

Uncovering the Risk–Return Relation in the Stock Market

Hui Guo^a and Robert F. Whitelaw^{b*}

July 21, 2003

^aResearch Department, Federal Reserve Bank of St. Louis (P.O. Box 442, St. Louis, Missouri 63166, E-mail: hui.guo@stls.frb.org); and ^bStern School of Business, New York University and NBER (44 W. 4th St., Suite 9-190, New York, NY 10012, E-mail: rwhitela@stern.nyu.edu). This paper formerly circulated under the title “Risk and Return: Some New Evidence”. We thank N.R. Prabhala for providing the implied volatility data and seminar participants at the Federal Reserve Bank of New York, the University of Illinois, Urbana-Champaign, the University of Texas, Austin, Vanderbilt University, the Australian Graduate School of Management, and the University of Queensland for helpful comments. Financial support from the C. V. Starr Center for Applied Economics is gratefully acknowledged by the first author. The views expressed in this paper are those of the authors and do not necessarily reflect the official positions of the Federal Reserve Bank of St. Louis or the Federal Reserve System.

UNCOVERING THE RISK–RETURN RELATION IN THE STOCK MARKET

Abstract

There is an ongoing debate in the literature about the apparent weak or negative relation between risk (conditional variance) and return (expected returns) in the aggregate stock market. We develop and estimate an empirical model based on the ICAPM to investigate this relation. Our primary innovation is to model and identify empirically the two components of expected returns—the risk component and the component due to the desire to hedge changes in investment opportunities. We also explicitly model the effect of shocks to expected returns on ex post returns and use implied volatility from traded options to increase estimation efficiency. As a result, the coefficient of relative risk aversion is estimated more precisely, and we find it to be positive and reasonable in magnitude. Although volatility risk is priced, as theory dictates, it contributes only a small amount to the time-variation in expected returns. Expected returns are driven primarily by the desire to hedge changes in investment opportunities. It is the omission of this hedge component that is responsible for the contradictory and counter-intuitive results in the existing literature.

1 Introduction

The return on the market portfolio plays a central role in the capital asset pricing model (CAPM), the financial theory widely used by both academics and practitioners. However, the intertemporal properties of stock market returns are not yet fully understood.¹ In particular, there is an ongoing debate in the literature about the relationship between stock market risk and return and the extent to which stock market volatility moves stock prices. This paper provides new evidence on the risk-return relation by estimating a variant of Merton's (1973) intertemporal capital asset pricing model (ICAPM).

In his seminal paper, Merton (1973) shows that the conditional excess market return, $E_{t-1}r_{M,t} - r_{f,t}$, is a linear function of its conditional variance, $\sigma_{M,t-1}^2$, (the risk component) and its covariance with investment opportunities, $\sigma_{MF,t-1}$, (the hedge component), i.e.,

$$E_{t-1}r_{M,t} - r_{f,t} = \left[\frac{-J_{WWW}}{J_W}\right]\sigma_{M,t-1}^2 + \left[\frac{-J_{WF}}{J_W}\right]\sigma_{MF,t-1}, \quad (1)$$

where $J(W(t), F(t), t)$ is the indirect utility function with subscripts denoting partial derivatives, $W(t)$ is wealth, and $F(t)$ is a vector of state variables that describe investment opportunities.² $\frac{-J_{WWW}}{J_W}$ is a measure of relative risk aversion, which is usually assumed to be constant over time. If people are risk averse, then this quantity should be positive.

Under certain conditions, Merton (1980) argues that the hedge component is negligible and the conditional excess market return is proportional to its conditional variance. Since Merton's work, this specification has been subject to dozens of empirical investigations, but these papers have drawn conflicting conclusions on the sign of the coefficient of relative risk aversion. In general, however, despite widely differing specifications and estimation techniques, most studies find a weak or negative relation. Examples include French, Schwert and Stambaugh (1987), Campbell (1987), Glosten, Jagannathan and Runkle (1993), Whitelaw (1994), and more recent papers, including Goyal and Santa-Clara (2003) and Lettau and Ludvigson (2003).

¹The expected stock market return was long considered to be constant until relatively recent work documenting the predictability of market returns (e.g., Fama and French (1989)). It is now well understood that time-varying expected returns are consistent with rational expectations. See Campbell and Cochrane (1999) and Guo (2003) for recent examples of this literature.

²Strictly speaking, equation (1) is the discrete time version of Merton's ICAPM (see Long [1974]). In addition, the equation holds for the aggregate wealth portfolio for which we use the market portfolio as a proxy.

The failure to reach a definitive conclusion on the risk-return relation can be attributed to two factors. First, neither the conditional return nor the conditional variance are directly observable; certain restrictions must be imposed to identify these two variables. Instrumental variable (IV) models and autoregressive conditional heteroscedasticity (ARCH) models are the two most commonly used identification methods. In general, empirical results are sensitive to the restrictions imposed by these models. For example, Campbell (1987) finds that the results depend on the choice of instrumental variables. Specifically, the nominal risk-free rate is negatively related to the expected return and positively related to the variance, and “these two results together give a perverse negative relationship between the conditional mean and variance for common stock” (Campbell (1987, p.391)). In the context of ARCH models, if the conditional distribution of the return shock is changed from normal to student-t, the positive relation found by French, Schwert and Stambaugh (1987) disappears (see Baillie and DeGennaro (1990)).

Second, there are no theoretical restrictions on the sign of the correlation between risk and return. Backus and Gregory (1993) show that in a Lucas exchange economy, the correlation can be positive or negative depending on the time series properties of the pricing kernel. This result suggests that the hedge component can be a significant pricing factor and can have an important effect on the risk-return relation. In general, the risk-return relation can be time-varying as observed by Whitelaw (1994). The theory, however, still requires a positive *partial* relationship between stock market risk and return. The more relevant empirical issue is to disentangle the risk component from the hedge component.

Scruggs (1998) presents some initial results on the decomposition of the expected excess market return into risk and hedge components. Assuming that the long-term government bond return represents investment opportunities, he estimates equation (1) using a bivariate exponential GARCH model and finds that the coefficient of relative risk aversion is positive and statistically significant. However, his approach has some weaknesses. For example, he assumes that the conditional correlation between stock returns and bond returns is constant, but Ibbotson Associates (1997) provide evidence that it actually changes sign over time in historical data. After relaxing this assumption, Scruggs and Glabadanidis (2003) fail to replicate the earlier results. Of course, this latter result does not imply a rejection of equation (1); rather, it challenges the assumption that bond returns are perfectly correlated with investment opportunities.

In contrast, we develop an asset pricing model based on Merton’s (1973) ICAPM and Campbell and Shiller’s (1988) log-linearization method and implement estimation using instrumental variables.³ Instead of working with the ex ante excess return, which is not directly observable, we decompose the ex post excess return into five components: the risk component and the hedge component, which together make up expected returns, revisions in these two components, which measure unexpected returns due to shocks to expected returns, and a residual component reflecting unexpected returns due to revisions in cash flow and interest rate forecasts. We explicitly model the volatility feedback effect,⁴ and we also control for innovations in the hedge component. Therefore, we explain part of the unexpected return on a contemporaneous basis and improve the efficiency of the estimation and the identification of the risk and hedge components of expected returns.

Another innovation relative to previous work is that we use non-overlapping monthly volatility implied by S&P 100 index option prices as an instrumental variable for the conditional market variance.⁵ Implied volatility is a powerful predictor of future volatility, subsuming the information content of other predictors in some cases (see, for example, Christensen and Prabhala (1998) and Fleming (1998)). Implied volatility is therefore an efficient instrumental variable and improves the precision of the estimation.

We get three important and interesting results from estimating the model with the implied volatility data. First, the coefficient of relative risk aversion is positive and precisely estimated. For example, in our favored specification the point estimate is 3.98 with a standard error of 1.45. Second, we find that expected returns are primarily driven by changes in investment opportunities, not by changes in stock market volatility. The two together explain 5.2% of the total variation in stock market returns, while the latter alone explains less than 1% of the variation. Moreover, the variance of the estimated hedge component of expected returns is approximately twice as big as that of the risk component. Third, the risk and hedge components are negatively correlated. Thus the omitted variables problem caused by estimating equation (1) without the hedge component can

³French, Schwert and Stambaugh (1987) argue that full information maximum likelihood estimators such as GARCH are generally more sensitive to model misspecification than instrumental variable estimators.

⁴Pindyck (1984,1988), Poterba and Summers (1986), French, Schwert and Stambaugh (1987), Campbell and Hentschel (1992) and Wu (2001) all emphasize the importance of the volatility feedback effect in detecting the risk-return relation.

⁵The implied volatility data is constructed by Christensen and Prabhala (1998) and is kindly provided to us by N. Prabhala.

cause a severe downward bias in the estimate of relative risk aversion.

One concern is that the implied volatility data only span the period November 1983 to May 1995 (139 observations). In order to check the robustness of our results, we also estimate the model with longer samples of monthly and quarterly data, in which the conditional market variance is estimated with lagged financial variables. The results from this empirical exercise are also more readily compared to those in the existing literature. Similar results are found in this longer dataset. For the monthly data, the point estimate of the coefficient of relative risk aversion is 1.46 with a standard error of 3.17. In spite of the longer sample, the standard error is higher due to the imprecision associated with estimating the conditional variance rather than using implied volatility. For the quarterly data, the estimate of relative risk aversion is 7.90 with a standard error of 3.50. In both cases, expected returns are driven primarily by changes in investment opportunities.

These analyses allow us to explain the counter-intuitive and contradictory evidence in the current literature. The primary issue is a classical omitted variables problem. Because the omitted variable, the hedge component, is large and negatively correlated with the included variable, the risk component, the coefficient is severely downward biased and can even be driven negative. In addition, the conditional variance is often measured poorly, thus generating large standard errors and parameter estimates that can vary substantially across specifications. Finally, controlling for the volatility feedback effect, i.e., the effect of the shock to the risk component on unexpected returns, increases the efficiency of our estimation, sometimes substantially. Interestingly, innovations in the hedge component do not seem to help in identifying the model and actually degrade its performance.

The remainder of the paper is organized as follows. Section 2 presents a log-linear model of stock returns that decomposes ex post returns. The data are discussed in Section 3, and we also provide an initial examination of conditional volatility. The main empirical investigation is conducted in Section 4. Section 5 concludes the paper.

2 Theory

2.1 A Log-Linear Asset Pricing Model

In this section, we derive an asset pricing model based on Merton's ICAPM and Campbell and Shiller's (1988) log-linearization method. The log-linear approximation provides both tractability and accuracy.

As in Campbell and Shiller (1988), the continuously compounded market return $r_{M,t+1}$ is defined as

$$r_{M,t+1} = \log(P_{M,t+1} + D_{M,t+1}) - \log(P_{M,t}), \quad (2)$$

where $P_{M,t+1}$ is the price at the end of period $t+1$ and $D_{M,t+1}$ is the dividend paid out during period $t+1$. Throughout this paper, we use upper case to denote the level and lower case to denote the log. In addition, the subscript M will be suppressed for notational convenience.

Using a first-order Taylor expansion around the steady state of the log dividend price ratio $\overline{d-p}$, equation (2) can be rewritten as a first-order difference equation for the stock price,

$$r_{t+1} \approx k + \rho p_{t+1} - p_t + (1 - \rho)d_{t+1}, \quad (3)$$

where

$$\begin{aligned} \rho &= \frac{1}{1 + \exp(\overline{d-p})}, \\ k &= -\log(\rho) - (1 - \rho) \log\left(\frac{1}{\rho} - 1\right), \end{aligned}$$

and ρ is set to be 0.997 as in Campbell, Lo and MacKinlay (1997, Chapter 7).

Solving equation (3) forward and imposing the transversality condition

$$\lim_{j \rightarrow \infty} \rho^j p_{t+j} = 0$$

the stock price becomes a function of future dividend flows and discount rates,

$$p_t = \frac{k}{1 - \rho} + \sum_{j=0}^{\infty} \rho^j [(1 - \rho)d_{t+1+j} - r_{t+1+j}]. \quad (4)$$

Equation (4) is simply an accounting identity, which also holds ex ante,

$$p_t = \frac{k}{1 - \rho} + E_t \sum_{j=0}^{\infty} \rho^j [(1 - \rho)d_{t+1+j} - r_{t+1+j}]. \quad (5)$$

Substituting equation (5) into equation (3), we decompose the ex post stock return into two parts, the expected return and the shocks to the return,

$$r_{t+1} = E_t r_{t+1} - \left[E_{t+1} \sum_{j=1}^{\infty} \rho^j r_{t+1+j} - E_t \sum_{j=1}^{\infty} \rho^j r_{t+1+j} \right] + \left[E_{t+1} \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - E_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} \right], \quad (6)$$

where Δd_{t+1+j} is dividend growth. Unexpected returns are themselves made up of two components, revisions in future expected returns and revisions in cash flow forecasts. For the excess market return, $e_{t+1} \equiv r_{t+1} - r_{f,t+1}$, where $r_{f,t+1}$ is the nominal risk-free rate, equation (6) can be rewritten as

$$e_{t+1} = E_t e_{t+1} - \left[E_{t+1} \sum_{j=1}^{\infty} \rho^j e_{t+1+j} - E_t \sum_{j=1}^{\infty} \rho^j e_{t+1+j} \right] - \left[E_{t+1} \sum_{j=1}^{\infty} \rho^j r_{f,t+1+j} - E_t \sum_{j=1}^{\infty} \rho^j r_{f,t+1+j} \right] + \left[E_{t+1} \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - E_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} \right]. \quad (7)$$

Merton's ICAPM (equation (1)) provides the model for expected excess returns

$$E_t e_{t+1} = \gamma \sigma_t^2 + \lambda_t \sigma_{MF,t}, \quad (8)$$

where $\frac{-J_{WWW}}{J_W} = \gamma$ and $\frac{-J_{WE}}{J_W} = \lambda_t$. γ is the constant relative risk aversion coefficient, and λ_t is a function of the state variables and is not necessarily constant over time. Substituting equation (8) into equation (7) and noting that $E_t e_{t+1+j} = E_t [E_{t+j} e_{t+1+j}]$ by iterated expectations, we get

$$e_{t+1} = \gamma \sigma_t^2 + \lambda_t \sigma_{MF,t} - \eta_{\sigma,t+1} - \eta_{F,t+1} - \eta_{f,t+1} + \eta_{d,t+1}, \quad (9)$$

where

$$\begin{aligned} \eta_{\sigma,t+1} &= E_{t+1} \sum_{j=1}^{\infty} \rho^j \gamma \sigma_{t+j}^2 - E_t \sum_{j=1}^{\infty} \rho^j \gamma \sigma_{t+j}^2, \\ \eta_{F,t+1} &= E_{t+1} \sum_{j=1}^{\infty} \rho^j \lambda_{t+j} \sigma_{MF,t+j} - E_t \sum_{j=1}^{\infty} \rho^j \lambda_{t+j} \sigma_{MF,t+j}, \\ \eta_{f,t+1} &= E_{t+1} \sum_{j=1}^{\infty} \rho^j r_{f,t+1+j} - E_t \sum_{j=1}^{\infty} \rho^j r_{f,t+1+j}, \\ \eta_{d,t+1} &= E_{t+1} \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - E_t \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j}. \end{aligned}$$

The first two terms in equation (9) capture the expected excess return. The third and fourth terms explicitly write out the unexpected return due to shocks to the risk component and hedge component of expected returns, respectively. The remaining terms are shocks to risk-free rate forecasts and cash flow forecasts.

2.2 Modeling the Risk and Hedge Components of Returns

The empirical implementation of equation (9) requires further specification of the risk and hedge components of returns. By imposing a specific time series model on these components, we can also reduce the shocks to these components, which are written in equation (9) as infinite sums, to more manageable closed-form terms.

First, consider the risk component of expected returns and the shock to this component. For periods during which we have implied volatility data, we assume that this variable is a reasonable observable proxy for the conditional variance. Thus, no further specification of the risk component is necessary. In order to simplify the expression for the shock to the risk component, we need to assume a time series process for the conditional variance. Poterba and Summers (1986) argue that an AR(1) process is appropriate, and, as we will show later, an AR(1) process fits the data well. Consequently, we model the evolution of the conditional variance process as

$$\sigma_{t+1}^2 = \alpha + \beta\sigma_t^2 + \varepsilon_{\sigma,t+1}. \quad (10)$$

This equation implies

$$\eta_{\sigma,t+1} = \frac{\rho\gamma}{1 - \rho\beta}\varepsilon_{\sigma,t+1} \quad (11)$$

(see the Appendix for details). Note that the unexpected return due to revisions to the risk component is just a simple linear function of the shock to the conditional variance. This term generates the volatility feedback effect in equation (9), i.e., returns are negatively related to contemporaneous innovations in the conditional variance.

For the longer sample period when implied volatility data is not available, we need a model for the conditional variance. We assume that the realized market variance, v_t^2 , is a linear function of its own lag and a set of state variables, X_{kt} , $k = 1, \dots, K$, i.e.,

$$v_t^2 = a_0 + a_1v_{t-1}^2 + \sum_{k=1}^K a_{2,k}X_{k,t-1} + \zeta_t \quad (12)$$

(13)

Discussion of the computation of the realized variance and the choice of state variables is postponed until Section 3.2. The fitted value from the estimation is used as a proxy for the conditional market variance,⁶ i.e.,

$$\begin{aligned}\hat{\sigma}_{t-1}^2 &= a_0 + a_1 v_{t-1}^2 + \sum_{k=1}^K a_{2,k} X_{k,t-1} \\ &= \omega_0 + \omega_1 Z_{t-1}\end{aligned}\tag{14}$$

where

$$Z_t = \begin{bmatrix} v_t^2 \\ X_t \end{bmatrix}$$

In order to calculate the innovation in the risk component we need to compute the shock to this conditional variance, which, in turn, requires specifying a process for the state variables. Following Campbell and Shiller (1988), among others, we assume that the state variables, X_{t+1} , follow a vector autoregressive (VAR) process with a single lag:⁷

$$X_{t+1} = A_0 + A_1 X_t + \varepsilon_{X,t+1}.\tag{15}$$

The joint process for the variance and the instrumental variables can be written as

$$Z_{t+1} = B_0 + B_1 Z_t + \varepsilon_{Z,t+1}\tag{16}$$

where

$$B_0 = \begin{bmatrix} \omega_0 \\ A_0 \end{bmatrix} \quad B_1 = \begin{bmatrix} \omega_1 \\ 0 \quad A_1 \end{bmatrix} \quad \varepsilon_{Z,t+1} = \begin{bmatrix} \zeta_t \\ \varepsilon_{X,t+1} \end{bmatrix}$$

Therefore

$$\eta_{\sigma,t+1} = \rho \gamma \omega_1 (I - \rho B_1)^{-1} \varepsilon_{Z,t+1}\tag{17}$$

⁶This type of specification has a long history in the literature. For example, French, Schwert and Stambaugh (1987) use a time series model of realized variance to model the conditional variance. Numerous papers since then have employed predetermined financial variables as additional predictors.

⁷Extending the VAR to longer lags is conceptually straightforward, but it adds nothing to the intuition from the model.

where I is a $(K+1)$ -by- $(K+1)$ identity matrix (see the Appendix for details). This result is essentially the multivariate analog of the result in equation (11); the innovation in the risk component is a function of the innovation in the variables that describe the conditional variance. Again, this term generates the well-known volatility feedback effect.

There are several ways to estimate the hedge component $(\lambda_t \sigma_{MF,t})$ in equation (9). Scruggs (1998) uses a bivariate exponential GARCH model, in which he assumes that λ_t is constant and the long-term government bond return is perfectly correlated with investment opportunities. Following Campbell (1996), we assume that the hedge component is a linear function of the same state variables, X_t , as used above to model the conditional variance,⁸ i.e.,

$$\lambda_t \sigma_{MF,t} = \phi_0 + \sum_{k=1}^K \phi_{1,k} X_{k,t} = \phi_0 + \phi_1 X_t. \quad (18)$$

Note that time-variation in λ_t is assumed to be captured by time-variation in the state variables; the coefficients on these state variables, ϕ_1 , are assumed to be constant.

This formulation needs some explanation since, in the stock return predictability literature, it is used to model total expected returns not just the component of expected returns due to hedging demands. Moreover, we are use the same variables to model both the risk and hedge components. The obvious danger is that we may mistakenly attribute part of the risk component to the hedge component, i.e., we will not be able to identify the two components separately. We avoid this problem by ensuring that our proxy for conditional volatility subsumes all the information about the risk component that is contained in the state variables in equation (18). In the case of the implied volatility data, this amounts to checking that the state variables have no marginal explanatory power for future volatility after controlling for the predictive power of implied volatility. In other words, as long as implied volatility is a sufficiently good proxy for the risk component, the state variables will only pick up the component of expected returns that is not related to risk, which is, by definition, the hedge component. Evidence to this effect is presented in Section 3. When we model conditional variance as a linear function of the state variables, the process of projecting realized variance on these variables guarantees that we have extracted all the (linear) information

⁸Using the same vector of state variables is without loss of generality since the vector can always be expanded to include all variables that are relevant for modeling either the risk or hedge components. That is, the coefficient on any particular variable may be zero in either of the equations.

about future volatility that they contain, and, again, the residual predictive power should be due only to the hedge component of expected returns.

One advantage of equation (18) is that it allows us to calculate the revision term for the hedge component, $\eta_{F,t+1}$ in equation (9), directly as in Campbell and Shiller (1988), Campbell (1991) and Campbell and Ammer (1993). By controlling for this component of returns, we can potentially increase the efficiency with of the estimation and the precision with which we estimate the coefficients. Specifically,

$$\eta_{F,t+1} = \rho\phi_1(I - \rho A_1)^{-1}\varepsilon_{X,t+1}, \quad (19)$$

where I is a K -by- K identity matrix (see the Appendix for details). As for the risk component, innovations to the hedge component are a relatively simple function of the shocks to the state variables.

After substituting equations (18), (19), and either (11) or (17), into equation (9), we obtain the two models that are estimated in this paper:

$$e_{t+1} = \gamma\sigma_t^2 + [\phi_0 + \phi_1 X_t] - \frac{\rho\gamma}{1 - \rho\beta}\varepsilon_{\sigma,t+1} - \rho\phi_1(I - \rho A_1)^{-1}\varepsilon_{X,t+1} - \eta_{f,t+1} + \eta_{d,t+1} \quad (20)$$

$$e_{t+1} = \gamma\sigma_t^2 + [\phi_0 + \phi_1 X_t] - \rho\gamma\omega_1(I - \rho B_1)^{-1}\varepsilon_{Z,t+1} - \rho\phi_1(I - \rho A_1)^{-1}\varepsilon_{X,t+1} - \eta_{f,t+1} + \eta_{d,t+1}. \quad (21)$$

They differ only by the proxies for conditional variance and the corresponding shocks to the risk component of expected returns. Both equations (20) and (21) captures the six components of excess market returns: expected returns due to the risk and hedge components, unexpected returns due to shocks to these components of expected returns, and shocks to cash flow and risk-free rate forecasts. The risk and risk revision terms are linear functions of the lagged conditional variance and the contemporaneous shock to the conditional variance, respectively. The hedge and hedge revision terms are written in terms of the lagged state variables and the shocks to the process that governs these variables, respectively. The shocks to the cash flow and risk-free rate forecasts are not written out explicitly, and they form the regression residual in the specification that we estimate.

3 Data Description

The models are estimated with two sets of data. The first dataset utilizes the volatility implied by S&P 100 index (OEX) option prices as a proxy for the conditional market variance. The second dataset adopts commonly used financial variables as instruments to estimate the conditional market variance.

3.1 Implied Volatility Data

The implied volatility data were constructed by Christensen and Prabhala (1998). They compute non-overlapping monthly implied volatility data for the S&P 100 index spanning the period November 1983 to May 1995, for a total of 139 observations. It is important to note that the S&P 100 index option contract expires on the third Saturday of each month. Christensen and Prabhala compute implied volatility based on a contract that expires in twenty-four days. The sampling month is thus different from the calendar month; moreover, some trading days are not included in any contract. For example, the implied volatility for October 1987 is calculated using the option price on September 23, 1987 for the option that expires on October 17, 1987. For November 1987, it is based on the option price on October 28, 1987 for the option that expires on November 21, 1987. Thus, trading days between October 17, 1987, and October 28, 1987, including the October 19, 1987 stock market crash, are not included in any contract. We will return to this point later.

The monthly excess market return and variance are constructed from daily excess market returns. We use daily value-weighted market returns (VWRET) from CRSP as daily market returns. The daily risk-free rate data are not directly available. Following Nelson (1991) and others, we assume that the risk-free rate is constant within each month and calculate the daily risk-free rate by dividing the monthly short-term government bill rate from Ibbotson Associates (1997) by the number of trading days in the month. The daily excess market return is the difference between the daily risk-free rate and the daily market return.

As in Christensen and Prabhala (1998), the realized monthly market variance is defined as

$$v_t^2 = \sum_{k=1}^{\tau_t} (e_{t,k} - \bar{e}_t)^2, \quad (22)$$

where τ_t is the number of days to expiration of the option contract in month t , $e_{t,k}$ is the daily

excess market return and \bar{e}_t is the mean return over the month, i.e., $\bar{e}_t = \frac{1}{\tau_t} \sum_{k=1}^{\tau_t} e_{t,k}$.⁹ The monthly excess market return is the sum of daily excess market returns,

$$e_t = \sum_{k=1}^{\tau_t} e_{t,k}. \quad (23)$$

Both the realized and implied variances of the S&P 100 index returns are larger than the realized market variance of the CRSP value-weighted portfolio because the S&P 100 is not a well-diversified portfolio. In our sample, the means of monthly realized and implied variances of the S&P 100 return are 0.0022 and 0.0020, respectively. The realized market variance of the value-weighted portfolio has a mean of 0.0014. Consequently, we scale the implied variance by 14/22, which equates the average realized variances of the two series. This rescaling has no effect on the statistical significance or explanatory power of the models estimated subsequently, since the implied variance enters the equations linearly. The point of the rescaling is simply to preserve the economic interpretation of the coefficient as the coefficient of relative risk aversion. Without rescaling, the higher variance of the S&P 100 would reduce this coefficient by exactly 14/22.

As mentioned above, the October 19, 1987 market crash and the following days are not included in any sampling month and the market variance returns to a normal level very quickly; thus, the crash has a small impact on the realized market variance in our sample.¹⁰ However, the implied market variance does jump after the market crash, although it also returns to a normal level very quickly.

The hedge component is estimated using two instrumental variables:¹¹ (1) the consumption-wealth ratio (CAY) (see Lettau and Ludvigson (2001)), and (2) the stochastically detrended risk-free

⁹Adjusting the variance for the realized mean daily return has no appreciable affect on the results. In other words, realized variance could also be computed as the sum of squared returns as is done in some of the literature.

¹⁰There may be other reasons to exclude the October 19, 1987 market crash from the sample. Schwert (1990b) shows it is unusual in many ways, and Seyhun (1990) argues that it is not explained by the fundamentals. It is not predicted by the option data used in this paper. We will return to this point in discussions of the monthly data.

¹¹An earlier version of the paper used a somewhat different set of four instrumental variables: (1) the yield spread between Baa-rated and Aaa-rated bonds, (2) the yield spread between 6-month commercial paper and 3-month Treasury bills, (3) the stochastically detrended risk-free rate, and (4) the dividend yield. The results are qualitatively similar, and the new specification is more parsimonious and better reflects the evolution of the literature on return predictability.

rate (RREL). The latter variable is defined as

$$\text{RREL}_t = r_{f,t} - \frac{1}{12} \sum_{k=1}^{12} r_{f,t-k}. \quad (24)$$

The risk-free rate is taken from Ibbotson Associates (1997), and the consumption-wealth ratio is computed and supplied by Martin Lettau.¹²

It is well known that we can predict stock market volatility with variables such as the nominal risk-free rate, the consumption-wealth ratio and lagged realized variance (see, for example, Campbell (1987), French, Schwert and Stambaugh (1987) and Lettau and Ludvigson (2003)), and we consider the same financial variables as used for estimating the hedge component.¹³ To test the information content of implied volatility, we regress the realized variance on its own lag, the instrumental variables and the scaled implied variance, V_t^2 , i.e.,

$$v_t^2 = a_0 + \sum_{k=1}^2 a_{1,k} X_{k,t-1} + a_2 v_{t-1}^2 + a_3 V_{t-1}^2 + \zeta_t. \quad (25)$$

If implied volatility is efficient, all the other variables should enter equation (25) insignificantly. We estimate equation (25), and restricted versions thereof, with GMM, and the parameter estimates and heteroscedasticity consistent standard errors are reported in Table 1.

We first exclude implied volatility in order to verify the predictive power of the other variables in our sample period. Both the consumption-wealth ratio and lagged variance enter significantly, with the signs of the coefficients consistent with the existing literature, and the explanatory power is substantial, with an R^2 close to 40%. The risk-free rate does not have any marginal explanatory power, but this may be specific to our sample period. When the implied variance is added to the specification, it is highly significant. Past variance and the consumption-wealth ratio are no longer significant at the 5% level,¹⁴ and the R^2 increases to almost 50%. Consequently, we conclude that implied volatility is an efficient instrumental variable for future stock market variance, and, as such,

¹²See his web site, <http://pages.stern.nyu.edu/~mlettau/>, for details.

¹³Other variables such as the commercial paper–Treasury spread have a small amount of marginal predictive power over and above CAY and RREL in our sample. However, the results are invariant to including these additional variables and we restrict the set of variables for parsimony.

¹⁴CAY does have some marginal explanatory power; it is significant at the 10% level. Note that throughout the paper significance levels are based on two-sided tests. To the extent that one-sided tests are appropriate in some cases, the relevant significance level is half of that reported.

it may well improve the estimation efficiency of the model. Of equal importance, this result implies that we will be able to separately identify the two components of expected returns. Specifically, the instrumental variables used for the hedge component will pick up little if any of the risk component because they have limited marginal explanatory power for future variance, after controlling for the predictive power of implied volatility.

In Table 1, we also report estimates from a regression of realized variance on the scaled implied volatility alone. If scaled implied volatility is a conditionally unbiased predictor of future variance, then the intercept in this regression should be equal to zero and the coefficient on implied volatility should be equal to one. However, if implied volatility is measured with some error (e.g., due to the failure of the Black-Scholes model or because the S&P 100 differs in an economically significant way from the value-weighted CRSP index), then the coefficient will be biased downwards and the intercept should be positive, even if the measurement error is mean zero. While the estimated coefficient is positive, it is significantly less than one, and the intercept is significantly positive, although it is small. Thus, while implied volatility may be informationally efficient relative to other variables it is not conditionally unbiased. As a result, it is likely that the coefficient on this variable will have a small downward bias in the later estimation of the model. Of some interest, there is a noticeable drop in the R^2 , from 49% to 41%, when the other three variables are dropped. However, including these additional variables in the model for conditional variance has no meaningful effect on the later estimation of the full model; therefore, for ease of exposition we ignore them.

3.2 Instrumental Variables Estimation of the Conditional Variance

The model is also estimated with postwar monthly data, spanning the period January 1959 to December 2000. The monthly variance of stock market returns is again computed from daily return data as in equation (22). In this case, however, we use calendar months. We use the daily market return data constructed by Schwert (1990a) before July 1962 and the daily value-weighted market return from CRSP thereafter. The risk-free rate is also from CRSP.

We assume that the conditional market variance is a linear function of its own lags and the same state variables: CAY and RREL. Including two lags, we estimate the regression

$$v_t^2 = a_0 + \sum_{i=1}^2 a_{1,i} v_{t-i}^2 + \sum_{k=1}^2 a_{2,k} X_{k,t-1} + \zeta_t. \quad (26)$$

Realized stock market variance shoots up to 0.0507 in October 1987 and returns to a more normal level soon thereafter. The crash has a confounding effect on the estimation of equation (26), which is reported in Table 2. Although only one lag of the market variance is statistically significant in both pre-crash (1/59-9/87) and post-crash (1/88-12/00) sub-samples (the first two regressions reported in Table 2, respectively), the second lag of market variance is also statistically significant in the full sample (the third regression). Not surprisingly, the R^2 of this third regression is also much lower (7% versus 29% and 42% in the subsamples), and the sum of the coefficients on the lagged variance terms is also much lower (0.27 versus 0.57 and 0.46). Basically, including the crash significantly degrades the predictive power of the regression over all the other months because this one observation dominates the sample in a OLS context. To reduce the impact of the October 1987 market crash, we arbitrarily set the realized stock market variance of October 1987 to 0.0094 basis points, the second largest realization in our sample.¹⁵ The corresponding results are shown in the fourth regression in Table 2. The coefficients and explanatory power look very similar across the subperiods and the full sample after this adjustment. The second lag of realized variance is statistically significant (at the 5% level) in the full sample. Nevertheless, this second lag adds little to the explanatory power, as demonstrated by the fifth regression, which excludes this term. The R^2 drops only 1%, from 33% to 32%. Of some interest, the consumption-wealth ratio is a significant predictor of future variance in all the regressions, entering with a negative coefficient. Consequently, our final specification for the conditional variance has the two state variables and a single lag of the realized variance, where the October 1987 variance is adjusted as described above.

Finally, we also use the same set two instrumental variables to model the hedge component of expected returns. Again, these variables should not pick up any of the risk component. The best linear combination is already included in the risk proxy, so they should have no marginal explanatory power for conditional volatility.

¹⁵A similar result could be achieved by dropping this observation from the sample.

4 Empirical Results

4.1 Econometric Strategy

For the sample for which we have implied volatility data, we simultaneously estimate equations (10), (15) and (20) using GMM, substituting the scaled implied variance for the conditional variance:

$$V_{t+1}^2 = \alpha + \beta V_t^2 + \varepsilon_{V,t+1} \quad (27)$$

$$X_{t+1} = A_0 + A_1 X_t + \varepsilon_{X,t+1} \quad (28)$$

$$e_{t+1} = \gamma V_t^2 + [\phi_0 + \phi_1 X_t] - \frac{\rho\gamma}{1 - \rho\beta} \varepsilon_{V,t+1} - \rho\phi_1 (I - \rho A_1)^{-1} \varepsilon_{X,t+1} + \epsilon_{t+1}. \quad (29)$$

For equations (27) and (28), we use the standard OLS moment conditions. The only subtlety is in formulating the moment conditions for equation (29). Note that theory does not imply that the terms $\varepsilon_{V,t+1}$ and $\varepsilon_{X,t+1}$ are orthogonal to the contemporaneous regression error ϵ_{t+1} ; therefore, these variables should not be used as instruments. Using a constant, V_t^2 and $X_t = \{CAY_t, RREL_t\}$ as instruments is sufficient to identify the free parameters. The two shocks are the fitted residuals from the first two equations, the parameters β and A_1 also come from these equations, and ρ is set to 0.997 (see Section 2). What then is the value of including these two terms at all? First, by reducing the amount of unexplained variation, they should improve the efficiency of the estimation and the accuracy with which the parameters of interest can be estimated. Second, these terms also depend on the parameters, and thus imposing the theoretical restrictions implied by the model also helps to pin down these parameters.

In order to understand what is driving our results relative to the existing literature and to understand the gains from imposing the additional restrictions on the revision terms, we also estimate restricted versions of the model that exclude various terms in equation (29). Specifically, we consider the following 5 cases:

$$1: \quad e_{t+1} = \phi_0 + \gamma V_t^2 + \epsilon_{t+1} \quad (30)$$

$$2: \quad e_{t+1} = \phi_0 + \gamma V_t^2 - \frac{\rho\gamma}{1 - \rho\beta} \varepsilon_{V,t+1} + \epsilon_{t+1} \quad (31)$$

$$3: \quad e_{t+1} = \phi_0 + \phi_1 X_t + \epsilon_{t+1} \quad (32)$$

$$4: \quad e_{t+1} = \gamma V_t^2 + [\phi_0 + \phi_1 X_t] + \epsilon_{t+1} \quad (33)$$

$$5: \quad e_{t+1} = \gamma V_t^2 + [\phi_0 + \phi_1 X_t] - \frac{\rho\gamma}{1 - \rho\beta} \varepsilon_{V,t+1} + \epsilon_{t+1} \quad (34)$$

Model 6 is the full model in equation (29). In each case, we use the set of instruments corresponding to the independent variables in the GMM estimation, i.e., one or more instruments from the set $[V_t^2, 1, X_t = \{CAY_t, RREL_t\}]$. For example, in Model 1, there two parameters to be estimated (ϕ_0, γ) and we use two instruments (V_t^2 and a constant). Model 2 is similar, except that the inclusion of the revision term should help to identify γ . In Model 3 we use a constant and the variables CAY and RREL as instruments, and the remaining models use the full set of instruments. Thus, all the models are exactly identified.

For the longer sample period, for which we do not have implied volatility data, we replace equation (27) with (26) with a single lag in order to estimate the conditional variance. The fitted value from this equation is used instead of the scaled implied variance in equation (29) and the revision to the risk term is defined in the corresponding way. The resulting system of equations is

$$v_t^2 = c + a_1 v_{t-1}^2 + \sum_{j=1}^2 b_j X_{j,t-1} + \zeta_t \quad (35)$$

$$X_{t+1} = A_0 + A_1 X_t + \varepsilon_{X,t+1} \quad (36)$$

$$e_{t+1} = \gamma \hat{\sigma}_t^2 + [\phi_0 + \phi_1 X_t] - \rho \gamma \omega_1 (I - \rho B_1)^{-1} \varepsilon_{Z,t+1} - \rho \phi_1 (I - \rho A_1)^{-1} \varepsilon_{X,t+1} + \epsilon_{t+1}, \quad (37)$$

where ω_1, B_1 and $\varepsilon_{Z,t+1}$ are defined in equations (14)-(16). Again all the equations are estimated simultaneously using GMM and restricted versions of equation (37) are also estimated as in models 1-5 above.

4.2 Implied Volatility Data

Table 3 reports results from the estimation of equations (27)-(29) using monthly data for the December 1983 to April 1995 period. The results for the conditional variance process, estimated using the implied volatility data, are shown in Panel A. The AR(1) coefficient is positive and significant, as expected, although the estimated persistence, and hence the R^2 , are not very large. Panel B shows the parameter estimates for the VAR(1) process for the state variables that are used to model the hedge component. Both the consumption-wealth ratio and the relative T-bill rate are quite persistent, but neither variable shows much predictive power for its counterpart.

The results from the estimation of the model itself are reported in Panel C. Recall that we estimate six different specifications—five restricted models given in equations (30)-(34) and the full specification given in equation (29). Model 1 is the standard risk-return model estimated in much

of the literature, i.e., a regression of returns on a measure of the conditional variance. However, in contrast to many existing results, we find a coefficient that is positive, albeit significant at only the 10% level, and reasonable in magnitude.¹⁶ If the hedge component is unimportant, the coefficient value of 2.2 represents an estimate of the coefficient of relative risk aversion of the representative agent; although, this estimate may be biased downwards slightly due to measurement error in the conditional variance, evidence of which is presented in Section 3.1. The absence of a hedge component also implies that the constant in the regression should be zero—a hypothesis that is rejected at the 10% significance level. Moreover, the R^2 of less than 1% is very small. Model 2 attempts to refine the specification by controlling for the effect of shocks to the risk component on unexpected returns, i.e., the volatility feedback effect. Adding this term leaves the parameters estimates essentially unchanged. Nevertheless, the estimation does provide corroborating evidence in that the R^2 jumps to 7.6%. In other words, shocks to the risk component explain 10 times more of the variation in returns than variation in the risk component of expected returns itself.

These initial results suggest two conclusions. First, sample period issues aside, improving the quality of the proxy for conditional variance, in our case using implied volatility, seems to help in recovering the theoretically justified positive risk-return relation. Ghysels, Santa-Clara and Valkanov (2003) present evidence consistent with this conclusion using functions of daily squared returns data to form a better measure of conditional variance. Second, while there is some evidence of a positive risk-return relation, statistical power is weak, and the model can be rejected. Thus, controlling for the hedge component of expected returns may be important.

Model 3 estimates the standard return predictability regression from the literature using our two state variables. In our case, however, we interpret this regression as an estimation of the hedge component of expected returns without controlling for the risk component. The signs of the coefficients, positive on the consumption-wealth ratio and negative on the relative T-bill rate, are consistent with the results in the literature, as is the R^2 of just over 3%. However, perhaps due to the short sample period, neither variable is individually significant. Nevertheless, the predictive

¹⁶The standard errors are computed via GMM and are asymptotic. Given the sample size, there is some question as to whether asymptotic standard errors are appropriate. This question is impossible to answer definitively, but we also compute small sample standard errors using a block bootstrap methodology. The resulting standard errors on the state variables are actually slightly lower than those reported in Table 3, but the standard errors on γ are somewhat higher, e.g., 2.5 versus 1.5 in model 4.

power is still substantially greater than that found for the risk component in model 1.

Under the ICAPM, both models 1 and 3 are misspecified since theoretically both the risk and hedge components should enter the model for expected returns. Model 4 combines these two terms, and the results are extremely positive. First and foremost, estimated risk aversion is now 4.4 and it is significant at the 1% level. The risk-return relation is extremely statistically significant and of a reasonable magnitude. The reversal of the weak and/or negative results in the literature is dramatic, and the explanation is clear. Model 1, and more generally similar specifications in the literature, suffer from a classical omitted variables problem, i.e., they do not control for the hedge component of expected returns. The effect of an omitted variable on the estimated coefficient of the included variable depends on the covariance of this variable with the included variable. In this case, if the covariance is negative, i.e., the risk and hedge components are negatively correlated, then the coefficient on conditional variance when this term is included alone, will be biased downwards. Second, including the risk component also helps in identifying the hedge component; the coefficient on the consumption-wealth ratio is now more significant. Third, the joint explanatory power of the two components exceeds the sum of the individual explained variations from the separate regressions; the R^2 increases to 5.2%. Finally, the results give us added confidence that we are correctly identifying the risk and hedge components. The hedge component is positively related to CAY, while the variance, if anything, is negatively related to CAY (see Table 1).

Can our identification of these components be improved by controlling for the effects of shocks to expected returns on contemporaneous unexpected returns? This question is addressed by models 5 and 6. In model 5, we add the shock to the risk component of expected returns. The results are not dramatically different from those of model 4, but a couple of observations are worth making. First, including the shock to the conditional variance can effect both the estimate of relative risk aversion and the hedge component. In this case, γ drops from 4.4 to 4.0, and the coefficients on both the state variables are closer to zero. Second, controlling for some of the variation of unexpected returns can increase the efficiency of the estimation. In this case, the R^2 increases from 5.0% to 11.5%, and the standard errors on the coefficients drop by between 3% and 18%.

Finally, model 6 also includes the shock to the hedge component in the regression, with some surprising results. First, the estimate of risk aversion rises substantially to 5.9, but the standard error is also substantially larger, and the coefficient is not significant. In fact, the standard errors

increases for all the coefficients, although most dramatically for γ . Second, the R^2 actually declines with the addition of this term; the shock to the hedge component causes a deterioration in the performance of the model. There are a number of possible explanations for these unexpected results. They could be indicative of the fact that the apparent predictive power of the state variables is spurious. If these variables do not have true predictive ability for expected returns then shocks to these variables will not explain unexpected returns either. However, this explanation is contradicted by the fact that inclusion of the state variables helps in identifying the risk component of expected returns in models 4 and 5. If their predictive power is spurious, this identification effect should not occur. A more likely explanation, perhaps, is that the model suffers from some form of misspecification. If the VAR(1) specification for the state variables is inadequate or the functional form of the hedge component is not linear in the state variables, estimation error in the shock to the hedge component could cause the reported results. A complete analysis of this issue is beyond the scope of the current paper; however, these intriguing results certainly warrant further investigation.

In order to better illustrate the omitted variables problem and to gain some economic understanding of the results, we construct the fitted risk and hedge components of expected returns using model 5. The two series are plotted in Figure 1 along with the NBER business cycle peaks and troughs (shaded bars represent recessions, i.e., the period between the peak of the cycle and the subsequent trough). Clearly, the hedge component is much more variable than the risk component of expected returns, with a sample variance that is almost twice as large. Moreover, much of the variance of the risk component is attributable to the spike in implied volatility following the crash in October 1987. The two series are negatively correlated, with a sample correlation of -0.33; thus, omitting the hedge component causes the coefficient on the risk component to be biased downwards. The magnitude of this bias depends on the covariance between the hedge component and the included variable (i.e., the conditional variance) relative to the variance of the included variable times its true coefficient (i.e., the true coefficient of relative risk aversion). In this sample, the covariance is -3.64 while the product of the estimated value of γ (from model 5) times the variance of the implied variance is 8.25. The bias is not sufficient to reverse the sign of the estimated coefficient in models 1 and 2, but it is substantial.

From an economic standpoint, the hedge component appears to exhibit some countercyclical variation (i.e., it increases over the course of recessions), but there is only a single recession in the

sample so this interpretation is extremely casual. The risk component exhibits little or no apparent business cycle patterns, although, as noted above, variation in this series is dominated by the crash and the subsequent increase in implied volatility. Of some interest, the hedge component is negative for substantial periods of time. There is no theoretical restriction on the sign of this component of expected returns, but it does imply that at these times the stock market serves as a hedge against adverse shifts in investment opportunities.

4.3 Instrumental Variables Estimation

The results of Section 4.2 go a long way in resurrecting the positive risk-return relation, but the analysis suffers from two problems: (i) the sample period is short, and (ii) it relies on implied volatility data that are not available in all periods or across all markets. Consequently, we now turn to an analysis that constructs conditional variance estimates from ex post variance computed using daily returns and conditioning variables that include lagged realized variance and our two state variables, the consumption wealth ratio and the relative T-bill rate.

Table 4 reports results for the estimation of the system in equations (35)-(37) using monthly data over the period February 1959 to December 2000. The estimation of the variance process is reported in Panel A, which is identical to the fifth regression of Table 2. Recall that realized variance in October 1987 is adjusted downwards as discussed in Section 3.2. Other than the substantial degree of explained variation, the key result is that the conditional variance is negatively and significantly related to the consumption-wealth ratio. Panel B reports the estimation of the VAR(1) on the two state variables. The results are comparable, although not identical, to those estimated over the shorter sample period in Table 3. Again, both variables exhibit strong persistence and cross-variable predictability is limited.

Panel C reports the major results of interest, i.e., the estimates from the 6 models in Section 4.1 and also estimated using the implied volatility data in Section 4.2. Models 1 and 2 contain only the risk component (the conditional variance) plus, in the latter case, the innovation in this component of expected returns. In the longer sample, the effects of the omitted variable, i.e., the hedge component, and measurement issues in the conditional variance are clear. The coefficient on the conditional variance is negative and the standard error is large. The negative coefficient is consistent with previous studies that have documented a negative risk-return relation over similar

sample periods (e.g., Whitelaw (1994)). The fact that the standard error is more than twice that in the shorter sample period (which uses implied volatility) is testament to the value of finding better proxies for the conditional variance.

When both components are estimated together, as in model 4, the anomalous negative coefficient on the conditional variance disappears. The estimate for γ is positive and reasonable in magnitude, but the standard error is still high. Moreover, the estimated risk component of expected returns is small; the R^2 of 4.4% in model 4 is no higher than that of model 3 which excludes the risk component. In contrast, the hedge component appears to be identified well, with the coefficients on both state variables significant at the 1% level. It is interesting to note that the hedge component, in contrast to the conditional variance, is positively related to the consumption-wealth ratio. Thus, we get the negative covariance between the hedge and risk components that exacerbates the omitted variable problem in models 1 and 2 and also aids in separate identification of the two components of expected returns.

Adding shocks to expected returns in models 5 and 6 has similar effects as documented for the shorter sample period in Table 3. Specifically, the shock to the risk component increases estimation efficiency and has a small effect on the estimated parameter values. Again, however, the innovation in the hedge component causes a deterioration in the results, at least with regard to the explanatory power of the model.

Figure 2 plots the fitted risk and hedge components (from model 5) over the longer sample period, with the NBER business cycle peaks and troughs marked as in Figure 1. Several observations are in order. First, as in Figure 1, the hedge component is more variable than the risk component of expected returns. In this case, however, the variances differ by a factor of approximately 55. This difference is partly attributable to the fact that we have adjusted the realized variance in October 1987 down as described in Section 3.2.¹⁷ Other factors include the dependence of conditional variance on the state variables, in addition to lagged realized variance, and the lower estimate of γ . Second, the correlation between the risk and hedge components of -0.27 is similar to that documented in Section 4.2. In this case, however, the relevant covariance of -3.11 exceeds γ

¹⁷The crash of October 1987, with its dramatic effect on both realized and implied volatilities, will inevitably have a disproportionate influence on the estimated risk component of expected returns. Nevertheless, our main results seem to be robust to the treatment of this unusual episode.

times the variance of the conditional variance, which equals 1.55; thus, the omitted variables bias drives the estimate of relative risk aversion negative in models 1 and 2.

From an economic standpoint, the countercyclical variation in the hedge component is now clear. The hedge component reaches its lowest values at business cycle peaks and its highest values at business cycle troughs. This result is not terribly surprising given the dependence of the hedge component on the consumption-wealth ratio, which is known to be a business cycle variable. Nevertheless, our decomposition allows us to give this variation a clear interpretation as variation in the ability of the stock market to hedge shifts in investment opportunities. Specifically, stocks appear to provide this hedge at the peak of the cycle, but not at the trough, when investors require compensation for holding an asset that covaries positively with investment opportunities. As before, the hedge component can take on a negative value for extended periods.

There are two important conclusions to be drawn from an analysis of the results in Tables 3 and 4. First, correcting for the omission of the hedge component in the model of expected returns restores the positive partial risk-return relation that has been so difficult to find in the literature. This omitted variables problem is especially severe because the hedge component is negatively correlated with the risk component of expected returns and is much more volatile. Second, superior proxies for conditional variance, such as implied volatility from option prices, also make identification of this relation much easier.

To complete the analysis we perform two final robustness checks. First, the results in Tables 3 and 4 are not directly comparable since they cover different sample periods. We reestimate the models in Table 4 using the sample period in Table 3 and report the results in Table 5. Second, the results should hold at all horizons, not just one month. Moreover, the predictive ability of variables such as the consumption-wealth ratio for returns tends to increase at longer horizons (Lettau and Ludvigson (2001)). Therefore, we reestimate the models in Table 4 using quarterly data over the sample period 1952Q3-2002Q3 and report the results in Table 6.

The Table 5 results confirm those of Table 4. Of course, Panel B is identical to that in Table 3, as is model 3. Of primary interest, the estimate of γ is negative over the shorter sample period when only the risk component is included, and the sign is reversed when controlling for the hedge component. Again, the estimate of relative risk aversion is reasonable in magnitude but statistically insignificant at conventional levels. As before, the innovation in the hedge component does not

improve the fit of the model. In other words, the qualitative nature of the results in Table 4 are preserved in the subsample covered by the implied volatility data in Table 3. Thus, we have confidence that the evidence in Table 3 is not sample specific.

The results at the quarterly horizon, reported in Table 6, also confirm those discussed earlier. The variance process in Panel A and the process for the state variables in Panel B exhibit the same features. More important, models 1 and 2 show that the standard risk-return regression generates a negative coefficient, as it does at the monthly frequency. However, models 3-6 do look somewhat different. First, the R^2 of the hedge component is higher (12.1% in model 3), with the majority of the additional explanatory power coming from an increase in the coefficient on the consumption-wealth ratio (from 0.46 in Table 4 to 1.77 in Table 6). Second, and more important, putting the risk and hedge components together, as in model 4, has a dramatic effect on both terms. That is, the effect of omitting either term is even more dramatic than in the monthly data. γ increases to 14.3, the coefficient on CAY more than doubles, and the R^2 climbs to 17.6%. The ability of the state variables in the hedge component to help identify the risk component is further evidence that their predictive power is not spurious. Third, the inclusion of the innovation of the risk component has a much larger effect. Moving from model 4 to model 5 results in γ dropping to a more economically reasonable value of 7.9, and the standard errors falling by about 40% on the variables of interest. Thus, controlling for innovations in expected returns and the resulting effect on unexpected returns can have an important impact on the inference drawn from the estimation. Finally, as before, adding the innovation in the hedge component appears to degrade estimation efficiency.

5 Conclusion

In this paper, we estimate a variant of Merton's (1973) intertemporal capital asset pricing model, and we find a positive relationship between stock market risk and return. The estimated coefficient of relative risk aversion is reasonable in magnitude (between 1.5 and 7.9 depending on the sample period and frequency); therefore, the power utility function appears to describe the data fairly well. The conflicting results found in previous studies are due, in large part, to the fact that they do not adequately distinguish the risk component of expected returns from the hedge component. Specifically, omitting the hedge component from the estimation causes a large downward bias in

the estimate of relative risk aversion due to the negative correlation between these series.

Although stock market volatility is positively priced, in most cases it only explains a small fraction of return variation. Expected returns are driven primarily by changes in the ability of the stock market to hedge shifts in investment opportunities. Most existing economic theories can explain neither why the investment opportunity set moves so dramatically nor the macroeconomic forces behind this variation. Some recent research tries to fill this gap. For example, Campbell and Cochrane (1999) address changing investment opportunities in a habit formation model. In their model, when consumption approaches the habit level, the agent becomes extremely risk averse and demands a large expected return. Guo (2003) studies an infinite horizon heterogeneous agent model in which only one type of agent holds stocks. If there are borrowing constraints and idiosyncratic labor income shocks, shareholders require a large equity premium when their borrowing constraints are close to binding. The investment opportunities are therefore determined by shareholders' liquidity conditions.¹⁸ In contrast, Whitelaw (2000) generates large changes in investment opportunities by modeling the underlying economy as a two-regime process. Because regimes are persistent, regime shifts represent large movements in investment opportunities with corresponding changes in required returns.

The focus of this paper is on understanding risk and expected returns at the market level in a time series context; however, a significant piece of the empirical asset pricing literature focuses on the cross-section of expected returns across individual securities or portfolios. Interestingly, the importance of hedging changes in the investment opportunity set at the aggregate level is also likely to have strong implications in the cross-section. In particular, if volatility is not the primary source of priced risk at the market level, then the dynamic CAPM will not hold and market betas will not be the correct proxies for expected returns in the cross-section. Clearly, this issue warrants further investigation from both an empirical and theoretical standpoint.

¹⁸Aiyagari and Gertler (1998) and Allen and Gale (1994) emphasize the liquidity effect on stock market volatility.

Appendix

Derivation of Equation (11)

The revision of the risk component is defined as:

$$\eta_{\sigma,t+1} = E_{t+1} \sum_{j=1}^{\infty} \rho^j \gamma \sigma_{t+j}^2 - E_t \sum_{j=1}^{\infty} \rho^j \gamma \sigma_{t+j}^2.$$

The model for the conditional variance from equation (10) is

$$\sigma_{t+1}^2 = \alpha + \beta \sigma_t^2 + \varepsilon_{\sigma,t+1}.$$

This equation implies that

$$E_t[\sigma_{t+j}^2] = \alpha \left(\frac{1 - \beta^j}{1 - \beta} \right) + \beta^j \sigma_t^2.$$

Thus,

$$\begin{aligned} \eta_{\sigma,t+1} &= \sum_{j=1}^{\infty} \rho^j \gamma (E_{t+1}[\sigma_{t+j}^2] - E_t[\sigma_{t+j}^2]) \\ &= \sum_{j=1}^{\infty} \rho^j \gamma \left[\alpha \left(\frac{1 - \beta^{j-1}}{1 - \beta} \right) + \beta^{j-1} \sigma_{t+1}^2 - \alpha \left(\frac{1 - \beta^j}{1 - \beta} \right) - \beta^j \sigma_t^2 \right] \\ &= \sum_{j=1}^{\infty} \rho^j \gamma \beta^{j-1} (\sigma_{t+1}^2 - \alpha - \sigma_t^2) \\ &= \frac{\rho \gamma}{1 - \rho \beta} \varepsilon_{\sigma,t+1}, \end{aligned}$$

which is equation (11).

Derivation of Equation (17)

The estimated conditional market variance (equation (14)) is

$$\hat{\sigma}_t^2 = \omega_0 + \omega_1 Z_t$$

where the state variables follow the process

$$Z_{t+1} = B_0 + B_1 Z_t + \varepsilon_{Z,t+1}$$

as in equation (16). From the derivation of equation (11) above, the risk component can be written as

$$\begin{aligned}
\eta_{\sigma,t+1} &= \sum_{j=1}^{\infty} \rho^j \gamma (E_{t+1}[\sigma_{t+j}^2] - E_t[\sigma_{t+j}^2]) \\
&= \sum_{j=1}^{\infty} \rho^j \gamma [E_{t+1}(\omega_0 + \omega_1 Z_{t+j}) - E_t(\omega_0 + \omega_1 Z_{t+j})] \\
&= \sum_{j=1}^{\infty} \rho^j \gamma \omega_1 [E_{t+1} Z_{t+j} - E_t Z_{t+j}].
\end{aligned}$$

The process for the state variables implies

$$E_t Z_{t+j} = (I - B_1^j)(I - B_1)^{-1} B_0 + B_1^j Z_t$$

Therefore

$$\begin{aligned}
\eta_{\sigma,t+1} &= \sum_{j=1}^{\infty} \rho^j \gamma \omega_1 [(I - B_1^{j-1})(I - B_1)^{-1} B_0 + B_1^{j-1} Z_{t+1} - (I - B_1^j)(I - B_1)^{-1} B_0 - B_1^j Z_t] \\
&= \sum_{j=1}^{\infty} \rho^j \gamma \omega_1 B_1^{j-1} [Z_{t+1} - B_1 Z_t - B_0] \\
&= \rho \gamma \omega_1 (I - \rho B_1)^{-1} \varepsilon_{Z,t+1}
\end{aligned}$$

which is equation (17).

Derivation of Equation (19)

The revision of the hedge component is defined as

$$\eta_{F,t+1} = E_{t+1} \sum_{j=1}^{\infty} \rho^j \lambda_{t+j} \sigma_{MF,t+j} - E_t \sum_{j=1}^{\infty} \rho^j \lambda_{t+j} \sigma_{MF,t+j}.$$

From equation (18)

$$\lambda_t \sigma_{MF,t} = \phi_0 + \phi_1 X_t.$$

Thus,

$$\begin{aligned}
\eta_{F,t+1} &= \sum_{j=1}^{\infty} \rho^j [E_{t+1}(\phi_0 + \phi_1 X_{t+j}) - E_t(\phi_0 + \phi_1 X_{t+j})] \\
&= \sum_{j=1}^{\infty} \rho^j \phi_1 [E_{t+1} X_{t+j} - E_t X_{t+j}]
\end{aligned}$$

The process for the state variables is

$$X_{t+1} = A_0 + A_1 X_t + \varepsilon_{X,t+1}.$$

(as in equation (15)), which implies

$$E_t X_{t+j} = (I - A_1^j)(I - A_1)^{-1} A_0 + A_1^j X_t$$

Therefore

$$\begin{aligned} \eta_{F,t+1} &= \sum_{j=1}^{\infty} \rho^j \phi_1 [(I - A_1^{j-1})(I - A_1)^{-1} A_0 + A_1^{j-1} X_{t+1} - (I - A_1^j)(I - A_1)^{-1} A_0 - A_1^j X_t] \\ &= \sum_{j=1}^{\infty} \rho^j \phi_1 A_1^{j-1} [X_{t+1} - A_1 X_t - A_0] \\ &= \rho \phi_1 (I - \rho A_1)^{-1} \varepsilon_{X,t+1} \end{aligned}$$

which is equation (19).

References

- [1] Aiyagari, S., and M. Gertler, 1998, Overreaction of asset prices in general equilibrium, C.V. Starr Center Working Paper, 98-25, New York University.
- [2] Allen, F., and D. Gale, 1994, Limited market participation and volatility of asset prices, *American Economic Review*, 84, 933-955.
- [3] Backus, D., and A. Gregory, 1993, Theoretical relation between risk premiums and conditional variances, *Journal of Business and Economic Statistics*, 11, 177-185.
- [4] Baillie, R., and R. DeGennaro, 1990, Stock returns and volatility, *Journal of Financial and Quantitative Analysis*, 25, 203-214.
- [5] Campbell, J., 1987, Stock returns and the term structure, *Journal of Financial Economics*, 18, 373-399.
- [6] Campbell, J., 1991, A variance decomposition for stock returns, *Economic Journal*, 101, 157-179.
- [7] Campbell, J., 1996, Understanding risk and return, *Journal of Political Economy*, 104, 298-345.
- [8] 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.
- [9] Campbell, J., and J. Cochrane, 1999, Force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy*, 107, 205-251.
- [10] Campbell, J., and L. Hentschel, 1992, No news is good news, *Journal of Financial Economics*, 31, 281-318.
- [11] Campbell, J., A. Lo and C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.
- [12] 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.

- [13] Christensen, B., and N. Prabhala, 1998, The relation between implied and realized volatility, *Journal of Financial Economics*, 50, 125-150.
- [14] Fama, E., and K. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics*, 25, 23-49.
- [15] Fleming, J., 1998, The quality of market volatility forecasts implied by S&P 100 index option prices, *Journal of Empirical Finance*, 5, 317-345.
- [16] French, K., G. Schwert, and R. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics*, 19, 3-30.
- [17] Ghysels, E., P. Santa-Clara, and R. Valkanov, 2003, There is a risk-return tradeoff after all, working paper, UCLA.
- [18] 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.
- [19] Goyal, A., and P. Santa-Clara, 2003, Idiosyncratic risk matters!, *Journal of Finance*, 58, 975-1007.
- [20] Guo, H., 2003, Limited stock market participation and asset prices in a dynamic economy, working paper, Federal Reserve Bank of St. Louis.
- [21] Ibbotson Associates, 1997, *Stocks, Bonds, Bills and Inflation 1997 Yearbook*, Ibbotson Associates, Chicago, Illinois.
- [22] Lettau, M., and S. Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815-849.
- [23] Lettau, M., and S. Ludvigson, 2003, Measuring and modeling variation in the risk-return tradeoff, working paper, New York University.
- [24] Long, J., 1974, Stock prices, inflation and the term structure of interest rates, *Journal of Financial Economics*, 1, 131-170.
- [25] Merton, R., 1973, An intertemporal capital asset pricing model, *Econometrica*, 41, 867-887.

- [26] Merton, R., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics*, 8, 323-361.
- [27] Nelson, D., 1991, Conditional heteroscedasticity in asset returns: A new approach, *Econometrica*, 59, 347-370.
- [28] Pindyck, R., 1984, Risk, inflation, and the stock market, *American Economic Review*, 74, 335-351.
- [29] Pindyck, R., 1988, Risk aversion and determinants of stock market behavior, *Review of Economics and Statistics*, 70, 183-190.
- [30] Poterba, J., and L. Summers, 1986, The persistence of volatility and stock market fluctuations, *American Economic Review*, 76, 1142-1151.
- [31] Schwert, W., 1990a, Indexes of United States stock prices from 1802 to 1987, *Journal of Business*, 63, 399-426.
- [32] Schwert, W., 1990b, Stock volatility and the crash of '87, *Review of Financial Studies*, 3, 77-102.
- [33] Scruggs, J., 1998, Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach, *Journal of Finance*, 53, 575-603.
- [34] Scruggs, J., and P. Glabadanidis, 2003, Risk premia and the dynamic covariance between stock and bond returns, *Journal of Financial and Quantitative Analysis*, 38, 295-316.
- [35] Seyhun, N., 1990, Overreaction or fundamentals: Some lessons from insiders' response to the market crash of 1987, *Journal of Finance*, 45, 1363-1388.
- [36] Whitelaw, R., 1994, Time variations and covariations in the expectation and volatility of stock market returns, *Journal of Finance*, 49, 515-541.
- [37] Whitelaw, R., 2000, Stock market risk and return: An equilibrium approach, *Review of Financial Studies*, 13, 521-547.

- [38] Wu, G., 2001, The determinants of asymmetric volatility, *Review of Financial Studies*, 14, 837-859.

	Const.	CAY _{t-1}	RREL _{t-1}	v _{t-1} ²	V _t ²	R ²
v _t ²	0.012** (0.005)	-0.028** (0.012)	0.126 (0.129)	0.530*** (0.190)		0.371
v _t ²	0.006* (0.004)	-0.015* (0.009)	-0.036 (0.074)	0.296 (0.217)	0.533*** (0.095)	0.492
v _t ²	0.036*** ^a (0.011)				0.772*** (0.114)	0.410

Table 1: The Efficiency of Implied Volatility

We estimate equation (25) over the period from December 1983 to May 1995. Heteroscedasticity consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5% and 10% levels are marked by ***, ** and *, respectively. ^a—the constant has been scaled by a factor of 100 for presentation purposes.

Subperiod		Const.	CAY _{t-1}	RREL _{t-1}	v _{t-1} ²	v _{t-2} ²	R ²
1/59-9/87	v _t ²	0.003** (0.001)	-0.007* (0.004)	-0.031 (0.052)	0.487*** (0.076)	0.085 (0.072)	0.287
1/88-12/00	v _t ²	0.015*** (0.004)	-0.037*** (0.010)	0.125 (0.145)	0.410*** (0.116)	0.046 (0.107)	0.414
1/59-12/00	v _t ²	0.012*** (0.004)	-0.028** (0.011)	-0.029 (0.057)	0.167** (0.082)	0.105** (0.046)	0.070
1/59-12/00 ^a	v _t ²	0.006*** (0.002)	-0.015*** (0.004)	-0.022 (0.048)	0.476*** (0.064)	0.118** (0.057)	0.334
1/59-12/00 ^a	v _t ²	0.006*** (0.002)	-0.015*** (0.004)	-0.032 (0.050)	0.540*** (0.057)		0.324

Table 2: Variance Predictability Regressions

We estimate equation (26) for various subsamples. Heteroscedasticity consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5% and 10% levels are marked by ***, ** and *, respectively. In the regression labeled ^a, the market variance of October 1987 is adjusted as discussed in Section 3.2.

Panel A: The Conditional Variance Process

	Const.	V_t^2	R^2
V_{t+1}^2	0.001*** (0.000)	0.403*** (0.145)	0.163

Panel B: The Process for the State Variables

	Const.	CAY _t	RREL _t	R^2
CAY _{t+1}	0.117*** (0.033)	0.700*** (0.084)	1.240* (0.639)	0.504
RREL _{t+1}	-0.002 (0.002)	0.005 (0.005)	0.845*** (0.056)	0.726

Panel C: The Model

	Const.	γ	CAY	RREL	R^2
Model 1	0.006* (0.003)	2.228* (1.270)			0.006
Model 2	0.006* (0.003)	2.228* (1.270)			0.076
Model 3	-0.188 (0.123)		0.504 (0.313)	-4.429 (3.827)	0.031
Model 4	-0.257* (0.131)	4.402*** (1.539)	0.668** (0.330)	-5.244 (3.763)	0.052
Model 5	-0.203* (0.108)	3.984*** (1.454)	0.530* (0.274)	-4.911 (3.651)	0.115
Model 6	-0.210* (0.113)	5.916 (3.866)	0.542* (0.284)	-5.145 (4.022)	0.096

Table 3: Estimation Results—Implied Volatility Data

For Panel A we estimate the AR(1) conditional variance process in equation (27) with the implied volatility data. For Panel B we estimate the VAR(1) process for the state variables in equation (28). For Panel C we estimate the model in equation (29) and various restricted versions thereof (see Section 4.1). In each case, the equations are estimated jointly via GMM over the period from December 1983 to April 1995. Heteroscedasticity consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5% and 10% levels are marked by ***, ** and *, respectively.

Panel A: The Variance Process

	Const.	CAY _t	RREL _t	v _t ²	R ²
v _{t+1} ²	0.006*** (0.002)	-0.015*** (0.004)	-0.032 (0.050)	0.540*** (0.057)	0.324

Panel B: The Process for the State Variables

	Const.	CAY _t	RREL _t	R ²
CAY _{t+1}	0.047*** (0.009)	0.877*** (0.022)	0.239 (0.340)	0.766
RREL _{t+1}	0.002*** (0.001)	-0.004** (0.002)	0.807*** (0.055)	0.670

Panel C: The Model

	Const.	γ	CAY	RREL	R ²
Model 1	0.007* (0.004)	-1.344 (3.269)			0.001
Model 2	0.007* (0.004)	-1.228 (2.980)			0.030
Model 3	-0.169*** (0.050)		0.460*** (0.130)	-5.516*** (2.065)	0.044
Model 4	-0.185*** (0.062)	1.623 (3.513)	0.496*** (0.155)	-5.415** (2.089)	0.044
Model 5	-0.184*** (0.059)	1.462 (3.167)	0.492*** (0.151)	-5.425*** (2.085)	0.069
Model 6	-0.225*** (0.069)	5.750 (3.940)	0.587*** (0.173)	-5.159** (2.027)	0.005

Table 4: Estimation Results—Instrumental Variables

For Panel A we estimate the conditional variance process in equation (35). For Panel B we estimate the VAR(1) process for the state variables in equation (36). For Panel C we estimate the model in equation (37) and various restricted versions thereof (see Section 4.1). In each case, the equations are estimated jointly via GMM over the period from February 1959 to December 2000 (503 observations). Heteroscedasticity consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5% and 10% levels are marked by ***, ** and *, respectively.

Panel A: The Variance Process

	Const.	CAY _t	RREL _t	v_t^2	R^2
v_{t+1}^2	0.012** (0.005)	-0.029** (0.012)	0.123 (0.129)	0.530*** (0.190)	0.370

Panel B: The Process for the State Variables

	Const.	CAY _t	RREL _t	R^2
CAY _{t+1}	0.117*** (0.033)	0.700*** (0.084)	1.240* (0.639)	0.504
RREL _{t+1}	-0.002 (0.002)	0.005 (0.005)	0.845*** (0.056)	0.726

Panel C: The Model

	Const.	γ	CAY	RREL	R^2
Model 1	0.009** (0.004)	-0.463 (2.730)			0.000
Model 2	0.009* (0.004)	-0.401 (2.350)			0.011
Model 3	-0.188 (0.123)		0.504 (0.313)	-4.429 (3.827)	0.031
Model 4	-0.254 (0.155)	3.613 (3.438)	0.663* (0.389)	-4.344 (3.783)	0.037
Model 5	-0.242* (0.144)	2.951 (2.769)	0.634* (0.364)	-4.359 (3.783)	0.048
Model 6	-0.277 (0.177)	4.837 (4.250)	0.717 (0.443)	-4.315 (3.804)	0.041

Table 5: Estimation Results—Instrumental Variables, 12/83-4/95 Subsample
 For Panel A we estimate the conditional variance process in equation (35). For Panel B we estimate the VAR(1) process for the state variables in equation (36). For Panel C we estimate the model in equation (37) and various restricted versions thereof (see Section 4.1). In each case, the equations are estimated jointly via GMM over the period from December 1983 to April 1995 (137 observations). Heteroscedasticity consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5% and 10% levels are marked by ***, ** and *, respectively.

Panel A: The Variance Process

	Const.	CAY _t	RREL _t	v_t^2	R^2
v_{t+1}^2	0.003*** (0.000)	-0.092*** (0.026)	-0.041 (0.075)	0.401*** (0.083)	0.307

Panel B: The Process for the State Variables

	Const.	CAY _t	RREL _t	R^2
CAY _{t+1}	-0.000 (0.001)	0.836*** (0.041)	0.141 (0.219)	0.692
RREL _{t+1}	-0.000 (0.000)	-0.008 (0.010)	0.716*** (0.081)	0.520

Panel C: The Model

	Const.	γ	CAY	RREL	R^2
Model 1	0.027** (0.013)	-2.553 (3.438)			0.005
Model 2	0.025** (0.011)	-1.972 (2.692)			0.053
Model 3	0.017*** (0.006)		1.772*** (0.460)	-5.422** (2.128)	0.121
Model 4	-0.044* (0.025)	14.346** (5.773)	3.809*** (0.968)	-3.747 (2.596)	0.176
Model 5	-0.017 (0.015)	7.897** (3.504)	2.893*** (0.615)	-4.500* (2.397)	0.190
Model 6	-0.099** (0.043)	27.466*** (10.194)	5.671*** (1.586)	-2.216 (3.250)	0.145

Table 6: Estimation Results–Instrumental Variables, Quarterly Data

For Panel A we estimate the conditional variance process in equation (35). For Panel B we estimate the VAR(1) process for the state variables in equation (36). For Panel C we estimate various versions of the model given in equations (30)-(34) and the full specification in equation (37). In each case, the equations are estimated jointly via GMM over the period 1952Q3-2002Q3 (201 observations). Heteroscedasticity consistent standard errors are reported in parentheses. Coefficients that are significant at the 1%, 5% and 10% levels are marked by ***, ** and *, respectively.

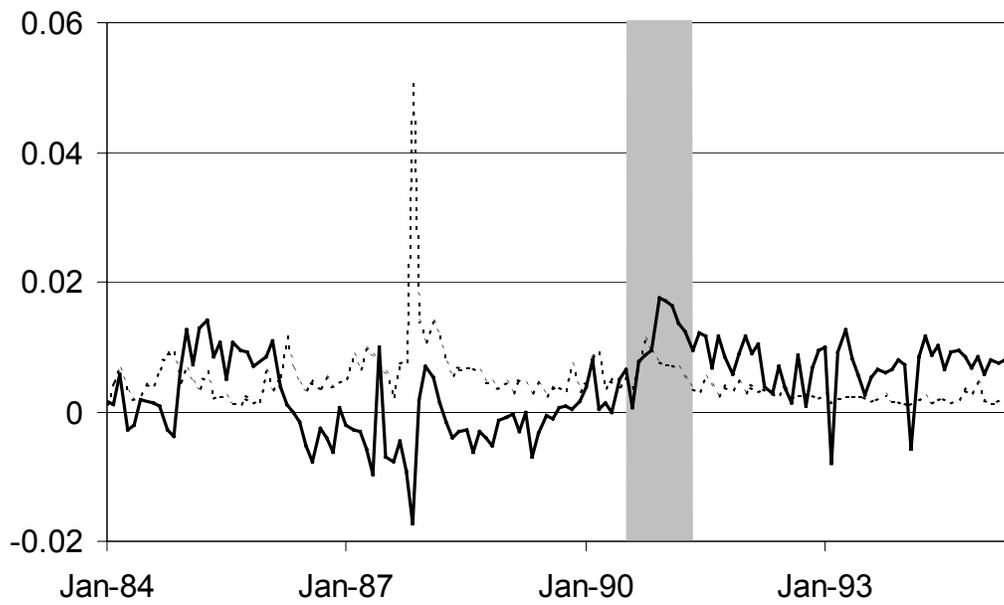


Figure 1: The Components of Expected Returns—Implied Volatility Data
 The risk (dashed line) and hedge (solid line) components of expected returns for the period 12/83 to 4/95, using the implied volatility data. The estimation results are reported in Table 3, Panel C, Model 5. Recessions are marked as shaded bars.

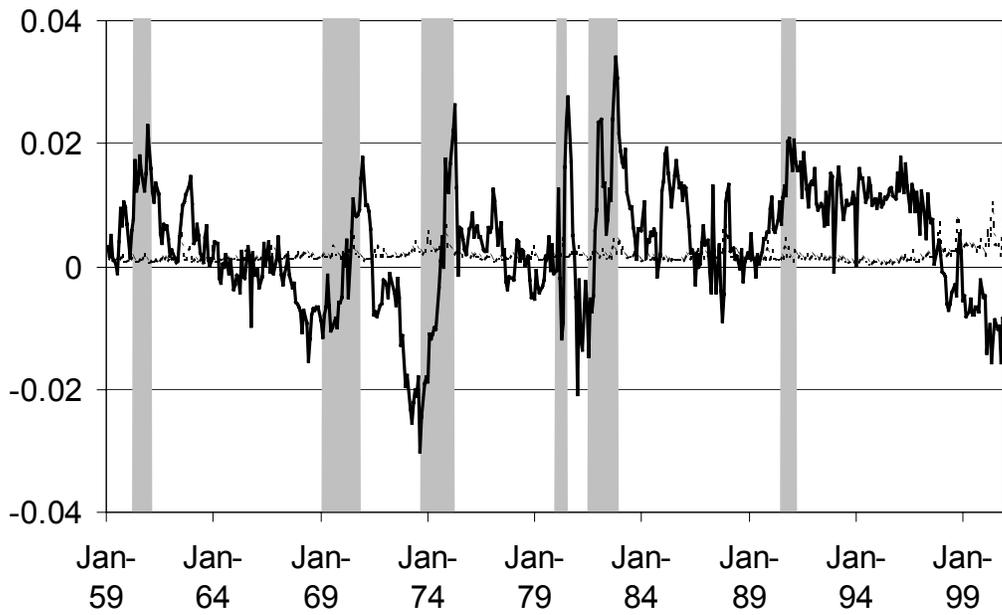


Figure 2: The Components of Expected Returns—Instrumental Variables
 The risk (dashed line) and hedge (solid line) components of expected returns for the period 2/59 to 12/00, using instrumental variables to estimate the conditional variance. The estimation results are reported in Table 4, Panel C, Model 5. Recessions are marked as shaded bars.