.4 Confidence Dynamics

As discussed in the previous section, the confidence measure V_t reflects the uncertainty of investors about future growth. Our objective is to specify a model where the confidence measure is time-varying, and is subject to occasional large increases. Our specification for V_t is motivated by theoretical literature on this issue and the empirical work. In terms of theoretical work, Veldkamp (2006) and Van Nieuwerburgh and Veldkamp (2006) present a model with endogenous learning, which features large discrete moves in the information about the future economy. These moves are broadly consistent with our model specification for confidence dynamics. The large, discrete moves in investors' uncertainty about future economy also obtain in the costly learning models due to lumpy information, as shown in Bansal and Shaliastovich (2009). Our time-series model is further motivated by strong empirical evidence for large positive moves in the confidence measure in the data, as discussed in Section 4.

Based on these arguments, we specify a following discrete-time jump-diffusion model for the confidence measure, which features persistence in the series and both Gaussian and jump-like innovations:

$$V_{t+1} = \sigma_v^2 + \nu(V_t - \sigma_v^2) + \sigma_w \sqrt{V_t} w_{t+1} + Q_{t+1}. \quad (2.14)$$

The parameters σ_v^2 is the mean value of V_t , ν captures its persistence while σ_w determines the volatility of the smooth Gaussian shock w_{t+1} . The non-Gaussian innovation in the variance process is denoted by Q_{t+1} . We model it as a compound Poisson jump,

$$Q_{t+1} = \sum_{i=1}^{N_{t+1}} J_{i,t+1} - \mu_j \lambda_t,$$

where N_{t+1} is the Poisson process, whose intensity $\lambda_t \equiv E_t N_{t+1}$ corresponds to the probability of having one jump in continuous-time model, while $J_{i,t+1}$ determines the distribution of the size of the jump. Parameter μ_j is the mean of jump size, so subtracting $\mu_j \lambda_t$ on the right-hand side of the above equation ensures that the jump innovation Q_{t+1} is conditionally mean zero⁴. In continuous time, it is possible to guarantee that V_t never falls below zero by placing an upper bound on the volatility parameter σ_w and considering only positive jumps. In the simulations, we consider positive jumps drawn from an exponential distribution, and verify that the conditions for non-negativity of V_t are satisfied.

To capture the dependence of jump probability on the level of variance, we assume that the arrival intensity λ_t is linear in V_t ,

$$\lambda_t = \lambda_0 + \lambda_1 V_t. \quad (2.15)$$

Positive value of the loading λ_1 implies that confidence jumps are more likely when the level of confidence measure V_t is high.

This specification of the dynamics of confidence measure is very similar to the models Broadie et al. (2007) and Eraker and Shaliastovich (2008) for the variance process in continuous time. Such model specification facilitates the analytical solution of the model.

2.5 Recency-Biased Expectation Formation

An important dimension of our model is that the investors do not observe the true expected growth and thus have to learn about it using the past history of the data and current information. To model a belief-updating process of the agents, let us consider a general specification, discussed in Hogarth and Einhorn (1992), where the prior estimate of investors is linearly adjusted by the impact of the recent news:

$$\hat{x}_t = \rho \hat{x}_{t-1} + K_t(\bar{x}_t - \rho \hat{x}_{t-1}). \quad (2.16)$$

⁴Indeed,

$$E_t(Q_{t+1}) = E_t(E_t(Q_{t+1}|N_{t+1})) = E_t(\mu_j N_{t+1}) - \mu_j \lambda_t = 0.$$

In this updating equation, \hat{x}_t is the investors' estimate of the true expected growth x_t , $\rho\hat{x}_{t-1}$ corresponds to the agents' prior belief about the expected state as of the last period, and K_t denotes the adjustment weight which measures the sensitivity of agents' expectations to the recent information in the average signal⁵. Special cases of this belief-updating setup, which correspond to different choices of K_t , include the Bayesian Kalman Filter case as well as the behaviorally motivated recency-biased expectation formation discussed in Kahneman and Tversky (1973) and De Bondt and Thaler (1990). For a broader survey of the behavioral literature, see Barberis and Thaler (2003).

For a particular choice of K_t , the learning specification corresponds to the Bayesian Kalman Filter. When confidence risk V_t is fluctuating, a Kalman Filter weight $0 < K_t < 1$ is time-varying and depends on the model parameters and the confidence measure V_t . In particular, as shown in the Appendix, the Kalman weight given to news today declines as V rises. In the model this weight would on average be 0.5, and vary from 0.2 to 0.75; Bayesian investors optimally choose to put less weight to the recent information when its quality is low (cross-sectional variance is high). It is worth noting that in high confidence risk states investors continue to put positive weight on the signals they receive, however, the Bayesian learning would quantitatively suppress the effects of the recent information and of confidence risks on asset prices by substantially decreasing K_t in periods of high confidence risk (high V_t). In contrast, recency-biased expectations overweigh recently received information. More precisely, empirical evidence in De Bondt and Thaler (1990) highlights two features about recency-biased expectation formation. First, expectations are extreme, that is, more recent information is overweighed by investors; this in our case corresponds to a weight K that is larger than the Kalman filter weight. Second, when uncertainty is high, the over-weighting of recent news increases; this in our case corresponds to significantly more overweighing news in times when V is high. We operationalize the De Bondt and Thaler (1990) recency-bias expectation formation by setting the weight that investors give to recent news to a constant K equal to 0.5 in equation (2.16). This calibration implies that investors considerably overweigh recent news in periods when confidence risk V is high, as the Bayesian weight in these periods is 0.2. In equilibrium, this overweighing of recent news leads to overreaction of asset prices and asset price volatility to confidence risk fluctuations. It is worth emphasizing that quantitatively, equilibrium asset prices would not react considerably to fluctuations in V if learning were Bayesian.

⁵In principle, agents can also use the history of consumption data in their update. However, as discussed in Appendix, at economically reasonable parameter values signals are much more informative about expected growth than consumption, so in our application we can safely disregard learning from consumption.

Under the recency bias specification, investor's expectation formation can be expressed in the following way:

$$\Delta c_{t+1} = \mu + \hat{x}_t + a_{c,t+1}, \quad (2.17)$$

$$\bar{x}_{t+1} = \rho \hat{x}_t + a_{x,t+1}, \quad (2.18)$$

$$\hat{x}_{t+1} = \rho \hat{x}_t + K a_{x,t+1}. \quad (2.19)$$

We emphasize here that unlike the optimal learning case, the weight K is now constant. In particular, the closer K is to one, the bigger is the role of the recency bias in expectation formation — in the limit when $K = 1$, the agents solely rely on recent information and entirely disregard past information. In our application, we set K to be 0.5; which also corresponds to the average value of the weight under the Kalman Filter and at the same time allows for a significant role of the recency bias for expectation formation.

Note that a_c and a_x in equations (2.17)-(2.19) correspond to short-run and long-run shocks, and satisfy the following equations:

$$\begin{aligned} a_{c,t+1} &= x_t - \hat{x}_t + \sigma_t \eta_{t+1}, \\ a_{x,t+1} &= \rho(x_t - \hat{x}_t) + \varphi_e \sigma_t \epsilon_{t+1} + (\bar{x}_{t+1} - x_t). \end{aligned} \quad (2.20)$$

Notably, these short and long-run shocks depend on the prediction error of the investors ($x_t - \hat{x}_t$). Further, as shown in the Appendix, the variance of the prediction error, denoted ω_t^2 , is directly related to the confidence measure:

$$\omega_t^2 = K V_t. \quad (2.21)$$

The variance of the prediction error increases when confidence risk goes up. It is critical to note that the recency bias expectation formation ensures that the variance of the prediction error sharply increases with V_t , as shown in the above equation. In contrast, under Bayesian learning, periods of high confidence risks are associated with lower K_t , which would dampen the role of confidence risks.

The equations (2.17) - (2.19), together with the time-series model for the confidence measure in (2.14) and consumption volatility in (2.7) fully describe the evolution of economy. In the next section, we incorporate preferences of the agents to solve for the equilibrium asset prices in the economy.

3 Model Solution

3.1 Discount Factor

We use the Euler equation of the agent in (2.4) to solve the model.

The equilibrium asset prices in the economy are derived using standard log-linearization of returns, discussed in Campbell and Shiller (1988). In Appendix we show that the equilibrium price-consumption ratio is linear in the expected growth state, consumption volatility and the confidence measure of the investors:

$$pc_t = B_0 + B_x \hat{x}_t + B_v V_t + B_\sigma \sigma_t^2,$$

where the expressions for the loadings are provided in Appendix.

The loading B_x measures the sensitivity of price-consumption ratio to the expected growth prospects. It is positive for $\psi > 1$, so that when the substitution effect dominates the income effect, prices rise following positive news about the expected growth rates. The sign of the loading B_v depends on the preference of the agent for the timing of the resolution of uncertainty. When the agent has a preference for early resolution of uncertainty ($\gamma > 1/\psi$), this loading is negative, so that lack of confidence of investors, i.e. high V_t , leads to a decline in equilibrium asset valuations. Similarly, as in a standard long-run risks model, positive news to consumption volatility decrease equilibrium price-consumption ratio ($B_\sigma < 0$), when agents have preference for early resolution of uncertainty. The magnitude of the consumption volatility loading, however, is smaller in absolute value in the model with confidence risks. Indeed, in the complete information case when true expected growth state is known, consumption volatility σ_t^2 explains all the conditional variation in short-run and long-run consumption shocks. On the other hand, with learning, the volatility of future consumption given investors' information also depends on the investors' confidence measure V_t , as shown in (2.20). Consequently, the contribution of the consumption volatility channel diminishes relative to standard models.

The equilibrium log discount factor can be expressed in terms of the underlying state variables and shocks in the economy:

$$\begin{aligned} m_{t+1} = & m_0 + m_x \hat{x}_t + m_v V_t + m_\sigma \sigma_t^2 \\ & - \gamma a_{c,t+1} - \lambda_x K a_{x,t+1} - \lambda_v \left(\sigma_w \sqrt{V_t} w_{t+1} + Q_{t+1} \right) - \lambda_\sigma \varphi_w \sigma_t w_{c,t+1}. \end{aligned} \tag{3.1}$$

The solutions for the discount factor loadings and prices of risks are pinned down by the model and preference parameters of the investors. Their expressions are provided in Appendix.

The innovations into the discount factor are important as they highlight the risks that investors face in the economy. The risk compensation for the immediate consumption risks $a_{c,t+1}$ is equal to the risk aversion coefficient γ . The price of the long-run risks $a_{x,t+1}$ is given by $\lambda_x K$. As in the standard long-run risks model, the risk compensations for these shocks continue to be positive⁶. The novel dimension of our paper is that the confidence risks ($\sigma_w \sqrt{V_t} w_{t+1} + Q_{t+1}$) are priced. Notably, confidence jump shocks Q_{t+1} are the source of the jump risk in the economy, even though there are no jump risks in the underlying consumption. When agents have preference for early resolution of uncertainty, the price of confidence risks λ_v is negative, as investors dislike positive shocks in the confidence measure (high uncertainty). It is worth noting that our specification of recency bias expectation formation affects the equilibrium market prices of risk; in particular, the market price of confidence risks is much smaller if investors use Bayesian learning. On the other hand, when investors have recency bias, they overreact to the current news in the average signal and do not reduce the weight to recent information, which enhances the role of confidence risks for the level and variation of the risk premium in the economy. Finally, the price of consumption volatility risks $\varphi_w \sigma_t w_{c,t+1}$ is given by λ_σ . It is negative as in a standard long-run risks model; however, due to learning, the magnitude of the risk compensation for consumption volatility risks decreases in absolute value, as discussed above.

Using the solution to the discount factor, we can derive the expressions for the equilibrium risk-free rates in the economy. Interest rates are linear in the expected growth state, the confidence level of the investors and consumption variance, where the expressions for price of zero coupon bonds with n months to maturity are given in the Appendix:

$$p_{t,n} = F_{0,n} + F_{x,n} \hat{x}_t + F_{v,n} V_t + F_{\sigma,n} \sigma_t^2. \quad (3.2)$$

In particular, real yields $r_{f,t}$ increase in the expected growth state, and decrease when the confidence of investors drops or consumption volatility increases.

3.2 Equity Prices

To obtain implications for the equity prices, we consider a dividend process of the form

$$\Delta d_{t+1} = \mu_d + \phi(\Delta c_{t+1} - \mu) + \varphi_d \sigma_t \eta_{d,t+1}, \quad (3.3)$$

where $\eta_{d,t+1}$ is a dividend shock independent from all other innovations in the economy. We continue to maintain the assumption that the average signal data is much more informative about the expected growth than consumption or dividend data, so

⁶Note that as investors cannot separate true short-run and long-run consumption innovations, the price of long-run risks decreases while the price of short-run risks goes up relative to complete information case; this is consistent with Croce et al. (2006).

investors learn about the expected state only from the average signals (see specification (2.17)-(2.19)).

The equilibrium price-dividend ratio is linear in the expected growth state and the confidence level of the investors:

$$pd_t = H_0 + H_x \hat{x}_t + H_v V_t + H_\sigma \sigma_t^2, \quad (3.4)$$

where the solutions for the loadings are provided in the Appendix. Similar to the valuation of consumption asset, equity prices increase in expected growth factor and decrease when the confidence measure is high or the consumption volatility is high. In particular, large positive moves in V_t endogenously translate into large jumps in asset valuations and returns. Indeed, the equilibrium log return on the dividend asset satisfies

$$\begin{aligned} r_{d,t+1} = & \mu_r + b_x \hat{x}_t + b_v V_t + b_\sigma \sigma_t^2 + \phi a_{c,t+1} + \kappa_{d,1} H_x a_{x,t+1} \\ & + \kappa_{d,1} H_v \left(\sigma_w \sqrt{V_t} w_{t+1} + Q_{t+1} \right) + \kappa_{d,1} H_\sigma \varphi_w \sigma_t w_{c,t+1} + \varphi_d \sigma_\epsilon \eta_{d,t+1}. \end{aligned} \quad (3.5)$$

As the return beta to confidence measure is negative ($H_v < 0$) when investors lose confidence about their estimate of expected growth, changes in the confidence measure are substantially magnified due to investors' concerns about the long-run growth, and can have large negative impacts on the equilibrium asset prices.

The recency bias and fluctuating confidence of investors channel can account for the large moves in asset prices in the data. Notably, the large negative moves in equilibrium returns obtain even though there are no corresponding large moves in the real consumption. Further, the conditional variance of returns is linear in consumption volatility and confidence measure, so that jumps in confidence measure also translate into simultaneous jumps in market variance.

The expected excess returns given the information of the agent depend on the consumption volatility and confidence measure states:

$$E(r_{t+1} | \mathcal{I}_t) - r_{f,t} = r_0 + r_v V_t + r_\sigma \sigma_t^2, \quad (3.6)$$

where the loadings r_0, r_v and r_σ are determined by the model parameters. The total equity premium can be further decomposed into the compensations for confidence risks and short-run, long-run and consumption volatility risks. The fluctuations in confidence risk drive variation in the risk premium. Hence, recency-biased expectation formation and confidence risk jumps will lead to the corresponding jumps in the risk premium, which enhances the time-variation in expected returns.

Notably, the confidence measure drives both the expected returns and the price-dividend ratio in (3.4), even if consumption volatility is constant. Hence, in the model,

as in the data, asset valuations predict future returns. We examine the quantitative implications of return predictability in Section 5.

4 Data

4.1 Evidence on Large Return Moves

The data on market returns and the risk-free rate correspond to monthly observations of broad value-weighted portfolio returns and average 1 month yields from CRSP from 1927 to 2007. Return data is adjusted by inflation to convert to a real series. Real consumption data come from the BEA tables of real expenditures on non-durable goods and services; the quarterly consumption data is available from mid 1947 to end of 2007, while monthly data are from 1959 to 2007. Real GDP series also comes from the BEA tables from 1947 to 2007, quarterly, while the index of industrial production series is taken from the FRED dataset from 1927 to end of 2007, monthly.

Table 1 presents summary statistics for market returns and the risk-free rate. The average equity premium in the sample is 6%, and the level of the real risk-free rate is about 0.6%. The magnitudes of the equity premium and the risk-free rate constitute the well-known equity premium and risk-free rate puzzles (see e.g. Mehra and Prescott, 1985). Further, the market return is very volatile, as its standard deviation is almost 19%, while the volatility of the risk-free rate in the sample is about 1%. The risk-free rate series is very persistent, with an autoregressive coefficient of 0.98. Summary statistics for the consumption growth rate are shown in the first column of Table 8. Over the long historical sample, consumption growth rate averages 2%, its volatility is 2% and it is mildly persistent with an autoregressive coefficient of 0.4.

The evidence from the higher-order moments of returns suggest that the distribution of market returns exhibits substantial heavy tails and negative skewness. Table 1 shows that the kurtosis of market returns is 9.7, while its skewness is -0.44 . The kurtosis and skewness of Normal distribution are 3 and 0, respectively, so that the excess kurtosis and negative skewness in the return series are indicative of abnormal and predominantly negative moves which thicken the left tail of the distribution of returns.

For a direct empirical evidence on large return moves in the data, we identify two standard deviation or above innovations in the series.⁷ Table 2 shows the number of identified large moves in returns, as well as the average magnitude of returns in those periods. On a monthly frequency, we observe 54 two standard deviation or

⁷Calculation of standardized innovations are based on AR(1)-GARCH(1,1) fit to the data.

above return moves over 80-year time-period. This translates into large moves at a frequency of once every 18 months. Further, 70% of these moves are negative, which explains the reasons for a negative skewness of returns in the data. The frequency of large negative moves below two standard deviations is almost two times higher than under the Normality assumption, which goes a long way to explain the heavy tails of returns in the data relative to the Normal distribution.⁸ Figure 1 show the time-series of monthly returns from 1927 to 2007 and identified periods with large return moves. Most prominent large move occurred in October 1987, when the monthly return dropped 6 standard deviations below its mean. Other examples include large moves in the beginning of 1930s, and a series of negative moves in the 1960s.

4.2 Confidence Data

In the model, investors can observe the confidence measure in the economy. To give quantitative content to the model, we directly measure investors' confidence using survey data. More specifically, we rely on the cross-section of forecasts from the Survey of Professional Forecasters (SPF) and construct confidence measure as a cross-sectional variance of the average forecast in the data. These calculations are consistent with the theoretical setup of the model and follow equation (2.13), assuming that each forecaster reports an individual expectation of the next-period series, so that the forecasts correspond to the signals in the model.

The Survey of Professional Forecasts started in the last quarter of 1968 as a joint project of the American Statistical Association and the National Bureau of Economic Research; in 1990 it was taken by the Federal Reserve Bank of Philadelphia. The data set contains quarterly forecasts on a variety of macroeconomic and financial variables made by the professional forecasters, who largely come from the business world and Wall Street, see Croushore (1993) for details and Zarnowitz and Braun (1993) for a comprehensive study of the survey.

We use the cross-section of individual forecasts from the SPF to calculate the average (consensus) forecast and the confidence measure for real GDP series for the period of 1968 Q4 to 2008 Q1⁹. As we observe the identification codes for each forecaster, we can pair the forecasts of the price index and nominal GDP to back out the implied forecast for the corresponding real series. Specifically, for each quarter t let $NGDP_{i,t}$ denote the next quarter forecasts of nominal GDP made by the forecaster i , while $P_{i,t}$ stand for the next quarter forecast of the price level. If n_t is the number

⁸It is worth noting that as one moves from monthly to quarterly and annual frequency, returns are smoothed out due to time-aggregation, and the frequency of large moves declines.

⁹Prior 1992, the GDP forecasts are for nominal GNP and GDP price index forecasts are for GNP deflator. In the data, the growth rates for the realized real GDP and GNP are very close to each other, with correlation in excess of 0.99.

of forecasts, then we calculate the average forecast for the log real GDP (*RGDP*) growth rate as

$$\overline{\Delta \log(RGDP)}_t = \frac{1}{n_t} \sum_{i=1}^{n_t} \left(\log \frac{NGDP_{i,t}}{NGDP_t} - \log \frac{P_{i,t}}{P_t} \right), \quad (4.1)$$

where $NGDP_t$ and P_t are the current values of the series known to the forecasters¹⁰.

We can further calculate the cross-sectional variance of the forecasts at each point in time, and compute the estimate of the uncertainty in the average forecasts, which corresponds to the confidence measure in the model. That is, based on the real GDP forecasts,

$$\begin{aligned} V_t &= \frac{1}{n_t} Var \left(\log \frac{RGDP_{i,t}}{RGDP_t} \right) \\ &= \frac{1}{n_t} \left(\frac{1}{n_t - 1} \sum_{i=1}^{n_t} \left(\log \frac{NGDP_{i,t}}{NGDP_t} - \log \frac{P_{i,t}}{P_t} - \overline{\Delta \log(RGDP)}_t \right)^2 \right). \end{aligned} \quad (4.2)$$

To make the inference robust to possible outliers and errors, we delete observations which are more than two standard deviations from the sample mean. We use a similar approach to construct the empirical confidence measures based on the forecasts of industrial production index.

David and Veronesi (2008) and Buraschi and Jitsov (2006) use very similar computations to obtain the uncertainty and disagreement measures in the economy, which rely on the cross-sectional dispersion in forecasts from the SPF. Anderson, Ghysels, and Juergens (2007) also associate the uncertainty with the dispersion in professional forecasters, and assign different weights across forecasts.

In the data, we find that the SPF average forecasts are persistent and very informative about future economy, even controlling for the past history of the data. The persistence of the average forecast of next-quarter real GDP growth is 0.72, relative to 0.25 in the underlying real GDP growth. Further, we assess the predictability of future macroeconomic series by the Kalman Filter estimate of the expected growth from the history of the series, and the SPF average forecast. We find that the addition of the survey forecast doubles the adjusted R^2 and makes the Kalman Filter estimate insignificant at 1 quarter horizon, while the R^2 increases from 4% to 26% as we add the survey average forecast to predict the real GDP growth 4 quarters ahead. The results for the other macroeconomic series, such as industrial production and corporate profit growth rates, are very similar.

¹⁰In the data, the average number of forecasts is 34.

In the top panel of Table 7 we report summary statistics for the square-root of the confidence measure scaled by the average volatility of the underlying series, based on the next quarter forecasts of real GDP and industrial production. These statistics are very similar for the two series. The level of the confidence measure in the data is about 1/15th of the volatility of the underlying series. This implies that the uncertainty of investors regarding their estimate of expected growth is on average quite small. Indeed, at the calibrated consumption parameters, the two standard deviation band for the investors' estimate of the expected growth is [1.8%, 2.2%], assuming the expected consumption growth of 2%. Our confidence measure is related to the sentiment indexes considered in the literature. For example, the correlation of the confidence measure based on real GDP forecasts with the University of Michigan consumer confidence index and the investment confidence index constructed in ? is about -50%. In general, our model provides a well-articulated equilibrium framework to analyze and interpret behaviorally-motivated investor sentiment.

Confidence measure fluctuates in time, and the standard deviation of the confidence measure scaled by the volatility of the underlying series is about 4%. We show the plot of the time-series of confidence measure on Figure 2. It is important to note on the graph that confidence measure in the data exhibits large, positive moves. We discuss the connection between the large moves in the confidence measure and returns in the next section, and statistically evaluate the evidence on fluctuations and large moves in the confidence measure in Section 5.

4.3 Predictability Evidence

In this section, we empirically evaluate the connection between large moves in returns and all the relevant macroeconomic variables, such as consumption and confidence measure. To do so, we construct a two standard deviation or above move indicators in returns and the corresponding macroeconomic series. We compute the dynamic relationship between these large move indicators in the data, which is summarized in Table 3. We also consider the link between the fluctuations in the confidence measure and the level and variation in asset prices and report our empirical findings in Table 4.

In the left panel of Table 3 we document the correlations between the large move indicators (zero-one variables) in returns and contemporaneous or future large move indicators in the macroeconomic series. The Table suggests that the contemporaneous correlation between large move indicators in returns and in consumption is essentially zero: the magnitudes of the correlations are -0.02 and -0.03 at monthly and quarterly frequency, respectively, and are statistically insignificant. Further, large moves in returns today do not anticipate jumps in future consumption, as the correlations of jump indicators in returns and jump indicators in future consumption 6 and 12 months

ahead are economically and statistically insignificant. This evidence indicates that large moves in returns do not reflect jumps in consumption contemporaneously or in the future. Very similar conclusions obtain when we consider industrial production or real GDP data; we omit the results to save space. On the other hand, the empirical evidence suggests that the large moves in the confidence measure in the data are significantly related to contemporaneous large moves in returns. It is important to note that the contemporaneous correlation of large move indicators in returns and in the confidence measure is 34% and is statistically significant. This evidence highlights the importance of confidence risks in asset prices. Further, the large moves in returns are not correlated with large moves in the confidence measure in the future, indicating that the large moves in the confidence measure are relatively short-lived. It is our objective to replicate this empirical correlation pattern in the model.

In the middle panel of Table 3 we report the correlation between the jump indicators in macroeconomic variables and the level of returns. In particular, we find that the correlations of returns with contemporaneous or future jump indicators in consumption are essentially zero; using industrial production or real GDP measures produces similar results. This evidence implies that periods with large moves in macro variables do not correspond to any systematic moves in returns. However, the contemporaneous correlation between the jump indicator in the confidence measure and the level of returns is equal to -0.32 and is highly significant. That is, returns are significantly lower in periods when the confidence measure experiences a large positive shock.

The right panel of Table 3 presents some evidence for a negative contemporaneous correlation between large return move indicators and the contemporaneous level of consumption growth, though, the correlation is close to zero for future consumption growth. This evidence suggests that large moves in returns are more likely when consumption growth is relatively low; however, it does not imply that jumps in returns are driven by large moves in consumption, as we discussed above.

Our main empirical findings is that in the data, there is no persuasive link between the large moves in real consumption and large moves in returns at the considered frequencies; however, the large moves in the confidence measure are significantly related to contemporaneous large moves in returns. In Table 4 we further show that the confidence measure has significant information about the asset valuations and the variation in asset prices. In the first panel of Table 4 we consider contemporaneous projections of the quarterly price-dividend ratio on the confidence measure and the conditional variance of real GDP. The loadings on the conditional variance and confidence measure are negative and statistically significant, which provides empirical support for the economic risk channels featured in our model. The results further imply that the fluctuations in the confidence measure have additional information about prices beyond the standard time-series volatility; we also verify that the confidence

measure has significant information about the future price-dividend ratio for several years ahead. Hence, the confidence risk is a potentially important channel that drives asset prices in the data.

We also consider the predictability of the variance of returns based on the implied volatility (VIX) index by the confidence and volatility measures. To highlight the tail properties of the return variance, we run the conditional quantile regression of the variance of returns on the confidence and GDP volatility measures. The results reported in second panel of Table 4 suggest that the confidence measure has important information about the right tale of the conditional distribution of variance of returns, while macroeconomic variance does not. That is, large upward moves in the variance of returns are related to the fluctuations in the confidence measure, rather than to the conditional volatility of the fundamentals of the economy. This evidence is consistent with Buraschi and Jitsov (2006), who show that the cross-sectional dispersion of forecasts from the survey data has information about the level and slope of the option smile and realized volatility of returns.

Overall, our empirical results indicate that fluctuations and large moves in the confidence measure contain important information for the asset prices in the data. We further discuss the important features of the confidence dynamics in the data, and present strong statistical support for the fluctuations and jump-like shocks in the confidence measure in the next Section.

5 Estimation of Confidence Dynamics

In this section, we focus on the empirical evidence for the confidence measure in the data, and present statistical support for the key features of our model specification of the confidence measure dynamics.

To quantitatively evaluate the evidence on fluctuation and large moves in the confidence measure, we fit the jump-diffusion model for the confidence measure presented in Section 2.4:

$$V_{t+1} = \sigma_v^2 + \nu(V_t - \sigma_v^2) + \sigma_w \sqrt{V_t} w_{t+1} + Q_{t+1}.$$

The shock w_{t+1} is the Gaussian innovation into the confidence measure, while Q_{t+1} is Poisson jump-like shock. The confidence measure jumps come from exponential distribution with mean μ_j . To capture the predictability of jumps, we allow the jump probability (intensity) to depend on the level of confidence, $\lambda_t = \lambda_0 + \lambda_1 V_t$. Positive estimate of λ_1 then suggests that large moves in the confidence measure are more likely to occur when the confidence measure V_t is high. For comparison, we also estimate a restricted model for the confidence measure where we shut off the jump channel, that is, where we set μ_j, λ_0 and λ_1 to zero. With these restrictions imposed,

there are no jumps in the confidence measure, so that the confidence measure follows a standard square-root process driven by Gaussian innovations w_{t+1} .

Table 5 shows the estimation results for the considered series. The unconditional mean of the confidence measure is quite small; indeed, σ_v is estimated at 0.12%, annualized, for real GDP and 0.18% for industrial production, which corresponds to 1/15 of the volatility of the underlying series (see also Table 7). The estimated persistence of the confidence measure is 0.7, so that the half-life of confidence shocks is about 6 months. This evidence suggests that the fluctuations in the confidence measure are very different from the variations in the aggregate volatility in the economy – the shocks to real GDP volatility are much more persistent, with the half life of 3 years. Hence, the confidence risk channel is distinct from the macroeconomic volatility risks in the economy.

The estimation results of the full model with confidence jumps suggest that Poisson jumps capture above 80% of the variation in the the confidence measure. The mean jump size is quite large: its estimated value is close to an unconditional level of confidence σ_v^2 . Further, we find that the probability of confidence jumps strongly depends on the level of the confidence measure: the jump intensity loading λ_1 is positive in the data and significant for both series. The fluctuations in the confidence measure are hard to explain using Gaussian model. The estimation of the square-root specification of the confidence measure dynamics suggests that, to capture a sizeable variation in the confidence measure, the scale parameter σ_w increases four-fold, relative to the model with confidence jumps. Nevertheless, the Gaussian model is misspecified, as it fails to capture the spikes and the variation in the confidence measures. The distribution of the extracted confidence shocks \hat{w}_{t+1} is heavy-tailed and positively-skewed: the sample kurtosis for these shocks is 20 for real GDP and 14 for industrial production, and sample skewness is 3 for both series. We additionally do a likelihood ratio test where the jump parameters μ_j , λ_0 and λ_1 are jointly equal to zero. As shown in the Table 5, the p-value for the test is indistinguishable from zero. This result suggests that the Gaussian model for the confidence measure is strongly rejected in favor of the model with confidence jumps.

Overall, the empirical evidence in the data provides a strong support for our model specification of the confidence measure dynamics, which features large positive moves in the series. We also show that these confidence jumps are more likely to happen when V_t is high. We match these key empirical features in the calibration of the confidence measure in the data, and discuss the asset-pricing implications of the confidence risks in the next Section.

6 Asset-Pricing Implications

6.1 Calibration

The model is calibrated on a monthly frequency. The baseline calibration values for the preference and endowment dynamics parameters, which are reported in Table 6, are very similar to the ones used in standard long-run risks literature (see e.g. Bansal and Yaron, 2004). Specifically, we let the subjective discount factor δ equal 0.9992. The risk aversion parameter is set at 10, and the intertemporal elasticity of substitution at 1.5. This configuration implies that the agent has a preference for early resolution of uncertainty, which has important implications for equilibrium prices, as we discussed in Section 3. As for the consumption dynamics, we set the persistence in the expected growth ρ at 0.975. The choice of φ_e and σ ensures that the model matches the historic volatility and persistence of consumption growth. Similar to Bansal, Kiku, and Yaron (2007a), we set the persistence of the consumption variance to $\nu_c = 0.995$, and calibrate the volatility of volatility parameter $\varphi_w = 5.19e - 04$. To calibrate dividend dynamics, we set the leverage parameter of the corporate sector ϕ to 2.75 and φ_d to 3 to match the properties of dividend growth rates in the data. We calibrate the model on monthly frequency and then time-aggregate to annual horizon. Table 8 shows that we can successfully match the mean, volatility, auto-correlations and variance ratios of the consumption dynamics in the data.

We calibrate confidence dynamics to match its unconditional moments in the data, as well as its conditional distribution. Consistent with the empirical evidence, we set the level of the confidence measure σ_v to be 1/15 of the volatility of consumption growth. The calibrated persistence coefficient $\nu = 0.91$ implies that the half-life of confidence shocks is 7 months, which is consistent with the estimated half-life of shocks to empirical confidence measure in the data (see Table 5). In the estimation of the confidence measure dynamics we found that most of the variation in the confidence measure in the data is coming from non-Gaussian shocks. Then, for parsimony of parameters, we set the volatility scale parameter σ_w to 0, so that the confidence measure in the model is driven purely by Poisson jumps. To calibrate the intensity of jumps $\lambda_0 + \lambda_1 V_t$, we set λ_0 to 0.18, as in the data, and divide it by 3 to convert to monthly frequency. Similarly, we calibrate the parameter $\lambda_1 = 20$ in the middle of range in the data and multiply by 40000 to adjust for the scaling of the confidence measure used in the estimation of the model. At the calibrated parameter values, the model-implied average frequency of jumps is one every 5 months. Finally, the jump size is set to match the distribution of the confidence measures in the data. Specifically, we set the mean jump μ_j to twice the unconditional level of confidence measure σ_v^2 to target skewness and kurtosis of the confidence measure in the data.

In simulations, we verify that the calibrated distribution of the confidence measure in the model can match the key features of the confidence measure in the data. We summarize the results in Table 7, where we report the moments of the square-root of confidence measure V_t scaled by the average volatility of the underlying series in the data and in the model. Such a normalization facilitates the comparison of the confidence measures based on different macroeconomic series; indeed, as the first panel of the Table suggests, the moments of the confidence measure based on real GDP and industrial production forecasts are very similar. The Table 7 shows that we can match well the unconditional distribution of the series in the data, as all the statistics in the data are close to their counterparts in the model, and are within the 5% – 95% confidence band. In particular, the calibration of the jump component can realistically match the non-Gaussian dimension of the distribution, as the skewness and kurtosis of confidence measure in simulations are close to the values in the data.

6.2 Jumps and Distribution of Returns

In our model, due to learning, the investors' confidence about their estimate of expected growth impacts their beliefs about future consumption, which influences equilibrium asset valuations and the risk premia in the economy. We illustrate this channel on Figure 3, where we show the distribution of the true expected growth given the information of the agents. In the standard long-run risks model (left panel), the true expected growth is observable by investors, so its distribution is degenerate and centered at the true value. On the other hand, in our model with learning and fluctuating confidence risks (right panel), the distribution of the true expected growth depends on investors' confidence measure, so that when the quality of signals is low, the confidence measure increases and the agents face high uncertainty about the true state. The fluctuations in confidence, magnified by the recency bias of investors, lead to the time-variation in agents' uncertainty about future consumption, even if the fundamental consumption volatility is constant.

Fluctuating confidence risks play an important role in accounting for the key features of the return distribution in the data. For illustration, we first shut off the conditional volatility of consumption growth, and focus on the implications of the fluctuations in investors' confidence for the asset prices. On Figure 4 we first show the unconditional distribution of consumption growth, which is Gaussian. On the middle panel, we plot the unconditional distribution of returns when the confidence measure is set constant, that is, when the fluctuating confidence risks are absent. As the consumption volatility is constant, the distribution of returns in this case is Gaussian as well. On the other hand, when the confidence measure fluctuates, large positive moves in the confidence measure endogenously translate into large negative jumps in returns, which can explain negative skewness and heavy-tails of the distribution of

returns in the data, even though consumption growth is Gaussian. The unconditional distribution of returns in this case, shown on the right panel of Figure 4, exhibits heavy tails and negative skewness, as in the data.

We summarize the model output for the return distribution in Table 9. As shown in the Table, in the full model with confidence risks and time-varying consumption volatility, the average return is 6.7%, and its volatility is almost 20%, which match the statistics in the data (see Table 1). The distribution of returns is heavy-tailed and negatively skewed. The kurtosis of market returns is 8.6, and its skewness is equal to -0.85; these values are close to the estimates in the data. Notably, the non-Gaussian features of the return distribution are due to the fluctuations and large moves in the confidence measure. Indeed, we verify that the magnitudes of the kurtosis and skewness of market returns are very similar even if the consumption volatility is set constant. Naturally, if both the confidence measure and consumption volatility are constant, the returns are Gaussian, as shown in Table 7. For comparison, we also report the model output for the return distribution in the standard long-run risks model. In this case, the true expected growth is known, so that the confidence risk channel is absent. The Table 9 suggests that in the standard model, even though returns exhibit somewhat heavier tails than Normal due to time-varying consumption volatility, the magnitude of the kurtosis of return distribution (4) is too small relative to the data (10). In addition, this specification cannot capture negative skewness of market returns in the data. These results provide additional evidence for the importance of confidence risks channel for the asset prices.

We verify that the model can broadly match the moments of the risk-free rate in the data. The mean risk-free rate is about 1% in the data and in the model. It is very persistent, with an autocorrelation coefficient of 0.98 in the data, and 0.91 in the model, respectively. The model-implied standard deviation of the risk-free rate is somewhat less than in the data, 0.4 relative to 1.1. The skewness and kurtosis of the risk-free rate distribution implied by the model are -1.4 and 6.8 , which are close to the estimates in the data.

In our model, asset prices exhibit negative jumps due to large positive moves in the confidence measure, while there are no jumps in consumption. Thus, our model can explain the puzzling evidence in the data for significant contemporaneous link between large moves in returns and in the confidence measure, and lack of connection between large move in returns and large moves in consumption, documented in Section 4.3. We present the quantitative results from the model in Table 10. Using the same large move identification approach, we find that the large moves in returns occur about once every 18 months in the model, which matches the frequency of large asset price moves in the data. Further, the model-implied correlation between large move indicators in returns and in consumption is 0, both contemporaneously and in the future. On the other hand, the correlation between large move indicators in returns

and in the confidence measure is 45%, which is close to the estimate of 34% in the data. The magnitudes of the correlation of large return move indicators with future large move indicators in confidence measure 6 and 12 month ahead are all zero, as in the data, as the confidence jump shocks are relatively short-lived. Next, our model can replicate zero correlation of consumption jump indicator and the level of returns in the data, both contemporaneously and in the future. Further, the correlation of returns with the jump indicator in confidence measure is -45% in the model, which is close to -32% in the data. As in the data, this correlation drops to zero using future confidence measure 6 and 12 months ahead. Finally, our model implies zero correlation of the jump indicator in returns and the level of consumption growth. As we discussed in Section 4.3, these correlations are somewhat negative in the data, though, they are close to zero for the future consumption. The correlation of large move indicators with the level of the confidence measure is also somewhat high in the model (38%), relative to 11% in the data.

Overall, the results suggest that our model with fluctuating confidence risks can account for the key features of the return distribution, and explain the connection between large moves in returns and macroeconomic series in the data. In the next section, we discuss the implications of the confidence risks for the risk compensation in the economy.

6.3 Predictability and Risk Compensation

As in a standard long-run risks specification, in our model short-run, long-run and consumption volatility risks are priced. The novel dimension of our model is that the confidence risks also receive risk compensation. Notably, confidence jump risks are priced even though there are no corresponding jumps in the consumption process. Table 11 reports prices of risks in the standard model with complete information, and in our model with learning and fluctuating confidence risks. As can be seen in the Table, the risk compensations for long-run and consumption volatility risks are smaller, while the compensation for short-run risks is somewhat higher in the model with learning and fluctuating confidence. For example, the price of risk for 1 standard deviation shock to the long-run growth is 8.5% in our model relative to 8.8% in the standard specification, while the price of consumption volatility shocks is -4.2% in our model relative to -9.1% in the standard specification. The compensation for the one standard deviation shock to the confidence measure is -6.6% in our model; naturally, in standard long-run risks specification the true expected growth state is known, so that the confidence risk channel is absent.

In Table 12 we show the model-implied equity premium for the dividend asset, as well as its decomposition into short-run, long-run, consumption volatility and confidence measure risks in the economy for different model specifications. For comparison,

the first panel presents the results for a standard long-run risks model, where investors observe expected growth process, and consumption volatility is time-varying. The model delivers a total risk premium of 5.1%. Most of the risk premium comes from the compensation for long-run and consumption volatility risks, 1.7% and 2.3%, respectively, while immediate consumption shocks contribute about 1.2% to the total compensation. These magnitudes are consistent with the long-run risks literature, see e.g. Bansal and Yaron (2004).

Our full model with confidence risks and time-varying consumption volatility generates an unconditional equity premium of about 5.3%. Confidence shocks and long-run risks shocks contribute 1.7% each to the total premium. It is important to note that as the calibrated confidence measure is driven by jump shocks, the compensation for confidence risks thus determines the compensation for jump risks in the economy. That is, the jumps in the market return demand 1.7%, or one-third, of the total equity premium in the economy. This magnitude is consistent with other studies, see e.g. Broadie et al. (2007), Singleton (2006) and Pan (2002), who use option prices data and other empirical approaches.

Notably, the compensation for the consumption volatility risks decreases relative to standard model from 2.3% to 0.8%. That means that quantitatively, the role of the consumption volatility channel for the asset valuations and risk premia in the economy diminishes once we introduce learning and fluctuating confidence risks. Without the time-varying consumption volatility risks, the total equity premium would fall by 0.8%, which corresponds to the risk compensation for the consumption volatility shocks. On the other hand, as shown in Table 12, without fluctuating confidence risks, the total equity premium would decrease to below 3%. This highlights the importance of confidence risks for the risk compensation in the economy.

Using the analytical results and calibrated model parameters, we can further calculate the magnitudes of the conditional risk premium using the empirical measure of investors' confidence based on the GDP forecasts, and the conditional volatility of real GDP. Specifically, the quarterly expected excess returns in the model are given by,

$$E(r_{t+1} + r_{t+2} + r_{t+3} | \mathcal{I}_t) - rf_{t,3} = A_0 + A_v V_t + A_\sigma \sigma_t^2,$$

for quarterly risk-free rate $rf_{t,3}$. The loadings A_0 , A_v and A_σ are pinned down by model and preference parameters. For V_t and σ_t^2 , we substitute the annualized confidence measure and conditional variance of real GDP, scaled to match the calibrated level of consumption volatility. This allows us to compute the quarterly risk premium in the sample implied by our model, and decompose it into the contributions for short-run, long-run, consumption volatility and confidence measure risks. We show the empirical results in the last panel of Table 12. The total equity premium is 5.6%, which agrees with the estimate in the sample. Most of the total risk premium is explained by the long-run shocks (1.9%) and confidence risks (2%), while the remaining 1% goes to

the immediate consumption shocks and 0.7% to consumption volatility shocks. These estimates are close to the unconditional values in the model.

The time-variation in the confidence of investors in equilibrium generates predictability of equity returns by price-dividend ratios, as both the expected excess returns and asset valuations in the model are time-varying with V_t and σ_t^2 . In Table 13 we report the results from the projections of future excess returns on price-dividend ratios in the data and in the model. In full specification, the model delivers the R^2 of 16% at 5 year horizon, relative to 18% in the data. Confidence risk channel plays a key role to match the predictability of returns in the data. Indeed, when consumption volatility is constant, the R^2 decreases only slightly to 14%. With constant confidence and constant consumption volatility, however, the small-sample R^2 goes down to 4%. Notably, the standard errors on slope coefficient and R^2 are quite large, and the small-sample slope estimates have a well-known downward bias. We also verify the predictability of future consumption growth by the current price-dividend ratio is consistent with the data, as shown in Table 14.

7 Conclusion

We develop a long-run risks type model in which investors learn about the unobserved expected growth using a cross-section of signals. The fluctuations in the quality of signals give rise to a novel learning-based risk channel in the model which we call confidence risks. This channel enhances the ability of the model to explain the joint distribution of the asset prices and consumption growth, time-variation in the equity risk premium, predictability of excess returns, and other key asset market facts.

To model the learning process of the agents, we specify a belief-updating model which allows us to incorporate a notion of the recency bias in forecast formation of investors. Our specification of the recency bias captures the intuition that as the information quality worsens, agents do not optimally decrease the weight attached to the new signals and therefore overweigh the impact of recent information on their forecasts. Therefore, in equilibrium, asset prices decline more sharply (i.e. overreact) when confidence risk rises relative to the case of optimal learning.

We construct a confidence measure in the data as a cross-sectional variance of the average forecast from the Survey of Professional Forecasters. While the level of the confidence measure is quite small, the series exhibits significant variation across time. We find that fluctuations in confidence can not be attributed to smooth Gaussian innovations only, and present statistical evidence for jump-like component in the dynamics of the series. Further, we show that large moves in the confidence measure in the data are significantly related to large moves in asset prices at frequencies of 18 months, while in the data there is no link between large moves in the real economy and

in returns at the considered frequencies. Similarly, confidence measure has significant information about the asset prices and the right tale of the distribution of the variance of returns, even controlling for the aggregate volatility in the economy.

We calibrate the model to match the key dimensions of consumption growth and the confidence measure in the data. The model with recency bias and fluctuating confidence of investors can explain the heavy-tailed and negatively skewed distribution of asset prices, even though fundamental consumption volatility is constant and there are no jumps in consumption. We document that the correlations of large moves in returns with large moves in consumption and confidence measure match the evidence in the data. At the calibrated model parameters, confidence jump risks contribute about one-third to the total equity premium, which highlights the importance of fluctuations in investors' confidence for the asset markets. The confidence jump risk channel can account for the predictability of excess returns and future consumption growth by the price-dividend ratios in the data. The results suggest that confidence jumps risk plays an important role to economically explain the asset prices and sources of risk in the economy.

A Kalman Filter Specification

A.1 Learning Dynamics

Let us first start with a general solution to the Kalman Filter when agents optimally learn about the unobserved state using the history of consumption and signal data.

Using the dynamics of the underlying economy in (2.5)-(2.6) and the specification of signals in (2.8), we obtain that the distribution of the states given the current information set and next-period confidence measure is conditionally Normal:

$$\begin{bmatrix} x_{t+1} \\ \Delta c_{t+1} \\ \bar{x}_{t+1} \end{bmatrix} \mid \mathcal{I}_t, V_{t+1} \sim N \left(\begin{bmatrix} \rho \hat{x}_t \\ \mu + \hat{x}_t \\ \rho \hat{x}_t \end{bmatrix}, \Sigma_{t+1} \right), \quad (\text{A.1})$$

where the variance-covariance matrix is given by,

$$\Sigma_{t+1} = \begin{bmatrix} \rho^2 \omega_t^2 + \varphi_e^2 \sigma_t^2 & \rho \omega_t^2 & \rho^2 \omega_t^2 + \varphi_e^2 \sigma_t^2 \\ \rho \omega_t^2 & \omega_t^2 + \sigma_t^2 & \rho \omega_t^2 \\ \rho^2 \omega_t^2 + \varphi_e^2 \sigma_t^2 & \rho \omega_t^2 & \rho^2 \omega_t^2 + \varphi_e^2 \sigma_t^2 + V_{t+1} \end{bmatrix}. \quad (\text{A.2})$$

The innovation representation of the system can then be written in the following way:

$$\Delta c_{t+1} = \mu + \hat{x}_t + a_{c,t+1}, \quad (\text{A.3})$$

$$\bar{x}_{t+1} = \rho \hat{x}_t + a_{x,t+1}, \quad (\text{A.4})$$

$$\hat{x}_{t+1} = \rho \hat{x}_t + K_{1,t+1} a_{c,t+1} + K_{2,t+1} a_{x,t+1}, \quad (\text{A.5})$$

where the Kalman Filter weights and the update for the filtering variance ω_t^2 satisfy standard equations

$$\begin{aligned} K_{t+1} &= \Sigma_{t+1}^{12} (\Sigma_{t+1}^{22})^{-1}, \\ \omega_{t+1}^2 &= \Sigma_{t+1}^{11} - \Sigma_{t+1}^{12} (\Sigma_{t+1}^{22})^{-1} \Sigma_{t+1}^{21}, \end{aligned} \quad (\text{A.6})$$

where the superscripts refer to the partitioning of Σ_{t+1} into four blocks, such that Σ_{t+1}^{11} is the (1,1) element of the matrix, Σ_{t+1}^{12} contain the elements from the first row and second and third columns, etc. The explicit solutions for the Kalman Filter weights satisfy

$$K_{1,t+1} = \frac{\rho \omega_t^2 V_{t+1}}{(\omega_t^2 + \sigma_t^2) V_{t+1} + (\varphi_e^2 \sigma_t^2 + (\varphi_e^2 + \rho^2) \omega_t^2) \sigma_t^2}, \quad (\text{A.7})$$

$$K_{2,t+1} = \frac{(\varphi_e^2 \sigma_t^2 + (\varphi_e^2 + \rho^2) \omega_t^2) \sigma_t^2}{(\omega_t^2 + \sigma_t^2) V_{t+1} + (\varphi_e^2 \sigma_t^2 + (\varphi_e^2 + \rho^2) \omega_t^2) \sigma_t^2}, \quad (\text{A.8})$$

while the evolution of the variance of the filtering error is given by

$$\omega_{t+1}^2 = V_{t+1} K_{2,t+1}. \quad (\text{A.9})$$

Notably, for the economically relevant parameter values, signals are much more informative about the expected growth than consumption data, the Kalman Filter weight on consumption news $K_{1,t+1}$ is much smaller than that on the average signal $K_{2,t+1}$. Hence, to simplify the exposition, we solve the model when the agent learns only from the average signals. In the case when investors do not look at consumption data and only update based on the average signal, $K_{1,t+1} = 0$ and $K_{2,t+1}$ simplifies to

$$K_{2,t+1} = \frac{\rho^2 \omega_t^2 + \varphi_e^2 \sigma_t^2}{\rho^2 \omega_t^2 + \varphi_e^2 \sigma_t^2 + V_{t+1}}. \quad (\text{A.10})$$

Further, the evolution of the estimate of expected growth can be rewritten in the following way:

$$\hat{x}_{t+1} = (1 - K_{2,t+1})\rho\hat{x}_t + K_{2,t+1}\bar{x}_{t+1}. \quad (\text{A.11})$$

In the approximate solution of the model that we feature in the paper, the weight $K_{2,t+1}$ is constant and equal to steady-state value. To solve for the steady state of the system, we plug the solution for filtering uncertainty in $w_t^2 = K_2 V_t$ into the above equation and solve a quadratic equation for the constant value of K_2 when the volatility processes V_t and σ_t^2 are set to their unconditional means.

A.2 Asset Prices

For completeness, we present a solution for the equilibrium asset prices in the general case with optimal learning.

In the general case, $K_{1,t+1}$ and $K_{2,t+1}$ are time-varying, and their solutions are non-linear functions of the consumption volatility σ_t^2 and confidence measure V_t (see expressions (A.7) and (A.8)).

We conjecture that log price-to-consumption ratio pc_t is linear in the expected growth state \hat{x}_t , and capture its non-linear dependence on the confidence measure, filtering uncertainty and consumption volatility states by the function $f(V_t, \omega_t^2, \sigma_t^2)$:

$$pc_t = B_x \hat{x}_t + f(V_t, \omega_t^2, \sigma_t^2). \quad (\text{A.12})$$

Using the log-linearization of returns (see Campbell and Shiller, 1988), we can write down the log-linearized return on consumption asset in the following way:

$$\begin{aligned} r_{c,t+1} &= \kappa_0 + \kappa_1 pc_{t+1} - pc_t + \Delta c_{t+1} \\ &= \kappa_0 + \mu + (B_x(\kappa_1 \rho - 1) + 1)\hat{x}_t - f_t + \kappa_1 f_{t+1} + (\kappa_1 B_x K_{1,t+1} + 1)a_{c,t+1} + \kappa_1 B_x K_{2,t+1} a_{x,t+1}. \end{aligned} \quad (\text{A.13})$$

Using Euler equation (2.4), we can directly solve for the loading B_x :

$$B_x = \frac{1 - \frac{1}{\psi}}{1 - \kappa_1 \rho}, \quad (\text{A.14})$$

while the volatility component f_t satisfies the recursive equation

$$\begin{aligned} f_t = & \log \delta + \kappa_0 + \left(1 - \frac{1}{\psi}\right)\mu \\ & + \frac{1}{\theta} \log E_t \exp \left(\theta \kappa_1 f_{t+1} + \frac{1}{2} \theta^2 \begin{bmatrix} \kappa_1 B_x K_{1,t+1} + 1 - \frac{1}{\psi} \\ \kappa_1 B_x K_{2,t+1} \end{bmatrix}' \Sigma_{cx,t+1} \begin{bmatrix} \kappa_1 B_x K_{1,t+1} + 1 - \frac{1}{\psi} \\ \kappa_1 B_x K_{2,t+1} \end{bmatrix} \right), \end{aligned} \quad (\text{A.15})$$

where

$$\Sigma_{cx,t+1} = \text{Var} \left(\begin{bmatrix} a_{c,t+1} \\ a_{x,t+1} \end{bmatrix} \right) = \begin{bmatrix} \omega_t^2 + \sigma_t^2 & \rho \omega_t^2 \\ \rho \omega_t^2 & \rho^2 \omega_t^2 + \varphi_e^2 \sigma_t^2 + V_{t+1} \end{bmatrix}. \quad (\text{A.16})$$

From here it follows that the log discount factor is given by,

$$\begin{aligned} m_{t+1} = & \theta \log \delta - (1 - \theta) \kappa_0 - \gamma \mu - \frac{1}{\psi} \hat{x}_t + (1 - \theta) f_t \\ & - (\gamma + (1 - \theta) \kappa_1 B_x K_{1,t+1}) a_{c,t+1} - (1 - \theta) \kappa_1 B_x K_{2,t+1} a_{x,t+1} - (1 - \theta) \kappa_1 f_{t+1}. \end{aligned} \quad (\text{A.17})$$

Using the discount factor above, we obtain the interest rates and equity prices in the economy. These solutions are non-linear functions of the volatility states, which have to be solved using the numerical methods.

B Preferred Model Specification

B.1 Discount Factor

The aggregate consumption volatility σ_t^2 follows a square-root process specified in (2.7), while the dynamics of confidence measure is given by a discrete-time jump-diffusion specification outlined in (2.14). The distribution of jump size $J_{i,t+1}$ is defined by its moment generating function,

$$l(y) \equiv E e^{y J_i}. \quad (\text{B.1})$$

For example, when jump size follows exponential distribution with mean jump μ_j ,

$$l(y) = (1 - \mu_j y)^{-1}. \quad (\text{B.2})$$

The log price-to-consumption ratio pc_t is linear in the states of the economy:

$$pc_t = B_0 + B_x \hat{x}_t + B_v V_t + B_\sigma \sigma_t^2. \quad (\text{B.3})$$

The solution for the loading B_x is given in (A.14). The loading on the confidence measure B_v satisfies non-linear equation

$$\frac{1}{2} \theta \kappa_1^2 \sigma_w^2 z^2 - (1 - \kappa_1 (\nu - \lambda_1 \mu_j)) z + \frac{1}{2} \theta B_x^2 K_2 ((1 - (1 - K_2) \kappa_1 \rho)^2 + \kappa_1 K_2) + \frac{\lambda_1}{\theta} (l(\theta \kappa_1 z) - 1) = 0, \quad (\text{B.4})$$

for $z = B_v + \frac{1}{2} \theta \kappa_1 B_x^2 K_2^2$, while B_σ solves a quadratic equation

$$\frac{1}{2} \theta \kappa_1^2 \varphi_w^2 B_\sigma^2 - (1 - \nu_c \kappa_1) B_\sigma + \frac{1}{2} \theta \left(\left(1 - \frac{1}{\psi}\right)^2 + \kappa_1^2 B_x^2 K_2^2 \varphi_e^2 \right) = 0. \quad (\text{B.5})$$

Finally, the log-linearization parameter, which is pinned down by the equilibrium level of the price-consumption ratio, satisfies the following non-linear equation:

$$\begin{aligned} \log \kappa_1 = & \log \delta + \left(1 - \frac{1}{\psi}\right) \mu + B_\sigma (1 - \kappa_1 \nu_c) \sigma^2 \\ & + (B_v (1 - \kappa_1) + \kappa_1 (1 - \nu) z) \sigma_v^2 + \frac{\lambda_0}{\theta} (l(\theta \kappa_1 z) - \theta \kappa_1 z \mu_j - 1). \end{aligned} \quad (\text{B.6})$$

As in Eraker and Shaliastovich (2008), in case of multiple roots for B_σ and B_v we choose the solution which is non-explosive as the variation in V_t or σ_t^2 is approaching zero.

Using the equilibrium solution to the price-consumption ratio, we can write down the expression for the discount factor in the following way:

$$\begin{aligned} m_{t+1} = & m_0 + m_x x_t + m_v V_t + m_\sigma \sigma_t^2 \\ & - \lambda_c a_{c,t+1} - \lambda_x K_2 a_{x,t+1} - \lambda_v \left(\sigma_w \sqrt{V_t} w_{t+1} + Q_{t+1} \right) - \lambda_\sigma \varphi_w \sigma_t w_{c,t+1}, \end{aligned} \quad (\text{B.7})$$

where the discount factor loadings and the prices of risks are pinned down by the dynamics of factors and preference parameters of the investors. Their solutions are given by,

$$\begin{aligned} m_x = & -\frac{1}{\psi}, \quad m_v = (1 - \theta) B_v (1 - \kappa_1 \nu), \quad m_\sigma = (1 - \theta) B_\sigma (1 - \kappa_1 \nu_c), \\ m_0 = & \theta \log \delta + (1 - \theta) \log \kappa_1 - \gamma \mu - m_v \sigma_v^2 - m_\sigma \sigma^2, \end{aligned} \quad (\text{B.8})$$

and

$$\lambda_x = (1 - \theta) \kappa_1 B_x, \quad \lambda_\sigma = (1 - \theta) \kappa_1 B_\sigma, \quad \lambda_v = (1 - \theta) \kappa_1 B_v. \quad (\text{B.9})$$

B.2 Asset Prices

Consider a log payoff tomorrow expressed as,

$$p_{n-1,t+1} = F_{0,n-1} + F_{x,n-1}\hat{x}_{t+1} + F_{v,n-1}V_{t+1} + F_{\sigma,n-1}\sigma_{t+1}^2 + F_{g,n-1}\Delta c_{t+1} + F_{d,n-1}\sigma_t\eta_{d,t+1}. \quad (\text{B.10})$$

Then, the solution for the coefficients in its log price today $p_{n,t}$ satisfies

$$\begin{aligned} F_{g,n} &= F_{d,n} = 0, \\ F_{x,n} &= m_x + F_{x,n-1}\rho + F_{g,n-1}, \\ F_{\sigma,n} &= m_\sigma + F_{\sigma,n-1}\nu_c + \frac{1}{2} \left((F_{g,n-1} - \lambda_c)^2 + \varphi_e^2(F_{x,n-1} - \lambda_x)^2 K_2^2 + \varphi_w^2(F_{\sigma,n-1} - \lambda_\sigma)^2 + F_{d,n-1}^2 \right), \\ F_{v,n} &= m_v + \frac{1}{2} (F_{g,n-1} - \lambda_c + \rho(F_{x,n-1} - \lambda_x)K_2)^2 K_2 + (q_{vx} + \lambda_v)\nu + \frac{1}{2} q_{vx}^2 \sigma_w^2 + \lambda_1(l(q_{vx}) - q_{vx}\mu_j - 1), \\ F_{0,n} &= m_0 + F_{0,n-1} + F_{g,n-1}\mu + F_{\sigma,n-1}\sigma^2(1 - \nu_c) + (q_{vx} + \lambda_v)\sigma_v^2(1 - \nu) + \lambda_0(l(q_{vx}) - q_{vx}\mu_j - 1) \end{aligned} \quad (\text{B.11})$$

for $q_{vx} = F_{v,n-1} - \lambda_v + \frac{1}{2}(F_{x,n-1} - \lambda_x)^2 K_2^2$.

Setting $F_{0,n-1} = F_{x,n-1} = F_{v,n-1} = F_{\sigma,n-1} = F_{g,n-1} = F_{d,n-1} = 0$ in the above recursion, we can obtain the solution to n -period real risk-free rate.

On the other hand, the price-dividend ratio is given by,

$$pd_t = H_0 + H_x\hat{x}_t + H_vV_t + H_\sigma\sigma_t^2, \quad (\text{B.12})$$

where the loadings satisfy the following equations:

$$\begin{aligned} H_x &= m_x + \kappa_{d,1}\rho H_x + \phi, \\ H_\sigma &= m_\sigma + \kappa_{d,1}H_\sigma\nu_c + \frac{1}{2} \left((\phi - \lambda_c)^2 + \varphi_e^2(\kappa_{d,1}H_x - \lambda_x)^2 K_2^2 + \varphi_w^2(\kappa_{d,1}H_\sigma - \lambda_\sigma)^2 + \varphi_d^2 \right), \\ H_v &= m_v + \frac{1}{2} (\phi - \lambda_c + \rho(\kappa_{d,1}H_x - \lambda_x)K_2)^2 K_2 + (q_{vx} + \lambda_v)\nu + \frac{1}{2} q_{vx}^2 \sigma_w^2 + \lambda_1(l(q_{vx}) - q_{vx}\mu_j - 1), \end{aligned} \quad (\text{B.13})$$

for $q_{vx} = \kappa_{d,1}H_v - \lambda_v + \frac{1}{2}(\kappa_{d,1}H_x - \lambda_x)^2 K_2^2$, and the log-linearization parameter

$$\begin{aligned} \log \kappa_{d,1} &= m_0 + \mu_d + \left(H_v(1 - \kappa_{d,1}\nu) + \frac{1}{2}(\kappa_{d,1}H_x - \lambda_x)^2 K_2^2(1 - \nu) \right) \sigma_v^2 \\ &\quad + H_\sigma(1 - \kappa_{d,1}\nu_c)\sigma^2 + \lambda_0(l(q_{vx}) - q_{vx}\mu_j - 1). \end{aligned} \quad (\text{B.14})$$

Tables and Figures

Table 1: **Return Summary Statistics**

	Mean	Vol	AR(1)	Skew	Kurtosis
Market Return	6.51 (2.27)	18.81 (1.60)	0.10 (0.05)	-0.44 (0.50)	9.68 (1.24)
Risk-free Rate	0.57 (0.42)	1.13 (0.14)	0.98 (0.01)	-0.80 (0.56)	6.91 (1.32)

Summary statistics for market return and the risk-free rate. Monthly data from 1927 to 2007. Standard errors are Newey-West adjusted with 12 lags. Mean and volatility are annualized, in percent.

Table 2: **Large Return Move Evidence**

	Negative 2std	Positive 2std
Number of observations	38	16
Average Return	-139.30	148.91

Number of large return moves, and the average return level in those periods. Large return moves correspond to negative and positive 2 standard deviations or above innovations in the series, calculated based on AR(1)-GARCH(1,1) fit.

Table 3: Large Move Correlations in Returns and Macro Series

Correlations:	Return Jump Ind., Factor Jump Ind.			Return, Factor Jump Ind.			Return Jump Ind., Factor		
	0m	6m	12m	0m	6m	12m	0m	6m	12m
<i>Monthly:</i>									
Consumption	-0.02 (0.03)	-0.01 (0.04)	0.02 (0.05)	0.00 (0.03)	-0.01 (0.03)	0.02 (0.04)	-0.12 (0.04)	-0.09 (0.03)	-0.01 (0.03)
<i>Quarterly:</i>									
Consumption	-0.03 (0.02)	-0.03 (0.03)	0.09 (0.12)	0.02 (0.04)	-0.04 (0.04)	-0.08 (0.08)	-0.08 (0.04)	0.08 (0.08)	-0.04 (0.07)
Confidence Measure	0.34 (0.17)	-0.04 (0.01)	-0.04 (0.01)	-0.32 (0.10)	-0.01 (0.04)	-0.05 (0.07)	0.11 (0.07)	0.04 (0.07)	0.06 (0.05)

Correlations of return moves with current and future moves in macroeconomic variables. The left panel shows the correlations of large return move indicator with current and future jump indicators in consumption and confidence measure 6 and 12 months ahead. The middle panel shows the correlations of the level of returns with current and future large move indicators in macroeconomic series, while the right panel depicts the correlations of the large return move indicator with current and future consumption and confidence measure 6 and 12 months ahead. Data on confidence measures are based on forecasts of real GDP from 1968 to 2007, on monthly consumption from 1959 to 2007, and on quarterly consumption from 1947 to 2007. Jump indicators correspond to 2 standard deviation or above move in a series, based on AR(1)-GARCH(1,1) fit.

Table 4: **Predictability of Asset Valuation and Return Variance**

	Confidence	Cond. Variance	R^2
pd	-5.31* [-2.38]	-0.13* [-5.16]	0.53
<i>Implied Variance:</i>			
projection	5527.45 [1.18]	243.50* [2.24]	0.16
25th quantile	1887.69 [1.42]	345.07* [10.19]	0.21
50th quantile	9034.78* [3.75]	337.20* [8.19]	0.24
75th quantile	18292.80* [3.07]	192.74 [1.22]	0.11

Regression of price-dividend ratio and the variance of returns on confidence measure and conditional variance of real GDP. Quarterly data on price-dividend ratio, returns and confidence measure are from 1968 to 2007, and implied volatility, based on the VIX index squared, is from 1990 to 2007. All series annualized in percent. T-statistics are in square brackets, and star superscript refers to significance at 1% level.

Table 5: Estimation of Confidence Measure Dynamics

	σ_v	ν	σ_w	μ_j	λ_0	λ_1	$p - value$
<i>Real GDP:</i>							
Square-Root	0.12 (0.03)	0.83 (0.10)	0.11 (0.01)				
Square-Root+Jump	0.12 (0.02)	0.73 (0.10)	0.03 (0.004)	0.01 (0.002)	0.18 (0.13)	45.73 (16.81)	0.00
<i>Industrial Production</i>							
Square-Root	0.18 (0.02)	0.66 (0.10)	0.16 (0.01)				
Square-Root +Jump	0.18 (0.02)	0.62 (0.09)	0.04 (0.005)	0.02 (0.004)	0.22 (0.12)	14.48 (5.68)	0.00

Estimation results for discrete-time jump diffusion model for the confidence measures. Square-root specification has only Normal innovations, while Square-root+Jump model features Gaussian shocks and Poisson jumps with time-varying arrival intensity and exponentially distributed jump size. Quarterly data on confidence measures based on real GDP and industrial production forecasts from 1969 to 2007, annualized in percent. P-value is computed for the Likelihood Ratio test that jump parameters μ_j, λ_0 and λ_1 are jointly equal to zero.

Table 6: **Model Parameter Calibration**

Parameter	Value
<i>Preference Parameters:</i>	
δ	0.9992
γ	10
ψ	1.5
<i>Consumption Dynamics:</i>	
μ	0.0017
ρ	0.975
σ	0.0064
ν_c	0.995
φ_w	5.19e-04
φ_e	0.038
ϕ	2.75
φ_d	3
<i>Confidence Dynamics:</i>	
σ_v	4.33e-04
ν	0.91
σ_w	0
μ_j	$3.59e - 07$
λ_0	0.18/3
λ_1	20×40000

Calibrated parameter values, monthly frequency.

Table 7: **Confidence Measure: Data and Model Calibration**

	Mean	Vol	AR(1)	Skew	Kurt
Real GDP	0.07 (0.02)	0.04 (0.02)	0.68 (0.07)	1.73 (0.39)	7.85 (2.04)
IP	0.06 (0.01)	0.04 (0.01)	0.60 (0.12)	1.19 (0.31)	4.06 (1.06)
Sim. Median	0.09	0.07	0.59	1.26	4.23
5%	0.05	0.03	0.39	0.80	2.77
95%	0.15	0.12	0.75	1.74	8.04

Summary statistics for square-root of confidence measure scaled by the average volatility of the underlying series. Data are based on quarterly observations of confidence measure based on forecasts of real GDP and industrial production from 1969 to 2007. Model estimates are calculated based on 100 simulations of 40 years of data. The square root of confidence measure scaled by the volatility of the consumption growth is sampled every third month.

Table 8: **Consumption Dynamics: Data and Model Calibration**

	Data		Model		
	Estimate	S.E.	5%	Median	95%
Mean	1.95	(0.32)	1.27	1.99	2.73
Vol	2.13	(0.52)	1.43	2.08	3.07
AR(1)	0.44	(0.13)	0.22	0.42	0.56
AR(2)	0.16	(0.18)	-0.13	0.14	0.38
AR(5)	-0.01	(0.10)	-0.21	0.03	0.23
VR(2)	1.58	(0.18)	1.22	1.42	1.56
VR(5)	2.23	(0.86)	1.28	1.95	2.72

Calibration of consumption dynamics. Data is annual real consumption growth for 1930-2006. Model is based on 100 simulations of 80 years of monthly consumption data aggregated to annual horizon, based on the full specification with fluctuating confidence and time-varying consumption volatility.

Table 9: **Model Output for Return Distribution**

	Mean	Vol	AR(1)	Skew	Kurt
<i>Complete Information, Time-Varying Vol:</i>					
Log return	6.31	14.02	-0.01	0.08	3.78
Risk-free rate	1.50	0.26	0.97	-0.38	3.16
<i>Constant confidence, Constant Vol:</i>					
Log return	4.22	10.95	0.01	0.00	3.00
Risk-free rate	1.71	0.23	0.97	0.00	3.00
<i>Time-varying confidence, Time-Varying Vol:</i>					
Log return	6.73	19.97	-0.06	-0.85	8.64
Risk-free rate	1.15	0.37	0.91	-1.35	6.82

Model-implied summary statistics for returns and risk-free rates. Based on 100 simulations of 80 years of data.

Table 10: **Model-Implied Large Move Correlations**

Correlations:	Return Jump Ind.,			Return,			Return Jump Ind.,		
	Factor Jump Ind.			Factor Jump Ind.			Factor		
	0m	6m	12m	0m	6m	12m	0m	6m	12m
Consumption	0.03 (0.10)	0.00 (0.08)	-0.01 (0.07)	0.01 (0.10)	0.00 (0.07)	0.00 (0.08)	0.00 (0.09)	0.00 (0.09)	0.00 (0.08)
Confidence Measure	0.39 (0.08)	0.00 (0.09)	0.00 (0.07)	-0.45 (0.07)	-0.01 (0.08)	0.00 (0.07)	0.38 (0.10)	0.15 (0.11)	0.02 (0.10)

Model-implied correlations of return moves with current and future moves in consumption and confidence measure. The left panel shows the correlations of large return move indicator with current and future jump indicators in consumption and confidence measure 6 and 12 months ahead. The middle panel shows the correlations of the level of returns with current and future large move indicators in macroeconomic series, while the right panel depicts the correlations of the large return move indicator with current and future consumption and confidence measure 6 and 12 months ahead. Statistics are calculated based on 100 simulations of 80 years of data.

Table 11: **Prices of Risk**

	Long-Run Growth	Short-Run Growth	Confidence Measure	Consumption Volatility
Complete Information	8.81	6.35		-9.06
Fluctuating Confidence	8.52	6.36	-6.06	-4.16

Prices of risk for 1 standard deviation shock, in per cent, in standard long-run risks model with complete information, and in the model with learning and fluctuating confidence risks.

Table 12: **Equity Premium Decomposition**

	Long-Run Growth	Short-Run Growth	Confidence Measure	Consumption Volatility	Total
<i>Model:</i>					
Complete Information Time-varying Vol.	1.67	1.15		2.30	5.11
Constant Confidence Constant Vol.	1.69	1.19			2.87
Time-varying Confidence Time-varying Vol.	1.66	1.19	1.69	0.75	5.25
<i>Data:</i>					
Based on Full model	1.94	0.96	1.96	0.74	5.56

Decomposition of equity risk premium into contributions from long-run, short-run, consumption volatility and confidence risks. Equity premium is annualized, in percent. Model-based decomposition in the data is based on calibrated parameter values and observed series of confidence measure and real GDP volatility from 1968 to 2007, quarterly.

Table 13: **Predictability of Excess Returns**

	1y		3y		5y	
	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
<i>Data:</i>						
Slope	-0.08	(0.05)	-0.25	(0.09)	-0.41	(0.11)
R^2	0.03	(0.04)	0.13	(0.07)	0.26	(0.07)
<i>Constant confidence,</i>						
<i>Constant Vol:</i>						
Sim. Slope	-0.03	(0.11)	-0.10	(0.30)	-0.19	(0.43)
R^2	0.01	(0.02)	0.01	(0.03)	0.02	(0.05)
<i>Time-varying confidence,</i>						
<i>Time-Varying Vol:</i>						
Sim. Slope	-0.39	(0.17)	-0.58	(0.26)	-0.68	(0.31)
R^2	0.14	(0.09)	0.16	(0.12)	0.16	(0.12)


Projection of future excess returns on price-dividend ratio in the data (first panel) and in the models with constant confidence and constant consumption volatility, and fluctuating confidence and time-varying consumption volatility (second panel). Monthly observations of equity returns, price-dividend ratios and risk-free rates from 1927 to 2007. Simulated slopes and R^2 together with 5% – 95% confidence band are calculated based on 100 simulations of 80 years of data.

Table 14: **Predictability of Consumption Growth**

	1y		3y		5y	
	Estimate	S.E.	Estimate	S.E.	Estimate	S.E.
<i>Data:</i>						
Slope	0.01	(0.01)	0.01	(0.02)	0.003	(0.01)
R^2	0.08	(0.06)	0.03	(0.02)	0.001	(0.03)
<i>Constant confidence,</i>						
<i>Constant Vol:</i>						
Sim. Slope	0.13	(0.02)	0.25	(0.06)	0.32	(0.10)
R^2	0.43	(0.10)	0.30	(0.12)	0.23	(0.12)
<i>Time-varying confidence,</i>						
<i>Time-Varying Vol:</i>						
Sim. Slope	0.05	(0.02)	0.08	(0.05)	0.10	(0.08)
R^2	0.14	(0.09)	0.09	(0.10)	0.07	(0.10)


Projection of future consumption growth on price-dividend ratio in the data (first panel) and in the models with constant confidence and constant consumption volatility, and fluctuating confidence and time-varying consumption volatility (second panel). Table reports slope coefficient β_1 and R^2 in the regressions $\sum_{j=1}^K \Delta c_{t+j} = \beta_0 + \beta_1 pd_t + error$, where K is from 1 to 5 years. Annual observations of real consumption growth and price-dividend ratios from 1930 to 2007. Simulated slopes and R^2 are calculated based on 100 simulations of 80 years of data.

Figure 1: **Time-Series of Returns**



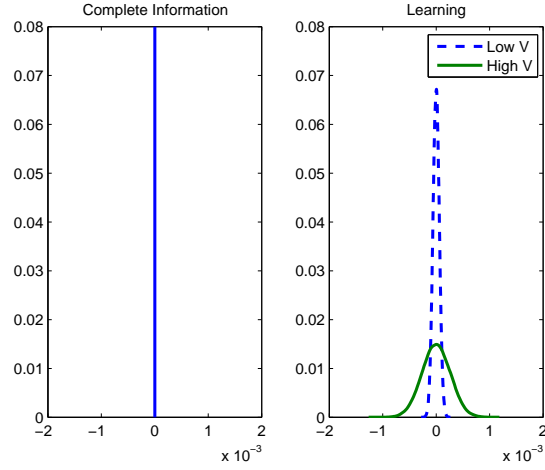
Time series of returns, monthly, from 1927 to 2007. Stars indicate periods with 2 standard deviations or above moves in returns.

Figure 2: **Time Series of Confidence Measure**




Square-root of confidence measure based on forecasts of next-quarter real GDP, annualized, in percent. Quarterly observations from the Survey of Professional Forecasts from 1969 to 2007.

Figure 3: **Distribution of True Expected Growth**



Distribution of true expected growth state, given information set of investors, in a standard long-run risks model with complete information (left panel), and model with learning and fluctuating confidence risks (right panel). Consumption volatility is constant, and Low and High V correspond to 0.25 and 0.75 quantile of calibrated distribution of confidence measure.

Figure 4: **Unconditional Distribution of Consumption and Returns**



Unconditional distribution of consumption growth (left panel) and returns in the model with constant confidence measure (middle panel) and in the model with fluctuating confidence risks (right panel). Consumption volatility is constant.

References

- Ai, Hengjie, 2007, Information about long-run risks: Asset pricing implications, forthcoming in *Journal of Finance*.
- Andersen, T., L. Benzoni, and J. Lund, 2002, An empirical investigation of continuous-time equity return models, *Journal of Finance* 57, 1239–1284.
- Anderson, Evan W., Eric Ghysels, and Jennifer L. Juergens, 2007, The impact of risk and uncertainty on expected returns, forthcoming in *Journal of Financial Economics*.
- Bansal, Ravi, Dana Kiku, and Amir Yaron, 2007a, A note on the economics and statistics of predictability: A long run risks perspective, working paper.
- Bansal, Ravi, Dana Kiku, and Amir Yaron, 2007b, Risks for the long run: Estimation and inference, Working paper.
- Bansal, Ravi, and Ivan Shaliastovich, 2009, Learning and asset-price jumps, NBER Working paper #14814.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.
- Barberis, Nicholas, and Richard Thaler, 2003, A survey of behavioral finance, in G. Constantinides, R. Stulz, and M. Harris, (eds.) *Handbook of the Economics of Finance* (North Holland, Amsterdam).
- Barro, Robert, 2006, Rare disasters and asset markets in the Twentieth century, *Quarterly Journal of Economics* 121.
- Brennan, M., 1998, The role of learning in dynamic portfolio decisions, *European Finance Review* 1, 295–306.
- Brennan, M.J., and Y.H. Xia, 2001, Stock price volatility and equity premium, *Journal of Monetary Economics* 47, 249–283.
- Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *Journal of Finance* 62, 1453–1490.
- Buraschi, A., and A. Jitsov, 2006, Model uncertainty and option markets with heterogeneous agents, *Journal of Finance* 61, 2841–2898.
- Campbell, John, and Robert Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *The Review of Financial Studies* 1, 195–228.

- Croce, Mariano, Martin Lettau, and Sydney Ludvigson, 2006, Investor information, long-run risk, and the duration of risky cash flows, working paper.
- Croushore, Dean, 1993, Introducing: The survey of professional forecasters, *Federal Reserve Bank of Philadelphia Business Review* Nov/Dec, 3–13.
- David, Alexander, 1997, Fluctuating confidence in stock markets: Implications for returns and volatility, *Journal of Financial and Quantitative Analysis* 32, 427–462.
- David, Alexander, and Pietro Veronesi, 2008, Inflation and earnings uncertainty and volatility forecasts: A structural form approach, working paper.
- De Bondt, Werner, and Richard Thaler, 1990, Do security analysts overreact?, *Stock Market Volatility* 80, 52–57.
- Detemple, J., 1986, Asset pricing in a production economy with incomplete information, *Journal of Finance* 41, 383–391.
- Epstein, Larry G., and Stanley Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.
- Eraker, Bjorn, M. Johannes, and N.G. Polson, 2003, The impact of jumps in returns and volatility, *Journal of Finance* 53, 1269–1300.
- Eraker, Bjørn, and Ivan Shaliastovich, 2008, An equilibrium guide to designing affine pricing models, *Mathematical Finance* 18, 519–543.
- Gabaix, Xavier, 2007, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, working paper.
- Genotte, G., 1986, Optimal portfolio choice under incomplete information, *Journal of Finance* 41, 733–746.
- Hansen, Lars, and Thomas Sargent, 2006, Fragile beliefs and the price of model uncertainty, working paper.
- Hansen, Lars Peter, John C. Heaton, and Nan Li, 2008, Consumption strikes back? Measuring long-run risk, *Journal of Political Economy*, forthcoming.
- Hogarth, Robin, and Hillel Einhorn, 1992, Order effects in belief updating: The belief-adjustment model, *Cognitive Psychology* 1, 1–55.
- Kahneman, Daniel, and Amos Tversky, 1973, On the psychology of prediction, *Psychological Review* 80, 237–251.

- Kreps, David M., and Evan L. Porteus, 1978, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica* 46, 185–200.
- Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.
- Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.
- Rietz, Thomas, 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.
- Singleton, Kenneth J., 2006, *Empirical Dynamic Asset Pricing - Model Specification and Econometric Assessment* (Princeton University Press).
- Van Nieuwerburgh, Stijn, and Laura Veldkamp, 2006, Learning asymmetries in real business cycles, *Journal of Monetary Economics* 53, 753–772.
- Veldkamp, Laura, 2006, Slow boom, sudden crash, *Journal of Economic Theory* 124, 230–257.
- Veronesi, Pietro, 1999, Stock market overreaction to bad news in good times: A rational expectations equilibrium model, *Review of Financial Studies* 12, 975–1007.
- Veronesi, Pietro, 2000, How does information quality affect stock returns?, *Journal of Finance* 55, 807–837.
- Zarnowitz, Victor, and P. Braun, 1993, Twenty-two years of the NBER-ASA quarterly economic outlook surveys: Aspects and comparisons of forecasting performance, in J.H. Stock and M.W. Watson, (eds.) *Business Cycles, Indicators, and Forecasting*, volume 28 of *NBER Studies in Business Cycles*, 11–84 (Chicago: University of Chicago Press).