The impact of risk and uncertainty on expected returns☆

Evan W. Anderson[*,a], Eric Ghysels[b], Jennifer L. Juergens[c]

[a]*Northern Illinois University, Department of Economics, Zulauf 515, DeKalb, IL 60115, USA*
[b]*Department of Finance, Kenan-Flagler Business School, and Department of Economics, University of North Carolina – Chapel Hill, Chapel Hill, NC 27599-3305, USA, and Research Department, Federal Reserve Bank of New York*
[c]*University of Texas at Austin, Red McCombs School of Business, Department of Finance, Austin, TX, 78712, USA*

Abstract

We study asset pricing in economies featuring both risk and uncertainty. In our empirical analysis, we measure risk via return volatility and uncertainty via the degree of disagreement of professional forecasters, attributing different weights to each forecaster. We empirically model the typical risk-return trade-off and augment these models with our measure of uncertainty. We find stronger empirical evidence for an uncertainty-return trade-off than for the traditional risk-return trade-off. Finally, we investigate the performance of a two-factor model with risk and uncertainty in the cross section.

Article history: Received 05 April 2007; received in revised format 6 November 2008; accepted 14 November 2008

Key words: Conditional volatility, Model Uncertainty, Disagreement, Factor Models

JEL classifications: G10, G12, G17, D81

---

1. Introduction

In this paper, we empirically investigate the relation between risk, uncertainty, and expected returns. The risk-return trade-off—one of the most empirically tested theoretical relationships in finance—states that the expected excess market return should vary positively and proportionally to market volatility. This relationship is so fundamental that it could well be described as the "first law of finance." Merton (1973)

☆This paper is forthcoming in the Journal of Financial Economics. We thank Lars Hansen, Peter Hansen, Chris Jones, Leonid Kogan, Jianjun Miao, Athanasios Orphanides, Simon Potter, Eric Renault, Robert Rich, Tom Sargent and Raman Uppal for very helpful comments. This paper was presented at Northern Illinois University, SAMSI Opening Workshop, 2006 EFA meetings, 2006 Winter Meetings of the Econometric Society, the SAMSI Conference on Model Uncertainty, 2006 NBER Summer Institute Asset Pricing Workshop, 2007 AFA meetings, UC Santa Cruz, University of Illinois at Urbana-Champaign, and University of Wisconsin-Madison. Early versions of the ideas in this paper were also presented at the Board of Governors, the FDIC, the Federal Reserve Bank of New York, Simon Fraser University and Stanford University. We thank seminar participants for helpful comments.

[*]Corresponding author. Tel.: +1 815 753 5748
E-mail address: ewanderson@niu.edu (E. Anderson)

derived this theoretical relationship in a continuous time model, often referred to as Merton's Intertemporal Capital Asset Pricing Model (or simply the ICAPM). More recently, studies suggest that uncertainty, in addition to risk, should matter for asset pricing. The focus of this paper is to examine the risk-return trade-off and the uncertainty-return trade-off using an innovative empirical measure to capture uncertainty in the economy.

The empirical evidence for a risk-return trade-off is mixed. Many studies have implemented the linear regression:

$$E_t r_{et+1} = \gamma V_t,$$

where r_{et+1} is the excess return of the market over a risk-free bond, γ is a risk aversion coefficient, and V_t is the conditional volatility of the market. The goal has been to find a significantly positive γ coefficient that captures the trade-off between risk and return. Baillie and DeGennaro (1990), French, Schwert, and Stambaugh (1987), and Campbell and Hentschel (1992) find a positive but mostly insignificant relation between the conditional variance and the conditional expected return. On the other hand, Campbell (1987), Nelson (1991), and Brandt and Kang (2004), among others, find a significantly negative relation. Glosten, Jagannathan, and Runkle (1993), Harvey (2001), and Turner, Startz, and Nelson (1989) find both a positive and a negative relation depending on the estimation method used. Finally, Ghysels, Santa-Clara, and Valkanov (2005) find a significant and positive relationship between the market return and conditional volatility using *Mixed Data Sampling*, or MIDAS, estimation methods.[1]

An important strand of recent research in finance contends that uncertainty, in addition to risk, should matter for asset pricing. When agents are unsure of the correct probability laws governing the market return, they demand a higher premium in order to hold the market portfolio. Following Knight (1921), Keynes (1937) described uncertainty by saying:

> By 'uncertain' knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth owners in the social system in 1970. About these

[1]When hedging demands are present additional terms affect the conditional expected value of the market. Guo and Whitelaw (2006) argue that the additional terms make the risk-return trade-off difficult to find because these additional terms can be correlated with conditional volatility.

matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.

We adopt the position that an event is *risky* if its outcome is unknown but the distribution of its outcomes is known, and an event is *uncertain* if its outcome is unknown and the distribution of its outcomes is also unknown.

Papers by Hansen and Sargent (1995, 2001, 2003, 2005, 2007), Hansen, Sargent, and Tallarini (1999), Anderson, Hansen, and Sargent (2003), Hansen, Sargent, Turmuhambetova, and Williams (2006), Chen and Epstein (2002), Maenhout (2004, 2006), Uppal and Wang (2003), Kogan and Wang (2002), and Liu, Pan, and Wang (2005), among many others, have shown how uncertainty affects optimal decisions and asset prices. So far the literature has been mostly theoretical. The main contribution of this paper is to investigate *empirically* the performance of asset pricing models when agents face uncertainty in addition to risk. Expanding on the framework provided by Merton (1973), we show that in the presence of uncertainty the traditional risk-return regression needs to be augmented because both risk and uncertainty carry a positive premium:

$$E_t r_{et+1} = \gamma V_t + \theta M_t,$$

where θ is a measure of aversion to uncertainty and M_t measures the amount of uncertainty in the economy. When there is no uncertainty, so that $M_t = 0$, or if agents are not averse to uncertainty, so that $\theta = 0$, Merton's original formulation is recovered.[2]

In the asset pricing context typically adopted by the literature and also in this paper, agents have a great deal of information about the volatility of returns but very little about mean returns. Therefore, it is assumed that the second and higher order central moments of all returns are known exactly, while there is uncertainty about mean returns. Consequently, asset returns are uncertain only because mean returns are not known.

To measure the degree of agents' uncertainty in mean returns we propose using the disagreement of professional forecasters. The predictions of forecasters are a reasonable measure of the universe of ideas to which agents in the economy are exposed. It is likely that agents, at least partly, base their beliefs on the predictions of professional forecasters. If all forecasters are in agreement about expected returns uncertainty is likely to be low. In contrast, if forecasters state very different forecasts, agents are likely to be unsure about mean returns, and uncertainty is high. Along the lines of Hansen and Sargent's work, we assume agents choose not to act like Bayesians and combine possible probability models because they are

[2]Kogan and Wang (2002) derive the same decomposition in a more restrictive setting, as discussed in Section 2.

not sure which probabilities to use. Instead, agents solve a robust control problem.

The relationship between the disagreement of professional forecasters and expected returns has been discussed in many recent papers without explicit links to uncertainty. Most of the existing literature measures disagreement with the dispersion of earnings forecasts made by financial analysts of individual stocks and studies the relationship between this measure and individual stock returns.[3] Unlike prior studies, we emphasize *aggregate* measures of disagreement instead of disagreement about *individual* stocks or portfolios. Theoretically, we show that disagreement (or uncertainty) matters for individual stocks only when the divergence of opinions about a stock is correlated with (aggregate) market disagreement. In our empirical analysis, we use data on forecasts of aggregate corporate profits rather than earnings forecasts of individual stocks, and we quantify uncertainty, M_t, as the dispersion of predictions of mean market return forecasts constructed from the forecasts of aggregate corporate profits. Our results show that assets that are correlated with uncertainty carry a substantial premium relative to assets that are uncorrelated with our uncertainty measure.

Our empirical analysis examines risk and uncertainty both in the time series and the cross section. In the time series, we find that uncertainty is a more important determinant of the expected market excess return than risk. The correlation between our estimated measure of uncertainty and the market excess return is 0.28 whereas the correlation of our measure of risk with the market excess return is only 0.15. We also investigate the empirical performance of risk and uncertainty in the cross section of stocks by constructing portfolios with varying degrees of risk and uncertainty, estimating the prices of risk and uncertainty, and testing whether risk and uncertainty have additional explanatory power over the Fama-French factors. We find that the price of uncertainty is significantly positive and helps explain returns in the cross section.

The paper is organized as follows. Section 2 describes an economy with uncertainty and derives a theoretical decomposition of excess returns into risk and uncertainty components. Section 3 describes how we measure risk from daily volatility, whereas Section 4 describes how we measure uncertainty from the dispersion of forecasts. Section 5 empirically investigates risk-return and uncertainty-return trade-offs for the market and Section 6 empirically examines the importance of risk and uncertainty for the cross

[3] A number of authors, including Anderson, Ghysels, and Juergens (2005) and Qu, Starks, and Yan (2003) find that more disagreement, as measured by the dispersion of earnings forecasts, implies higher expected returns. In particular, Anderson, Ghysels, and Juergens (2005) observe that the dispersion factors (portfolios that are long in high dispersion stocks and are short in low dispersion stocks) are positively related to expected returns and have explanatory power beyond traditional Fama-French and momentum factors. Similarly, Qu, Starks, and Yan (2003) observe a positive relation between expected returns and a factor for disagreement, constructed from the annual volatility of a firm's earnings dispersion. Others, including Diether, Malloy, and Scherbina (2002) and Johnson (2004), find that higher dispersion stocks have lower future returns.

section of returns. Section 7 concludes.

2. The theoretical impact of risk and uncertainty on returns

We decompose asset returns into risk and uncertainty components. The decomposition is derived in a general equilibrium model, similar to that proposed by Merton (1973), populated with identical agents who have power utility functions, can invest in many risky assets and a risk-free bond, and most importantly, are concerned about model misspecification as in Hansen and Sargent (2001). Our setup closely follows approaches taken by Hansen and Sargent, except that we break the link between uncertainty and risk, and allow concerns for robustness to vary over time in ways that are not related to risk. Similar to Uppal and Wang (2003), our approach allows different assets to have different degrees of uncertainty and, following Maenhout (2004), we scale these concerns for robustness by the value function.

The decomposition of returns into risk and uncertainty components has been previously obtained by Kogan and Wang (2002). Their analysis assumes (1) utility functions that are defined over all real numbers, (2) a two-period discrete time model, and (3) asset returns and potential misspecifications that are confined to normally distributed models. The power utility function—central to Merton's analysis—is not defined for negative numbers and cannot be used in their framework. Only utility functions which are defined for negative consumption, such as the exponential and quadratic utility functions, can be used. Our formulation (1) allows agents to have power utility functions, (2) sets the analysis in the context of an infinite horizon continuous time dynamic equilibrium model, and (3) specifies asset returns and potential misspecifications as Brownian motions which allow for a broader class of asset returns and alternative models at discrete frequencies. Hence, unlike the Kogan and Wang setup, our analysis is consistent with the standard continuous time infinite horizon Merton model.

2.1. Infinite horizon continuous time equilibrium

Consider a state vector x which agents believe approximately follows the process

$$dx_t = a_t\, dt + \Lambda_t\, dB_t, \tag{1}$$

where B_t is a vector of independent standard Brownian motions; and $a_t = a(x_t)$ and $\Lambda_t = \Lambda(x_t)$ are functions of the current state. Agents perceive that the instantaneous risk-free rate is $\rho_t = \rho(x_t)$, and they can invest in a set of assets. They perceive that the price of the kth asset, P_{kt}, approximately follows the process

$$dP_{kt} = d_{kt} P_{kt}\, dt + \zeta_{kt} P_{kt}\, dB_t, \tag{2}$$

where $d_{kt} = d_k(x_t)$ is a scalar and $\zeta_{kt} = \zeta_k(x_t)$ is a row vector. Let d_t and P_t be vectors whose kth elements are d_{kt} and P_{kt}, respectively, and ζ_t be a matrix whose kth row is ζ_{kt}. The first asset is interpreted as the market. The wealth y_t of an agent approximately follows the process

$$dy_t = (\psi_t' \lambda_t y_t + \rho_t\, y_t - c_t)\; dt + \psi_t' \zeta_t y_t\, dB_t, \tag{3}$$

where $\lambda_t = d_t - \rho_t$ is the expected excess return of the available assets over the risk-free bond, ψ_t is a vector of portfolio weights whose kth element gives the fraction of wealth (possibly greater than one or less than zero) invested in the kth asset, and c_t is consumption. Wealth approximately, and not necessarily exactly, follows Eq. (3) because of the ambiguity of the asset process in Eq. (2). We call the processes in Eqs. (1), (2), and (3) the reference model.

Agents consider the possibility that the conditional mean of the state is $a_t - \Delta_t g_t$ rather than a_t and the conditional expected return of assets is $d_t - \eta_t g_t$ rather than d_t. Here $g_t = g(x_t, y_t)$ is a vector of the same dimension as B_t, $\Delta_t = \Delta(x_t)$ is a matrix of the same dimension as Λ_t, and $\eta_t = \eta(x_t)$ is a matrix of the same dimension as ζ_t. Agents believe (and indeed they are correct) that the reference model correctly specifies the conditional variances of the state (Λ_t), the conditional variances of the assets (ζ_t), and the risk-free rate (ρ_t). Yet, they worry that the underlying state, the price of assets, and the evolution of the wealth are given by

$$dx_t = (a_t - \Delta_t g_t)\; dt + \Lambda_t\, dB_t \tag{4a}$$

$$dP_{kt} = (d_{kt} - \eta_{kt} g_t)\, P_{kt}\, dt + \zeta_{kt} P_{kt}\, dB_t, \qquad \forall k \tag{4b}$$

$$dy_t = (\psi_t' \lambda_t y_t - \psi_t' y_t \eta_t g_t + \rho_t\, y_t - c_t)\; dt + \psi_t' \zeta_t y_t\, dB_t \tag{4c}$$

instead of Eqs. (1), (2), and (3). In Eq. (4b), η_{kt} is the kth row of η_t. We assume agents have full knowledge of the matrices Δ_t and η_t but do not know the value of the vector g_t.

Rather than acting as Bayesians and using a distribution for g_t, agents solve a robust control problem which provides the value of g_t as a function of the exogenous state and wealth. They consider a worst-case specification for g_t that is constrained to be close to the reference model, which is achieved by penalizing deviations from the reference model with the quadratic term $(g_t' g_t)/2\phi_t$, where $\phi_t = \phi(x_t, y_t)$ is a function that can depend on the exogenous state and wealth. The functions Δ and η allow some perturbations of x and y to be penalized more heavily than others. For example, consider a model in which both x and B (as well as g) are two-dimensional, there is only one risky asset which is the market, and for some t

$$\Delta_t = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \qquad \eta_t = \begin{bmatrix} 0 & 100 \end{bmatrix}. \tag{5}$$

In this case a higher penalty is imposed for perturbing the first element of x than the market return. In particular, perturbing the first element of x by 0.001 has the same penalty as perturbing the market return by 0.1. The second element of x is presumed to be known exactly so that under no circumstances will agents consider perturbations in it since the second element of $\Delta_t g_t$ is zero for any finite g_t. In this way, Δ and η allow us to capture the notion that agents may have more or less doubts about the conditional means of some variables compared to others.

In work by Hansen and Sargent, ϕ_t is taken to be constant; and Δ and η are linked to volatility so that $\Delta_t = \Lambda_t$ and $\eta_t = \zeta_t$ for all t. Hansen and Sargent suggest this is reasonable because it is more difficult to learn about conditional means in the presence of high volatility.[4] We do not restrict Δ and η to necessarily be tied to volatility, and we allow for the possibility that they depend more flexibly on the state. For example, there may be some state variables that have a high conditional variance, but agents have very little doubt about their conditional mean. In addition, doubts may vary over time in interesting ways that are not linked to conditional variances. During major political changes and crises agents may be willing to consider more perturbations in all variables than in stable times.[5] For the reasons discussed in Maenhout (2004) we let ϕ depend on the exogenous state and wealth (also see below for more details).[6]

Agents maximize their objective:

$$E_0 \int_0^\infty \exp(-\delta t) \left[\frac{c_t^{1-\gamma}}{1-\gamma} + \frac{1}{2\phi_t} g_t' g_t \right] dt, \tag{6}$$

by choosing adapted consumption and portfolio holdings; and they minimize their objective by choosing an adaptive process for g subject to the constraints in Eqs. (4a), (4b), and (4c), where δ is a time discount rate and E_0 denotes the expectation with respect to information available at time zero. At any date, the first component of the objective is the utility obtained from consumption where γ is a risk aversion parameter which is greater than zero and not equal to one. The second component penalizes deviations from the reference model.

[4]For example, see Anderson, Hansen, and Sargent (2003), Hansen, Sargent, and Tallarini (1999), and Hansen, Sargent, Turmuhambetova, and Williams (2006) .

[5]Figures presented in later sections show that our measure of uncertainty tends to be small when an incumbent president is reelected.

[6]Uppal and Wang (2003) allow the parameter ϕ to vary across assets and state variables, though they require their parameters to be time-invariant. Our model could be viewed as a generalization of their model in which the uncertainty parameters are allowed to vary over time.

The agent's value function, denoted $J(y, x)$, satisfies the Hamilton-Jacobi equation

$$\begin{aligned} 0 = \max_{\psi, c} \min_{g} \Bigg[& \frac{c^{1-\gamma}}{1-\gamma} + \frac{1}{2\phi} g' g - \delta J \\ & + J_x' (a - \Delta g) + J_y (\psi' \lambda y - \psi' y \eta g + \rho y - c) + \\ & \frac{1}{2} \mathrm{tr} (\Lambda \Lambda' J_{xx}) + \frac{1}{2} J_{yy} \psi' \zeta \zeta' \psi y^2 + \psi' \zeta \Lambda' J_{xy} y \Bigg], \end{aligned} \quad (7)$$

where we drop the t subscripts and the subscripts on J denote differentiation. In the limit as ϕ approaches zero at every date, the functional Eq. (7) becomes the usual Hamilton-Jacobi equation studied by Merton (1973). The additional terms present are the same terms present in Hansen and Sargent's formulation except that Δ and η are flexible functions of the state and ϕ is a flexible function of the state and wealth. The minimizing choice of g is

$$g = \phi \Delta' J_x + \phi \eta' \psi y J_y, \quad (8)$$

which illustrates how specifications of ϕ, Δ, and η endogenously determine the perturbations of conditional means that agents consider. The optimal choice of the fraction of wealth to invest in the market, ψ, satisfies the first-order condition:

$$J_y \lambda y - J_y y \eta g + J_{yy} \zeta \zeta' \psi y^2 + \zeta \Lambda' J_{xy} y = 0. \quad (9)$$

In equilibrium, market clearing requires that all agents invest in the market and no other asset so that (since the market is the first asset) $\psi_1 = 1$ and $\psi_k = 0$ when $k > 1$. Substituting in the right-hand side of Eq. (8) for g, imposing the equilibrium conditions on ψ, and rearranging terms allows us to write Eq. (9) as

$$\lambda = \gamma \varsigma + \phi J_y y \varrho + \phi \eta \Delta' J_x - \zeta \Lambda' \frac{J_{xy}}{J_y}, \quad (10)$$

where ς and ϱ denote the first columns of the matrices $\zeta \zeta'$ and $\eta \eta'$. The kth element of ς is the covariance between the market (the first asset) and the kth asset. Likewise, the kth element of ϱ represents the "covariance" between the uncertainty in the market and the uncertainty in the kth asset. We consider the specification of ϕ proposed by Maenhout (2004):

$$\phi(x, y) = \frac{\theta}{(1-\gamma) J(x, y)}, \quad (11)$$

where θ is a time-invariant constant. With this specification formula (10) simplifies to

$$\lambda = \gamma \varsigma + \theta \varrho + \theta \eta \Delta' \frac{J_x}{(1-\gamma) J} - \zeta \Lambda' \frac{J_{xy}}{J_y} \quad (12)$$

since $J_y y / [(1-\gamma) J] = 1$. The term

$$\theta \eta \Delta' \frac{J_x}{(1-\gamma) J} - \zeta \Lambda' \frac{J_{xy}}{J_y} \quad (13)$$

comes from the hedging component of optimal portfolios. To simplify the analysis we make sufficient assumptions for this hedging component to be zero. We first assume the noise driving the market is orthogonal to the noise driving the state (so that $\zeta_t \Lambda_t' = 0$ for all t) and second the uncertainty in the market is unrelated to the uncertainty in the state (so that $\eta_t \Delta_t' = 0$ for all t). It follows that

$$\lambda = \gamma \varsigma + \theta \varrho. \tag{14}$$

The assumption that the noise and uncertainty underlying the market and the state are not related is somewhat implausible and is primarily made for convenience given available data. It is one source of possible misspecification that agents are worried about. The reference model, assuming orthogonality as captured by Eq. (14), is a good approximation to the more general model in Eq. (12) when γ is close to one, and both $\zeta_t \Lambda_t'$ and $\eta_t \Delta_t'$ are close to matrices of zeros for all t.

2.2. Discrete time approximation

Although the theoretical model is based on continuous time inputs, in our empirical work we estimate the model in discrete time both for the market (Section 5) and for the cross section of returns (Section 6). To do the estimation, we need to relate the true expected excess return for stocks to λ. We make the assumption that λ is the expected excess return on stocks which is appropriate if the reference model is correct. For the market, we consider a discrete time approximation to Eq. (14) in which the quarterly excess return of the market over a risk-free bond between periods t and $t+1$, denoted r_{et+1}, satisfies

$$E_t r_{et+1} = \gamma V_t + \theta M_t, \tag{15}$$

where $V_t = \varsigma_{1t}$ is the (conditional) variance and $M_t = \varrho_{1t}$ the (conditional) uncertainty of the market. More generally, for any asset k we define

$$\beta_{vk} = \frac{\varsigma_{kt}}{V_t} \qquad \beta_{uk} = \frac{\varrho_{kt}}{M_t} \tag{16}$$

and assume β_{vk} and β_{uk} are constant over time. This assumption implies restrictions on the exogenous processes for the state vector and asset prices and allows us to think of β_{vk} and β_{uk} as respectively being regression coefficients of the risk in asset k on market risk and of the uncertainty in asset k on market uncertainty.[7] We therefore estimate

$$E_t r_{kt+1} = \beta_{vk} \gamma V_t + \beta_{uk} \theta M_t, \tag{17}$$

[7]We make the assumption that the betas are time-invariant because for our estimations in later sections we only have 35 years of quarterly data.

where r_{kt+1} is an excess return and where, similar to the market, we have assumed the reference model is correct. The above equation provides the underpinnings for the empirical cross-sectional analysis covered in Section 6.

In reality, the reference model may indeed not be correct and concerns about misspecification may be justified. As discussed above, the reference model may be misspecified because we assumed that the noise and uncertainty underlying the market and the state are not related which rules out hedging demands. Another reason the reference model may be misspecified is because we have assumed β_{vk} and β_{uk} are time-invariant. There is no economic reason why the betas cannot evolve over time in interesting ways. The reference model also could be misspecified in many other ways.

We deal with the possible misspecification of the reference model in several ways. First, for some of our estimations, we use quasi-maximum likelihood (QMLE) which allows us to obtain consistent estimates in the presence of certain types of misspecification. Second, when estimating (15) and (17) we include additional constant terms. If there is non-zero unconditional mean misspecification then the constant terms would be significant. Third, in some of our specifications we allow uncertainty to affect the quarterly conditional volatility of asset returns. Although uncertainty should not affect volatility in our continuous time model, it is plausible that uncertainty affects quarterly volatility because model misspecification might appear as additional noise at the quarterly frequency.[8]

3. Measuring risk with volatility

In the next two sections, we describe the empirical implementation of the model presented in Section 2. In this section, we deal with estimating risk, which is fairly standard, while Section 4 covers the estimation of uncertainty. A key issue is how we empirically distinguish risk and uncertainty, which is discussed below.

3.1. Empirically distinguishing risk from uncertainty

A key contribution of our work is how we empirically identify risk with volatility and uncertainty with the disagreement among professional forecasters, which proxies for the confidence of agents in their beliefs about mean returns. Following the work of Merton (1980) and Foster and Nelson (1996), it is well known that precise estimation of volatility can be achieved through sampling returns over arbitrarily short time intervals. It is therefore not unreasonable to assume that volatility is known, although in practice prominent high-frequency data characteristics such as leptokurtosis, intra-daily deterministic patterns, and

[8]Recall that even if the reference model is false, it is by assumption a good description of reality so that the additional constant terms and the additional noise should be small in magnitude.

market microstructure features such as price discreteness, nonsynchronous trading, and bid-ask spread further contaminate the data used in empirical research. Merton also shows that, in contrast to volatility, estimation of the drift component only depends on the span of the data, not the sampling frequency. Only a longer span of data yields more precise estimation. In practice, long data spans uncontaminated by some sort of structural breaks are next to impossible to find. Hence, the estimation of the drift component remains extremely difficult. We therefore take the view that future asset returns (and future values of other state variables) are *risky* because they might deviate from their conditional means, and they are *uncertain* because their true conditional means are not known. Under this view, uncertainty is limited to uncertainty in first moments and everything about higher order central moments is assumed to be perfectly known.

Consider the following example: assume agents believe that the excess return on the market is distributed normal with some unknown mean μ and known standard deviation σ : $r_e \sim N\left(\mu, \sigma^2\right)$. In this situation, we view σ^2 as market *risk* whereas *uncertainty* in the market is the agent's beliefs about the variance of μ. In other words, uncertainty is a measure of the confidence of an agent in her beliefs about μ and can be thought of as an approximation to the mean-squared error of her beliefs, $E(\mu - \hat{\mu})^2$, where $\hat{\mu}$ is her best approximation of μ.

We do not take a hard-line view on the separation of risk and uncertainty when it comes to empirical identification. Risk is empirically identified as asset return volatility (see next section for details). In our empirical applications for uncertainty, it does not matter if agents know the distribution of the drift μ. If they know the distribution of μ we think of the variance of μ as uncertainty. If agents do not know the distribution of μ then we think of agents as measuring uncertainty with their best approximation to the variance.[9] Regardless of whether agents know the distribution of μ, we assume they choose to treat the uncertainty in μ differently from the variance of returns because they may be more (or less) averse to situations in which they have little confidence in μ than to situations in which the variance of returns is large. As argued by Hansen and Sargent, a reasonable strategy when facing uncertain distributions is to solve a version of the robust control problem described in Section 2.1.

3.2. Estimating conditional volatility

Volatility is a prevailing feature of financial markets. Since volatility displays strong persistence it is predictable, which has led to a large literature on volatility risk models and many ways to estimate volatility. When we confine our attention to models exclusively based on returns data, a natural choice

[9]Because of our practical views, some of what we call uncertainty may indeed be risk as defined by Knight (1921) and Keynes (1937). However, it seems reasonable to us that the uncertainty in μ is of an order of magnitude larger than the risk in μ and that from a practical perspective, calling everything that is unknown about the true value of μ uncertainty is a reasonable approximation.

is autoregressive conditional heteroskedasticity (ARCH)-type models. Our models are estimated at the quarterly frequency, implying quarterly ARCH models; however, this is rather unappealing as volatility would be estimated quite imprecisely. Instead, we adopt an approach which allows us to estimate volatility at a quarterly frequency more precisely by exploiting daily returns data. In recent work, Ghysels, Santa-Clara, and Valkanov (2005) suggest that volatility can be modeled with a *Mixed Data Sampling*, or MIDAS, approach. The key to modeling conditional variances is parsimony, summarizing in a convenient way the temporal dynamics which yield predictions of future volatility.[10] In this section, we review one parsimonious flexible functional form for measuring the conditional volatility of the market, which we call (market) risk.

Ghysels, Santa-Clara, and Valkanov (2006) and Ghysels, Sinko, and Valkanov (2007) suggest that a discretized beta distribution is a flexible functional form that can conveniently capture many plausible patterns of time-series decay. The discretization is based on the standard continuous beta probability density function which is

$$w[x] = \frac{(x-a)^{\alpha-1}(d-x)^{\chi-1}}{B(\alpha,\chi)(d-a)^{\alpha+\chi-1}}, \tag{18}$$

where B is the beta function and α, χ, a, and d are parameters. The discretized beta distribution we use is

$$w_i = \frac{(i-a)^{\alpha-1}(d-i)^{\chi-1}}{B(\alpha,\chi)(d-a)^{\alpha+\chi-1}}\left[\sum_{j=1}^{n}\frac{(j-a)^{\alpha-1}(d-j)^{\chi-1}}{B(\alpha,\chi)(d-a)^{\alpha+\chi-1}}\right]^{-1} \tag{19}$$

$$= \frac{(i-a)^{\alpha-1}(d-i)^{\chi-1}}{\sum_{j=1}^{n}(j-a)^{\alpha-1}(d-j)^{\chi-1}}, \tag{20}$$

with n values ($i = 1, 2, \ldots n$) that receive positive probability. We require $a \leq 1$ and $d \geq n$. Note that a potentially large set of weights is tightly parameterized via a small set of parameters. Ghysels, Santa-Clara, and Valkanov (2006) and Ghysels, Sinko, and Valkanov (2007) discuss how the discretized beta distribution captures many different weighting schemes associated with time-series memory decay patterns observed in volatility dynamics and other persistent time-series processes by varying parameters. They also observe that setting $\alpha = 1$ yields downward sloping weighting schemes typically found in models of volatility predictions. By construction, Eq. (20) is a well-formed probability density function since $\sum_{i=1}^{n} w_i = 1$, and we interpret the w_i's as weights. This convenient scheme is used in our empirical work, both in a time-series context to parameterize risk and, as we describe in the next section, in a cross-sectional setting to specify uncertainty.

[10] We focus exclusively on MIDAS volatility forecasts as it allows us to use high-frequency data to forecast directly a low-frequency volatility realization. The benefits of this have been shown in various recent applications, including Ghysels, Santa-Clara, and Valkanov (2006), Alper, Fendoglu, and Saltoglu (2008), and Becker, Clements, and White (2007), among others.

To measure risk we construct a measure of conditional variance, which is obtained by weighting prior daily squared (demeaned) returns. More specifically, the weight on the ith prior lag is

$$l_i(\omega) = \frac{(s+1-i)^{\omega-1}}{\sum_{j=1}^{s}(s+1-j)^{\omega-1}},$$

where s is the maximum number of lags.[11] The functional form of these weights is determined by a discretized beta distribution with $\alpha = 1$, $\chi = \omega$, and $d = s+1$. The value of a does not matter since $\alpha = 1$. The single free parameter ω models the decay pattern of the weight function and the top plot in Fig. 1 provides an example of the weights.[12] The resulting conditional variance is equal to

$$V_t = \sigma^2 \mathrm{vol}_t(\omega), \tag{21}$$

where σ^2 is a time-invariant constant, and

$$\mathrm{vol}_t(\omega) = s\sum_{i=1}^{s} l_i(\omega)\left(r_{et,i} - \frac{1}{s}\sum_{j=1}^{s} r_{et,j}\right)^2 + \\ 2s\sum_{i=1}^{s-1}\sqrt{l_i(\omega)l_{i+1}(\omega)}\left(r_{et,i} - \frac{1}{s}\sum_{j=1}^{s} r_{et,j}\right)\left(r_{et,i+1} - \frac{1}{s}\sum_{j=1}^{s} r_{et,j}\right) \tag{22}$$

is the component of the conditional variance which is determined from the volatility of daily excess returns. Here $r_{et,1}$ is the market excess return on the last trading day of quarter t, $r_{et,2}$ is the market excess return on the second to last trading day of quarter t, and $r_{et,i}$ is the market excess return on the ith to last trading day of quarter t. (Since s corresponds to roughly a year, the ith to last trading day of quarter t will occur in quarters $t-1$, $t-2$, or $t-3$ when i is sufficiently large.) As we go backward in time, i increases. Note that vol_t depends on the parameter ω since the weights l_i depend on ω. The second component of vol_t allows for the effect of serial correlation in daily returns on quarterly volatility. Such a correction did not appear in the original formulation of MIDAS volatility estimators.

In Section 5 we estimate ω and σ^2 from data on excess market returns. Ghysels, Santa-Clara, and Valkanov (2005) argue that σ^2 should equal one and they fix it at one for their results. We estimate σ^2 to allow for this part of the model to be misspecified. In Section 6 we estimate ω using information from the cross section.

[11] The value of s determines how many daily lags are used to predict future volatility. We set s to be roughly the number of trading days in a year. Since the number of trading days per year varies slightly throughout our sample and we prefer s be constant for all dates, we set s to be the minimum number of trading days in the previous 12 months available throughout our sample.

[12] The weights displayed in Fig. 1 are obtained from the empirical estimates, discussed in Section 4.

4. Measuring uncertainty with disagreement

The data used in constructing our uncertainty measure are forecasts on macroeconomic and financial variables from the Survey of Professional Forecasters (henceforth SPF).[13] The SPF is an attractive survey for this measurement because it provides a long time series of data (the data set begins in 1968) and it provides forecasts at many different horizons. Each quarter participants are asked for forecasts of the levels of variables for the previous quarter, this quarter, and the next four quarters.[14] The forecasters selected for the SPF come primarily from large financial institutions. The series we use from the SPF are forecasts of output (forecasts of Gross National Product, GNP, before 1992Q1 and Gross Domestic Product, GDP, after), the output deflator (forecasts of the GNP deflator before 1992Q1 and the GDP deflator thereafter), and corporate profits after taxes.

The number of forecasters participating in the SPF varies through time. The average (median) number of forecasts between 1968 and 2003 is 39.5 (36). In the early years, the number occasionally increased to greater than 100 forecasters, but it began to decline nearly monotonically throughout the 1970s and 1980s. After the Federal Reserve Bank of Philadelphia (FRB-Philadelphia) took over the SPF in 1990, the average (median) number of forecasters each quarter is 36 (35), with a low of 29 and a high of 52.[15] Not all forecasts are usable because some are incomplete. At a given date, the predictions of a forecaster are included only if she provides predictions of corporate profits after taxes, output and the output deflator for the current quarter, the next quarter, two quarters ahead, and three quarters ahead. Across all dates, we are able to use a median of 26 forecasts, a minimum of nine, and a maximum of 74.

4.1. Imputing asset return forecasts

We now discuss how we impute forecasts of the real market return from forecasts of nominal corporate profits and the price level by using the Gordon growth model. In Appendix A we describe how we compute forecasts of the expected real return on a nominally risk-free bond from the nominal risk-free rate and forecasts of the price level.

---

[13]See the Web-page *http://www.phil.frb.org/econ/spf/index.html* and the comprehensive overview in Croushore (1993) for more information about the survey. There are many other papers that make use of data from the SPF, most of which evaluate the quality of the forecasts (see, for example, Zarnowitz, 1985; and Braun and Zarnowitz, 1993).

[14]Data on forecasts four quarters ahead are sparse in the initial years of the survey. Data on forecasts for the previous quarter are included because the actual final values for last quarter may not be known perfectly at the time the survey is published. The survey also includes annual and longer horizon forecasts. Despite the fact that surveys of professional forecasters are more reliable than other surveys, Ghysels and Wright (2009) report that responses in the SPF data appear to have some evidence of staleness with respect to the reference reporting date of the survey.

[15]In the second and third quarters of 1990, there are extremely low numbers of respondents, corresponding to the transfer of the survey from the American Statistical Association (ASA) / National Bureau of Economic Research (NBER) to the FRB-Philadelphia. To avoid having a missing data point, they included a 1990Q2 survey with the 1990Q3 survey. The total number of respondents was nine.

The Gordon growth model (or dividend discount model) is a widely used method of stock valuation linking the current stock price, the current level of the dividend, the expected growth rate of dividends, and the capitalization rate.[16] We use aggregate corporate profit forecasts rather than individual stock earnings forecasts as inputs to the Gordon growth model. Let π_{t+1} be forecasts of aggregate real corporate profits and q_t be the actual (not forecasted) market value of all domestic corporations at time t.[17] For us, the Gordon growth model amounts to assuming that forecaster i's constructed prediction of the gross return on the market is

$$E_{it} r_{mt+1} = E_{it}\left[\frac{\pi_{t+1}}{q_t}\right] + \xi_{it}, \tag{23}$$

where ξ_{it} is forecaster i's predicted gross growth rate of real corporate profits over a long horizon. In Appendix A we describe how we construct approximate forecasts of the real level of corporate profits and the growth rate of real corporate profits.

Table 1 shows that for the market, the Gordon growth model gives a reasonable approximation of the unconditional mean return. For the period between 1968 and 2003 the average of the median forecasts of the gross market return computed from the Gordon growth model is 1.0230 (with ξ_t being the forecasted average return from the last period to *three* quarters ahead—a horizon of four), which slightly overestimates the actual average real market return, 1.0168, but is statistically close to it (using the second column which gives the standard deviation of the actual market return).[18] Table 1 also shows that our constructed average median forecasts of a real risk-free rate are qualitatively similar to the actual average real return on a nominally risk-free bond.[19]

4.2. Uncertainty and disagreement

As discussed in Section 2, the conditional expected excess return of assets in the agents' reference model is approximately λ_t; however, it may be misspecified and the expected excess return could be λ_t −

[16]Wiese (1930) and Williams (1938) were among the first to apply present value theory to common stocks; however, their models suffered from the assumptions about the magnitude and timing of dividend payouts. Gordon (1962) popularized the model by assuming a constant growth rate of dividends into the future and a terminal price for the holding period. Anderson, Ghysels, and Juergens (2005), Brav, Lehavy, and Michaely (2005), Crombez (2001), Gebhardt, Lee, and Swaminathan (2001), and Guay, Kothari, and Shu (2003), among others, have utilized short-term earnings and long-term earnings growth forecasts of investment analysts as inputs to the Gordon growth model. Jagannathan, McGrattan, and Scherbina (1996) have used variations of the Gordon growth model, related to Campbell and Shiller (1988), in resolving the equity premium puzzle.

[17]Ideally, we would like forecasts of corporate profits without any seasonal adjustment, but in the SPF forecasters are asked to predict deseasonalized corporate profits.

[18]If one looks at a longer horizon of actual data than is reported in Table 1 the approximation is much better. The actual average market return between the second quarter of 1926 and the third quarter of 2006 is 1.0227, a difference of only 0.0003 from the average median forecast. This might be viewed as evidence that the Gordon growth model is approximating very well the large sample level of the market return and suggests that forecasters might be looking at a long horizon (and using past information) when making forecasts. The behavior of forecasters is consistent with the implication of Merton's model that the longest possible horizon should be used when estimating mean returns.

[19]The exact formula used to impute forecasts of a real risk-free bond appears in Appendix A in Eq. (45).

$\eta_t g_t$. The matrix $\eta_t\eta_t'$ represents how confident agents are in their beliefs about the expected returns. The $(1,1)$ element of $\eta_t\eta_t'$ is a measure of market uncertainty and the other elements in the first column of $\eta_t\eta_t'$ represent the covariance of the uncertainty in other assets with market uncertainty. In this section we present an overview of a method for measuring the first column of the matrix $\eta_t\eta_t'$. As discussed in Section 2, this is the column that affects equilibrium returns.

In our empirical approach we adopt the view that the $(1,1)$ element of the matrix $\eta_t\eta_t'$ (agents' beliefs about the mean-squared error of the expected market excess return λ_{1t}) can be proxied by a weighted variance of the predictions of the market return stated by professional forecasters.[20] Under this perspective, if all forecasters are in agreement, then market uncertainty is likely to be small. Consequently, the first element of the optimal endogenous perturbation $\eta_t g_t$ will be small so that only outcomes relatively close to the reference model matter to the agents in our economy. In contrast, if forecasters state very different forecasts, then agents are unsure that their reference model is correct and the first element of $\eta_t g_t$ will be large so that outcomes relatively far from the reference model matter to the agents in our economy.

To further illustrate, consider a situation where agents have very little confidence in the reference model and acquire the predictions of n forecasters. Let the vector y be the collection of all forecasts. Agents believe that approximately

$$y = \mathbf{1}\mu + v, \tag{24}$$

where μ is the true mean of the market return, $\mathbf{1}$ is an n dimensional vector of ones, v is a normally distributed vector with mean zero and covariance matrix $\sigma^2 R$. Assume agents know R perfectly but have no prior information about the value of σ^2.[21] They estimate μ and σ^2 from the predictions of forecasters using maximum likelihood, neglecting any prior information they might have about μ and σ^2. Maximum likelihood estimates are a weighted mean and a weighted variance

$$\hat{\mu} = \left(\mathbf{1}'R^{-1}\mathbf{1}\right)^{-1}\mathbf{1}'R^{-1}y$$
$$\hat{\sigma}^2 = \frac{1}{n}\left(y - \mathbf{1}\hat{\mu}\right)' R^{-1}\left(y - \mathbf{1}\hat{\mu}\right)$$

of forecasts. The value of $\hat{\sigma}^2/n$ can be taken to be an approximation of the amount of uncertainty of agents.[22]

[20] In our formulation, there is a separation between forecasters and agents. Agents are all alike and assign the same approximate probabilities while forecasters are heterogeneous.

[21] It is important to emphasize that this is only an approximation from the point of view of agents. The vector v may not be normally distributed and agents might not know the value of R perfectly.

[22] When computing uncertainty in later sections we do not divide analogs of $\hat{\sigma}^2$ by the number of forecasters, n, because the number of forecasters agents pay attention to is likely different from the number of forecasters in the SPF.

To measure uncertainty, we use a parsimonious flexible functional form that determines the weights across forecasts as we have very little information about the actual mechanisms by which agents compute uncertainty. Hence, we adopt a reduced-form approximation rather than a fully specified structural approach, which is impossible to implement with the data and information we have at our disposal. To capture uncertainty we need to apply weights cross sectionally across different forecasts. We again use the beta specification as it is parsimonious, i.e., with only two parameters we can determine which part of the distribution of predictions matter for asset pricing. The beta weighting scheme adapts easily to a cross-sectional application among forecasters because setting $\alpha = \chi$ (and hence requiring an even smaller set of parameters) yields various bell-shaped weighting schemes. Ghysels, Santa-Clara, and Valkanov (2006) and Ghysels, Sinko, and Valkanov (2007) give detailed discussions of the advantages of using flexible functional forms for capturing volatility. In Appendix C we give a brief example that illustrates one benefit of using flexible functional forms for computing cross-sectional variances. One problem we face when estimating cross-sectional variances is that a few outlying observations can have a large effect on estimates. In later sections, our weighting schemes allow us to optimally ignore extreme forecasts.

We construct a measure of the cross-sectional variance that allows for the possibility of ignoring extreme forecasts. We pick one series, call it x, and rank the forecasts each period of x from low to high. (For us, x is usually the forecast of the market return.) The weight on the ith lowest forecast is

$$w_{it}(\nu) = \frac{i^{\nu-1}\,(f_t + 1 - i)^{\nu-1}}{\sum_{j=1}^{f_t} j^{\nu-1}\,(f_t + 1 - j)^{\nu-1}},$$

where f_t forecasts are available at time t and ν is a parameter. This is the discretized beta distribution described in Section 3 with $\alpha = \nu$, $\chi = \nu$, $a = 0$, and $d = f_t + 1$. Instead of letting the first power parameter equal one (as when we computed conditional volatility), and letting the second power parameter, ω, determine the decay pattern, we set both power parameters of the beta distribution equal to each other and estimate the single common parameter as a free parameter ν. This specification forces the weights to be symmetric.[23] The disagreement or uncertainty is then measured by a beta-weighted variance of forecasts of x :

$$\text{unc}_t(\nu) = \sum_{i=1}^{f_t} w_{it}(\nu) \left[x_{it+1|t} - \sum_{j=1}^{f_t} w_{jt}(\nu) x_{jt+1|t} \right]^2 . \tag{25}$$

The uncertainty component is thus constructed as a weighted variance of predictions on a single financial/macroeconomic variable where the weights are determined by a discretized beta distribution. In Section 5 we estimate ν from quarterly data on excess market returns and in Section 6 we estimate ν using

[23]See, for instance, the bottom plot in Fig. 1 for an example of the weights.

information from the cross section. Our estimates typically entail ignoring the extremes and placing all the weight on the center of the distribution. The flexible weights we use may cause us to underestimate the true cross-sectional variance since forecasts that are far from the median receive little weight, even if they happen to be valid forecasts. For our purposes this concern is somewhat mitigated because, as described in the next section, in addition to estimating parameters determining flexible weights, we estimate a parameter that scales cross-sectional variances.

While the above may be a plausible behavioral description of how agents compute uncertainty, one may wonder whether it is optimal. In particular, is the uncertainty of a fully rational agent necessarily linked to disagreement? The answer depends crucially on how beliefs, information, and models are distributed among forecasters. In Appendix B we describe an environment in which there is a direct link. In this environment, uncertainty is always proportional to disagreement and it is reasonable to view M_t as equaling a time-invariant constant times uncertainty. Although the environment we describe makes reasonable assumptions, in reality uncertainty probably is not always proportional to disagreement. To the extent disagreement approximates the amount of uncertainty in the economy, the approach taken in this paper is reasonable. In this paper we refer to M_t as uncertainty, but it is important to remember that M_t is at best approximately proportional to the amount of uncertainty in the economy.[24]

How should the other elements of the first column of $\eta_t\eta_t'$ be measured? The other elements represent the covariances of the uncertainty in other assets with market uncertainty. Data on the covariances of disagreement across stocks are difficult to obtain. Consequently we devise a method for computing the other elements of the first column of $\eta_t\eta_t'$ without actually observing agents' beliefs about other stocks. The method exploits the fact that the model in Section 2 entails that the covariance of the uncertainty in any asset with the market should affect expected returns, and is described in detail in Section 6.

Our method of measuring uncertainty does not attempt to identify g_t directly with the data. Instead, as sketched above, we measure the first column of $\eta_t\eta_t'$ and let the model described in Section 2 determine the endogenous worst case, summarized by g_t, about which agents worry.

5. The impact of risk and uncertainty on the market return—empirical results

In this section, we estimate the amount of market risk and market uncertainty in the economy and investigate the relative importance of risk and uncertainty for the expected market return. The estimates

[24] In an ideal world, we would not only have mean forecasts across forecasters for each stock, but we would also have some measure of dispersion of each forecasters' beliefs about the mean. Graham and Harvey (2003), using survey data from Chief Financial Officers (CFOs) about the expected risk premium, are able to obtain a distribution of beliefs for each individual respondent. In the future, work along the lines of Graham and Harvey (2003) may eventually yield better measures of uncertainty.

allow us to construct an index which measures the amount of uncertainty in the economy.

We consider again a version of the decomposition suggested in Section 2:

$$E_t(r_{et+1}) = b + \gamma V_t + \theta M_t, \tag{26}$$

where we also include a constant term b, even though according to the reference model this term should be zero. The constant is added to allow for the possibility that the reference model is misspecified. Including a constant also guarantees that the empirical regressions have well-behaved residuals, and as discussed later, testing the statistical significance of b can be used as a model validation test. We assume that market excess returns are normally distributed with time-varying volatility:

$$r_{et+1} \sim N\left[b + \gamma V_t + \theta M_t, V_t\right]. \tag{27}$$

The assumption of normality is made only for the purpose of estimation, yielding a quasi-maximum likelihood (QMLE) setting. As discussed in Sections 3 and 4 we measure risk, V_t, with $\sigma^2 \mathrm{vol}_t(\omega)$ and uncertainty, M_t, with $\mathrm{unc}_t(\nu)$. To estimate the parameters b, τ, θ, ω, and ν, we maximize the (quasi-)likelihood of quarterly excess returns based on:

$$r_{et+1} \sim N\left[b + \tau \mathrm{vol}_t(\omega) + \theta\, \mathrm{unc}_t(\nu), \sigma^2 \mathrm{vol}_t(\omega)\right], \tag{28}$$

where $\tau = \gamma\sigma^2$.

It is important to remember that the agents inside our model are worried that the reference model is false, and if their worries are justified, our empirical regressions are misspecified as well. The reference model may be misspecified because we have ignored the hedging component and excess returns may not really be conditionally normally distributed. In addition to including the constant b we take into account the possibility that (27) is misspecified in several other ways; we report quasi-maximum likelihood standard errors; and we include additional constant terms [not present in Ghysels, Santa-Clara, and Valkanov (2005)] that potentially could pick up some aspects of model misspecification in V_t. Finally, we estimate an alternative specification in which quarterly volatility partly depends on the amount of uncertainty in the economy.

5.1. The evidence for risk-return and uncertainty-return trade-offs

We begin by investigating if there is a risk-return trade-off in our data set in the absence of uncertainty. The combination of quarterly returns and daily returns yields the MIDAS setup. Quasi-likelihood estimates of the parameters in Eq. (28) which determine risk are displayed in the first three estimations in Table 2. We see that according to a t-test and a likelihood ratio test, estimates of τ are not significant, though estimates of $\log \omega$ are highly significant.[25] The results suggest that, in our data set, MIDAS does provide a better measure of conditional volatility than current realized volatility, although there is no evidence of a risk-return trade-off.[26] Further evidence that MIDAS captures volatility can be provided by examining the relationship between vol_t and realized volatility. We define realized volatility as

$$Q_t = z_t \sum_{i=1}^{z_t} \left(r_{et,i} - \frac{1}{z_t} \sum_{j=1}^{z_t} r_{et,j} \right)^2 + 2z_t \sum_{i=1}^{z_t - 1} \left(r_{et,i} - \frac{1}{z_t} \sum_{j=1}^{z_t} r_{et,j} \right) \left(r_{et,i+1} - \frac{1}{z_t} \sum_{j=1}^{z_t} r_{et,j} \right), \tag{29}$$

where z_t is the number of days in a quarter t.[27] Table 3 shows that future realized volatility (Q_{t+1}) is more highly correlated with vol_t than it is with current realized volatility (Q_t). This confirms that vol_t does provide a better measure of conditional volatility than current realized volatility.

Our implementation differs from the implementation in Ghysels, Santa-Clara, and Valkanov (2005) in that we estimate σ^2, rather than fix it at one. If the reference model is correctly specified then σ^2 should equal one since we designed vol_t to be the conditional variance of the market and our model says that $\gamma \mathrm{vol}_t$ should be the conditional mean of the market excess return, in the absence of uncertainty. However, if the reference model is misspecified then σ^2 need not equal one. We find estimates of σ^2 are significantly greater than one in models that perform poorly but are close to one in models that perform well. This provides further evidence that the poorly performing models are misspecified. Another diagnostic test of the model is if Jensen's time series alpha, denoted as b in the table, is close to zero. We find Jensen's alpha is not significantly different from zero in Table 2, supporting the reference model. The standard errors are computed with QMLE robust standard errors.

We now investigate whether an uncertainty-return trade-off exists. Quasi-likelihood estimates of the parameters appearing in Eq. (28) are displayed in Table 2 where unc_t is a beta-weighted variance of market return forecasts. In the fourth specification we include uncertainty but measure it with an unweighted (or flat weighted) variance, obtained by setting $\log \nu = 0$. In this case the estimate of θ is not significant and there is very little improvement to the log likelihood without uncertainty. Thus, including uncertainty with flat weights does not improve much upon specifications in which uncertainty is left out. In the 5th, 6th, and 7th estimations we estimate θ and $\log \nu$ along with other parameters. In these regressions, unc_t is allowed to be a non-degenerate beta-weighted variance. Estimates of θ and $\log \nu$ are significant (by likelihood ratio

[25]By a likelihood ratio (LR) test for specifications one and two, the parameter $\log \omega$ is significant, with a p-value of 0.006. By an LR test for specifications two and three, the parameter τ is not significant with a p-value of 0.702. The null hypotheses in both cases are that the parameters are zero.

[26]Appendix D discusses why our results differ from Ghysels, Santa-Clara, and Valkanov (2005), who find evidence of a risk-return trade-off using MIDAS.

[27]Note that Q_t is similar to vol_t except the weights are uniform and z_t does not necessary equal s. Since Q_t is realized volatility within a quarter, z_t corresponds to the number of days in a quarter.

tests and t-tests) and there is a large improvement to the log-likelihood.[28]

It is also interesting to note that estimates of σ^2 are not significantly different from one and estimates of the constant b are not significantly different from zero. Both of these results confirm the predictions of the reference model. Further informal evidence for an uncertainty-return trade-off is provided in Panel B of Table 3. In particular, the correlation between our estimated measure of uncertainty in the last regression and the excess return is 0.28. In comparison the correlation of our measure of risk with the excess return is only 0.15.

5.2. An index of uncertainty

It is common to talk about volatility, to examine plots of volatility estimates, and to interpret them. Such plots are often examined with the idea that risk is high during certain times. A by-product of our empirical analysis is that we can do a similar analysis with uncertainty. Therefore, we now discuss the empirical properties of estimated uncertainty. We let the index of uncertainty be the series unc_t [exp(2.731)] estimated in the last regression in Table 2. We plot the index in Fig. 2 along with plots of excess returns and volatility. We provide some simple statistics in Table 3 and estimates of autoregressions in Table 9. Table 3 also presents correlations of unc_t with the Fama-French factors. As noted above, the correlation between our estimated measure of uncertainty and the market excess return is 0.28. We also find that unc_t is significantly positively correlated with the Small Minus Big (SMB) factor (0.24), but insignificantly negatively correlated to the High Minus Low (HML) factor (-0.07).

It is well known that volatility is highly persistent. In our data, Table 9 shows that quarterly volatility is positively and significantly related to its first three lags, that is three quarters. Uncertainty is also persistent but not as much as volatility. Uncertainty is positively and significantly related to its first two lags, or half a year.

Panel A of Table 9 shows there is not a significant relationship between uncertainty and lagged volatility. We see from Table 3, there is very little contemporaneous correlation between uncertainty and volatility which suggests that the actual conditional variance (past volatility) has some, but not much, impact on the beliefs of agents about uncertainty. Uncertainty is not highly correlated with future volatility. Past uncertainty does not predict future volatility and vice versa. Hence, volatility and uncertainty appear as nearly orthogonal processes. From Table 3 we see that unweighted uncertainty is slightly more related to future volatility than optimally weighted uncertainty. Perhaps the fringes of forecasts matter for volatility (i.e.,

[28]By a likelihood ratio (LR) test for specifications three and seven, the parameters θ and $\log \nu$ are jointly significant, with a p-value of 0.01. By an LR test for specifications four and seven, the parameter $\log \nu$ is significant, with a p-value of 0.002. However, when $\log \nu$ is fixed at zero, by an LR test in specifications three and four, the parameter θ is not significant, with a p-value of 0.893. The null hypotheses in all cases are that the parameters are zero.

maybe agents are heterogeneous and extreme forecasts correspond to the beliefs of noise traders) but not for expected returns. However, this effect is not strong.

In Fig. 3 we graph uncertainty with several different events. We see that uncertainty is often large just before the onset and just before the end of a recession. Whenever uncertainty has been unusually large, the market excess return in the following quarter has also been large. Two of the lowest readings of uncertainty occurred when incumbent presidents were reelected (Nixon in 1972 and Clinton in 1996).

5.3. Robustness of empirical results with alternative specifications

The results so far are important as they tell us that the our empirical measure of uncertainty carries a substantial premium and is a more important determinant of the market excess return than risk. In this subsection we verify the robustness of this finding, to scrutinize it further. Yet, at this point it is perhaps also important to realize that much of the literature has focused on risk, while our empirical findings suggest it might be more important to think about uncertainty in empirical asset pricing.

5.3.1. Uncertainty about output and corporate profits

In Table 4 we consider the uncertainty-return trade-off when uncertainty is measured by the beta-weighted variance in variables other than the market return forecasts. One could make an argument that the dispersion of alternative variables should affect uncertainty and thus excess returns. For example, uncertainty in future output could reflect underlying structural uncertainty in the economy that perhaps should be priced. In Panel A, we consider the uncertainty in constructed real output growth forecasts, while in Panel B we consider the uncertainty in growth rate of corporate profits at many different horizons. The uncertainty in real output forecasts does not have a significant effect on excess returns. At long horizons (three and four) the uncertainty in corporate profit forecasts does have a significant effect, but this effect is mitigated at shorter horizons. For corporate profits model uncertainty measures, the Jensen alpha time-series estimate is significant at the shortest horizon, while it is insignificant at longer horizons. Since our market return forecasts are constructed from a combination of short-term and long-term corporate profit forecasts, the results in Panel B suggest that the underlying driving force for our earlier results comes from long-term corporate profit forecasts and not short-term forecasts.

5.3.2. Normal weights

We now briefly consider some alternative specifications to investigate if our results crucially depend on measuring uncertainty with a symmetric beta distribution. In Panel A of Table 5 we measure uncertainty with a normal weighted variance in which the weights are

$$w_{it}(\xi) = \frac{\exp\left(\frac{-\left(i-\frac{f_t+1}{2}\right)^2}{\xi^2}\right)}{\sum_{j=1}^{f_t} \exp\left(\frac{-\left(j-\frac{f_t+1}{2}\right)^2}{\xi^2}\right)}, \tag{30}$$

where ξ is a parameter. The results in Panel A are very similar to the results in specifications six and seven of Table 2. The coefficients on uncertainty are virtually identical. The estimated normal weights place positive weights on the same parts the distribution of forecasts that the beta weights do. There is strong evidence for an uncertainty-return trade-off even with a different specification of the cross-section weights. The estimates of Jensen's alpha also remain insignificant in all cases.

5.3.3. *Non-symmetric weights*

In Panel B of Table 5 we measure uncertainty with weights that are not restricted to be symmetric. We consider non-symmetric weights because forecasters might have significantly different beliefs than agents. The bias of forecasters might lead agents to be concerned about the uncertainty in pessimistic (or alternatively optimistic) forecasts. To allow for the possibility of bias, we use beta weights in which the weights are

$$w_{it}(\alpha, \chi) = \frac{i^{\alpha-1}(f_t + 1 - i)^{\chi-1}}{\sum_{j=1}^{n} j^{\alpha-1}(f_t + 1 - j)^{\chi-1}}, \tag{31}$$

where α and χ are free parameters. This is the discretized beta distribution described in Section 3 with $a = 0$ and $d = f_t+1$. Allowing the weights to be non-symmetric lets the agents' perceived uncertainty depend on any part of the distribution of forecasts. If agents pay more attention to the variance in worst-case forecasts then χ should be greater than α; alternatively, if agents focus on optimistic forecasts then χ should be less than α. We find in Table 5 that the estimates of α and χ are not significantly different from each other. Since the estimated value of χ is slightly greater than α, the estimated weights slightly emphasize the variance of pessimistic forecasts over optimistic forecasts; however, there is not compelling evidence to suggest that non-symmetric weights more precisely measure perceived uncertainty than symmetric weights.[29]

5.3.4. *Fixed weights*

Results for several different types of fixed weighting schemes are displayed in Table 6. In Panel A we fix, rather than estimate, $\log \nu$ at many different values and find that uncertainty has a significant effect

[29]This is not a test of whether or not agents are worried about worst-case outcomes. According to our approach agents choose the worst case based on their perceived amount of uncertainty (see Eq. (8)). The issue in this paragraph is simply when computing perceived uncertainty, should the weights across forecasts be symmetric or should they place more weight on the low or high end of the distribution.

on market excess returns when $\log \nu$ is not small. In Panel B, uncertainty is measured with a truncated variance where the lowest p percent of forecasts and the highest p percent of forecasts are discarded each quarter. We see that as long as p is not small, uncertainty has a significant effect on market excess returns. In Panel C, we Winsorize forecasts each quarter by replacing the lowest p percent of forecasts and the highest p percent of forecasts with lowest and highest forecasts in the middle $1 - 2p$ percent of forecasts. We again see that as long as p is not small, uncertainty has a significant effect on market excess returns.[30] We also see that the likelihoods using the best settings of p with truncated variances or Winsorization are virtually identical to the likelihoods with the optimal setting of $\log \nu$. Because fewer parameters are being estimated the standard errors for the uncertainty-return trade-off can appear to be much smaller when $\log \nu$ or p is fixed. One advantage of our weighting scheme is that it allows a researcher to estimate the weights and it provides a truer picture of standard errors. We conclude that our results are robust to alternative fixed weight specifications as long as extreme forecasts are down-weighted, removed, or replaced.

5.3.5. Long-term horizon of three quarters

In Panel C of Table 5, we measure uncertainty by the beta-weighted variance of constructed market return forecasts when the long-term horizon is three quarters rather than four quarters. Setting the long-term horizon at three does better than setting the long-term horizon at four.[31] The Gordon growth model requires a long-term horizon forecast and it is most natural to let the long-term horizon be four because that is the longest horizon for which data are plentiful. It is slightly puzzling that a horizon of three performs better empirically than a horizon of four. Perhaps a horizon of four is too far ahead for forecasters to accurately report their beliefs. In this paper, we choose to emphasize a horizon of four rather than three but it is important to note that our results become stronger if we use a horizon of three. Finally, it is also worth noting again that the estimates of Jensen's alpha remain insignificant in all cases.

5.3.6. Allowing uncertainty to affect volatility

To allow for additional ways in which the reference model could be misspecified we estimate a specification in which

$$r_{et+1} \sim N\left[b + \tau \mathrm{vol}_t(\omega) + \theta\, \mathrm{unc}_t(\nu),\ \sigma_v^2 \mathrm{vol}_t(\omega) + \sigma_u^2 \mathrm{unc}_t(\nu)\right], \tag{32}$$

where we restrict $\sigma_v \geq 0$ and $\sigma_u \geq 0$. The term $\sigma_u^2 \mathrm{unc}_t(\nu)$ in the variance of r_{et+1} embodies the notion that if the reference model is false then at the quarterly frequency the expected excess return should have

[30]When $\log \nu$ is extremely large and p is close to 50, the significance can start to break down because only the very middle part of the distribution of forecasts is considered. For example, when p is 40, only the middle 20% of forecasts are used to compute uncertainty.

[31]Given our results in Panel B of Table 4 this is perhaps not surprising. Corporate profit growth forecasts at a horizon of three are more related to excess returns than corporate profit growth forecasts at a horizon of four.

a higher variance than is found in previous daily returns and the variance should depend on the magnitude of $\text{unc}_t(\nu)$. In Table 7 we show that estimates of σ_u^2 are not statistically significant. Moreover the mean of estimates of $\sigma_u^2 \text{unc}_t(\nu)$ is typically at least an order of magnitude smaller than the mean of estimates of $\sigma_v^2 \text{vol}_t(\omega)$ and sometimes much smaller. In Panel A, we estimate $\sigma_u^2, \log \nu$, and θ simultaneously, which because of singularities, leads to high standard errors for most parameters, including θ. In Panel B, we fix the value of $\log \nu$ at several different values and show that θ is significant, provided $\log \nu$ is not small, and σ_u^2 is never significant. In all of the specifications, allowing σ_u^2 to be greater than zero has almost no effect on the likelihood. We conclude that it is reasonable to assume that σ_u^2 is close to zero and in the rest of the paper we keep σ_u^2 fixed exactly at zero.

5.3.7. Subsamples

As an additional robustness check, we break our sample into four subsamples of equal length and allow for time-varying aversion to uncertainty. Because this introduces several new parameters we do not also consider a risk-return trade-off. We estimate different θ's for each subsample simultaneously with $b, \sigma^2, \log \omega$, and $\log \nu$. The latter parameters are required to have the same value throughout the sample. In particular we estimate

$$r_{et+1} \sim N\left[b + \sum_{k \in \Omega} \theta_k \, \mathbf{1}_{t,k} \, \text{unc}_t(\nu), \; \sigma^2 \text{vol}_t(\omega)\right],$$

where $\Omega = \{$1968:4-1977:2, 1977:3-1986:1, 1986:2-1994:4, 1995:1-2003:3$\}$ and where the indicator function $\mathbf{1}_{t,k}$ is one if $t \in k$ and zero otherwise. The results in Table 8 show that estimates of θ_k are positive in all four subperiods and significantly different from zero in three of the four subperiods. When we allow for time-varying aversion to uncertainty, the increase in the likelihood, as compared to the likelihoods in Table 2, is small. We cannot reject the hypothesis that all of the θ's are equal and aversion to uncertainty is constant over time.

5.3.8. Are prices driving our results?

One possible concern about our empirical measure of model uncertainty is the implementation of a modified Gordon growth model to construct aggregate forecasts of market returns. Since the first term of the model employs forecasted aggregate corporate profits scaled by a price variable, the critique exists that prices are driving our results. However, this is not the case. Panel B of Table 4 shows that the underlying driving force for our results is the second term, disagreement in long-term forecasts of profits, which is not scaled by a price variable. In line four of Table 4, Panel B, we see that disagreement in long-term forecasts of profits does nearly as well as disagreement about our constructed market return. Moreover, it can be shown that the first term of the model is quantitatively small.[32] Though, the first term does have some predictive power and we include it to adhere to the structure of our theoretical model.

6. Risk, uncertainty, and the cross section of stocks—empirical results

In the previous section, we showed that market uncertainty matters for expected market returns. In this section, we investigate whether market risk and uncertainty matter for the cross section by (1) studying the returns on portfolios with varying degrees of exposure to risk and uncertainty and (2) testing if exposure to risk and uncertainty can explain the returns on many other often studied portfolios.

The model presented in Section 2 implies that the conditional expected excess return of any asset k is

$$E_t r_{kt+1} = \beta_{vk}\gamma V_t + \beta_{uk}\theta M_t, \tag{33}$$

where β_{vk} and β_{uk} are regression coefficients of the risk in asset k on market risk and of the uncertainty in asset k on market uncertainty.[33] In Section 5, we estimated Eq. (33) for the market excess return using flexible functional forms for risk $V_t = \sigma^2 \text{vol}_t(\omega)$ and uncertainty $M_t = \text{unc}_t(\nu)$. In particular we estimated the nonlinear regression

$$r_{et+1} = b + \tau \text{vol}_t(\omega) + \theta \text{unc}_t(\nu) + \epsilon_{et+1}, \tag{34}$$

where $\tau = \gamma\sigma^2$ and the conditional variance of ϵ_{et+1} is $\sigma^2 \text{vol}_t(\omega)$. In most of this section we keep b, τ, θ, ω, and ν fixed at their QMLE estimates, $\hat{b} = -0.012$, $\hat{\tau} = 0.120$, $\hat{\theta} = 1453.191$, $\hat{\omega} = \exp(2.704) = 14.939$, and $\hat{\nu} = \exp(2.731) = 15.346$ from Section 5 (these are the estimates displayed in specification 7, Table 2), though we do present some results in which ω and ν are estimated entirely from the cross section using no information (directly) from the market.

Taking unconditional expected values of Eq. (33) yields an expected return-beta formulation of our model

$$E r_{kt+1} = \beta_{vk}\lambda_v + \beta_{uk}\lambda_u,$$

where the prices of risk and uncertainty are

$$\lambda_v = E\ \gamma V_t \tag{35}$$

$$\lambda_u = E\ \theta M_t. \tag{36}$$

[32] This result is appealing as it seems reasonable that there is more disagreement about long-term corporate profits than short-term corporate profits.

[33] We make the assumption throughout this paper that β_{vk} and β_{uk} are constant over time. As noted earlier, this assumption implies restrictions on the exogenous processes for the state vector and asset prices.

This formulation is straightforward but does not take advantage of the information about β_{vk} in the contemporaneous correlation of r_{et+1} and r_{kt+1}.

An alternative expected return-beta formulation which exploits the contemporaneous correlation recovers β_{vk} and β_{uk} from a time-series regression of r_{kt+1} on $[r_{et+1} - b - \theta \text{unc}_t(\nu)]$ and $\theta \text{unc}_t(\nu)$. To verify the alternative, we write Eq. (33) for asset k as

$$r_{kt+1} = \beta_{vk} \tau \text{vol}_t(\omega) + \beta_{uk} \theta \text{unc}_t(\nu) + \epsilon_{kt+1} \tag{37}$$

and decompose the error term ϵ_{kt+1} into two components:

$$\epsilon_{kt+1} = \beta_{vk} \epsilon_{et+1} + \varrho_{kt+1}. \tag{38}$$

The first component $\beta_{vk}\epsilon_{et+1}$ depends on market noise and the second component ϱ_{kt+1} is idiosyncratic noise for asset k which is uncorrelated with market noise, ϵ_{et+1}, market risk, and market uncertainty.[34] The time t expected values of both error terms, ϵ_{et+1} and ϱ_{kt+1}, are zero. Substituting (38) into (37) and rearranging yields

$$r_{kt+1} = \beta_{vk} \left[\tau \text{vol}_t(\omega) + \epsilon_{et+1}\right] + \beta_{uk} \theta \text{unc}_t(\nu) + \varrho_{kt+1}. \tag{39}$$

Since ϱ_{kt+1} is orthogonal to $\text{vol}_t(\omega)$, ϵ_{et+1}, and $\text{unc}_t(\nu)$, it follows that a population regression of r_{kt+1} on $[r_{et+1} - b - \theta \text{unc}_t(\nu)]$ and $\theta \text{unc}_t(\nu)$ yields estimates of the coefficients β_{vk} and β_{uk}.[35] In our estimation of Eq. (39), we also include a constant term, denoted a_k. We can think of this representation as a two-factor model with the factors being a measure of market risk and a measure of market uncertainty.

In Section 6.1 we construct portfolios that are designed to have large and small values of β_{vk} and β_{uk} and study their returns. In Section 6.2 we use Generalized Method of Moments (GMM) to estimate β_{vk} and β_{uk} for 130 portfolios that have been studied in recent research and investigate the prices of risk and uncertainty. In Section 6.3 we use GMM to estimate a stochastic discount factor formulation of our model which allows us to easily estimate market risk, market uncertainty, and their impact on the cross section of returns simultaneously.

6.1. The return on uncertainty and risk portfolios

Individual stock returns are regressed on $[r_{et+1} - \hat{b} - \hat{\theta} \text{unc}_t(\hat{\nu})]$ and $\hat{\theta} \text{unc}_t(\hat{\nu})$ to yield estimates of β_{vk} and β_{uk} from Eq. (39). In order to efficiently estimate the betas, we use a rolling sample regression method approach for each firm in the sample. Rolling sample regressions are run for each firm where at least 20

[34] The assumption that β_{vk} is constant over time allows us to make this decomposition.

[35] Note that Eq. (34) guarantees that $[r_{et+1} - b - \theta \text{unc}_t(\nu)] = \tau\, \text{vol}_t(\omega) + \epsilon_{et+1}$. Also note that since $(\beta_{vk}\epsilon_{et+1} + \varrho_{kt+1})$ and ϵ_{t+1} are orthogonal to $\text{unc}_t(\nu)$, it follows that ϱ_{kt+1} is orthogonal to $\text{unc}_t(\nu)$.

quarters of returns are available, and estimates of β_{vk} and β_{uk} are collected from each firm in each quarter. From the over 25,000 firms in the Center for Research in Security Prices (CRSP) universe between fourth quarter 1969 and fourth quarter 2003, only 14,252 meet the requirements for rolling sample regressions.

Portfolios of individual stocks are formed in two ways. First, we investigate portfolios sorted only on sensitivities to uncertainty, $\hat{\theta}\text{unc}_t(\hat{v})$. In order to form portfolios, stocks are ranked each quarter according to the coefficient β_{uk} on uncertainty. Stocks are then sorted into quintiles based on the exposure to uncertainty in each quarter. Within each quintile, stocks are value-weighted relative to the other firms in the quintile, which are then cumulated to form the portfolio return. For the five portfolios over the 121 quarters of the sample, the average returns to portfolios sorted on sensitivities to uncertainty range from 1.7% to 3.6% per quarter.

Even though the standard errors for all of the portfolios are large and the returns cannot be statistically distinguished due to our small sample size, it is still interesting to investigate whether there is any evidence that firms with large exposure to uncertainty have higher returns. Excluding the quintile of stocks with the very lowest exposure to uncertainty, stocks with a larger exposure to uncertainty have higher returns on average. The returns on stocks with the most exposure to uncertainty are especially large. While portfolio returns are not statistically distinguishable from one another, the economic differences are important: high uncertainty portfolios on average have returns 200 basis points greater than low uncertainty portfolios per quarter.[36]

We also form portfolios by sorting on sensitivities to risk and uncertainty. Similar to the method described above, we first rank stocks according to β_{uk} and form three portfolios. For each of these portfolios, we then rank stocks by β_{vk} and sort into three portfolios. The resulting nine portfolios have varying exposure to risk and uncertainty, and we examine summary statistics for these portfolios in Table 10. Even though our standard errors for all of the portfolios are again large and the returns cannot be statistically distinguished due to our small sample size, it is again interesting to investigate whether firms with large exposure to risk and uncertainty have higher returns. Regardless of the level of risk exposure, the average returns on portfolios are increasing in exposure to uncertainty with one exception: medium-uncertainty, low-risk stocks have a higher return than high-uncertainty, low-risk stocks. Regardless of the level of uncertainty exposure, the average returns on portfolios are increasing in exposure to risk with two exceptions: medium-risk stocks have lower returns that low-risk stocks when the stocks have either a low or medium exposure to uncertainty.

[36]More detailed descriptions of the results for cross-sectional uncertainty portfolios are available on request.

We need to caution that although the portfolio returns are not statistically distinguishable from one another, we interpret the point estimates as if they are meaningfully different. When we compare low-risk/low-uncertainty portfolios to high-risk/high-uncertainty portfolios, there exists an approximate 200 basis point spread of quarterly returns. This represents a significant economic magnitude as far as point estimates go.

Finally, we also provide in the lower panel of Table 10 summary statistics on the weights of each portfolio over the time period relative to the CRSP universe.[37] We observe that stocks with low exposure to risk are generally larger than stocks with high exposure to risk, while stocks with moderate uncertainty exposure tend to be larger on average than stocks in either the low- or high-uncertainty portfolios.

6.2. Estimating an expected return-beta representation

In this section, we use GMM to estimate an expected return-beta representation and investigate if risk and uncertainty can help explain the returns on 130 portfolios which have been the subject of many previous studies. The 130 portfolios include 25 portfolios sorted on size and book-to-market, 25 portfolios sorted on size and short-term reversal, 25 portfolios sorted on size and momentum, 25 portfolios sorted on size and long-term reversal, ten portfolios sorted on earnings-to-price, ten portfolios sorted on dividend-to-price, and ten portfolios sorted on cash flow-to-price. Data for all the portfolios are obtained from Kenneth French's Web site. Real excess returns are formed by subtracting the nominal risk-free rate and adjusting for inflation with the Consumer Price Index (CPI).[38]

Our two-factor risk and uncertainty model implies that the expected excess return on any portfolio, Er_{kt+1}, should be linearly related to sensitivities to market risk β_{vk} and market uncertainty β_{uk},

$$Er_{kt+1} = \beta_{vk}\lambda_v + \beta_{uk}\lambda_u + \iota_k, \tag{40}$$

where ι_k is a pricing error for asset k, which according to our model should be zero. Here, β_{vk} and β_{uk} are coefficients in a time-series regression of r_{kt+1} on a constant a_k, market risk, and market uncertainty:

$$r_{kt+1} = a_k + \beta_{vk}\dot{V}_t + \beta_{uk}\dot{M}_t + \varrho_{kt+1}, \tag{41}$$

where

$$\begin{aligned}\dot{V}_t &= \hat{\tau}\text{vol}_t(\hat{\omega}) + \hat{\epsilon}_{et+1}\\ \dot{M}_t &= \hat{\theta}\text{unc}_t(\hat{\nu}),\end{aligned}$$

[37]Because the sample contains approximately 60% of the number of stocks in the CRSP universe, the weights do not sum to 100%.

[38]Summary statistics are available upon request. For the CPI we use the Consumer Price Index For All Urban Consumers, All Items, seasonally adjusted.

and where the hats denote parameters fixed at QMLE estimates. We investigate whether the price of risk, λ_v, and the price of uncertainty, λ_u, are significant; and if the pricing errors, ι_k, for all assets are jointly close to zero.

We estimate the cross-sectional relationships in Eq. (40) and the time-series regressions in Eq. (41) simultaneously using GMM. For each asset k we have the moment conditions

$$E\begin{bmatrix} r_{kt+1} - a_k - \beta_{vk}\dot{V}_t - \beta_{uk}\dot{M}_t \\ \left(r_{kt+1} - a_k - \beta_{vk}\dot{V}_t - \beta_{uk}\dot{M}_t\right)\dot{V}_t \\ \left(r_{kt+1} - a_k - \beta_{vk}\dot{V}_t - \beta_{uk}\dot{M}_t\right)\dot{M}_t \\ r_{kt+1} - \beta_{vk}\lambda_v - \beta_{uk}\lambda_u \end{bmatrix} = 0. \tag{42}$$

The moment conditions for all assets combined yield a system for which we can estimate the scalars λ_v and λ_u (which do not vary across assets) as well as a_k, β_{vk}, and β_{uk} (which vary across assets). When there are n assets there are $4n$ moment conditions and $(3n + 2)$ parameters.[39]

The GMM estimation of our joint system involves setting $(3n + 2)$ linear combinations of the sample moments equal to zero. More formally, GMM sets

$$a_T g_T = 0,$$

where a_T is a $(3n+2)$ by $4n$ matrix and g_T is a $4n$ by 1 vector of sample means corresponding to the moment conditions in Eq. (42) for all assets. Similar to Cochrane (2005), we specify the matrix a_T so that GMM estimates of the time-series parameters (a_k, β_{vk}, and β_{uk} for all assets) are identical to their least squares estimates and GMM estimates of the cross-sectional parameters (the scalars λ_v and λ_u) are identical to their Generalized Least Squares (GLS) estimates. Unlike Cochrane (2005), the covariance matrix of all assets is used as the weighting matrix for the GLS estimates of λ_v and λ_u rather than the covariance matrix of the residuals from the time-series regressions.[40] Even though our GMM estimates are identical to least squares and GLS estimates, estimating our system with GMM is convenient because GMM allows us to easily produce asymptotic standard errors for λ_v and λ_u which take into account that the time-series and

[39]Typically in factor models the factors are contemporaneous with returns. However, our theoretical model says that the dispersion of forecasts of time t returns stated at $t-1$ [labeled $\hat{\theta}\text{unc}_{t-1}(\hat{\nu})$] should be related to returns at time t not time $t-1$. Therefore it is correct to include $\hat{\theta}\text{unc}_{t-1}(\hat{\nu})$ as a factor at time t and it would be incorrect to include $\hat{\theta}\text{unc}_{t}(\hat{\nu})$ as a factor at time t. A similar argument can be applied to our measure of predicted volatility, $\hat{\tau}\text{vol}_{t-1}(\hat{\omega})$, since it is designed to predict returns (volatility) at time t but not $t-1$. Though, as we discuss above, we usually take an alternative route and use the contemporaneous factor $[\hat{\tau}\text{vol}_{t-1}(\hat{\omega}) + \hat{\epsilon}_{et}]$ for volatility. For uncertainty there is no counterpart and we have to use non-contemporaneous factors. When we include Fama and French factors below, we do use them contemporaneously because this is the usual way to use the Fama-French factors and we want to give the Fama-French factors the best possible chance of mattering.

[40]In our approach, the GLS weighting matrix does not depend on the time-series estimates and thus is the same regardless of which factors are included, though using the weighting matrix advocated by Cochrane (2005) would not lead to a substantial change in our results.

cross-sectional regressions are estimated simultaneously.[41]

We find in Table 11 that the price of risk is negative and not significant. While our sample has a quarterly frequency, this result is in line with the analysis of Ang, Hodrick, Xing, and Zhang (2006) who examine the pricing of aggregate volatility risk in a monthly cross-sectional analysis of stock returns. In contrast to the price of risk, the price of uncertainty is relatively large and positive. In estimating the prices of risk and uncertainty, we use standard reduced-form econometric techniques even though our model has predictions for these values which are displayed in Eqs. (35) and (36). The price of uncertainty, λ_u, should be the unconditional expectation of $\hat{\theta}\mathrm{unc}_t(\hat{\nu})$ and the price of risk should be the unconditional expectation of $\hat{\tau}\mathrm{vol}_t(\hat{\omega})$. We see that in model one of Table 11 the estimate of λ_u, 0.027, is very close to the sample mean of $\hat{\theta}\mathrm{unc}_t(\hat{\nu})$ which is 0.025. The standard error of the estimate of λ_u tells us that we can not reject its value being 0.025, confirming a prediction of our model. Our estimate of the price of risk, λ_v, -0.011, is also statistically close to the sample mean of $\hat{\tau}\mathrm{vol}_t(\hat{\omega})$ which is 0.001.[42]

In Table 11 we also present results for the Capital Asset Pricing Model (CAPM) and the Fama-French factors (including factors for momentum, short-term reversal, and long-term reversal). We use GMM to estimate various versions of the joint time-series and cross-sectional system where in some specifications we include the market excess return, the HML, the SMB, the UMD (Momentum), the STR (Short-Term Reversal), and the LTR (Long-Term Reversal) factors, in addition to risk and uncertainty factors.[43] The results for specification six in Table 11 show that including additional factors has only a very small effect on estimates of the price of uncertainty. The prices of all of the Fama-French factors are significantly positive except for the price of SMB which is insignificantly negative. In specifications five and six the estimates of the standard errors of the price of uncertainty are smaller than estimates of the standard errors for the price of any other factor and estimates of its *t*-statistics are larger than the estimates for any other factor. In specifications two and seven we measure risk with $\hat{\tau}\mathrm{vol}_t(\hat{\omega})$ rather than with $[r_{et+1} - b - \theta\mathrm{unc}_t(\hat{\nu})]$ and find that the price of risk is not significantly altered.

One way to evaluate the performance of models is to look at pricing errors. The *J*-stats presented in Table 11 provide a measure of how big the pricing errors are for all of the assets and the corresponding *p*-values tell us how likely it is to see pricing errors at least this large. Our results show that with probability one, we should see pricing errors as large as observed for all our models. However, given the large number

[41]To compute the spectral density matrix at frequency zero, a key ingredient of GMM standard errors, we use the method of Newey and West (1987) with eight lags.

[42]Note that the sample mean of $\hat{\tau}\mathrm{vol}_t(\hat{\omega})$ is equal to the sample mean of $[r_{et+1} - b - \theta\mathrm{unc}_t(\hat{\nu})]$.

[43]The data for these factors are obtained from Kenneth French's Web page. The factors are adjusted for inflation with the CPI. The exact specification for the GMM estimation is reported in the legend of Table 11.

of moment conditions and the fact that J-stats require the pseudo-inversion of a term which involves the spectral density matrix at frequency zero, S, the results for the J-stats should be viewed with caution. Since we pre-specify a_T, estimates of standard errors of parameters do not require an inversion of S and can be viewed as being more reliable.[44]

One drawback of the approach taken in this section is that the standard errors reported in Tables 11 do not take into account that the parameters ω and ν are pre-estimated from the market. In the next section, the same 130 portfolios with a stochastic discount factor representation are examined that allows us to jointly estimate the betas, the lambdas, and the nonlinear parameters ω and ν using information from the cross section without imposing pre-estimated market values.

6.3. Estimating a stochastic discount factor representation

We now estimate a stochastic discount factor representation of our model. We form the stochastic discount factor

$$s_{t+1} = a + s_v \left[\hat{\tau}\mathrm{vol}_t(\omega)\right] + s_u \left[\hat{\theta}\mathrm{unc}_t(\nu)\right]$$

and use the implication of our model that

$$E s_{t+1} r_{jt+1} = 1$$

for any asset return r_{jt+1}, at any date, to estimate a, s_v, and s_u, as well as ω and ν.[45] As discussed by Cochrane (2005), estimates of s_v and s_u are directly related to estimates of λ_v and λ_u in the expected return-beta formulation of our model, Eq. (40), but answer the question whether risk and uncertainty can help explain the return on assets given the other factors rather than the question of whether risk and uncertainty are priced. In our context, the stochastic discount factor representation is especially convenient because it is easily amendable to estimating the nonlinear parameters ω and ν from the cross section. For asset returns, we use data on gross returns rather than excess returns and emphasize the measure $\hat{\tau}\mathrm{vol}_t(\omega)$ of risk though we present results for both of our specifications of risk.[46] We discuss results for the fixed weighting matrix proposed by Hansen and Jagannathan (1997) [hereafter referred to as the HJ weighting matrix] and the optimal GMM weighting matrix discussed by Hansen (1982). For comparison purposes, we also study specifications of the stochastic discount factor where the Fama-French factors enter linearly.

[44] When computing the spectral density matrix we did not subtract sample means. Asymptotically, subtracting sample means should not matter if our models are correct. However, in finite samples it does matter and would drastically change our reported J-stats and p-values. It would, however, have no noticeable effect on our estimates of standard errors.

[45] The parameters τ and θ are fixed at $\hat{\tau}$ and $\hat{\theta}$ throughout to achieve a convenient scaling of risk and uncertainty. This has no effect on the fit of the model.

[46] We emphasize $\hat{\tau}\mathrm{vol}_t(\omega)$ as a measure of risk because it is not possible to estimate ω and ν together when risk is measured with $[r_{et+1} - b - \theta\mathrm{unc}_t(\nu)]$.

In Table 12, we estimate the contributions of risk and uncertainty to the stochastic discount factor in explaining the returns on the same 130 portfolios studied in Section 6.2. Estimates of the contribution of uncertainty, s_u, are fairly constant across specifications and independent of the weighting matrix employed. For the HJ weighting matrix, estimates of s_u are significant: (1) in the two-factor model both when ν is estimated and when it is fixed, and (2) in the presence of the Fama-French factors when ν is fixed but not when ν is reestimated. For the optimal weighting matrix, estimates of s_u are significant in all specifications regardless of which factors are present or if ν is reestimated.

When we reestimate the nonlinear parameters, ω and ν, we find that estimates are similar to the market estimates presented in Section 5. Estimates of $\log \nu$ range from 2.230 to 2.717, which are very close to its market estimate of 2.731. Estimates of $\log \omega$ range from 4.116 to 4.491, which, given the large standard errors, are statistically close to the market estimate of 2.704. This provides additional evidence confirming our estimated weighting scheme that emphasizes the dispersion of non-extreme forecasts and the MIDAS weighting scheme which places more weight on recent daily volatility.

In Panel A of Table 12, we provide results for the Hansen-Jagannathan distance (hereafter the HJ-dist) which measures how far a candidate stochastic discount factor is from a stochastic discount factor which can unconditionally exactly price all assets. We find that all models perform poorly on this criteria and the standard errors of the HJ-dist indicate that the performances of the models are indistinguishable. There is not a significant drop in the HJ-dist when uncertainty is added to the Fama-French factors or the CAPM. The Fama-French factors and the CAPM do not perform any better; there is not a significant drop in the HJ-dist when the Fama-French factors are added to the CAPM or, in results available upon request, when the market is added to a constant stochastic discount factor.

Given the large number of moment conditions used and the fact that the optimal weighting matrix requires the inversion of the spectral density matrix at frequency zero, S, the results for the optimal weighting matrix should be viewed with caution. The results for the HJ weighting matrix, including the HJ-dist and the standard errors for parameters, do not require inverting S and are not as problematic.

6.4. Summary

Uncertainty by no means provides a complete explanation of the cross section of stock returns, but there is evidence that uncertainty matters. In particular, for portfolios often studied in the literature we find that: (1) estimates of the price of uncertainty are very significant in all specifications considered and are consistent with market estimates; (2) in the two-factor model, uncertainty significantly contributes to the stochastic discount factor for both the optimal weighting matrix and the HJ weighting matrix regardless of whether the nonlinear parameters are fixed or reestimated; (3) in the presence of the Fama-French factors, uncertainty significantly contributes to the stochastic discount factor when the optimal weighting matrix is used (regardless of whether the nonlinear parameters are fixed or reestimated); however, with the HJ weighting matrix, uncertainty significantly contributes to the stochastic discount factor when the nonlinear parameters are fixed but not when they are reestimated; and (4) estimates of the nonlinear parameter ν are very significant and similar to its market estimate for both the optimal weighting matrix and the HJ weighting matrix.

On the other hand, there is not a significant drop in the HJ-dist when uncertainty is added to the CAPM or when uncertainty is added to the Fama-French factors. The problems we face in the cross section are similar to the problems all models face in explaining the cross section. For example, when Fama-French factors are added to the market return there is not a significant drop in the HJ-dist. Similar negative results for asset pricing models are abundant in the literature. Hansen and Jagannathan (1997) show that the CAPM and consumption-based models are not much of an improvement over a constant discount factor. Lewellen, Nagel, and Shanken (2006) show that although many models perform well on the 25 portfolios sorted on size and book-to-market, the same models perform poorly on other assets.

7. Conclusions

The growing recent literature on uncertainty and its impact on asset pricing has relatively few empirical implementations. Although uncertainty is difficult to measure, we suggest a reasonable proxy for the amount of uncertainty in the economy is the degree of disagreement of professional forecasters. In contrast to prior literature, which has focused on disagreement about individual stocks, our emphasis is on aggregate measures of disagreement. Furthermore, we offer an alternative explanation for why disagreement is priced, namely that economic agents interpret disagreement as model uncertainty.

Moreover, we provide a measure of uncertainty constructed with a flexible weighting scheme, which can accommodate assigning more or less weight to extreme forecasts. Our estimates of the optimal weights entail ignoring the extremes and placing nearly all of the weight on the center of the distribution. We find that uncertainty is empirically significantly related to market returns only when unequal weighting schemes are implemented. Flat weighted measures of uncertainty are not highly correlated with the market return and do not have a significant effect in regressions.

Uncertainty seems to be different from risk and seems to have a different effect on returns than risk. Uncertainty is highly correlated with the market excess return whereas risk is not. Uncertainty has a very weak correlation with risk and past uncertainty has no predictive ability of future risk or vice versa. We find stronger empirical evidence for an uncertainty-return trade-off than for the traditional risk-return trade-off. Further, our measure of uncertainty does not seem to encompass risk.

Our results are generally not sensitive to the measure of uncertainty we construct as long as extreme forecasts are removed, replaced, or down-weighted. We find similar results if aggregate corporate profit forecasts are used instead of constructed aggregate market return forecasts. Uncertainty aversion is significant across subperiods of our sample and whenever uncertainty has been unusually large, the market excess return the subsequent quarter has also been large. Interestingly, two of the lowest values of uncertainty occurred when Presidents Nixon and Clinton were reelected.

We also investigate the importance of uncertainty for the cross section and find that the price of uncertainty is significantly positive. Moreover, uncertainty contributes to the explanation of the returns on other assets in the presence of the Fama-French factors. While uncertainty is by no means a complete explanation of the cross section of stock returns, there is evidence that it matters as we find that estimates of the price of uncertainty are very significant.

Appendix A. More on imputing asset return forecasts

Some of our analysis requires predictions on variables which do not directly appear in the SPF. In this appendix, we briefly describe how we construct forecasts of the real rate of growth of a variable for which only nominal forecasts of levels are available and then show how we construct forecasts of the expected real return on a nominally risk-free bond. Finally, we provide more details on how we implement the Gordon growth model.

The constructed gross quarterly forecasted rate of real growth, according to forecaster i, in the nominally forecasted variable X between quarters m and n is

$$\left(\frac{E_{it}X_n \; E_{it}P_m}{E_{it}X_m \; E_{it}P_n}\right)^{\frac{1}{n-m}}, \tag{43}$$

where P_s is the price level at time s and, for any variable Y, $E_{it}Y_s$ is forecaster i's time t prediction of the value of Y that will be realized at period s. In general, this is only an approximation since usually

$$E_{it}\left[\left(\frac{X_n P_m}{X_m P_n}\right)^{\frac{1}{n-m}}\right] \neq \left(\frac{E_{it}X_n \; E_{it}P_m}{E_{it}X_m \; E_{it}P_n}\right)^{\frac{1}{n-m}} \tag{44}$$

even when $t = m$ and $n = m + 1$. For forecasts of the price level we use forecasts of the output deflator since forecasts of the CPI only became available in the fourth quarter of 1991.[47]

We approximate forecaster i's prediction of the gross real return on a nominally risk-free bond with

$$E_{it}r_{bt+1} = \frac{R_{b+1}P_t}{E_{it}P_{t+1}}, \tag{45}$$

where R_{bt+1} is the gross nominal return on the bond (which is known at time t) and P_t is the time t value of the output deflator.[48]

We face a difficult timing issue when implementing the Gordon growth model in that forecasts in the SPF are given in the middle of a quarter. For example forecasts made during the first quarter of 2001 had to be returned to the Federal Reserve Bank of Philadelphia no later than February 12, 2001. In the 2001Q1 survey, forecasters were asked to provide predictions for the previous quarter (2000Q4), the current quarter (2001Q1), the next quarter (2001Q2), two quarters (2001Q3), three (2001Q4), and four quarters ahead (2002Q1). Since some information about the values of the variables in the first quarter may be learned in January it would be inappropriate to view the forecasts for the current quarter as being forecasts stated during $t = 2000Q4$ of $t + 1 = 2001Q1$ values. One could view the forecasts for next quarter as stated during $t = 2001Q1$ of $t + 1 = 2001Q2$ values. However, this neglects the short-term information in the current quarter forecasts. Consequently, when implementing the Gordon growth model, we interpret the sum of forecasts *stated* for the current quarter's and next quarter's corporate profits (deflated by forecasts of the price level), divided by two, as effectively being forecasts stated during $t = 2001Q1$ of $t + 1 = 2001Q2$ corporate profits.[49]

For the long-term growth rate in the Gordon growth model, ξ_{it}, we use forecaster i's predicted growth rate of corporate profits over the longest horizon available in the SPF. In the early years of the survey forecasts four quarters ahead are very sparse; therefore, we usually let the forecast horizon be from last quarter to three quarters ahead. We refer to this as a horizon of four. For instance, in the first quarter of 1975, we consider the forecasted growth rate from the fourth quarter of 1974 to the fourth quarter of 1975.

Appendix B. Uncertainty and disagreement

We describe an environment in which disagreement is directly related to uncertainty. We assume forecasters have prior information about the market return and every period observe a vector of information that is related to the market return. We provide conditions under which the amount of uncertainty in the economy is always proportional to the amount of disagreement.

In order to illustrate the relationship between uncertainty and disagreement we take a strong stand on the types of models forecasters are using. We assume each forecaster's uncertainty is limited to uncertainty in the mean of the market return. Assume

[47] We use the CPI when we deflate the actual level of variables.

[48] In general this is an approximation because usually

$$\frac{R_{b+1}P_t}{E_{it}P_{t+1}} \neq R_{b+1}P_t \; E_{it}\left[\frac{1}{P_{t+1}}\right].$$

[49] This assumption does not have a large effect on our results. If we implemented the Gordon growth model literally and ignored current quarter stated corporate profit forecasts, our results are essentially the same.

there are n forecasters, where n is large, and that before observing a vector of observations at time t forecaster i believes that the *mean* of the market return at time t is approximately μ_{it-1}. The confidence of forecaster i in this belief is measured with

$$P_{it-1} = E_{it-1}\left[(\mu_{it-1} - \mu_{t-1})^2\right], \tag{46}$$

where μ_{t-1} is the true mean of the market return and E_{it-1} denotes the expectation with respect to information available to forecaster i at time $t-1$.[50] We call P_{it-1} the uncertainty of forecaster i. We assume P_{it-1} (but not μ_{it-1}) is constant across forecasters. We call $P_{t-1} = P_{it-1}$ the amount of uncertainty in the economy at the end of period $t-1$. The true mean of the market return evolves over time:

$$\mu_t = A_{t-1}\mu_{t-1} + \iota_t, \tag{47}$$

where ι_t is an unobserved scalar standard normal random variable with mean zero and variance Q_{t-1}. Forecasters know the values of A_{t-1} and Q_{t-1} at time $t-1$.

Each period, all forecasters simultaneously observe a different vector of random variables. Forecaster i observes the vector

$$z_{it} = G_{t-1}\mu_{t-1} + w_t + w_{it}, \tag{48}$$

where μ_{t-1} is the true mean of the market return; w_t is a vector of independent normal random variables with mean zero and covariance matrix H_{t-1}; and w_{it-1} is a vector of normal random variables with mean zero and covariance matrix K_{t-1}. We assume G_{t-1}, H_{t-1} and K_{t-1} are constant across forecasters and that their values are known at time $t-1$. We assume w_{it} is independent of w_s for any s, independent of w_{is} when $s \neq t$, and independent of w_{js} when $j \neq i$ for any s. The vector w_t includes information that is common to all forecasters and the vector w_{it} includes information that is specific to forecaster i. Forecaster i does not observe z_{jt} or μ_{jt-1} for $j \neq i, \mu_{t-1}$, w_t, or w_{it} for any i. He only observes z_{it}. We assume K_{t-1} is positive semi-definite, H_{t-1} is positive semi-definite, and the sum $K_{t-1} + H_{t-1}$ is positive definite.

After seeing z_{it}, forecaster i's belief about the mean of the market return for the next period (time $t+1$) is

$$\mu_{it} = A_{t-1}\left[1 + P_{t-1}G'_{t-1}(H_{t-1}+K_{t-1})^{-1}G_{t-1}\right]^{-1}\left(\mu_{it-1} + P_{t-1}G'_{t-1}(H_{t-1}+K_{t-1})^{-1}z_{it}\right) \tag{49}$$

and a measure of his confidence in this belief is

$$P_t = A^2_{t-1}P_{t-1}\left[1 + G'_{t-1}(H_{t-1}+K_{t-1})^{-1}G_{t-1}P_{t-1}\right]^{-1} + Q_{t-1}. \tag{50}$$

These formulas are a special case of the updating equations for the Kalman filter. Here P_t is the amount of uncertainty at the end of period t.

Let the amount of disagreement before forecasters observe z_{it} be denoted $\mathcal{D}_{t-1}$. This is measured as the variance of μ_{it-1} across forecasters. Since in this simple example all forecasters have the same amount of uncertainty, an equally weighted variance is sensible. After observing z_{it} the amount of disagreement about the return at time $t+1$, which is the variance of μ_{it} across forecasters, is

$$\begin{aligned}\mathcal{D}_t = A^2_{t-1}&\left[1 + P_{t-1}G'_{t-1}(H_{t-1}+K_{t-1})^{-1}G_{t-1}\right]^{-1}\cdot \\ &\left(\mathcal{D}_{t-1} + P_{t-1}G'_{t-1}(H_{t-1}+K_{t-1})^{-1}K_{t-1}(H_{t-1}+K_{t-1})^{-1}G_{t-1}P_{t-1}\right)\cdot \\ &\left[1 + P_{t-1}G'_{t-1}(H_{t-1}+K_{t-1})^{-1}G_{t-1}\right]^{-1}.\end{aligned} \tag{51}$$

This formula is valid because when n is large the sample variance of z_{it} across forecasters is K_{t-1}, with probability one. If n is finite then it is an approximation. Define the scalars

$$\phi_{dt-1} = \frac{\mathcal{D}_{t-1}}{P_{t-1}} \tag{52}$$

$$\phi_{ft-1} = \frac{G'_{t-1}(H_{t-1}+K_{t-1})^{-1}K_{t-1}(H_{t-1}+K_{t-1})^{-1}G_{t-1}}{G'_{t-1}(H_{t-1}+K_{t-1})^{-1}G_{t-1}} \tag{53}$$

to be respectively the ratios of the amount of a priori disagreement to uncertainty and a measure of the ratio of the amount of idiosyncratic observation noise to the total observation noise. The interpretation of ϕ_{ft-1} is valid when G_{t-1}, H_{t-1}, and K_{t-1} are all scalars, in which case

$$\phi_{ft-1} = \frac{K_{t-1}}{H_{t-1}+K_{t-1}}, \tag{54}$$

and heuristic otherwise.

Consider an example in which $Q_t = 0$ for all t. In this case if ϕ_{dt-1} and ϕ_{ft-1} are equal to each other, call their common value ϕ_{t-1}, then $\mathcal{D}_{t-1} = \phi_{t-1}P_{t-1}$ and it can be shown that

$$\mathcal{D}_t = \phi_{t-1}P_t, \tag{55}$$

[50] In this appendix we recycle notation. The definitions of symbols apply only for this appendix.

when n is large, so that the ratio of uncertainty to disagreement is the same at time $t-1$ and time t, with probability one.

More generally, if ϕ_{ft} is equal to the same time-invariant constant, ϕ, at all dates

$$\phi_{ft} = \phi, \qquad \forall t \tag{56}$$

and the initial amount disagreement at time zero, $\mathcal{D}_0$, equals ϕP_0, so that $\phi_{d0} = \phi$, then it follows that the the ratio of uncertainty to disagreement is always constant:

$$\mathcal{D}_t = \phi P_t, \qquad \forall t \tag{57}$$

when n is large, with probability one. Over time, since $Q = 0$, eventually both $\mathcal{D}_t$ and P_t converge to zero under weak assumptions on the other parameter values.

More generally when Q_t is not necessarily zero at all dates, there are conditions on ϕ_{dt} and ϕ_{ft} which guarantee that the ratio of uncertainty to disagreement is always constant. The conditions require that $\phi_{dt} = \tau_t \phi_{ft}$ be constant over time for a *particular* sequence of time-varying constants $\{\tau_t\}$. If the parameters A_t, Q_t, G_t, H_t, and K_t are constant over time then under weak assumptions, P_t converges to a positive number, $\mathcal{P}_\infty$ and $\mathcal{D}_t$ also converges to a positive number, $\mathcal{D}_\infty$. At the limit the ratio of uncertainty to disagreement is necessarily constant over time.

If K_{t-1} is not positive definite it is possible that ϕ_{ft-1} is zero which would make the link between disagreement and uncertainty not very useful. For example, if K_{t-1} is a matrix of zeros then eventually μ_{it} would be identical across agents and there would be no disagreement, even if there is a large amount of uncertainty. In this case all forecasters are eventually alike, so that even if each forecaster has a large amount of uncertainty there is no disagreement.

In this appendix, we have provided conditions in two different examples that guarantee uncertainty is proportional to disagreement. We have shown how with the accumulation of new information it is possible that the proportionality is preserved. In reality the beliefs of forecasters may respond to new information in more complicated ways than we have described and the dispersion of models across forecasters may be more heterogeneous.

Appendix C. An illustration of the importance of flexible weights

Consider the following example. Let there be 30 forecasters. Assume in the population of forecasts, each forecast is distributed normal with mean 0.02 and variance 0.00010. If we randomly generate 30 forecasts from this distribution and take the sample variance we would usually get a number close to 0.00010. For example, from one set of 30 draws, we find the estimated cross-sectional variance to be 0.00009. Now what if in addition to these 30 forecasts there is one irrational forecaster who believes that the excess return is x? We examine what happens to the estimated cross-sectional variance when the data consist of the 30 randomly generated forecasts from the population and the one extreme forecast. If the extreme forecast is close to 0.02 then it will not have a large effect on our estimated variances. For example, for the same set of 30 draws discussed above, if $x = 0.05$ then the estimated cross-sectional variance becomes 0.00011. In this case the one extreme observation has a noticeable but not a large effect. If $x = 0.10$ then the estimated cross-sectional variance becomes 0.00029. In this case, one extreme forecast causes the estimate of the cross-sectional variance to increase by almost three times. If $x = 0.20$ then the estimated cross-sectional variance is 0.00113, representing a ten-fold increase. To deal with this problem we use beta weights to compute weighted variances of forecasts.[51] This flexible weighting scheme can assign more or less weight to extreme forecasts.

Appendix D. Comparisons with Ghysels, Santa-Clara, and Valkanov (2005)

Using MIDAS, Ghysels, Santa-Clara, and Valkanov (2005) find there is a significant positive relation between risk and return in the stock market. They show that this finding is robust to asymmetric specifications of the variance process, and to controlling for variables associated with the business cycle. The risk-return relationship with MIDAS volatility estimates also holds in many subsamples and for other countries (see, for instance, Leon, Naveb, and Rubio, 2007; and Maheu and McCurdy, 2007). Given these empirical findings, it provides a good benchmark reduced-form regression to introduce uncertainty.

Ghysels, Santa-Clara, and Valkanov (2005) focus on monthly returns whereas we devote our attention to quarterly sampling frequencies because the professional forecast data used to construct our measure of uncertainty, M_t, are only available quarterly.[52] We expect, however, that the focus on the quarterly sampling frequency weakens empirical evidence of the risk-return trade-off, and the results confirm this. Moreover, we consider a different time period (1969-2003) and our definition of a quarter refers to a calendar quarter (matching forecasts) whereas the definition of a quarter in Ghysels, Santa-Clara, and Valkanov (2005) corresponds to a fixed number of trading days which are not directly related to calendar quarters.

[51] A formal statistical argument for computing weighted variances can be made. For example in the behavioral model described in Section 4.2, it is optimal to measure uncertainty with a weighted variance.

[52] The bulk of the analysis in Ghysels, Santa-Clara, and Valkanov (2005) focuses on monthly horizons though they do provide quarterly regressions between 1964 and 2000 of a specification similar to Eq. (28), without uncertainty, and find a significant and positive relation between risk and return.

There are several additional differences in our implementation which do not have a large effect on our results. As explained in the text, we estimate σ^2 (rather than set it equal to one), allow for serial correlation in daily returns in Eq. (22) and subtract sample means in Eq. (22). [Some of the results in Ghysels, Santa-Clara, and Valkanov (2005) subtracted sample means.] We also use the beta weights advocated by Ghysels, Santa-Clara, and Valkanov (2006) and Ghysels, Sinko, and Valkanov (2007) rather than the normal weights (Almon lags) used by Ghysels, Santa-Clara, and Valkanov (2005).

References

Alper, C. E., Fendoglu, S., Saltoglu, B., 2008. Forecasting stock market volatilities using MIDAS regressions: An application to the emerging markets. MPRA Paper No. 7460, University Library Of Munich, Germany.

Anderson, E. W., Ghysels, E., Juergens, J. L., 2005. Do heterogeneous beliefs matter for asset pricing? Review of Financial Studies 18, 875–924.

Anderson, E. W., Hansen, L. P., Sargent, T. J., 2003. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. Journal of the European Economic Association 1, 68–123.

Ang, A., Hodrick, R., Xing, Y., Zhang, X., 2006. The cross-section of volatility and expected returns. Journal of Finance 61, 259–299.

Baillie, R. T., DeGennaro, R. P., 1990. Stock returns and volatility. Journal of Financial and Quantitative Analysis 25, 203–214.

Becker, R., Clements, A., White, S., 2007. Does implied volatility provide any information beyond that captured in model-based volatility forecasts? Journal of Banking and Finance 31, 2535–2549.

Brandt, M., Kang, Q., 2004. On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach. Journal of Financial Economics 72, 217–259.

Braun, P., Zarnowitz, V., 1993. Twenty-two years of the NBER-ASA quarterly economic outlook surveys: Aspects and comparisons of forecasting performance. In: Stock, J., Watson, M. (Eds.), Business Cycles, Indicators, and Forecasting. University of Chicago Press, Chicago, pp. 11–93.

Brav, A., Lehavy, R., Michaely, R., 2005. Using expectations to test asset pricing models. Financial Management 34, 31–64.

Campbell, J. Y., 1987. Stock returns and the term structure. Journal of Financial Economics 18, 373–399.

Campbell, J. Y., Hentschel, L., 1992. No news is good news: An asymmetric model of changing volatility in stock returns. Journal of Financial Economics 31, 281–318.

Campbell, J. Y., Shiller, R. J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. Review of Financial Studies 1, 195–228.

Chen, Z., Epstein, L. G., 2002. Ambiguity, risk and asset returns in continuous time. Econometrica 70, 1403–1443.

Cochrane, J. H., 2005. Asset Pricing, Revised ed. Princeton University Press, Princeton, NJ.

Crombez, J., 2001. Updating expected returns based on consensus forecasts. Unpublished working paper, Ghent University.

Croushore, D., 1993. Introducing: The survey of professional forecasters. Federal Reserve Bank of Philadelphia Business Review, 3–13.

Diether, K. B., Malloy, C. J., Scherbina, A., 2002. Differences of opinion and the cross-section of stock returns. Journal of Finance 57, 2113–2141.

Foster, D. P., Nelson, D. B., 1996. Continuous record asymptotics for rolling sample variance estimators. Econometrica 64, 139–174.

French, K. R., Schwert, W., Stambaugh, R. F., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–29.

Gebhardt, W. R., Lee, C. M., Swaminathan, B., 2001. A test of efficiency of a given portfolio. Journal of Accounting Research 39, 135–176.

Ghysels, E., Santa-Clara, P., Valkanov, R., 2005. There is a risk-return tradeoff after all. Journal of Financial Economics 76, 509–548.

Ghysels, E., Santa-Clara, P., Valkanov, R., 2006. Predicting volatility: Getting the most out of return data sampled at different frequencies. Journal of Econometrics 131, 59–95.

Ghysels, E., Sinko, A., Valkanov, R., 2007. MIDAS regressions: Further results and new directions. Econometric Reviews 26, 53–90.

Ghysels, E., Wright, J., 2009. Forecasting professional forecasters. The Journal of Business and Economic Statistics, forthcoming.

Glosten, L. R., Jagannathan, R., Runkle, D. E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. Journal of Finance 48, 1779–1801.

Gordon, M. J., 1962. The Investment, Financing, and Valuation of the Corporation. Irwin Publishers, Homewood, Illinois.

Graham, J., Harvey, C., 2003. Expectations of equity risk premia volatility and asymmetry from a corporate finance perspective. Unpublished working paper, Duke University.

Guay, W. R., Kothari, S. P., Shu, S., 2003. Properties of implied cost of capital using analysts' forecasts. Unpublished working paper, MIT Sloan School of Management.

Guo, H., Whitelaw, R. F., 2006. Uncovering the risk-return relation in the stock market. Journal of Finance 61, 1433–1463.

Hansen, L. P., 1982. Large sample properties of generalized method of moments estimators. Econometrica 50, 1029–1054.

Hansen, L. P., Jagannathan, R., 1997. Assessing specification errors in stochastic discount factor models. Journal of Finance 52, 557–590.

Hansen, L. P., Sargent, T. J., 1995. Discounted linear exponential quadratic Gaussian control. IEEE Transactions on Automatic Control 40, 968–971.

Hansen, L. P., Sargent, T. J., 2001. Acknowledging misspecification in macroeconomic theory. Review of Economic Dynamics 4, 519–535.

Hansen, L. P., Sargent, T. J., 2003. Robust control of forward-looking models. Journal of Monetary Economics 50, 581–604.

Hansen, L. P., Sargent, T. J., 2005. Robust estimation and control under commitment. Journal of Economic Theory 124, 258–301.

Hansen, L. P., Sargent, T. J., 2007. Recursive robust estimation and control without commitment. Journal of Economic Theory 136, 1–27.

Hansen, L. P., Sargent, T. J., Tallarini, Jr., T. D., 1999. Robust permanent income and pricing. Review of Economic Studies 66, 873–907.

Hansen, L. P., Sargent, T. J., Turmuhambetova, G. A., Williams, N., 2006. Robust control and model misspecification. Journal of Economic Theory 128, 45–90.

Harvey, C. R., 2001. The specification of conditional expectations. Journal of Empirical Finance 8, 573–638.

Jagannathan, R., McGrattan, E., Scherbina, A., 1996. The declining U.S. equity premium. Federal Reserve Bank of Minneapolis Quarterly Review 24, 3–19.

Johnson, T. C., 2004. Forecast dispersion and the cross section of expected returns. Journal of Finance 59, 1957–1979.

Keynes, J. M., 1937. The General Theory of Employment. Quarterly Journal of Economics 51, 209–223.

Knight, F. H., 1921. Risk, Uncertainty and Profit. Houghton Mifflin Company, Boston, MA.

Kogan, L., Wang, T., 2002. A simple theory of asset pricing under model uncertainty. Unpublished working paper, MIT and The University of British Columbia.

Leon, A., Naveb, J. M., Rubio, G., 2007. The relationship between risk and expected return in Europe. Journal of Banking and Finance 31, 495–512.

Lewellen, J., Nagel, S., Shanken, J., 2006. A skeptical appraisal of asset-pricing tests. Unpublished working paper, Dartmouth College, Stanford University, and Emory University.

Liu, J., Pan, J., Wang, T., 2005. An equilibrium model of rare-event premia and its implication for option smirks. Review of Financial Studies 18, 131–164.

Maenhout, P. J., 2004. Robust portfolio rules and asset prices. Review of Financial Studies 17, 951–983.

Maenhout, P. J., 2006. Robust portfolio rules and detection error probabilities for a mean-reverting risk premium. Journal of Economic Theory 128, 136–163.

Maheu, J. M., McCurdy, T. H., 2007. Components of market risk and return. Journal of Financial Econometrics 5, 560–590.

Merton, R. C., 1973. An intertemporal capital asset pricing model. Econometrica 41, 867–887.

Merton, R. C., 1980. On estimating the expected return on the market: An exploratory investigation. Journal of Financial Economics 8, 323–361.

Nelson, D. B., 1991. Conditional heteroskedasticity in asset returns: A new approach. Econometrica 59, 347–370.

Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.

Qu, S., Starks, L. T., Yan, H., 2003. Risk dispersion of analyst forecasts and stock returns. Unpublished working paper, University of Texas at Austin.

Turner, C. M., Startz, R., Nelson, C. R., 1989. A Markov model of heteroskedasticity, risk, and learning in the stock market. Journal of Financial Economics 25, 3–22.

Uppal, R., Wang, T., 2003. Model misspecification and underdiversification. Journal of Finance 58, 2465–2486.

Wiese, R., September 8, 1930. Investing for true value. Barron's, 5.

Williams, J. B., 1938. The Theory of Investment Value. Fraser Publishing, Burlington, VT.

Zarnowitz, V., 1985. Rational expectations and macroeconomic forecasts. Journal of Business and Economic Statistics 3, 293–312.

Table 1.
Data summary and description. This table provides a number of different statistics for the actual values and constructed forecasts of the real quarterly (gross) market return, r_{mt}, and the real quarterly (gross) return on a nominally risk-free bond, r_{bt}. The "actuals" columns provide the sample means and the sample standard deviations for the actual returns. The "median forecasts" columns provide the sample means and the sample standard deviations for the median of each period's constructed forecasts of returns. The "forecast errors" column provides the square root of the sample mean squared error of the median forecasts. The forecast data start with forecasts made in the fourth quarter of 1968 and end with forecasts made in the third quarter of 2003, collected from the Survey of Professional Forecasters. The actual data set runs from the first quarter of 1969 to the fourth quarter of 2003. The quarterly actual asset pricing data was computed from the monthly actual asset pricing data on Kenneth French's Web site and deflated by the CPI (obtained from Federal Reserve Economic Data, FRED, series CPIAUCSL) to obtain real returns. The market value of domestic corporations, used to compute the constructed market return forecasts, are from the Flow of Funds Accounts of The United States, prepared by the Board of Governors of the Federal Reserve System.

	Actuals		*Median forecasts*		*Forecast errors*
Variable	Mean	Std.	Mean	Std.	Root mean squared error
r_{mt}	1.0168	0.0901	1.0230	0.0179	0.0917
r_{bt}	1.0034	0.0064	1.0051	0.0045	0.0050

Table 2.

Risk-return and uncertainty-return trade-offs. This table displays estimates of several versions of the nonlinear regression

$$r_{et+1} = b + \tau \, \mathrm{vol}_t(\omega) + \theta \, \mathrm{unc}_t(\nu) + \epsilon_{t+1}$$

of quarterly excess returns r_{et+1} on the measure of risk, $\mathrm{vol}_t(\omega)$, specified in Eq. (22), and the measure of uncertainty, $\mathrm{unc}_t(\nu)$, specified in Eq. (25). Uncertainty is measured by the beta-weighted variance of market return forecasts. The estimates of b represent time-series Jensen alpha estimates. The variance of the error term, ϵ_{t+1}, is $\sigma^2 \mathrm{vol}_t(\omega)$ where σ^2 is a constant which we estimate. The measures vol_t and unc_t are based on information available in the previous quarter (the quarter before $t + 1$). Quasi-likelihood standard errors are listed under the estimates in parentheses. A dash, '-', indicates that the parameter was fixed at zero. The data for r_{et+1} are quarterly from 1969:1 to 2003:4. The forecast data and the daily data used to compute unc_t and vol_t are from 1968:4 to 2003:3.

Specification	b	τ	θ	$\log \omega$	$\log \nu$	σ^2	Log likelihood
1	0.012	-	-	-	-	1.277	147.297
	(0.007)					(0.160)	
2	0.011	-	-	2.780	-	1.582	151.111
	(0.006)			(0.446)		(0.237)	
3	0.009	0.812	-	2.768	-	1.577	151.184
	(0.009)	(1.759)		(0.448)		(0.240)	
4	0.007	0.742	4.626	2.764	-	1.576	151.193
	(0.011)	(1.840)	(34.170)	(0.450)		(0.240)	
5	-0.012	-	1540.556	-	2.708	1.179	152.867
	(0.010)		(658.146)		(0.564)	(0.148)	
6	-0.012	-	1455.415	2.705	2.730	1.459	155.800
	(0.009)		(677.966)	(0.515)	(0.548)	(0.229)	
7	-0.012	0.120	1453.191	2.704	2.731	1.458	155.802
	(0.010)	(1.713)	(678.866)	(0.515)	(0.549)	(0.230)	

Table 3.

Properties of uncertainty and volatility. This table displays quarterly statistics of realized volatility Q, the estimated $\text{vol}_t(\hat{\omega})$ (Eq. (22)) series, with $\hat{\omega} = 14.939$ and the estimated $\text{unc}_t(\hat{\nu})$ (Eq. (25)) series with $\hat{\nu} = 15.346$. Panel A reports means and standard deviations and Panel B reports correlations. Panel C reports correlations at the quarterly frequency among unc, the market excess return, and the Fama-French factors.

Panel A: Means and standard deviations of vol and unc

	Mean	Standard deviation
Q	0.006592	0.007634
$\text{vol}(\hat{\omega})$	0.005876	0.005428
unc(1)	0.000345	0.000233
$\text{unc}(\hat{\nu})$	0.000017	0.000016

Panel B: Correlations of market excess returns with vol and unc

	r_{et+1}	Q_{t+1}	Q_t	$\text{vol}_t(\hat{\omega})$	$\text{unc}_t(1)$	$\text{unc}_t(\hat{\nu})$
r_{et+1}	1.000	-0.397	0.128	0.154	0.175	0.283
Q_{t+1}		1.000	0.202	0.312	0.051	0.004
Q_t			1.000	0.748	0.145	0.081
$\text{vol}_t(\hat{\omega})$				1.000	0.211	0.075
$\text{unc}_t(1)$					1.000	0.662
$\text{unc}_t(\hat{\nu})$						1.000

Panel C: Correlations of unc, the excess market return, and the Fama-French factors

	$\text{unc}_t(\hat{\nu})$	r_{et+1}	$r_{\text{hml}\,t+1}$	$r_{\text{smb}\,t+1}$	$r_{\text{umd}\,t+1}$	$r_{\text{str}\,t+1}$	$r_{\text{ltr}\,t+1}$
$\text{unc}_t(\hat{\nu})$	1.000	0.283	-0.073	0.240	-0.122	0.157	0.084
r_{et+1}		1.000	-0.482	0.478	-0.227	0.313	-0.146
$r_{\text{hml}\,t+1}$			1.000	-0.179	-0.092	-0.077	0.489
$r_{\text{smb}\,t+1}$				1.000	-0.358	0.383	0.237
$r_{\text{umd}\,t+1}$					1.000	-0.514	-0.151
$r_{\text{str}\,t+1}$						1.000	0.071
$r_{\text{ltr}\,t+1}$							1.000

Table 4.

The effect of uncertainty in output and the effect of uncertainty in corporate profits. This table displays estimates of regressions similar to those in Table 2, but with different measures for uncertainty. Quasi-likelihood standard errors are listed under the estimates in parentheses. In Panel A uncertainty is measured by the beta-weighted variance of constructed forecasts of real output growth between last quarter and different horizons. In Panel B uncertainty is measured by the beta-weighted variance of corporate profit growth forecasts between last quarter and different horizons. If the horizon is 1 (respectively 2, 3, or 4) then uncertainty in the growth between last quarter and this quarter (the next quarter, two quarters ahead, or three quarters ahead, respectively) is considered.

Panel A: The effect of uncertainty in constructed real output growth forecasts

Horizon	b	θ	$\log\omega$	$\log\nu$	σ^2	Log likelihood
1	0.008	166.650	2.675	0.319	1.540	151.626
	(0.007)	(173.054)	(0.528)	(0.361)	(0.246)	
2	0.010	123.087	2.745	0.074	1.570	151.180
	(0.008)	(402.416)	(0.483)	(0.651)	(0.242)	
3	0.017	-4653.506	2.808	1.452	1.583	151.462
	(0.009)	(6298.592)	(0.448)	(0.438)	(0.234)	
4	0.020	-69343.288	2.737	3.404	1.543	152.305
	(0.009)	(84256.607)	(0.414)	(1.552)	(0.221)	

Panel B: The effect of uncertainty in constructed real corporate profit growth forecasts

Horizon	b	θ	$\log\omega$	$\log\nu$	σ^2	Log likelihood
1	0.020	-2.922	2.838	-29.102	1.583	151.887
	(0.006)	(0.837)	(0.436)	(4.4×10^7)	(0.227)	
2	0.003	263.542	2.831	2.630	1.574	152.181
	(0.008)	(182.651)	(0.457)	(0.494)	(0.243)	
3	-0.008	930.048	3.009	2.760	1.537	156.436
	(0.008)	(234.035)	(0.379)	(0.238)	(0.218)	
4	-0.010	1551.778	2.759	2.796	1.480	155.492
	(0.009)	(807.618)	(0.481)	(0.591)	(0.228)	

Table 5.

Alternative specifications of the uncertainty regressions. This table displays estimates of regressions similar to those in Table 2 except the specification of unc_t is different. Quasi-likelihood standard errors are listed under the estimates in parentheses. A dash, '-', indicates that the parameter was fixed at zero. In Panel A, uncertainty is measured with a symmetric normal weighted variance of the same constructed market return forecast as in Table 2. In Panel B, non-symmetric cross-sectional weights are allowed and uncertainty is measured with a beta-weighted variance of the same constructed market return forecast as in Table 2 with two free parameters α and χ. In Panel C, uncertainty is measured by a beta-weighted variance of constructed market return forecasts when the long-term horizon is three quarters rather than four quarters. The specification numbers for each row correspond to the specification numbers in Table 2.

Panel A: Normal weighted variance

Specification	b	τ	θ	$\log\omega$	$\log\xi$	σ^2	Log likelihood
6	-0.011	-	1546.979	2.699	-2.113	1.458	155.763
	(0.009)		(654.824)	(0.519)	(0.239)	(0.230)	
7	-0.012	0.121	1544.776	2.698	-2.113	1.457	155.764
	(0.010)	(1.706)	(655.849)	(0.519)	(0.239)	(0.230)	

Panel B: Non-symmetric cross-sectional weights

Specification	b	τ	θ	$\log\omega$	$\log\alpha$	$\log\chi$	σ^2	Log likelihood
6	-0.012	-	2281.821	2.743	3.096	3.241	1.458	156.321
	(0.009)		(1386.986)	(0.550)	(0.761)	(0.788)	(0.238)	
7	-0.012	-0.085	2290.205	2.743	3.099	3.244	1.458	156.322
	(0.010)	(1.716)	(1397.800)	(0.554)	(0.759)	(0.786)	(0.238)	

Panel C: Uncertainty in the constructed market return with a long-term horizon of three

Specification	b	τ	θ	$\log\omega$	$\log\nu$	σ^2	Log likelihood
6	-0.009	-	899.286	2.977	2.713	1.519	156.630
	(0.008)		(231.007)	(0.387)	(0.219)	(0.218)	
7	-0.009	0.046	899.091	2.977	2.713	1.519	156.763
	(0.009)	(1.734)	(230.739)	(0.391)	(0.221)	(0.220)	

Table 6.

Fixed weighting schemes. This table displays estimates of regressions similar to those in Table 2 except that the weights for unc_t are fixed at many different values. Quasi-likelihood standard errors are listed under the estimates in parentheses. If there is no standard error present then the parameter is fixed and not estimated. In this case the value of the parameter in the estimate column is the value at which it is fixed. In Panel A, we fix $\log \nu$ at several different values. In Panel B we measure uncertainty with a truncated variance in which the lowest p and highest p percent of forecasts are discarded and flat weights are used on the middle $(1 - 2p)$ percent of forecasts. In Panel C we measure uncertainty with Winsorized forecasts in which the lowest p and highest p percent of forecasts are replaced with the lowest and highest forecasts in the middle $(1 - 2p)$ percent of forecasts. Note that in Panels B and C, $\log \nu$ does not appear in our specification and hence can not be estimated.

Panel A: Fixed $\log \nu$

Specification	b	τ	θ	$\log \omega$	$\log \nu$	σ^2	Log likelihood
1	-0.004	0.259	144.161	2.754	1.000	1.547	152.344
	(0.011)	(1.767)	(97.594)	(0.453)		(0.237)	
2	-0.009	0.131	356.786	2.759	1.500	1.518	153.731
	(0.011)	(1.748)	(146.648)	(0.462)		(0.233)	
3	-0.012	0.102	701.305	2.750	2.000	1.488	155.026
	(0.010)	(1.739)	(214.382)	(0.476)		(0.228)	
4	-0.011	0.142	1779.780	2.680	3.000	1.454	155.716
	(0.010)	(1.696)	(432.869)	(0.529)		(0.230)	
5	-0.004	0.340	2877.538	2.632	4.000	1.464	154.609
	(0.009)	(1.653)	(697.124)	(0.564)		(0.237)	
6	0.002	0.495	3583.219	2.674	5.000	1.503	153.290
	(0.009)	(1.658)	(1071.177)	(0.525)		(0.238)	

Panel B: Truncated variance

Specification	b	τ	θ	$\log \omega$	p	σ^2	Log likelihood
1	-0.006	0.201	135.773	2.782	10	1.546	152.765
	(0.011)	(1.777)	(75.561)	(0.454)		(0.232)	
2	-0.008	0.119	217.263	2.783	15	1.530	153.501
	(0.010)	(1.755)	(95.362)	(0.452)		(0.233)	
3	-0.012	0.064	395.563	2.725	20	1.479	155.096
	(0.011)	(1.729)	(133.198)	(0.455)		(0.221)	
4	-0.013	0.121	608.248	2.746	25	1.474	155.618
	(0.010)	(1.757)	(178.836)	(0.503)		(0.231)	
5	-0.012	0.064	956.099	2.799	30	1.483	155.924
	(0.010)	(1.767)	(248.391)	(0.463)		(0.224)	
6	-0.002	0.311	1836.949	2.625	40	1.473	154.094
	(0.009)	(1.659)	(483.508)	(0.622)		(0.248)	

Panel C: Winsorization

Specification	b	τ	θ	$\log \omega$	p	σ^2	Log likelihood
1	-0.003	0.290	74.634	2.757	10	1.553	152.105
	(0.011)	(1.793)	(57.550)	(0.456)		(0.233)	
2	-0.005	0.104	121.159	2.768	15	1.539	152.878
	(0.010)	(1.765)	(64.195)	(0.464)		(0.235)	
3	-0.009	0.089	199.377	2.752	20	1.511	153.945
	(0.010)	(1.752)	(81.381)	(0.447)		(0.225)	
4	-0.011	0.148	323.817	2.724	25	1.481	154.998
	(0.011)	(1.750)	(112.623)	(0.528)		(0.236)	
5	-0.010	0.091	486.001	2.841	30	1.515	154.985
	(0.010)	(1.803)	(154.000)	(0.445)		(0.227)	
6	-0.001	0.182	1173.149	2.693	40	1.508	153.301
	(0.009)	(1.719)	(404.200)	(0.580)		(0.247)	

Table 7.

The impact of uncertainty on volatility. This table displays estimates of regressions in which uncertainty is permitted to have an effect on quarterly volatility. We run several versions of the nonlinear regression

$$r_{et+1} = b + \tau \text{vol}_t(\omega) + \theta \,\text{unc}_t(\nu) + \epsilon_{t+1}$$

of quarterly excess returns r_{et+1} on a constant, the measure of volatility, $\text{vol}_t(\omega)$, specified in Eq. (22), and the measure of uncertainty, $\text{unc}_t(\nu)$, specified in Eq. (25). The variance of the error term, ϵ_{t+1}, is

$$\sigma_v^2 \text{vol}_t(\omega) + \sigma_u^2 \text{unc}_t(\nu)$$

where σ_u^2 and σ_v^2 are constants which we estimate. Quasi-likelihood standard errors are listed under the estimates in parentheses. If there is no standard error present then the parameter is fixed at the listed value. In Panel A we estimate $\log \nu$ along with other parameters, while in Panel B we fix $\log \nu$.

Panel A: Impact when $\log \nu$ is estimated

Specification	b	τ	θ	$\log \omega$	$\log \nu$	σ_v^2	σ_u^2	Log likelihood
1	0.011	-	-	-	0.924	1.051	9.105	148.053
	(0.006)				(0.383)	(0.219)	(7.301)	
2	0.010	-	-	3.215	1.306	1.172	23.693	152.967
	(0.006)			(0.475)	(0.584)	(0.260)	(18.992)	
3	0.009	0.491	-	3.187	1.315	1.170	23.726	152.997
	(0.009)	(1.781)		(0.498)	(0.582)	(0.265)	(19.112)	
4	0.008	0.202	13.496	3.155	-	1.191	5.220	152.402
	(0.011)	(1.905)	(33.321)	(0.599)		(0.287)	(3.608)	
6	-0.011	-	1196.519	2.874	2.553	1.389	21.063	155.822
	(0.011)		(2073.468)	(1.495)	(1.519)	(0.648)	(153.534)	
7	-0.012	0.122	1196.632	2.869	2.555	1.389	20.856	155.824
	(0.012)	(1.729)	(2070.585)	(1.477)	(1.518)	(0.646)	(152.610)	

Panel B: Impact when $\log \nu$ is fixed

Specification	b	τ	θ	$\log \omega$	$\log \nu$	σ_v^2	σ_u^2	Log likelihood
1	0.008	0.202	13.496	3.155	0.000	1.191	5.220	152.402
	(0.011)	(1.905)	(33.321)	(0.599)		(0.287)	(3.608)	
2	-0.003	0.092	143.135	3.150	1.000	1.155	15.640	153.930
	(0.012)	(1.793)	(98.851)	(0.513)		(0.265)	(9.029)	
3	-0.008	0.084	331.960	3.101	1.500	1.209	22.958	154.834
	(0.011)	(1.768)	(160.395)	(0.481)		(0.279)	(16.644)	
4	-0.010	0.083	651.765	3.038	2.000	1.286	27.866	155.534
	(0.011)	(1.762)	(248.298)	(0.533)		(0.295)	(32.160)	
5	-0.011	0.142	1779.780	2.680	3.000	1.454	0.000	155.716
	(0.012)	(1.832)	(965.816)	(1.377)		(0.429)	(207.629)	
6	-0.004	0.340	2877.538	2.632	4.000	1.464	0.000	154.609
	(0.009)	(1.640)	(754.125)	(0.646)		(0.279)	(222.246)	
7	0.002	0.495	3583.219	2.674	5.000	1.503	0.000	153.290
	(0.010)	(1.675)	(2795.576)	(0.934)		(0.341)	(869.256)	

Table 8.

Time-varying uncertainty aversion. This table displays estimates of regressions in which uncertainty aversion is permitted to be time-varying. We run the nonlinear regression

$$r_{et+1} = b + \left(\sum_{k \in \Omega} \theta_k \, \mathbf{1}_{t,k} \, \mathrm{unc}_t(\nu) \right) + \epsilon_{t+1}$$

of quarterly excess returns r_{et+1} on a constant and the measure of uncertainty, $\mathrm{unc}_t(\nu)$, specified in Eq. (25). Here

$$\Omega = \{1968{:}4\text{-}1977{:}2,\ 1977{:}3\text{-}1986{:}1,\ 1986{:}2\text{-}1994{:}4,\ 1995{:}1\text{-}2003{:}3\}$$

and uncertainty aversion assumes four different values—one value for each of the periods in Ω. The variance of the error term, ϵ_{t+1}, is $\sigma^2 \mathrm{vol}_t(\omega)$ where σ^2 and ω are constants which we estimate. Quasi-likelihood standard errors are listed under the estimates in parentheses.

b	$\theta_{1968:4-1977:2}$	$\theta_{1977:3-1986:1}$	$\theta_{1986:2-1994:4}$	$\theta_{1995:1-2003:3}$	$\log \omega$	$\log \nu$	σ^2	Log likelihood
-0.013	764.199	1332.698	1505.093	2179.550	2.645	2.660	1.427	156.540
(0.010)	(1083.007)	(571.077)	(644.735)	(952.561)	(0.594)	(0.469)	(0.242)	

Table 9.

Time series properties of uncertainty and volatility. The table displays estimates of regressions

$$y_t = b + \sum_{i=1}^{n} \psi_i \ \mathrm{vol}_{t-i}(\hat{\omega}) + \sum_{i=1}^{m} \varphi_i \ \mathrm{unc}_{t-i}(\hat{\nu}) + e_t$$

on past predictors of volatility $\mathrm{vol}_{t-i}(\hat{\omega})$ and measures of past uncertainty $\mathrm{unc}_{t-i}(\hat{\nu})$ where the variance of the error term, e_t, is assumed constant over time. Here $\hat{\omega} = 14.939$ and $\hat{\nu} = 15.346$. We vary the dependent variables and the values of n and m. Ordinary least squares standard errors are listed under the estimates in parentheses. In Panel A we set $y_t = \mathrm{unc}_t(\hat{\nu})$ and consider regressions of uncertainty on past predictors of volatility and past uncertainty. In Panel B we set $y_t = Q_t$ and consider regressions of realized volatility on past predictors of volatility and past uncertainty. In Panel C we set $y_t = \mathrm{vol}_t(\hat{\omega})$ and consider regressions of predictors of volatility on past predictors of volatility and past uncertainty. We report estimates of the coefficients $\{\psi\}_{i=1}^{n}$ and $\{\varphi\}_{i=1}^{m}$ for various values of n and m.

Panel A: Regressions of uncertainty on past predictors of volatility and uncertainty

b	ψ_1	ψ_2	ψ_3	φ_1	φ_2	σ^2	Log likelihood
0.000	0.00043	-	-	0.288	-	2.3×10^{-10}	1325.435
(0.000)	(0.00039)			(0.119)		(0.9×10^{-10})	
0.000	0.00044	-	-	0.211	0.238	2.2×10^{-10}	1329.742
(0.000)	(0.00038)			(0.106)	(0.066)	(0.8×10^{-10})	

Panel B: Regressions of realized volatility on past predictors of volatility and uncertainty

b	ψ_1	ψ_2	ψ_3	φ_1	φ_2	σ^2	Log likelihood
0.004	0.440	-	-	-10.264	-	5.323×10^{-5}	479.700
(0.001)	(0.088)			(32.786)		(3.455×10^{-10})	
0.003	0.388	0.033	0.174	-	-	5.230×10^{-5}	480.914
(0.001)	(0.092)	(0.064)	(0.087)			(3.470×10^{-5})	

Panel C: Regressions of predictors of volatility on past predictors of volatility and uncertainty

b	ψ_1	ψ_2	ψ_3	φ_1	φ_2	σ^2	Log likelihood
0.004	0.290	-	-	-0.620	-	2.733×10^{-5}	525.365
(0.001)	(0.065)			(18.734)		(0.628×10^{-5})	
0.002	0.190	0.077	0.339	-	-	2.359×10^{-5}	535.438
(0.001)	(0.064)	(0.057)	(0.130)			(0.501×10^{-5})	

Table 10.

Summary statistics on risk and uncertainty sorted portfolios. This table presents summary statistics on portfolios sorted on sensitivity to risk and uncertainty. Portfolios are constructed from rolling sample regressions of Eq. (39), where regressions are rolled forward each quarter throughout the life of the stock. Only firms with at least 20 quarters of return data are used in the sample (N = 14,252). Once sensitivities are obtained, firms are sorted first into three portfolios based on those sensitivities to uncertainty in each quarter and then sorted again into three portfolios based on sensitivities to risk. Portfolios are constructed by value-weighting the stocks within the portfolio each quarter. The sample span ranges from fourth quarter 1973 through fourth quarter 2003; however, in order to construct rolling samples, data are used from fourth quarter 1968. In addition to summary statistics on the returns to the portfolios presented in Panel A, summary statistics on the weights of each portfolio are described in Panel B. The weights do not sum to 100% since the firms analyzed are only a fraction of the entire CRSP universe in each quarter.

Panel A: Portfolio returns

		Average returns			**Standard deviation**		
		Uncertainty			Uncertainty		
		Low	*Med*	*High*	*Low*	*Med*	*High*
	Low	1.01633	1.02382	1.01902	0.05776	0.06551	0.10581
Volatility	*Med*	1.01535	1.01829	1.02806	0.08951	0.09110	0.13635
	High	1.02070	1.02281	1.04302	0.15469	0.13446	0.20693

Panel B: Portfolio weights

		Average weights			**Standard deviation**		
		Uncertainty			Uncertainty		
		Low	*Med*	*High*	*Low*	*Med*	*High*
	Low	0.10887	0.11474	0.08006	0.06575	0.05342	0.04463
Volatility	*Med*	0.12195	0.15562	0.09344	0.06061	0.04251	0.04752
	High	0.05740	0.10065	0.04511	0.03133	0.04579	0.01949

Table 11.

GMM estimates of the prices of factors. This table displays GMM estimates of the prices of factors for various versions of the joint time series and cross-sectional system

$$\begin{aligned} r_{kt+1} &= a_k + \beta_k' f_{t+1} + \varrho_{kt+1}, \qquad\qquad k = 1 \dots n, \\ Er_{kt+1} &= \beta_k' \lambda + \iota_k \end{aligned}$$

where a_k is a time-series pricing error, ι_k is a cross-sectional pricing error, β_k is a vector of regression coefficients, f_{t+1} is a vector of factors, and λ is a vector of prices. (Herein we describe the estimation generally for any possible set of factors. See Eq. (42) in the text for a simpler description in the special case in which only risk and uncertainty factors are present.) In some of our specifications we include a market risk factor, an alternative measure of the market risk factor, a market uncertainty factor, the market excess return, the growth (HML) factor, the size (SMB) factor, the momentum (UMD) factor, the short-term reversal (STR) factor, and the long-term reversal (LTR) factor:

$$\begin{aligned} \beta_k &= \begin{bmatrix} \beta_{vk} & \bar{\beta}_{vk} & \beta_{uk} & \beta_{mk} & \beta_{\text{hml}\,k} & \beta_{\text{smb}\,k} & \beta_{\text{umd}\,k} & \beta_{\text{str}\,k} & \beta_{\text{ltr}\,k} \end{bmatrix}', \\ \lambda &= \begin{bmatrix} \lambda_v & \bar{\lambda}_v & \lambda_u & \lambda_m & \lambda_{\text{hml}} & \lambda_{\text{smb}} & \lambda_{\text{umd}} & \lambda_{\text{str}} & \lambda_{\text{ltr}} \end{bmatrix}', \\ f_{t+1} &= \begin{bmatrix} \hat{\tau}\text{vol}_t(\hat{\omega}) + \hat{\epsilon}_{et+1} & \hat{\tau}\text{vol}_t(\hat{\omega}) & \hat{\theta}\text{unc}_t(\hat{\nu}) & r_{mt+1} & r_{\text{hml}\,t+1} & r_{\text{smb}\,t+1} & r_{\text{umd}\,t+1} & r_{\text{str}\,t+1} & r_{\text{ltr}\,t+1} \end{bmatrix}'. \end{aligned}$$

The moment conditions for asset k are

$$E \begin{bmatrix} r_{kt+1} - a_k - \beta_k' f_t \\ \left(r_{kt+1} - a_k - \beta_k' f_t\right) \otimes f_t \\ r_{kt+1} - \beta_k' \lambda \end{bmatrix} = 0.$$

The moment conditions for all assets are combined and GMM estimates of the prices of factors are listed below for the fixed weighting matrix described in Section 6.2. GMM standard errors are listed in parentheses below estimates and are computed using the method of Newey and West (1987) with eight lags. Parameters without standard errors are fixed at zero. (When λ_x is fixed at zero for some factor x, we remove the corresponding regression coefficient β_{xk} from the vector β_k for each k.) Estimates of a_k and β_k are not displayed but are available upon request. The nonlinear parameters ω and ν are fixed at their QMLE estimates of 14.939 and 15.346 throughout this table. The asset return data are quarterly from 1969:1 to 2003:4 and consist of real excess returns for 130 portfolios which include 25 portfolios sorted on size and book-to-market, 25 portfolios sorted on size and short-term reversal, 25 portfolios sorted on size and momentum, 25 portfolios sorted on size and long-term reversal, ten portfolios sorted on earnings-to-price, ten portfolios sorted on dividend-to-price, and ten portfolios sorted on cash flow-to price.

Specification	λ_{vol}	$\bar{\lambda}_{vol}$	λ_{unc}	λ_m	λ_{hml}	λ_{smb}	λ_{umd}	λ_{ltr}	λ_{str}	J-stat	p-value
1	-0.01106	-	0.02717	-	-	-	-	-	-	16.42010	1.00000
	(0.00698)		(0.00293)								
2	-	-0.00001	0.02713	-	-	-	-	-	-	16.42078	1.00000
		(0.00001)	(0.00291)								
3	-	-	-	0.01611	-	-	-	-	-	16.40791	1.00000
				(0.00669)							
4	-	-	-	0.01612	0.01938	-0.00361	0.02367	0.01221	0.02252	16.38972	1.00000
				(0.00668)	(0.00548)	(0.00551)	(0.00473)	(0.00525)	(0.00486)		
5	-	-	0.02708	0.01611	0.01939	-0.00353	0.02327	0.01212	0.02255	16.39684	1.00000
			(0.00298)	(0.00667)	(0.00549)	(0.00549)	(0.00477)	(0.00509)	(0.00490)		
6	-0.01096	-	0.02708	-	0.01939	-0.00353	0.02327	0.01212	0.02255	16.39684	1.00000
	(0.00694)		(0.00298)		(0.00549)	(0.00549)	(0.00477)	(0.00509)	(0.00490)		
7	-	-0.00001	0.02709	-	0.01938	-0.00352	0.02328	0.01229	0.02259	16.41269	1.00000
		(0.00001)	(0.00296)		(0.00550)	(0.00549)	(0.00478)	(0.00502)	(0.00490)		

Table 12.

GMM estimates of the stochastic discount factor. This table estimates

$$E s_{t+1} r_{jt+1} = 1 \qquad j \ldots n$$

for various formulations of the stochastic discount factor:

$$s_{t+1} = a + s' f_{t+1}$$

where

$$s = \begin{bmatrix} s_v & \bar{s_v} & s_u & s_m & s_{\mathrm{hml}} & s_{\mathrm{smb}} & s_{\mathrm{umd}} & s_{\mathrm{str}} & s_{\mathrm{ltr}} \end{bmatrix}',$$

$$f_{t+1} = \begin{bmatrix} \hat{\tau}\mathrm{vol}_t(\hat{\omega}) + \epsilon_{et+1} & \hat{\tau}\mathrm{vol}_t(\omega) & \hat{\theta}\mathrm{unc}_t(\nu) & r_{mt+1} & r_{\mathrm{hml}\,t+1} & r_{\mathrm{smb}\,t+1} & r_{\mathrm{umd}\,t+1} & r_{\mathrm{str}\,t+1} & r_{\mathrm{ltr}\,t+1} \end{bmatrix}'.$$

The assets considered are the same 130 portfolios used in Table 11 though the returns are gross real returns rather than real excess returns. (The factors are real excess returns.) GMM standard errors are listed in parentheses below estimates and are computed using the method of Newey and West (1987) with eight lags. If there is no standard error present then the parameter is fixed at the listed value. A dash, '-', indicates that the parameter was fixed at zero. In Panel A, the fixed weighting matrix proposed by Hansen and Jagannathan (1997) is employed and we report the HJ-dist and its standard error. In Panel B, the optimal GMM weighting matrix is employed and we report the *J*-stat and its *p*-value.

Panel A: The Hansen and Jagannathan (1997) weighting matrix

Spec.	a	s_v	$\bar{s_v}$	s_u	s_m	s_{hml}	s_{smb}	s_{umd}	s_{ltr}	s_{str}	$\log\omega$	$\log\nu$	HJ $_{\mathrm{dist}}$
1	2.319	-0.975	-	-52.384	-	-	-	-	-	-	2.704	2.731	6.856
	(0.334)	(1.586)		(16.188)									(0.571)
2	2.027	-	4493.637	-53.443	-	-	-	-	-	-	2.704	2.731	6.851
	(0.376)		(3557.785)	(16.072)									(0.574)
3	1.084	-	-	-	-4.797	-	-	-	-	-	2.704	2.731	6.948
	(0.049)				(1.368)								(0.568)
4	1.864	-	-	-	-11.330	-11.994	7.425	-10.865	-3.472	-11.980	2.704	2.731	6.871
	(0.166)				(2.220)	(2.940)	(2.881)	(2.139)	(3.563)	(3.235)			(0.577)
5	3.031	-	-	-51.368	-7.891	-11.462	9.219	-10.714	-1.554	-11.100	2.704	2.731	6.782
	(0.394)			(16.997)	(2.366)	(3.824)	(3.784)	(2.464)	(5.345)	(3.832)			(0.580)
6	3.126	-7.891	-	-59.259	-	-11.462	9.219	-10.714	-1.554	-11.100	2.704	2.731	6.782
	(0.402)	(2.366)		(17.281)		(3.824)	(3.784)	(2.464)	(5.345)	(3.832)			(0.580)
7	2.549	-	6845.720	-55.725	-	-6.706	5.370	-9.322	-1.152	-15.790	2.704	2.731	6.791
	(0.410)		(5460.552)	(16.509)		(3.963)	(3.299)	(2.660)	(4.929)	(6.630)			(0.583)
8	1.857	-	8380.267	-52.586	-	-	-	-	-	-	4.491	2.694	6.822
	(0.384)		(4399.430)	(35.915)							(1.271)	(0.578)	(0.579)
9	3.132	-	-	-40.718	-7.858	-11.473	9.342	-10.712	-1.383	-11.297	2.704	2.423	6.781
	(0.404)			(30.969)	(2.440)	(3.873)	(3.974)	(2.521)	(5.440)	(3.962)		(0.663)	(0.580)
10	2.470	-	10703.046	-46.665	-	-5.212	7.024	-9.166	-1.466	-18.179	4.393	2.496	6.754
	(0.426)		(6143.905)	(32.043)		(4.056)	(3.894)	(2.973)	(5.246)	(7.154)	(1.143)	(0.621)	(0.589)

Panel B: Optimal GMM weighting matrix

Spec.	a	s_v	$\bar{s_v}$	s_u	s_m	s_{hml}	s_{smb}	s_{umd}	s_{ltr}	s_{str}	$\log\omega$	$\log\nu$	*J*-stat	*p*-value
1	2.337	-1.664	-	-51.896	-	-	-	-	-	-	2.704	2.731	16.414	0.000
	(0.012)	(0.086)		(0.372)										
2	2.080	-	3594.745	-53.478	-	-	-	-	-	-	2.704	2.731	16.420	0.000
	(0.028)		(346.567)	(0.304)										
3	1.102	-	-	-	-5.110	-	-	-	-	-	2.704	2.731	16.400	0.000
	(0.004)				(0.067)									
4	1.913	-	-	-	-11.698	-12.336	8.166	-10.612	-3.511	-12.747	2.704	2.731	16.391	0.000
	(0.010)				(0.121)	(0.192)	(0.170)	(0.153)	(0.136)	(0.217)				
5	3.068	-	-	-49.192	-8.516	-12.684	8.341	-11.403	-1.426	-11.454	2.704	2.731	16.393	0.000
	(0.021)			(0.663)	(0.226)	(0.335)	(0.264)	(0.194)	(0.371)	(0.259)				
6	3.171	-8.516	-	-57.708	-	-12.683	8.341	-11.403	-1.426	-11.454	2.704	2.731	16.395	0.000
	(0.022)	(0.226)		(0.638)		(0.335)	(0.264)	(0.194)	(0.371)	(0.259)				
7	2.539	-	7380.880	-56.989	-	-6.704	5.589	-9.484	-1.130	-15.151	2.704	2.731	16.413	0.000
	(0.028)		(397.807)	(0.858)		(0.259)	(0.345)	(0.186)	(0.346)	(0.454)				
8	1.736	-	9249.235	-51.730	-	-	-	-	-	-	4.262	2.717	16.420	0.000
	(0.025)		(356.096)	(0.968)							(0.065)	(0.014)		
9	3.206	-	-	-32.700	-8.383	-12.481	9.207	-11.624	-1.627	-11.809	2.704	2.230	16.361	0.000
	(0.023)			(1.103)	(0.212)	(0.350)	(0.267)	(0.216)	(0.296)	(0.341)		(0.031)		
10	2.480	-	10737.625	-46.165	-	-5.603	6.842	-9.982	-1.773	-19.144	4.116	2.471	16.410	0.000
	(0.026)		(300.410)	(0.946)		(0.257)	(0.305)	(0.221)	(0.203)	(0.400)	(0.066)	(0.020)		

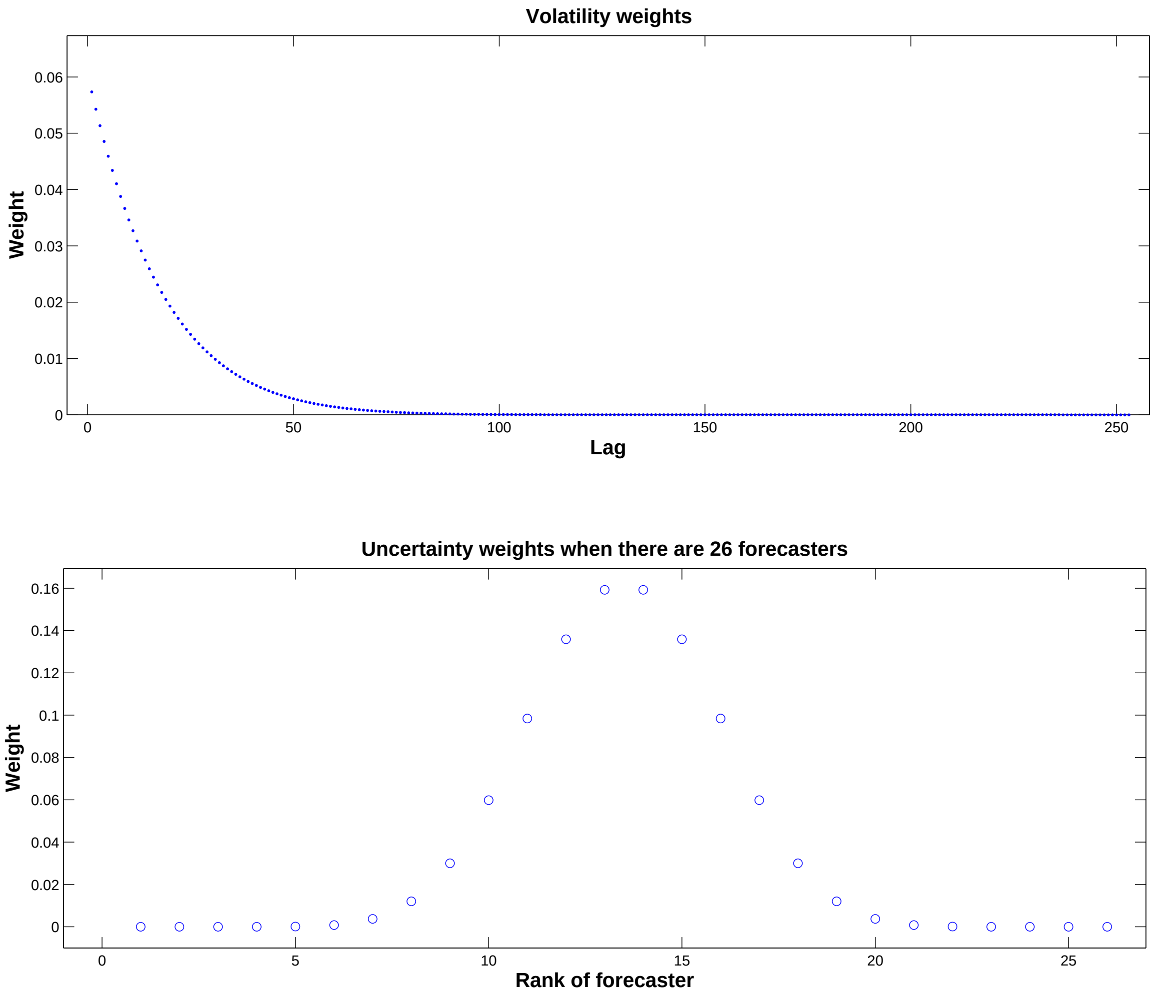

Fig. 1. Volatility and uncertainty weights. QMLE estimates of the parameters appearing in Table 2 are used to compute the weights on daily lagged volatility and the weights across forecasters. The weights appear in our measure of volatility, $\text{vol}_t(\omega)$, specified in Eq. (22), and our measure of uncertainty, $\text{unc}_t(\nu)$, specified in Eq. (25). The top graph displays the weights on lagged daily volatility when $\omega = 14.939$ and the bottom graph displays the weights on forecasters when $\nu = 15.346$. These are the estimates of ω and ν reported in specification seven of Table 2. In the top graph, the x-axis represents lagged trading days and the y-axis represent weights. The weight on daily volatility on the last day of the current quarter (the time t quarter) corresponds to $x = 1$ and is a little less than 0.06. The bottom graph displays the weights on forecasters for a quarter in which there are 26 available forecasters ($f_t = 26$). The weights on the lowest and highest indexed forecasters are nearly zero and the weights on the 13th and 14th indexed forecasters are about 0.16.

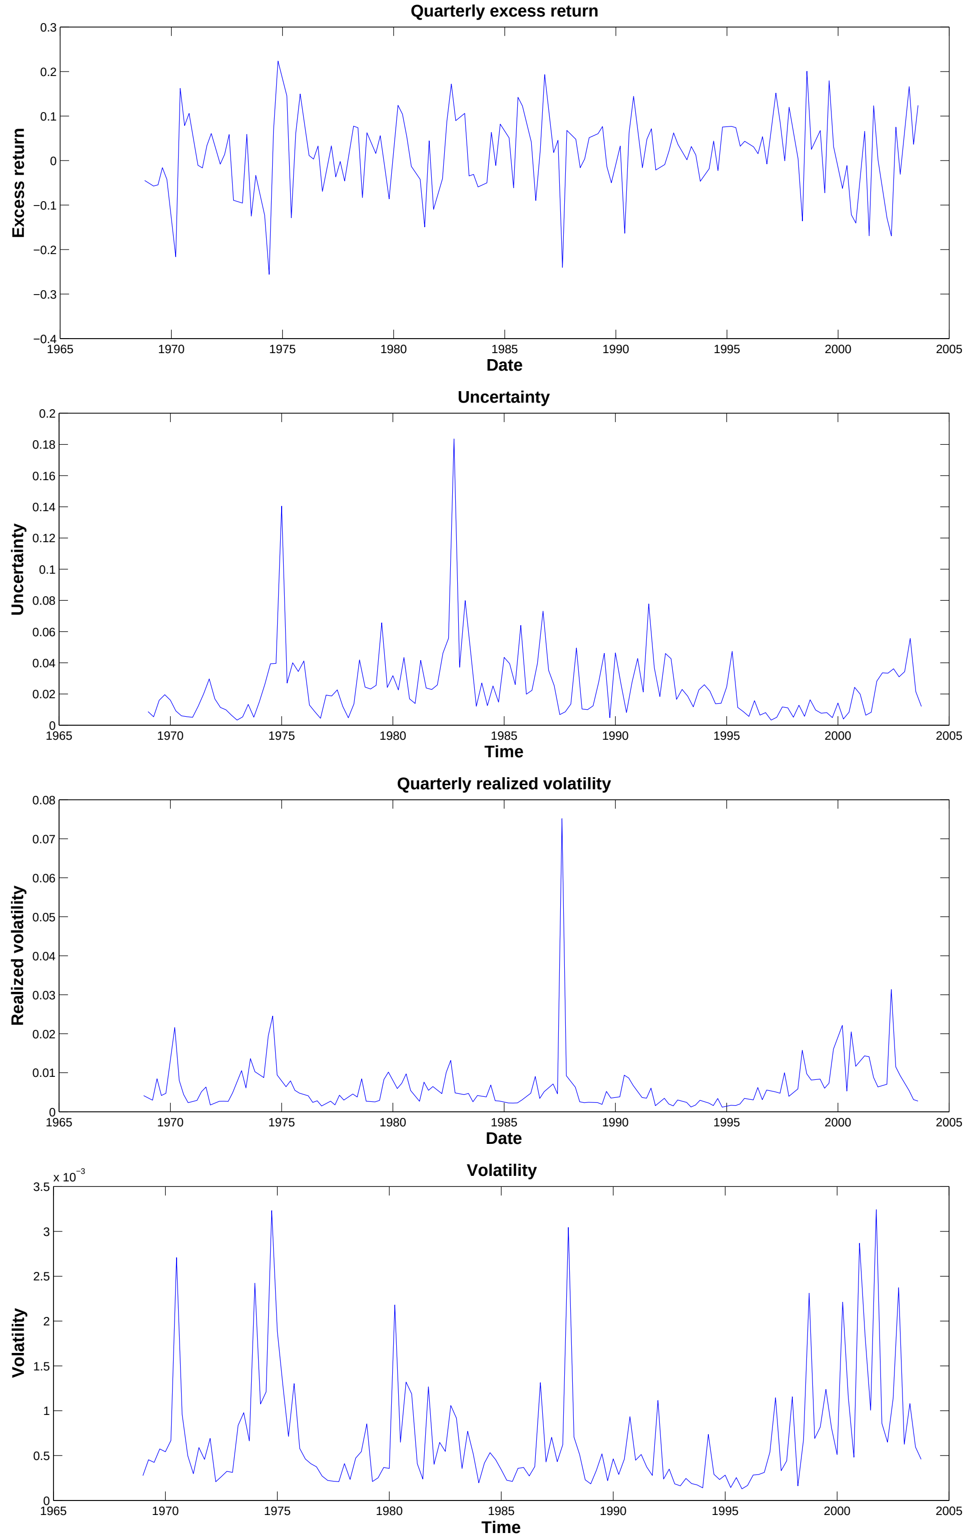

Fig. 2. Time series plots. The first figure displays a plot of the quarterly excess return r_{et} and the second figure displays a plot of uncertainty, $\theta \text{unc}_{t-1}(\nu)$, where $\theta = 1453.191$ and $\nu = 15.346$ are set at their QMLE estimates from specification seven of Table 2. The third figure displays a plot of quarterly realized volatility, Q_t, and the fourth figure displays a plot of volatility, $\tau \text{vol}_{t-1}(\omega)$, where $\tau = 0.120$ and $\omega = 14.939$ are set at their QMLE estimates from specification seven of Table 2.

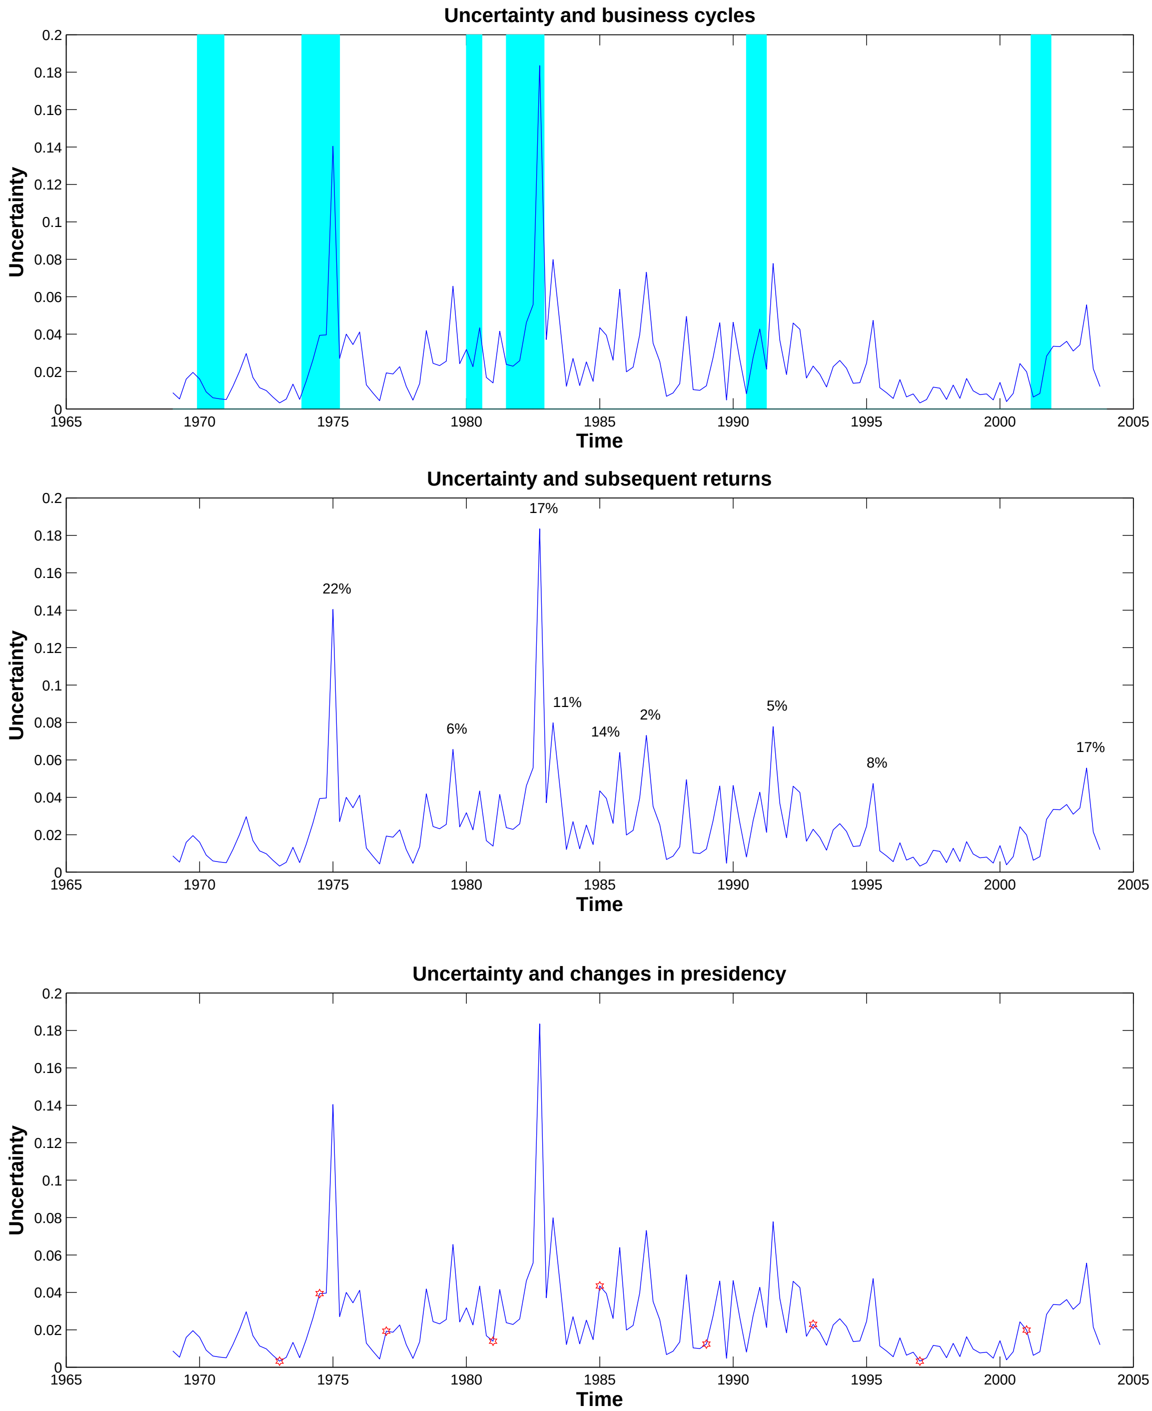

Fig. 3. Uncertainty and events. All of the figures plot $\theta\text{unc}_{t-1}(\nu)$ with other events where $\theta = 1453.191$ and $\nu = 15.346$ are set at their QMLE estimates from specification seven of Table 2. The top figure includes recessions (as defined by the NBER) in the shaded regions. The middle figure displays the quarterly excess returns (r_{et}) at the peaks of uncertainty [$\theta\text{unc}_{t-1}(\nu)$]. The bottom figure indicates changes in presidency with a circle. The circle corresponds to the quarter that a new president assumes office. The plotted uncertainty curve is based on forecasts stated during the previous quarter, which is the quarter that the president was elected (with the exception of President Ford).