The Impact of Risk and Uncertainty on Expected Returns

EVAN W. ANDERSON, ERIC GHYSELS, and JENNIFER L. JUERGENS*

August 19, 2005

ABSTRACT

We study asset pricing when agents face model uncertainty and empirically demonstrate that model uncertainty matters for asset pricing. We measure the amount of model uncertainty in the economy with the disagreement of professional forecasters, attributing different weights to each professional forecaster. The weighting scheme is estimated via a method that is inspired by recent work on MIDAS regressions. We run regressions representing the typical risk-return trade-off, where risk is represented by conditional volatility and augment these regressions with a measure of model uncertainty. We find stronger empirical evidence for a model uncertainty-return trade-off than for the traditional risk-return trade-off.

*Evan Anderson: Northern Illinois University, Department of Economics, Zulauf 515, Dekalb, IL 60115, phone: (919) 604-0067, email: ewanderson@niu.edu. Eric Ghysels: Department of Finance, Kenan-Flagler Business School, and Department of Economics, University of North Carolina – Chapel Hill, Chapel Hill, NC 27599-3305, phone: (919) 966-5325, email: eghysels@unc.edu. Jennifer L. Juergens: Arizona State University, W.P. Carey School of Business, Department of Finance, Tempe, AZ 85287, phone: (480) 965-3131, email: Jennifer.Juergens@asu.edu. We would like to thank Lars Hansen, Peter Hansen, Athanasios Orphanides, Simon Potter, Robert Rich and Tom Sargent for very helpful comments. Early versions of this paper were presented at the Board of Governors, the FDIC, the Federal Reserve Bank of New York, Northern Illinois University, Simon Fraser University and Stanford University. We thank seminar participants for helpful comments.

One of the most studied theoretical relationships in empirical finance states that the expected excess return of the market over a risk-free bond should vary positively and proportionally to the conditional volatility of the market return. Merton (1973) derived this theoretical relationship in a continuous time model in which all agents have power preferences when hedging concerns are negligible. This risk-return trade-off is so fundamental in financial economics that it could well be described as the "first fundamental law of finance." However, the empirical evidence for this relationship is mixed. Many researchers have failed to find a significant relationship between excess returns and conditional volatility. 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. Some researchers have even found the risk-return tradeoff to be negative. Campbell (1987), Nelson (1991), 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 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.

An important strand of recent research in finance contends that model uncertainty, in addition to risk, should matter for asset pricing. When agents are unsure of the correct model governing the market return they demand a higher premium in order to hold the market portfolio. The contribution of this paper is to empirically investigate the performance of asset pricing models when agents are uncertain of the correct model. We begin by theoretically showing, following work by Hansen, Sargent, and Tallarini (1999); Hansen and Sargent (2001); Anderson, Hansen, and Sargent (2003), Maenhout (2004); and Uppal and Wang (2003), that when agents have a concern for model misspecification the traditional risk-return regression is augmented so that both risk and uncertainty carry a positive premium. When hedging concerns are negligible and uncertainty is held fixed the expected excess return is proportional to risk. When hedging returns are negligible and risk is held fixed, the expected excess return is proportional to uncertainty. Next, we turn to the main focus of the paper and empirically investigate the importance of the uncertainty-return trade-off in conjunction with the risk-return trade-off. We find stronger empirical evidence for a uncertainty-return trade-off than for the traditional risk-return trade-off.

One of main features of our paper is to measure the amount of model uncertainty with data on disagreement of professional forecasters. We construct an empirical measure (a time series process) of model uncertainty, denoted unc_t. The relationship between the disagreement of professional forecasters and expected returns has been discussed in many recent papers, without a link to model uncertainty. A number of authors, including Anderson, Ghysels, and Juergens (2005) and Qu, Starks, and Yan (2003) find that more disagreement implies higher expected returns. In particular, Anderson, Ghysels, and Juergens (2005) observe that the disagreement factors (portfolios that long high disagreement stocks and short low disagreement 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 disagreement. Others, including Diether, Malloy, and Scherbina (2002) and Johnson (2004), find that higher dispersion stocks have lower future returns.

Several different rationales have been proposed to explain the negative effect of disagreement on expected returns. Anderson, Ghysels, and Juergens (2005) took the disagreement of forecasters about the future values of variables as an indication of heterogeneity in the beliefs of agents and showed how the disagreement is priced in a heterogeneous agents model with micro-foundations. Diether, Malloy, and Scherbina (2002) rationalize their findings that higher dispersion stocks have lower future returns with arguments from the short-sale constraints literature, in particular Miller (1977). They argue dispersion proxies for differences of opinions among traders where the only the most optimistic opinions are reflected, thereby driving up current prices. Johnson (2004) offers an alternative explanation to the findings of Diether, Malloy, and Scherbina (2002). In his model, levered firms may be able to reduce the cost of capital by increasing idiosyncratic risk of earnings volatility and subsequently the dispersion of earnings forecasts. Johnson views dispersion as a manifestation of idiosyncratic risk relating to the opacity in the underlying value of a stock.

This paper suggests an alternative explanation for why disagreement is priced. We interpret the disagreement among forecasters as arising because forecasters are using different models.[1] If forecasters are using different models then the economy as a whole faces a situation involving model uncertainty– there are a variety of models which could possibly correctly describe the world. We assume the agents inside our model think that each forecaster's predictions constitute one model that is possibly true. The predictions of all forecasters are weighted to form the beliefs of agents. This model uncertainty, measured via our empirical measure unc_t, is priced and is the channel that prices disagreement.

The existing literature tends to measures disagreement (which for us is equivalent to uncertainty) with flat weighted variances. We show that disagreement measured by flat weighted variances has no relationship with the expected market excess return. Disagreement only matters when extreme forecasts are disregarded. Extreme forecasts are forecasts whose predictions are far from other predictions. To construct unc_t, we measure disagreement with a flexible weighting scheme across forecasts which can accommodate assigning more or less weight to extreme forecasts. The weighting scheme is inspired by the recent work on MIDAS regressions. Parameters determining the weights are estimated by quasi-maximum likelihood in conjunction with other parameters. We find the parameters determining the optimal weights to be significant. Our estimates entail ignoring the extremes and placing all of the weight on the center of the distribution.

Most of the existing literature which studies the relationship between asset prices and disagreement focuses on individual stocks or portfolios representing a subset of the market. The existing literature measures disagreement with the dispersion of earnings forecasts made by financial analysts. In this paper we focus on the market excess return and use data on forecasts of aggregate corporate profits rather than earning forecasts. Focusing on aggregate forecasts allows us to use a longer time series of data than is used in most studies of disagreement.

This paper is organized as follows. Section I describes an economy with model uncertainty and derives a decomposition of excess returns into risk and model uncertainty components. Section II describes the flexible functional forms that will be used in this paper to measure volatility and uncertainty. Section III gives an overview of the data used in this paper including the data on beliefs. Section IV investigates risk-return and model uncertainty-return trade-offs. Section V concludes.

I. The theoretical impact of risk and uncertainty on excess returns

This section derives a decomposition of excess returns into risk and model uncertainty components. We consider a setup similar to Merton (1973) except that agents are worried about model misspecification as in Hansen and Sargent (2001). Our setup closely follows Hansen and Sargent except that we do not tie model uncertainty to conditional volatility.[2] Similar to Uppal and Wang (2003), we allow concerns for robustness to vary with states. Our approach also follows Maenhout (2004) in that we scale concerns for robustness by the value function. We consider a general equilibrium model in which all agents are alike and have power utility functions. Agents are allowed to invest in the market and a risk-free bond, though in equilibrium we require that agents fully invest all of their wealth in the market and do not hold bonds. Unlike previous papers we allow concerns for robustness to vary over time in interesting ways that are not linked to conditional volatility.

There is an underlying 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 that the price of the market, P_t, approximately follows the process

$$dP_t = d_t P_t \, dt + \zeta_t P_t \, dB_t \tag{2}$$

where $d_t = d(x_t)$ is a scalar and $\zeta_t = \zeta(x_t)$ is a row vector. 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 market return in excess of the risk-free bond, ψ_t is the fraction of wealth (possibly greater than one or less than zero) invested in the excess return, and c_t is consumption. Wealth approximately, and not necessarily exactly, follows equation (3) because the price of the market only approximately, and not necessarily exactly, follows the process in equation (2). We will call the processes in equations (1), (2), and (3) the reference model.

Agents believe that the reference model provides a reasonable approximation for the processes which govern the state, market return, and wealth though they worry that the approximation may be misspecified. In particular, they worry about whether or not a_t and d_t are the correct specifications of the conditional means. 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 market return 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, $\Delta_t = \Delta(x_t, y_t)$ is a matrix of the same dimension as Λ_t, and $\eta_t = \eta(x_t, y_t)$ is a vector 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) and the market return (ζ_t). Agents believe (and they also are correct) that the risk-free rate is ρ_t. In summary, they worry that the underlying state, market return and the evolution of the wealth are given by

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

$$dP_t = (\lambda_t - \eta_t g_t)\, P_t\, dt + \zeta_t P_t\, dB_t \tag{4b}$$

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

instead of by equations (1), (2), and (3). They are uncertain about the conditional mean of their wealth because they are uncertain about the conditional expected market return. We allow the functions Δ and η to depend on agent's wealth although in some important special cases they will only depend on the exogenous state x. The function g is an endogenous function of the exogenous state and wealth that will be determined below.

Agents consider a worst case specification for g_t that is constrained to be close to the reference model. We capture the requirement that g_t is close to the reference model by penalizing deviations from the reference model with the quadratic term

$$\frac{1}{2\theta} g_t' g_t \tag{5}$$

where θ is a scalar time-invariant non-negative parameter.[3] 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 and for some t

$$\Delta_t = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \qquad \eta_t = \begin{bmatrix} 0 & 100 \end{bmatrix}. \tag{6}$$

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 1.0 has the same penalty as perturbing the market return by 0.1. In this example the second element of x is presumed to be known exactly so that under no circumstances will agents consider perturbations in it. So, Δ and η allow us to capture the notion that agents may have more or less doubts about the conditional means of some variables than others.

In work by Hansen and Sargent 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 being tied to volatility and allow for the possibility that they depend on the state more flexibly. For example, there may be some state variables which have a high conditional variance but that 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. For example, during the gasoline crisis in the mid 1970's agents may have been willing to consider more perturbations in all variables than they were in mid 1990's. For the reasons discussed in Maenhout (2004) we will let Δ and η depend on wealth (see below for more details).[5]

The objective of agents is taken to be

$$\int_0^\infty \exp\left(-\delta t\right)\ \left[\frac{c_t^{1-\gamma}}{1-\gamma} + \frac{1}{2\theta} g_t' g_t\right] \tag{7}$$

where δ is the time discount rate. At any date, the first component of the objective is the utility obtained from consumption where γ is a parameter which is greater than zero and not equal to one. The second component penalizes deviations from the reference model, and is added rather than subtracted because we view the agents as minimizing the time zero expected value of equation (7) with respect to g in addition to maximizing it with respect to c and portfolio choices subject to equations (4a), (4b), and (4c).

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

$$J = \max_{\psi,c} \min_{g} \Big\{ \frac{c^{1-\gamma}}{1-\gamma} + \frac{1}{2\theta} g'g - \delta J + J_x' \left[a - \Delta g\right] + J_y \left[\psi\lambda\, y - \psi y \eta g + \rho\, y - c\right] + \frac{1}{2}\mathrm{tr}\left[\Lambda\Lambda' J_{xx}\right] + \frac{1}{2} J_{yy}\psi^2 \zeta\zeta' y^2 + \mathrm{tr}\left[\psi\zeta\Lambda' J_{xy} y\right] \Big\} \tag{8}$$

where we drop the t subscripts and the subscripts on J denote differentiation. In the limit as θ approaches zero, the functional equation (8) becomes the usual Hamilton-Jacobi equation studied by Merton (1973) and many subsequent researchers. 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 wealth. The minimizing choice of g is

$$g = \theta\Delta' J_x + \theta\eta' J_y \psi y. \tag{9}$$

Equation (9) shows how specifications of θ, Δ and η endogenously determine the perturbations of conditional means that agents worry about. The maximizing choice of the fraction of wealth to invest in the market, ψ, satisfies the first-order condition:

$$J_y \lambda y - J_y y \eta g + J_{yy} \psi \zeta \zeta' y^2 + \mathrm{tr}\left[\zeta \Lambda' J_{xy} y\right] = 0. \tag{10}$$

In equilibrium we require $\psi = 1$. Substituting in the right hand side of equation (9) for g, imposing $\psi = 1$ and rearranging terms allows us to write equation (10) as

$$\lambda = \gamma \zeta \zeta' + \theta \eta \eta' J_y y + \theta \eta \Delta' J_x - \mathrm{tr}\left[\zeta \Lambda' \frac{J_{xy}}{J_y}\right]. \tag{11}$$

In this paper we will consider specifications of Δ and η of the form

$$\Delta(x, y) = \frac{\bar{\Delta}(x)}{\sqrt{(1-\gamma)J(x,y)}} \qquad \eta(x, y) = \frac{\bar{\eta}(x)}{\sqrt{(1-\gamma)J(x,y)}} \tag{12}$$

which closely follow specifications in Maenhout (2004). Recall J is the agent's value function and we have previously assumed $\gamma \neq 1$. The functions $\bar{\Delta}$ and $\bar{\eta}$ can depend on the exogenous state x but not the agent's wealth y. The scaling by $1 - \gamma$ times the value function is the same scaling in Maenhout (2004) though Maenhout restricts $\bar{\Delta}$ and $\bar{\eta}$ to be constant. With this specification formula (11) simplifies to

$$\lambda = \gamma \zeta \zeta' + \theta \bar{\eta} \bar{\eta}' + \theta \bar{\eta} \bar{\Delta}' \frac{J_x}{J} - \mathrm{tr}\left[\zeta \Lambda' \frac{J_{xy}}{J_y}\right] \tag{13}$$

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

In our empirical work, appearing in later sections, we will consider a discrete time approximation to equation (13) 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

$$E_t r_{et+1} = \gamma V_t + \theta M_t + H_t \tag{14}$$

where

$$V_t = \zeta_t \zeta_t', \tag{15}$$

$$M_t = \bar{\eta}_t \bar{\eta}_t', \tag{16}$$

$$H_t = \theta \bar{\eta}_t \bar{\Delta}_t' \frac{J_x(x_t, y_t)}{J(x_t, y_t)} - \text{tr}\left[\zeta_t \Lambda_t' \frac{J_{xy}(x_t, y_t)}{J_y(x_t, y_t)}\right]. \tag{17}$$

We call V_t the conditional variance, M_t the conditional model uncertainty, and H_t the hedging component. Equation (14) illustrates the affect of model uncertainty on expected returns. Equation (14) is the analog of the well-studied equation from Merton (1973) which links expected returns to volatility and hedging concerns. When there is no model uncertainty, so that $\theta = 0$, we recover Merton's original formulation. Concerns about risk and model uncertainty affect the hedging component as well as V_t and M_t. The value of θ can affect H_t directly and indirectly through the value function and its derivatives.

The stylized approach of most studies of the risk-return trade-off consists of ignoring model uncertainty and hedging, and running the following regression:

$$E_t r_{et+1} = b + \gamma V_t \tag{18}$$

where b is a constant which according to the model should be zero. The goal of many researchers has been to find a significantly positive γ coefficient that captures the trade-off between risk and return, a relationship henceforth referred to as Merton's ICAPM or simply the ICAPM.[6] There is an abundant literature on estimating volatility V_t and many different measures have been used in the ICAPM. The main challenge for us is how to measure and estimate the M_t component in the presence of model uncertainty. This is the focus of the remainder of the paper, after a brief review of flexible functional forms. Flexible functional forms are a crucial ingredient of both the measure of volatility we use and our measure of model uncertainty.

II. Flexible weights

In recent work, Ghysels, Santa-Clara, and Valkanov (2005) suggested that the risk-return trade-off can be modeled with a mixed data sampling, or MIDAS, regression approach. Namely, they ran monthly risk-return regressions, yet used also daily data to model the conditional variance. The mixed use of daily and monthly data was achieved via a parsimonious parameterization of the conditional variance. The key to modeling conditional variances is indeed parsimony, summarizing in a convenient way the temporal dynamics which yield predictions of future volatility. To model uncertainty the issues are different, yet the key is also parsimony. In our empirical work we will use data reflecting predictions of professional forecasters. Predictions are plentiful and the challenge is to summarize them to reflect the fundamental uncertainty in the economy. Our approach makes extensive use of flexible functional forms similar to those suggested in a time series volatility prediction context by Ghysels, Santa-Clara, and Valkanov (2005). In particular, we use the Beta distribution and the normal distribution in the cross-section. The weighting scheme represents the probability that agents assign to each plausible model, described by the predictions of forecasters (discussed in detail in Subsection B). The remainder of this section generically describes the weights we use, leaving discussions of their applications to later sections.

It is noted by Ghysels, Santa-Clara, and Valkanov (2004) that a discretized Beta distribution is a flexible functional form that is convenient to capture many plausible patterns of time series decay. The discretization is based on the standard continuous Beta probability density function which is

$$w[x] = \frac{(x-a)^{\alpha-1}(d-x)^{\beta-1}}{B(\alpha,\beta)(d-a)^{\alpha+\beta-1}} \tag{19}$$

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

$$\begin{aligned} w_i &= \frac{(i-a)^{\alpha-1}(d-i)^{\beta-1}}{B(\alpha,\beta)(d-a)^{\alpha+\beta-1}} \left[\sum_{j=1}^{n} \frac{(j-a)^{\alpha-1}(d-j)^{\beta-1}}{B(\alpha,\beta)(d-a)^{\alpha+\beta-1}} \right]^{-1} && (20) \\ &= \frac{(i-a)^{\alpha-1}(d-i)^{\beta-1}}{\sum_{j=1}^{n}(j-a)^{\alpha-1}(d-j)^{\beta-1}} && (21) \end{aligned}$$

with n values ($i = 1, 2, \ldots n$) that receive positive probability. We require $a \leq 1$ and $d \geq n$. In a time series application n could be the number of lags used in a volatility prediction. Note that a potentially large set of weights is tightly parameterized via a small set of parameters. Ghysels, Santa-Clara, and Valkanov (2004) and Ghysels, Sinko, and Valkanov (2003) discuss how, by varying parameters, the discretized Beta distribution can capture many different weighting schemes associated with time series memory decay patterns observed in volatility dynamics and other persistent time series processes.

They also observe that setting $\alpha = 1$ yields downward sloping weighting schemes typically found in models of volatility predictions. In the cross-sectional applications in this paper, the Beta tightly parameterizes the distribution of forecasts and helps determine which part of the distribution of model predictions matter for asset pricing. The Beta weighting scheme adapts easily to cross-sectional application because setting $\alpha = \beta$ (and hence requiring an even smaller set of parameters) yields various bell-shaped weighting schemes. By construction this is a well-formed pdf since $\sum_{i=1}^{n} w_i = 1$ and we will interpret the w_i's as weights. This convenient scheme is used in our empirical work, both in a time series context to parameterize volatility and in a cross-sectional setting to specify uncertainty.

While Ghysels, Santa-Clara, and Valkanov (2004) use the Beta weights there are other convenient flexible functional forms that are also explored in our empirical work. One such functional form is a discretized normal distribution in which

$$w_i = \frac{\exp\left(\frac{-(i-\mu)^2}{\xi^2}\right)}{\sum_{j=1}^{n} \exp\left(\frac{-(j-\mu)^2}{\xi^2}\right)} \tag{22}$$

for $i = 1, 2, \ldots n$ where μ and ξ are parameters. It is sometimes more convenient to write this discretization as

$$w_i = \frac{\exp\left(-a_1 i - a_2 i^2\right)}{\sum_{j=1}^{n} \exp\left(-a_1 j - a_2 j^2\right)} \tag{23}$$

where a_1 and a_2 are parameters. These weights, called Almon lags, were used in a time series context by Ghysels, Santa-Clara, and Valkanov (2005).

III. Data on beliefs

Our approach requires an empirical measure of the beliefs of agents about the amount of uncertainty they have about the reference model. The beliefs of agents about this uncertainty will be determined from the survey responses of professional forecasters. Economists can infer a lot of indirect information about beliefs by observing choices but direct data on beliefs is primarily limited to surveys. The accuracy of survey responses has long been a major concern and is one reason why survey data has received mixed attention in economics.[7] The typical survey respondent may not put much effort into accurately completing surveys. Respondents typically receive the same reward (often no reward) no matter how carefully and accurately they complete their surveys. Respondents may not even bother to read or understand directions. The drawbacks of surveys are not as strong when the respondents are professional forecasters. Professional forecasters likely will make an effort to understand directions and take great care in accurately reporting their beliefs (via predictions) about the economy. In some surveys forecasters are screened before they can participate and are rewarded for the ex-post accuracy of their responses.

There are several data sets available which give the predictions of professional forecasters on macroeconomic and financial variables. In this paper we focus on data from the Survey of Professional Forecasters (henceforth SPF). There are many other papers that make use of data from the Survey of Professional Forecasters, most of which evaluate the quality of the predictions [see, for example, Zarnowitz (1985) and Braun and Zarnowitz (1993)]. We use the forecasts to represent the beliefs of agents in economic models and consider asset pricing applications. We use the dispersion of forecasts as a proxy for the amount of model uncertainty that agents have about the reference model described in Section I. The SPF is an attractive survey because it provides a long time series of data (the data begins in 1968) and it provides predictions at many different horizons. Each quarter participants are asked for predictions of the levels for the previous quarter, this quarter, next quarter, two quarters ahead, three quarters ahead and four quarters ahead.[8] The forecasters selected for the SPF come largely from Wall Street, commercial banks, and other institutions. The series we use from the SPF are forecasts of output (before 1992Q1 these consist of forecasts of GNP and after of GDP), the output deflator (before 1992Q1 these consist of forecasts of the GNP deflator and after of the GDP deflator), and Corporate Profits After Taxes. Appendix A describes the details of the data.

Some of our analysis requires predictions on variables which do not directly appear in the SPF. For example, we need forecasts of the real rate of quarterly corporate profit growth but the SPF only provides forecasts of the level of nominal corporate profits and forecasts of the level of prices. Let $E_{it}Y_q$ be forecaster i's time t prediction of the value of Y that will be realized at period q. We *approximate* the gross quarterly forecasted rate of real growth, according to forecaster i, in the nominal variable X between quarters m and n as

$$\left(\frac{E_{it}X_n \; E_{it}P_m}{E_{it}X_m \; E_{it}P_n}\right)^{\frac{1}{n-m}} \tag{24}$$

where P is the GDP (GNP) deflator. See Appendix A for more details.

Our analysis also requires us to infer forecasts of the real market return and forecasts of the real return on a nominally risk-free bond. We determine forecasts of the expected market return from predictions of nominal corporate profits and the price level by using the Gordon growth model. We approximate forecaster i's prediction of the return on the market as

$$E_{it} r_{mt+1} = E_{it} \left[\frac{\pi_{t+1}}{q_t} \right] + \xi_{it}. \tag{25}$$

where π_{t+1}, is aggregate real corporate profits, q_t is the market value of firms at time t, and ξ_{it} is forecaster i's predicted growth rate of real corporate profits over a long horizon. See Appendix A for more details. We approximate forecaster i's prediction of the real return on a nominally risk-free bond with

$$E_{it} r_{bt+1} = \frac{R_{b+1} P_t}{E_{it} P_{t+1}} \tag{26}$$

where R_{bt+1} is the nominal return on the bond (which is known at time t) and P_t is the GDP (GNP) deflator. Again see Appendix A for more details.

Table I shows that for the real market, the Gordon growth model gives a reasonable approximation of the unconditional mean return. For the period between 1968 and 2003 the average median forecast of the market return computed from the Gordon growth model (with ξ_t being the forecasted average return from the last period to *three* quarters ahead – a horizon of four) slightly overestimates the actual average market return. Table I also shows that the average median forecasts computed using the formula in equation (26) are very close to the actual average real return on a nominally risk-free bond.

IV. Volatility and model uncertainty-return trade-offs

In this section, we report estimates of the ICAPM with risk and model uncertainty. The estimates allow us to appraise the relative importance of risk and model uncertainty. The estimates in this section also allow us construct an index reflecting the amount of model uncertainty in the economy.

The risk-return trade-off has been the subject of much empirical research. Most papers have used an ARCH-type model, see e.g. French, Schwert, and Stambaugh (1987), Campbell and Hentschel (1992), Nelson (1991), Glosten, Jagannathan, and Runkle (1993), Harvey (2001). Recently, Ghysels, Santa-Clara, and Valkanov (2005) suggested estimating conditional volatility with MIDAS estimation methods. In this paper we will also measure risk using the approach of Ghysels, Santa-Clara, and Valkanov (2005).

To measure uncertainty we face a different challenge. We have data reflecting various predictions about future variables that affect expected returns such as predictions of the market return, corporate profits, etc. Our approach summarizes the cross-sectional variation among forecasters with a parametric specification that allows us to compute a measure of model uncertainty in the economy.

We consider again a version of the decomposition suggested in Section I

$$\begin{aligned} E_t(r_{et+1}) &= b &+& \quad \underbrace{\gamma\, V_t} &+& \quad \underbrace{\theta\, M_t} \\ &= \text{Constant} &+& \quad \text{Risk} &+& \quad \text{Model Uncertainty} \end{aligned} \tag{27}$$

where we have set the hedging component H_t to be zero.[9] Sufficient conditions for the hedging component to be zero are that for all t the noise driving the market is uncorrelated with the noise driving the state (so that $\zeta_t \Lambda_t' = 0$ for all t) and the uncertainty in the market is unrelated to the uncertainty in the state (so that $\eta_t \Delta_t' = 0$ for all t).

It is probably unreasonable to assume the noise and uncertainty underlying the market and the state are not related. An alternative justification of equation (27) is that it is an approximation to equation (14). If γ is close to one, $\zeta_t \Lambda_t'$ is close to zero for all t, and $\eta_t \Delta_t'$ is close to a matrix of zeros for all t then the approximation is very good. When we estimate this model we use quasi-maximum likelihood because the *econometrician's model*, which will be a formal version of equation (27) discussed below, may be misspecified. One reason it may be misspecified is that the hedging term is set to zero. Another reason the model may be misspecified is that we take the reference model as the data generating processes. The agents inside our model are worried the reference model is misspecified. If it is misspecified then our regressions are misspecified too.

The most critical step of our approach, and one of the main innovations of this paper, is quantifying M_t. We take advantage of the flexible functional forms described in Section II to specify M_t. We link M_t to a weighted variance of forecasts given by professional forecasters.

In the remainder of this section we provide the details of this approach as well as the empirical results. In Subsection A following Ghysels, Santa-Clara, and Valkanov (2005), we review a measure of volatility and then devote most of our attention to investigating if there is a model uncertainty-return trade-off in subsection B. An index reflecting the amount of model uncertainty in the economy is discussed in Subsection C.

A. Volatility

The MIDAS estimator forecasts the conditional variance with a weighted average of lagged daily squared returns (or alternatively lagged *demeaned* squared returns) using a flexible functional form to parameterize the weight given to each lagged daily squared return (or lagged squared demeaned excess return). Ghysels, Santa-Clara, and Valkanov (2005) estimate the coefficients of the conditional variance process jointly with b and γ from the expected return equation with quasi-maximum likelihood and show that a parsimonious weighting scheme with only two parameters works quite well.

More specifically, the MIDAS estimator of the conditional variance of excess returns, V_t, is based on prior *daily* squared demeaned return data. The weight on the ith lag is

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

where s is the maximum number of lags which we fix at the number of trading days in a year. The functional form of these weights is determined by the discretized Beta distribution described in Section II with $\alpha = 1$, $\beta = \omega$, and $d = s$. The value of a does not matter since $\alpha = 1$. The single free parameter ω models the decay pattern of the weight function. The top plot in Figure 1 for an example of the weights.[10] The resulting conditional variance is equal to

$$V_t = \sigma^2 \text{vol}_t(\omega) \tag{28}$$

where σ^2 is a time-invariant constant and

$$\text{vol}_t(\omega) = s \sum_{i=1}^{s} l_i(\omega) \left(r_{et}^{s-i,s-i+1} - \frac{1}{s} \sum_{j=1}^{s} r_{et}^{s-j,s-j+1} \right)^2 + \\ 2s \sum_{i=1}^{s-1} \sqrt{l_i(\omega) l_{i+1}(\omega)} \left(r_{et}^{s-i,s-i+1} - \frac{1}{s} \sum_{j=1}^{s} r_{et}^{s-j,s-j+1} \right) \left(r_{et}^{s-i-1,s-i} - \frac{1}{s} \sum_{j=1}^{s} r_{et}^{s-j,s-j+1} \right) \tag{29}$$

is the component of the conditional variance which is determined from the volatility of daily excess returns. Here $r_{et}^{s-i,s-i+1}$ is the daily return between trading days $s-i$ and $s-i+1$ which occur between periods $t-1$ and t. Notice that vol_t implicitly depends on the parameter ω since the weights l_i depend on ω. The second component allows for the effect on quarterly volatility of serial correlation in daily returns. Such a correction did not appear in the original formulation of Ghysels, Santa-Clara, and Valkanov (2005).

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.

To estimate the parameters, we maximize the (quasi-)likelihood of quarterly excess returns based on:

$$r_{et+1} \sim N\left(b + \gamma\, V_t, V_t\right) \tag{30}$$

which we implement as

$$r_{et+1} \sim N\left[b + \tau \,\text{vol}_t\left(\omega\right), \sigma^2 \text{vol}_t\left(\omega\right)\right] \tag{31}$$

where $\tau = \gamma\sigma^2$. The combination of quarterly returns and daily squared returns yields the MIDAS setup. Using this setup, Ghysels, Santa-Clara, and Valkanov (2005) find there is a significant positive relation between risk and return in the stock market. This finding is robust to asymmetric specifications of the variance process, and to controlling for variables associated with the business cycle. It also holds in many subsamples. Given these empirical findings it is a good benchmark reduced form regression to introduce model uncertainty. Ghysels, Santa-Clara, and Valkanov (2005) focused on monthly returns whereas we devote our attention to quarterly sampling frequencies because the professional forecast data used to construct our measure of model uncertainty, unc_t, is only available quarterly. We expect, however, that the focus on the quarterly sampling frequency weakens empirical evidence of the risk-return trade-off, and the results reported below confirm this.

Quasi-likelihood estimates of the parameters appearing in equation (31) are displayed in the first three estimations in Panel B of Table II. As a benchmark, the first estimation in Panel A provides quasi-likelihood estimates of the specification

$$r_{et+1} \sim N\left[b, \sigma^2\right] \tag{32}$$

in which there is no risk-return trade-off and the excess return is i.i.d. normal. We see that according to a t-test and a likelihood ratio test, estimates of τ are not significant. However estimates of $\log \omega$ are extremely significant, both according to a t-test and a likelihood ratio test. Hence, these results suggest that, in our data set, although there is no evidence of a risk-return trade-off, MIDAS does provide a very convenient approach for capturing volatility.

Further evidence that MIDAS captures volatility can be provided by examining the relationship between vol_t and realized volatility. We define realized volatility as

$$Q_t = q_t \sum_{i=1}^{q_t} \left(r_{et}^{q_t-i,q_t-i+1} - \frac{1}{q_t} \sum_{j=1}^{q_t} r_{et}^{q_t-j,q_t-j+1} \right)^2 + \\ 2q_t \sum_{i=1}^{q_t-1} \left(r_{et}^{q_t-i,q_t-i+1} - \frac{1}{q_t} \sum_{j=1}^{q_t} r_{et}^{q_t-j,q_t-j+1} \right) \left(r_{et}^{q_t-i-1,q_t-i} - \frac{1}{q_t} \sum_{j=1}^{q_t} r_{et}^{q_t-j,q_t-j+1} \right) \quad (33)$$

where q_t is the number of days in a quarter t.[11] Table III shows that vol_t is more highly correlated with future realized volatility (Q_{t+1}) than past realized volatility (Q_t). This suggests that vol_t indeed does provide a better measure of conditional volatility in equation (30) than an alternative version which measure conditional volatility with realized volatility (Q_t).

Our implementation differs from the implementation in Ghysels, Santa-Clara, and Valkanov (2005) in that in this paper we estimate σ^2, rather than fixing 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 $b + \gamma \text{vol}_t$ should be the conditional mean of the market excess return. However if the econometrician's 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 the constant term, b, is close to zero. We find the constant term is not significantly different from zero in Table II, which confirms the predictions of the theoretical model.

Our results differ from Ghysels, Santa-Clara, and Valkanov (2005). The bulk of their analysis focuses on monthly return horizons. However, to construct measures of model uncertainty we are confined to quarterly data. Ghysels, Santa-Clara, and Valkanov (2005) do provide quarterly regressions between 1964 and 2000 of a specification similar to equation (31), and show there is evidence of a risk-return trade-off. There are several reasons the results in this paper are different. The most important reasons are that Ghysels, Santa-Clara, and Valkanov (2005) considered a different time period than the time period considered in this paper and that our definition of a quarter refers to a calendar quarter (matching forecasts) whereas the definition of quarter in Ghysels, Santa-Clara, and Valkanov (2005) corresponds to a fixed number of trading days which are not directly related to calendar quarters.[12]

B. Model uncertainty

We now extend our analysis to include an empirical measure of model uncertainty so that there is a possibility of a model uncertainty-return trade-off in addition to a risk-return trade-off. We proxy for the contribution of model uncertainty to the excess return with, $\theta\ \mathrm{unc}_t$, where θ is a time-invariant constant and unc_t measures the disagreement among forecasters about the growth rate of a single variable. We focus in particular on measuring unc_t with the disagreement about the market return where market return forecasts are constructed from the Gordon growth model. We also consider measuring unc_t with the disagreement in constructed GDP and corporate profit growth forecasts.

Why is it reasonable to measure model uncertainty with the disagreement of professional forecasters? In other words, using the notation of Section I, why is it reasonable that the disagreement of professional forecasters at time t is related to $\eta_t\eta_t'$? Recall that η_t effectively determines the penalty for considering perturbations to the market return from the reference model. For example, the penalty for considering a perturbation of η_t is 1.0. We suggest that when the disagreement among professional forecasters is large, η_t should be large. In our environment all agents are alike. Imagine each agent listening to the statements of professional forecasters. If all forecasters are in agreement then agents have very little doubt that the reference model is correct and η_t is small. If forecasters state very different forecasts then agents are unsure which forecast is correct and want to consider a variety models so η_t is large.

We consider a quasi-likelihood estimator similar to (31) of the following form:

$$r_{et+1} \sim N\left[b + \tau \mathrm{vol}_t(\omega) + \theta\ \mathrm{unc}_t(\nu), \sigma^2 \mathrm{vol}_t(\omega)\right]. \tag{34}$$

where we interpret $M_t = \text{unc}_t(\omega)$. In equation (34) we explicitly indicate the dependence of unc_t on a parameter ν. We next discuss how we construct unc_t and the role of the parameter ν.

As indicated in Section II, we will use the Beta distribution to tightly parameterize how forecasts determine unc_t. Instead of letting the first power parameter equal to 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. The model 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. To construct the weights we proceed as follows: we pick one series, call it x, and rank the forecasts each period of x from low to high. The weight on the ith lowest forecast is

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

where f_t forecasts are available at time t and ν is a parameter. This is the discretized Beta distribution described in Section II with $\alpha = \nu$, $\beta = \nu$, $a = 0$ and $d = f_t$. Requiring $\alpha = \beta$ forces the weights to be symmetric.[13] The disagreement or model uncertainty is then measured by a Beta-weighted variance of forecasts of x :

$$\text{unc}_t(\nu) = \sum_{i=1}^{f_t} w_{it}(\nu)\left[x_{it+1|t} - \sum_{j=1}^{f_t} w_{jt}(\nu) x_{jt+1|t}\right]^2. \tag{35}$$

Quasi-likelihood estimates of the parameters appearing in equation (34) are displayed in Panel B of Table II. In the fourth regression in Panel B we include model uncertainty but measure model uncertainty with an unweighted (or flat weighted) variance which is obtained by setting $\log\nu = 0$. We see in this case the estimate of θ is not significant and there is very little improvement to the log likelihood without uncertainty. Including uncertainty with flat weights does not improve much upon specifications in which uncertainty is left out. In the 5th, 6th and 7th estimations in Panel B we estimate θ and $\log\nu$ along with other parameters. In these regressions unc_t is a non-degenerate Beta-weighted variance. Estimates of θ and $\log\nu$ are significant (by likelihood ratio tests and t-tests) and there is a large improvement to the log-likelihood. Including Beta-weighted uncertainty significantly improves the fit. 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 our theoretical model.

Further informal evidence for a model uncertainty-return trade-off is provided in Table III – in particular the correlation between our estimated measure of model uncertainty in the last regression in Panel B and the excess return is 0.28. In comparison the correlation of our measure of volatility with the excess return is only 0.15. The visual evidence in the joint plots in Figure 2 yield additional insights into the nature of the relationship between uncertainty and excess returns. We see that when uncertainty is high, excess returns also tend to be high. When uncertainty is low however there is not a strong relationship between uncertainty and excess returns.

Figure 3 graphs the likelihood function in the last regression of Panel B as a function of $\log \nu$ and $\log \omega$. We see that the optimal value of $\log \nu$ is sharply determined. Large values of $\log \omega$ give low likelihoods but many small values give essentially the same likelihood.

In the second regression in Panel A of Table II we see that estimates of θ and $\log \nu$ are jointly significant according to a likelihood ratio test when the errors are assumed to be homoskedastic. According to a t-test, $\log \nu$ but not θ is significant at conventional levels. This suggests that our model uncertainty-return trade-offs are robust to the specification of volatility.

In Table IV 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 the alternative variables we analyze should affect model uncertainty and thus excess returns. In Panel A we consider the uncertainty in constructed real GDP growth forecasts and in Panel B we consider the uncertainty in growth rate of corporate profits at many different horizons. We see that the uncertainty in real GDP forecasts does not have a significant effect on excess returns. At long horizons (three and four) the uncertainty in corporate profits forecasts does have a significant effect but at shorter horizons (one and two) it has essentially no effect. Our constructed market return forecasts are based on a combination of short and long term corporate profits forecasts. The results in Panel B suggest that the underlying driving force for our results comes from long term corporate profits forecasts.

We now briefly consider some alternative specifications. In Panel A of Table V we measure uncertainty with a symmetric normal weighted variance in which the weights are

$$w_{it}(\xi) = \frac{\exp\left(\frac{-\left(i-\frac{f_t+1}{2}\right)^2}{\xi^2}\right)}{\sum_{j=1}^{f_t}\exp\left(\frac{-\left(j-\frac{f_t+1}{2}\right)^2}{\xi^2}\right)} \tag{36}$$

where ξ is a parameter. This is the first discretized normal specification given in Section II with

$$\mu = \frac{f_t + 1}{2} \tag{37}$$

available at time t. This specification of μ forces the weights to be symmetric. The results in Panel A are very similar to the results in specifications six and seven of Table II. 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 a model uncertainty-return trade-off even with a different specification of the cross-section weights.

In Panel B of Table V we measure uncertainty with weights that are not restricted to being symmetric. We use Beta weights in which the weights are

$$w_{it}(\alpha, \beta) = \frac{i^{\alpha-1}(f_t - i)^{\beta-1}}{\sum_{j=1}^{n} j^{\alpha-1}(f_t - j)^{\beta-1}} \tag{38}$$

where α and β are free parameters. This is the discretized Beta distribution described in Section II with $a = 0$ and $d = f_t$. Allowing the weights to be non-symmetric lets the perceived agents' perceived model uncertainty depend on any part of the distribution of forecasts. If agents are worried about worse case forecasts then β should be greater than α likewise if agents focus on rosy forecasts then β should be less than α. The estimates α and β are not significantly different from each other. Since the estimate value of β is slightly greater than α, the estimated weights slightly emphasize pessimistic forecasts over optimistic forecasts. There is not compelling to evidence to suggest that agents are pessimistic and that non-symmetric weights more precisely measure perceived variances.

In Panel C of Table V we measure uncertainty by the Beta-weighted variance of constructed market return forecasts when the long term horizon is three periods rather than four periods. Setting the long term horizon at three does better than setting the long term horizon at four.[14] The Gordon growth model requires a long term horizon forecast and it is most natural to the let the long term horizon of four because that is the longest horizon for which data is 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 is important to note that all of results would become stronger if we used a horizon of three. In particular, there is much stronger evidence for a model-uncertainty trade-off when the long term horizon is three rather than four.

C. An index of model uncertainty

This section discusses a few of the empirical properties of estimated uncertainty. We let the index of model uncertainty be the uncertainty series, $\text{unc}_t(15.346)$, estimated in the last regression in Panel B of Table II. We plot the index in Figure 4 along with plots of excess returns and volatility. We provide some simple statistics in Table III and estimates of autoregressions in Table VI.

It is well known that volatility is highly persistent. In our data, Table VI 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 VI shows there is not a significant relationship between uncertainty and lagged volatility. We see from Table III, there is very little contemporaneous correlation between uncertainty and volatility. This suggests that the actual conditional variance (past volatility) has some, but not much, impact on the beliefs of agents about model uncertainty.

Model uncertainty is not highly correlated with future volatility. Past model uncertainty does not predict future volatility and vice versa. Hence, volatility and model uncertainty appear as orthogonal processes. From Table III we see that unweighted model uncertainty is slightly more related to future volatility than the optimally weighted model uncertainty. Maybe the fringes of forecasts matter for volatility (i.e. maybe they are noise traders) but not for expected returns. Though, this effect is not strong.

V. Conclusions

An important strand of recent research in finance contends that model uncertainty, in addition to risk, should matter for asset pricing. Model uncertainty is, however, difficult to measure. We suggest a reasonable proxy for the amount of model uncertainty in the economy is the disagreement of professional forecasters. We find stronger empirical evidence for a model uncertainty-return trade-off than for the traditional risk-return trade-off.

We construct our measure of uncertainty using a flexible weighting scheme across forecasts which can accommodate assigning more or less weight to extreme forecasts. The weighting scheme is inspired by the recent work on MIDAS regressions. Our estimates of the weighting scheme entail ignoring the extremes and placing all of the weight on the center of the distribution. Without the cross-sectional weighting scheme introduced in the paper, we find that flat weighted model uncertainty is not related to excess returns. It is not highly correlated with the market return and it does not have a significant effect in regressions. In contrast, the model uncertainty measure obtained from our cross-sectional weighting is found to be empirically significantly related to returns.

We find slight but not significant evidence for pessimism: Symmetric cross-sectional weights can not be rejected. Our model does better when the constructed return forecasts are based on a long term horizon of three rather than four. However we emphasize a long term horizon of four because that is the longest horizon for which data is plentiful and the Gordon growth model requires a long term forecast. It is perhaps puzzling that uncertainty in macro variables such as GDP do not affect excess returns. But giving the low correlation of the realization of most macro variables with the realization of financial variables, perhaps this should not be surprising.

Our results thus provide empirical support to recent research in finance which contends that model uncertainty, in addition to risk, should matter for asset pricing.

Appendix A. Data on beliefs

This Appendix describes the details of the data and is organized into several subsections. Subsection A describes the Survey of Professional Forecasters. Subsection B discusses how we compute growth rate forecasts from level forecasts and how we compute real forecasts from nominal forecasts. Subsection C explains the computations of asset return forecasts.

Appendix A. The Survey of Professional Forecasters

The Survey of Professional Forecasters (henceforth SPF) began in the fourth quarter of 1968 as a joint project between the American Statistical Association (ASA) and the National Bureau of Economic Research (NBER). In the first quarter of 1990, the ASA/NBER discontinued the project. The Federal Reserve Bank of Philadelphia (FRB-Philadelphia) reinstated the SPF in the third quarter of 1990.[15] The SPF provides a long time series of data. Each quarter participants are asked for predictions of the levels for the previous quarter, this quarter, next quarter, two quarters ahead, three quarters ahead and four quarters ahead.[16]

The number of forecasters participating in the SPF has varied through time. The average (median) number of forecasts between 1968 and 2004 is 39.5 (36). In the early years, the number occasionally increased to greater than 100 forecasters, but began to decline nearly monotonically throughout the 1970s and 1980s. After the 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.[17] Not all of the forecasts are usable because some are incomplete. Across all dates, we were able to use a median of 26 forecasts, a min of 9, and a max of 74.

Since the survey began, several series have been added, a few have been dropped, and some have been replaced with similar series. In the early 1990's, (1992Q1) output forecasts switched from being forecasts of GNP to being forecasts of GDP. At the same time forecasts of the GNP deflator were replaced by forecasts of the GDP deflator. The switch coincided with the substitution of GNP by GDP undertaken by the Bureau of Economic Analysis. Forecasts of real consumption expenditures and the consumer price index were both added to the survey in the third quarter of 1981.

The series we use from the SPF are Nominal GDP (NGDP), GDP deflator (PGDP), and Corporate Profits After Taxes (CPROF). We discard forecasts that were incomplete at a particular date. In order for a forecaster's forecasts to be included at a particular date it is necessary that he provide forecasts for NGDP, PGDP and CPROF for this quarter, next quarter, two quarters ahead, and three quarters ahead. Forecasts were not dropped if forecasts four quarters ahead were not provided.

Appendix B. Computing quarterly real growth rate forecasts

We compute implied growth rate forecasts from forecasts of levels since the SPF only provides forecasts of levels. We do this as follows. Let $E_{it}x_q$ be forecaster i's time t prediction of the value of x that will be realized at period q.

One needs to make a distinction between real and nominal variables in the SPF. Let us start with the case of real variables. Consider forecaster i's gross quarterly forecasted rate of real growth in x, where x is forecasted in real terms between quarters m and n. This will be approximated by

$$\left(\frac{E_{it}x_n}{E_{it}x_m}\right)^{\frac{1}{n-m}}. \tag{A1}$$

If quarter m is in the future then this is only an approximation for the quarterly growth forecast between quarters m and n since in general

$$E_{it}\left[\left(\frac{x_n}{x_m}\right)^{\frac{1}{n-m}}\right] \neq \left(\frac{E_{it}x_n}{E_{it}x_m}\right)^{\frac{1}{n-m}}. \tag{A2}$$

If m is in the past (so that $m \leq t$ and the value of x_m is perfectly known) and $n = m+1$ then these constructed growth rates are exact and the inequality in equation (A2) becomes an equality.

For some of our variables we need to construct implied real growth rates from nominal forecasts. For example, we need forecasts of the real rate of corporate profit growth but in the SPF only forecasts of the nominal level of corporate profits are provided. We compute approximate real forecasts from nominal forecasts and forecasts of the price level. 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

$$\left(\frac{E_{it}X_n \; E_{it}P_m}{E_{it}X_m \; E_{it}P_n}\right)^{\frac{1}{n-m}} \tag{A3}$$

where P_q is the price level at time q. In general this is only an approximation since usually

$$E_{it}\left[\left(\frac{X_nP_m}{X_mP_n}\right)^{\frac{1}{n-m}}\right] \neq \left(\frac{E_{it}X_n \; E_{it}P_m}{E_{it}X_m \; E_{it}P_n}\right)^{\frac{1}{n-m}} \tag{A4}$$

even when $t = m$ and $n = m + 1$. (It will be exact if $m < t$ and $n = m + 1$ so that the values of all variables are known perfectly,) For forecasts of the price level we use forecasts of the GDP (GNP) price deflator since forecasts of the CPI only became available in the fourth quarter of 1991.[18] Under the assumption that money is neutral the approximate real growth rate forecast will be exact when $m \leq t$ and $n = m + 1$.[19] Of course in reality money is not neutral but data limitations force us to use this approximation anyway.

Appendix C. Computing asset return forecasts

In this section we discuss how we compute forecasts of the real market return and forecasts of the real return on a nominally risk-free bond. We determine forecasts the expected market return from predictions of nominal corporate profits and the price level by using the Gordon growth model. We determine forecasts of a nominally risk-free bond from forecasts of the price level.

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. 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 (2004), 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.

In this paper we use corporate profits forecasts rather than earnings forecasts as inputs to the Gordon growth model. Let π_t be aggregate corporate profits and q_t the market value of firms at time t.[20] For us, the Gordon growth model amounts to assuming that forecaster i's constructed prediction of the return on the market is

$$E_{it} r_{mt+1} = E_{it} \left[\frac{\pi_{t+1}}{q_t} \right] + \xi_{it}. \tag{A5}$$

where ξ_{it} is forecaster i's predicted growth rate of corporate profits over a long horizon.

We face a difficult timing issue when implementing equation (A5). 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 ahead (2001Q3), three quarters ahead (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 $t = 2001Q1$ and 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.[21]

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

We also need to compute the expected real return on a nominally risk-free bond. We approximate forecaster i's prediction of the real return on a nominally risk-free bond with

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

where R_{bt+1} is nominal return on the bond (which is known at time t) and P_t is the time t value of the GDP (GNP) deflator. In general this is an approximation because usually

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

References

Anderson, Evan W., Eric Ghysels, and Jennifer L. Juergens, 2005, Do heterogeneous beliefs matter for asset pricing?, *Review of Financial Studies* 18, 875–924.

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

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

Brandt, Michael, and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Economies* 72, 217–259.

Braun, Phillip, and Victor Zarnowitz, 1993, Twenty-two years of the NBER-ASA quarterly economic outlook surveys: Aspects and comparisons of forecasting performance, in J.K. Stock, and M.W. Watson, ed.: *Business Cycles, Indicators, and Forecasting* (University of Chicago Press: Chicago).

Brav, Alon, Reuven Lehavy, and Roni Michaely, 2004, Using expectations to test asset pricing models, Working Paper, Duke University.

Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economies* 18, 373–399.

———, and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economies* 31, 281–318.

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

Chen, Zengjing, and Larry G. Epstein, 2002, Ambiguity, risk and asset returns in continuous time, *Econometrica* 70, 1403–1443.

Crombez, John, 2001, Updating expected returns based on consensus forecasts, Working Paper, Ghent University.

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

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina, 2002, Differences of opinion and the cross-section of stock returns, *Journal of Finance* 57, 2113–2141.

French, Kenneth R., William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economies* 19, 3–29.

Gebhardt, William R., Charles M.C. Lee, and Bhaskaran Swaminathan, 2001, A test of efficiency of a given portfolio, *Journal of Accounting Research* 39, 135–176.

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2004, Predicting volatility: Getting the most out of return data sampled at different frequencies, Working Paper, Forthcoming in *Journal of Econometrics*.

———, 2005, There is a risk-return tradeoff after all, *Journal of Financial Economics* 76, 509–548.

Ghysels, Eric, Arthur Sinko, and Rossen Valkanov, 2003, MIDAS regressions: Further results and new directions, Working paper, UNC and UCLA.

Glosten, Lawrence R., Ravi Jagannathan, and David E. Runkle, 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, Myron J., 1962, *The Investment, Financing, and Valuation of the Corporation* (Homewood, Illinois, Irwin Publishers).

Guay, Wayne R., S. P. Kothari, and Susan Shu, 2003, Properties of implied cost of capital using analysts' forecasts, Working Paper, MIT Sloan School of Management.

Hansen, Lars P., and Thomas J. Sargent, 2001, Acknowledging misspecification in macroeconomic theory, *Review of Economic Dynamics* 4, 519–535.

———, and Thomas D. Tallarini Jr., 1999, Robust permanent income and pricing, *Review of Economic Studies* 66, 873–907.

Hansen, Lars P., Thomas J. Sargent, Gauhar A. Turmuhambetova, and Noah Williams, 2004, Robustness and model misspecification, Working Paper, University of Chicago, New York University and Princeton University.

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

Jagannathan, Ravi, Ellen McGrattan, and Anna Scherbina, 1996, The declining us equity premium, *Federal Reserve Bank of Minneapolis Quarterly Review* 24, 3–19.

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

Keane, Michael P., and David E. Runkle, 1990, Testing the rationality of price forecasts: New evidence from panel data, *American Economic Review* 80, 714–735.

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

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

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

Miller, Edward M., 1977, Risk, uncertainty, and divergence of opinion, *Journal of Finance* 32, 1151–1168.

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

Qu, Shisheng, Laura T. Starks, and Hong Yan, 2003, Risk dispersion of analyst forecasts and stock returns, Working Paper, University of Texas at Austin.

Rudin, Jeremy R., 1992, What do private agents believe about the time-series properties of gnp?, *Canadian Journal of Economics* 25, 369–391.

Turner, Christopher M., Richard Startz, and Charles R. Nelson, 1989, A markov model of heteroskedasticity, risk, and learning in the stock market, *Journal of Financial Economies* 25, 3–22.

Uppal, Raman, and Tan Wang, 2003, Model misspecification and underdiversification, *Journal of Finance* 58, 2465–2486.

Wiese, Robert, 1930, Investing for true value, *Barron's* 5.

Williams, John Burr, 1938, *The Theory of Investment Value* (Fraser Publishing).

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

———, and Louis A. Lambros, 1987, Consensus and uncertainty in economic prediction, *Journal of Political Economy* 95, 591–620.

Notes

[1]In general, the disagreement among forecasters could arise because forecasters have different models, different information, different initial beliefs (priors), or different forecasting techniques. In this paper we interpret the differences as arising from different models and assume forecasters have the same information and the same initial beliefs.

[2]Many other recent papers have discussed model uncertainty including Hansen, Sargent, Turmuhambetova, and Williams (2004), Liu, Pan, and Wang (2005) and Chen and Epstein (2002).

[3]Introducing the parameter θ effectively over-parameterizes the problem for us. We could explain the same behavior by always requiring $\theta = 1$ and scaling Δ and η by a time-invariant constant. However, allowing θ to not equal one is convenient for interpreting our results.

[4]For example see Hansen, Sargent, and Tallarini (1999); Anderson, Hansen, and Sargent (2003); and Hansen, Sargent, Turmuhambetova, and Williams (2004).

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

[6]Some investigations of the ICAPM do not ignore the hedging term but almost all empirical investigations ignore the model uncertainty component.

[7]See Keane and Runkle (1990), Rudin (1992), and Zarnowitz and Lambros (1987) for further discussion.

[8]The survey also includes annual and longer horizon forecasts.

[9]Note that we also include a constant term b even though according to the model this term should be zero. The constant is added so that the empirical regressions have well-behaved residuals. As discussed later, testing the statistical significance of b can also be used as model validation test.

[10]The weights displayed in Figure 1 are those obtained from empirical estimates, discussed later.

[11]Note that Q_t is similar to vol_t except the weights are uniform and q does not necessary equal s. Since the weights used to compute Q_t are uniform, Q_t does not depend on ω. Since Q_t is realized volatility within a quarter, q corresponds to the number of days in a quarter. The goal of vol_t is to

predict future volatility. To do this it usually useful to use data form many previous quarters so that s typically will be much larger than q_t.

[12]As previously mentioned, there are also several other differences between the implementation in this paper and the implementation in Ghysels, Santa-Clara, and Valkanov (2005). In this paper we estimate σ^2 (rather than setting it equal to one), allow for serial correlation in daily returns in equation (29), and use a different data source for returns. These differences do not have a large effect our results. If we used the same implementation as Ghysels, Santa-Clara, and Valkanov (2005) we would get essentially the same results.

[13]See for instance the bottom plot in Figure 1 for an example of the weights.

[14]Given our results in Panel B of Table IV 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 three.

[15]See the web-page *http://www.phil.frb.org/econ/spf/index.html* and a comprehensive overview (Croushore 1993) for more information about the survey.

[16]Data on forecasts four quarters ahead is 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. Respondents are given preliminary estimates of last quarter's values and most respondents report these estimates as their forecasts.

[17]There are some extreme low numbers in the second and third quarters of 1990 and they correspond to the transfer of the survey from the ASA/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.

[18]When we deflate the *actual* level of variables we do use the CPI.

[19]To see this let $m \leq t$ and $n = m + 1$. Also define the real variable

$$f = \frac{X_n P_m}{X_m P_n}.$$

It follows that $X_m E_{it} P_n f = P_m E_{it} X_n$ since $m \leq t$. Now since money is neutral the random variable P_n is independent of the random variable f so $X_m E_{it} P_n E_{it} f = P_m E_{it} X_n$ and

$$E_{it} f = \frac{P_m \; E_{it} X_n}{X_m \; E_{it} P_n}.$$

The left hand side of this equation is equal to the left hand side of equation (A4) and the right hand side of this equation is equal to the right hand side of equation (A4) when $m \leq t$ and $n = m + 1$.

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

[21]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 profits forecasts, our results are essentially the same.

Table I

Data Summary and Description

In each row, we list a number of different statistics for the actuals and constructed forecasts of a single variable. In the row beginning with the label r_{mt}, the "actuals" columns provide statistics for the actual real market return. In the row beginning with the label r_{bt}, the "actuals" columns provide statistics for the actual real return on the nominally risk-free bond. Ex_{t+1} is the unconditional actual expected value of x measured with the sample mean of x. Sx_t is the unconditional actual standard deviation of x measured with the sample standard deviation of x. E $\text{med}_t\mu_{xit+1|t}$ is the unconditional expected value of the median forecasts. S $\text{med}_t\mu_{xit+1|t}$ is the unconditional standard deviation of the median forecast. med $S_t\mu_{xit+1|t}$ is the unconditional median of the conditional standard deviations of the forecasted means. $\sqrt{E\left[\left(x_{t+1} - \text{med}_t\mu_{xit+1|t}\right)^2\right]}$ is the square root of the unconditional expected squared forecast error. The forecast data starts with forecasts made in the fourth quarter of 1968 and ends with forecasts made in the third quarter of 2003. The actual data runs from the first quarter of 1969 to the fourth quarter of 2003. Daily and monthly nominal actual asset pricing data is from Kenneth French's web site. They are deflated by the CPI from FRED II to obtain real returns. Flow of funds data, used to compute the constructed market return forecasts, is from the Federal Reserve Board.

	Actuals		*Forecasts*			*Forecast errors*
Variable	Ex_{t+1}	Sx_{t+1}	E $\text{med}_t\mu_{xit+1\|t}$	S $\text{med}_t\mu_{xit+1\|t}$	med $S_t\mu_{xit+1\|t}$	$\sqrt{E\left[\left(x_{t+1} - \text{med}_t\mu_{xit+1\|t}\right)^2\right]}$
r_{mt}	1.0168	0.0901	1.0230	0.0179	0.0173	0.0917
r_{bt}	1.0034	0.0064	1.0051	0.0045	0.0025	0.0050

Table II
Risk-Return and Uncertainty-Return trade-offs

This table displays estimates of several versions of the nonlinear regression

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

of quarterly excess returns r_{et+1} on a constant b, a measure of volatility $\text{vol}_t(\omega)$, specified in equation (29), and a measure of uncertainty $\text{unc}_t(\nu)$, appearing in (35). The estimates are based on the quasi-likelihood estimator appearing in (31). In Panel A the variance of the error term e_{t+1} is σ^2 where σ^2 is a constant which we estimate. In Panel B, the variance of the error term, e_{t+1}, is $\sigma^2 \text{vol}_t(\omega)$ 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 $t+1$). The standard errors of the parameters are quasi-maximum likelihood standard errors and are listed under the variables in parenthesis. If there is no standard error present then the variable was fixed and not estimated. In this case the value of the variable in the estimate column is the value at which it is fixed. The data is quarterly from 1968:04 to 2003:03.

Panel A: Homoskedastic errors

b	θ	$\log \nu$	σ^2	Log Likelihood
0.0134	0	0	0.00790	140.20
(0.0075)			(0.00103)	
-0.0140	1480.	2.65	0.00727	146.06
(0.0115)	(823.79)	(0.748)	(0.000991)	

Panel B: Heteroskedastic error

Specification	b	τ	θ	$\log \omega$	$\log \nu$	σ^2	Log Likelihood
1	0.012	0.000	0.000	0.000	0.000	1.277	147.297
	(0.007)					(0.160)	
2	0.011	0.000	0.000	2.780	0.000	1.582	151.111
	(0.006)			(0.446)		(0.237)	
3	0.009	0.812	0.000	2.768	0.000	1.577	151.184
	(0.009)	(1.759)		(0.448)		(0.240)	
4	0.007	0.742	4.626	2.764	0.000	1.576	151.193
	(0.011)	(1.840)	(34.170)	(0.450)		(0.240)	
5	-0.012	0.000	1540.556	0.000	2.708	1.179	152.867
	(0.010)		(658.146)		(0.564)	(0.148)	
6	-0.012	0.000	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 III
Properties of Uncertainty and Volatility

This table displays quarterly statistics of realized volatility Q, the estimated $\text{vol}_t(14.939)$ series and the estimated $\text{unc}_t(15.346)$ series. Panel A reports means and standard deviations and Panel B reports correlations.

Panel A: Means and standard deviations of vol and unc

	Mean	Standard Deviation
Q	0.006592	0.007634
vol(14.939)	0.005876	0.005428
unc(1)	0.000345	0.000233
unc(15.346)	0.000017	0.000016

Panel B: Correlations of excess returns with vol and unc

	r_{e+1}	Q_{t+1}	Q_t	$\text{vol}_t(14.939)$	$\text{unc}_t(1)$	$\text{unc}_t(15.346)$
r_{e+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
$\text{vol}_t(14.939)$				1.000	0.211	0.075
$\text{unc}_t(1)$					1.000	0.662
$\text{unc}_t(15.346)$						1.000

Table IV

Uncertainty Regressions on GDP and Corporate Profits

This table displays estimates of the same regression as in Panel B of Table II except the variables used to measure uncertainty are different. In Panel A uncertainty is measured by the Beta-weighted variance of constructed forecasts of real GDP growth between last quarter and different horizons and in Panel B uncertainty is measured by the Beta-weighted variance of corporate profits 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 (respectively next quarter, two quarters ahead, or three quarters ahead) is considered.

Panel A: Uncertainty in constructed real GDP 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: 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)	(44540660.719)	(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 V
Alternative Specifications of the Uncertainty Regressions

This table displays estimates of the same regression as in Panel B of Table II except the specification of unc_t is different. In panel A uncertainty is measured with the a symmetric normal weighted variance of the same constructed market return forecast as in Table II. In Panel B, non-symmetric cross-sectional weights are allowed and uncertainty is measure with a Beta-weighted variance of the same constructed market return forecast as in Table II with two free parameters α and β. In Panel C uncertainty is measured by Beta-weighted variance of constructed market return forecasts when the long term horizon is three periods rather than four periods. The specifications numbers for each row correspond to the specification numbers in Table II.

Panel A: Normal Weighted Variance

Specification	b	τ	θ	$\log\omega$	$\log\xi$	σ^2	Log Likelihood
6	-0.011	0.000	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	$\log\omega$	$\log\alpha$	$\log\beta$	σ^2	Log Likelihood
6	-0.006	0.000	2110.001	3.089	3.689	3.886	1.543	157.360
	(0.008)		(1890.912)	(0.463)	(1.289)	(1.245)	(0.244)	
7	-0.006	-0.021	2111.886	3.090	3.690	3.887	1.543	157.360
	(0.009)	(1.626)	(1909.424)	(0.470)	(1.298)	(1.255)	(0.246)	

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	0.000	899.286	2.977	2.713	1.519	156.763
	(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 VI
Time Series Properties of Uncertainty and Volatility

This table displays estimates of regressions

$$y_t = b + \sum_{i=1}^{n} \psi_i \, \text{vol}_{t-i}(14.939) + \sum_{i=1}^{m} \varphi_i \, \text{unc}_{t-i}(15.346) + e_{t+1}$$

on past predictors of volatility $\text{vol}_{t-i}(\omega)$ and measures of past uncertainty $\text{unc}_{t-i}(\nu)$ where the variance of the error term, e_{t+1}, is assumed constant. We vary the dependent variables and the values of n and m. The standard errors of the parameters are listed under the variables in parenthesis. If there is no standard error present then the variable was fixed and not estimated. In this case the value of the variable in the estimate column is the value at which it is fixed. In Panel A we set $y_t = \text{unc}_t(15.346)$ 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 volatility and past uncertainty. In Panel C we set $y_t = \text{vol}_t(14.939)$ and consider regressions of predictors of volatility on past predictors 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	0.000	0.000	0.288	0.000	0.00000000023	1325.435
(0.000)	(0.00039)			(0.119)		(0.00000000009)	
0.000	0.00044	0.000	0.000	0.211	0.238	0.00000000022	1329.742
(0.000)	(0.00038)			(0.106)	(0.066)	(0.00000000008)	

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	0.000	0.000	-10.264	0.000	0.00005323	479.700
(0.001)	(0.088)			(32.786)		(0.00003455)	
0.003	0.388	0.033	0.174	0.000	0.000	0.00005230	480.914
(0.001)	(0.092)	(0.064)	(0.087)			(0.00003470)	

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	0.000	0.000	-0.620	0.000	0.00002733	525.365
(0.001)	(0.065)			(18.734)		(0.00000628)	
0.002	0.190	0.077	0.339	0.000	0.000	0.00002359	535.438
(0.001)	(0.064)	(0.057)	(0.130)			(0.00000501)	

Figure 1. Time series weights (volatility) and cross-sectional weights (model uncertainty)

Quasi-likelihood estimates of the parameters appearing in Table II are used to compute volatility $\text{vol}_t(\omega)$, specified in equation (29), and a measure of uncertainty $\text{unc}_t(\nu)$, appearing in (35). The estimates are based on the quasi-likelihood estimator appearing in (31). The top graph displays the weights on lagged daily volatility when $\omega = 14.939$ and the bottom graph displays the weights on forecasters when $\nu = 15.346$. In the top graph, the x-axis represents lagged trading days within a quarter and the y-axis represent weights. So the weight on daily volatility on the last day of the quarter corresponds to $x = 1$ and is a little less than 0.1. In our data set the number of forecasters varies over time (see the discussion in Section A). The bottom graph displays the weights on forecasters for a quarter in which there are 26 available forecasters ($f_t = 26$). So 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.

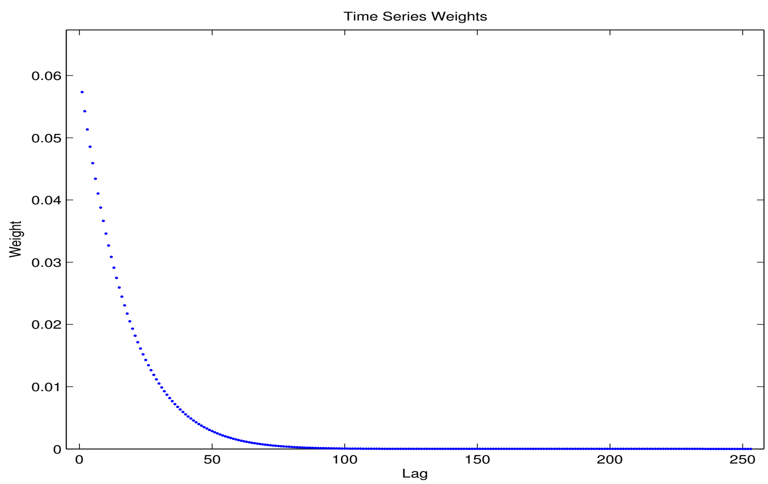

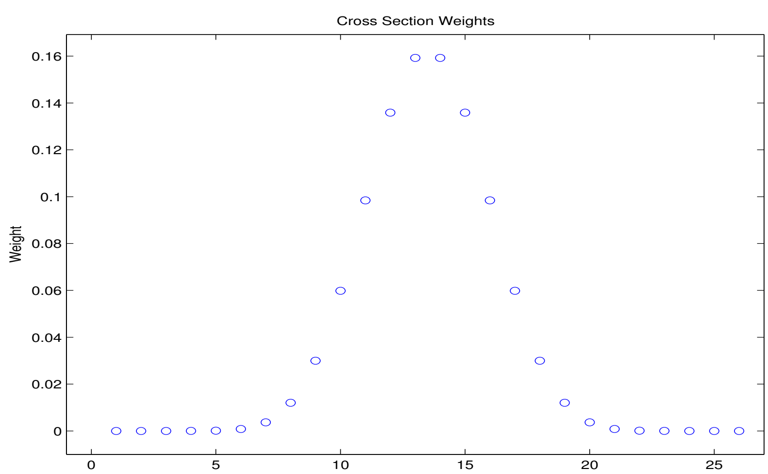

Figure 2. Uncertainty and Returns

The top graph displays a quarterly scatter plot of the actual real market excess return (x-axis) r_{et+1} and the unweighted (or flat-weighted) lagged variance of market return forecasts (y-axis) $\text{unc}_t(1)$. The middle graph displays a quarterly scatter plot of the actual real market excess return (x-axis) r_{et+1} and the Beta-weighted lagged variance of market return forecasts (y-axis) $\text{unc}_t(15.346)$. The bottom graph displays a quarterly scatter plot of the actual real market excess return (x-axis) r_{et+1} and a Beta-weighted variance of market return forecasts (y-axis) $\text{unc}_t(121.5)$.

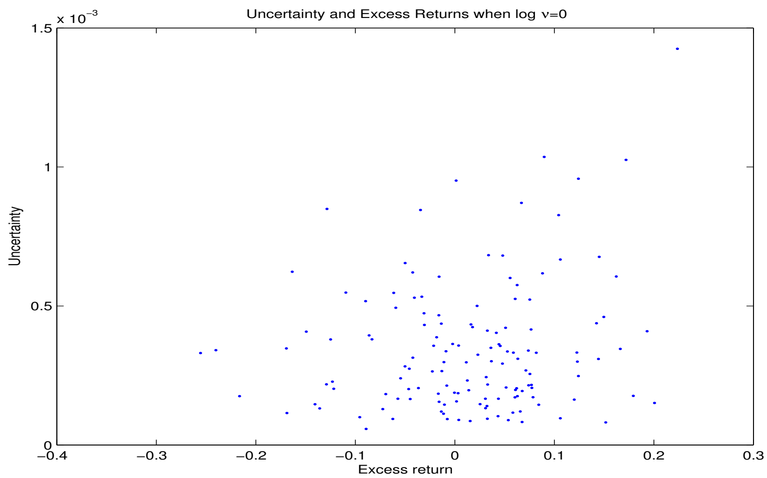

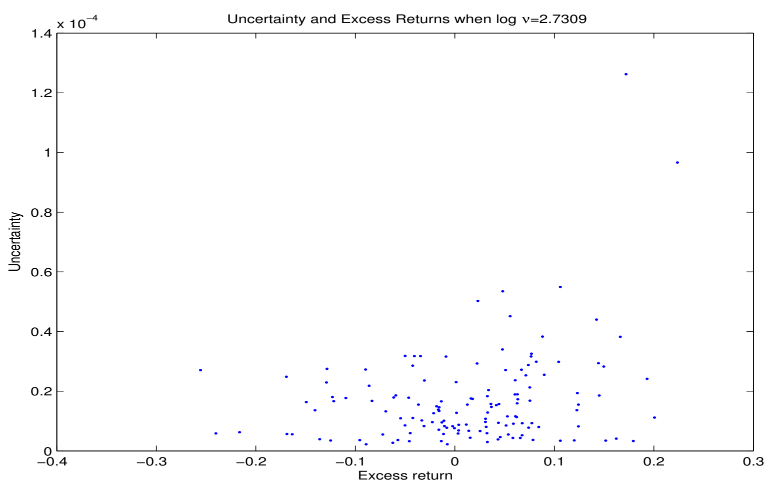

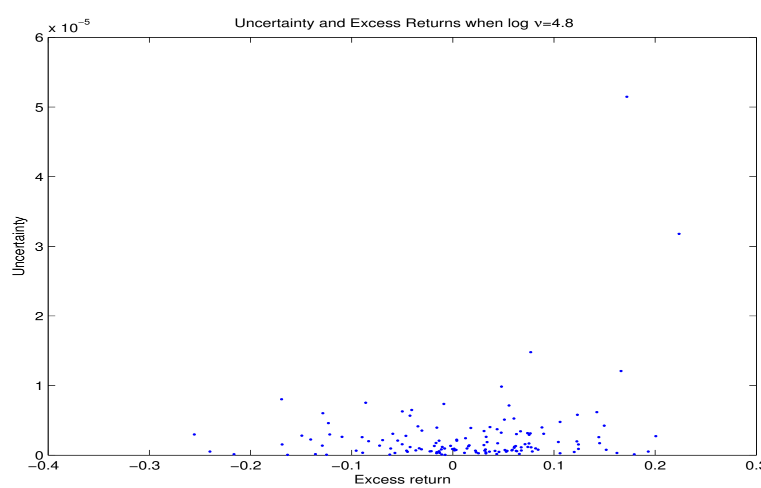

Figure 3. Likelihood

This figure displays the log-likelihood as a function of $\log \nu$ and $\log \omega$ for the fourth specification in Panel D of Table II.

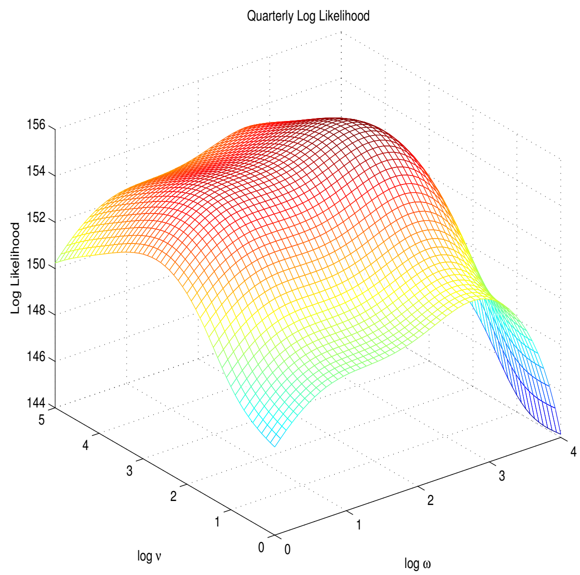

Figure 4. Time Series Plots

In the top row the left figure displays a plot of the quarterly excess return and the right figure displays a quarterly plot of uncertainty, $\mathrm{unc}_t(15.346)$. In the bottom row the left figure displays a plot of quarterly realized volatility, $\mathrm{vol}_t(1)$ and the right figure displays a plot of the predicted volatility for the following quarter, $\mathrm{vol}_t(14.939)$.