

THE TIME VARIATION IN RISK APPETITE AND UNCERTAINTY *

Geert Bekaert[†] Eric C. Engstrom[‡] Nancy R. Xu[§]

February 10, 2021

Abstract

We formulate a dynamic no-arbitrage asset pricing model for equities and corporate bonds, featuring time variation in both risk aversion and economic uncertainty. The joint dynamics among cash flows, macroeconomic fundamentals and risk aversion accommodate both heteroskedasticity and non-Gaussianity. The model delivers measures of risk aversion and uncertainty at the daily frequency. We verify that equity variance risk premiums are very informative about risk aversion, whereas credit spreads and corporate bond volatility are highly correlated with economic uncertainty. Our model-implied risk premiums outperform standard instruments for predicting asset excess returns. Risk aversion is substantially correlated with consumer confidence measures, and in early 2020 reacted more strongly to new Covid cases than did an uncertainty proxy.

JEL Classification: C1, G10, G12, G13.

Keywords: Risk aversion, Economic uncertainty, Dynamic asset pricing model, VIX, Variance risk premium, Sentiment, Covid crisis.

*The views expressed in this document do not necessarily reflect those of the Federal Reserve System, its Board of Governors, or staff. All errors are the sole responsibility of the authors. We would like to thank Haoxiang Zhu (the editor), the Associate Editor, three anonymous referees, Gurdip Bakshi (discussant), Ing-Haw Cheng (discussant), Anna Cieslak (discussant), Drew Creal (discussant), Thomas Dangl (discussant), Rob Engle, Bob Hodrick, Mark Kamstra, Lisa Kramers, Thomas Maurer (discussant), Bradley Paye (discussant), Helene Rey, Nick Roussanov, Gill Segal (discussant), and Paul Tetlock for many insightful discussions and seminar/conference participants at Columbia Business School, Baruch College, London Business School, University of Zurich, Luxembourg School of Finance, Chicago Fed, 2019 EFA, 2019 CICF, 2019 EFMA, 2019 FIRS, 2019 MFA, 2019 European Winter Finance Summit, 2019 AFA, 31st Australasian Finance and Banking Conference, 2018 China International Risk Forum, 2018 Northern Finance Association Annual Meeting, “Machine Learning and Finance: The New Empirical Asset Pricing” hosted by University of Chicago Booth, 2018 North American Summer Meeting of the Econometric Society, 11th Annual SoFiE Conference, Federal Reserve Board’s Conference on Risk, Uncertainty, and Volatility, and Columbia Women in Economics. Corresponding author: Nancy Xu (nancy.xu@bc.edu; +1(617)552-2713), Boston College, Carroll School of Management, 140 Commonwealth Avenue, Chestnut Hill, MA, 02467, USA. Indices and Online Appendices are available at <https://www.nancyxu.net>.

[†]Finance Division, Columbia Business School, 3022 Broadway, New York, NY 10027, U.S.A. Email: gb241@gsb.columbia.edu.

[‡]Research and Statistics Division, Federal Reserve Board, 20th Street and Constitution Avenue N.W., Washington, DC 20551, U.S.A. Email: eric.c.engstrom@frb.gov.

[§]Corresponding author. Finance Department, Carroll School of Management, Boston College, 140 Commonwealth Avenue, Chestnut Hill, MA 02467, U.S.A. Email: nancy.xu@bc.edu.

1 Introduction

Many economic models combine assumptions regarding the preferences of economic agents with assumptions regarding the data generating process for consumption growth or productivity shocks to derive implications for financial asset prices. A large class of models (see e.g. Bansal, Kiku, Shaliastovich, and Yaron (2014)) relies on time variation in economic uncertainty as the main mechanism to generate variation over time in financial risk premiums, while assuming the risk aversion of households to be time invariant. Another class of models featuring habit-forming utility, starting with Campbell and Cochrane (1999), stresses time-varying risk aversion as the main driver of financial market risk premiums.

In this article, we separately identify time-varying uncertainty in fundamentals, using macro data, and time-varying aggregate risk aversion (or its inverse, which we call “risk appetite”), using both macro data and financial asset prices, through the lens of a dynamic asset pricing model.¹ To do so, we build on the habit models of Campbell and Cochrane (1999), Menzly, Santos, and Veronesi (2004) and Wachter (2006), but in contrast to those models we allow stochastic risk aversion to have a component that is uncorrelated with fundamentals. The non-fundamental component may reflect economic news that is imperfectly correlated with realized measures of aggregate activity, or consumer sentiment regarding the economy, a hypothesis we formally test. However, it may also reflect pure mood swings (e.g., Kamstra, Kramer, and Levi (2003)’s weather-induced swings), or unmodeled institutional factors, such as risk constraints faced by financial institutions (e.g., He and Krishnamurthy (2013); Adrian and Shin (2013)), that end up affecting aggregate risk aversion.

To develop the risk aversion measure in an internally consistent manner, we must solve for asset prices as a function of preferences, consumption growth and cash flow dynamics. We employ two prominent risky asset classes, corporate bonds and equities. To give the macroeconomic and cash flow-based fundamentals a maximal chance of fitting asset price dynamics, we use, inter alia, monthly data on industrial production, which is helpful in identifying cyclical variation, and model fundamental and cash flow shocks (earnings for equities, loss rates for corporate bonds) using non-Gaussian distributions with time-varying second and higher order moments. Concretely, we use the Bad Environment-Good Environment (BEGE, henceforth) framework developed in Bekaert and Engstrom (2017), where shocks to key state variables are modeled as the sum of two centered gamma distributions with time-varying shape parameters. These shape parameters drive changes in “bad” (“good”) volatility, associated with negative (positive) skewness, respectively. Despite the fact that the model accommodates state variables with time-varying non-Gaussian shocks, our formulation admits (quasi) closed-form solutions for asset prices within the affine class. Our modeling framework is quite different from the model in Bekaert and Engstrom (2017), which appends a simplified BEGE model for consumption growth to the non-linear price of risk model of Campbell and Cochrane (1999), preventing closed-form equilibrium solutions. Our modeling of macroeconomic uncertainty also delivers an uncertainty index as a by-product, contributing to a recent cottage industry for developing

¹Pflueger, Siriwardane, and Sunderam (2020) aim to identify the product of “uncertainty” and “risk aversion,” which they term risk appetite, and show that it is an important determinant of real rate variation.

indices of macroeconomic uncertainty (e.g. Jurado, Ludvigson, and Ng (2015)).

To identify the model parameters and stochastic risk aversion, we go beyond using only information in historical realized returns, which are known to be noisy. In particular, we use both realized variances and option-implied variances in the estimation of the model parameters. A large empirical literature (see e.g. Andersen, Bollerslev, Diebold, and Labys (2003)) shows that realized variances can be measured fairly precisely and provide accurate forecasts of future return variances. Moreover, conditional return variances are an exact function of the relevant state variables (including risk aversion) in our pricing framework (see Joslin, Le, and Singleton (2013) for a similar observation in a term structure model). There is also a large literature on inferring risk and risk preferences from option prices, which we discuss in more detail in Section 2.² Option-implied volatility, such as the famous VIX index in the equity market, reflects both the physical return distribution, including the probability of crashes, and risk aversion. The risk aversion of agents creates a demand for insurance against potential losses, making (out-of-money) put options relatively more expensive than call options. Such expensive put options are the source of the consistent presence of a positive variance risk premium (often empirically measured as the difference between the VIX index-squared and the physical conditional return variance) (see Bekaert and Hoerova (2016); Bakshi and Madan (2006) for formal arguments). Option data should also be informative about conditional risk premiums, which are difficult to observe from the data. Martin (2017) uses option-implied variances to provide bounds on equity premiums, and several articles (see Bollerslev and Todorov (2011); Liu, Pan, and Wang (2004); Santa-Clara and Yan (2010)) suggest that compensation for rare events (“jumps”) accounts for a large fraction of equity risk premiums.

An important output of our model and contribution of this work is a measure of time-varying aggregate risk aversion that consistently helps price assets in the context of our structural asset pricing model and is easily tracked over time, even at high frequencies. To accomplish this, we exploit the model implication that asset prices and variances are an exact function of the uncertainty and risk aversion state variables. While we filter the uncertainty state variables from macroeconomic data, we use a method of moments estimation for the preference parameters, which exploits the model implication that risk aversion is a linear function of a set of observable financial variables, such as credit spreads and equity risk-neutral variances. The measure should be a useful model-based complement to “sentiment indices” developed in the behavioral finance literature (see, e.g., Lemmon and Portniaguina (2006); and Baker and Wurgler (2006)), or practitioner indices developed by financial institutions (see Coudert and Gex (2008) for a survey). We hope that our measurement of risk aversion will be useful in other areas of economics as well. For example, in monetary economics, recent research suggest a potential link between loose monetary policy and greater risk appetite of market participants, spurring a literature that explores what structural economic factors drive risk aversion changes (see, e.g., Rajan (2006); Adrian and Shin (2009); Bekaert, Hoerova, and Lo Duca (2013)). In international finance, Miranda-Agrippino and Rey (2020) and Rey (2013) suggest that global risk aversion is

²A number of these articles develop time-varying risk aversion measures motivated by models that really assume a “constant” risk aversion coefficient and hence are inherently inconsistent (see, e.g., Bollerslev, Gibson, and Zhou (2011); Faccini, Konstantinidi, Skiadopoulos, and Sarantopoulou-Chiourea (2019)).

a key transmission vector by which US monetary policy is “exported” to foreign countries and is a major source of asset return comovements across countries (see also Xu (2019)).

Our main results are as follows: First, we find significant time variation in the volatilities and higher-order moments of the fundamentals, especially in real activity. The time variation in uncertainty is dominated by strongly countercyclical “bad” volatility. Moreover, we find that macroeconomic uncertainty is informative about uncertainty regarding risky asset cash flows, both for the equity and corporate bond markets. Nonetheless, the volatility of corporate bond loss rates shows independent time variation.

Second, the extracted risk aversion process loads most significantly on equity risk-neutral variances (with a positive sign) and realized variances (with a negative sign), consistent with the literature finding the variance premium a good proxy for aggregate risk aversion. This finding is consistent with recent work in the consumption-based asset pricing literature, showing the variance premium to be very informative for identifying equilibrium models featuring complex data generating processes for the fundamentals (see Drechsler and Yaron (2011); Bollerslev, Tauchen, and Zhou (2009); Bekaert and Engstrom (2017)). Nevertheless, corporate bond market variables – the credit spread and realized corporate bond variance – also account for almost 35 percent of the measured variation in risk aversion. Moreover, our measure of risk aversion sometimes deviates materially from the signal provided by the variance premium. In particular, the residual from a regression of risk aversion on the variance risk premium shows meaningful countercyclical variation. The risk aversion process is much more rapidly mean reverting than would be implied by habit models, which is consistent with the results in Martin (2017).

Third, economic uncertainty is highly correlated with corporate bond volatility and, especially, with credit spreads, suggesting that these financial measures are good predictors of macroeconomic turbulence. In addition, our economic uncertainty index predicts output negatively and significantly. Because equity risk premium variation is dominated by changes in risk aversion, but the conditional variance of equity returns also loads strongly on macroeconomic uncertainty, our results help explain the failure of a large literature in finance (starting with French, Schwert, and Stambaugh (1987)) to find a robust link between future equity returns and the conditional variance of equity returns, while assuming a constant price of variance risk. In addition, the model-implied equity premium is always above and very highly correlated with the lower bound provided in Martin (2017).

Fourth, to aid with the interpretation and validation of our risk aversion measure, we conduct several exercises. We present the correlation of the risk aversion measure with macroeconomic news data to verify its relation to alternative measures of real activity. Among 7 news announcements, we find industrial production news to be most important determinant but it still only accounts for a small part of the variation in risk aversion, consistent with our model findings. We also relate risk aversion to 16 alternative sentiment/confidence measures, most of which do not rely on asset prices. Even when those external measures are orthogonalized with respect to economic uncertainty, our risk aversion proxy is highly correlated with them and risk aversion is most correlated with measures focusing on consumer sentiment/confidence. The highest correlation occurs with the Sentix investor sentiment measure designed “to reflect investors’ emotions fluctuating between fear and greed.” In addition, we analyze the behavior

of risk aversion during the Covid crisis. We find that, controlling for economic news, our high frequency proxy to risk aversion reacts more to information regarding the volume of new cases of infection, than does our high frequency proxy to economic uncertainty.

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 presents the estimation results for the fundamentals and cash flow dynamics, and Section 4 for risk aversion. In Section 5, we investigate how our measures of risk aversion and macroeconomic uncertainty correlate with (and predict) macroeconomic activity and asset price changes, and examine their relation with extant indices. We also examine the link between risk aversion and various consumer and investor confidence measures, finishing with a study of the behavior of risk aversion and uncertainty during the Covid crisis. Concluding remarks are in Section 6.

2 Modeling Risk Appetite and Uncertainty

In this section, we first define our concept of risk aversion. We then build a dynamic model with stochastic risk aversion and macroeconomic factors affecting the cash flows processes of two main risky asset classes, corporate bonds and equity. The state variables are described in Section 2.2 and the pricing kernel in Section 2.3.

2.1 Risk Aversion

An ideal measure of risk aversion would be model-free and would not confound time variation in economic uncertainty with time variation in risk aversion. There are many attempts in the literature to approximate this ideal measure, but invariably various modeling and statistical assumptions are necessary to identify risk aversion. For example, in the options literature, a number of articles (Ait-Sahalia and Lo (2000); Rosenberg and Engle (2002); Jackwerth (2000); Bakshi, Kapadia, and Madan (2003); Britten-Jones and Neuberger (2000); Bliss and Panigirtzoglou (2004); Bakshi and Wu (2010); Faccini, Konstantinidi, Skiadopoulou, and Sarantopoulou-Chiourea (2019)) appear at first glance to infer risk aversion from equity options prices in a model-free fashion, but it is generally the case that the utility function is assumed to be of a particular form and/or to depend only on stock prices.³

Our approach is to start from a utility function defined over both consumption (“fundamentals”) and a potential “non-fundamental” factor. Our measure of risk aversion is then the coefficient of relative risk aversion implied by the utility function. We specify a consumption process accommodating time variation in economic uncertainty and use the utility framework to price all assets consistently, given general processes for their cash flows.

Consider a period utility function in the hyperbolic absolute risk aversion (HARA) class:

$$U\left(\frac{C}{Q}\right) = \frac{\left(\frac{C}{Q}\right)^\gamma}{1 - \gamma} \quad (1)$$

where γ is the curvature parameter, C is consumption and Q is a process that will be shown

³This is also true in the recent debate about the claim of recoverability of physical probabilities from option prices, which, if true, identifies risk aversion as well (Ross (2015); Carr and Wu (2016)).

to drive time-variation in risk aversion. Essentially, when Q is high, consumption delivers less utility and marginal utility increases. We assume:

$$Q = \frac{C}{C - H} = f(C) \quad (2)$$

where H is an exogenous reference level or process ($C > H$), for example the “external” habit stock as in Campbell and Cochrane (1999) (CC henceforth) or a subsistence level. Critically, the H process can vary through time but is “exogenous” to the agent’s optimization problem, as in the well-known “catching up with the Jones’s” preferences (see also Abel (1990)). This excludes internal habit models. Note that Q is a negative function of consumption. If Q were simply an exogenous process, risk aversion is equal to γ and does not vary over time (see Abel (1990) for such “multiplicative” habit models).

The coefficient of relative risk aversion for this class of models is given by

$$RRA = \frac{CU''(C)}{U'(C)} = \gamma Q \quad (3)$$

and is thus proportional to Q . We use the terms “risk aversion” and “risk appetite” as each other’s inverse.⁴

For pricing assets, it is helpful to derive the log pricing kernel which is the intertemporal marginal rate of substitution in a dynamic economy. We assume an infinitely lived agent, facing a constant discount factor of β , and the HARA period utility function in Equation (1). The log pricing kernel, m_{t+1} , is then given by

$$m_{t+1} = \ln(\beta) + \ln \left[\frac{U^\theta(C_{t+1})}{U^\theta(C_t)} \right] = \ln(\beta) - \gamma \Delta c_{t+1} + \gamma \Delta q_{t+1} \quad (4)$$

where we use t to indicate time, lower case letters to indicate logs of uppercase variables, and Δ to represent the difference operator. For all gross returns R^i , it is true that $E_t [\exp(m_{t+1})R^i_{t+1}] = 1$.

There are a variety of approaches to model Q . In the external habit model of CC, Q_t is the inverse of the surplus ratio. CC models q_t exogenously as a slow-moving, persistent process, but restrict the correlation between shocks to q_t and Δc_t to be *perfect*. That is, risk aversion is fully driven by consumption shocks. Importantly, there is no time variation in economic uncertainty in their model as the consumption growth process is homoskedastic. The “moody investor” economy in Bekaert, Engstrom, and Grenadier (2010) is also a special case. In that model, q_t is also exogenously modeled, but has its own shock; that is, there are preference shocks not correlated with fundamentals. Another special case is the model in Brandt and Wang (2003), in which the risk aversion process specifically depends on inflation in addition to consumption growth.

We specify a stochastic process for q (risk aversion), which is partly but not fully driven by macroeconomic fundamentals (consumption growth) and features an independent preference

⁴Gai and Vause (2006) and Pflueger, Siriwardane, and Sunderam (2020) however use the term “risk appetite” to indicate the price of risk, that is the product of “risk aversion” and “the amount of risk” (which would be the volatility of consumption growth in a power utility model).

shock. Shocks to risk aversion that are independent of macroeconomic fundamentals may arise in a variety of ways. The experimental literature (see e.g. Cohn, Engelmann, Fehr, and Maréchal (2015)) shows that the subjective willingness to take risk is indeed lower during a recession, which is simulated by “priming” people with a stock market crash (versus boom), and that this risk aversion is rooted in emotions of fear.⁵ Thus, bad economic news can increase risk aversion but it is unlikely that the aggregate component of this type of counter-cyclical risk aversion is perfectly correlated with “measured” aggregate consumption growth.

In addition, the wealth of richer people conceivably decreases proportionally more than that of poorer people in bad times (because more of their wealth is tied up in risky asset classes). Thus, changes in the wealth distribution across individuals may cause changes in aggregate wealth-weighted risk aversion and even induce counter-cyclical risk aversion. Of course, we cannot exclude risk aversion changes that are driven by other sources, ranging from reactions to political speeches to pure mood swings.

Finally, it goes without saying that our risk aversion process is identified within the context of a particular rational expectations model, and thus alternative interpretations of our results are possible. We provide more discussion and external validation in Section 5.3.

2.2 Economic Environment: State Variables

2.2.1 Macroeconomic Factors

In typical asset pricing models, agents have utility over consumption, but it is well known that consumption growth and asset returns show very little correlation. Instead, we use industrial production as our main macroeconomic factor, with its availability at the monthly level an additional advantage. We extract two macro risk factors from industrial production, “good” uncertainty, denoted by p_t , and “bad” uncertainty, denoted by n_t , thereby contributing to the recent macroeconomic literature on the measurement of “real” uncertainty (see e.g. Jurado, Ludvigson, and Ng (2015)) and its effects on the real economy (see e.g. Bloom (2009)).

Specifically, under our model the change in log industrial production index, θ_t , has time-varying higher-order moments governed by two state variables: p_t and n_t . These two factors additionally affect the conditional mean of growth, which also has an autoregressive component:

$$\theta_{t+1} = \bar{\theta} + \rho (\theta_t - \bar{\theta}) + m_p(p_t - \bar{p}) + m_n(n_t - \bar{n}) + u_{t+1}, \quad (5)$$

where the growth shock is decomposed into two independent centered gamma shocks,

$$u_{t+1} = \sigma_p \omega_{p:t+1} - \sigma_n \omega_{n:t+1}. \quad (6)$$

The shocks follow centered gamma distributions with time-varying shape parameters,

$$\omega_{p:t+1} \sim \tilde{\Gamma}(p_t, 1), \quad \omega_{n:t+1} \sim \tilde{\Gamma}(n_t, 1), \quad (7)$$

⁵There is an active literature on the neural basis of risk taking behavior in a financial context, see e.g. Kuhnen and Knutson (2005) for an early paper.

where $\tilde{\Gamma}(x, 1)$ denotes a centered gamma distribution with shape parameter x and a unit scale parameter. The shape factors, p_t and n_t , follow autoregressive processes,

$$p_{t+1} = \bar{p} + \rho_p(p_t - \bar{p}) + \sigma_{pp}\omega_{p;t+1}, \quad (8)$$

$$n_{t+1} = \bar{n} + \rho_n(n_t - \bar{n}) + \sigma_{nn}\omega_{n;t+1}, \quad (9)$$

where ρ_x denotes the autoregressive term of process x_{t+1} , σ_{xx} the sensitivity to shock $\omega_{x;t+1}$, and \bar{x} the long-run mean. We denote the macroeconomic state variables as, $\mathbf{Y}_t^{mac} = [\theta_t \ p_t \ n_t]'$.⁶

Because macro risks are also allowed to affect expected growth, our model can accommodate cyclical effects (e.g., lower conditional means in bad times), or the uncertainty effect described in Bloom (2009). The conditional higher moments of output growth are linear functions of the bad and good uncertainties. For example, the conditional variance and the conditional unscaled skewness are as follows,

$$\begin{aligned} \text{Conditional Variance:} \quad E_t \left[(u_{t+1})^2 \right] &= \sigma_\rho^2 p_t + \sigma_n^2 n_t, \\ \text{Conditional Unscaled Skewness:} \quad E_t \left[(u_{t+1})^3 \right] &= 2\sigma_\rho^3 p_t - 2\sigma_n^3 n_t. \end{aligned}$$

This reveals the sense in which p_t represents “good” and n_t “bad” volatility: p_t (n_t) increases (decreases) the skewness of industrial production growth.

The state variables and shocks derived from industrial production growth serve as key macro determinants for consumption growth and cash flows.

2.2.2 Cash Flows and Cash Flow Uncertainty

To model the cash flows for equities and corporate bonds, we focus attention on two variables that exhibit strong cyclical movements, namely earnings growth (see e.g. Longstaff and Piazzesi (2004)) and corporate defaults (see e.g. Gilchrist and Zakrajšek (2012)).

Corporate Bond Loss Rate To model corporate bond pricings, we must model the possibility of defaults. Suppose a portfolio of one-period nominal bonds has a promised payoff of $C = \exp(c)$ at $(t + 1)$, but will in fact only pay an unknown fraction $F_{t+1} < 1$ of that amount. Therefore, the nominal payoff for a one-period zero-coupon defaultable corporate bond at period $t + 1$ is $C = F_{t+1} \exp(c + \ln(F_{t+1})) = \exp(c - l_{t+1})$. Thus, l_{t+1} is defined as $\ln(F_{t+1}) = -\ln(1 - L_{t+1})$ where L_{t+1} (i.e., $1 - F_{t+1}$) is the aggregate corporate loss rate, which can be computed as the default rate times one minus the recovery rate. We provide more detail on the pricing of defaultable bonds in the asset pricing section (Section 2.3).

⁶This model was selected among 8 models for the growth rate of industrial production. Specifically, the different models consider variations in the conditional mean process: (1) with an autoregressive term, (2) with 2 volatilities-in-mean terms, (3) with 2 past volatility shocks, (4) with an autoregressive term and 2 volatilities-in-mean terms; (5)–(8) are (2)–(4) with the good volatility long-run mean equal to 500. Model (8) exhibits the lowest AIC and BIC criteria and thus is applied in the rest of the analysis. In addition, the same model but with constant ρ_t is rejected.

The dynamics for the aggregate corporate bond log loss rate, l_t , are modeled as follows:

$$l_{t+1} = l_0 + \rho_{ll}l_t + m_{lp}p_t + m_{ln}n_t + \sigma_{lp}\omega_{p;t+1} + \sigma_{ln}\omega_{n;t+1} + u_{t+1}^l, \quad (10)$$

$$u_{t+1}^l = \sigma_{llp}\omega_{lp;t+1} \quad \sigma_{lln}\omega_{ln;t+1}, \quad (11)$$

$$\omega_{lp;t+1} \quad \tilde{\Gamma}(lp_t, 1), \quad \omega_{ln;t+1} \quad \tilde{\Gamma}(\bar{ln}, 1), \quad (12)$$

with $\bar{ln} > 0$, and where the law of motion for cash flow uncertainty is,

$$lp_{t+1} = \bar{lp} + \rho_{lp}(lp_t - \bar{lp}) + \sigma_{lp}\omega_{lp;t+1}. \quad (13)$$

The conditional mean depends on an autoregressive term and the good and bad macro uncertainty state variables p_t and n_t . The total disturbance of loss rate is governed by three independent heteroskedastic shocks: the good and bad environment macro shocks $\tilde{f}\omega_{p;t+1}, \omega_{n;t+1}g$ and the (orthogonal) loss rate shock u_{t+1}^l . The loss rate shock follows a typical BEGE process, but we only allow ω_{lp} 's shape parameter to be time-varying, so that only the volatility factor associated with the positively-skewed loss rate shock varies over time.⁷

This dynamic system allows macroeconomic uncertainty to affect both the conditional mean and conditional variance of the loss rate process. However, it also allows the loss rate to have an autonomous autoregressive component in its conditional mean (making l_t a state variable) and accommodates heteroskedasticity not spanned by macroeconomic uncertainty. This ‘‘financial’’ cash flow uncertainty has a time-varying component, denoted by lp_t , and a constant component denoted by \bar{ln} . If σ_{llp} and $\sigma_{lp|p}$ are positive, as we would expect, the loss rate and its volatility are positively correlated; that is, in bad times with a high incidence of defaults, there is also more uncertainty about the loss rate, and because the gamma distribution is positively skewed, the (unscaled) skewness of the process increases. We also expect the sensitivities to the good (bad) economic environment shocks, σ_{lp} (σ_{ln}) to be negative (positive): intuitively, defaults should decrease (increase) in relatively good (bad) times.

The conditional variance of the loss rate is $\sigma_{lp}^2 p_t + \sigma_{ln}^2 n_t + \sigma_{llp}^2 lp_t + \sigma_{lln}^2 \bar{ln}$, and its conditional unscaled skewness is $2 \left(\sigma_{lp}^3 p_t + \sigma_{ln}^3 n_t + \sigma_{llp}^3 lp_t - \sigma_{lln}^3 \bar{ln} \right)$. We denote the financial state variables as, $Y_t^{fin} = \begin{bmatrix} l_t & lp_t \end{bmatrix}^0$.

Earnings, Consumption and Dividends Log earnings growth, g_t , is defined as the change in log real earnings of the aggregate stock market. It is modeled as follows:

$$g_{t+1} = g_0 + \rho_{gg}g_t + \begin{matrix} ' \\ g;mac \end{matrix} Y_t^{mac} + \begin{matrix} ' \\ g;fin \end{matrix} Y_t^{fin} + \sigma_{gp}\omega_{p;t+1} + \sigma_{gn}\omega_{n;t+1} + \sigma_{glp}\omega_{lp;t+1} + \sigma_{gln}\omega_{ln;t+1} + u_{t+1}^g, \quad (14)$$

$$u_{t+1}^g = \sigma_{gg}\omega_{g;t+1}, \quad \omega_{g;t+1} \quad N(0, 1). \quad (15)$$

⁷We experimented with 5 other models including letting only ln_t follow a BEGE process, letting both ln_t and lp_t follow a BEGE process (with or without restricting the parameters to be the same across ln_t and lp_t) and, finally, a version of the last two models where ln_t and lp_t also enter the loss rate mean. Our final loss rate model outperforms other models based on standard model selection criteria. Details on alternative models are available upon request.

The conditional mean is governed by an autoregressive component and the three macro factors; the time variation in the conditional variance comes from the good and bad uncertainty factors, and the loss rate uncertainty factor. The earnings shock is assumed to be Gaussian and homoskedastic, because we fail to reject the null that the residuals series, after controlling for the heteroskedastic fundamental shocks, is Gaussian and homoskedastic.

We model consumption as stochastically cointegrated with earnings so that the consumption-earnings ratio becomes a relevant state variable. Define $\kappa_t = \ln\left(\frac{C_t}{E_t}\right)$. The model for κ_t is completely analogous to the model for g_t (i.e., replacing g by κ in Equations (14)–(15)). Similarly to earnings growth, there is an autonomous conditional mean component but the heteroskedasticity of κ_t is spanned by macroeconomic and financial uncertainties. As with log earnings growth, we fail to reject Gaussianity and homoskedasticity of u_{t+1} .

Finally, the log dividend payout ratio, η_t , is expressed as the log ratio of dividends to earnings. Recent evidence in Kostakis, Magdalinos, and Stamatogiannis (2015) shows that the monthly dividend payout ratio is stationary. We model η_t analogously to κ_t and g_t , again assuming a Gaussian and homoskedastic residual shock (as justified by the data). Using η_{t+1} and g_{t+1} , dividend growth Δd_{t+1} , is given by $\eta_{t+1} = \eta_t + g_{t+1}$.

2.2.3 Pricing Kernel State Variables

We now discuss the real pricing kernel components, consumption growth, changes in risk aversion and the inflation process needed to price nominal cash flows.

Consumption Growth By definition, log real consumption growth, $\Delta c_{t+1} = \ln\left(\frac{C_{t+1}}{C_t}\right) = g_{t+1} + \Delta\kappa_{t+1}$. Therefore, consumption growth is spanned by the previously defined state variables and shocks, and it inherits an intricate shock distribution with time-varying higher order moments, including a time-varying “volatility of volatility” (Bollerslev, Tauchen, and Zhou (2009)), that may be spiky and skewed, mimicking the jumps in consumption growth volatility in Drechsler and Yaron (2011).

Risk Aversion The state variable capturing risk aversion, $q_t = \ln\left(\frac{C_t}{C_t - H_t}\right)$, follows,

$$\begin{aligned} q_{t+1} &= q_0 + \rho_{qq}q_t + \rho_{qp}p_t + \rho_{qn}n_t \\ &\quad + \sigma_{qp}\omega_{p;t+1} + \sigma_{qn}\omega_{n;t+1} + \sigma_{qq}\omega_{q;t+1} + \sigma_q \omega_{;t+1} + u_{t+1}^q, \end{aligned} \quad (16)$$

$$u_{t+1}^q = \sigma_{qq}\omega_{q;t+1}, \omega_{q;t+1} \sim \tilde{\Gamma}(q_t, 1). \quad (17)$$

The risk aversion disturbance loads on the consumption growth shocks, and features an orthogonal preference shock. Thus, given the distributional assumptions on these shocks, the model-implied conditional variance is $\sigma_{qp}^2 p_t + \sigma_{qn}^2 n_t + \sigma_{qq}^2 q_t + \sigma_{gg}^2 + \sigma^2$, and the conditional unscaled skewness $2(\sigma_{qp}^3 p_t + \sigma_{qn}^3 n_t + \sigma_{qq}^3 q_t)$. With $\sigma_{qq} = 0$ and certain restrictions on σ_{qp} , σ_{qn} , σ_{gg} and σ_q , the model implies a perfect conditional correlation between risk aversion and real activity, as in the Campbell-Cochrane (1999) model.

We model the pure preference shock also with a demeaned gamma distributed shock, so that its variance and (unscaled) skewness are proportional to its own level. Controlling for cur-

rent business conditions, when risk aversion is high, so is its conditional variability and unscaled skewness. The plausibility of this assumption is illustrated, for example, by the pattern that option-implied volatilities, which are partially driven by risk aversion, are much more volatile in stressful times. The higher moments of risk aversion are perfectly spanned by macroeconomic uncertainty on the one hand and pure risk aversion itself (q_t) on the other hand. Key identifying assumptions are that q_t does not affect the macro variables and $u_{q;t+1}$ represents a pure preference shock. The conditional mean is modeled as before: an autonomous autoregressive component and dependence on p_t and n_t .

Inflation Because the assets we are pricing are quoted in nominal terms, we close the model with a specification for inflation. The conditional mean of inflation depends on an autoregressive term and the three macro factors \mathbf{Y}_t^{mac} . The conditional variance and higher moments of inflation are proportional to the good and bad uncertainty factors $\bar{f}_{p_t, n_t} g$. The inflation innovation u_{t+1} is assumed to be Gaussian and homoskedastic. There is no feedback from inflation to the macro variables:

$$\pi_{t+1} = \pi_0 + \rho \pi_t + \omega_{p;t+1} \mathbf{Y}_t^{mac} + \sigma_p \omega_{p;t+1} + \sigma_n \omega_{n;t+1} + u_{t+1}, \quad (18)$$

$$u_{t+1} = \sigma \omega_{p;t+1}, \omega_{n;t+1} \sim N(0, 1). \quad (19)$$

2.2.4 Matrix Representation

The dynamics of all state variables introduced above can be written compactly in matrix notation. We define the macro factors $\mathbf{Y}_t^{mac} = [\theta_t \ p_t \ n_t]^\theta$ and other state variables $\mathbf{Y}_t^{other} = [\pi_t \ l_t \ g_t \ \kappa_t \ \eta_t \ lp_t \ qt]^\theta$. Among the ten state variables, the industrial production growth θ_t , the inflation rate π_t , the loss rate l_t , earnings growth g_t , the log consumption-earnings ratio κ_t and the log divided payout ratio η_t are observable, while the other four state variables, $\bar{f}_{p_t, n_t}, lp_t, qt, g$ are latent. There are eight independent centered gamma and Gaussian shocks in this economy. The system can be formally described as follows (technical details are relegated to the Online Appendix):

$$\mathbf{Y}_{t+1} = \mathbf{A} \mathbf{Y}_t + \boldsymbol{\Sigma} \boldsymbol{\epsilon}_{t+1}, \quad (20)$$

where constant matrices, \mathbf{A} (10 × 10), $\boldsymbol{\Sigma}$ (10 × 9), are implicitly defined, $\mathbf{Y}_t = [\mathbf{Y}_t^{mac} \ \mathbf{Y}_t^{other}]^\theta$ (10 × 1) is a vector comprised of the state variable levels, and $\boldsymbol{\epsilon}_{t+1} = [\omega_{p;t+1} \ \omega_{n;t+1} \ \omega_{l;t+1} \ \omega_{lp;t+1} \ \omega_{ln;t+1} \ \omega_{g;t+1} \ \omega_{p;t+1} \ \omega_{n;t+1} \ \omega_{q;t+1}]^\theta$ (9 × 1) is a vector comprised of all the independent shocks in the economy.

Note that, among the nine shocks, five shocks are gamma distributed—the good uncertainty shock ($\omega_{p;t+1}$), the bad uncertainty shock ($\omega_{n;t+1}$), the right-tail loss rate shock ($\omega_{lp;t+1}$), the left-tail loss rate shock ($\omega_{ln;t+1}$), and the risk aversion shock ($\omega_{q;t+1}$). The remaining four shocks are standard homoskedastic Gaussian shocks (i.e., $N(0, 1)$). Importantly, given our preference structure, the state variables driving the time variation in the higher order moments of these shocks are the only ones driving the time variation in asset risk premiums and their higher order moments. Economically, we therefore rely on time variation in risk aversion – as in the

classic Campbell-Cochrane model and its variants (see e.g. Bekaert, Engstrom, and Grenadier (2010); Wachter (2006)) – and time variation in economic uncertainty – as in the Bansal-Yaron (2004) model – to explain risk premiums. The model’s implications for conditional asset return variances are critical in identifying the dynamics of risk aversion (see also Joslin, Le, and Singleton (2013)).

Our specific structure admits conditional non-Gaussianity yet generates affine pricing solutions. The model is tractable because the moment generating functions of gamma and Gaussian distributed variables can be derived in closed form, delivering exponentiated affine functions of the state variables. In particular,

$$E_t[\exp(\mathbf{1}'\mathbf{Y}_{t+1})] = \exp\left[\mathbf{1}'\mathbf{S}_0 + \frac{1}{2}\mathbf{1}'\mathbf{S}_1\boldsymbol{\Sigma}^{other}\mathbf{S}_1' + \mathbf{f}_S(\mathbf{1})\mathbf{Y}_t + S_2(\mathbf{1})\bar{\ln}\right], \quad (21)$$

where \mathbf{S}_0 (10×1) is a vector of drift coefficients; \mathbf{S}_1 (10×4) is a selection matrix of 0s and 1s which picks out the Jensen’s inequality terms of the four Gaussian shocks; $\boldsymbol{\Sigma}^{other}$ (4×4) represents the covariance of the Gaussian shocks. The matrix $\mathbf{f}_S(\mathbf{1})$ (the scalar $S_2(\mathbf{1})$) is a non-linear function of $\mathbf{1}$, involving the feedback matrix and the scale parameters of the gamma-distributed variables; see the Online Appendix for more details.

2.3 Asset Pricing

In this section, we present the model solutions. The Online Appendix contains detailed proofs and derivations.

2.3.1 The pricing kernel and asset prices

The log pricing kernel can be expressed as follows:

$$m_{t+1} = m_0 + \mathbf{m}_2'\mathbf{Y}_t + \mathbf{m}_1'\boldsymbol{\Sigma}\mathbf{1}_{t+1}, \quad (22)$$

where m_0 , \mathbf{m}_1 (10×1), \mathbf{m}_2 (10×1) are constant scalars or vectors that are implicitly defined using Equations (14)–(17). The real pricing kernel in our model follows an affine process as well. Assuming complete markets, this kernel prices any cash flow pattern spanned by our state variable dynamics.

To price nominal assets, we define the nominal pricing kernel, \tilde{m}_{t+1} , which is a simple transformation of the log real pricing kernel, m_{t+1} ,

$$\tilde{m}_{t+1} = m_{t+1} - \pi_{t+1} = \tilde{m}_0 + \tilde{\mathbf{m}}_2'\mathbf{Y}_t + \tilde{\mathbf{m}}_1'\boldsymbol{\Sigma}\mathbf{1}_{t+1}, \quad (23)$$

where \tilde{m}_0 , $\tilde{\mathbf{m}}_1$ (10×1) and $\tilde{\mathbf{m}}_2$ (10×1) are implicitly defined. The nominal risk free rate, $\tilde{r}f_t$, is defined as $-\ln fE_t[\exp(\tilde{m}_{t+1})]g$ which can be expressed as an affine function of the state vector.

To price defaultable nominal bonds, we assume that a one period nominal bond portfolio faces a fractional (logarithmic) loss of l_t . Given the structure assumed for l_t in Equation (10),

the log price-coupon ratio of the one-period defaultable bond portfolio is

$$pc_t^1 = \ln f E_t [\exp(\tilde{m}_{t+1} - l_{t+1})] g \quad (24)$$

$$= b_0^1 + \mathbf{b}_1^{1'} \mathbf{Y}_t, \quad (25)$$

where b_0^1 and $\mathbf{b}_1^{1'}$ are implicitly defined. It is straightforward to show that a portfolio of zero-coupon nominally defaultable corporate bonds, maturing in N periods, has a price that is affine in the state variables. The assumed zero-coupon structure of the payments before maturity implies that the unexpected returns to this portfolio are exactly linearly spanned by the shocks to \mathbf{Y}_t .

Equity is a claim to the dividend stream. The Online Appendix shows that the price-dividend ratio is the sum of an infinite number of exponential affine functions of the state vector, with the coefficients following simple difference equations.

2.3.2 Asset Returns

Given that the log price-coupon ratio of a defaultable corporate bond can be expressed as an exact affine function of the state variables, it immediately implies that the log nominal return (before maturity), $\tilde{r}_{t+1}^{cb} = pc_{t+1} - pc_t$, can be represented in closed-form. For equities, the log nominal equity return is derived as follows, $\tilde{r}_{t+1}^{eq} = \ln \left(\frac{PD_{t+1+1}}{PD_t} \frac{D_{t+1}}{D_t} \exp(\pi_{t+1}) \right)$. It is therefore a non-linear but known function of the state variables, which we approximate by a linear function (See the Online Appendix for details).

To account for the approximation error, we allow for two asset-specific homoskedastic shocks that are orthogonal to the state variable innovations. As a result, log nominal asset returns approximately satisfy the following factor model,

$$\tilde{r}_{t+1}^i = \tilde{\xi}_0^i + \tilde{\boldsymbol{\gamma}}_1^{i'} \mathbf{Y}_t + \tilde{\boldsymbol{\gamma}}^{i'} \boldsymbol{\Sigma} \boldsymbol{\varepsilon}_{t+1} + \varepsilon_{t+1}^i, \quad (26)$$

where \tilde{r}_{t+1}^i is the log nominal asset return i from t to $t+1$, $\tilde{\xi}_0^i = feq, cbg$; $\tilde{\boldsymbol{\gamma}}_1^i$ (10×1) is the loading vector on the state vector; $\tilde{\boldsymbol{\gamma}}^i$ (10×1) is the loading vector on the state variable shocks, and ε_{t+1}^i is a homoskedastic error term with unconditional volatility σ_j .

Rather than exploiting the model restrictions on prices, we exploit the restrictions the economy imposes on asset returns, physical variances and risk-neutral variances. Given Equation (26) and the pricing kernel, the model implies that one period expected log excess returns are given by:

$$RP_t^i - E_t(\tilde{r}_{t+1}^i) - \tilde{r}f_t = \sum_{w=p;n;lp;q} \left\{ \sigma_w(\tilde{\boldsymbol{\gamma}}^i) + \ln \left[\frac{1 - \sigma_w(\tilde{\boldsymbol{m}}_1 + \tilde{\boldsymbol{\gamma}}^i)}{1 - \sigma_w(\tilde{\boldsymbol{m}}_1)} \right] \right\} w_t + C(RP^i) \quad (27)$$

Here $C(RP^i)$ is a constant defined in the Online Appendix and (as before), $\tilde{\boldsymbol{m}}_1$ and $\tilde{\boldsymbol{\gamma}}^i$ are vectors containing the sensitivities of the log nominal pricing kernel and the log nominal asset returns to the state variable shocks, respectively. The symbol $\sigma_w(\mathbf{x})$ ($w = p, n, lp, q$) represents linear functions of state variables' sensitivities to the good uncertainty shock ($\omega_{p;t+1}$), the bad uncer-

tainty shock ($\omega_{n;t+1}$), the right-tail loss rate shock ($\omega_{lp;t+1}$), and the risk aversion shock ($\omega_{q;t+1}$). Expected excess returns thus vary through time and are affine in p_t, n_t, lp_t (macroeconomic and cash flow uncertainties) and q_t (aggregate risk aversion).

The signs of state variable coefficients are also intuitive. For instance, because $\tilde{\mathbf{m}}_1 = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & \gamma & \gamma & 0 & 0 & \gamma \end{bmatrix}^\theta$ and $\Sigma_{\bullet 9} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{qq} \end{bmatrix}^\theta$,⁸ $\sigma_q(\tilde{\mathbf{m}}_1) = \tilde{\mathbf{m}}_1^\theta \Sigma_{\bullet 9} = \gamma \sigma_{qq} > 0$, where $\gamma > 0$ follows from the concavity of the utility function and $\sigma_{qq} > 0$ implies positive skewness of risk aversion in Equation (16). It immediately implies that an asset with a negative return sensitivity to the risk aversion shock exhibits a higher risk premium when risk aversion is high. That is, for such an asset, $\sigma_q(\tilde{\mathbf{r}}^i) < 0$; then, it can be easily shown that $\sigma_q(\tilde{\mathbf{r}}^i) + \ln \left[\frac{1 - q(\tilde{\mathbf{m}}_1 + \tilde{\mathbf{r}}^i)}{1 - q(\tilde{\mathbf{m}}_1)} \right] \sigma_q(\tilde{\mathbf{r}}^i) - \frac{q(\tilde{\mathbf{r}}^i)}{1 - q(\tilde{\mathbf{m}}_1)} > 0$.

The physical conditional return variance, $VAR_t^i = VAR_t(\tilde{\mathbf{r}}_{t+1}^i)$, and the one-period risk-neutral conditional return variance, $VAR_t^{i:Q} = VAR_t^Q(\tilde{\mathbf{r}}_{t+1}^i)$, are obtained as follows:

$$VAR_t^i = \sum_{w=p;n;lp;q} \left(\sigma_w(\tilde{\mathbf{r}}^i) \right)^2 w_t + C(P^i); \quad (28)$$

$$VAR_t^{i:Q} = \sum_{w=p;n;lp;q} \left(\frac{\sigma_w(\tilde{\mathbf{r}}^i)}{1 - \frac{\sigma_q(\tilde{\mathbf{r}}^i)}{\sigma_q(\tilde{\mathbf{m}}_1)}} \right)^2 w_t + C(Q^i); \quad (29)$$

where $C(P^i)$ and $C(Q^i)$ are constants defined in the Online Appendix. The expected variances under both the physical measure and the risk-neutral measures are time-varying and affine in p_t, n_t, lp_t and q_t .

Note that the functions in Equation (29) are affine transformations from the ones in Equation (28), using the “ $\sigma(\tilde{\mathbf{m}}_1)$ ” functions. Under normal circumstances, we would expect that the relative importance of risk aversion (q_t) increases under the risk neutral measure. In Equation (29), this intuition can be formally established as $\sigma_q(\tilde{\mathbf{m}}_1)$ is positive given our parameter choices. As derived above, $\sigma_q(\tilde{\mathbf{m}}_1) = \gamma \sigma_{qq}$ is strictly positive; therefore, as long as $1 > 1 - \gamma \sigma_{qq} > 0$, the risk neutral variance should load more heavily on q_t than does the physical variance. The same risk transfer intuition does not necessarily hold for bad uncertainty n_t , because it does not only affect risk aversion (which it should affect with a positive coefficient), but also affects consumption growth through its effect on earnings growth and the consumption earnings ratio. In this case, given that $\Sigma_{\bullet 2} = \begin{bmatrix} \sigma_n & 0 & \sigma_{nn} & \sigma_n & \sigma_{ln} & \sigma_{gn} & \sigma_n & \sigma_n & 0 & \sigma_{qn} \end{bmatrix}^\theta$, the risk transfer coefficient $\sigma_n(\tilde{\mathbf{m}}_1) = \tilde{\mathbf{m}}_1^\theta \Sigma_{\bullet 2}$, reduces to $\sigma_n - \gamma(\sigma_{gn} + \sigma_n) + \gamma \sigma_{qn}$. While a negative σ_{gn} (earnings growth loading negatively on bad uncertainty) and positive σ_{qn} suggest that risk neutral variances load more heavily on bad uncertainty, it is conceivable that consumption smoothing induces a positive σ_n , which could potentially undo this effect.

3 Estimation of Macroeconomic and Cash Flow Dynamics

There are 10 state variables in the model, but only four latent state variables drive risk premiums and conditional physical and risk neutral variances in the model: two economic un-

⁸Matrix $\Sigma_{\bullet j}$ is the j -th column of the shock coefficient matrix in the state variable process, or Σ in Equation (20).

certainty variables, p_t and n_t , cash flow uncertainty, lp_t , and risk aversion, q_t . While the total number of model parameters is large, most parameters describe the dynamics of the macroeconomic and cash flow state variables. Moreover, there is no feedback from risk aversion to other state variables. Thus, it is possible to estimate all parameters governing the exogenous macroeconomic factors and cash flow processes directly from macroeconomic and cash flow data, without using financial asset prices. While the richness of the assumed dynamics gives a maximal chance to uncertainty variables to drive asset price dynamics, our approach ensures that we do not impart unrealistic dynamics to the macro and cash flow environment. We first discuss the estimation of macroeconomic factors, then of the cash flow dynamics.

3.1 The Macroeconomic Factors

Our output variable is the change in log real industrial production, θ_t , where the monthly real industrial production index is obtained from the Federal Reserve Bank at St. Louis (from January 1947 to February 2015). The system for θ_t , described in Equations (5)–(9), is estimated using Bates (2006)’s approximate MLE procedure, which delivers both parameters and filtered state variables. We collect the three state variables in $\hat{\mathbf{Y}}^{mac}_t = [\theta_t \hat{p}_t \hat{n}_t]^\theta$, where a hat superscript is used to indicate estimated variables or matrices. Similarly, we denote the filtered shocks, $\hat{\boldsymbol{\lambda}}^{mac}_t = [\hat{\omega}_{p;t} \hat{\omega}_{n;t}]^\theta$.

Table 1: The Dynamics of the Macro Factors

This table reports parameter estimates of the industrial production growth process using the monthly log growth data θ_{t+1} from January 1947 to February 2015 (source: FRED). The model involves two latent state variables: “good” economic uncertainty p_t and “bad” economic uncertainty n_t . The model is estimated using the MLE-filtration methodology described in Bates (2006). The full dynamic processes of θ_{t+1} , p_{t+1} and n_{t+1} are described in Equations (5)–(9) in Section 2. Standard errors are displayed in parentheses. Note that the effective loading of θ_{t+1} on $\omega_{n,t+1}$ is -0.00174 and the estimate of σ_{θ_n} is 0.00174. Bold (italic) coefficients have <5% (10%) p-values.

	θ_{t+1}	p_{t+1}	n_{t+1}
mean	0.00002 (0.00045)	500 (fix)	16.14206 (2.14529)
ρ	0.13100 (0.03094)	0.99968 (0.01918)	0.91081 (0.01350)
m_p	0.00001 (0.00034)		
m_n	-0.00020 (0.00002)		
$\omega_{p,t+1}$ loading	0.00011 (0.00001)	0.55277 (0.07073)	
$\omega_{n,t+1}$ loading	-0.00174 (0.00014)		2.17755 (0.15027)

The parameter estimates for the industrial production growth process are reported in Table 1. Industrial production growth features slight positive auto-correlation and high realizations of “bad” volatility decrease its conditional mean significantly. The p_t process is extremely persistent (almost a unit root) and nearly Gaussian, forcing us to fix its unconditional mean

at 500 (for such values, the skewness and kurtosis of shocks to p_t are effectively zero). The shape parameter n_t has a much lower mean, featuring gamma-distributed shocks $\omega_{n,t+1}$ with an unconditional skewness coefficient of 0.50 ($\frac{2}{16.14}$) and an excess kurtosis coefficient of 0.37 ($\frac{6}{16.14}$). It is also less persistent than the p_t process.

We graph the conditional mean of θ_t and the p_t and n_t processes in Figure 1 together with NBER recessions. The strong countercyclicality of the n_t process and the procyclicality of the conditional mean of the growth rate of industrial production are apparent from the graph. We also confirm the cyclicity by running a regression of the three processes (conditional mean, p_t , and n_t) on a constant and a NBER dummy. The NBER dummy obtains a highly statistically significant positive (negative) coefficient for the n_t (conditional mean) equation. The coefficient is in fact positive in the p_t equation as well, but not statistically significant. In fact, the n_t regression with simply a NBER dummy features an adjusted R^2 of almost 45%.

The conditional variance of industrial production and its conditional unscaled skewness are dominated by n_t and therefore highly countercyclical. Thus, exposure to such macroeconomic uncertainty may render asset risk premiums and variances countercyclical as well. The bottom plot in Figure 1 graphs the conditional variance with a 90% confidence interval that embeds parameter uncertainty. The parameter uncertainty is determined by drawing 1,000 parameter sets from the asymptotic distribution of the parameter estimates and then re-apply the Bates filter to obtain alternative conditional variance estimates. The 90% intervals are quite tight as the median relative size of the standard error to the conditional variance is only 18%.⁹

3.2 Cash Flow Dynamics

Next, we must estimate the latent cash flow uncertainty factor lp_t , which determines the time variation in the conditional variance of the log corporate bond loss rate. The log corporate bond loss rate (l) requires data on default rates and recovery rates for the US corporate bond market. We obtain data on 3-month average all-corporate bond default rates from Moody's and monthly recovery rates spanning November 1980 to February 2015 from the Federal Reserve Board. We use 6 month moving averages of these raw data to compute the log loss rate representative for each month. The estimation of the loss rate process uses data from January 1982 to February 2015.

The dynamics of the variables are described in Equations (10)–(13). We again use Bates (2006)' approximate MLE to estimate the model parameters. Unlike the BEGE structure for real shocks, for the idiosyncratic loss rate shock, only the right-tail shock (i.e., the adverse tail) is heteroskedastic. We denote the estimated right-tail loss rate shape parameter as \hat{lp}_t , and the loss rate shocks as $\hat{\omega}_{lp,t+1}$ and $\hat{\omega}_{ln,t+1}$.

As shown in Table 2, the loss rate process is persistent with the autocorrelation coefficient close to 0.83. The p_t -process does not significantly affect the loss rate process, either through the conditional mean or through shock exposures. However, the $\omega_{n,t}$ shock has a statistically significant effect on the loss rate process; moreover n_t affects the loss rate's conditional mean with a statistically significant positive coefficient. The time-varying part of the conditional variance,

⁹This computation does not take filter uncertainty into account.

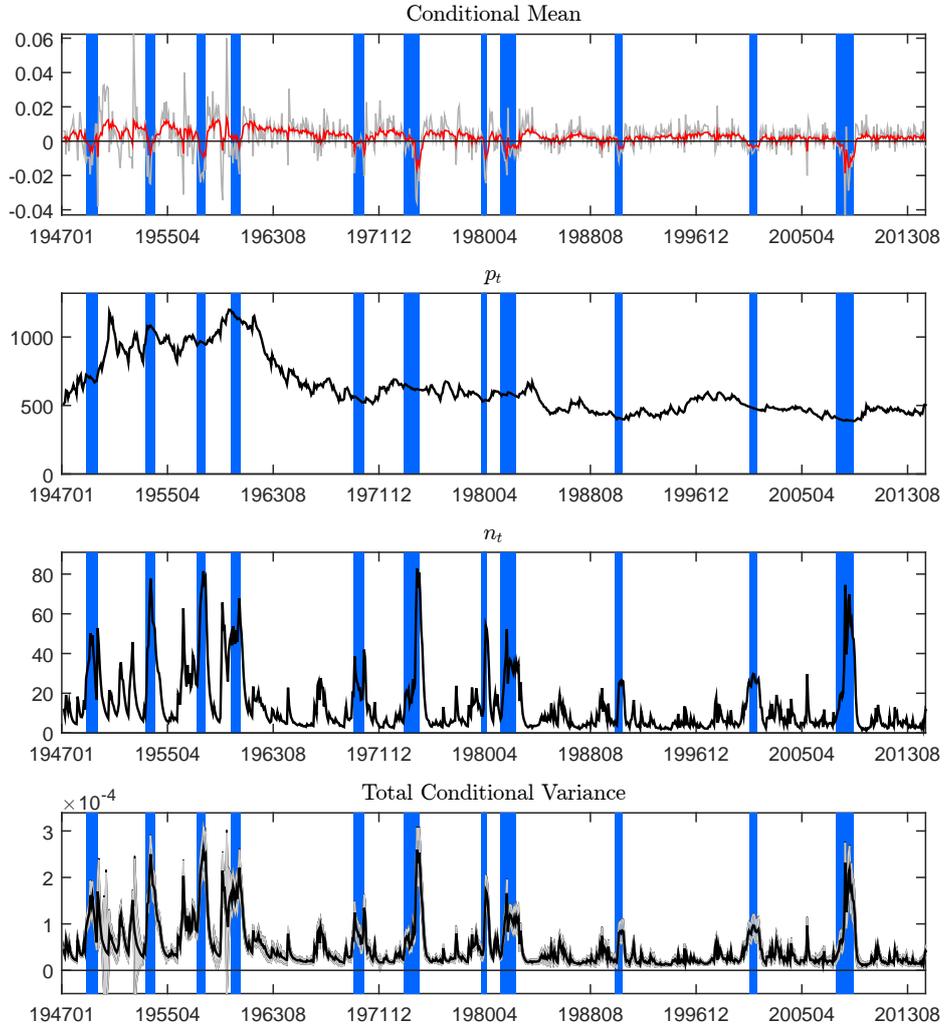


Figure 1: Macroeconomic State Variables Filtered From Industrial Production Growth

From top to bottom: conditional mean (red) and actual monthly log growth rates (gray); good uncertainty state variable, p_t ; bad uncertainty state variable, n_t ; total conditional variance of θ_{t+1} , $\sigma_{\theta_p}^2 p_t + \sigma_{\theta_n}^2 n_t$, and a 90% confidence interval reflecting parameter uncertainty (obtained through bootstrapping). The plot covers the estimation period from January 1947 to February 2015, and the estimation results are displayed in Table 1. The shaded regions are NBER recession months from the NBER website.

Table 2: The Dynamics of the Corporate Loss Rate

This table reports parameter estimates of the corporate loss rate process $l_t = \ln(L_t)$ using monthly data from June 1984 to February 2015. We obtain L_{t+1} using the identity $L_{t+1} = DEF_{t+1} (1 - RECOV_{t+1})$, where the default rate DEF_t and recovery rate $RECOV_t$ are proxied by 6-month moving averages of 3-month average all-corporate bond default rates (source: Moody's) and monthly all-corporate bond recovery rates (source: FRED), respectively. The full dynamic processes of l_{t+1} and the cash flow uncertainty state variable lp_{t+1} are described in Equations (10)–(13) in Section 2. The conditional mean part of l_{t+1} is estimated by projection first, and then the variance equation by Bates (2006)'s approximate MLE-filtration. Standard errors are displayed in parentheses. Bold (italic) coefficients have <5% (10%) p-values.

Mean:			
l_0	ρ_{ll}	m_{lp}	m_{ln}
-0.0009	0.8306	1.95E-06	1.44E-04
(0.0017)	(0.0241)	(3.57E-06)	(2.23E-05)
Shock Sensitivities:			
σ_{lp}	σ_{ln}	σ_{llp}	σ_{lln}
-4.36E-06	0.0005	0.0006	<i>1.08E-04</i>
(7.37E-06)	(0.0001)	(0.0001)	(5.78E-05)
Shape Parameter Dynamics:			
\bar{l}_p	ρ_{lp}	$\sigma_{lp lp}$	\bar{l}_n
5.2153	0.8556	1.8615	103.58
(0.2566)	(0.0126)	(0.1809)	(1.2566)

lp_t , is persistent with an autoregressive coefficient of 0.86. The idiosyncratic shocks to the loss rate process also exhibit substantial excess kurtosis (unconditional kurtosis = 1.15) and positive skewness (unconditional skewness = 0.90). The gamma shock generating negative skewness, which has a time-invariant shape parameter, is nearly Gaussian, with the shape parameter exceeding 100, so that while it contributes to the variance of the loss rate, there is no meaningful negative skewness associated with this shock.

In Figure 2, we first plot the loss rate process l . The loss rate clearly spikes around recessions, from an overall average of 0.6% to 2.1% on average in recessions (the maximum value is 5.6% during February 2009). The conditional mean of the loss rate in fact inherits the countercyclicality of the loss rate itself, given the loss rate's high persistence and its positive dependence on n_t . Our model fits the positive skewness of the loss rate process through the positively skewed u^l shocks and the positive dependence on ω_n .

Next, in the second and third plots of Figure 2, we show the conditional higher-order moments of the loss rate process, including the lp_t process and the total conditional variance. While lp_t is overall countercyclical, it appears to peak a few months after recessions. The conditional variance in the third panel ($Var_t(l_{t+1})$) also appears countercyclical, which is the combined result of a countercyclical lp_t process and a strongly countercyclical n_t process (σ_{ln} being positive). In fact, a regression of lp_t on a constant and a NBER dummy, yields a NBER coefficient of 6.78 with a t statistic of 3.03, but the t statistic increases to 8.87 when regressing the total variance on the NBER dummy.

We decompose the total conditional variance of the loss rate in its contributions coming from shocks associated with lp_t , p_t and n_t in the fourth plot of Figure 2. The dominant sources

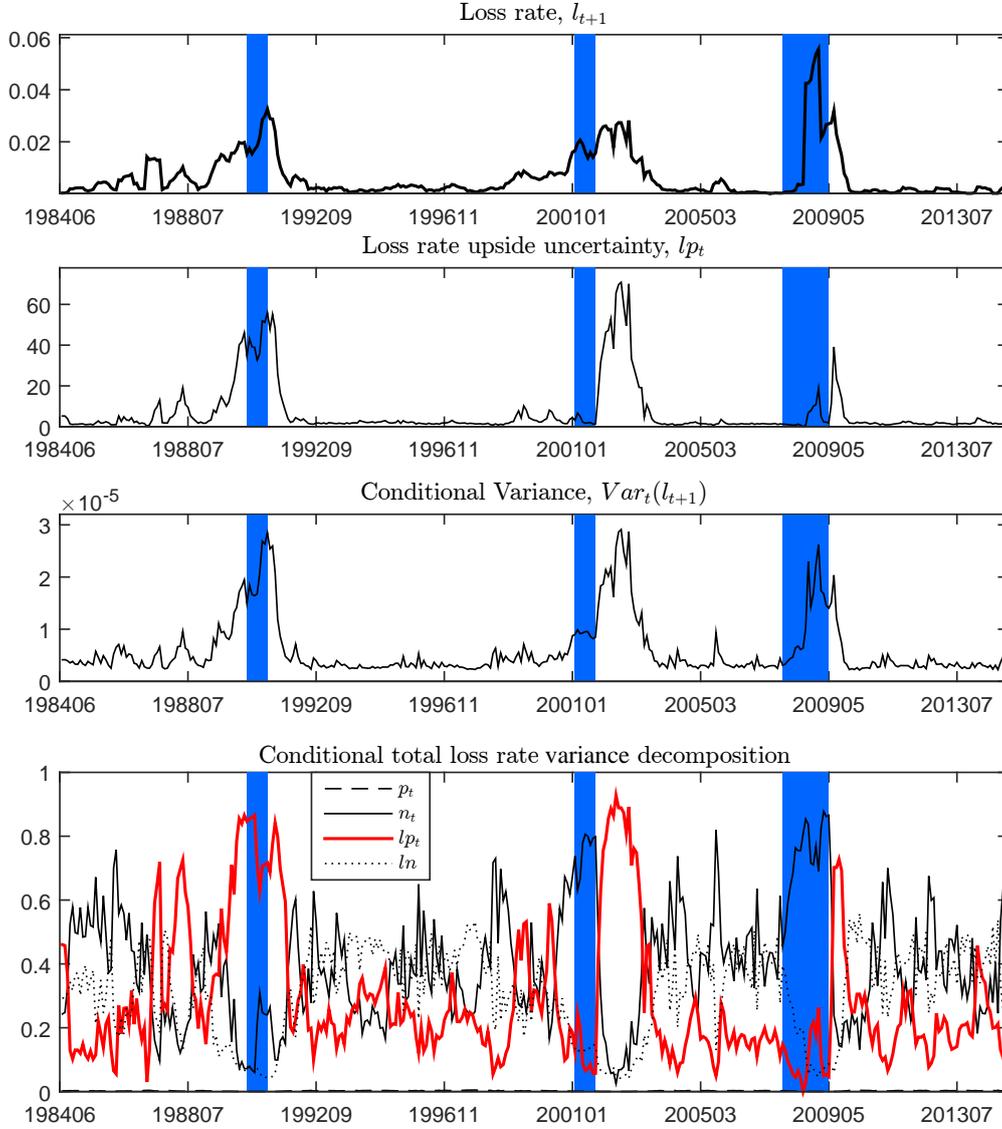


Figure 2: The Dynamics and Properties of the Corporate Loss Rate

From top to bottom: loss rate l_{t+1} , the cash flow uncertainty state variable lp_t , the total conditional variance, and the total variance decomposition. From Equations (10)–(13), the dynamics of the loss rate total disturbance are determined by four independent gamma shocks (ω_p , ω_n , ω_{lp} and ω_{ln}), and therefore its conditional variance components are additive; the fourth plot depicts the fractions of the total conditional variance explained by each of the four shocks at each point of time. The loss rate estimation uses the longest sample available from June 1984 to February 2015, and the estimation results are shown in Table 2. The shaded regions are NBER recession months from the NBER website.

of variation are lp_t (accounting for 29% of the total variance on average) and n_t (accounting for 40%). The relative importance of lp_t drops slightly in recessions while that of n_t increases, but peaks when the economy starts recovering, reaching as high as 93%. The p_t process has a negligible effect on the loss rate variance. Clearly, the loss rate variance has substantial independent variation not spanned by macroeconomic uncertainty.

With the loss rate process estimated, the dynamics of the other cash flow state variables (earnings growth, the consumption earnings ratio and the payout ratio) follow straightforwardly. We simply use linear projections of those variables onto the previously identified state variables and shocks. The data we use for these variables are standard and we relegate a discussion of the data sources and the empirical results to the Online Appendix. Noteworthy results are the strong cyclicity of earnings growth (see also Longstaff and Piazzesi (2004)), primarily reflected in the positive dependence on industrial production growth and negative dependence on the loss rate, and the countercyclicity of the conditional means of the consumption-earnings and the dividend-earnings ratios. The latter is likely a natural result of consumption and dividend smoothing, relative to highly cyclical earnings.

4 Estimation of Risk Aversion

The remaining task is to identify the structural kernel parameters, including the risk aversion process parameters, and filter the latent risk aversion process. Our approach here is unusual in that we simultaneously estimate the structural parameters while spanning risk aversion with observable financial instruments, delivering a risk aversion process that can be measured at high frequencies. We first lay out the estimation strategy and methodology and then discuss the results.

4.1 Estimation Strategy

To retrieve risk aversion from the model and data on corporate bonds and equities, we exploit the fact that, under the null of the model, asset prices, risk premiums and variances are an exact function of the state variables, including risk aversion. It thus follows that risk aversion should be spanned by a set of asset prices and risk variables. Given our desire to generate a high frequency risk aversion index, we select these instruments to be observable at high frequencies and to reflect risk and return information for our two asset classes. In particular, we postulate

$$q_t = \mathbf{z}_t' \mathbf{z}_t, \quad (30)$$

where \mathbf{z}_t is a vector of 6 observed asset prices (and a vector of ones), including (1) the term spread (the difference between the 10-year Treasury yield and the 3-month Treasury yield, where the yield data is obtained from the Federal Reserve Bank of St. Louis); (2) the credit spread (the difference between Moody's Baa yield and the 10-year Treasury bond yield); (3) a "detrended" dividend yield or earnings yield (the difference between the raw dividend yield and a moving average term that takes the 5 year average of monthly dividend yields, starting one year before, or $DY5yr_t = DY_t - \sum_{i=1}^{60} DY_{t-12-i}$ where DY_t denotes the ratio of 12-month trailing dividends

and the equity market price);¹⁰ (4) the realized equity return variance; (5) the risk-neutral equity return variance; and (6) the realized corporate bond return variance. Realized return variances rely on return data. Daily equity returns are the continuously compounded value-weighted nominal market returns with dividends from CRSP; the daily corporate bond market return is the continuously compounded log change in the daily Dow Jones corporate bond total return index (source: Global Financial Data). The monthly realized variance is the sum of the squared daily equity or corporate bond returns within the same month. The monthly return (r^{eq}) is the sum of daily returns within the same month. We use the square of the month-end VIX index (divided by 120000) as the one-period risk-neutral conditional variance of equity returns ($QVAR^{eq}$) which is obtained from the Chicago Board Options Exchange (CBOE) and is only available from the end of January 1990. We use the VXO index prior to 1990, also from CBOE, going back to June 1986.

The instruments make economic sense. The term spread may reflect information about the macro-economy (see e.g. Harvey (1988)) and was also included in the risk appetite index of Bekaert and Hoerova (2016). The credit spread and cash flow yields contain direct price information from the corporate bond and equity market respectively and thus partially reflect information about risk premiums. Ideally, we would include information on both risk-neutral and physical variances for both equities and corporate bonds, but we do not have data on the risk neutral corporate bond return variance. We use the realized variance for both markets, rather than an estimate of the physical conditional variance, because realized variances are effectively observed, whereas conditional variances must be estimated. Given a loading vector β , the risk aversion process can be computed daily from observable data.

We report some properties of these financial instruments in the Online Appendix and offer a summary here. First, all of the instruments are highly persistent. This high autocorrelation is the main reason we use a stochastically detrended dividend (AR(1)=0.982) or earnings yield (AR(1)=0.984) series rather than the actual dividend or earnings yield series.¹¹ Second, the various instruments are positively correlated but the correlations never exceed 85%. Perhaps surprisingly, the term spread is also positively correlated with the other instruments, even though it is generally believed that high term spreads indicate good times, whereas the yield and variance instruments would tend to be high in bad times. Third, 4 of the instruments show significant positive skewness. This is consistent with our assumption that risk aversion is positively skewed through its gamma distributed shock (see Equation (17)), and we need the linear spanning model to be consistent with the assumed dynamics for risk aversion. The term spread, and earnings yields are significantly negatively skewed so that a negative weight on one of them could also induce positive skewness in risk aversion, but their skewness coefficients are much smaller in magnitude.

To identify the risk aversion process and the parameters in the spanning condition, Equation (30), we exploit the restrictions the model imposes on return risk premiums (equities and

¹⁰We create an analogous detrended earnings yield variable using earnings data.

¹¹The dividend yield shows a secular decline over part of the sample that induces much autocorrelation. This decline is likely due to the introduction of a tax policy favoring repurchases rather than dividends as a means of returning cash to shareholders, and therefore not likely informative about risk aversion (see e.g. Boudoukh, Michaely, Richardson, and Roberts (2007)).

corporate bonds), physical conditional variances (equities and corporate bonds) and risk neutral variances (for equities only). In terms of measurement, risk premiums are the monthly excess returns minus the one-month Treasury bill rate from the end of the prior month (source: CRSP). We project the monthly realized variance onto the lagged risk-neutral variance and the lagged realized variance to obtain the monthly one-period physical conditional variance $PVAR^{eq}$, analogous to Bekaert, Hoerova, and Lo Duca (2013); the conditional corporate bond variance under the physical measure ($PVAR^{cb}$) is the projection of monthly realized variance onto the lagged realized variance and the lagged credit spread.

Our procedure implies that our risk aversion estimate is forced to satisfy the properties of risk aversion implied by the model: it is an element of the pricing kernel, which must, in turn, correctly price asset returns and be consistent with properties of return volatility under both the physical and risk-neutral measures. To accomplish this formally, we adopt a GMM procedure.

4.2 Estimation Methodology

The estimation is a GMM system in which we use as instruments the same variables that are used to span risk aversion (\mathbf{z}_t). For the GMM estimation, the sample spans the period from June 1986 to February 2015 (T=345 months). Apart from the parameters, we must also identify γ , the curvature parameter,¹² and the scale parameter of the preference shock, σ_{qq} . Note that the level of risk aversion is also driven by the q_t process, so that γ may not be well-identified. Therefore, we impose $\gamma = 2$. The GMM system thus has 8 unknown parameters,

$$\Theta = [\chi_0, \chi_{tsprd}, \chi_{csprd}, \chi_{CF5yr}, \chi_{rvareq}, \chi_{qvareq}, \chi_{rvarcb}, \sigma_{qq}],$$

where the notation is obvious, and $CF5yr$ refers to either a detrended dividend or earnings yield (“DY5yr” or “EY5yr”). Before the moment conditions can be evaluated, we must identify the state variables and their shocks, the pricing kernel, and the return shocks. The estimation process consists of six steps: for each candidate $\hat{\Theta} = [\hat{\gamma}, \hat{\sigma}_{qq}]$ vector,

1. Identify the implied risk aversion series given the loading choices, $\hat{q}_t = \hat{\gamma}' \mathbf{z}_t$. Consistent with the theoretical habit motivation for $q_t = \ln\left(\frac{C_t}{C_t - H_t}\right)$ (i.e., $\frac{C_t}{C_t - H_t} > 1$), and the statistical assumption for q_t (i.e., the shape parameter of the ω_q shock > 0), we impose a lower boundary of 10^{-8} on q_t during the estimation, which turns out to be non-binding.

2. Identify the state variable levels (\mathbf{Y}_t) and shocks (Σ_{t+1}). The parameters of the state variable processes, $f_{\mathbf{Y}_t^{mac}}, \mathbf{Y}_t^{fin}, \pi_t, g_t, \kappa_t, \eta_t g$, are pre-determined (see Section 3). To identify the risk aversion shock $\hat{\omega}_{q,t+1}$, we first project \hat{q}_{t+1} on $\hat{q}_t, \hat{p}_t, \hat{n}_t, \hat{\omega}_{p,t+1}, \hat{\omega}_{n,t+1}, \hat{\omega}_{g,t+1}$ and $\hat{\omega}_{\cdot,t+1}$ to obtain the residual term \hat{u}_{t+1}^q , and then divide it by $\hat{\sigma}_{qq}$ (see Equation (17)). Given $\hat{\gamma}$, a full set of state variables levels, $\hat{\mathbf{Y}}_t$, and eight independent shocks, $\hat{\boldsymbol{\epsilon}}_{t+1}$ including $\hat{\omega}_{t+1}^q$, are now identified.

3. Identify the nominal pricing kernel. Given \hat{q}_t, γ , inflation and consumption growth as

¹²Given our focus on risk premiums and volatility dynamics, the discount factor “ β ” is not identified. When using the short rate to tie down its value, we estimate its value to be around 0.98. Albuquerque, Eichenbaum, Luo, and Rebelo (2016) develop a model where variation in the discount factor plays a key role. In principle, we cannot exclude that our risk aversion shocks represent time variation in the discount factor, but we view this as very unlikely, given our external validation results discussed in Section 5.3.

the sum of real log earnings growth and the change in the log consumption-earnings ratio (i.e., $g_t + \Delta\kappa_t$), the monthly nominal kernel is obtained:

$$\widehat{m}_{t+1} = \ln(\beta) - \gamma\Delta c_{t+1} + \gamma(\widehat{q}_{t+1} - \widehat{q}_t) - \pi_{t+1}.$$

Constant matrices related to the log nominal kernel— \widetilde{m}_0 , \widetilde{m}_1 , \widetilde{m}_2 (as in the affine representation of the kernel; see Equation (23))—are implicitly identified.

4. Estimate the return loadings. We project log nominal asset returns on the 10 × 1 state variable vector \widehat{Y}_t and the 9 × 1 shock vector $\widehat{\boldsymbol{\varepsilon}}_{t+1}$:

$$\widetilde{r}_{t+1}^i = \xi_0^i + \boldsymbol{\beta}_i' \widehat{Y}_t + \boldsymbol{\gamma}_i' \widehat{\boldsymbol{\varepsilon}}_{t+1} + \varepsilon_{t+1}^i, \quad (31)$$

where \widetilde{r}_{t+1}^i is the log nominal return for asset i , $\widehat{\boldsymbol{\varepsilon}}_{t+1}$ and $\widehat{\boldsymbol{\varepsilon}}_{t+1}$ are identified previously, and the asset-specific approximation error shock ε_{t+1}^i (see Equation (26)) has mean 0 and variance σ_i^2 .

5. Obtain the model-implied endogenous moments. We derive three moments for the asset returns: 1) the expected excess return implied by the model (using the pricing kernel), RP^i ; 2) the physical (conditional expected) return variance, VAR^i , which only depends on the return definition in Equation (31) and 3) the risk neutral conditional variance, $VAR^{i:Q}$, which also uses the pricing kernel. The expressions for these variables are derived in Equations (27)–(29).

6. Obtain the moment conditions “ $(\Theta; \Psi_t)$ ”. Given data on asset returns and options, we use the derived moments to define 7 error terms that can be used to create GMM orthogonality conditions. There are three types of errors we use in the system. First, neither risk premiums nor physical conditional variances are observed in the data, but we use the restriction that the observed returns/realized variances minus their expectations under the null of the model ought to have a conditional mean of zero:

$$\mathbf{1}(\Theta; \Psi_t) = \begin{bmatrix} \left(\widetilde{r}_{t+1}^{eq} \quad \widetilde{r}f_t \right) \quad \widehat{RP}_t^{eq} \\ RVAR_{t+1}^{eq} \quad \widehat{VAR}_t^{eq} \\ \left(\widetilde{r}_{t+1}^{cb} \quad \widetilde{r}f_t \right) \quad \widehat{RP}_t^{cb} \\ RVAR_{t+1}^{cb} \quad \widehat{VAR}_t^{cb} \end{bmatrix}, \quad (32)$$

where \widetilde{r}_{t+1}^i is the realized nominal return from t to $t + 1$, $\widetilde{r}f_t$ is the nominal short rate, and $RVAR_{t+1}^i$ is the realized variance from t to $t + 1$; Ψ_t denotes the information set at time t . Because the risk neutral variance can be measured from options data, we use the error:

$$\varepsilon_2(\Theta; \Psi_t) = \left[QVAR_t^{eq} \quad \widehat{VAR}_t^{eq:Q} \right], \quad (33)$$

where $QVAR_t^{eq}$ is the ex-ante risk-neutral variance of r_{t+1}^{eq} . We assume that $\varepsilon_2(\Theta; \Psi_t)$ reflects model and measurement error, orthogonal to Ψ_t . Finally, we also construct two moment conditions to identify σ_{qq} , exploiting the model dynamics for u_{t+1}^q (i.e., the shock to the risk aversion

process as in Equation (16)):

$$\mathbf{v}_3(\Theta; \Psi_t) = \begin{bmatrix} (\hat{u}_{t+1}^q)^2 & (\hat{\sigma}_{qq})^2 \hat{q}_t \\ (\hat{u}_{t+1}^q)^3 & 2(\hat{\sigma}_{qq})^3 \hat{q}_t \end{bmatrix} \quad (34)$$

Let $\mathbf{v}_{1,2}(\Theta; \Psi_t) = [\mathbf{v}_1(\Theta; \Psi_t)^\top \ \varepsilon_2(\Theta; \Psi_t)]$. Under our assumptions these errors are mean zero given the information set, Ψ_t . We can therefore use them to create the usual GMM moment conditions. Given our previously defined set of instruments, \mathbf{z}_t (7 × 1, including a vector of 1's), we define the moment conditions as:

$$E[\mathbf{g}_t(\Theta; \Psi_t, \mathbf{z}_t)] = E \begin{bmatrix} \underbrace{\mathbf{v}_{1,2}(\Theta; \Psi_t)}_{5 \times 1} \underbrace{\mathbf{z}_t}_{7 \times 1} \\ \underbrace{\mathbf{v}_3(\Theta; \Psi_t)}_{2 \times 1} \end{bmatrix} = \underbrace{\mathbf{0}}_{37 \times 1}. \quad (35)$$

Note that to keep the set of moment conditions manageable, we only use two moment conditions for the identification of σ_{qq} . Denote $\mathbf{g}_t(\Theta; \Psi_t, \mathbf{z}_t)$ (37 × 1) as the vector of errors at time t , and $\mathbf{g}_T(\Theta; \Psi, \mathbf{z})$ (37 × 1) the sample mean of $\mathbf{g}_t(\Theta; \Psi_t, \mathbf{z}_t)$ from $t = 1$ to $t = T$. Then, the GMM objective function is,

$$J(\Theta; \Psi, \mathbf{z}) = T \mathbf{g}'_T(\Theta; \Psi, \mathbf{z}) \mathbf{W} \mathbf{g}_T(\Theta; \Psi, \mathbf{z}),$$

where \mathbf{W} is the weighting matrix. We use the standard GMM procedure, first using an identity weighting matrix, yielding a first stage set of parameters $\hat{\Theta}_1$. We then compute the optimal weighting matrix as the inverse of the spectral density at frequency zero of the orthogonality conditions, $\hat{\mathbf{S}}_1$, using 5 Newey and West (1987) lags:

$$\hat{\mathbf{S}}_1 = \sum_{j=-5}^5 \frac{1}{5} \frac{j|j|}{5} \hat{E}[\mathbf{g}_t(\hat{\Theta}_1; \Psi_t, \mathbf{z}_t) \mathbf{g}_{t-j}(\hat{\Theta}_1; \Psi_{t-j}, \mathbf{z}_{t-j})']. \quad (36)$$

Then, the inverse of $\hat{\mathbf{S}}$ is shrunk towards the identity matrix with a shrinkage parameter of 0.1 in obtaining the second-step weight matrix \mathbf{W}_2 :

$$\mathbf{W}_2 = 0.9 \hat{\mathbf{S}}^{-1} + 0.1 \mathbf{I}_{37 \times 37}, \quad (37)$$

where $\mathbf{I}_{37 \times 37}$ is a identity matrix of dimension 37 × 37. This gives rise to a second-round $\hat{\Theta}_2$ estimator. To ensure that poor first round estimates do not affect the estimation, we conduct one more iteration with shrinkage, compute $\hat{\mathbf{S}}_2(\hat{\Theta}_2)$, and produce a third-round GMM estimator, $\hat{\Theta}_3$. Lastly, the asymptotic distribution for the third-step GMM estimation parameter is, $\sqrt{T}(\hat{\Theta}_3 - \Theta_0) \xrightarrow{d} N(0, \mathbf{Avar}(\hat{\Theta}_3))$, where $\mathbf{Avar}(\hat{\Theta}_3) = (\mathbf{G}'_T(\hat{\Theta}_3) \hat{\mathbf{S}}_2^{-1} \mathbf{G}_T(\hat{\Theta}_3))^{-1}$ and where \mathbf{G}_T denotes the gradient of \mathbf{g}_T .

Because the estimation involves several steps and is quite non-linear in the parameters, we increase the chance of finding the true global optimum by starting from 24,000 different starting values for $\hat{\Theta}$ drawn randomly from a large set of possible starting values for each parameter. The global optimum is defined as the parameter estimates generating the lowest minimum objective

function value.

4.3 Risk Aversion Estimation Results

Table 3 reports the parameter estimates in the spanning relation. The system estimates 8 parameters with 37 moment conditions. The test of the over-identifying restrictions fails to reject at the 5% level but rejects at the 10% level. (We investigate the fit of the model along various dimensions in more detail later.) Except for the term spread, all instruments are significant at the 10% level or better. The positive coefficient on the risk neutral and the negative coefficient on the physical realized equity return variances is consistent with the idea that the variance risk premium (the difference between the two) may be quite informative about risk aversion in financial markets (see also Bekaert and Hoerova (2016)). To translate the coefficients into statistics of economic importance, we also report a variance decomposition, reporting the ratio of the covariance of the estimated coefficient times the instrument with risk aversion over the variance of risk aversion (these statistics sum up to 100%). Jointly the realized and risk neutral variance account for 65% of the total risk aversion variation; the credit spread and the realized corporate bond variance for about 35%. The implied risk aversion process shows a 0.45 correlation with the NBER indicator and is thus highly counter-cyclical.

Table 3: Risk Aversion Spanning Parameters

This table presents the GMM estimation results for risk aversion, $q_t = \mathbf{z}'\mathbf{z}_t$, using equity market and corporate bond market asset moments. The GMM system also consistently estimates σ_{qq} , and has a total of 8 unknown parameters. The p-value of Hansen’s over-identification test (J test) is calculated from the asymptotic χ^2 distribution with the degree of freedom being 29 (37-8). Variance decomposition, or “VARC”, of a linear variable z_t is obtained by $\chi_z \frac{cov(q_t, z_t)}{var(q_t)}$ 100% (the sum=100%). Efficient standard errors are shown in parentheses. Bold (italic) coefficients have <5% (10%) p-values. The sample period is 1986/06 to 2015/02 (345 months).

A. Estimation Results							
	Constant	χ_{tsprd}	χ_{csprd}	χ_{EY5yr}	χ_{rvareq}	$\chi_{qvarreq}$	χ_{rvarcb}
Est	0.050	-0.753	7.166	0.763	-16.984	54.038	118.248
SE	(0.014)	(0.566)	(1.030)	(0.291)	(0.490)	(1.753)	(10.826)
VARC		-0.90%	23.30%	2.18%	-34.12%	98.93%	10.62%
B. Specifications							
$\rho(q_t, NBER_t)$:	0.454			Hansen’s J:		41.1254	
SE:	(0.043)			p-value:		0.0671	

In Table 4, we estimate the dynamic properties of the risk aversion process according to Equation (16). All the parameters are estimated by OLS, except for the σ_{qq} parameter, which is delivered by the GMM estimation. The process shows moderate persistence (an autocorrelation coefficient of 0.74), and the conditional mean also shows significant positive loadings on p_t and n_t . However, q_t and n_t account for 84% and 16% of the variation in the conditional mean, respectively. Risk aversion shocks do not load significantly on the macroeconomic shocks and therefore most of their variation is driven by the risk aversion specific shock. These results suggest that economic models that impose a very tight link between aggregate fundamentals and risk aversion, such as pure habit models (Campbell and Cochrane (1999)) are missing important

Table 4: Structural Risk Aversion Parameters

This table presents the model-implied risk aversion process parameters. In the first panel, parameter estimates are obtained either from simple projection or from the GMM estimation. The second and third panels report the variance decomposition results of the conditional mean and shock structure of \hat{q}_{t+1} , respectively. In the second panel, VARC of a linear variable x with coefficient β_x is as follows, $VARC = \beta_x \frac{cov(\hat{y}, x)}{var(\hat{y})}$, where $\hat{y} = \hat{E}_t(\hat{q}_{t+1})$; VARC in the third panel uses $\hat{y} = \hat{q}_{t+1} - \hat{E}_t(\hat{q}_{t+1})$. Robust and efficient standard errors are shown in parentheses. Bold (italic) coefficients have $<5\%$ (10%) p-values. The sample period is 1986/06 to 2015/02 (345 months).

$$\begin{aligned}\hat{q}_{t+1} &= q_0 + \rho_{qq}\hat{q}_t + \rho_{qp}\hat{p}_t + \rho_{qn}\hat{n}_t + \sigma_{qp}\hat{\omega}_{p,t+1} + \sigma_{qn}\hat{\omega}_{n,t+1} + \sigma_{qg}\hat{\omega}_{g,t+1} + \sigma_{q\kappa}\hat{\omega}_{\kappa,t+1} + u_{t+1}^q, \\ u_{t+1}^q &= \sigma_{qq}\omega_{q,t+1}, \\ \omega_{q,t+1} &= \tilde{\Gamma}(q_t, 1).\end{aligned}$$

Structural Risk Aversion Parameters, q_{t+1}									
	Projection								GMM
	Constant	p_t	n_t	q_t	$\omega_{p,t+1}$	$\omega_{n,t+1}$	$\omega_{g,t+1}$	$\omega_{\kappa,t+1}$	$\omega_{q,t+1}$
Est	-0.0503	0.0003	0.0036	0.7387	<i>0.0004</i>	0.0004	-0.0002	-0.0040	0.1417
(SE)	(0.0538)	(0.0001)	(0.0007)	(0.0351)	(0.0002)	(0.0022)	(0.0069)	(0.0069)	(0.0020)
Conditional Mean Variance Decomposition (75% of Total Variance)									
VARC		p_t	n_t	q_t					
		-0.01%	16.21%	83.81%					
Shock Structure Variance Decomposition (25% of Total Variance)									
VARC					$\omega_{p,t+1}$	$\omega_{n,t+1}$	$\omega_{g,t+1}$	$\omega_{\kappa,t+1}$	$\omega_{q,t+1}$
					0.86%	0.00%	-0.01%	0.17%	99.14%

variation in actual risk aversion. In addition, risk aversion is much less persistent than the risk aversion implied by these models; the autocorrelation coefficient of the surplus ratio process in the CC model is 0.99 at the monthly level, compared to 0.74 for q_t . This result is not pre-ordained as many of the financial instruments spanning risk aversion are highly persistent, with the earnings yield being most persistent. In CC, the dividend yield is a sufficient statistic for risk aversion. Our results suggest that, in the context of our model, a measure of risk aversion that depends solely on the dividend yield would not fare well with respect to the moments that we fit in the GMM step.

Table 5 examines in more detail how well the estimated dynamic system fits critical asset price moments in the data. The moments are reported in monthly units; for example, the monthly equity premium produced by the model is 80 basis points. All model moments are within two standard errors of the data moments and most are within one standard error of the data moment.¹³ The model over-estimates the equity premium but is still close to within one standard error of the data moment. The corporate bond risk premium is 10 basis points higher than the data moment. The model implied variance moments are all quite close to their empirical counterparts. Finally, the table also reports the model-implied variance and unscaled

¹³Bootstrapped standard errors for the five asset price moments (equity risk premium, equity physical variance, equity risk-neutral variance, corporate bond risk premium, and corporate bond physical variance) use different block sizes to accommodate different serial auto-correlations, to ensure that the sampled blocks are approximately i.i.d.. Following Politis and Romano (1995) and Politis and White (2004), we look for the smallest integer after which the correlogram appears negligible, where the significance of the autocorrelation estimates is tested using the Ljung-Box Q Test (Ljung and Box (1978)).

skewness of the risk aversion innovation, $\sigma_{qq}^2 q_t$ and $2\sigma_{qq}^3 q_t$ (respectively).

The model endogenously generates the implied risk neutral variance, which can be compared with the actual risk neutral variance. The two series are 87.9% correlated, which represents a remarkable fit, but the model does fail to match some distinct spikes of the VIX in several crisis periods (see Figure F.2 in the Online Appendix).

Table 5: Fit of Moments

This table evaluates the fit of conditional moments of equity and corporate bond returns. Column “Model” reports the averages of the relevant model-implied conditional moments. The “Empirical Averages” represent the sample averages of the excess returns (for “Mom 1” and “Mom 4”), the sample averages of empirical conditional variances (for “Mom 2”, “Mom 3”, and “Mom 5”). For “Mom 6” and “Mom 7”, “Risk Aversion Innovation” refers to u_{t+1}^q in Equation (16); the variance and unscaled skewness rows compare the average model-implied conditional moments with the unconditional moments. Block bootstrapped standard errors are shown in parentheses; we allow the block size to vary for different moments to accommodate different levels of persistence: block sizes=[0 6 15 1 10] for Mom 1 to Mom 5, respectively. Asymptotic standard errors are reported for Mom 6 and Mom 7. Bold numbers denote a distance of less than 1.645 standard errors from the corresponding empirical point estimates. The sample period is 1986/06 to 2015/02 (345 months).

	Moment	Model	Empirical Average	Boot.SE/SE
Mom 1	Equity Risk Premium	0.00800	0.00530	(0.00246)
Mom 2	Equity Physical Variance	0.00325	0.00286	(0.00051)
Mom 3	Equity Risk-neutral Variance	0.00393	0.00397	(0.00049)
Mom 4	Corporate Bond Risk Premium	<i>0.00488</i>	0.00388	(0.00050)
Mom 5	Corporate Bond Physical Variance	0.00023	0.00024	(0.00003)
Mom 6	Risk Aversion Innovation Variance	0.00783	0.00843	(0.00163)
Mom 7	Risk Aversion Innovation Unscaled Skewness	0.00222	0.00164	(0.00078)

4.4 Robustness

Online Appendix G considers several robustness checks to our main estimation. We consider different values for gamma (1.1; 3.5 or estimated), consider setting p_t constant at 500 (given its minor role in asset pricing dynamics), and consider q_t loading on only p_t and n_t rather than all macro-shocks. It turns out that γ is estimated to be 2.124. While some of the spanning parameters change across different specifications, the resulting risk aversion process is highly correlated with the one analyzed in this article. The one exception is the model with $\gamma=3.5$, where the risk aversions process is only 72% correlated with the reported one, but this model is rejected by Hansen’s J-test and fails to fit the key corporate bond moments. We conclude that our current estimation is robust to slight specification variations.

5 Risk Aversion, Uncertainty and Asset Prices

In this section, we first characterize the link between risk aversion and macroeconomic uncertainty, and asset prices. Then, we present external validation evidence of our risk aversion measure.

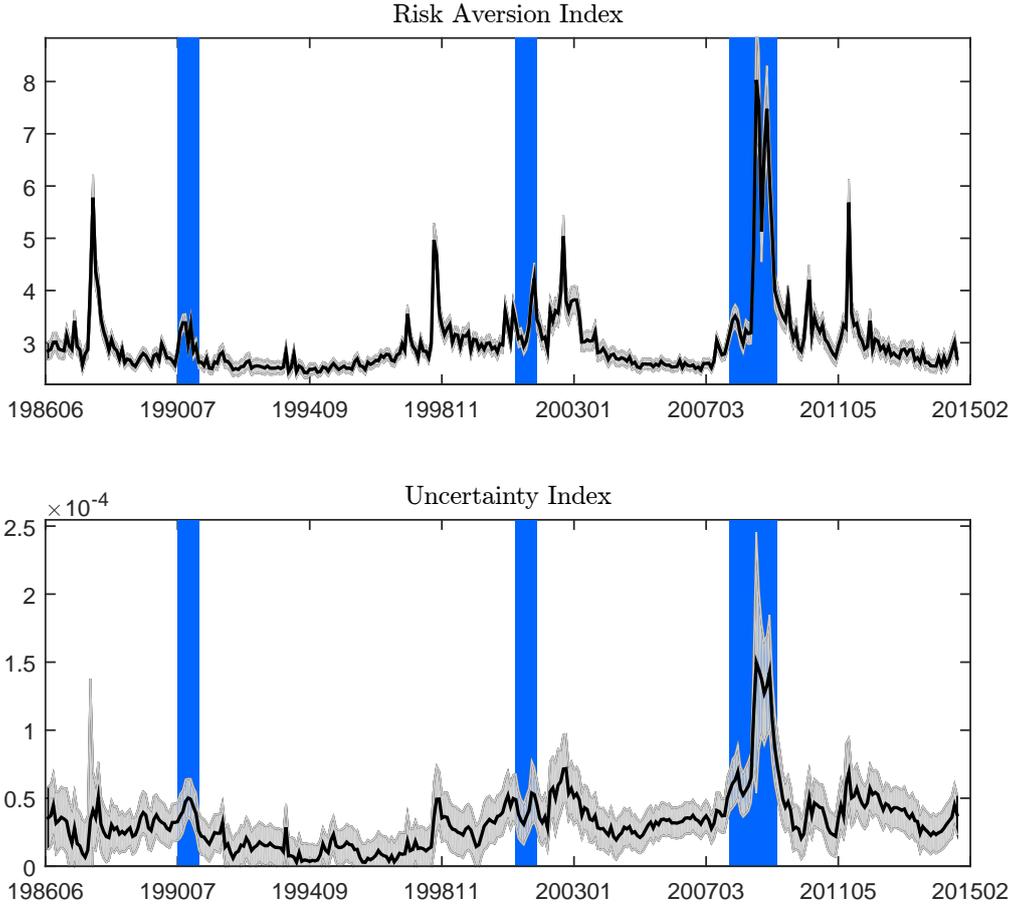


Figure 3: Risk Aversion Index and Uncertainty Index

The risk aversion index is denoted as $ra_t^{BEX} = \gamma \exp(q_t)$ and the uncertainty index is denoted as unc_t^{BEX} . Both q_t and unc_t^{BEX} are linear functions of a set of financial instruments as in Equations (30) and (39), respectively. The utility curvature parameter γ is 2. The correlation between the two series is 81.70%. The gray region denotes 90% confidence intervals around the point estimates. These intervals are computed using the asymptotic variance-covariance matrix of the spanning coefficients and the Delta method. The shaded regions are NBER recession months from the NBER website.

5.1 Risk Aversion, Uncertainty and Asset Returns

Figure 3 graphs the risk aversion process (top plot), ra_t^{BEX} , which in our model equals $ra_t^{BEX} = \gamma \exp(q_t)$. The gray lines around the estimates represent a 90% confidence interval, reflecting the sampling error in the coefficient estimates.¹⁴

Clearly, these intervals are extremely tight. The weak countercyclicality of the process is immediately apparent with risk aversion spiking in all three recessions, but also showing distinct peaks in other periods. The highest risk aversion of 8.03 is reached at the end of October in 2008, at the height of the Great Recession. However, the risk aversion process also peaks in the October 1987 crash, the August 1998 crisis (Russia default and LTCM collapse), after the TMT

¹⁴We use the asymptotic covariance matrix from the GMM estimation and the Delta method to obtain these intervals.

bull market ended in August 2002 and in August 2011 (Euro area debt crisis).

How important is risk aversion for asset prices? In our model, the priced state variables for risk premiums and variances are those entering the conditional covariance between asset returns and the pricing kernel and therefore are limited to risk aversion q_t , the macroeconomic uncertainty state variables, p_t and n_t and the loss rate variability lp_t . In Table 6, we report the loadings of risk premiums and variances on the 4 state variables. To help interpret these coefficients, we scaled the projection coefficients by the standard deviation of the state variables so that they can be interpreted as the response to a one standard deviation move in the state variable. For the equity premium, by far the most important state variable is q_t which has an effect more than 10 times larger than that of n_t . The effects of p_t and lp_t are trivially small. The economic effect of a one standard deviation change in q_t is large representing 51 basis points at the monthly level (this is a bit lower than the average monthly equity premium). For the corporate bond premium, n_t and q_t are again the most important state variables. A one standard deviation increase in n_t increases the monthly corporate bond risk premium by 4 basis points, representing more than 10% of the average monthly premium. The effect of q_t is about three times as large as the effect of n_t .

Table 6: Model-Implied Coefficients of Moments on State Variables $\hat{f}p_t, n_t, lp_t, qtg$

This table reports the decomposition of model-implied conditional moments by the four state variables, $\hat{f}p_t, n_t, lp_t, qtg$: coefficients and variability decomposition. The closed-form solution of each conditional moment is shown in Section 2 (see the Online Appendix for detailed derivations). For interpretation and reading purposes, the coefficients are multiplied by standard deviations of the corresponding state variables of the same column and then multiplied by 10000. The variance decomposition (VARC) is reported in a bold italic font and is calculated by $coefficient \frac{Cov(x_t, Mom_t)}{Var(Mom_t)}$ where $x \in \hat{f}p, n, lp, qtg$ and Mom is from Mom 1 to Mom 5. The four VARCs in the same row add up to 100%.

		Upside Economic Uncertainty	Downside Economic Uncertainty	Cash Flow Uncertainty	Risk Aversion
Conditional Moment		p_t	n_t	lp_t	q_t
Mom 1	Equity Risk Premium	0.1506	2.7861	-0.0853	50.5085
	VARC	-0.004%	3.256%	-0.033%	96.782%
Mom 2	Equity Physical Variance	0.0486	2.8394	0.0556	6.1202
	VARC	-0.063%	27.466%	0.137%	72.460%
Mom 3	Equity Risk-neutral Variance	0.0487	2.8065	0.0556	11.9204
	VARC	-0.022%	14.350%	0.083%	85.589%
Mom 4	Corporate Bond Risk Premium	0.0626	3.5398	0.1151	12.1779
	VARC	-0.032%	17.668%	0.164%	82.200%
Mom 5	Corporate Bond Physical Variance	0.0004	0.1483	0.0394	0.0181
	VARC	-0.073%	84.807%	8.136%	7.130%

The coefficients for variances are somewhat harder to interpret, but n_t and q_t remain the most important state variables with the former (latter) more important for corporate bond (equity) variances. The one variable for which q_t is only the third most important variable is the corporate bond physical variance, which reacts more strongly to n_t , and lp_t . Recall that lp_t measures the idiosyncratic component in corporate loss rates but that loss rates load very significantly on our business cycle variable.

Because the relationship between asset prices and state variables is affine, we also compute a variance decomposition. That is, we compute, for $x \in \{p, n, v, qg\}$, coefficient on x_t $\frac{Cov(x_t, Mom_t)}{Var(Mom_t)}$ where Mom represents an asset price moment like the equity risk premium, or corporate bond physical variance. These variance proportions add up to one. In the model, 97% of the equity risk premium’s variance is driven by risk aversion; only 72.5% of the corporate bond risk premium is driven by risk aversion, while 27.5% is accounted for by “bad” macroeconomic uncertainty. The physical equity variance is predominantly driven by risk aversion (72.5%) while 85% of the corporate bond return’s physical variance is driven by bad macroeconomic uncertainty. Nevertheless, macroeconomic uncertainty also accounts for 27.5% of the variance of the physical equity variance. It would be logical that the risk neutral variance would load more on risk aversion and less on macroeconomic uncertainty than the physical variance and this is indeed the case, with risk aversion accounting for 85.5% of the variance of the risk neutral variance.

Bekaert, Hoerova, and Lo Duca (2013) argue that the variance risk premium houses much information about risk aversion. Is this true in our model? To answer this question, we compute the model-implied variance risk premium as the difference between the risk neutral variance and the physical variance. A projection on the 4 state variables reveals that 96.8% of the variance of the variance risk premium is accounted for by risk aversion. Conversely, regressing risk aversion on the variance premium, the coefficient is 170.50 with a t-stat of 49.08, and the R^2 is 87.5%. Through the lens of our model, the variance premium is clearly a good proxy for risk aversion. However, they are not identical. As Cheng (2019) discusses in detail, the estimates of the variance risk premium occasionally go negative while our risk aversion process, by construction, never does. Moreover, the residual from a regression of risk aversion on the variance risk premium shows meaningful time variation which is highly correlated with the credit spread and earnings yield. This residual is also statistically significantly countercyclical. These properties of our risk aversion processes are robust across specification variation as shown in Online Appendix G.

We next verify that the model-implied risk premiums indeed predict realized excess returns. We test this in Table 7. We regress realized cumulative excess returns in the equity and corporate bond markets at various horizons on the corresponding model-implied risk premium estimates. The R^2 ’s increase with horizon, topping out at 14.63% at the 12-month horizon for corporate bonds, but do not exceed 0.43% for equity returns. All coefficients are statistically significant at the 10% level. The one month risk premiums are also more than 47.6% correlated with the NBER recession indicator, and thus countercyclical.

Given the vast literature on return predictability, it is informative to contrast the predictive power of our model implied premiums with the predictive power of the usual predictors used in the literature. We do this exercise out of sample as the literature has shown huge biases due to in sample over-fitting (Welch and Goyal (2008)) and parameter instability (Kojien and Van Nieuwerburgh (2011)). We consider five empirical models, depending on the predictors used: 1) earnings yield, 2) earnings yield, term spread and credit spread, 3) and 4) analogous with the dividend yield replacing the earnings yield, 5) physical uncertainty and variance risk premium estimate. For equity (corporate bond) returns, we use the physical uncertainty derived from

Table 7: Predicting Excess Returns Using Model-Implied Risk Premiums

This table evaluates the k-month return predictability using model-implied k-month risk premiums. The k-month excess returns are $\frac{1}{k} \sum_{i=1}^k \tilde{r}_{t+i} - \tilde{r}f_{t+i-1}$. The model-implied k-month risk premiums are $\frac{1}{k} \sum_{i=1}^k E_t(RP_{t+i-1})$, where RP_{t+i-1} denotes the model-implied one-month ahead expected excess returns of $t+i$. Given the model solution, the expectation of future risk premiums, $E_t(RP_{t+i-1})$, is obtained using the law of iterated expectations for $i > 1$. Hodrick (1992) standard errors are reported in parentheses, and adjusted R^2 s are in %. Bold (italic) coefficients have <5% (10%) p-values.

		Regression Estimates of b_k in							
		$\tilde{r}f_{t+i-1} = a_k + b_k \frac{1}{k} \sum_{i=1}^k E_t(RP_{t+i-1}) + \epsilon_{t+k}$							
		Equity:				Corporate Bond:			
		1m	3m	6m	12m	1m	3m	6m	12m
b_k		<i>0.7598</i>	<i>0.5094</i>	0.5644	0.4247	1.3802	1.6708	1.4752	1.5166
		(0.4167)	(0.2817)	(0.2356)	(0.2146)	(0.4127)	(0.2743)	(0.2241)	(0.2017)
R^2		0.24%	0.24%	0.43%	0.30%	3.18%	9.87%	11.43%	14.63%

equity (corporate bond) returns as before. We then generate out-of-sample predictions for the risk premiums by starting the sample after five years of data and then running rolling samples to generate predictions from the five-year point to one month ahead. With those competing risk premium estimates in hand, we then run simple horse races over the full sample by estimating:

$$\tilde{r}_{t+1} - \tilde{r}f_t = a \text{ Mod}(t) + (1 - a) \text{ Emp Mod}(t, i) + e_{t+1}, \text{ for } i = 1, 2, 3, \quad (38)$$

where $\text{Mod}(t)$ represents the one-month-ahead model-implied risk premium.

The results for the “a”-coefficients are reported in Table 8. The implied risk premiums from the model clearly outperform the empirical models for both equity and corporate bond returns with the “a”-coefficients being well over 0.50, varying between 0.81 and 1.01. All “a”-coefficients are highly statistically different from zero. We conclude that our model captures the predictable variation better than the fitted values extracted from standard instruments used in the literature. While it is true that the model risk premiums are not truly out-of-sample, the exercise imposes the structural parameter stability and numerous restrictions implied by the model. The poor performance of the empirical models involving the earnings and dividend yields may be surprising relative to an older return predictability literature, but direct regressions reveal that the equity yield variables have no statistically significant predictive power for our sample period.

Finally, Figure 4 compares the model-implied equity premium with the lower bound for the equity premium proposed and estimated by Martin (2017). Martin (2017) shows that the equity premium can be bounded by an index of option prices, closely related to but not identical to the VIX. Our estimates are larger but show very similar variation compared to Martin (2017)’s bound. In fact, the correlation between the two series in the overlapping sample is 95%. This is not surprising given our previous results. Risk aversion is highly correlated with the variance risk premium and is also the main determinant of the equity risk premium in the model (see Table 6). These results also provide economic confirmation of the empirical finding that the variance risk premium robustly predicts stock returns, but the conditional variance in the stock

Table 8: Out-of-Sample Predictability

This table evaluates the relative importance of model-implied risk premium estimates and empirical risk premium estimates in predicting future excess returns. “Mod” represents the model-implied risk premiums whose dynamics are fully spanned by $\tilde{r}_{p_t, n_t, l_{p_t}, q_t}g$. The empirical risk premium estimates are obtained out-of-sample (using 5-year of data); “Emp Mod (i)” (i=1,2,3,4,5) corresponds to predictor set (1) $\tilde{r}_{EY_{5yr}g}$, (2) $\tilde{r}_{fsprd, csprd, EY_{5yr}g}$, (3) $\tilde{r}_{DY_{5yr}g}$, (4) $\tilde{r}_{fsprd, csprd, DY_{5yr}g}$, (5) $\tilde{r}_{PVAR, VRPg}$. The table reports the optimal combination of model-implied and empirical risk premium estimates that minimizes the sum of squared residuals. Least Square standard errors are shown in parentheses. Bold (italic) coefficients have $<5\%$ (10%) p-values (against zero).

Least-Square Estimates of a in						
\tilde{r}_{t+1}	$\tilde{r}_t = a_i$	Mod(t) + (1	a_i	Emp Mod(t, i) +	e_{t+1}	
Emp	Mod	(1)	(2)	(3)	(4)	(5)
<i>Equity</i>						
a_i		0.8228	0.9266	0.9067	0.9658	0.8086
		(0.1079)	(0.0943)	(0.0495)	(0.0327)	(0.0802)
R^2		0.7%	1.1%	1.3%	1.7%	1.8%
<i>Corporate Bond</i>						
a_i		0.9351	0.8272	1.0114	0.8294	0.8108
		(0.1294)	(0.0798)	(0.1307)	(0.0827)	(0.0756)
R^2		0.9%	1.8%	0.9%	1.6%	2.5%

market fails to predict returns or predicts returns with a negative sign (see Bekaert and Hoerova (2014)).

5.2 Interpreting Economic Uncertainty

Because of its dependence on financial instruments, we can compute risk aversion even at a daily level. In contrast, economic uncertainty, the conditional variance of industrial production growth, is a function of both p_t and n_t , $\sigma_p^2 p_t + \sigma_n^2 n_t$ (see Table 1), and is filtered at the monthly level. Here, we use financial instruments to approximate macro uncertainty.

In Table 9, we show the coefficients from a regression of uncertainty on the financial instruments used to span risk aversion and two additional instruments, the detrended dividend yield and realized variances of speculative bond returns. We obtain monthly realized speculative corporate bond return (source: FRED, “ICE BofAML US High Yield Total Return Index”) variances using the same methodology as for overall corporate returns. Because the daily index only starts in February 1990, we use an empirical model to fill in the missing data from June 1986 to January 1990.¹⁵

The R^2 is 50% and uncertainty loads significantly on all instruments except for the realized equity and speculative bond return variances. Unlike the risk aversion process, uncertainty loads very strongly on both credit spreads and the physical corporate bond variance. The term spread also has a significant negative effect on uncertainty (and no effect on risk aversion). This makes sense as flattening yield curves are associated with future economic downturns. The

¹⁵The empirical model for imputing daily realized speculative corporate bond return variances before 1990 is explained in the Online Appendix D.

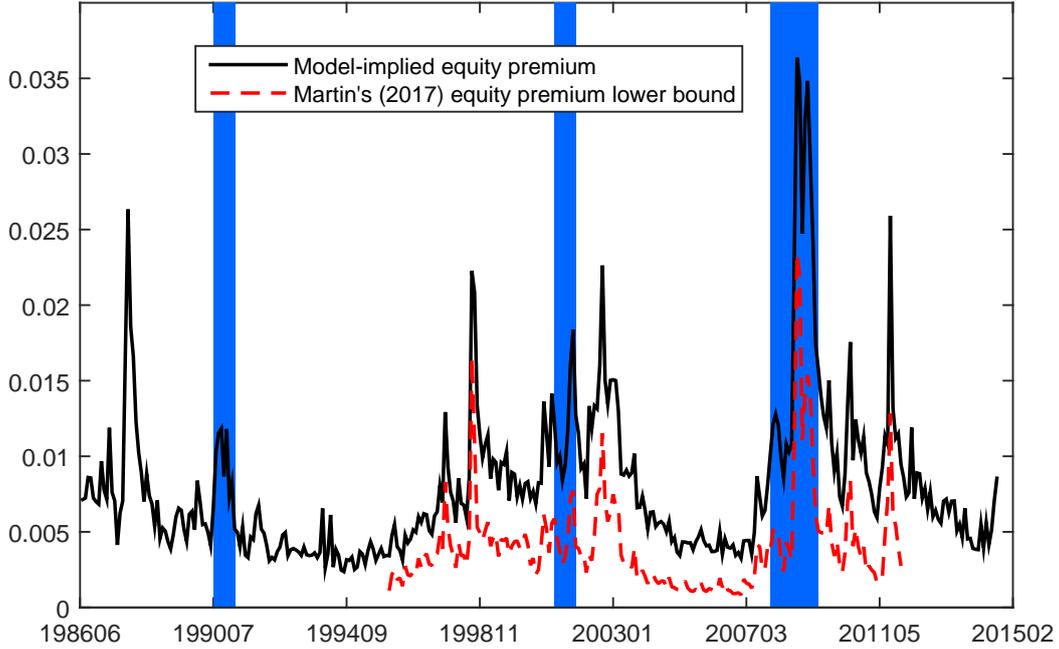


Figure 4: Model-Implied Equity Premium and Martin’s (2017) Equity Premium Lower Bound

The solid black line depicts the model-implied monthly equity premium, as formulated in Equation (27). The dashed red line depicts the “epbound” series constructed by Martin (2017), representing the lower bound to the equity premium as given by the right-hand side of the inequality in Equation (15) of his paper. This series is available from 1996/01 to 2012/01, and is downloadable from <http://personal.lse.ac.uk/martiniw/>. Note that Martin (2017) reports annualized lower bound estimates at the daily frequency; the dashed red line in this plot takes the end-of-month values and divides them by 12 to match with the monthly numbers in our analysis. The shaded regions are NBER recession months from the NBER website. The two series are 94.78% correlated.

table also reports regressions from the two components in macroeconomic uncertainty, bad and good uncertainty, onto the instruments. Clearly, the variation in macroeconomic uncertainty is dominated by the bad component and the coefficients for the bad component projection are very similar to those of total uncertainty. We also report the results from a variance decomposition applied to the fitted values of the regression. The credit spread explains almost 63% of the explained uncertainty variation. The dividend and earnings yield variables likely offset one another partially with one contributing a positive, the other a negative amount to the total variation but jointly the equity yield variables still explain close to 20%. Finally, the risk neutral equity variance and the physical corporate bond return variance each contribute about 12-14% of the explained variation of uncertainty.

From this analysis, we create an uncertainty index representing the part of economic uncertainty that is explained by the financial instruments:

$$unc_t^{BEX} = unc_t Z_t. \tag{39}$$

In the bottom plot of Figure 3, we graph the uncertainty proxy with a confidence inter-

Table 9: Projecting Macroeconomic Uncertainty on Financial Instruments

This table presents regression results of the monthly macroeconomic uncertainties (estimated from industrial production growth; see Table 1) on monthly financial instruments; some are used to span the time-varying risk aversion. “Total” indicates the total industrial production growth conditional variance, or $\sigma_{\theta p}^2 p_t + \sigma_{\theta n}^2 n_t$; “Upside”, $\sigma_{\theta p}^2 p_t$; “Downside”, $\sigma_{\theta n}^2 n_t$. “ 10^3 ” at the top means that the coefficients and their SEs are multiplied by 1000 for reporting convenience. VARC reports the variance decomposition. Robust and efficient standard errors are shown in parentheses. Adjusted R^2 s are reported. Bold (italic) coefficients have $<5\%$ (10%) p-values. The sample period is 1986/06 to 2015/02 (345 months).

	(10^3) Total	VARC	(10^3) Upside	VARC	(10^3) Downside	VARC
Constant	<i>-0.009</i> (0.005)		0.006 (0.000)		-0.015 (0.005)	
χ_{tsprd}	-0.577 (0.112)	-2.33%	-0.004 (0.002)	2.70%	-0.573 (0.112)	-2.47%
χ_{csprd}	2.024 (0.246)	62.69%	-0.016 (0.004)	6.52%	2.040 (0.246)	62.32%
χ_{DY5yr}	2.343 (0.456)	41.57%	-0.162 (0.007)	139.79%	2.505 (0.456)	44.74%
χ_{EY5yr}	-0.609 (0.189)	-22.57%	0.048 (0.003)	-55.56%	-0.657 (0.189)	-24.28%
χ_{rvareq}	-0.257 (0.620)	-3.76%	-0.002 (0.010)	-0.03%	-0.255 (0.621)	-3.65%
$\chi_{qvarreq}$	<i>1.190</i> (0.669)	13.25%	0.066 (0.010)	5.20%	<i>1.124</i> (0.670)	12.20%
χ_{rvarcb}	17.792 (5.927)	13.67%	-0.056 (0.092)	0.37%	17.848 (5.935)	13.49%
$\chi_{rvarcbSPEC}$	-2.233 (5.564)	-2.51%	-0.108 (0.087)	1.01%	-2.125 (5.571)	-2.35%
R^2	50.20%		70.80%		50.60%	

val obtained from the asymptotic variance-covariance matrix of \mathbf{unc} in Equation (39). The correlation between actual uncertainty and risk aversion is 60%; when we use the proxy the correlation increases to 82%. Obviously, most of the time crisis periods feature both high uncertainty and high risk aversion. There are exceptions however. For example, the October 1987 crash happened during a time of relatively low economic uncertainty. It also appears that at the end of the 90s, macro-uncertainty secularly increases, consistent with the Great Moderation ending around that time (see also Baele, Bekaert, Cho, Inghelbrecht, and Moreno (2015)). Note that the uncertainty index is measured with more error than is the risk aversion index.

Bloom (2009) has argued that uncertainty, extracted from data on the VIX and realized stock return variances, precedes bad economic outcomes. Segal, Shaliastovich, and Yaron (2015) show that a measure of “bad” macroeconomic uncertainty predicts economic growth negatively. We regress future real industrial production growth at various horizons on our uncertainty index — its financial proxy and the actual one — and the risk aversion process. In addition, we use the squared VIX (or QVAR in our notation). The results are in Table 10. We use Hodrick (1992) standard errors to accommodate the overlap in the data. Panel A shows univariate results. All indices predict growth with a negative sign at the one month, one quarter and one year horizons.

Table 10: On the Predictive Power of Risk Aversion and Uncertainty for Future Output Growth

This table reports the coefficient estimates of the following predictive regression,

$$\frac{1}{k} \sum_{\tau=1}^k \theta_{t+\tau} = a_k + \mathbf{b}'_k \mathbf{x}_t + \omega_{t+k},$$

where $\frac{1}{k} \sum_{\tau=1}^k \theta_{t+\tau}$ represents the future k-month industrial production growth from $t + 1$ to $t + k$, and \mathbf{x}_t represents a vector of current-month predictors: (1) our financial instrument proxy of economic uncertainty, unc^{BEX} , (2) our risk aversion, ra^{BEX} , (3) the risk-neutral conditional variance (the squared month-end VIX (VXO before 1990) index divided by 120000), $QVAR$, and (4) the true total macroeconomic uncertainty filtered from industrial production growth unc^{true} (Table 1). The coefficients are scaled by the standard deviation of the predictor in the same column for interpretation purposes. Hodrick (1992) standard errors are reported in parentheses, and adjusted R^2 s are in %. Bold (italic) coefficients have <5% (10%) p-values.

	unc^{BEX}	ra^{BEX}	$QVAR$	unc^{true}
A. Univariate				
1m	-0.0028 (0.0004) 20.6%	-0.0021 (0.0004) 11.1%	-0.0016 (0.0006) 6.5%	-0.0023 (0.0006) 13.1%
3m	-0.0027 (0.0004) 37.9%	-0.0021 (0.0004) 21.9%	-0.0017 (0.0005) 15.3%	-0.0023 (0.0005) 26.5%
12m	-0.0014 (0.0003) 17.7%	-0.0007 (0.0003) 4.3%	-0.0006 (0.0002) 3.7%	-0.0009 (0.0003) 6.5%
B. Multivariate (1)				
				R^2
1m	-0.0034 (0.0007)	0.0007 (0.0006)		21.1%
3m	-0.0031 (0.0006)	0.0005 (0.0005)		38.3%
12m	-0.0025 (0.0005)	0.0014 (0.0004)		23.3%
C. Multivariate (2)				
				R^2
1m	-0.0031 (0.0005)		0.0004 (0.0004)	20.9%
3m	-0.0028 (0.0005)		0.0001 (0.0004)	37.9%
12m	-0.0017 (0.0004)		<i>0.0005</i> (0.0002)	18.8%
D. Multivariate (3)				
				R^2
1m	-0.0025 (0.0005)		-0.0005 (0.0007)	20.9%
3m	-0.0022 (0.0004)		-0.0007 (0.0004)	39.1%
12m	-0.0016 (0.0003)		0.0003 (0.0003)	18.0%

Our financial instrument uncertainty index generates the highest R^2 by far. This suggests that it is indeed macro uncertainty predicting output growth, with the VIX having much lower predictive power in univariate regressions. The actual macroeconomic uncertainty (Column “ unc^{true} ”) exhibits substantially more predictive power than the VIX (Column “ $QVAR$ ”), but still substantially less than the combination of financial instruments most correlated with it (Column “ unc^{BEX} ”). This is likely due to the important role played by the credit spread in unc^{BEX} ; with the credit spread known to predict future economic activity (see De Santis (2018), and the references therein).

This result is confirmed in multivariate regressions. In Panels B through D of Table 10, we pit the financial instrument uncertainty proxy versus risk aversion (Panel B), the squared VIX (Panel C) and actual economic uncertainty (Panel D). In every single case, unc^{BEX} is highly statistically significant at all horizons, whereas the coefficients on the other variables mostly turn insignificant and often become positive.¹⁶

Uncertainty measures have become very popular in the macroeconomic literature. Jurado, Ludvigson, and Ng (2015) use the weighted sum of the conditional volatilities of 132 financial and macroeconomic series, with the bulk of them being macroeconomic. They have three versions of the measure depending on the forecasting horizon, but we focus on the one month horizon, which is most consistent with our model. The correlation with our economic uncertainty index, which only uses industrial production data, is highly significant and substantial at 81%.

Macroeconomic uncertainty may be correlated with political uncertainty, which has recently been proposed as a source of asset market risk premiums (Pástor and Veronesi (2013)). Baker, Bloom, and Davis (2016) create a policy uncertainty measure, based on newspaper coverage frequency. The index shows a highly significant correlation of 0.34 with our uncertainty index.

Finally, we also examine the correlation between lp_t , the idiosyncratic variance component of corporate bond loss rates, with financial instruments, but the R^2 in such a regression is only 9% (see the Online Appendix, Table F.3).

5.3 External Validation of the Risk Aversion Measure

Ultimately, our risk aversion proxy is a latent pricing kernel variable that helps the model fit corporate bond and equity risk premiums, variance dynamics and the risk neutral equity variance in an internally consistent fashion. We cannot, however, exclude that other models with alternative latent variables fit the data equally well. In addition, models outside of expected utility frameworks, such as prospect theory with probability weighting (see e.g. Baele, Driessen, Ebert, Londono, and Spalt (2019)), or models featuring biased expectations or beliefs (see e.g. Lochstoer and Muir (2020)) may provide plausible alternative explanations for the data. To hopefully increase the reader’s comfort with “ Q ” actually measuring aggregate risk aversion, we

¹⁶The positive and significant coefficient at the 12 month horizon for risk aversion is surprising. However, if we replace industrial production growth with consumption growth (to better mimic the economic model), the coefficient on risk aversion is negative and insignificant (see Table F.2 in the Online Appendix). The coefficient on economic uncertainty remains very significant and negative, supporting our finding that uncertainty dominates risk aversion in predicting economic growth.

now provide several external validation exercises.

First, while risk aversion features a pure preference shock in our model, it is motivated by a habit framework, and we therefore expect it to be consistent with the typical habit intuition. Following Wachter (2006), we create a “fundamental” risk aversion process from consumption data and the parameter estimates of Campbell and Cochrane (1999). Recall that the curvature of the utility function is a negative affine function of the log “consumption surplus ratio,” which in turns follows a heteroskedastic autoregressive process with shocks perfectly correlated with consumption growth. This “habit” risk aversion has a 0.21 correlation with our risk aversion measure, which is significantly different from zero. Work by Bekaert, Engstrom, and Grenadier (2010) and Martin (2017) also suggests the existence of more variable risk aversion in financial markets, imperfectly correlated with fundamentals.

A salient implication of the habit framework is that bad economic shocks should increase risk aversion. Even if true, it is unlikely that realized monthly or quarterly consumption growth data capture all relevant news. Much relevant economic news and events hitting markets every day are not captured in the actual economic data until much later, if at all. We therefore calculate daily measures of macro shocks (actual data minus survey expectations) around 7 macro announcements, industrial production, the unemployment rate, GDP, the CPI, balance of payments, consumer confidence and manufacturing confidence. We link our model implied risk aversion measure to their end-of-month cumulative shocks as a more direct measure of salient economic news. Models in the Campbell-Cochrane tradition predict negative links with industrial production and GDP news, and a positive link with unemployment rate news. Table ?? in the Online Appendix shows that these are indeed the three variables that show a statistically significant link with risk aversion, however, GDP growth news has an anomalous positive sign. Positive confidence news decreases risk aversion but the effect is insignificant, as are the effects of the balance of payments and inflation (undermining somewhat the Brandt and Wang (2003) model linking risk aversion to inflation). When we run a multivariate regression, the three economic activity measures remain statistically significant and no other variables are significant, with the coefficient signs remaining unchanged. Among the 7 macro news shocks, industrial production shocks alone account for 50% of the risk aversion variation explained by macro news shocks, with unemployment accounting for 33% and GDP news only accounting for 12% of explained variation. Overall, the reaction of our risk aversion to macroeconomic announcements is mostly in line with the habit intuition. However, importantly, the R^2 contributed by these macro shocks in all these regressions is quite low, which is of course consistent with our main finding that the bulk of the variation in risk aversion is not driven by macro fundamentals.

Second, and stepping outside of the paradigm of habit-based utility functions, the behavioral finance literature suggests that the sentiment of retail investors may drive asset prices and cause non-fundamental price swings. We now analyze in depth the relationship between our risk aversion measure and alternative sentiment/confidence measures that, importantly, do not rely on asset prices. The various measures are listed in Table 11. The expected sign is reported in the last column. In Panel A, we examine several measures that measure the sentiment of consumers, mostly through surveys, such as Conference Board, the University of Michigan (see also Lemmon and Portniaguina (2006) and Qiu and Welch (2004)), OECD and Reuter’s sur-

veys. Such surveys tend to measure the confidence or sentiment of consumers regarding the economic outlook and may therefore be directly related to their overall risk aversion, without reference to asset prices. Table 11 shows that all the 4 confidence/sentiment survey measures show significant and negative correlations. The strongest correlation is with Reuter’s IPSOS consumer sentiment measure. Their index is a composite index of 11 questions regarding the overall and current economic and financial situation that is run monthly via online polls. We also use two variants of a text-based measure from Da, Engelberg, and Gao (2014), who create a risk aversion measure based on the volume of internet searches for words such as “recession” and “bankruptcies”. However, their month-end indices are weakly correlated with our measure, perhaps confirming that sentiment goes beyond pure economic news. In the next to last column, we orthogonalize the various sentiment measures with respect to our economic uncertainty measure and compute the risk aversion index’s correlation with the residual. The correlations go down in magnitude, but remain negative and significant for all 4 confidence measures.

In Panel B of Table 11, we use confidence measures aimed at investors rather than consumers. Here, asset prices may indirectly influence the measures. The Yale and the American Association of Individual Investors (AAII) surveys essentially gauge the percentage of individual investors who are bullish or bearish about the stock market. The Sentix sentiment index measures investor emotion (fear, greed, etc.) using weekly surveys of more than 5,000 private and institutional investors in 14 financial markets. All measures show the correct sign, and all are significant, with the Sentix measure being particularly highly correlated with our risk aversion index (at -0.66). Again, the correlations decrease when the measures are orthogonalized with respect to economic uncertainty, but they remain significantly correlated with our risk aversion measure, with the exception of the AAI-bullish percentage. In Panel C, the OECD business confidence index is -0.36 correlated with our risk aversion index, whereas the news based sentiment measure of Shapiro, Sudhof, and Wilson (2020) is -0.49 correlated with risk aversion in Panel D. The latter measure computes the average tone of economic news articles (therefore, positive economic news is associated with positive values of the index).

We also conduct a multivariate analysis, computing the first principal component (PCA) of the standardized and orthogonalized consumer- and investor-based sentiment measures. As Panel F of Table 11 indicates, the consumer PCA receives a coefficient of 0.20; the investor PCA a coefficient of 0.16, both highly statistically significant. The consumer PCA accounts for 58.2% of the predictable variation; the investor PCA for 41.8% of the variation. The adjusted R^2 is 30.5% so that a linear function of these two PCA’s is more than 55% correlated with our risk aversion measure. Using the business confidence in either PCA measure actually worsens the fit (see Online Appendix, Table F.6).

Measures of confidence, especially when extracted from questions regarding future economic outcomes, may not necessarily be revealing about the mood and sentiment of consumers, and investors. In fact, Barsky and Sims (2012), using an analysis of the predictive content of consumer confidence for economic activity far in the future, find that confidence innovations largely reflect genuine news about future productivity, which does not show up in current macroeconomic data. They find a relatively minor role for the standard “animal spirits,” which they interpret as expectational errors (excessive optimism or pessimism about

Table 11: External Validation: Risk Aversion and Extant Consumer, Investor, Business and News Sentiment Measures

Panels A–E assemble a list of 16 widely-used sentiment and confidence measures, and then presents the correlations between our risk aversion index ra^{BEX} and these measures at the monthly (end-of-month) frequency using the longest overlapping sample. Column ρ reports correlations with the raw sentiment/confidence measures, and Column ρ^{Orth} reports correlations with measures orthogonalized by economic uncertainty (obtained from Table 1); bold correlation coefficients have <5% p-values. Column “Sign” indicates the expected sign of correlation, given the constructions of these measures. We thank Zhi Da for providing the updated data for the FEARS index. Panel F reports the contemporaneous regression results of ra^{BEX} on a consumer sentiment PCA and an investor sentiment PCA using adjusted (standardized, orthogonalized, sign-corrected) sentiment/confidence measures from Panels A and B. More details are provided in Tables F.5 and F.6 of the Online Appendix.

	Source	ρ	ρ^{Orth}	Sign
A: Survey-based consumer sentiment				
1	Conference Board consumer confidence	-0.280	-0.186	-
2	University of Michigan sentiment index, Surveys of Consumers	-0.359	-0.225	-
3	OECD consumer confidence	-0.427	-0.151	-
4	Reuter/IPSOS consumer sentiment	-0.526	-0.394	-
5	Da, Engleberg, and Gao (2014)’s FEARS25	-0.130	-0.104	+
6	Da, Engleberg, and Gao (2014)’s FEARS30	-0.130	-0.102	+
B: Survey-based investor sentiment				
7	Yale “crash” confidence (%believe in no crash)	-0.498	-0.283	-
8	Yale valuation confidence (%believe the market is not too high)	0.359	0.248	+
9	AII bullish percentage	-0.114	-0.070	-
10	AII bearish percentage	0.321	0.206	+
11	Sentix investor sentiment	-0.657	-0.423	-
C: Survey-based business sentiment				
12	OECD business confidence	-0.363	-0.225	-
D: News-based sentiment				
13	Shapiro, Sudhof, and Wilson (2020)	-0.490	-0.314	-
E: Price or macro data-based measures				
14	Baker and Wurgler (2006)’s orthogonalized sentiment	-0.161	-0.142	-
15	Credit Suisse First Boston Risk Appetite Index	-0.491	-0.282	-
16	Wachter (2006)’s Habit risk aversion	0.208	0.172	+
F: Regressing ra^{BEX} on Consumer and Investor PCAs ($R^2=30.5\%$)				
	Constant	ConsumerPC1: 1 6	InvestorPC1: 7 11	
Coef.	2.839	0.203	0.157	
(SE)	(0.078)	(0.059)	(0.060)	
VARC		58.2%	41.8%	

growth rates), but may of course also reflect true temporary mood swings. Importantly, the experimental literature (Cohn, Engelmann, Fehr, and Maréchal (2015)) suggests that positive (negative) news can invoke decreases (increases) in risk aversion. Thus, an increase in consumer confidence could still reflect a change in risk aversion, potentially even consistent with a wider interpretation of a habit model, where positive economic news should reduce risk aversion. We feel that our collective evidence is largely consistent with variation in Q reflecting changes in aggregate risk aversion. First, the strongest correlation is observed for the SentiX survey explicitly designed to reflect “investor’s emotions which fluctuate between fear and greed.” (<https://www.sentiX.de/index.php/en/item/sntm.html>) Other measures that generate higher correlations such as the Michigan survey and Reuter’s IPSOS consumer sentiment also feature questions more likely to elicit emotional responses, than predictions about the direction of the economy. Second, it is comforting to see economic news sentiment featuring such high correlation with our measure, consistent with the experimental evidence, and with models incorporating habit.

A well-known sentiment index in the academic literature is the one created by Baker and Wurgler (2006). The index is based on the first principal component of six (standardized) sentiment proxies including: the closed-end fund discount, the NYSE share turnover, the number and the average first-day returns of IPOs, the share of equity issues in total equity and debt issues, and the dividend premium (the log-difference of the average market-to-book ratios of payers and nonpayers). High values mean positive sentiment so we expect a negative correlation with our risk aversion indicator, and indeed the correlation is significantly negative but still relatively small at -0.16 . Hence, our risk aversion index correlates more with pure consumer sentiment indices than with a sentiment index derived from financial variables.

In addition, many financial services companies create their own risk appetite indices. As a well-known example, we obtain data on the Credit Suisse First Boston Risk Appetite Index. The indicator draws on the correlation between risk appetite and the relative performance of safe assets (proxied by seven to ten-year government bonds) and risky assets (equities and emerging market bonds). The underlying assumption is that an increasing risk preference shifts the demand from less risky investments to assets associated with higher risks, thus pushing their prices up relative to low-risk assets (and vice versa). The indicator is based on a cross-sectional linear regression of excess returns of 64 international stock and bond indices on their risk, approximated by their past 12-month volatility. The slope of the regression line represents the risk appetite index. The index shows a -0.49 correlation with our index and is thus highly correlated with our concept of risk aversion.

5.4 Risk Aversion, Uncertainty, and Crises

We first analyze the behavior of our risk aversion measure and uncertainty proxies during the Covid crisis. Being simple affine functions of financial instruments, we can compute both variables at the daily level throughout the March 2, 2020 to June 23, 2020 Covid crisis period. The start date is determined by the first Covid death in the US, as before that date the daily data on US Covid cases were very erratic. We download data from OWID to compute the daily

logarithmic change in Covid cases. Naturally, the pandemic is associated with feelings of fear, anxiety and uncertainty, fed by reams of bad news regarding the spread of the disease worldwide and its devastating consequences. A higher incidence of Covid cases should be plausibly associated with higher overall risk aversion. Of course, the spread of the disease is also accompanied by economic devastation, which may directly increase risk aversion through a pure habit channel, and plausibly increase economic uncertainty as well. We verify how our financial instrument proxy to uncertainty and risk aversion react to Covid case increases in Table 12. Importantly, we control for economic news, by using the sentiment measure of Shapiro, Sudhof, and Wilson (2020). Recall that this daily measure is positive (negative) when economic news sentiment is positive (negative). For ease of interpretation, we standardize both variables, so that the coefficients indicate the risk response to a one standard deviation increase in the independent variable.

Table 12 shows first that both independent variables are highly statistically significant in both regressions, with the adjusted R^2 slightly higher for the uncertainty regression (52% versus 48%). Importantly, our high frequency risk aversion reacts more to information regarding the volume of new cases of infection, than does our high frequency proxy to economic uncertainty, with the response being twice as large. In contrast, uncertainty reacts more strongly to economic news than does risk aversion, consistent with our risk aversion measure being driven more by the non-fundamental shock. In all, 86% of the explained variation in the risk aversion regression comes from Covid news, whereas only 36% does in the uncertainty regression, where economic news dominates. While indirect, this evidence is plausibly consistent with variation in Q indeed being related to changes in risk aversion.

Table 12: Risk and the Covid Crisis

We regress our daily risk variables (the risk aversion and the financial proxy to economic uncertainty) on log daily percentage changes in U.S. COVID-19 cases and the daily standardized economic news sentiment index (Shapiro, Sudhof, and Wilson (2020)). The sample spans 80 (trading) days from March 2, 2020 to June 23, 2020; the starting date is set on March 2, 2020 (the first trading day after the first death case in the U.S. was confirmed by CDC on February 29) to avoid extreme case increases during the early days. “Z” indicates standardized variables; “VARC” indicates variance decomposition. Bold correlation coefficients have <5% p-values.

DV:	Daily risk aversion ra^{BEX}			Daily uncertainty unc^{BEX}		
	Original	Z	VARC	Original	Z	VARC
Cases % Chg	22.473 (2.848)	5.184 (0.657)	86%	2.401 (0.455)	3.328 (0.630)	36%
Economic News Sentiment, Z	-1.001 (0.353)	-0.231 (0.081)	14%	-0.409 (0.056)	-0.567 (0.078)	64%
Constant	3.644 (0.446)	-0.503 (0.103)		3.137 (0.071)	-0.323 (0.099)	

With our daily risk measures in hand, we can also compare their behavior during the current Covid crisis and the Great Financial Crisis (GFC, henceforth). Figure 5 shows risk aversion and the daily proxy to uncertainty from January 2, 2020 to June 23, 2020 for the Covid crisis and from September 2, 2008 till March 31, 2009 for the GFC. Focusing first on uncertainty,

note that the long-run volatility is about 1.88%. During the Covid crisis, uncertainty almost doubled to over 4% in March and April, 2020 before dropping to below 3% in May. Using industrial production data to filter p_t and n_t during the Covid crisis, it turns out that true economic uncertainty in May, 2020, given the devastating drop in output, increased to 7.5%, a number never reached during the GFC. The financial proxy to uncertainty stayed elevated at slightly higher levels and for a longer time period, from October 2008 to April 2009, often exceeding 4%.

As to risk aversion, with a long term level of around 3, risk aversion was actually slightly below its long term level in January and February 2020 but then skyrocketed in March, reaching a high of 26.36 on March 16, when the Federal Reserve cut the Federal Funds rate to 0. Average risk aversion over March 2020 was about 10, and then dropped to an average of 4.85 in April. The steep increases in risk aversion early on in the crisis also occurred during the GFC with risk aversion peaking on October 10, 2008 (after Lehman Brothers collapsed) at 23.76. Risk aversion averaged 13.4 during October 2008, dropped to an average of 11.5 in December and then hovered around 7 till March 2009. In contrast, during the Covid crisis risk aversion has fallen more quickly and more steeply, averaging 3.9 in May 2020. More summary statistics on the behavior of the two risk variables during various months in the two crises are reported in the Online Appendix.

6 Conclusion

We formulate a no-arbitrage model where fundamentals such as industrial production, consumption earnings ratios, corporate loss rates, etc. follow dynamic processes that admit time-variation in both conditional variances and the shape of the shock distribution. The agent in the economy takes this time-varying uncertainty into account when pricing equity and corporate bonds, but also experiences preference shocks that are less than perfectly correlated with fundamentals. The state variables in the economy that drive risk premiums and higher order moments of asset prices involve risk aversion, good and bad economic uncertainty and the conditional variance of loss rates on corporate bonds. We use equity and corporate bond returns, physical equity and corporate bond return variances and the risk neutral equity variance to estimate the model parameters and simultaneously derive a risk aversion spanning process. Risk aversion is spanned by 6 financial instruments, namely the term spread, credit spread, a detrended earnings yield, realized and risk-neutral equity return variances, and the realized corporate bond return variance.

We find that risk aversion loads significantly and positively on the risk neutral equity variance and the realized corporate bond variance, and negatively on the realized equity return variance. Risk aversion is much less persistent than the risk aversion process implied by standard habit models. It is the main driver of the equity premium and the equity return risk neutral variance. It also accounts for 72% of the conditional variance of equity returns with the remainder accounted for by bad macro uncertainty. For corporate bonds, bad economic uncertainty plays a relatively more important role. It accounts for 18% of the risk premium variation and 85% of the corporate bond physical variance. Hence, different asset markets reflect differential

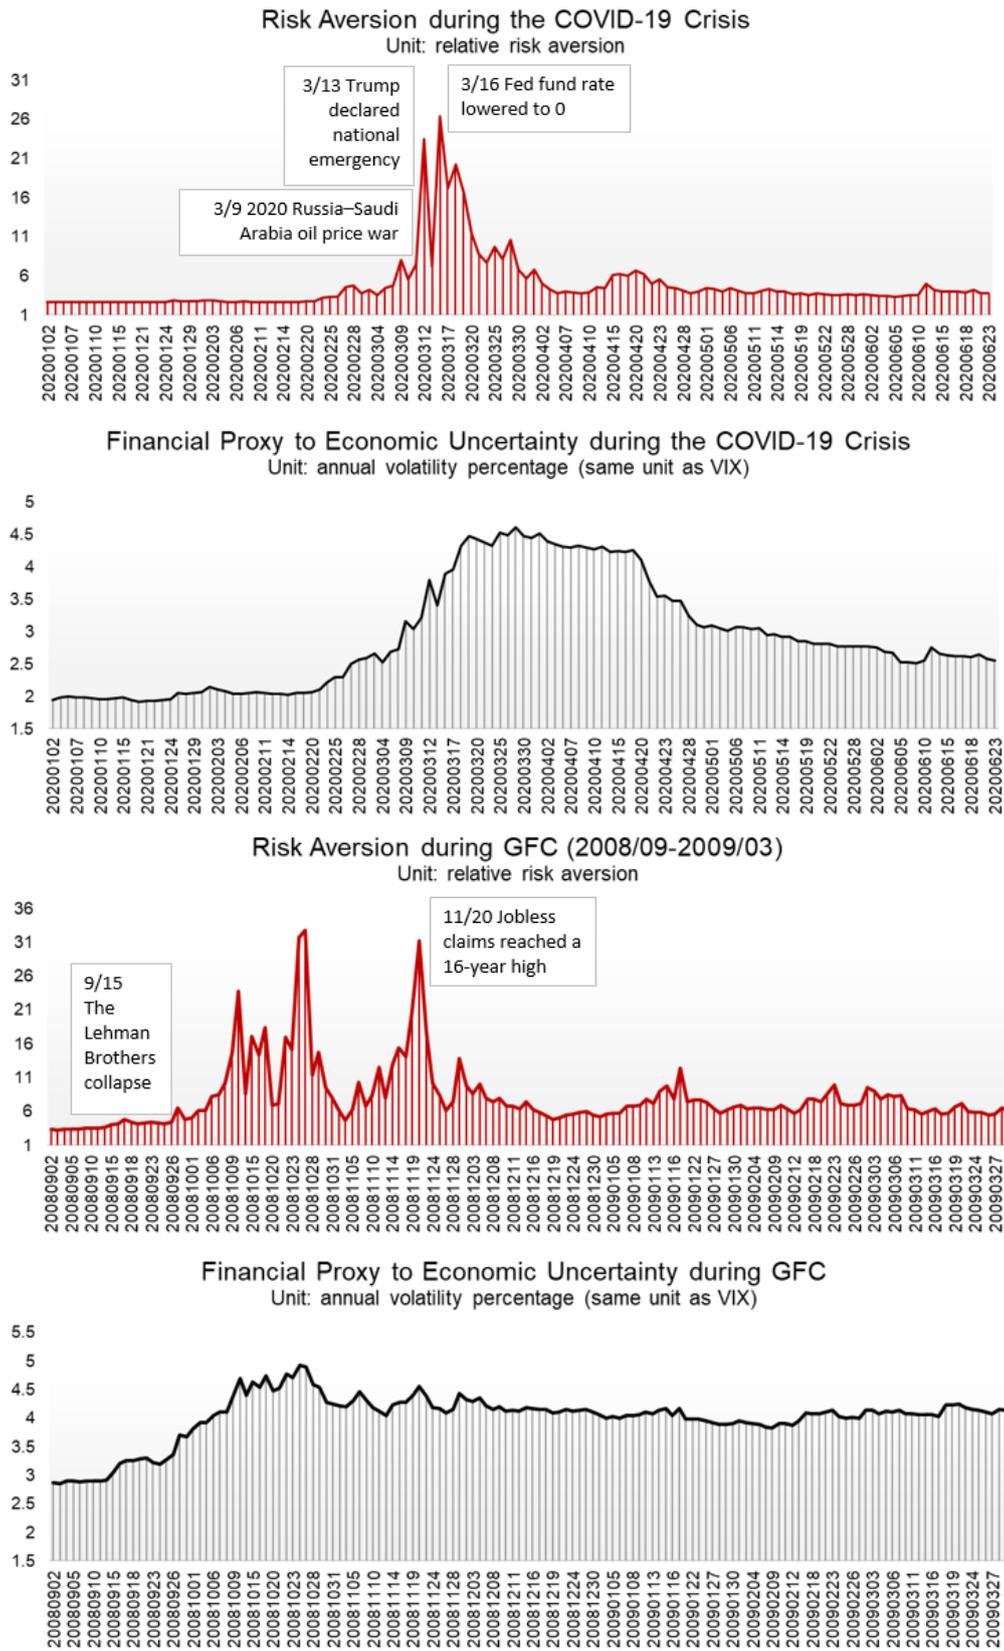


Figure 5: Risk aversion and economic uncertainty at daily frequencies around the COVID-19 crisis (top two) and the Global Financial Crisis (bottom two).

information about risk appetite versus economic uncertainty. Our model-implied risk premiums beat standard predictors of equity and corporate bond returns in an out-of-sample horse race.

Our risk aversion measure is highly correlated with the variance risk premium in equity markets, but also shows strong correlation with existing confidence/sentiment indices — especially indices measuring consumer confidence. It shows the strongest correlation with a sentiment measure focused on investor emotions. We also detect several empirical links confirming the habit model intuition beyond the strong link with measured consumption growth data. Our measure significantly reacts to industrial production news shocks, for instance, and is highly correlated with economic news sentiment.

Because the spanning instruments represent financial data, we can track the risk aversion index at higher frequencies. We also construct a financial proxy to economic uncertainty (the conditional variance of industrial production growth) which can be obtained at the daily frequency as well. The financial proxy to economic uncertainty predicts output growth negatively and significantly and is a much stronger predictor of output growth than is the VIX. In an out-of-sample analysis on the Covid crisis, risk aversion strongly reacts to Covid case increases and more so than does the uncertainty proxy. Our risk aversion and uncertainty indices are available on our websites and we plan to update them regularly, which could potentially be useful for both academic researchers and practitioners.

Our work also has implications for the dynamic asset pricing literature. To explain asset return dynamics in different asset classes, both changes in risk aversion and economic uncertainty must be accommodated. In addition, aggregate risk aversion must contain a relatively non-persistent, variable component. This variation also causes substantial variation in economically important variables such as the conditional equity premium, which is in line with recent estimates in Martin (2017). Bretscher, Hsu, and Tamoni (2019) show that risk aversion significantly affect the impact of uncertainty shocks on equity prices. Given that in our model this variation arises from an externality in preferences, it is conceivable that it is economically inefficient. More research on the determinants of risk aversion changes is clearly warranted.

Finally, we only used risky asset classes to create the risk appetite index, omitting Treasury bonds, arguably an additional important asset class. In principle, given a process for inflation our model should also price Treasury bonds. In fact, Cremers, Fleckenstein, and Gandhi (2020) claim that an implied volatility measure computed from Treasury bonds predicts the level and volatility of macroeconomic indicators better than stock market implied indicators do. However, a problem with considering Treasuries as determining general risk aversion is that they are often viewed as the benchmark “safe” assets and are subject to occasional flights-to safety (see Baele, Bekaert, Inghelbrecht, and Wei (2020)). This makes it ex-ante unlikely that a simple pricing model such as ours can jointly price the three assets classes. In our current model, interest rates are excessively volatile, for example. We therefore defer incorporating Treasury bonds to future work.

References

Abel, A. B., 1990. Asset prices under habit formation and catching up with the joneses. *American Economic Review* 80, 38–42.

- Adrian, T., Shin, H. S., 2009. Money, liquidity, and monetary policy. *American Economic Review* 99, 600–605.
- Adrian, T., Shin, H. S., 2013. Procyclical leverage and value-at-risk. *Review of Financial Studies* 27, 373–403.
- Ait-Sahalia, Y., Lo, A. W., 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9–51.
- Albuquerque, R., Eichenbaum, M., Luo, V. X., Rebelo, S., 2016. Valuation risk and asset pricing. *The Journal of Finance* 71, 2861–2904.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., Labys, P., 2003. Modeling and forecasting realized volatility. *Econometrica* 71, 579–625.
- Baele, L., Bekaert, G., Cho, S., Inghelbrecht, K., Moreno, A., 2015. Macroeconomic regimes. *Journal of Monetary Economics* 70, 51–71.
- Baele, L., Bekaert, G., Inghelbrecht, K., Wei, M., 2020. Flights to safety. *The Review of Financial Studies* 33, 689–746.
- Baele, L., Driessen, J., Ebert, S., Londono, J. M., Spalt, O. G., 2019. Cumulative prospect theory, option returns, and the variance premium. *The Review of Financial Studies* 32, 3667–3723.
- Baker, M., Wurgler, J., 2006. Investor sentiment and the cross-section of stock returns. *Journal of Finance* 61, 1645–1680.
- Baker, S. R., Bloom, N., Davis, S. J., 2016. Measuring economic policy uncertainty. *Quarterly Journal of Economics* 131, 1593–1636.
- Bakshi, G., Kapadia, N., Madan, D., 2003. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies* 16, 101–143.
- Bakshi, G., Madan, D., 2006. A theory of volatility spreads. *Management Science* 52, 1945–1956.
- Bakshi, G., Wu, L., 2010. The behavior of risk and market prices of risk over the nasdaq bubble period. *Management Science* 56, 2251–2264.
- Bansal, R., Kiku, D., Shaliastovich, I., Yaron, A., 2014. Volatility, the macroeconomy, and asset prices. *Journal of Finance* 69, 2471–2511.
- Barsky, R. B., Sims, E. R., 2012. Information, animal spirits, and the meaning of innovations in consumer confidence. *American Economic Review* 102, 1343–77.
- Bates, D. S., 2006. Maximum likelihood estimation of latent affine processes. *Review of Financial Studies* 19, 909–965.
- Bekaert, G., Engstrom, E., 2017. Asset return dynamics under habits and bad environment–good environment fundamentals. *Journal of Political Economy* 125, 713–760.
- Bekaert, G., Engstrom, E., Grenadier, S. R., 2010. Stock and bond returns with moody investors. *Journal of Empirical Finance* 17, 867–894.
- Bekaert, G., Hoerova, M., 2014. The vix, the variance premium and stock market volatility. *Journal of Econometrics* 183, 181–192.
- Bekaert, G., Hoerova, M., 2016. What do asset prices have to say about risk appetite and uncertainty? *Journal of Banking & Finance* 67, 103–118.
- Bekaert, G., Hoerova, M., Lo Duca, M., 2013. Risk, uncertainty and monetary policy. *Journal of Monetary Economics* 60, 771–788.
- Bliss, R. R., Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. *Journal of Finance* 59, 407–446.
- Bloom, N., 2009. The impact of uncertainty shocks. *Econometrica* 77, 623–685.
- Bollerslev, T., Gibson, M., Zhou, H., 2011. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of Econometrics* 160, 235–245.

- Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. *Review of Financial Studies* 22, 4463–4492.
- Bollerslev, T., Todorov, V., 2011. Tails, fears, and risk premia. *Journal of Finance* 66, 2165–2211.
- Boudoukh, J., Michaely, R., Richardson, M., Roberts, M. R., 2007. On the importance of measuring payout yield: Implications for empirical asset pricing. *Journal of Finance* 62, 877–915.
- Brandt, M. W., Wang, K. Q., 2003. Time-varying risk aversion and unexpected inflation. *Journal of Monetary Economics* 50, 1457–1498.
- Bretschler, L., Hsu, A., Tamoni, A., 2019. The real response to uncertainty shocks: the risk premium channel. Georgia Tech Scheller College of Business Research Paper .
- Britten-Jones, M., Neuberger, A., 2000. Option prices, implied price processes, and stochastic volatility. *Journal of Finance* 55, 839–866.
- Campbell, J. Y., Cochrane, J. H., 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 205–251.
- Carr, P., Wu, L., 2016. Analyzing volatility risk and risk premium in option contracts: A new theory. *Journal of Financial Economics* 120, 1–20.
- Cheng, I.-H., 2019. The vix premium. *Review of Financial Studies* 32, 180–227.
- Cohn, A., Engelmann, J., Fehr, E., Maréchal, M. A., 2015. Evidence for countercyclical risk aversion: an experiment with financial professionals. *American Economic Review* 105, 860–85.
- Coudert, V., Gex, M., 2008. Does risk aversion drive financial crises? testing the predictive power of empirical indicators. *Journal of Empirical Finance* 15, 167–184.
- Cremers, M., Fleckenstein, M., Gandhi, P., 2020. Treasury yield implied volatility and real activity. *Journal of Financial Economics* .
- Da, Z., Engelberg, J., Gao, P., 2014. The sum of all FEARS investor sentiment and asset prices. *Review of Financial Studies* 28, 1–32.
- De Santis, R. A., 2018. Unobservable systematic risk, economic activity and stock market. *Journal of Banking & Finance* 97, 51–69.
- Drechsler, I., Yaron, A., 2011. What’s vol got to do with it. *The Review of Financial Studies* 24, 1–45.
- Faccini, R., Konstantinidi, E., Skiadopoulos, G., Sarantopoulou-Chiourea, S., 2019. A new predictor of us real economic activity: The s&p 500 option implied risk aversion. *Management Science* 65, 4927–4949.
- French, K. R., Schwert, G. W., Stambaugh, R. F., 1987. Expected stock returns and volatility. *Journal of Financial Economics* 19, 3–29.
- Gai, P., Vause, N., 2006. Measuring investors’ risk appetite. *International Journal of Central Banking* 2.
- Gilchrist, S., Zakrajšek, E., 2012. Credit spreads and business cycle fluctuations. *American Economic Review* 102, 1692–1720.
- Harvey, C. R., 1988. The real term structure and consumption growth. *Journal of Financial Economics* 22, 305–333.
- He, Z., Krishnamurthy, A., 2013. Intermediary asset pricing. *American Economic Review* 103, 732–70.
- Hodrick, R. J., 1992. Dividend yields and expected stock returns: Alternative procedures for inference and measurement. *Review of Financial Studies* 5, 357–386.
- Jackwerth, J. C., 2000. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13, 433–451.
- Joslin, S., Le, A., Singleton, K. J., 2013. Why gaussian macro-finance term structure models are (nearly) unconstrained factor-VARs. *Journal of Financial Economics* 109, 604–622.

- Jurado, K., Ludvigson, S. C., Ng, S., 2015. Measuring uncertainty. *American Economic Review* 105, 1177–1216.
- Kamstra, M. J., Kramer, L. A., Levi, M. D., 2003. Winter blues: A SAD stock market cycle. *American Economic Review* 93, 324–343.
- Koijen, R. S., Van Nieuwerburgh, S., 2011. Predictability of returns and cash flows. *Annu. Rev. Financ. Econ.* 3, 467–491.
- Kostakis, A., Magdalinos, T., Stamatogiannis, M. P., 2015. Robust econometric inference for stock return predictability. *Review of Financial Studies* 28, 1506–1553.
- Kuhnen, C. M., Knutson, B., 2005. The neural basis of financial risk taking. *Neuron* 47, 763–770.
- Lemmon, M., Portniaguina, E., 2006. Consumer confidence and asset prices: Some empirical evidence. *Review of Financial Studies* 19, 1499–1529.
- Liu, J., Pan, J., Wang, T., 2004. An equilibrium model of rare-event premia and its implication for option smirks. *Review of Financial Studies* 18, 131–164.
- Ljung, G. M., Box, G. E., 1978. On a measure of lack of fit in time series models. *Biometrika* 65, 297–303.
- Lochstoer, L. A., Muir, T., 2020. Volatility expectations and returns .
- Longstaff, F. A., Piazzesi, M., 2004. Corporate earnings and the equity premium. *Journal of Financial Economics* 74, 401–421.
- Martin, I., 2017. What is the expected return on the market? *Quarterly Journal of Economics* 132, 367–433.
- Menzly, L., Santos, T., Veronesi, P., 2004. Understanding predictability. *Journal of Political Economy* 112, 1–47.
- Miranda-Agrippino, S., Rey, H., 2020. Us monetary policy and the global financial cycle. *The Review of Economic Studies* 87, 2754–2776.
- Newey, W. K., West, K. D., 1987. Hypothesis testing with efficient method of moments estimation. *International Economic Review* pp. 777–787.
- Pástor, L., Veronesi, P., 2013. Political uncertainty and risk premia. *Journal of Financial Economics* 110, 520–545.
- Pflueger, C., Siriwardane, E., Sunderam, A., 2020. Financial market risk perceptions and the macroeconomy. *The Quarterly Journal of Economics* 135, 1443–1491.
- Politis, D. N., Romano, J. P., 1995. Bias-corrected nonparametric spectral estimation. *Journal of Time Series Analysis* 16, 67–103.
- Politis, D. N., White, H., 2004. Automatic block-length selection for the dependent bootstrap. *Econometric Reviews* 23, 53–70.
- Qiu, L., Welch, I., 2004. Investor sentiment measures. National Bureau of Economic Research .
- Rajan, R. G., 2006. Has finance made the world riskier? *European Financial Management* 12, 499–533.
- Rey, H., 2013. Dilemma not trilemma: the global financial cycle and monetary policy independence .
- Rosenberg, J. V., Engle, R. F., 2002. Empirical pricing kernels. *Journal of Financial Economics* 64, 341–372.
- Ross, S., 2015. The recovery theorem. *Journal of Finance* 70, 615–648.
- Santa-Clara, P., Yan, S., 2010. Crashes, volatility, and the equity premium: Lessons from S&P 500 options. *Review of Economics and Statistics* 92, 435–451.
- Segal, G., Shaliastovich, I., Yaron, A., 2015. Good and bad uncertainty: Macroeconomic and financial market implications. *Journal of Financial Economics* 117, 369–397.
- Shapiro, A. H., Sudhof, M., Wilson, D., 2020. Measuring news sentiment .
- Wachter, J. A., 2006. A consumption-based model of the term structure of interest rates. *Journal of Financial Economics* 79, 365–399.
- Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 21, 1455–1508.
- Xu, N. R., 2019. Global risk aversion and international return comovements. Working Paper .