

Hard Times

John Y. Campbell, Stefano Giglio, and Christopher Polk¹

First draft: April 2010

This draft: May 2011

¹Campbell: Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138, and NBER. Email john_campbell@harvard.edu. Phone 617-496-6448 Giglio: Department of Economics, Littauer Center, Harvard University, Cambridge MA 02138. Email sgiglio@fas.harvard.edu. Polk: Department of Finance, London School of Economics, London WC2A 2AE, UK. Email c.polk@lse.ac.uk. We are grateful to Tuomo Vuolteenaho for conversations and data analyses that helped to motivate and shape this paper, and to seminar participants at the Stanford SITE 2010 conference and at Princeton for comments.

Abstract

This paper shows that the stock market downturns of 2000-2002 and 2007-09 have very different proximate causes. The early 2000's saw a large increase in the discount rates applied to corporate profits by rational investors, while the late 2000's saw a decrease in rational expectations of future profits. In each case the downturn reversed the trends of the previous boom. We reach these conclusions using a vector autoregressive model of aggregate stock returns and valuations, estimated both freely and imposing the cross-sectional restrictions of the intertemporal capital asset pricing model (ICAPM). As stock returns are very noisy, exploiting an economic model such as the ICAPM to extract information about future corporate profits from realized returns can potentially improve precision. We show that the contrast between the two downturns is also visible in a simple graphical data analysis. Our findings imply that the 2007-09 downturn was particularly serious for rational long-term investors, whose losses were not offset by improving stock return forecasts to the same extent as they were earlier in the decade.

JEL classification: G12, G14, N22

1 Introduction

During the past 15 years the US stock market has experienced two long booms, in each case followed by a sharp downturn. From the end of March 1994 through the end of March 2000, the S&P 500 index rose 221% in current dollars and 177% after adjustment for inflation. In the following two years (from March 2000 to September 2002), it declined 39% (42% in real terms). Similarly, from September 2002 to September 2007 the S&P 500 rose 75% (51%) and from September 2007 to March 2009 declined 44% (45%).

How should we interpret these dramatic fluctuations? Adopting the perspective of a rational investor or stock market analyst, should we think of the stock market booms as reflecting good news about future corporate profits, discounted at a constant rate as in traditional “random walk” models of stock prices? Or were stock prices driven up by declines in the discount rates that rational investors applied to corporate cash flows? And when the booms ended, did prices fall because rational investors became pessimistic about profits, or because they discounted future profits more heavily?²

Answers to these questions are important for several reasons. They tell us about the proximate causes of stock market fluctuations, and they allow us to track the rational outlook for the stock market over time. If the hard times experienced by stock market investors in 2000-02 and 2007-09 were due to lower expected corporate profits, then they were permanent in the sense that rational investors had no reason to expect stock prices to rebound to previous levels. On the other hand, if those hard times were driven by an increase in discount rates, or equivalently expected future returns, then it was rational to expect stock prices to recover over time, and in this sense the hard times were temporary.

Related to this, the optimal consumption response of a long-term investor can be different depending on the nature of a stock market downturn. A long-term investor with an intertemporal elasticity of substitution (IES) less than one should cut consumption less when discount rates increase than when expected profits decline; the reverse is true for a long-term investor with IES greater than one, who should save aggressively when discount rates increase.

²An increase in the discount rates applied by rational investors can occur for several reasons: an increase in aggregate risk; an increase in the risk aversion of rational investors; or a transfer of aggregate risk from irrational to rational investors, as in models with noise traders who have fluctuating sentiment and sell stocks to rational investors when they become pessimistic.

In this paper we argue that the downturns of 2000–02 and 2007–09 have very different proximate causes. In 2000–02, stock prices fell primarily because discount rates increased, while in 2007–09 cash flow prospects worsened, with discount rates playing little role until late 2008. Similarly, the preceding booms were driven primarily by discount rates in the 1990’s and by a mix of cash flows and discount rates in the mid-2000’s.

We reach this conclusion using a structured econometric approach with three main ingredients: first, a vector autoregressive (VAR) model of aggregate stock returns, valuation ratios, and other relevant financial variables; second, the approximate accounting identity of Campbell and Shiller (1988); and third, the cross-sectional restrictions of the intertemporal capital asset pricing model (ICAPM) of Merton (1973) and Campbell (1993), as implemented empirically by Campbell (1996), Campbell and Vuolteenaho (2004), and Campbell, Polk, and Vuolteenaho (2010). Relative to these earlier papers, our contribution here is to estimate the aggregate VAR imposing the cross-sectional restrictions of the ICAPM, thereby reducing uncertainty about the components of stock market fluctuations under the assumption that the ICAPM is correct.

We impose these restrictions because forecasting the equity premium with a pure time-series-based approach is a difficult task. Consequently, exploiting the economic logic of a cross-sectional asset-pricing model can help sharpen forecasts if the model imposed does a reasonably good job describing patterns in average returns. We join others in arguing that imposing such economically reasonable guidelines can be useful in forecasting subsequent excess market returns.³

Of course, the VAR methodology used in the above tests relies on specific assumptions about the data-generating process. Though the assumptions we make are reasonable, we also show that our findings about the proximate causes of the 2000–2002 and 2007–09 downturns, and their link to the cross section of equity returns, are consistent with a much simpler, less theoretically structured analysis of financial and macroeconomic data.

³Campbell and Shiller (1988a, 1988b) and Fama and French (1989) argue that high stock prices should imply a low equity premium. Merton (1980) argues that the equity premium should usually be positive because of risk aversion. Polk, Thompson, and Vuolteenaho (2006) argue that the cross-sectional pricing of risk should be consistent with the time-series pricing of risk, and assume the CAPM to make that comparison. Campbell and Thompson (2008) argue that imposing the restrictions of steady-state valuation models improves forecasting ability.

Beyond simply forecasting the equity premium, our results provide insight into the process by which the market prices the cross-section of equities. The model we impose argues that value stocks (relative to growth stocks) do better on average but worse during those stock market downturns that are permanent, in the sense that those downturns reflect hard times today due to expectations of lower corporate profits in the future. Our empirical success confirms that this economic model was a useful description of the recent US stock market experience.

Other work has used implications from the cross section to derive new equity premium predictors. For example, Polk, Thompson, and Vuolteenaho (2006) point out that if the CAPM is true, a high equity premium implies low prices for stocks with high betas. Relative valuations of high-beta stocks can therefore be used to predict the market return. Though their CAPM-based equity premium predictor does well in the pre-1963 subsample, it performs poorly in the post-1963 subsample, perhaps not surprising given the poor performance of the CAPM in that period. Unlike Polk, Thompson, and Vuolteenaho (2006), not only do we use an asset-pricing model (the ICAPM) that has had better empirical success in the post-1963 sample, we estimate a time series model that is fully restricted to be consistent with cross-sectional pricing.

Additionally, our results suggest that tests of the Campbell and Vuolteenaho (2004) implementation of the ICAPM that jointly estimate both the VAR coefficients and the pricing parameters together will be favorable to that model. Not only will the model's pricing performance improve but the integrity of the resulting news terms may not be dramatically sacrificed. Though a joint estimation approach will twist the VAR coefficients away from the OLS estimates used by Campbell and Vuolteenaho, in order to better fit the more precisely measured cross-sectional pricing implications, the resulting equity premium forecasts should still perform well out of sample.

Our final contribution is to expand the set of variables included in the Campbell and Vuolteenaho VAR. We specifically add the default yield spread, as shocks to this variable should contain information about future corporate profits. Consistent with this intuition, our restricted VAR chooses to include the default spread as an important component of aggregate cash-flow news. Interestingly, though the key variable of Campbell and Vuolteenaho, the small-stock value spread, continues to be an important component of market news, its role does not seem as critical in our structured econometric approach. This helps to address concerns about the sensitivity of the results in Campbell and Vuolteenaho (2004) and Campbell, Polk,

and Vuolteenaho (2010) to the inclusion of the small-stock value spread.

A precursor to our paper is Ranish (2009). This paper also argues that cash-flow news was relatively important in the downturn of 2007-09, but it does so using high-frequency data and does not seek to use the restrictions of asset pricing models to improve the precision of the return decomposition. More recently, Lettau and Ludvigson (2011) have used cointegration analysis of aggregate consumption and major components of wealth to distinguish permanent and transitory movements in wealth. Consistent with our findings, Lettau and Ludvigson argue that the 2000-02 downturn was almost entirely driven by transitory shocks, while the 2007-09 downturn reflected both permanent and transitory shocks.

The remainder of the paper is organized as follows. Section 2 explains our methodology for identifying the components of stock returns. Section 3 discusses the data, and Section 4 presents our VAR estimates with and without ICAPM restrictions. Section 5 contrasts the two boom-bust cycles of the late 1990's and early 2000's and the mid to late 2000's. Section 6 concludes.

2 Identifying the Components of Stock Returns

2.1 Cash-flow and discount-rate shocks

Campbell and Shiller (1988) provide a convenient framework for analyzing cash-flow and discount-rate shocks. They develop a loglinear approximate present-value relation that allows for time-varying discount rates. Linearity is achieved by approximating the definition of log return on a dividend-paying asset, $r_{t+1} \equiv \log(P_{t+1} + D_{t+1}) - \log(P_t)$, around the mean log dividend-price ratio, $(\bar{d}_t - \bar{p}_t)$, using a first-order Taylor expansion. Above, P denotes price, D dividend, and lower-case letters log transforms. The resulting approximation is $r_{t+1} \approx k + \rho p_{t+1} + (1 - \rho)d_{t+1} - p_t$, where ρ and k are parameters of linearization defined by $\rho \equiv 1/(1 + \exp(\bar{d}_t - \bar{p}_t))$ and $k \equiv -\log(\rho) - (1 - \rho)\log(1/\rho - 1)$. When the dividend-price ratio is constant, then $\rho = P/(P + D)$, the ratio of the ex-dividend to the cum-dividend stock price. The approximation here replaces the log sum of price and dividend with a weighted average of log price and log dividend, where the weights are determined by the average relative magnitudes of these two variables.

Solving forward iteratively, imposing the “no-infinite-bubbles” terminal condition that $\lim_{j \rightarrow \infty} \rho^j (d_{t+j} - p_{t+j}) = 0$, taking expectations, and subtracting the current dividend, one gets

$$p_t - d_t = \frac{k}{1 - \rho} + \mathbf{E}_t \sum_{j=0}^{\infty} \rho^j [\Delta d_{t+1+j} - r_{t+1+j}], \quad (1)$$

where Δd denotes log dividend growth. This equation says that the log price-dividend ratio is high when dividends are expected to grow rapidly, or when stock returns are expected to be low. The equation should be thought of as an accounting identity rather than a behavioral model; it has been obtained merely by approximating an identity, solving forward subject to a terminal condition, and taking expectations. Intuitively, if the stock price is high today, then from the definition of the return and the terminal condition that the dividend-price ratio is non-explosive, there must either be high dividends or low stock returns in the future. Investors must then expect some combination of high dividends and low stock returns if their expectations are to be consistent with the observed price.

Campbell (1991) extends the loglinear present-value approach to obtain a decomposition of returns. Substituting (1) into the approximate return equation gives

$$\begin{aligned} r_{t+1} - \mathbf{E}_t r_{t+1} &= (\mathbf{E}_{t+1} - \mathbf{E}_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (\mathbf{E}_{t+1} - \mathbf{E}_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j} \quad (2) \\ &= N_{CF,t+1} - N_{DR,t+1}, \end{aligned}$$

where N_{CF} denotes news about future cash flows (i.e., dividends or consumption), and N_{DR} denotes news about future discount rates (i.e., expected returns). This equation says that unexpected stock returns must be associated with changes in expectations of future cash flows or discount rates. An increase in expected future cash flows is associated with a capital gain today, while an increase in discount rates is associated with a capital loss today. The reason is that with a given dividend stream, higher future returns can only be generated by future price appreciation from a lower current price.

If the decomposition is applied to the returns on the investor’s portfolio, these return components can be interpreted as permanent and transitory shocks to the investor’s wealth. Returns generated by cash-flow news are never reversed subsequently, whereas returns generated by discount-rate news are offset by lower returns in the future. From this perspective it should not be surprising that conservative long-term investors are more averse to cash-flow risk than to discount-rate risk.

2.2 VAR methodology

An important issue is how to measure the shocks to cash flows and to discount rates. One approach, introduced by Campbell (1991), is to estimate the cash-flow-news and discount-rate-news series using a vector autoregressive (VAR) model. This VAR methodology first estimates the terms $E_t r_{t+1}$ and $(E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j}$ and then uses realization of r_{t+1} and equation (2) to back out the cash-flow news. Because of the approximate identity linking returns, dividends, and stock prices, this approach yields results that are almost identical to those that are obtained by forecasting cash flows explicitly using the same information set, provided that the information set includes the dividend yield and sufficient lags of the forecasting variables. Replacing the dividend yield with an alternative smooth valuation ratio, such as the smoothed earnings-price ratio or book-price ratio, also generates similar results whether returns or cash flows are forecast. Thus the choice of variables to enter the VAR is the important decision in implementing this methodology.⁴

When extracting the news terms in our empirical tests, we assume that the data are generated by a first-order VAR model

$$z_{t+1} = a + \Gamma z_t + u_{t+1}, \quad (3)$$

where z_{t+1} is a m -by-1 state vector with r_{t+1} as its first element, a and Γ are m -by-1 vector and m -by- m matrix of constant parameters, and u_{t+1} an i.i.d. m -by-1 vector of shocks. Of course, this formulation also allows for higher-order VAR models via a simple redefinition of the state vector to include lagged values.

Provided that the process in equation (3) generates the data, $t + 1$ cash-flow and discount-rate news are linear functions of the $t + 1$ shock vector:

$$\begin{aligned} N_{DR,t+1} &= e1' \lambda u_{t+1}, \\ N_{CF,t+1} &= (e1' + e1' \lambda) u_{t+1}. \end{aligned} \quad (4)$$

Above, $e1$ is a vector with first element equal to unity and the remaining elements equal to zeros. The VAR shocks are mapped to news by λ , defined as $\lambda \equiv \rho \Gamma (I - \rho \Gamma)^{-1}$. $e1' \lambda$ captures the long-run significance of each individual VAR shock to

⁴Chen and Zhao (2009) discuss the sensitivity of VAR decomposition results to alternative specifications. Campbell, Polk, and Vuolteenaho (2010), Cochrane (2008), and Engsted, Pedersen, and Tanggaard (2010) clarify the conditions under which VAR results are robust to the decision whether to forecast returns or cash flows.

discount-rate expectations. The greater the absolute value of a variable's coefficient in the return prediction equation (the top row of Γ), the greater the weight the variable receives in the discount-rate-news formula. More persistent variables should also receive more weight, which is captured by the term $(I - \rho\Gamma)^{-1}$.

2.3 Imposing the ICAPM

Campbell (1993) derives an approximate discrete-time version of Merton's (1973) intertemporal CAPM. The model's central pricing statement is based on the first-order condition for an investor who holds a portfolio p of tradable assets that contains all of her wealth. Campbell assumes that this portfolio is observable in order to derive testable asset-pricing implications from the first-order condition.

Campbell considers an infinitely lived investor who has the recursive preferences proposed by Epstein and Zin (1989, 1991), with time discount factor δ , relative risk aversion γ , and elasticity of intertemporal substitution ψ . Campbell assumes that all asset returns are conditionally lognormal, and that the investor's portfolio returns and its two components are homoskedastic. The assumption of lognormality can be relaxed if one is willing to use Taylor approximations to the true Euler equations, and the model can be extended to allow changing variances, something we plan to do in future research.

Campbell derives an approximate solution in which risk premia depend only on the coefficient of relative risk aversion γ and the discount coefficient ρ , and not directly on the elasticity of intertemporal substitution ψ . The approximation is accurate if the elasticity of intertemporal substitution is close to one, and it holds exactly in the limit of continuous time (Schroder and Skiadas 1999) if the elasticity equals one. In the $\psi = 1$ case, $\rho = \delta$ and the optimal consumption-wealth ratio is conveniently constant and equal to $1 - \rho$.

Under these assumptions, the optimality of portfolio strategy p requires that the risk premium on any asset i satisfies

$$E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_{i,t}^2}{2} = \gamma \text{Cov}_t(r_{i,t+1}, r_{p,t+1} - E_t r_{p,t+1}) + (1 - \gamma) \text{Cov}_t(r_{i,t+1}, -N_{p,DR,t+1}), \quad (5)$$

where p is the optimal portfolio that the agent chooses to hold and $N_{p,DR,t+1} \equiv$

$(E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{p,t+1+j}$ is discount-rate or expected-return news on this portfolio.

The left hand side of (5) is the expected excess log return on asset i over the riskless interest rate, plus one-half the variance of the excess return to adjust for Jensen's Inequality. This is the appropriate measure of the risk premium in a lognormal model. The right hand side of (5) is a weighted average of two covariances: the covariance of return i with the return on portfolio p , which gets a weight of γ , and the covariance of return i with negative of news about future expected returns on portfolio p , which gets a weight of $(1 - \gamma)$. These two covariances represent the myopic and intertemporal hedging components of asset demand, respectively. When $\gamma = 1$, it is well known that portfolio choice is myopic and the first-order condition collapses to the familiar one used to derive the pricing implications of the CAPM.

Campbell and Vuolteenaho (2004) rewrite equation (5) to relate the risk premium to betas with cash-flow news and discount-rate news. Using $r_{p,t+1} - E_t r_{p,t+1} = N_{p,CF,t+1} - N_{p,DR,t+1}$ to replace the portfolio covariance with news covariances, and then multiplying and dividing by the conditional variance of portfolio p 's return, $\sigma_{p,t}^2$, we have

$$E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_{i,t}^2}{2} = \gamma \sigma_{p,t}^2 \beta_{i,CF,t} + \sigma_{p,t}^2 \beta_{i,DR,t}. \quad (6)$$

Here the cash-flow beta $\beta_{i,CF}$ is defined as

$$\beta_{i,CF} \equiv \frac{\text{Cov}(r_{i,t}, N_{CF,t})}{\text{Var}(r_{M,t}^e - E_{t-1} r_{M,t}^e)}, \quad (7)$$

and the discount-rate beta $\beta_{i,DR}$ as

$$\beta_{i,DR} \equiv \frac{\text{Cov}(r_{i,t}, -N_{DR,t})}{\text{Var}(r_{M,t}^e - E_{t-1} r_{M,t}^e)}. \quad (8)$$

Note that the discount-rate beta is defined as the covariance of an asset's return with *good* news about the stock market in form of *lower-than-expected* discount rates, and that each beta divides by the total variance of unexpected market returns, not the variance of cash-flow news or discount-rate news separately. This implies that the cash-flow beta and the discount-rate beta add up to the total market beta,

$$\beta_{i,M} = \beta_{i,CF} + \beta_{i,DR}. \quad (9)$$

Equation (6) delivers the prediction that “bad beta” with cash-flow news should have a risk price γ times greater than the risk price of “good beta” with discount-rate news, which should equal the variance of the return on portfolio p .

In our empirical work, we assume that portfolio p is fully invested in a value-weighted equity index. This assumption implies that the risk price of discount-rate news should equal the variance of the value-weighted index. The only free parameter in equation (6) is then the coefficient of relative risk aversion, γ .

3 Data and Econometrics

Our estimation method involves specifying a set of state variables for the VAR, together with a set of test assets on which we impose the ICAPM conditions. We first describe the data, then our econometric approach for imposing the restrictions of the asset pricing model.

3.1 VAR data

Our full VAR specification includes five variables, four of which are the same as in Campbell and Vuolteenaho (2004). Because of data availability issues, we replace the term yield series used in that paper with a new series, as described below. To those four variables, we add a default yield spread series. The data are all quarterly, from 1929:2 to 2010:4.

The first variable in the VAR is the excess log return on the market, r_M^e , the difference between the log return on the Center for Research in Securities Prices (CRSP) value-weighted stock index and the log risk-free rate. The risk-free data are constructed by CRSP from Treasury bills with approximately three month maturity.

The second variable is the price-earning ratio (PE) from Shiller (2000), constructed as the price of the S&P 500 index divided by a ten-year trailing moving average of aggregate earnings of companies in the S&P 500 index. Following Graham and Dodd (1934), Campbell and Shiller (1988b, 1998) advocate averaging earnings over several years to avoid temporary spikes in the price-earnings ratio caused by cyclical declines in earnings. We avoid any interpolation of earnings in order to en-

sure that all components of the time- t price-earnings ratio are contemporaneously observable by time t . The ratio is log transformed. We have explored replacing PE with the price-to-dividend ratio, PD , measured using average dividends over the past year; our results remain qualitatively similar for our full specification.

Third, the term yield spread (TY) is obtained from Global Financial Data. In Campbell and Vuolteenaho (2004), TY was computed as the yield difference between ten-year constant-maturity taxable bonds and short-term taxable notes. Since the series used to construct it were discontinued in 2002, we compute the TY series as the difference between the log yield on the 10-Year US Constant Maturity Bond (IGUSA10D) and the log yield on the 3-Month US Treasury Bill (ITUSA3D).

Fourth, the small-stock value spread (VS) is constructed from the data on the six “elementary” equity portfolios made available by Professor Kenneth French on his web site. These elementary portfolios, which are constructed at the end of each June, are the intersections of two portfolios formed on size (market equity, ME) and three portfolios formed on the ratio of book equity to market equity (BE/ME). The size breakpoint for year t is the median NYSE market equity at the end of June of year t . BE/ME for June of year t is the book equity for the last fiscal year end in $t - 1$ divided by ME for December of $t - 1$. The BE/ME breakpoints are the 30th and 70th NYSE percentiles.

At the end of June of year t , we construct the small-stock value spread as the difference between the $\log(BE/ME)$ of the small high-book-to-market portfolio and the $\log(BE/ME)$ of the small low-book-to-market portfolio, where BE and ME are measured at the end of December of year $t - 1$. For months from July to May, the small-stock value spread is constructed by adding the cumulative log return (from the previous June) on the small low-book-to-market portfolio to, and subtracting the cumulative log return on the small high-book-to-market portfolio from, the end-of-June small-stock value spread. The construction of this series follows Campbell and Vuolteenaho (2004) closely.

The fifth and last variable in our VAR is the default spread (DEF), defined as the difference between the log yield on Moody’s BAA and AAA bonds. The series is obtained from the Federal Reserve Bank of St. Louis. We add the default spread to the Campbell and Vuolteenaho (2004) VAR specification partially because that variable is known to track time-series variation in expected excess returns on the market portfolio (Fama and French, 1989), but mostly because shocks to the default spread should to some degree reflect news about aggregate default probabilities. Of

course, news about aggregate default probabilities should in turn reflect news about the market's future cash flows.

Table 1 reports descriptive statistics on these variables. The lower panel of the table shows some quite strong correlations among the VAR explanatory variables, for example a positive correlation of 0.65 between the value spread and the default spread and a negative correlation of -0.60 between the log price-earnings ratio and the default spread. These correlations complicate the interpretation of individual VAR coefficients when all the variables are included in the VAR.

3.2 Test asset data

Our main set of test assets is the six elementary ME and BE/ME sorted portfolios, described in the previous section. We impose the ICAPM conditions on the returns of these six assets and on the return of the market portfolio, the CRSP value-weighted stock index.

All the test portfolios are highly correlated with the market return. When we estimate the model, we impose the ICAPM equations on the difference between the return of each test asset and the return of the market; in this way, we remove part of the correlation between the errors of the moment conditions, which is computationally convenient. The condition on the market return is instead imposed in excess of the risk-free rate.

3.3 Estimation methodology

This section details the estimation technique used to solve for the restricted model, using Hansen, Heaton, and Yaron's (1996) continuously updated Generalized Method of Moments (GMM).

We use the notation K for the dimension of the VAR and I for the number of test assets. The restricted model gives us $R = K(K + 1) + I + 1$ orthogonality conditions. $K(K + 1)$ of these estimate the intercepts and dynamic coefficients of the VAR, and $I + 1$ orthogonality conditions are imposed by the ICAPM on the test assets and the market portfolio. There is one free parameter in the ICAPM, the coefficient of relative risk aversion, so there are I overidentifying restrictions.

The VAR restrictions impose, for each equation k , that the error at $t + 1$ is uncorrelated with each of the state variables measured at time t . They also impose a zero unconditional mean on the innovation vector.

The ICAPM conditions are derived as follows. Starting from equation (5), we assume lognormality, substitute in the market portfolio for portfolio p and use the VAR conditions to get

$$E_t(R_{i,t+1} - R_{f,t+1}) = \gamma \text{Cov}_t(r_{i,t+1}, e1'(z_{t+1} - a - \Gamma z_t)) + (1 - \gamma) \text{Cov}_t(r_{i,t+1}, -e1'\lambda(z_{t+1} - a - \Gamma z_t)). \quad (10)$$

We can condition down using the fact that $E_t[u_{t+1}] = 0$, so that we obtain

$$E(R_{i,t+1} - R_{f,t+1}) = \gamma E(r_{i,t+1} e1'(z_{t+1} - a - \Gamma z_t)) - (1 - \gamma) E(r_{i,t+1} e1'\lambda(z_{t+1} - a - \Gamma z_t)). \quad (11)$$

We use this orthogonality condition for the market portfolio, but rewrite the I orthogonality conditions for the test assets in excess of the market return, rather than the riskfree rate:

$$E[R_{t+1}^i - R_{t+1}^M - (r_{t+1}^i - r_{t+1}^M) e1'(\gamma - (1 - \gamma)\rho\Gamma(I - \rho\Gamma)^{-1})(z_{t+1} - a - \Gamma z_t)] = 0. \quad (12)$$

This is useful for the numerical estimation because it removes a large amount of the correlation between the errors of the moment conditions.

When $K = 5$, we have a large number of parameters to estimate. We therefore have to restrict our continuously updated GMM estimation procedure in several ways in order to achieve convergence to acceptable parameter values for every subsample, a property that we need for out-of-sample analysis. First, we place an upper bound on the risk aversion coefficient γ . This results in an estimate that in the full sample and some subsamples actually hits the bound, while in others it converges below it. In analysis not shown, we find that the results are robust to various specifications for this upper bound. We also impose a lower bound on γ of one, which is never binding.

Second, we impose stationarity on the estimated VAR by requiring that the absolute value of the maximum eigenvalue of the transition matrix Γ is less than or equal to a value $\bar{\lambda} < 1$. We explore the cases $\bar{\lambda} = 0.98$ and $\bar{\lambda} = 0.99$ and report the former in the online Appendix. Again, results appear to be robust to this choice.

In order to define convergence of the estimator, we use a tolerance level of $1e-5$ on the objective function value (whose order of magnitude is around $1e-2$), $1e-4$ on the values of the parameters, $1e-3$ on γ , and $1e-4$ on the maximum eigenvalue constraint and on the response to discount rate news shocks.

Finally, since in some cases the search function seems to converge to local minima, we start the estimation from several different points, where the VAR parameters are the ones of the unrestricted OLS estimate, while γ varies from 1 to the upper bound. This method seems to converge well to the global minimum.

4 Alternative VAR Estimates

In Tables 2 through 4, we estimate three alternative VAR systems. For comparison with previous work, we begin in Table 2 with a simple two-variable VAR without restrictions, including only the market return and log price-earnings ratio as state variables. Then in Tables 3 and 4 we include all five state variables, first without restrictions and then imposing the restrictions of the ICAPM described in the previous section.

4.1 A two-variable VAR system

Table 2 reports results that are familiar from previous research using this methodology. The top panel shows that the market return is predicted negatively by the log price-smoothed earnings ratio (with a partial regression coefficient of -0.046 and a standard error of 0.016), which itself follows a persistent AR(1) process. This implies that discount rate news is quite volatile and explains most of the variance of the market return.

One way to see the extent to which discount-rate news is an important component of the market return is to calculate the coefficients mapping state variable shocks into news terms, as we do in the middle panel of Table 2. If we orthogonalize the state variable shocks, using a Cholesky decomposition with the market return ordered first, the “structural” market return shock gets credit for the movement in the price-earnings ratio that normally accompanies a market return shock, while the structural shock to the price-earnings ratio is interpreted as an increase in the price-earnings ratio without any change in the market return, that is, a negative shock to earnings with no change in price. The first shock has a discount-rate effect that is over four times larger than its cash-flow effect. The second shock carries both bad cash-flow news and offsetting good discount-rate news to keep the stock price constant.

Another way to see the importance of discount-rate news is to calculate the standard deviations of discount-rate and cash-flow news. We do this in the bottom panel of Table 2. Discount-rate news is more than twice as volatile as cash-flow news, consistent with results reported by Campbell (1991) and others. There is only a weak correlation of 0.13 between the two news terms.

We have explored what happens when we impose the restrictions of the ICAPM on this two-variable VAR system. The predictability of the market return from the price-earnings ratio diminishes (the partial regression coefficient is only 30% of its previous value), and therefore the volatility of discount-rate news diminishes. The estimated system implies that cash-flow and discount-rate news have similar volatilities and a large positive correlation; that is, almost all stock market fluctuations are attributed to a roughly equal mix of the two types of shocks, as if the market overreacts to cash-flow news. The estimate of risk aversion is a modest 2.2, and the overidentifying restrictions of this model are very strongly rejected.

These unpromising results are driven by the fact that in our full sample, the value spread is negatively correlated with the price-earnings ratio as shown in Table 1. During the Great Depression, the value spread was wide and the price-earnings ratio was low, while the postwar period has been characterized by a lower value spread and a higher average price-earnings ratio. Given this fact, a model that only includes the price-earnings ratio as a predictor variable implies that value stocks have high discount-rate betas (since on average they do well when the price-earnings ratio rises, and this predicts low future stock returns). Since the discount-rate beta has a low price of risk in the ICAPM, the implied value premium is actually lower than it would be in the simple CAPM; equivalently, the model implies that value stocks have a negative CAPM alpha. To mitigate this effect, the restricted model reduces the predictability of stock returns (but does not eliminate it altogether), and estimates a relatively low coefficient of relative risk aversion, thus a relatively small difference between the risk prices for cash-flow and discount-rate betas. The poor fit of the model to the cross-section of stock returns implies that the ICAPM restrictions can be statistically rejected.

4.2 Five-variable VAR systems

In Table 3 we include all five state variables in an unrestricted VAR. Consistent with previous research, the term spread predicts the market return positively while the value spread predicts it negatively; however, the predictive coefficients on these variables are not precisely estimated. The default spread has an imprecisely estimated negative coefficient, probably a symptom of multicollinearity among the explanatory variables as the default spread and the value spread have a correlation of 0.65 in Table 1. Discount-rate news is considerably more volatile than cash-flow discount-rate news, just as in the unrestricted two-variable model of Table 2. The volatility of aggregate cash-flow news is 0.043 while the volatility of discount-rate news is 0.100, more than twice as large. The correlation between these two components of the market shock is a relatively small -0.07.

While the unrestricted five-variable VAR has similar properties to the unrestricted two-variable VAR, the same is not true when we impose the ICAPM restrictions on the system. Table 4 reports a restricted five-variable VAR. The restrictions strengthen the predictive power of the term spread and the default spread for the market return, and weaken the predictive power of the value spread. The partial regression coefficient on TY is 0.012 with a standard error of 0.006. DEF forecasts the log excess market return in the presence of the other five variables with a coefficient of -0.051 whose standard error is 0.020. The two news terms are now each estimated to be more volatile than the market return itself, and strongly negatively correlated, implying that booms and busts are typically due to cash-flow news overwhelming discount-rate news or vice versa, but not typically due to an equal mix of both types of shocks. The coefficient of risk aversion is estimated to be around 5, and the overidentifying restrictions of the ICAPM are marginally rejected at the 4.4% significance level.

Shocks to the default spread play a much more significant role in the determination of cash-flow news in the restricted VAR relative to the unrestricted estimates. For the reduced-form mapping, the coefficient on the default spread is now -0.93, compared to the unrestricted estimate of -0.06. For the structural mapping, the coefficient on the default spread is -0.21 instead of -0.01. In other words, the restricted VAR interprets a widening default spread as a sign of deteriorating profit prospects at the aggregate level, and explains the value premium in part as the result of the sensitivity of value stock returns to surprise increases in the default spread. This finding seems reasonable, given the result in Fama and French (1993) that both small stocks and value stocks covary more with a default risk factor than large stocks and growth stocks

do. Of course, unlike Fama and French, we restrict the price of risk for exposure to this factor to be consistent with the ICAPM.

Interestingly, though the key variable of Campbell and Vuolteenaho, the small-stock value spread, continues to be an important component of market news, its role does not seem as critical in our structured econometric approach. This helps to address concerns about the sensitivity of the results in Campbell and Vuolteenaho (2004) and Campbell, Polk, and Vuolteenaho (2010) to the inclusion of the small-stock value spread. In results not shown, dropping VS from the restricted five-variable VAR has little effect on the qualitative conclusions drawn from the VAR. The discount-rate and cash-flow news terms appear very similar in the restricted four- and five-variable VAR systems, although it is true that the ICAPM restrictions are more strongly rejected at the 1% level in the four-variable system.

The online Appendix to this paper examines the robustness of these estimates to two other variations in the methodology. The first variation is that we impose a slightly tighter bound on the maximum eigenvalue of the transition matrix Γ . The restricted VAR estimates that result from this tighter bound deliver a risk aversion estimate that is approximately 40% larger, with volatilities of the two news terms that are about 35% smaller. The model is rejected slightly more strongly, at the 2.8% significance level, and the state variables enter both the reduced-form and structural mapping functions in much the same fashion. The second methodological variation allows the cash-flow and discount-rate beta estimates of the test assets to differ across the pre-1963 and the post-1963 subperiods. The resulting risk aversion coefficient is almost three times the value in Table 4, and the transition matrix estimate also differs to some degree. For example, in the return forecasting regression of the VAR, the coefficient on the lagged excess market return is now negative instead of positive. However, in neither case was the resulting estimate statistically significant from zero. Moreover, the news terms appear to behave similarly to those from Table 4. Again, these two news series are more volatile than the market return itself, and strongly negatively correlated.

5 Understanding Recent US Financial History

5.1 The VAR approach

What account do these alternative VAR models give of US financial history? In Figure 1, we report exponentially smoothed news series over the full sample period from 1929:2 through 2010:4. The smoothing parameter is 0.08 per quarter, corresponding to a half-life of two years. Our three models are organized vertically, in each case with cash-flow news on the left and the negative of discount-rate news on the right. Increases in each news series imply an increase in stock prices driven by cash-flow or discount-rate changes. For each model, the two smoothed news series add to the smoothed unexpected excess return on the stock market, which varies somewhat across models since different models imply different expected returns.

The three models give a fairly similar explanation of the large movements in the stock market over this eighty-year period. The Great Depression was a prolonged period of negative cash-flow news that lasted until World War II, together with a sharp increase in discount rates—equivalently, a decline in investor sentiment—in the very early 1930's. This was followed by a profit boom in the late 1940's through to the early 1960's, but discount rates remained high in this period (sentiment remained depressed), dampening the effect on stock prices. Profits were weak in the 1980's and early 1990's, and stronger if somewhat erratic during the last two decades of the sample. Declining discount rates (improving sentiment) drove stock prices up from the late 1970's through the year 2000.

In Figure 2, we look more closely at the period since 1995. All three models show that declining discount rates (improving sentiment) drove the stock market up during the late 1990's, and then reversed in the early 2000's. All three models also show a profit boom in the mid-2000's followed by a collapse in 2007-08. The restricted model shows a particularly strong hump shape in cash-flow news over this period.

There is less consistency across models about cash-flow news in the 1990's, which is estimated to be modestly positive in the two-variable model but not in the five-variable models. The models also give different accounts of the trough of the crisis in the fall of 2008 and the subsequent recovery. The unrestricted models interpret the sharp decline in the price-earnings ratio in late 2008 as a piece of negative discount-rate news; they say that investor sentiment, which had remained modestly positive

over the previous four years, collapsed at that point and recovered in 2009. The restricted model, on the other hand, attributes the stock market decline in the fall of 2008 to the arrival of extremely bad cash-flow news, offset to some degree by declining discount rates. As we shall see in the next section, this interpretation results from the emphasis placed by the restricted model on the default spread, which widened dramatically in the fall of 2008.

Summarizing these results, our VAR models tell us that the two boom-bust cycles of the 1990's and the 2000's were quite different in their proximate causes. The technology boom and bust that occurred in the late 1990's and early 2000's was primarily driven by discount-rate news. The credit boom and bust of the mid to late 2000's saw an extended profit boom followed by negative cash-flow news at the onset of the financial crisis. Discount rates remained low, contributing to high stock prices during the boom, and did not drive stock prices down until late 2008 at the earliest.

5.2 Out-of-sample return forecasts

Another way to understand the differences between the technology boom-bust and the credit boom-bust is to use our three VAR models to generate out-of-sample return forecasts during the period 1994–2010. We estimate each model on an expanding sample and, in Figure 3, plot the resulting out-of-sample forecasts. We also plot realized returns, rescaled to have the same standard deviation as the restricted five-variable forecasts in order to improve the readability of the figure.

The striking pattern in Figure 3 is that return forecasts were considerably lower during the boom of the late 1990's than they were during the boom of the mid-2000's, and they increased more strongly and rapidly during the downturn following the year 2000 than they did in 2007-08. Only at the very end of 2008 and in 2009 did return forecasts increase meaningfully. These differences are noticeable in all the models, but are stronger in the five-variable models than in the two-variable model, and strongest in the five-variable model with theoretical restrictions imposed. The implication is that the stock market downturn of the early 2000's was mitigated, for long-term investors, by an increase in expected future stock returns. This was much less the case in 2007-08.

This exercise also allows us to judge the success of our models at predicting excess stock returns in this period. The mean squared error of in-sample forecasts is 0.0088

over this period for the unrestricted two-variable and five-variable models, and 0.0086 for the restricted five-variable model. Out of sample, the mean squared errors are 0.0096 for the unrestricted two-variable model, 0.0102 for the unrestricted five-variable model, and 0.0100 for the restricted five-variable model. The restricted model is particularly successful at forecasting returns in the early part of the sample, and relatively less successful during the recovery from the credit crisis; thus shortening the evaluation period by dropping earlier data tends to favor the unrestricted models over the restricted model.

This evaluation period is both short and unusual, and we do not propose to pick models based on these numbers. However the results do suggest that theoretical restrictions can improve the out-of-sample performance of VAR forecasting models for stock returns, particularly when those models include many variables and therefore numerous free parameters.

5.3 Recent booms and busts in event time

The differences between the technology and credit booms and busts can be appreciated without using VAR methodology. Figure 4 plots several key aggregate inputs to our analysis. To aid comparison of the technology and credit episodes, we plot variables in event time, where the event is the stock market peak: 2000Q1 for the technology boom, and 2007Q3 for the credit boom. The horizontal axis is labelled in years relative to the market peak, and vertical lines are drawn every half year (two quarters).

One can see from Figure 4 that though, by construction, returns increased leading up to the peak and then decreased, there are clear differences in the source of that variation across these two key episodes in recent market history. For the tech boom, the PE ratio increased considerably during the time leading up to the peak, while during the credit boom, the market's PE was essentially flat. Since Campbell and Shiller (1988b) and others document that discount-rate news dominates cash-flow news in moving the ratios of prices to accounting measures of stock market value, movements in PE should be thought of as reflecting good news about market discount rates. In contrast, the market's smoothed earnings (E) grew strongly during the years of the credit boom, while during the tech boom there was a much smaller increase in E . In fact, any increase in E prior to the tech peak occurred entirely within a year of the market top. Of course, movements in this variable are naturally associated

with market cash-flow news.

Post event, the plots show that PE moved dramatically downward for both the tech and credit busts. However, during the credit bust this movement in PE is associated with a strong downward movement in E , while during the tech bust, aggregate earnings actually increased in the first three quarters after the market peak. Similar conclusions can be drawn from examining the price-to-dividend ratio, PD , and that ratio's cash-flow component, dividends over the last year (D). Movements in PD are very similar to movements in PE . Prices scaled by dividends rise and then fall around the tech boom. However D is essentially flat during the six years surrounding this episode. In stark contrast, D rises sharply throughout the three years preceding the credit boom and then begins to taper off in the months subsequent to the peak. The slowdown in the growth of D corresponds with a rapid decline in the PD ratio. These movements in simple aggregates are consistent with our claim that the tech boom and bust was primarily a discount-rate event while the credit boom and bust was primarily a cash-flow event.

Figure 4 also confirms the usefulness of examining the cross-section of asset returns for information about market aggregates. In particular, the movement in the value spread shows a striking difference across the two periods. During the tech boom, VS increases leading up to the market peak and then after the peak starts to decline, though this decline does take over a year to begin. In stark contrast, the VS decreases in the time leading up to the credit peak and then begins to rise sharply after the market begins to decline in late 2007. This response is exactly what one would expect if cross-sectional pricing followed the ICAPM of Merton (1973), Campbell (1996), Campbell and Vuolteenaho (2004), and Campbell, Polk, and Vuolteenaho (2010). Specifically, the price gap between expensive and cheap stocks should narrow during times when shocks to market cash-flow news are positive or when shocks to market discount-rate news are negative. Conversely, this price gap should widen during times when shocks to market cash-flow news are negative or when shocks to market discount-rate news are positive. We find it comforting that this straightforward prediction of the ICAPM can be seen clearly in the data, indicating that our conclusions do not hinge on the details of a VAR specification.

When we do estimate a VAR imposing the restrictions of the ICAPM, the restricted system places greater emphasis on movements in the default spread. Figure 4 shows that the default spread was relatively stable during the tech boom and bust, but spiked upwards in late 2008 and early 2009, between one and two years after the

market peak. This increase in the default spread, together with the simultaneous collapse in smoothed earnings also visible in Figure 4, accounts for the fact that the restricted VAR system perceives highly negative cash-flow news in the fall of 2008 as shown at the bottom right of the figure. Because cash-flow news and discount-rate news are negatively correlated in the restricted model, the restricted VAR estimate of discount-rate news is positive in the same period, illustrated in the bottom left of the figure.

5.4 Other macroeconomic evidence

The contrast between the technology and credit booms and busts is also visible in other macroeconomic time series not included in our VAR analysis. Figure 5 shows several of these series. The top panel of the figure is the cumulative log stock return, and the second panel reports seasonally adjusted logs of industrial production, vehicle miles travelled in the US, and the Ceridian-UCLA Pulse of Commerce Index which is based on diesel fuel consumption in the US trucking industry. All of these series show much greater declines in the credit bust than in the technology bust.

The third panel of Figure 5 shows the logs of business and nonbusiness bankruptcies. Both types of bankruptcy decline sharply in 2006 with a change in US bankruptcy law, but were flat during the technology bust and rapidly increasing during the credit bust.

Given that these time series are not visible to our VAR systems, it is striking how they confirm our VAR finding, one version of which (the unrestricted five-variable system) is recapitulated in the bottom two panels of the figure, that cash-flow news plays a greater role in the credit bust than in the technology bust.

Finally, we note that two other recent papers, Cochrane (2011) and Lettau and Ludvigson (2011), have used movements in aggregate consumption to analyze recent stock market fluctuations. Cochrane's Figure 9 shows that aggregate consumption declined far more dramatically in the credit bust than in the technology bust. This is what one would expect if consumers are forward-looking, have a low intertemporal elasticity of substitution, and perceived increasing rates of return on equity investments in the technology bust but not in the credit bust. Lettau and Ludvigson use cointegration analysis of consumption and wealth to argue that temporary wealth shocks drove down stock prices in 2000-02, while both temporary and perma-

ment shocks were important in 2007-09. Their finding, while obtained using a very different methodology, is consistent with ours.

6 Conclusion

Over the last three decades, financial economists have dramatically changed their interpretation of stock market movements. A wave of research has challenged the traditional paradigm in which the equity premium is constant, excess stock returns are unforecastable, and stock price fluctuations solely reflect news about corporate profits. Although there is still an active debate about the extent of predictability in stock returns (see for example Campbell and Thompson 2008, Cochrane 2008, and Goyal and Welch 2008), many financial economists have adopted a new paradigm in which a significant fraction of the variation in market returns reflects information about future expected returns.

The new paradigm implies that market returns are a very noisy proxy for corporate fundamentals. They often reflect temporary valuation movements instead of shifts in aggregate profitability, so it is difficult to learn about changes in corporate fundamentals simply from raw realized returns. We turn to asset pricing theory to provide a better understanding of the fundamentals hidden in stock market returns than can be achieved by purely statistical methods. In particular, we use the theoretical restrictions of the ICAPM to jointly estimate both a time-series model for the aggregate market return and a cross-sectional model of average stock returns. Out-of-sample tests confirm the usefulness of our theory-driven approach.

Our analysis implies that bad news about future corporate profits was much more important in the stock market downturn of 2007–2009 than in the previous downturn of 2000–2002. The earlier downturn was driven primarily by a large increase in expected future stock returns. Although the 2007–2009 proportional decline in stock prices was only slightly greater than the 2000–2002 decline, it had more serious implications for long-term rational investors, because there was a smaller increase in expected future returns to reassure investors that stock prices were likely to recover over time. Our model implies that high returns in late 2009 and 2010 reflect unexpected positive shocks. In this sense, times were particularly hard at the bottom of the recent downturn in March 2009.

References

- Campbell, John Y., 1991, “A Variance Decomposition for Stock Returns”, *Economic Journal* 101:157–179.
- Campbell, John Y., 1993, “Intertemporal Asset Pricing Without Consumption Data”, *American Economic Review* 83:487–512.
- Campbell, John Y., 1996, “Understanding Risk and Return”, *Journal of Political Economy* 104:298–345.
- Campbell, John Y., Christopher Polk, and Tuomo Vuolteenaho, 2010, “Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns” *Review of Financial Studies* 23:305–344.
- Campbell, John Y. and Robert J. Shiller, 1988a, “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors”, *Review of Financial Studies* 1:195–228.
- Campbell, John Y. and Robert J. Shiller, 1988b, “Stock Prices, Earnings, and Expected Dividends”, *Journal of Finance* 43:661–676.
- Campbell, John Y. and Robert J. Shiller, 1998, “Valuation Ratios and the Long-Run Stock Market Outlook”, *Journal of Portfolio Management* 24(2):11–26.
- Campbell, John Y. and Sam Thompson, 2008, “Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?”, *Review of Financial Studies* 21:1509–1531.
- Campbell, John Y. and Tuomo Vuolteenaho, 2004, “Bad Beta, Good Beta”, *American Economic Review* 94:1249–1275.
- Chen, Long and Xinlei Zhao, 2009, “Return Decomposition”, *Review of Financial Studies* 22:5213–5249.
- Cochrane, John H., 2008, “The Dog That Did Not Bark: A Defense of Return Predictability”, *Review of Financial Studies* 21:1533–1575.
- Cochrane, John H., 2011, “Discount Rates”, NBER Working Paper No. 16972, forthcoming *Journal of Finance*.

- Davis, James L., Eugene F. Fama, and Kenneth R. French, “Characteristics, Covariances, and Average Returns: 1929 to 1997”, *Journal of Finance* 55:389–406.
- Engsted, Tom, Thomas Q. Pedersen, and Carsten Tanggaard, 2010, “Pitfalls in VAR Based Return Decompositions: A Clarification”, unpublished paper, University of Aarhus.
- Epstein, Lawrence and Stanley Zin, 1989, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework”, *Econometrica* 57:937–69.
- Epstein, Lawrence and Stanley Zin, 1991, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Investigation”, *Journal of Political Economy* 99:263–86.
- Fama, Eugene F. and Kenneth R. French, 1989, “Business Conditions and Expected Returns on Stocks and Bonds”, *Journal of Financial Economics* 25:23–50.
- Fama, Eugene F. and Kenneth R. French, 1993, “Common Risk Factors in the Returns on Stocks and Bonds,” *Journal of Financial Economics* 33, 3–56.
- Goyal, Amit and Ivo Welch, 2008, “A Comprehensive Look at the Empirical Performance of Equity Premium Prediction”, *Review of Financial Studies* 21:1455–1508.
- Graham, Benjamin and David L. Dodd, 1934, *Security Analysis*, first edition, McGraw Hill, New York.
- Hansen, Lars Peter, John Heaton, and Amir Yaron, 1996, “Finite-Sample Properties of Some Alternative GMM Estimators”, *Journal of Business and Economic Statistics* 14:262–280.
- Lettau, Martin and Sydney C. Ludvigson, 2011, “Shocks and Crashes”, NBER Working Paper No. 16996.
- Merton, Robert C., 1973, “An Intertemporal Capital Asset Pricing Model”, *Econometrica* 41:867–87.
- Polk, Christopher, Sam Thompson, and Tuomo Vuolteenaho, 2006, “Cross-Sectional Forecasts of the Equity Premium”, *Journal of Financial Economics* 81:101–141.

Ranish, Ben, 2009, “Was This Stock Market Crash Different?”, unpublished paper, Harvard University.

Schroder, Mark and Costis Skiadas, 1999, “Optimal Consumption and Portfolio Selection With Stochastic Differential Utility”, *Journal of Economic Theory* 89:68–126.

Table 1: **Descriptive Statistics**

The table reports the descriptive statistics of the VAR state variables over the full sample period 1929:2-2010:4, 327 quarterly data points. r_M^e is the excess log return on the CRSP value-weight index. PE is the log ratio of the S&P 500's price to the S&P 500's ten-year moving average of earnings. TY is the term yield spread in percentage points, measured as the yield difference between the log yield on the ten-year US constant maturity bond and the log yield on the three-month US treasury. VS is the small-stock value-spread, the difference in the log book-to-market ratios of small value and small growth stocks. The small-value and small-growth portfolios are two of the six elementary portfolios constructed by Davis et al. (2000). DEF is the default yield spread in percentage points between the log yield on Moody's BAA and AAA bonds.

| Variable | Mean | Median | Std. | Min | Max |
|---------------------|---------|--------|--------|--------|--------|
| r_M^e | 0.012 | 0.027 | 0.108 | -0.439 | 0.639 |
| PE | 2.923 | 2.916 | 0.384 | 1.508 | 3.910 |
| TY | 1.459 | 1.419 | 1.050 | -1.650 | 3.748 |
| VS | 1.639 | 1.509 | 0.362 | 1.180 | 2.685 |
| DEF | 1.078 | 0.847 | 0.685 | 0.324 | 5.167 |
| <i>Correlations</i> | r_M^e | PE | TY | VS | DEF |
| r_M^e | 1.000 | 0.077 | 0.051 | -0.036 | -0.168 |
| PE | 0.077 | 1.000 | -0.240 | -0.368 | -0.601 |
| TY | 0.051 | -0.240 | 1.000 | 0.321 | 0.402 |
| VS | -0.036 | -0.368 | 0.321 | 1.000 | 0.650 |
| DEF | -0.168 | -0.601 | 0.402 | 0.650 | 1.000 |

Table 2: **Unrestricted VAR Parameter Estimates, Two Variables**

The table shows the OLS parameter estimates for a first-order VAR model including a constant, the log excess market return (r_M^e) and the price-earnings ratio (PE). In the upper panel, each set of two rows corresponds to a different dependent variable. The first two columns report coefficients on the two explanatory variables, and the remaining column shows the R^2 . Standard errors are in parentheses. The middle panel of the table shows the coefficients mapping state variable shocks into news terms for both a reduced-form VAR and a structural VAR where r_M^e is ordered first and PE second. The bottom panel of the table reports the correlation matrix of the shocks with shock standard deviations on the diagonal. The sample period for the dependent variables is 1929:2-2010:4, providing 327 quarterly data points.

| <i>VAR estimate</i> | r_M^e | PE | R^2 |
|-------------------------------|-------------------|-------------------|-------|
| r_M^e | -0.011 (0.055) | -0.046 (0.016) | 0.027 |
| PE | 0.061 (0.051) | 0.963 (0.015) | 0.932 |
| Error to N_{CF} | 0.938 | -0.873 | |
| Error to $-N_{DR}$ | 0.062 | 0.873 | |
| Structural Error to N_{CF} | 0.020 | -0.036 | |
| Structural Error to $-N_{DR}$ | 0.087 | 0.036 | |
| <i>News terms corr/std</i> | N_{CF} | $-N_{DR}$ | |
| | N_{CF} | 0.133 | 0.133 |
| | $-N_{DR}$ | 0.133 | 0.094 |

Table 3: **Unrestricted VAR Parameter Estimates, Five Variables**

The table shows the OLS parameter estimates for a first-order VAR model including a constant, the log excess market return (r_M^e), the price-earnings ratio (PE), the term yield spread (TY), the small-stock value spread (VS), and the default yield spread (DEF). In the upper panel, each set of two rows corresponds to a different dependent variable. The first five columns report coefficients on the five explanatory variables, and the remaining column shows the R^2 . Standard errors are in parentheses. The middle panel of the table shows the coefficients mapping state variable shocks into news terms for both a reduced-form VAR and a structural VAR where the state variables are ordered as in the columns. The bottom panel of the table reports the correlation matrix of the shocks with shock standard deviations on the diagonal. The sample period for the dependent variables is 1929:2-2010:4, providing 327 quarterly data points.

| <i>VAR estimate</i> | r_M^e | PE | TY | VS | DEF | R^2 |
|-------------------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------|
| r_M^e | -0.019 (0.056) | -0.057 (0.019) | 0.007 (0.006) | -0.033 (0.022) | -0.004 (0.014) | 0.040 |
| PE | 0.059 (0.053) | 0.962 (0.018) | 0.006 (0.006) | -0.028 (0.020) | 0.004 (0.013) | 0.932 |
| TY | 0.114 (0.298) | 0.057 (0.102) | 0.781 (0.033) | 0.023 (0.115) | 0.201 (0.074) | 0.709 |
| VS | 0.057 (0.046) | 0.007 (0.016) | -0.004 (0.005) | 0.952 (0.018) | 0.021 (0.011) | 0.943 |
| DEF | -0.278 (0.155) | -0.023 (0.053) | 0.007 (0.017) | 0.130 (0.060) | 0.837 (0.038) | 0.817 |
| Error to N_{CF} | 0.930 | -1.018 | 0.004 | -0.150 | -0.055 | |
| Error to $-N_{DR}$ | 0.070 | 1.018 | -0.004 | 0.150 | 0.055 | |
| Structural Error to N_{CF} | 0.015 | -0.035 | 0.002 | -0.017 | -0.012 | |
| Structural Error to $-N_{DR}$ | 0.091 | 0.035 | -0.002 | 0.017 | 0.012 | |
| <i>News terms corr/std</i> | N_{CF} | $-N_{DR}$ | | | | |
| | N_{CF} | 0.043 | -0.070 | | | |
| | $-N_{DR}$ | -0.070 | 0.100 | | | |

Table 4: **Restricted VAR Parameter Estimates, Five Variables**

The table shows the restricted, continuously-updated GMM parameter estimates for a first-order VAR model including a constant, the log excess market return (r_M^e), the price-earnings ratio (PE), the term yield spread (TY), the small-stock value spread (VS), and the default yield spread (DEF) where the resulting news terms are restricted to price the cross-section of average returns as given by the six ME and BE/ME sorted portfolios from Professor Kenneth French's website. Each set of two rows corresponds to a different dependent variable. The first five columns report coefficients on the two explanatory variables, and the remaining column shows the R^2 , computed as $1 - \text{MSE}/(\text{Variance of the dependent variable})$. Standard errors are in parentheses. The second panel of the table shows the coefficients mapping state variable shocks into news terms for both a reduced-form VAR and a structural VAR where the state variables are ordered as in the columns. The third panel reports the estimated value for the risk aversion parameter gamma, the p-value for the J-test of overidentifying restrictions, and the correlation matrix of the shocks with shock standard deviations on the diagonal. The sample period for the dependent variables is 1929:2-2010:4, providing 327 quarterly data points.

| <i>VAR estimate</i> | r_M^e | PE | TY | VS | DEF | R^2 |
|-------------------------------|-------------------|-------------------|-------------------|-------------------|-------------------|--------|
| r_M^e | 0.125 (0.080) | -0.057 (0.017) | 0.012 (0.006) | -0.004 (0.022) | -0.051 (0.020) | -0.048 |
| PE | 0.215 (0.088) | 0.963 (0.015) | 0.011 (0.006) | 0.001 (0.023) | -0.036 (0.025) | 0.927 |
| TY | -0.063 (0.286) | 0.035 (0.097) | 0.783 (0.047) | 0.010 (0.110) | 0.183 (0.082) | 0.713 |
| VS | -0.040 (0.067) | 0.017 (0.018) | -0.004 (0.005) | 0.939 (0.017) | 0.042 (0.011) | 0.941 |
| DEF | -0.837 (0.359) | -0.009 (0.032) | -0.015 (0.016) | 0.040 (0.044) | 0.983 (0.062) | 0.793 |
| Error to N_{CF} | 1.604 | -1.811 | 0.075 | -0.589 | -0.934 | |
| Error to $-N_{DR}$ | -0.604 | 1.811 | -0.075 | 0.589 | 0.934 | |
| Structural Error to N_{CF} | 0.173 | 0.023 | 0.038 | -0.135 | -0.206 | |
| Structural Error to $-N_{DR}$ | -0.062 | -0.023 | -0.038 | 0.135 | 0.206 | |
| <i>News terms corr/std</i> | N_{CF} | $-N_{DR}$ | | Gamma | 4.993 | |
| | N_{CF} | -0.935 | | | | |
| | $-N_{DR}$ | 0.258 | J-stat p-value | | 0.044 | |

Figure Descriptions

Figure 1: This figure plots the cash-flow news and the negative of discount-rate news, smoothed with a trailing exponentially-weighted moving average. Each row of graphs plots news terms extracted from the corresponding VARs estimated in Tables 2 through 4 respectively. The decay parameter is set to 0.08 per quarter, and the smoothed news series are generated as $MA_t(N)=0.08N_t+(1-0.08)MA_{t-1}(N)$. The sample period is 1929:2-2010:4.


Figure 2: This figure plots the cash-flow news and the negative of discount-rate news, smoothed with a trailing exponentially-weighted moving average. Each row of graphs plots news terms extracted from the corresponding VARs estimated in Tables 2 through 4 respectively. The decay parameter is set to 0.08 per quarter, and the smoothed news series are generated as $MA_t(N)=0.08N_t+(1-0.08)MA_{t-1}(N)$. The sample period is 1995:1-2010:4.

Figure 3: This figure plots the 1993-2010 out-of-sample equity premium forecasting performance of the VARs estimated as in Tables 2 through 4. The models are estimated on an expanding window basis and then used to predict quarterly excess log returns on the CRSP value-weight index. For comparison, we also plot realized excess equity returns, rescaled to have the same standard deviation as the VAR forecast from Table 4.






Figure 4: This figure plots the evolution in event time for the key aggregate variables in our analysis for both the tech boom of 1997-2002 and the credit boom of 2005-2010. The event for each period is the market peak (tech: 2000Q1, credit: 2007Q3). The variables we plot include the excess return on the market, the small stock value spread (VS), the default yield spread (DEF), the smoothed price-earnings ratio (PE), the market's smoothed earnings (E), the smoothed price-to-dividend ratio (P/D), and the market's smoothed dividends (E).

Figure 5: This figure plots over 1995-2010 the cumulative log stock return (top panel), the seasonally adjusted logs of industrial production, vehicle miles travelled in the US, and the Ceridian-UCLA Pulse of Commerce Index based on diesel fuel consumption in the US trucking industry (second panel), the logs of business and nonbusiness bankruptcies (third panel), and the cash-flow and discount-rate news series implied by our restricted five-variable VAR system (fourth and fifth panels).


Smoothed Ncf








Smoothed -Ndr



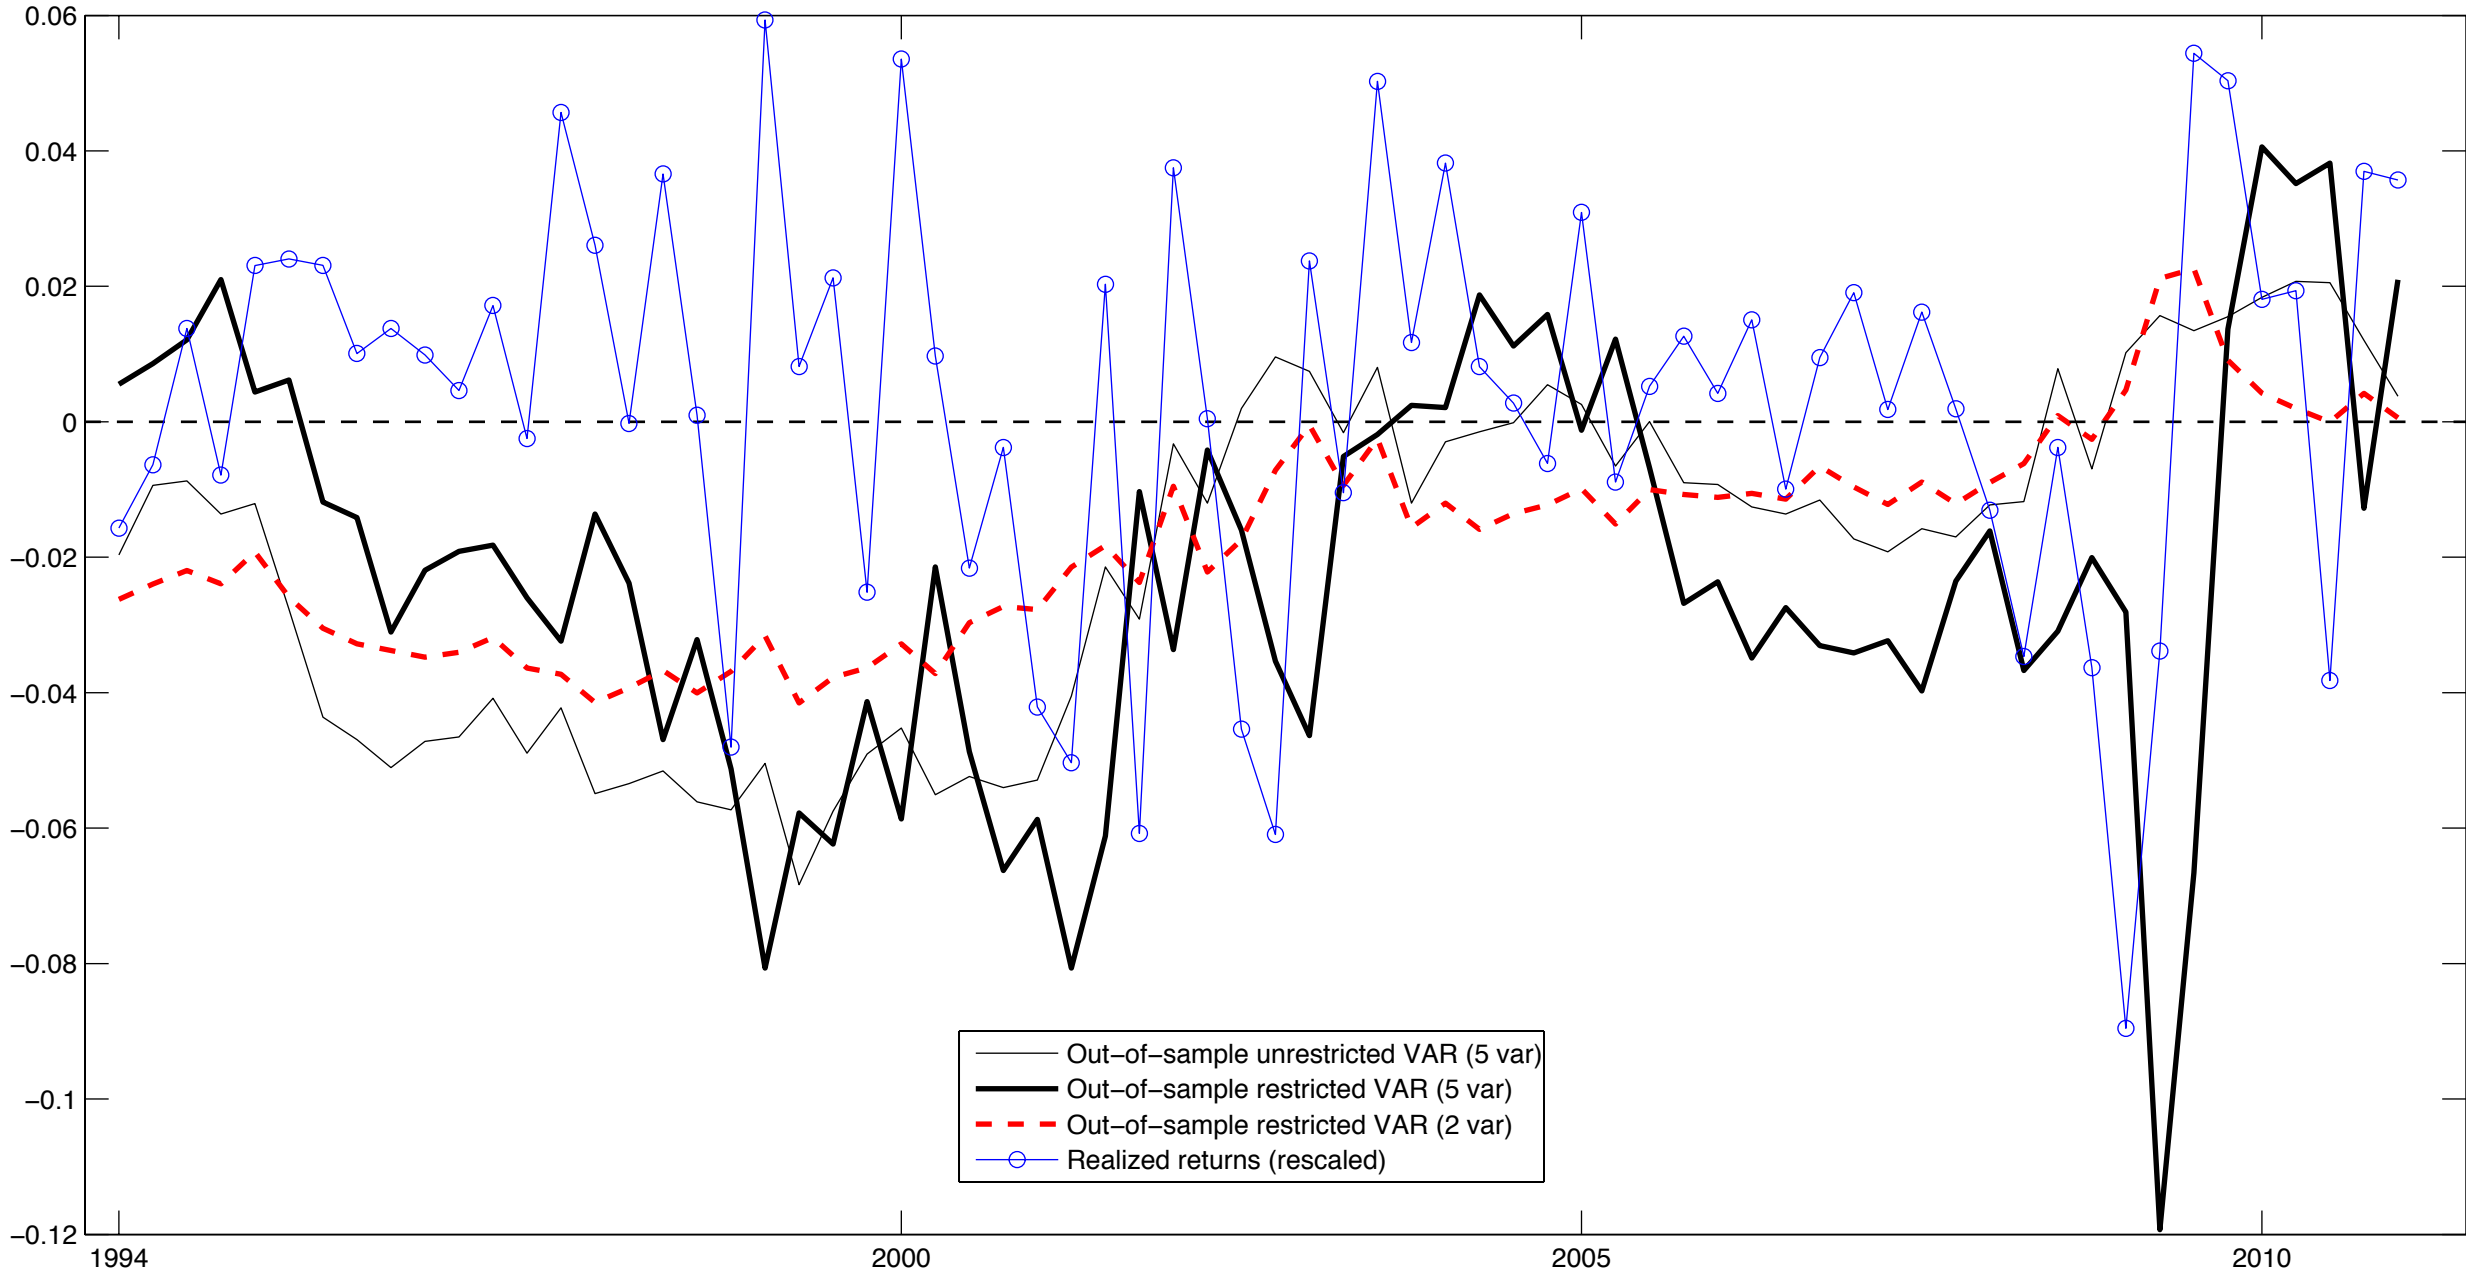
Smoothed Ncf





Smoothed -Ndr



Quarterly return forecast out of sample, 1994–2010





— Tech episode (1997–2002): peak 2000q1
 - - - Credit episode (2005–2010): peak 2007q3

