

Time-Varying Risk Aversion and the Risk-Return Relation

Daniel R. Smith^a and Robert F. Whitelaw^{b*}

This version: June 19, 2009

PRELIMINARY and INCOMPLETE

Abstract

Time-varying risk aversion is the economic mechanism underlying several recent theoretical models that appear to match important features of equity return data at the market level. In this paper, we estimate a time-varying risk-return relation using only market level data that allows for feedback from both news about volatility and news about risk-aversion. Allowing for feedback effects dramatically improves our ability to explain variation in returns, and the estimated model exhibits statistically and economically significant variation in the price of risk. Consistent with theoretical intuition, the price of risk varies counter-cyclically, with risk aversion increasing substantially over the course of economic contractions. Interestingly, it is variation in the price of risk, not in the quantity of risk, that is the dominant component of the equity risk premium. This phenomenon may partially account for the puzzling results that often arise in estimating models with an assumed constant price of risk.

^aFaculty of Business Administration, Simon Fraser University (8888 University Drive, Burnaby, BC, Canada, V5A 1S6, Email: drsmith@sfu.ca, <http://www.sfu.ca/~drsmith>), ^bStern School of Business, New York University and NBER (44 W. 4th St., Suite 9-190, New York, NY 10012, Email: rwhitela@stern.nyu.edu). Thanks to seminar participants at the London School of Economics, HEC Lausanne and INSEAD for helpful comments.

I. Introduction

A number of prominent papers use time-varying risk aversion in a representative agent setting to explain various features of asset prices. Perhaps the best known example is the use of models in which the representative agent has preferences that exhibit habit formation, of varying forms, in order to explain stock return predictability or the equity premium.¹ The intuition is that a negative shock to consumption pushes the representative agent closer to the habit level, thereby increasing risk aversion and the required return on risky assets. A similar effect can be generated in models with heterogeneous agents whose utility functions take the same form but whose risk aversion varies. In these models assets are priced by a representative agent whose risk aversion is a weighted average of those of the heterogeneous investors, where the weights depend on relative wealth.² Those with lower risk aversion hold a larger fraction of their wealth in risky assets; thus, positive returns on these assets tilt the wealth distribution towards these less risk averse investors and decrease the risk aversion of the representative agent.

The key features of all these models is that shocks to the underlying economy cause changes in risk aversion which subsequently feed through to future expected returns. Time-varying risk aversion is thus a critical element in the determination of the time-varying equity risk premium. However, there have been few direct empirical investigations of the role of time-varying risk aversion in determining the risk-return relation at the market level.

Using options market data, there have been a number of attempts to estimate risk aversion by exploiting the difference between the risk neutral and physical distributions of returns.³ But, the limited time series sample of option prices limits the value of these methodologies for examining low frequency (i.e., business cycle frequency) movements in risk aversion and the risk-return relation.

Empirical models motivated by the dynamic CAPM or Merton's (1973) ICAPM with a time-varying price of risk, i.e., time-varying risk aversion, have been estimated in both a

¹See, for example, Abel (1990), Constantinides (1990), and Campbell and Cochrane (1999).

²See, for example, Chan and Kogan (2002) and Weinbaum (2008).

³See, for example, Bliss and Panizgirtzoglou (2004), and Bollerslev, Gibson and Zhou (2005).

U.S. and international setting, most often using a cross-section of returns.⁴ The consensus is that the price of risk varies over time, but again there has been little effort to examine the implications for the market risk premium.⁵

A natural approach to this issue is to exploit the relation between expected returns and volatility at the market level in order to estimate risk aversion and the equity risk premium in a single, consistent specification. In particular, the dynamic CAPM, or equivalently a specialization of Merton's (1973) ICAPM in which demand for hedging shifts in the investment opportunity set, should they exist, is negligible, the expected return on the aggregate wealth portfolio is linearly related to its conditional variance via the relative risk aversion of the representative agent. It is this linear relation that has been estimated in the vast literature on the risk-return relation in the stock market. Until recently, however, the empirical results have been inconclusive in the sense that papers have failed to consistently identify and estimate a reasonable risk aversion coefficient, albeit under the assumption that risk aversion is constant.⁶ The primary reasons for these failures are clear. Both the expected return and conditional variance are unobservable; therefore, the primary technique has been to regress realized returns on an estimate of the conditional variance. Poor conditional variance estimates may severely hamper the regression. More importantly, even under the most optimistic circumstances, the R-squared of this regression is not likely to exceed 5% at a monthly frequency. In other words, the vast majority of the variation in returns is from unexpected returns, with variation in expected returns explaining little, if any, of realized returns.⁷

More recent work has been able to overcome this problem by exploiting further the economic structure of the problem (Guo and Whitelaw (2006), Smith (2008)). Specifically, if expected returns are proportional to the conditional variance and the conditional variance is persistent, then shocks to the conditional variance will generate unexpected returns of

⁴See, for example, Harvey (1989, 1991) and De Santis and Gerrard (1997).

⁵An exception is a recent paper by Brandt and Wang (2007).

⁶See, for example, French, Schwert and Stambaugh (1987), Glosten, Jagannathan and Runkle (1993), and Whitelaw (1994), among many others.

⁷There is still some controversy about whether expected returns at the market level are time-varying and predictable at all. See, for example, the papers in the recent special issue of the *Review of Financial Studies*, 2008, Vol. 24, No. 4.

the opposite sign via the shock to the discount rate, i.e., the expected return, an effect that is known as volatility feedback. Exploiting the existence of this feedback effect potentially allows for a much more efficient estimation of the parameter of interest, relative risk aversion, because this parameter governs not only the expected return but also the unexpected return due to shocks to expected returns. Guo and Whitelaw (2006) and Smith (2008) both report estimates of relative risk aversion that are positive, reasonable in magnitude and tightly estimated, at least in comparison to much of the earlier literature. Both papers assume that risk aversion is constant, but in principle the same insight can be used to estimate a potentially time-varying coefficient of relative risk aversion; the one additional subtlety being that this type of model will generate a risk aversion feedback effect on top of the volatility feedback effect. In other words, shocks to risk aversion imply shocks to expected returns, which in turn imply unexpected returns of the opposite sign. This mechanism could allow for a more precise estimation of risk aversion in the same way that volatility feedback does for the constant risk aversion parameter.

It is this empirical exercise that we undertake in this paper. We parameterize both the conditional variance of market returns and the coefficient of relative risk aversion as linear functions of a set of state variables. For risk aversion we use three financial variables that have been shown to predict returns and that are related to the business cycle—the dividend yield, the credit spread in the corporate bond market, and the term spread in the Treasury market. To model the conditional variance we use the same variables plus the lagged realized variance. The state variables themselves are assumed to follow a VAR. The market return is composed of three terms—the expected return and the unexpected return due to shocks to expected cash flows and shocks to expected returns. In this setting, this latter term is a function of shocks to risk aversion and shocks to the conditional variance, and this joint feedback effect is a quadratic function of the innovations in the state variables.

We estimate the full system and special cases thereof, e.g., with constant risk aversion or no feedback effects, via GMM. Over our preferred sample period, 1952-2005, the key results are threefold. First, the constant risk aversion models, with or without feedback, show limited ability to fit the data and generate risk aversion estimates that are negative and statistically indistinguishable from zero. Second, the model that allows for time-varying

risk aversion, but without feedback effects, also shows limited explanatory power and shows little if any convincing evidence of time-varying risk aversion. Finally, in contrast, the full specification exhibits not only substantially better goodness of fit but also statistically and economically significant variation in risk aversion.

Examining the fitted values for risk aversion more closely demonstrates that this quantity exhibits substantial countercyclical variation. The average value during contractions exceeds that during expansions by more than 2. More dramatically, the increase in risk aversion during the course of a contraction, i.e., between the peak and trough of the cycle, exceeds 4 even though these phases last less than 12 months on average. These results are consistent with the economic intuition underlying equilibrium models with time-varying risk aversion. A further interesting result is that it is variation in risk aversion rather than variation in risk, i.e., the conditional variance of returns, that dominates the equity risk premium. In fact, the latter is negatively correlated with the risk premium even though this premium is defined as the product of the price of risk and the conditional variance.

The remainder of the paper is organized as follows. Section II first extends extant volatility feedback models to include time-varying risk aversion and the associated feedback effects. We then present the econometric methodology for estimating the parameters of interest. Section III introduces the data and presents some preliminary descriptive statistics. In Section IV, we conduct the main empirical analysis with a special focus on the properties of the time-varying risk aversion. Section V concludes.

II. Methodology

A. Time-Varying Risk Aversion and Feedback Effects

In a dynamic CAPM setting, or in Merton's (1973) ICAPM when the effects of changing investment opportunities on expected returns are negligible, the conditional expected excess return on the aggregate wealth (market) portfolio will be proportional to the conditional variance of this return, where the constant of proportionality depends on the risk aversion of the representative agent. This relation between risk and return at the market level has been

the subject of a large literature, spanning perhaps three decades. However, the vast majority of this research assumes a relation that is constant over time. In contrast, the starting point of our analysis is a specification that allows for a time-varying price of risk. While this specification is motivated by the increased attention that has recently been devoted to equilibrium models that incorporate time-varying risk aversion, we do not estimate or test a specific equilibrium model. Instead, we specify a reduced form model that captures the essential nature of these equilibrium approaches.

In particular, we model the conditional expected excess return on the market portfolio m as a linear function of the conditional variance of the market returns where the price of risk γ is allowed to vary through time:

$$E_t(r_{m,t+1}) = \gamma_t Var_t(r_{m,t+1}) = \gamma_t \sigma_{m,t}^2 \quad (1)$$

We further model the price of risk and the conditional variance as linear functions of conditioning information Z_t :

$$\sigma_{m,t}^2 = \sigma_0 + \sigma_1' Z_t \quad (2)$$

and

$$\gamma_t = \gamma_0 + Z_t' \gamma_1 \quad (3)$$

Z_t is the demeaned version of the state vector X_t , which follows a VAR(1) process:

$$X_{t+1} = (I - \Phi)\mu + \Phi X_t + v_{t+1}. \quad (4)$$

with mean $E(X_{t+1}) = \mu$. Therefore, the demeaned state vector $Z_{t+1} = X_{t+1} - \mu$ evolves as

$$Z_{t+1} = \Phi Z_t + v_{t+1}. \quad (5)$$

Working with a zero mean state variable dramatically simplifies the algebra with no loss of generality given the existence of constants in equations (2)-(3).

The feedback effect can be developed formally using the log-linear approximation to the intertemporal budget equation suggested by Campbell and Shiller (1988), Campbell (1991), and Campbell and Ammer (1993). Specifically, returns on the market portfolio are defined as

$$1 + R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t} \quad (6)$$

yielding a continuously compounded return of

$$k_{t+1} = \log(P_{t+1} + D_{t+1}) - \log(P_t) \quad (7)$$

which is approximated by

$$k_{t+1} \approx c + \rho p_{t+1} - p_t + (1 - \rho)d_{t+1} \quad (8)$$

where $p_{t+1} = \log(P_{t+1})$, $d_{t+1} = \log(D_{t+1})$, $\rho = \frac{1}{1 + \exp(d-p)}$, where $\exp(d-p)$ is the average log-dividend-price ratio, and $c = -\log(\rho) - (1 - \rho) \log\left(\frac{1}{\rho} - 1\right)$. Campbell, Lo and Mackinlay (1997, Chap. 7) note that the historical annual average log dividend-price ratio is about 4% and suggest using $\rho = 0.96$ for annual data, $\rho = 0.997$ for monthly data, and $\rho = 0.9998$ for daily data. Solving (8) forward and imposing a transversality condition, $\lim_{j \rightarrow \infty} \rho^j p_{t+j} = 0$, allows us to write the following accounting identity relating current excess returns (r_{t+1}) to changes in future expected returns, expected dividends (Δd_{t+j+1}) and expected risk-free rates ($r_{f,t+j+1}$):

$$r_{t+1} - E_t(r_{t+1}) = \eta_{d,t+1} - \eta_{r,t+1}, \quad (9)$$

where

$$\eta_{r,t+1} = (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+j+1} \quad (10)$$

$$\eta_{d,t+1} = (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+j+1} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{f,t+j+1}. \quad (11)$$

We will refer to the term $\eta_{d,t+1}$, which contains shocks to both future dividend growth and future risk-free rates, as news about dividends or news about cash flows.

In Appendix A we show that when both the price of risk and the conditional variance are linear functions of a common set of state variables the feedback effect becomes a quadratic function of the innovations in the state variables:

$$\eta_{r,t+1} = Av_{t+1} + Bvec(v_{t+1}v'_{t+1} - \Sigma_v) \quad (12)$$

where

$$A = \rho\sigma_0\gamma'_1(I - \rho\Phi)^{-1} + \rho\gamma_0\sigma'_1(I - \rho\Phi)^{-1} \quad (13)$$

$$B = \rho(\gamma'_1 \otimes \sigma'_1)(I - \rho(\Phi \otimes \Phi))^{-1} \quad (14)$$

$$\Sigma_v = E[v_{t+1}v'_{t+1}]. \quad (15)$$

The intuition behind this result is relatively straightforward. The first term captures the linear feedback effects from both risk aversion and the conditional variance that are analogous to the standard volatility feedback effect when the price of risk is constant. The first piece of the coefficient vector, A , corresponds to risk aversion feedback. A positive shock to risk aversion causes subsequent expected returns to be higher, thus reducing current prices, and generating a contemporaneous negative unexpected return. Shocks to the state vector, v_{t+1} , are translated into risk aversion shocks via the coefficient vector, γ_1 . The magnitude of these effects on current returns are determined by how shocks to risk aversion propagate through time, i.e., the process for the state vector, Φ , and the “discount” factor, ρ . The second piece of A is the corresponding effect due to shocks to the conditional variance, which also move expected returns and thus contemporaneous unexpected returns.

The second term captures the quadratic effects of shocks to the state vector. Since expected returns are the product of two linear functions of the state vector, Z_t , they depend on the squares and the cross-products of the elements of this vector. Again, the coefficient vector, B , depends on how shocks translate into risk aversion and conditional variance and on how those shocks propagate through time. The relevant shocks themselves are the cross-products of the individual shocks less their expected values. For simplicity we assume that the covariance matrix of the shocks to the state vector, Σ_v , is constant over time, i.e., that the state variables do not display heteroscedasticity. This assumption is somewhat at odds with the notion that it is the time-varying variance of market returns that drives, at least in part, time-variation in the market risk premium. Of course, heteroscedasticity in returns could come from unmodeled heteroscedasticity in news about cash flows. From an empirical perspective, as we will see later, this quadratic interaction term is relatively unimportant.

Combining equation (12) with the decomposition of returns in equation (9) and expected returns as given in equations (1)-(3) and adding a constant, we get the full specification for returns:

$$r_{t+1} = \lambda + (\sigma_0 + \sigma_1' Z_t) \cdot (\gamma_0 + \gamma_1' Z_t) - Av_{t+1} - Bvec(v_{t+1}v_{t+1}' - \Sigma_v) + \eta_{d,t+1}, \quad (16)$$

where the coefficient vectors, A and B , are defined in (13)-(14). The model implies that the constant term, λ , should be zero. However, a non-zero constant may capture a risk premium associated with the extent to which the return on the market hedges shifts in investment

opportunities. Also, note that the news about dividends, $\eta_{d,t+1}$, is the unmodeled residual in this specification. The volatility of this term, i.e., the residual volatility in returns, is a measure of the goodness of fit of the model.

B. Parameter Estimation

We will denote the full vector of model parameters θ :

$$\theta = [\lambda, \gamma_0, \gamma_1', \sigma_0, \sigma_1', \mu, \text{vec}(\Phi)', \text{vech}(\Sigma_v)']'. \quad (17)$$

The parameters from the VAR for the state variables in equation (4), $[\mu, \Phi]$, are estimated from the standard OLS moment conditions for the VAR via GMM. We also need to estimate the covariance matrix of the residuals from this VAR, Σ_v , and we again employ GMM, using the natural moments conditions and the fitted residuals from the estimated VAR.

To estimate the parameters of the conditional variance specification in equation (2), we project the realized variance in each month on the demeaned state vector:

$$V_{t+1}^2 = \sigma_0 + \sigma_1' Z_t + v_{V,t+1} \quad (18)$$

We use lagged realized variance as one element of the state vector; thus, equation (18) is subsumed by the specification of the VAR for the state variables in equation (4) and the associated parameters are elements of the larger VAR parameter set. As before, the GMM moment conditions are those from OLS. Following French, Schwert and Stambaugh (1987), among others, the monthly realized variance is computed from daily returns within the month as

$$V_t^2 = \sum_{s \in t} r_s^2 + 2 \sum_{s \in t} r_s r_{s+1} \quad (19)$$

where r_s denotes daily returns, and $s \in t$ refers to all days in month t . The second term adjusts for serial correlation in returns induced, for example, by nonsynchronous trading, and the results are insensitive to the inclusion of this adjustment. Subtracting the estimated mean daily return within the month from the individual daily returns also has no qualitative effect on the results.

The final set of parameters of the model, and those of primary interest, $\lambda, \gamma_0, \gamma_1'$, are estimated from the return specification in equation (16). Specifically, we use the moment conditions associated with maximizing the likelihood function of the news about cash flows.

We estimate the model parameters jointly by stacking the moment conditions associated with the individual sets of parameters:

$$g_t = \left(\begin{array}{c} \frac{\partial \log f(\eta_{d,t+1})}{\partial \theta_\gamma} \\ \left(\begin{array}{c} u_{t+1} \\ v_{t+1} \end{array} \right) \otimes \left(\begin{array}{c} 1 \\ Z_t \end{array} \right) \\ \text{vech}(v_{t+1}v_{t+1}' - \Sigma) \end{array} \right) \quad (20)$$

where $\theta_\gamma = [\lambda, \gamma_0, \gamma_1']$. The news about future dividends, $\eta_{d,t+1}$, comes from a simple rearrangement of equation (16)

$$\eta_{d,t+1} = r_{t+1} - \lambda - (\gamma_0 + \gamma_1'Z_t) \cdot (\sigma_0 + \sigma_1'Z_t) + Av_{t+1} + Bvec(v_{t+1}v_{t+1}' - \Sigma) \quad (21)$$

This system of moment conditions for the full specification is exactly identified. When we examine special cases of this model, e.g., a constant price of risk or suppressing the feedback terms, we modify the definition of θ_γ to maintain exact identification.

Given that we do not perform over-identifying restrictions tests, the adequacy of the various models can be measured in two ways. First, as mentioned above, we add a constant to the return specification, while all the models imply that this constant should be zero. Second, it is easy to compute the fitted residuals, $\eta_{d,t+1}$, at the estimated parameter values. The variance of this term measures the fit of the model. We thus compute and report the pseudo- R^2 statistic:

$$R^{2*} \equiv 1 - \frac{\text{Var}(\eta_{d,t+1})}{\text{Var}(r_{t+1})} \quad (22)$$

which represents the fraction of return variance explained by the model.

III. Data

Our primary analysis spans the period 1952-2005. We use monthly market returns and risk-free rates from CRSP to compute excess returns. We also use daily returns from the same source and over the same period to compute realized monthly variance as described above.

The three state variables are (1) the dividend yield (computed over the prior) 12 months, (2) the term spread—the ten-year over the three-month yield spread in the Treasury market—and (3) the default spread—the yield spread between Baa- and Aaa-rated corporate bonds. The use of these three financial variables has a long history in the literature, including for the prediction of future equity returns and volatilities. Data on these variables comes from the Federal Reserve Economic Data (FRED) maintained by the Federal Reserve Bank of St. Louis.⁸

Figure 1 plots the three financial state variables and realized volatility over the period 1935-2005. These plots partly motivate the choice of the sample period. It is reasonably clear that all three state variables, but particularly the interest rate series, are subject to a structural break at some point in the post-WWII period. A natural breakpoint is 1952, since it was the Fed-Treasury Accord of 1951 that restored the independence of the Federal Reserve.

NBER business cycle peaks and troughs are also shown on the plots as solid vertical lines. All three variables exhibit cyclical variation, a necessary condition for them to pick up cyclical variation in risk aversion, should it exist. There are also good economic reasons to think that these variables might proxy for variation in the price of risk. Part of the default spread is presumably a risk premium for default risk, which should depend on risk aversion. Similarly, the term spread captures the risk premium associated with long-term versus short-term securities. Finally, the dividend yield will pick up variations in the discount rates on equity cash flows associated with changing risk aversion.

The time series plots in Figure 1 also highlight a few other notable features of the data. First, there is a large spike in realized volatility associated with the stock market crash of October 1987. Realized variance in this month is 0.0671, almost 20 standard deviations away from the average realized variance. This single outlier has a significant influence on the estimation results. For example, realized volatility appears significantly less persistent if the crash is included in the sample. Consequently, we choose to Winsorize the data by replacing the observation in October 1987 with the second highest observation in the sample—0.0189 in July 2002. Second, there is a marked decline in dividend yields over the sample period.

⁸See <http://research.stlouisfed.org/fred2/>.

If this decline is due to a decline in risk aversion and an associated drop in discount rates, then our model will pick up this feature of the data. On the other hand, if this decline is due to other factors, it will degrade the fit of our model.

Finally, both yield spreads exhibit unusual behavior during the Volcker Fed experiment of the late 1970s and early 1980s. This phenomenon may also add noise to the results. As a check on the importance of these latter two phenomena, we also examine results for two subsamples, 1952-1978 and 1979-2005. The breakpoint is at the middle of the full sample, but it also coincides with the beginning of the Fed experiment.

Table I shows descriptive statistics for the data. Panel A reports the means, standard deviations, autocorrelations and cross-correlations for log excess returns, realized volatility, and the three financial state variables—the dividend yield (DY), the default spread (DEF), and the term spread (TERM). Returns and volatilities are not annualized; they are reported on a monthly basis. For ease of interpretation we report the realized standard deviation rather than the realized variance. Finally, we report statistics for both the original and the Winsorized conditional volatility series.

It is no surprise that the Winsorized volatility series has a lower mean and standard deviation than its original counterpart. It is also more persistent, with higher autocorrelations at both lags 1 and 12. However, the two series are highly correlated. The correlations between both series and excess returns are negative. Of course, these negative correlations are not inconsistent with a positive risk-return tradeoff. Realized returns and realized volatility both consist of two components, a component due to expectations at the beginning of the month and a shock, or unexpected, component. The risk-return tradeoff refers to the positive correlation between the expected components of the two series. Given this effect, the correlations between the unexpected components should be negative due to the volatility feedback effect—shocks to volatility increase expected returns and thus generate a contemporaneous negative unexpected return. The negative correlations in Panel A simply suggest that this latter effect may dominate the data, and that modeling this effect may be important for identifying the risk-return tradeoff.

The state variables are all highly persistent, with first order autocorrelations exceeding

0.95. Interestingly, in spite of what appear to be common cyclical components in Figure 1, these variables are not highly correlated. These low correlations are important in that the variables may provide independent information about fluctuations in both the conditional variance and the price of risk.

Panel B of Table I reports the extent to which the same 6 series fluctuate with the business cycle. We report means and standard deviations conditional on NBER expansions and contractions. For the three financial state variables we also report the average change in the variable between the peak and trough, and trough and peak, of the cycles.

As expected, returns are both lower and more variable during contractions. Confirming this latter phenomenon, realized volatility is also higher during economic downturns. Interestingly, realized volatility is also more variable in these same periods, an effect that is especially pronounced for the Winsorized series. This empirical phenomenon is consistent with the equilibrium stochastic volatility model in Danielsson, Shin and Zigrand (2009).

The state variables exhibit pronounced cyclical patterns. Dividend yields and default spreads are both higher in contractions, while the reverse is true for the term spread. This latter effect is slightly counter-intuitive because one normally associates flatter term structures with higher inflation periods and thus expansions; however, it is the timing of the movements in the slope of the term structure that generates this effect. This timing is illustrated nicely by the changes in the variables over the business cycle. All three variables increase during contractions (from peak to trough) and decrease during expansions (from trough to peak). In the case of the term spread, this decrease during expansions happens later in the cycle, leaving the average term spread during expansions higher than it is during contractions.

IV. Empirical Results

Table II shows estimates for the realized variance process given in equation (18). Specifically, we regress the monthly realized variance on the lagged realized variance and the three financial state variables (DY, DEF and TERM) and on the three financial variables alone. The top third of the table reports results for the full sample period (1952-2005), with the results

for the two equal subsamples reported below. Heteroscedasticity-consistent t-statistics are reported in parentheses below the corresponding estimate.

The basic conclusions from the estimation are straightforward. Lagged realized variance is the primary predictor, accounting for the majority of the explanatory power in the regressions and having the highest significance levels. The three financial variables are also statistically significant, with a single exception noted below. The variance is positively related to the default spread, i.e., returns are volatile when default spreads are high, and negatively related to the term spread, i.e., returns are volatile when the term structure is flat or inverted. The relation between the dividend yield and variance is less clear. The coefficient is negative and significant in the full sample and the second subsample, but positive and insignificant in the first subsample. The general decline in the level of the dividend yield during the sample period may account for this instability.

The fitted annualized, conditional standard deviation of returns from the full sample estimation is plotted in Figure 2. The range of fitted values seems reasonable. High frequency movements are determined largely by the lagged realized variance, with spikes in the conditional variance associated with spikes in this series as soon in Figure 1. Lower frequency movements are determined by both the lower frequency fluctuations in the lagged realized variance and by those in the three financial state variables. Conditional volatility appears to be higher during business cycle contractions, a feature that we will examine in somewhat more detail when we look closely at the equity risk premium.

Table III shows the estimation results for the VAR(1) of the state variables given in equation (4). As in Table II, we run the regressions for the full sample and two subsamples and report heteroscedasticity-consistent t-statistics in parentheses. Note that realized variance is included as a state variable; therefore, the top row of each panel duplicates the results for the realized variance dynamics reported in Table II.

As expected, the predominant feature of the data is the persistence of the individual variables. In every case, the lagged dependent variable is the primary predictor. Nevertheless, there are a number of statistically significant coefficient estimates scattered among the off-diagonal elements of the coefficient matrix. For example, both the lagged dividend yield and

lagged realized variance are positively related to the default spread.

We report the key estimation results in Table IV. While all the parameters are estimated jointly by stacking the moment conditions as in equation (20), Table IV reports the coefficients in the specification of the price of risk that are identified by the moments from the return equation (16), with heteroscedasticity-consistent standard errors in parentheses. We also report the estimate for the added constant, λ . We estimate the full specification and three special cases thereof. In the first column, the price of risk is assumed to be constant, and the volatility feedback term is not included. The second column maintains the constant price of risk but adds the volatility feedback term. Column three adds a possibly time-varying price of risk but ignores feedback, both via volatility and via shocks to the price of risk. Finally, the last column allows for a time-varying price of risk with full feedback.

The results are quite striking. The constant price of risk specifications in the first two columns replicate well known results from the literature. Risk aversion is hard to pin down, and the point estimate is negative, albeit statistically insignificant. This failure to identify a positive risk-return tradeoff is the puzzle that has attracted the attention of numerous researchers. The amount of variation in returns that is explained by the model without feedback is minimal. Adding volatility feedback increases explained variation to 12.2%, suggesting the possibility that modeling feedback might be useful in pinning down the price of risk. In both specifications the added constant term, λ , is statistically significant at conventional levels, which can be viewed as a rejection of the models. Of course, this constant could be compensation for the ability, or lack thereof, of equity returns to hedge shifts in investment opportunities. A positive constant is necessary to match the mean positive excess return given that the negative price of risk implies a negative risk premium if taken in isolation.

Adding time-variation in the price of risk without feedback effects in column three generates little improvement in the model over the constant price of risk model (without feedback) in the first column. The constant in the price of risk specification, γ_0 , is still negative and insignificant. This constant represents the average price of risk since the state variables are demeaned. None of the individual coefficients on the demeaned state variables are statistically significant. A test of their joint significance has a p-value of 0.1831, so the hypothesis

that the coefficient vector is equal to zero cannot be rejected. In fact, the hypothesis that variance risk is unpriced, i.e., the coefficients including γ_0 are jointly zero, cannot be rejected either, with a p-value of 0.2709. Finally, the degree of explained variation is unimpressive and the constant, λ , is significant at the 10% level in a one-sided test.

In contrast to the preceding specifications, the full specification in the last column exhibits dramatically different features. The coefficients on both the dividend yield and the term spread in the time-varying price of risk equation are statistically significant, and a test of joint significance has a p-value of 0.0005. In other words, the model provides convincing statistical evidence of time-varying risk aversion. The constant, γ_0 , in the price of risk equation is now positive, with a value of 1.09, although the standard error is still high (2.59). For certain standard utility functions, the price of risk is simply the coefficient of relative risk aversion of the representative agent. Again recognizing that the state variables are mean zero, an average coefficient of relative risk aversion of approximately 1 might be considered somewhat low but not totally unreasonable. In addition, the constant, λ , is much smaller and statistically insignificant, suggesting that, at least by this metric, the data do not reject the model.

Perhaps the most surprising result is that over 80% of the variation in returns is explained by this full specification. That is, the model suggests that the majority of unexpected returns are due to shocks to expected returns either via shocks to the conditional variance or shocks to risk aversion. This amount is larger than has been found in previous return decomposition studies, and there are a number of plausible explanations. First by splitting expected returns into two components, we allow the shocks to the state variables to enter twice, with different coefficients, in the shocks to expected returns as in equation (13). Second, the vast majority of the literature has employed expected return specifications that are linear in the state variables, and state variables that follow a VAR. By modeling both the price of risk and the conditional variance as linear functions, expected returns become quadratic in the same state variables. This nonlinearity adds a nonlinear component to the expected return shocks as in equation (14). These shocks are tightly parameterized by the model, but they do offer a degree of flexibility not seen in previous models. Third, our estimation explicitly incorporates the affects of the parameters not just on expected returns, but also on unexpected returns

via shocks to expected returns, as indicated by the theory.

One obvious way to disentangle these potential explanations is to examine a variance decomposition of our model. Specifically, the model implies that returns are made up of five components—expected returns, linear risk aversion feedback, linear variance feedback, the quadratic interaction feedback, and the unmodeled residual variation in news about cash flows. Table V shows the covariance matrix of the four modeled components at the estimated parameter values for the full specification from Table IV. For ease of interpretation, all the variances and covariances are expressed as percentages of the unconditional variance of returns. Hence the number 4.53 in the top righthand corner of the table indicates that fitted expected returns have a variance equal to 4.53% of the total variance of returns.

The major implications of the table are clear. First and foremost, the variance of the linear risk aversion feedback term is the dominant element, and therefore it is risk aversion feedback that accounts for much of the explained variation documented in Table IV. Consequently, the fact that the methodology explicitly incorporates the feedback effect in parameter estimation is the primary explanation for the explanatory power of the model. Second, the variance of expected returns is larger than one would get with a standard linear, or even nonlinear specification. For example, a simple linear model using lagged realized variance and the three financial state variables as explanatory variables generates an R^2 of only 1.7% over the same sample. Adding an additional six quadratic terms to capture the nonlinear nature of the fitted expected return in our specification fails to bring the R^2 over 3%. Thus, incorporating shocks to expected returns as a determinant of unexpected returns suggests greater variation in expected returns than modeling expected returns in isolation.

The pseudo R^2 number in Table IV and the variance decomposition in Table V are not directly comparable to the variance decomposition results in the earlier literature (Campbell (1991), Campbell and Ammer(1993)). These earlier papers model expected returns in a standard regression setting, then examine the relative variance of unexpected returns due to shocks to expected returns relative to the total variance of unexpected returns. In the notation of our paper, they examine the ratio

$$\frac{Var(\eta_{r,t+1})}{Var(r_{t+1} - E_t(r_{t+1}))} \quad (23)$$

For our full specification this ratio is 79.9%, i.e., innovations in expected returns account for 79.9% of the variation in unexpected returns with the residual variation accounted for by news about cash flows and the covariance between these two shocks. Modeling expected returns as a linear function of lagged realized variance and the three state variables generates a corresponding ratio of 71.5%.

A second issue is how is the explanatory power divided among the three financial state variables—DY, DEF and TERM—that define risk aversion, and hence shocks to risk aversion. Given the complex interaction among these variables over time implied by the VAR, perhaps the easiest way to look at this question is to examine the effects of shocks to these variables in an impulse-response analysis. Fortunately, while the effect of a shock on the future paths of the state variables may be complex, the effect on contemporaneous returns is immediate, with no ongoing impact on unexpected returns. Thus, we compute the effect of a one standard deviation shock to each of the four state variables in turn, setting the other shocks to zero. The linear risk aversion and volatility feedback effects and the sum of these effects is reported in Table VI. The key variable is clearly the dividend yield. A one standard deviation shock to this variable generates an almost 4% decline in contemporaneous returns due to the associated increase in risk aversion. The other effects are at least 15 times smaller in magnitude. The importance of the dividend yield is not completely surprising. Due to the persistent in dividends, shocks to returns and shocks to dividend yields are highly negatively correlated. The model exploits this negative correlation in the feedback effect. However, the negative correlation is not solely responsible for the results. First, even in models without feedback risk aversion is positively related to the dividend yield, although the coefficient is smaller and statistically insignificant (see column 3 of Table IV. Second, the model also identifies a significant effect of the term spread on risk aversion. The negative correlation between shocks to TERM and returns is much smaller, as is the associated response in Table VI, yet the coefficient is statistically significant.

A final issue is whether the explanatory is consistent across the sample or is driven by data during certain subperiods. To address this concern, Figure 3 plots the pseudo R^2 of the model over 5-year rolling periods. Recall that the overall R^2 from Table IV is 81.2%. The date on the x axis refers to the date at the center of the corresponding 5-year period.

The explanatory power of the model is extremely consistent. There are some small jumps, such as when the crash of October 1987 enters and leaves the 5-year period, but the only notable feature of the data is the relative decline in explanatory power at the end of the sample. However, even in this period the R^2 remains above 50%. One possible explanation is a degradation in the information quality associated with the dividend yield. As illustrated in Figure 1, the time series behavior of this series may have changed, perhaps starting at the beginning of the stock market run up in the 1990s.

Looking more closely at the estimation results, Figure 4 plots the fitted time-varying price of risk from the models with (solid line) and without (dashed line) feedback effects, along with the NBER business cycle peaks and troughs. The model with feedback effects produces a price of risk that is relatively reasonable in magnitude. The spike in estimated risk aversion during the July 1981 to November 1982 recession, where it reaches values close to 15, is clearly caused by the contemporaneous high levels of both the default and term spreads. To the extent that these underlying movements are caused by Fed policy and not the underlying economics of risk aversion that we attempt to model, they may be spurious. This type of problem is endemic to any approach using state variables that may pick up ancillary factors, and it suggests the use of multiple variables, as we do, to “diversify” away some of these effects.

A second slightly troubling aspect of the results are the periods of negative fitted prices of risk in the expansions ending in December 1969 and March 2001. Of course, these are fitted values with associated standard errors, so one should not interpret the point estimates too literally. They may also be a symptom of some overfitting. For example, realized variance and estimated conditional variance increased during the latter part of the technology bubble that coincides with the second of these periods. Given the large positive returns to holding equity in this same period, it is perhaps not surprising that the model thinks that variance risk was negatively priced. In addition, as noted above, much of the explanatory power of the model comes from the linear risk aversion feedback term. What matters for this effect is not the level of risk aversion (be it positive or negative), but the changes in risk aversion. As an illustration of this point, the pseudo R^2 during periods in which risk aversion is negative is still 74.8%.

While the fitted price of risk with feedback effects may not be perfect, it clearly has better properties than the same fitted series for the model without feedback. This latter series is almost always negative, although it does appear to pick up the same cyclical features, which is not surprising given that the signs of the coefficients are consistent across the last two columns of Table IV.

Figure 5, combining the price of risk in Figure 4 and the conditional volatility in Figure 2, plots the annualized equity risk premium in percent. This quantity takes on negative values when the price of risk is negative, and it exhibits features of both the price of risk and conditional volatility, although the former appears to be dominant.

While the cyclical behavior of the equity risk premium, the conditional volatility and the price of risk are apparent in the figures, it is worthwhile to look more closely at these series. Panel A of Table VIII provides basic descriptive statistics for these three fitted variables, based on the parameter estimates in Table IV from the full specification. Note that for ease of interpretation we report statistics for the conditional standard deviation rather than for the conditional variance, which is the quantity that appears in the specification of the risk premium.

As noted above the mean price of risk is close to 1, and it exhibits substantial time-variation, with a standard deviation of over 3. This time-variation in risk aversion is key to the economic story of the paper, but this degree of variation relative to the mean also accounts for the anomalous negative fitted values discussed above. The price of risk is extremely persistent, as it must be given the persistence of the three state variables of which it is a linear function, but this also coincides with economic intuition about persistent risk aversion. The conditional volatility is less persistent, a result of including lagged realized variance as a predictor. Interestingly, the two elements of the risk premium are negatively correlated in the sample. Both are counter-cyclical, as we shall discuss below, but their counter-cyclical variation clearly does not coincide. This negative correlation will determine the type and degree of the bias in a mis-specified model with a constant price of risk. In general, a negative correlation will result in a constant coefficient estimate that is biased downwards and a positive intercept. While the magnitude of the correlation is perhaps not sufficient to fully explain the empirical results for the constant price of risk models in Table

IV, it does work in the correct direction. A second interesting phenomenon is that the risk premium is slightly negatively correlated with the conditional volatility, but strongly positively correlated with the price of risk. Clearly, the price of risk is the dominant element in determining the risk premium, which again explains the potential problems associated with ignoring time-variation in this quantity when formulating empirical models of the equity risk premium. This result is also consistent with the earlier evidence that it is shocks to risk aversion that explain most of unexpected returns.

The cyclical behavior of the price of risk, conditional volatility, and risk premium are documented in Panel B of Table VIII relative to NBER business cycle dates. All three series are counter-cyclical, with higher means and standard deviations during contractions, with the price of risk and the risk premium exhibiting the more pronounced behavior. This counter-cyclicity is exactly the driving force behind various equilibrium models that seek to explain stock market predictability and other phenomena in the data. Thus, our reduced form approach can be seen as corroborating the basic economic intuition behind these models. The behavior can be seen most dramatically in the changes between the peaks and troughs of the cycle. From peak to trough, risk aversion increases dramatically in economic terms, with the reverse effect occurring between trough and peak. The conditional volatility effect works in the opposite direction, but it is variation in the price of risk that dominates the risk premium. Thus, on average, the risk premium increases over 1.2% during recessions and drops by over 1.4% during expansions. Recall that these are monthly risk premia, corresponding to annualized numbers of approximately 15% and 17%, respectively. Not only does the model imply that risk aversion varies over time, but also that the economic effect of this variation on the equity premium is economically very substantial.

V. Conclusion

Motivated in part by equilibrium models that use the mechanism of time-varying risk aversion to match features of equity return data, we estimate a reduced form model of market returns that accommodates this time-variation. In order to pin down the parameters of the model, we rely on feedback effects from both shocks to conditional variance and risk aversion that

generate contemporaneous shocks to returns. Our full specification, which models this novel risk aversion feedback, appears to fit the data well. In particular, it captures more than 80% of the variation in returns. The estimated price of risk is significantly time-varying and the dominant component of the equity risk premium. Moreover, it exhibits strong counter-cyclical variation consistent with theory.

There is still much work to be done in understanding the equity risk premium and its components, but incorporating time-varying risk aversion seems to be moving us in the correct direction!

Appendix A: Derivation of equation (12)

The instrumental variables X_t follow a VAR(1) process

$$X_{t+1} = (I - \Phi)\mu + \Phi X_t + v_{t+1} \quad (24)$$

which we find useful to express as a demeaned process $Z_{t+1} = X_{t+1} - \mu$. It is useful to express the VAR(1) dynamics of Z_t in equation (5) as

$$Z_{t+j} = \Phi^j Z_t + \sum_{k=1}^j \Phi^{j-k} v_{t+k}$$

We are now in a position to write the conditional mean as

$$\begin{aligned} E_t(r_{m,t+j+1}) &= E_t[(\sigma_0 + \sigma'_1 \Phi^j Z_t + \sigma'_1 \sum_{k=1}^j \Phi^{j-k} v_{t+k}) \times \\ &\quad (\gamma_0 + Z'_t (\Phi^j)' \gamma_1 + \sum_{k=1}^j v'_{t+k} (\Phi^{j-k})' \gamma_1)] \\ &= \sigma_0 \gamma_0 + \sigma_0 Z'_t (\Phi^j)' \gamma_1 + \sigma'_1 \Phi^j Z_t \gamma_0 + \sigma'_1 \Phi^j Z_t Z'_t (\Phi^j)' \gamma_1 + \\ &\quad \sigma'_1 \sum_{k=1}^j \Phi^{j-k} \Sigma_v (\Phi^{j-k})' \gamma_1 \end{aligned}$$

and the updated expectation given information revealed at time $t + 1$ is given by

$$\begin{aligned} E_{t+1}(r_{m,t+j+1}) &= \sigma_0 \gamma_0 + \sigma_1 Z'_t (\Phi^j)' \gamma_1 + \sigma_0 v'_t (\Phi^{j-1})' \gamma_1 + \sigma'_1 \Phi^j Z_t \gamma_0 + \\ &\quad \sigma'_1 \Phi^j Z_t Z'_t (\Phi^j)' \gamma_1 + \sigma'_1 \Phi^{j-1} v_{t+1} \gamma_0 + \\ &\quad \sigma'_1 \Phi^{j-1} v_{t+1} v'_{t+1} (\Phi^{j-1})' \gamma_1 + \sigma'_1 \sum_{k=2}^j \Phi^{j-k} \Sigma_v (\Phi^{j-k})' \gamma_1 \end{aligned}$$

which implies that

$$\begin{aligned} (E_{t+1} - E_t)(r_{m,t+j+1}) &= \gamma'_1 \Phi^{j-1} v_{t+1} \sigma_0 + \sigma'_1 \Phi^{j-1} v_{t+1} \gamma_0 + \\ &\quad \sigma'_1 \Phi^{j-1} v_{t+1} v'_{t+1} (\Phi^{j-1})' \gamma_1 - \sigma'_1 \Phi^{j-1} \Sigma_v (\Phi^{j-1})' \gamma_1 \\ &= [(\sigma'_0 \otimes (\gamma'_1 \Phi^{j-1})) + (\gamma'_0 \otimes (\sigma'_1 \Phi^{j-1}))] v_{t+1} + \\ &\quad (\gamma'_1 \otimes \sigma'_1) (\Phi^{j-1} \otimes \Phi^{j-1}) \text{vec}(v_{t+1} v'_{t+1} - \Sigma_v) \\ &= \gamma'_1 \Phi^{j-1} v_{t+1} \sigma_0 + \sigma'_1 \Phi^{j-1} v_{t+1} \gamma_0 + \\ &\quad (\gamma'_1 \otimes \sigma'_1) (\Phi^{j-1} \otimes \Phi^{j-1}) \text{vec}(v_{t+1} v'_{t+1} - \Sigma_v) \end{aligned}$$

where we use $vec(ABC) = (C' \otimes A)vec(B)$. We finally can express the feedback effect as

$$\begin{aligned}
\eta_{r,t+1} &= \sum_{j=1}^{\infty} \rho^j (E_{t+1} - E_t)(r_{t+j+1}) \\
&= \rho \gamma'_1 (I - \rho \Phi)^{-1} v_{t+1} \sigma_0 + \rho \sigma'_1 (I - \rho \Phi)^{-1} v_{t+1} \gamma_0 + \\
&\quad \rho (\gamma'_1 \otimes \sigma'_1) (I - \rho (\Phi \otimes \Phi))^{-1} vec(v_{t+1} v'_{t+1} - \Sigma_v) \\
&= \rho (\sigma'_0 \otimes (\gamma'_1 (I - \rho \Phi)^{-1}) + \gamma'_0 \otimes (\sigma'_1 (I - \rho \Phi)^{-1})) v_{t+1} + \\
&\quad \rho (\gamma'_1 \otimes \sigma'_1) (I - \rho (\Phi \otimes \Phi))^{-1} vec(v_{t+1} v'_{t+1} - \Sigma_v)
\end{aligned}$$

When the price of risk is constant

$$E_{t+j+1} r_{m,t+j+1} = \gamma \sigma_{m,t+j+1}^2 \quad (25)$$

and the solution simplifies to

$$(E_{t+1} - E_t) r_{m,t+j+1} = \gamma \sigma'_1 \Phi^{j-1} v_{t+1} \quad (26)$$

Solving forward we obtain

$$\eta_{r,t+1} = \gamma \rho \sigma'_1 (I - \rho \Phi)^{-1} v_{t+1} \quad (27)$$

References

- Abel, A., 1990, "Asset Pricing under Habit Formation and Catching up with the Joneses," *American Economic Review* 80, 38-42.
- Bliss, R., and N. Panigirtzoglou, 2004, "Option-Implied Risk Aversion Estimates," *Journal of Finance* 59, 407-446.
- Bollerslev, T., M. Gibson, H. Zhao, 2005, "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option Implied and Realized Volatilities," working paper.
- Brandt, M., and K. Wang, 2003, "Time-Varying Risk Aversion and Unexpected Inflation," *Journal of Monetary Economics* 50, 1457-1498.
- Brandt, M., and L. Wang, 2007, "Measuring the Time-Varying Risk-Return Relation from the Cross-Section of Equity Returns," working paper.
- Brunnermeier, M., and S. Nagel, 2005, "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals Asset Allocation," working paper.
- Campbell, J., 1991, "A Variance Decomposition for Stock Returns," *Economic Journal* 101, 157-179.
- Campbell, J., and J. Ammer, 1993, "What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns," *Journal of Finance* 48, 3-37.
- Campbell, J., and J. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy* 107, 205-251.
- Campbell, J., A. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ).
- Campbell, J., and R. Shiller, 1988, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies* 1, 195-227.
- Constantinides, G., 1990, "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy* 98, 519-543.
- Chan, L., and L. Kogan, 2002, "Catching up with the Joneses: Heterogeneous Preferences

and the Dynamics of Asset Prices,” *Journal of Political Economy* 110, 1255-1285.

Cho, S., 2008, “What Drives Stochastic Risk Aversion,” working paper.

De Santis, G., and B. Gerard, 1997, “International Asset Pricing and Portfolio Diversification with Time-Varying Risk,” *Journal of Finance* 52, 1881-1912.

French, K., W. Schwert, and R. Stambaugh, 1987, “Expected Stock Returns and Volatility,” *Journal of Financial Economics* 19, 3-30.

Glosten, L., R. Jagannathan, and D. Runkle, 1993, “On the Relation Between the Expected Value and the Variance of the Nominal Excess Return on Stocks,” *Journal of Finance* 48, 1779-1801.

Guo, H., Z. Wang, and J. Yang, 2006, “Does Aggregate Risk Aversion Change Countercyclically over Time? Evidence from the Stock Market,” working paper.

Guo, H., and R. Whitelaw, 2006, “Uncovering the Risk-Return Relation in the Stock Market,” *Journal of Finance* 61, 1433-1463.

Harvey, C., 1989, “Time-Varying Conditional Covariances in Tests of Asset Pricing Models,” *Journal of Financial Economics* 24, 289-317.

Harvey, C., 1991, “The World Price of Covariance Risk,” *Journal of Finance* 46, 111-157.

Merton, R., 1973, “An Intertemporal Capital Asset Pricing Model,” *Econometrica* 41, 867-887.

Smith, D., 2008, “Risk and Return in Stochastic Volatility Models: Volatility Feedback Matters!” working paper.

Weinbaum, D., 2008, “Investor Heterogeneity, Asset Pricing and Volatility Dynamics,” working paper.

Whitelaw, R., 1994, “Time Variations and Covariations in the Expectation and Volatility of Stock Market Returns,” *Journal of Finance* 49, 515-541.

Table I: Descriptive Statistics

Panel A: Full Sample (1952-2005)								
	Mean	SD	AC(1)	AC(12)	Volatility	Correlation		
						Volatility*	DEF	TERM
Returns	0.48%	4.28%	0.07	0.03	-0.30	-0.26		
Volatility	3.84%	2.02%	0.50	0.19		0.97		
Volatility*	3.82%	1.86%	0.56	0.24				
DY	3.25%	1.11%	0.99	0.87			0.29	-0.14
DEF	0.93%	0.42%	0.97	0.68				0.27
TERM	1.38%	1.17%	0.95	0.44				

Panel B: Business Cycle Dependence						
	Expansion		Contraction		Average Change	
	Mean	SD	Mean	SD	Peak-Trough	Trough-Peak
Returns	0.55%	4.00%	0.05%	5.67%		
Volatility	3.63%	1.89%	5.09%	2.32%		
Volatility*	3.61%	1.69%	5.09%	2.32%		
DY	3.11%	1.04%	4.11%	1.14%	0.03%	-0.54%
DEF	0.89%	0.36%	1.22%	0.59%	0.35%	-0.36%
TERM	1.42%	1.17%	1.17%	1.13%	1.94%	-1.95%

Panel A presents descriptive statistics (means, standard deviations, autocorrelations, cross-correlations) for realized returns, realized volatility (both unadjusted and Winsorized, with the latter denoted by a *), and the three financial variables—the dividend yield (DY), the default spread (DEF) and the term spread (TERM)—over the full sample (1952-2005). Panel B presents means and standard deviations for the same variables for expansions and contractions as dated by the NBER. This panel also shows the average change in the financial variable from the peak to the trough of the business cycle and vice versa.

Table II: Parameter Estimates for Realized Volatility Dynamics

Data	Const	V_{t-1}^2	DY	DEF	TERM	R^2
1952-2005	0.0018	0.4326	-0.0002	0.0007	-0.0002	0.2474
	(24.4738)	(12.1117)	(-3.2072)	(3.5054)	(-2.8300)	
	0.0018		-0.0004	0.0014	-0.0003	0.0757
	(22.1183)		(-4.8758)	(6.3521)	(-4.3951)	
1952-1978	0.0015	0.5174	0.0001	0.0009	-0.0005	0.3899
	(18.6061)	(10.5378)	(1.2988)	(2.8551)	(-3.7739)	
	0.0015		0.0004	0.0022	-0.0009	0.1776
	(16.0672)		(3.0461)	(6.5707)	(-7.0751)	
1979-2005	0.0021	0.3145	-0.0006	0.0014	-0.0002	0.1928
	(17.6746)	(5.9021)	(-3.8985)	(3.4244)	(-1.9730)	
	0.0021		-0.0008	0.0020	-0.0003	0.1047
	(16.8278)		(-5.6279)	(5.1144)	(-2.7453)	

We report parameter estimates and T-ratios, in parentheses, of the regression of monthly realized variance, V_t^2 , on the vector of lagged predictor variables and lagged variance, given in equation (18), for the full sample, 1952-2005, and two equal subsamples.

Table III: Estimates of Φ

	V_t^2	DY $_t$	DEF $_t$	TERM $_t$
1952-2005				
V_{t+1}^2	0.4326 (7.4344)	-0.0002 (-2.4773)	0.0007 (3.6238)	-0.0002 (-2.6323)
DY $_{t+1}$	0.0661 (0.0141)	0.9905 (141.6617)	-0.0218 (-1.0061)	-0.0113 (-1.7965)
DEF $_{t+1}$	9.0899 (2.7906)	0.0114 (3.1952)	0.9593 (53.6242)	-0.0072 (-1.3133)
TERM $_{t+1}$	-2.6496 (-0.4093)	-0.0064 (-0.5487)	0.1528 (2.0568)	0.9382 (40.8508)
1952-1978				
V_{t+1}^2	0.5175 (4.1375)	0.0001 (1.4425)	0.0009 (2.9485)	-0.0005 (-3.3036)
DY $_{t+1}$	1.5654 (0.1539)	0.9789 (72.6555)	-0.0322 (-0.6852)	-0.0214 (-1.2399)
DEF $_{t+1}$	12.5369 (1.8419)	0.0101 (2.3992)	0.9727 (40.8962)	-0.0131 (-1.5614)
TERM $_{t+1}$	-19.3831 (-1.6532)	0.0482 (2.9581)	0.3108 (4.7313)	0.8843 (33.8612)
1979-2005				
V_{t+1}^2	0.3148 (6.6320)	-0.0006 (-3.3543)	0.0014 (3.3602)	-0.0002 (-1.8719)
DY $_{t+1}$	-0.3443 (-0.0778)	0.9936 (88.4876)	-0.0163 (-0.4834)	-0.0057 (-0.7493)
DEF $_{t+1}$	8.2582 (2.3562)	0.0332 (3.9440)	0.8973 (32.9256)	-0.0072 (-1.0262)
TERM $_{t+1}$	-2.1232 (-0.2578)	-0.0256 (-0.9704)	0.1492 (1.3700)	0.9406 (31.0060)

We report the estimated autoregressive parameters (Φ) and T-ratios, in parentheses, for the VAR(1) given in equation (4) with realized stock return variance and the cyclical financial predictor variables—the dividend yield (DY), the default spread (DEF), and the term spread (TERM). Results are reported for the full sample, 1952-2005 and two equal subsamples.

Table IV: GMM Parameter Estimates: 1952-2005

Parameter	No Feedback	Feedback	No Feedback	Feedback
λ	0.0080 (0.0037)	0.0108 (0.0057)	0.0080 (0.0053)	0.0020 (0.0043)
γ_0	-1.7728 (2.1897)	-3.5396 (3.3912)	-1.7813 (3.3812)	1.0923 (2.5906)
γ_{DY}			0.6589 (1.3148)	2.0666 (1.0166)
γ_{DEF}			2.1596 (3.1647)	2.1711 (3.1013)
γ_{TERM}			0.7977 (0.8614)	1.8017 (0.8042)
R^{2*}	0.5%	12.2%	1.4%	81.2%
$\gamma'_1 = 0$			4.8502 [0.1831]	17.827 [0.0005]
$[\gamma_0, \gamma'_1] = 0$			5.1640 [0.2709]	18.3528 [0.0011]

We report parameter estimates for the model given in equation (16), and special cases thereof, for the full sample, estimated using GMM and the moment conditions given in equation (20). Heteroscedasticity-consistent standard errors are in parentheses. R^{2*} is the fraction of return variance explained by the model. We also report test statistics for testing two hypothesis—(1) that the price of risk is constant, i.e., that the coefficients on the predictor variables are jointly equal to zero, and (2) that the price of risk is zero—with corresponding p-values in square brackets.

Table V: Variance Decomposition

	Expected	Risk	Feedback	
	Return	Aversion	Volatility	Quadratic
Expected Return	4.53	-0.25	0.02	0.17
Risk Aversion Feedback	-0.25	85.06	-4.91	1.19
Volatility Feedback	0.02	-4.91	1.07	-0.04
Quadratic Feedback	0.17	1.19	-0.04	0.39

We report the variances of and covariances between the four elements on the righthand side of equation (16)—the expected return, risk aversion feedback, volatility feedback, and the quadratic interaction between the latter two feedback effects. The quantities are expressed as a percentage of the variance of returns. Parameter estimates are from Table IV for the full sample.

Table VI: Impulse-Response

	V^2	DY	DEF	TERM
Risk Aversion Feedback	0.08	-3.98	0.20	-0.26
Volatility Feedback	-0.23	0.26	-0.22	0.23
Total	-0.16	-3.72	-0.02	-0.03

We report the contemporaneous effect on unexpected monthly returns, via risk aversion feedback and volatility feedback, due to a one standard deviation shock to the four state variables in turn, setting the other shocks to zero. Parameter estimates are from Table IV for the full sample. The feedback terms are given in equation (12).

Table VII: GMM Parameter Estimates: Subsamples

Parameter	No Feedback	Feedback	No Feedback	Feedback
1952-1978				
λ_0	0.0072 (0.0051)	0.0028 (0.0080)	0.0101 (0.0057)	0.0047 (0.0059)
γ_0	-2.0594 (3.8070)	0.9713 (5.6439)	-4.3260 (4.4793)	-0.6581 (4.3265)
γ_{dy}			4.9633 (2.7047)	3.4474 (1.5912)
γ_{qs}			4.0925 (5.5647)	4.7036 (5.6715)
γ_{ts}			2.7290 (2.5191)	3.3866 (1.7942)
σ_d	0.0402 (0.0021)	0.0402 (0.0022)	0.0389 (0.0018)	0.0145 (0.0012)
1979-2005				
λ_0	0.0094 (0.0056)	0.0154 (0.0060)	0.0067 (0.0102)	-0.0021 (0.0073)
γ_0	-1.8984 (2.7470)	-5.2608 (3.3119)	-0.3601 (5.5745)	2.9416 (3.9189)
γ_{dy}			1.0353 (2.4519)	2.3876 (1.9221)
γ_{qs}			-0.5330 (5.8764)	-0.5670 (4.9680)
γ_{ts}			0.4004 (1.0496)	1.1558 (1.0296)
σ_d	0.0450 (0.0032)	0.0383 (0.0127)	0.0449 (0.0032)	0.0215 (0.0014)

We report parameter estimates for the model given in equation (16), and special cases thereof, for two equal subsamples, estimated using GMM and the moment conditions given in equation (20). Heteroscedasticity-consistent standard errors are in parentheses. R^{2*} is the fraction of return variance explained by the model.

Table VIII: Estimated Price of Risk, Volatility and Risk Premium

Panel A: Full Sample						
	Mean	SD	AC(1)	AC(12)	Correlation Conditional Vol.	Risk Premium
γ_t	1.09	3.37	0.98	0.62	-0.17	0.86
Conditional Vol.	4.12%	1.07%	0.70	0.41		-0.02
Risk Premium	0.35%	0.85%	0.85	0.43		

Panel B: Business Cycle Dependence						
	Expansion		Contraction		Average Change	
	Mean	SD	Mean	SD	Peak-Trough	Trough-Peak
γ_t	0.75	3.16	3.10	3.89	4.31	-5.41
Conditional Vol.	4.01%	0.98%	4.73%	1.36%	-0.35%	0.63%
Risk Premium	0.25%	0.73%	0.90%	1.22%	1.23%	-1.43%

Panel A presents descriptive statistic (mean, standard deviation, autocorrelations, and cross-correlations) for the estimated price of risk (γ_t), the conditional volatility, and the equity risk premium for the full sample period (1952-2005) based on the parameter estimates from Tables II and IV. Panel B presents means and standard deviations for the same variables for expansions and contractions as dated by the NBER. This panel also shows the average change in the value of the series from the peak to the trough of the business cycle and vice versa.

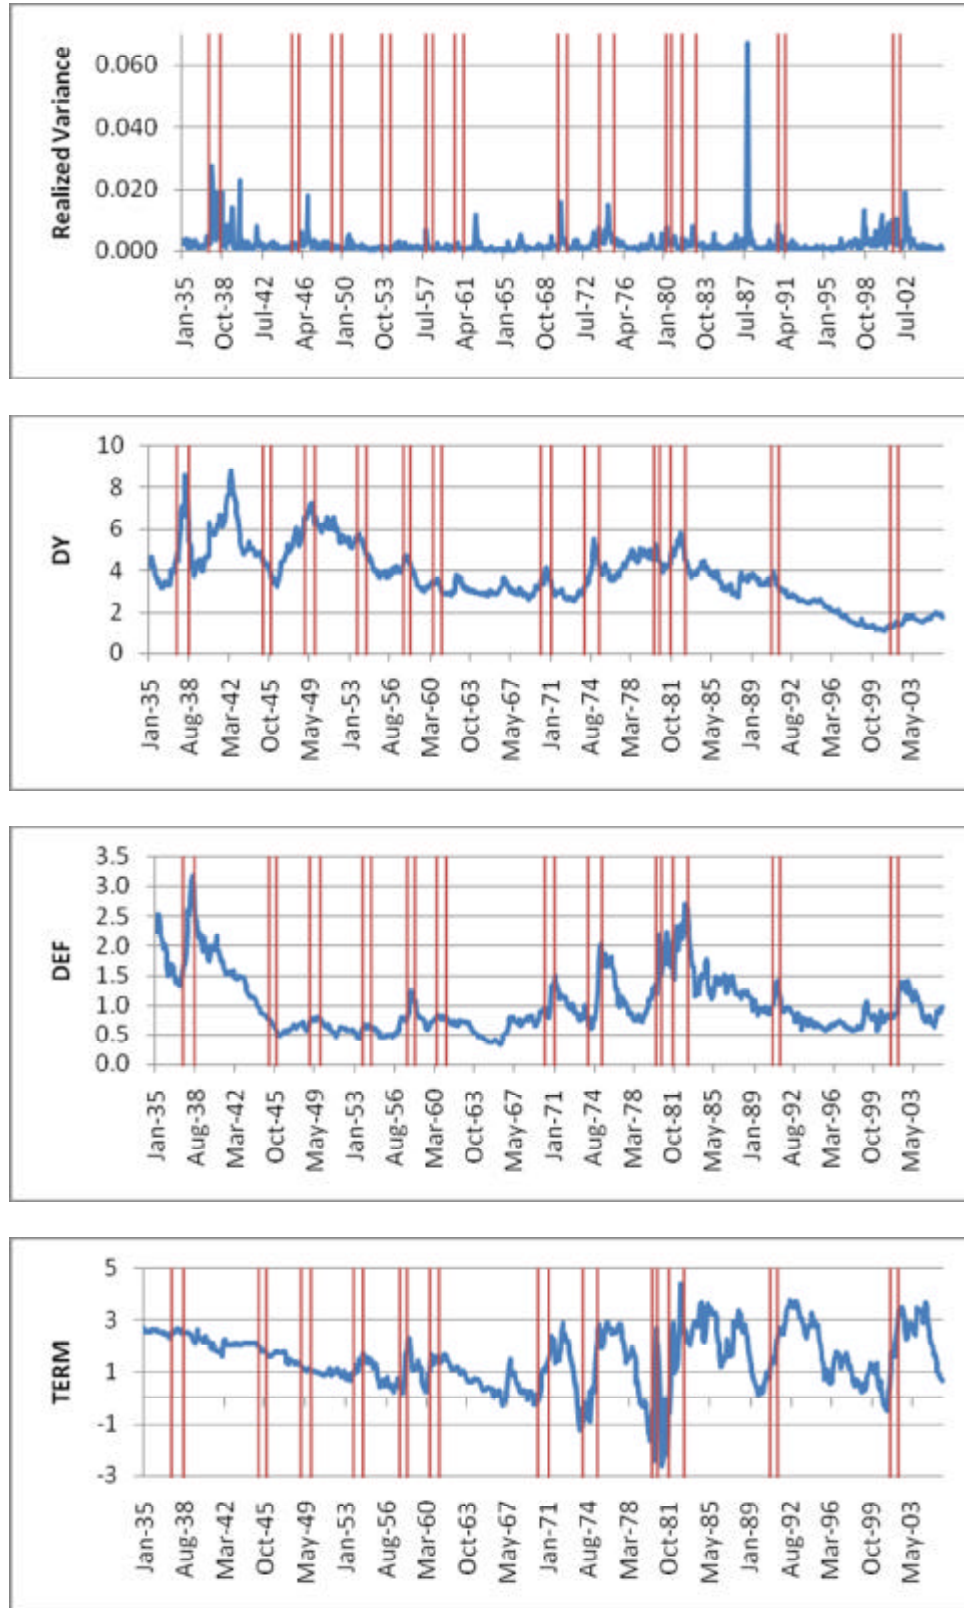


Figure 1: Time Series Plots of Realized Variance and Financial State Variables, 1935-2005
 Vertical lines mark NBER business cycle peaks and troughs.

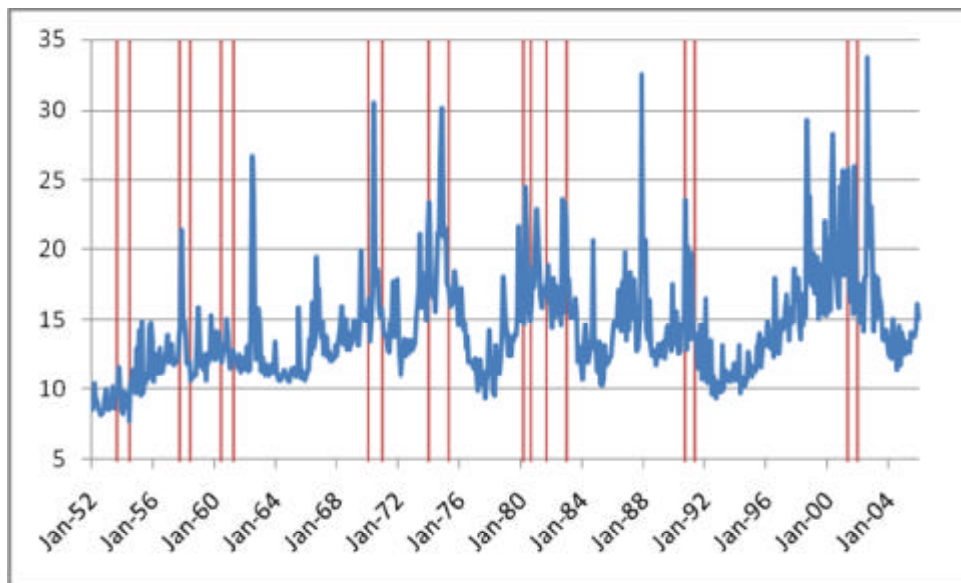


Figure 2: Estimated Annualized Conditional Volatility

Estimated annualized conditional standard deviation of returns (in percent), 1952-2005. Parameter estimates are from Table II. Specification is given in equation (17). Vertical lines mark NBER business cycle peaks and troughs.

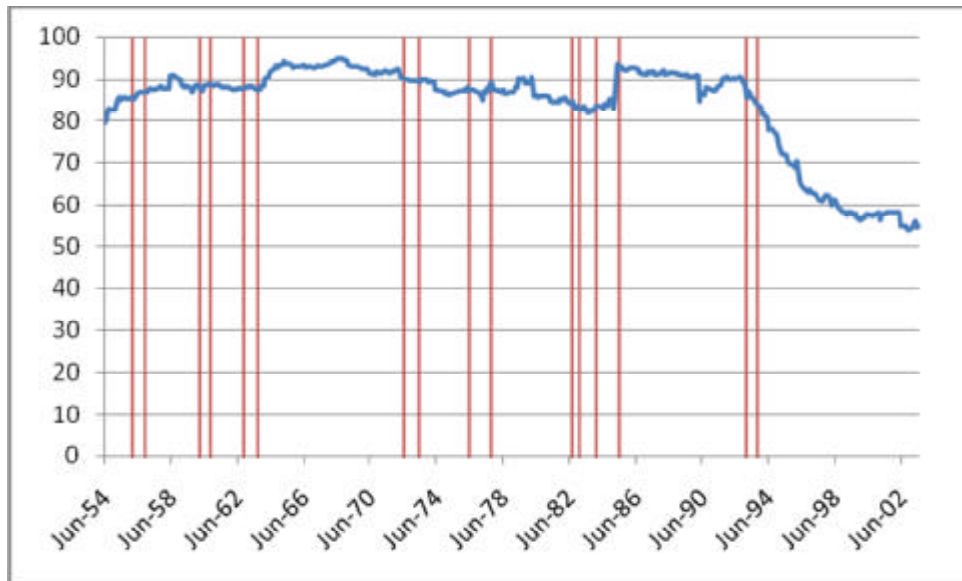


Figure 3: Rolling Pseudo R-squared

5-year rolling pseudo r-squared (in percent), defined in equation (21), 1952-2005. The date on the x axis refers to the month at the center of the rolling window. Parameter estimates are from Table IV. Specification is given in equation (15). Vertical lines mark NBER business cycle peaks and troughs.

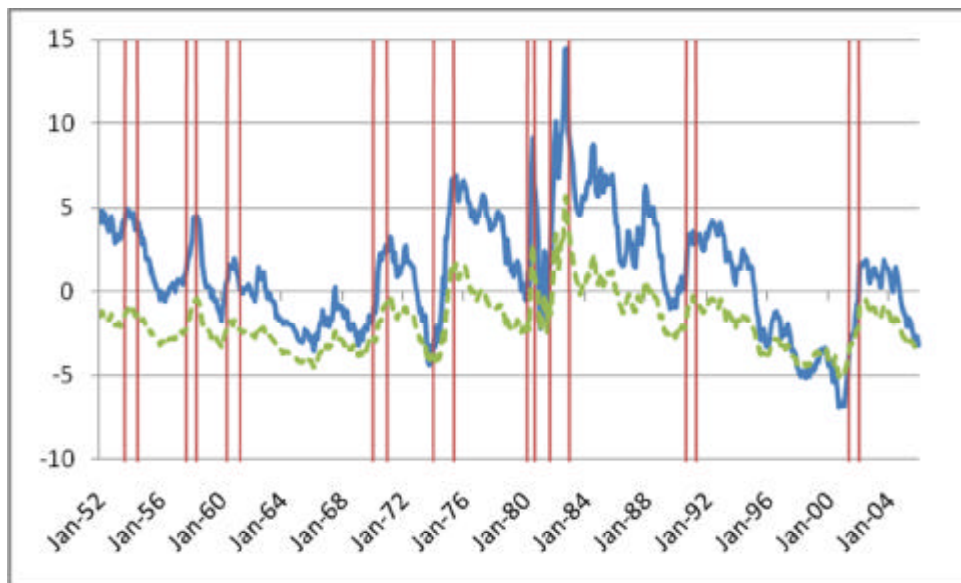


Figure 4: Estimated Price of Risk

Estimated price of risk, 1952-2005, with feedback (solid blue line), without feedback (dashed green line). Parameter estimates are from Table IV. Specification is given in equation (15). Vertical lines mark NBER business cycle peaks and troughs.

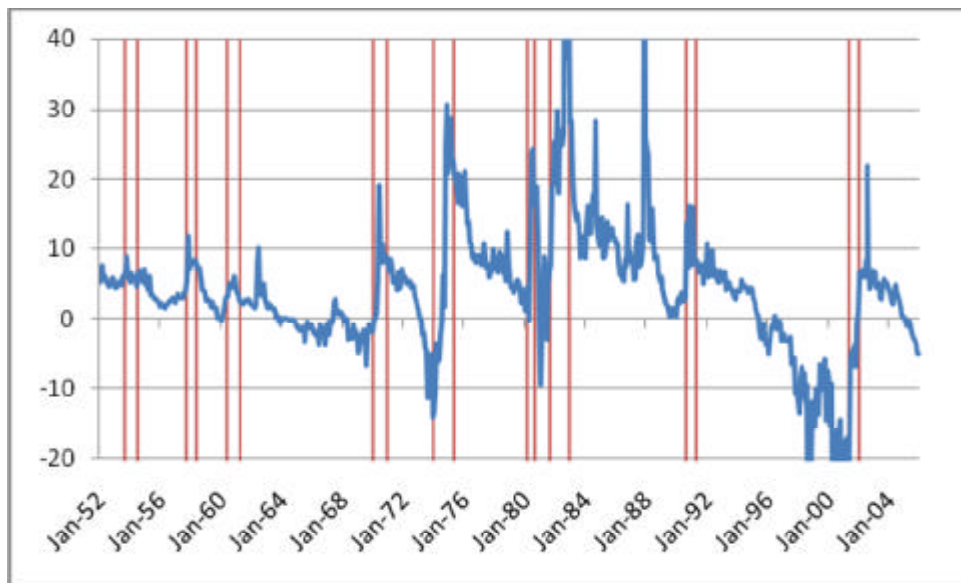


Figure 5: Estimated Annualized Risk Premium (%)

Estimated annualized risk premium (in percent), 1952-2005. Parameter estimates are from Table IV. Specification is given in equation (15). Vertical lines mark NBER business cycle peaks and troughs.