Hard Times

John Y. Campbell
Department of Economics, Littauer Center, Harvard University and NBER

Stefano Giglio
Booth School of Business, University of Chicago and NBER

Christopher Polk
Department of Finance, London School of Economics

We show that the stock market downturns of 2000–2002 and 2007–2009 have very different proximate causes. The early 2000s saw a large increase in the discount rates applied to profits by rational investors, while the late 2000s saw a decrease in rational expectations of future profits. We reach these conclusions by using a VAR model of aggregate stock returns and valuations, estimated both without restrictions and imposing the cross-sectional restrictions of the intertemporal capital asset pricing model (ICAPM). Our findings imply that the 2007–2009 downturn was particularly serious for rational long-term investors, whose losses were not offset by improving stock return forecasts as in the previous recession. (*JEL* G12, N22)

During the past 15 years, the U.S. 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, and how do they compare with the fluctuations of stock market prices we observed in the last century? 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? Alternatively, were stock prices instead driven up by declines in the discount rates that rational investors applied to corporate cash flows? Then, when the booms ended, did prices fall because rational investors became pessimistic about profits, or because they discounted future profits more heavily?[1]

Answers to these questions are important for several reasons. The answers 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 2000-2002 and 2007-2009 were due to lower expected corporate profits, then those conditions 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.

The nature of a stock market downturn also determines the optimal consumption response of a long-term investor. A long-term investor with an elasticity of intertemporal substitution (EIS) of 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 an EIS greater than one, who should save aggressively when discount rates increase.

In this paper we argue that the downturns of 2000–2002 and 2007–2009 have very different proximate causes. In 2000–2002, stock prices fell primarily because discount rates increased, while in 2007–2009 cash flow prospects worsened, with discount rates playing only a minor role until late 2008. Similarly, the preceding booms were driven primarily by discount rates in the 1990s and by a mix of cash flows and discount rates in the mid-2000s. Looking back to the history of booms and busts in the U.S. since 1929, we find only a few other episodes driven mainly by cash flow news, namely the onset of the Great Depression and the recession of 1937–1938. These, like the current crisis, were particularly hard times. Most other episodes, instead, were driven mainly by discount rate news (with or without a delayed response of cash flow news), with much less severe consequences for a long-term investor.

We reach these conclusions 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;

---

[1] An 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.

second, the approximate accounting identity of Campbell and Shiller (1988a); 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 novel contribution is to estimate the aggregate VAR jointly with 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 the ICAPM restrictions as additional moments in a generalized method of moments (GMM) estimation of the VAR system 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.[2] A formal test based on the standard Diebold and Mariano (1995) out-of-sample test statistic, adjusted using the methodology of Clark and West (2006, 2007), confirms that our novel approach of imposing ICAPM restrictions is indeed useful in improving out-of-sample predictability.

Of course, the VAR methodology used in the above tests relies on specific assumptions about the data-generating process. However, our robustness tests indicate that our main conclusions are relatively insensitive to most aspects of the estimation methodology. Moreover, we show that our use of cross-sectional theory appears to discipline the analysis and reduce the danger of overfitting the data as our main conclusions are relatively insensitive to the particular variables included in the VAR. Furthermore, although we argue that our VAR implementation is reasonable, we also show that our findings about the proximate causes of the 2000–2002 and 2007–2009 downturns are consistent with a much simpler, less theoretically structured analysis of financial and macroeconomic data.

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 do better than growth stocks on average, but underperform during those stock market downturns that are permanent, in the sense that they reflect expectations of lower

---

[2] Campbell 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.

corporate profits in the future. Our empirical success confirms that this economic model was a useful description of the recent U.S. stock market experience. These results are particularly interesting as much if not all of this recent experience is subsequent to the samples of Campbell (1996), Campbell and Vuoltenaho (2004), and Campbell, Polk, and Vuolteenaho (2010).

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. Although 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 restricted to be consistent with cross-sectional pricing. In recent work, Kelly and Pruitt (2011) propose a statistical methodology to aggregate the cross-section of valuation ratios to improve the prediction of market returns. However, they do not impose theory-motivated restrictions in the estimation.

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 perform well out of sample.

Our final contribution is to expand the set of variables included in the Campbell and Vuolteenaho (2004) 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 (that imposes ICAPM conditions in the estimation) chooses to include the default spread as an important component of aggregate cash-flow news.[3] Interestingly, although the key variable of Campbell and Vuolteenaho, the small-stock value spread, continues to be an important component of market

---

[3] Previous research including Fama and French (1989) focuses on the ability of the default spread to forecast the equity premium, ignoring the implications of the estimated data-generating process for the relation between shocks to the default spread and aggregate cash-flow news. In specifications that include only the term spread, Fama and French find that the default spread forecasts aggregate stock returns with a positive sign. In richer specifications that include other variables such as the price-earnings ratio, our joint estimation procedure finds that the partial regression coefficient forecasting returns is negative, implying an intuitive negative relation between aggregate cash-flow news and shocks to the default yield.

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). Ranish also argues that cash-flow news was relatively important in the downturn of 2007–2009, but he 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 2000–2002 downturn was almost entirely driven by transitory shocks, while the 2007–2009 downturn reflected both permanent and transitory shocks.

Our paper also relates to the macroeconomics literature that studies the nature and origins of business cycles. At least since Burns and Mitchell (1946), economists have tried to understand whether all U.S. business cycles can be attributed to a common set of shocks, or whether different cycles are driven by different shocks. While Burns and Mitchell focus on graphing the behavior of different variables during the cycle, Blanchard and Watson (1986) and Claessens et al. (2008) use a statistical approach to summarize the cyclical behavior of the variables studied. We provide a counterpart to both exercises in our paper, for stock market cycles rather than business cycles. In Section 5 we conduct a graphical analysis of the history of booms and busts in the stock market during the twentieth century. In Sections 1–4, we decompose, using econometric methods, the main stock market cycles into their cash-flow and discount-rate components, guided in this by the insights of intertemporal asset pricing theory.

The remainder of the paper is organized as follows. Section 1 explains our methodology for identifying the components of stock returns. Section 2 discusses the data and our econometric methods. Section 3 presents our VAR estimates, both with and without ICAPM restrictions. Section 4 contrasts the two boom-bust cycles of the late 1990's and early 2000's and the mid to late 2000's. Section 5 compares these cycles with other stock market fluctuations that have occurred since 1929. Section 6 concludes. An Online Appendix (Campbell, Giglio, and Polk 2012) presents various robustness exercises, which we summarize in the text.

1. Identifying the Components of Stock Returns

1.1 Cash-flow and discount-rate shocks

Campbell and Shiller (1988a) 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, $\overline{(d_t - p_t)}$, 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 $\rho \equiv 1/(1+ \exp(\overline{d_t - p_t}))$ 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 $\lim_{j \to \infty} \rho^j(d_{t+j} - p_{t+j}) = 0$, taking expectations, and subtracting the current dividend, one gets:

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

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 log-linear present-value approach to obtain a decomposition of returns. Substituting (1) into the approximate return equation gives:

$$r_{t+1} - E_t r_{t+1} = (E_{t+1} - E_t) \sum_{j=0}^{\infty} \rho^j \Delta d_{t+1+j} - (E_{t+1} - E_t) \sum_{j=1}^{\infty} \rho^j r_{t+1+j},$$
$$= N_{CF, t+1} - N_{DR, t+1}, \qquad (2)$$

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.

1.2 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 the realization of r_{t+1} and Equation (2) to back out 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 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:

$$z_{t+1} = a + \Gamma z_t + u_{t+1}, \tag{3}$$

where z_{t+1} is a m-by-1 state vector with r_{t+1} as its first element, a and Γ are an m-by-1 vector and an 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:

---

[4] Chen and Zhao (2009) discuss the sensitivity of the 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 of whether to forecast returns or cash flows.

$$N_{DR,\,t+1} = e1'\lambda u_{t+1},$$
$$N_{CF,\,t+1} = (e1'+e1'\lambda)u_{t+1}. \tag{4}$$

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 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 $(I - \rho\Gamma)^{-1}$.

1.3 Imposing the ICAPM

Campbell (1993) derives an approximate discrete-time version of Merton's (1973) ICAPM. 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 (1993) 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 homoscedastic. 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 tackle in separate work (Campbell et al. 2011).

Campbell (1993) 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:

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

8

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 the 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 risk-less 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:

$$
E_t[r_{i, t+1}] - r_{f, t+1} + \frac{\sigma_{i, t}^2}{2} = \gamma \sigma_{p, t}^2 \beta_{i, CF_p, t} + \sigma_{p, t}^2 \beta_{i, DR_p, t}. \tag{6}
$$

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

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

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

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

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

$$
\beta_{i, p} = \beta_{i, CF} + \beta_{i, DR}. \tag{9}
$$

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.

2. 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.

2.1 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-earnings 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 ensure that all components of the time-t price-earnings ratio are contemporaneously observable by time t. The ratio is log transformed. In the Appendix we explore alternative ways to construct PE and using the price-dividend ratio instead (Tables A8, A16, A17, and A18).

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

Table 1
Descriptive statistics

Variable	Mean	Median	Std.	Min	Max
R_m	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
correlations	R_m	PE	TY	VS	DEF
R_m	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

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 is the excess log return on the CRSP value-weighted 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 U.S. 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.

the 10-year U.S. Constant Maturity Bond (*IGUSA10D*) and the log yield on the 3-month U.S. Treasury bill (*ITUSA3D*).

Fourth, the small-stock value spread (*VS*) is constructed from the data on the six "elementary" equity portfolios made available by 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

11

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.

2.2 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 price a parsimonious cross-section to ensure that the numerical estimation is manageable and that test asset portfolios are reasonably diversified in the early part of the sample. 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. We also impose, separately, that the model matches the equity premium exactly.

2.3 Estimation methodology

This section details the estimation technique we use for the restricted model that jointly imposes time-series and cross-sectional orthogonality conditions in a GMM estimation. We use Hansen, Heaton, and Yaron's (1996) continuously updated (CUE) GMM as Newey and Smith (2004) highlight the finite-sample advantages of this method and other generalized empirical likelihood estimators over standard GMM.

Bansal et al. (2012) have also recently used a version of this GMM estimator. However, they employ a simplified version of the CUE GMM estimator, in which the covariance of the test-asset returns with the news terms is taken as known (as opposed to being estimated with error) and the weighting matrix is obtained by discarding the off-diagonal elements of the

variance-covariance matrix of moment residuals. In this paper, by contrast, we employ the fully correctly specified CUE GMM. We preserve the full information content of the variance-covariance matrix of moment residuals and we take into account all terms that need to be estimated. The price of this is that our estimation method is numerically more involved, and requires a few additional restrictions to reduce the instability of the estimates, as described below.

We use the notation K for the dimension of the VAR and I for the number of test assets. The restricted model gives us $R = K(K+1)+I$ orthogonality conditions. $K(K+1)$ of these estimate the intercepts and dynamic coefficients of the VAR, and I orthogonality conditions are imposed by the ICAPM on the test assets. There is one free parameter in the ICAPM, the coefficient of relative risk aversion, so there are $I - 1$ 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. First, we substitute the market for portfolio p in Equation (5), obtaining:

$$E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_{i,t}^2}{2} = \gamma \text{Cov}_t(r_{i,t+1}, r_{m,t+1} - E_t r_{m,t+1}) \qquad (10)$$
$$+ (1 - \gamma)\text{Cov}_t(r_{i,t+1}, -N_{m,DR,t+1}).$$

Then, we use the aggregate VAR (which contains the market return) to rewrite:

$$r_{m,t+1} - E_t r_{m,t+1} = e1'(z_{t+1} - a - \Gamma z_t),$$
$$-N_{m,DR,t+1} = -N_{DR,t+1} = -e1'\lambda(z_{t+1} - a - \Gamma z_t),$$

so that we obtain:

$$E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_{i,t}^2}{2} = \gamma \text{Cov}_t(r_{i,t+1}, e1'(z_{t+1} - a - \Gamma z_t)) \qquad (11)$$
$$+ (1 - \gamma)\text{Cov}_t(r_{i,t+1}, -e1'\lambda(z_{t+1} - a - \Gamma z_t)).$$

Finally, given lognormality, a first-order linear approximation around $E_t R_{i,t+1} = 1$ and $R_{f,t+1} = 1$ results in[5]:

$$E_t[r_{i,t+1}] - r_{f,t+1} + \frac{\sigma_{i,t}^2}{2} \simeq E_t(R_{i,t+1}) - R_{f,t+1}.$$

---

[5] In particular, lognormality implies $E_t[r_{i,t+1}] + \frac{\sigma_{i,t}^2}{2} = ln[E_t(R_{i,t+1})]$, and, combining it with the risk free rate, $E_t[r_{i,t+1}] + \frac{\sigma_{i,t}^2}{2} - r_{f,t+1} = ln(E_t R_{i,t+1}) - ln(R_{f,t+1})$. For $E_t R_{i,t+1}$ and $R_{f,t+1}$ close to 1, a first-order approximation gives $ln(E_t R_{i,t+1}) - ln(R_{f,t+1}) \simeq [E_t, (R_{i,t+1}) - 1] - [R_{f,t+1} - 1] = E_t, (R_{i,t+1}) - R_{f,t+1}$.

13

Therefore, we can rewrite the asset pricing equation as:

$$
\begin{aligned}
\mathrm{E}_t(R_{i,t+1} - R_{f,t+1}) = {} & \gamma \mathrm{Cov}_t(r_{i,t+1}, e1'(z_{t+1} - a - \Gamma z_t)) \\
& + (1 - \gamma)\mathrm{Cov}_t(r_{i,t+1}, -e1'\lambda(z_{t+1} - a - \Gamma z_t)).
\end{aligned}
\tag{12}
$$

We can condition down using the fact that $\mathrm{E}_t[u_{t+1}] = 0$, so that we obtain:

$$
\begin{aligned}
\mathrm{E}(R_{i,t+1} - R_{f,t+1}) = {} & \gamma \mathrm{E}(r_{i,t+1} e1'(z_{t+1} - a - \Gamma z_t)) \\
& - (1 - \gamma)\mathrm{E}(r_{i,t+1} e1'\lambda(z_{t+1} - a - \Gamma z_t)).
\end{aligned}
\tag{13}
$$

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 risk-free rate:

$$
\mathrm{E}[R_{i,t+1} - R_{m,t+1} - (r_{i,t+1} - r_{m,t+1})e1'(\gamma - (1-\gamma)\rho\Gamma(I - \rho\Gamma)^{-1})(z_{t+1} - a - \Gamma z_t)] = 0.
\tag{14}
$$

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 impose some restrictions to our continuously updated GMM estimation procedure in order to achieve convergence to acceptable parameter values for every subsample, a property that we need for out-of-sample analysis.

First, we impose that our model matches the equity premium (i.e., we use Equation (13) for $i = Market$ as a constraint in our estimation) as opposed to a moment condition. In theory, we could add this equation as an additional moment condition in the GMM estimation. However, we find that, as the Appendix reports, when we do so the estimator gives a low enough weight to the market equity premium condition that we obtain an unreasonably large predicted value for the equity premium. To prevent this from occurring, in our baseline estimate we impose that Equation (13) is matched exactly for $i = Market$, and report the case where we add it to the moment conditions in Appendix Table A5.

We also place an upper bound on the risk aversion coefficient γ at 15. This results in an estimate that in some subsamples actually hits the bound, while in others it converges below it. As reported in Appendix Tables A3 and A4, imposing larger bounds yields similar results, although when the bound is very large the estimate for γ tends to hit the bound. This makes it difficult to choose a particular value for the bound. As our baseline case, then, we choose a relatively low bound that yields a reasonable value for risk aversion, and such that the estimate lies below the bound. We also impose a lower bound on γ of one, which is never binding.

Finally, 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 Appendix Table A1. Again, results appear to be robust to this choice.

The GMM problem for the restricted model can then be written as follows:

$$min\left(\frac{1}{T}\sum_t g_t(a,\Gamma,\gamma)\right)' V_T(a,\Gamma,\gamma)^{-1}\left(\frac{1}{T}\sum_t g_t(a,\Gamma,\gamma)\right),$$

$$s.t. maxeig\,\Gamma \le \bar{\lambda},$$

$$1 \le \gamma \le 15,$$

$$E[R_{m,t+1} - R_{f,t+1} - r_{m,t+1}e1'(\gamma - (1-\gamma)\rho\Gamma(I-\rho\Gamma)^{-1})(z_{t+1} - a - \Gamma z_t)] = 0,$$

where the vector g_t includes the $K(K+1)$ orthogonality conditions for the VAR plus the cross-sectional conditions for the I size/book-to-market sorted portfolios, and V_T is the continuously updated variance-covariance matrix of the residuals g_t.

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 the constraint on γ, and 1e-4 on the other constraints.

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

3. Alternative VAR Estimates

We now present the estimates of three alternative VAR systems. For comparison with previous work, we begin with a simple two-variable VAR without restrictions, including only the market excess return and log price-earnings ratio as state variables. Then we include all five state variables, first without restrictions and then imposing the restrictions of the ICAPM described in the previous section.

3.1 Two-variable VAR system

Table 2 reports results that are familiar from previous research using this methodology. The table 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

Table 2
Unrestricted VAR estimate, 2 variables

VAR estimate		R_m		PE	R^2
R_m		−0.011		−0.046	0.027
		(0.055)		(0.016)	
PE		0.061		0.963	0.932
		(0.051)		(0.015)	
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	

News terms corr/std		N_{CF}		$-N_{DR}$	γ	6.557
N_{CF}		0.041		0.133		
-N_{DR}		0.133		0.094		

Betas		Small			Large		
		Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow		0.262	0.227	0.230	0.181	0.173	0.203
Discount Rate		1.079	1.004	1.109	0.775	0.828	0.985
				$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted		0.009	0.005	0.007	−0.001	−0.001	0.003	0.023
Realized		0.006	0.015	0.022	−0.002	0.002	0.011	0.018
Error		−0.002	0.010	0.015	−0.001	0.003	0.008	−0.005

Betas (early sample)		Small			Large		
		Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow		0.228	0.231	0.251	0.166	0.177	0.243
Discount Rate		1.088	1.096	1.273	0.775	0.939	1.163
				$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted		0.009	0.010	0.015	−0.003	0.001	0.012	0.034
Realized		0.013	0.016	0.026	−0.002	0.004	0.017	0.025
Error		0.004	0.007	0.011	0.001	0.003	0.005	−0.009

Betas (late sample)		Small			Large		
		Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow		0.317	0.225	0.200	0.208	0.170	0.144
Discount Rate		1.073	0.859	0.850	0.776	0.655	0.709
				$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted		0.008	0.002	0.001	0.001	−0.002	−0.003	0.016
Realized		0.002	0.014	0.019	−0.001	0.001	0.006	0.013
Error		−0.007	0.012	0.018	−0.001	0.003	0.009	−0.003

The table shows the results obtained with a first-order VAR model including a constant, the log excess market return *(R_m)* and the price-earnings ratio *(PE)*. The upper panel reports the estimates of the transition matrix of the VAR (standard errors in parentheses) and the R^2 of each regression. It also reports the coefficients mapping state variable shocks into news terms for both a reduced-form VAR and a structural VAR where R_m is ordered first and *PE* second. Finally, the upper panel reports the correlation matrix of the shocks with shock standard deviations on the diagonal, and the risk aversion parameter γ implied by the ICAPM model estimated as in Campbell and Vuolteenaho (2004) using the six size/book-to-market sorted portfolios. The lower panel reports cash-flow and discount-rate news betas for the six portfolios, the predicted and realized mean return of each portfolio in excess of the market, as well as the equity premium. Betas and excess returns are reported for the full sample (1929:2–2010:4), as well as for the early (1929:2–1963:2) and late (1963:3–2010:4) subsamples.

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 next in 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. 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.

Table 2 next computes the cash-flow and discount-rate betas of the six *ME* and *BE*/*ME* sorted portfolios, for the full sample, and separately for the early (1929:2–1963:2) and late (1963:3–2010:4) samples. It also reports an estimate of the risk-aversion parameter that best fits the cross-sectional asset pricing equations, and the predicted and realized mean returns of each test asset obtained using that estimate. The next section discusses these results and compares them with those obtained using the 5-variable VAR.

We have explored what happens when we impose the restrictions of the ICAPM via GMM 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 2.1, 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

17

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.

3.2 Unrestricted five-variable VAR system

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. However, although some of the partial regression coefficients in the return regression are statistically insignificant, we can reject the null hypothesis that all five coefficients are jointly equal to zero.

We find that discount-rate news is considerably more volatile than cash-flow 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.

Following Table 2, Table 3 reports the cash-flow and discount-rate betas of the six ME and BE/ME sorted portfolios. We find here the pattern pointed out by Campbell and Vuolteenaho (2004). In the early period, value stocks have both higher cash-flow beta and higher discount-rate beta, and therefore are overall riskier than growth stocks. In the modern sample, they have (slightly) higher cash-flow betas, but noticeably lower discount-rate betas. These facts account for the failure of the CAPM in the modern sample. We can contrast these results with those in Table 2, which shows that the two-variable VAR is not rich enough to capture such patterns in cash flow and discount rate betas across portfolios in the modern sample.

We also compute the implied risk aversion parameter γ obtained by regressing cross-sectionally the six test assets and the market return on their respective betas, imposing the constraint that the risk-free rate is the zero-beta rate and that the risk premium on discount rate news is the variance of the market return (as predicted by the ICAPM). We obtain an estimate of γ just above 8, slightly higher than that obtained using the two-variable VAR.

Table 3
Unrestricted VAR estimate, 5 variables

VAR estimate	R_m	PE	TY	VS	DEF	R^2
R_m	−0.019	−0.057	0.007	−0.033	−0.004	0.040
	(0.056)	(0.019)	(0.006)	(0.022)	(0.014)	
PE	0.059	0.962	0.006	−0.028	0.004	0.933
	(0.053)	(0.018)	(0.006)	(0.020)	(0.013)	
TY	0.114	0.057	0.781	0.023	0.201	0.709
	(0.298)	(0.102)	(0.033)	(0.115)	(0.074)	
VS	0.057	0.007	−0.004	0.952	0.021	0.943
	(0.046)	(0.016)	(0.005)	(0.018)	(0.011)	
DEF	−0.278	−0.023	0.007	0.130	0.837	0.817
	(0.155)	(0.053)	(0.017)	(0.060)	(0.038)	
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	

News terms corr/std	N_{CF}		$-N_{DR}$		γ	8.061
N_{CF}	0.043		−0.070			
$-N_{DR}$	−0.070		0.100			

Betas	Small			Large		
	Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow	0.177	0.193	0.230	0.109	0.151	0.208
Discount Rate	1.167	1.041	1.111	0.845	0.849	0.985
			$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted	0.007	0.007	0.011	−0.003	0.001	0.008	0.022
Realized	0.006	0.015	0.022	−0.002	0.002	0.011	0.018
Error	−0.001	0.007	0.010	0.001	0.001	0.003	−0.004

Betas (early sample)	Small			Large		
	Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow	0.251	0.271	0.323	0.163	0.210	0.308
Discount Rate	1.082	1.075	1.220	0.782	0.915	1.114
			$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted	0.011	0.014	0.023	−0.005	0.003	0.020	0.039
Realized	0.013	0.016	0.026	−0.002	0.004	0.017	0.025
Error	0.002	0.002	0.003	0.003	0.000	−0.002	−0.015

Betas (late sample)	Small			Large		
	Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow	0.064	0.077	0.091	0.028	0.062	0.056
Discount Rate	1.297	0.987	0.940	0.939	0.747	0.786
			$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted	0.004	0.003	0.003	−0.001	0.000	0.000	0.010
Realized	0.002	0.014	0.019	−0.001	0.001	0.006	0.013
Error	−0.003	0.011	0.016	0.000	0.001	0.006	0.003

The table shows the results obtained with a first-order VAR model including a constant, the log excess market return (R_m), the price-earnings ratio (PE), the term yield spread (TY), the small-stock value spread (VS), and the default yield spread (DEF). The upper panel reports the estimates of the transition matrix of the VAR (standard errors in parentheses) and the R^2 of each regression. It also reports the coefficients mapping state variable shocks into news terms for both a reduced-form VAR and a structural VAR where R_m is ordered first and PE second. Finally, the upper panel reports the correlation matrix of the shocks with shock standard deviations on the diagonal, and the risk aversion parameter γ implied by the ICAPM model estimated as in Campbell and Vuolteenaho (2004) using the six size/book-to-market sorted portfolios. The lower panel reports cash-flow and discount-rate news betas for the six portfolios, the predicted and realized mean return of each portfolio in excess of the market, as well as the equity premium. Betas and excess returns are reported for the full sample (1929:2–2010:4), as well as for the early (1929:2–1963:2) and late (1963:3–2010:4) subsamples.

19

Table 3 next shows the predicted and realized excess mean returns of each test asset in excess of the market return, by analogy with Equation (14). The predicted return is obtained by multiplying the cash-flow and discount-rate betas presented above with the risk premia estimated in the restricted OLS regression. We can see that in the full sample, the equity premium is over-predicted slightly (2.2% per quarter as opposed to the realized value of 1.8%), but the cross-section of value and growth stocks is fit relatively well (and better than using the two-variable VAR).

3.3 Restricted five-variable VAR system

Table 4 reports the restricted five-variable VAR. The restrictions strengthen the predictive power of the default spread for the market return, and weaken the predictive power of the value spread. *DEF* forecasts the log excess market return, controlling for the other four variables including the *PE* ratio, with a coefficient of -0.034 whose standard error is 0.024. This result points towards a role for the default spread in capturing a worsening in earning prospects. This is particularly interesting when compared with previous research (e.g., Fama and French 1989) that finds a positive univariate relation between the default spread and expected returns. The two news terms are now estimated to be more volatile, and 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.

We estimate the coefficient of risk aversion to be around 10. This estimate is arguably within a reasonable range. In terms of overall model fit, we find that the overidentifying restrictions of the ICAPM are rejected at the 1% significance level. While this rejection indicates that the model we impose is unable to meet the serious challenge of fully capturing both the time-series and the cross-sectional dimensions of the data, we argue that useful information can be learned by studying the properties of the time series that best fit the cross-sectional patterns, through the lens of the economic model. In fact, as we show in Section 4.3, the restricted model performs very well in predicting returns out of sample: A formal statistical test confirms that imposing ICAPM restrictions improves out-of-sample predictability. Thus, while statistically rejected in sample, the economic restrictions imposed through the ICAPM contain valuable information about asset prices.

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 -0.28, compared to the unrestricted estimate of -0.06. For the structural mapping, the coefficient on the default spread is -0.06 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

Table 4
Restricted VAR estimate, 5 variables

VAR estimate	R_m	PE	TY	VS	DEF	R^2
R_m	0.100	−0.037	0.005	0.003	−0.034	−0.020
	(0.084)	(0.017)	(0.007)	(0.025)	(0.024)	
PE	0.186	0.990	0.002	0.007	−0.015	0.929
	(0.087)	(0.015)	(0.006)	(0.024)	(0.023)	
TY	−0.108	0.037	0.801	0.002	0.191	0.713
	(0.301)	(0.096)	(0.048)	(0.111)	(0.087)	
VS	0.001	0.024	−0.007	0.944	0.040	0.942
	(0.066)	(0.017)	(0.005)	(0.017)	(0.013)	
DEF	−0.639	−0.021	−0.013	0.039	0.961	0.803
	(0.358)	(0.037)	(0.017)	(0.061)	(0.080)	
Error to N_{CF}	0.985	−1.542	0.037	−0.288	−0.279	
Error to $-N_{DR}$	0.015	1.542	−0.037	0.288	0.279	
Structural Error to N_{CF}	0.010	−0.035	0.020	−0.050	−0.062	
Structural Error to $-N_{DR}$	0.100	0.035	−0.020	0.050	0.062	

News terms corr/std	N_{CF}	$-N_{DR}$	γ	10.136
N_{CF}	0.089	−0.577		
$-N_{DR}$	−0.577	0.133	J-stat p-value	0.002

Betas	Small			Large		
	Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow	0.069	0.180	0.280	−0.019	0.133	0.263
Discount Rate	1.265	1.038	1.043	0.947	0.837	0.900
			$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted	0.006	0.017	0.029	−0.008	0.009	0.026	0.017
Realized	0.006	0.015	0.022	−0.002	0.002	0.011	0.018
Error	0.000	−0.002	−0.008	0.007	−0.007	−0.015	0.001

Betas (early sample)	Small			Large		
	Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow	0.373	0.433	0.547	0.190	0.337	0.515
Discount Rate	0.905	0.851	0.930	0.695	0.713	0.826
			$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted	0.025	0.035	0.057	−0.013	0.014	0.049	0.061
Realized	0.013	0.016	0.026	−0.002	0.004	0.017	0.025
Error	−0.012	−0.018	−0.031	0.011	−0.011	−0.032	−0.036

Betas (late sample)	Small			Large		
	Growth	Neutral	Value	Growth	Neutral	Value
Cash Flow	−0.450	−0.249	−0.172	−0.375	−0.213	−0.163
Discount Rate	1.878	1.354	1.232	1.375	1.046	1.025
			$E[R_i\text{-}R_m]$			$E[R_m\text{-}R_f]$
Predicted	−0.007	0.005	0.010	−0.005	0.005	0.009	−0.014
Realized	0.002	0.014	0.019	−0.001	0.001	0.006	0.013
Error	0.009	0.009	0.009	0.004	−0.004	−0.003	0.027

The table shows the results obtained with a first-order VAR model including a constant, the log excess market return (R_m), the price-earnings ratio (PE), the term yield spread (TY), the small-stock value spread (VS), and the default yield spread (DEF). The upper panel reports the estimates of the transition matrix of the VAR (standard errors in parentheses) and the R^2 of each regression. It also reports the coefficients mapping state variable shocks into news terms for both a reduced-form VAR and a structural VAR where R_m is ordered first and PE second. Finally, the upper panel reports the correlation matrix of the shocks with shock standard deviations on the diagonal, and the risk aversion parameter γ implied by the ICAPM model estimated as in Campbell and Vuolteenaho (2004) using the six size/book-to-market sorted portfolios. The lower panel reports cash-flow and discount-rate news betas for the six portfolios, the predicted and realized mean return of each portfolio in excess of the market, as well as the equity premium. Betas and excess returns are reported for the full sample (1929:2–2010:4), as well as for the early (1929:2–1963:2) and late (1963:3–2010:4) subsamples.

21

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, although the key variable of Campbell and Vuolteenaho (2004), 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 Appendix Table A8, we can see that dropping *VS* from the restricted five-variable VAR has little effect on the qualitative conclusions drawn from the VAR: the news terms are (even more) strongly negatively correlated and very volatile, and the default spread similarly becomes extremely important for predicting returns. The discount-rate and cash-flow news terms appear very similar in the restricted four- and five-variable VAR systems.

The lower part of Table 4 shows that the patterns of cash-flow and discount-rate betas across the test assets appear reinforced in the restricted VAR. It is clear that imposing the ICAPM model on the test assets affects the news terms in a way that strongly emphasizes a higher exposure of value stocks to cash-flow news and a lower exposure to discount-rate news. The errors in the moment conditions are slightly larger than in the unrestricted case, and concentrated in large stocks. Note that we impose that Equation (13) holds exactly for the market portfolio. However, we compute our predicted values for the excess returns by calculating the two betas and using those together with γ to arrive at a predicted value for the equity premium (which corresponds to using the right-hand side of Equation (12)). In theory, Equations (12) and (13) are completely equivalent. However, there will be small differences between the two when the orthogonality conditions of the VAR are not matched exactly, so that the assumptions needed to condition down the covariances are not fully satisfied in the sample. In Appendix Table A4, we show that the results are essentially unchanged if we impose the restriction of the GMM estimation in terms of Equation (12) directly.

3.4 Robustness of the estimates

The Online Appendix examines the robustness of these estimates to variations in the methodology and the set of variables employed. First, in Appendix Table A1 we impose a slightly tighter bound on the maximum eigenvalue of the transition matrix Γ. Second, as discussed above, we run our estimator imposing different upper bounds on risk aversion γ. The results are reported in Appendix Tables A2 ($\gamma \leq 50$) and A3 ($\gamma \leq 200$). Third, in Appendix Table A4 we impose the restriction on the equity premium in the form of Equation

(12), i.e., using covariances, as opposed to expectations. Fourth, we impose the restriction on the equity premium as a moment condition rather than as a constraint in the estimation; the results are reported in Appendix Table A5. Fifth, we allow the cash-flow and discount-rate beta estimates of the test assets to differ across the pre-1963 and the post-1963 subperiods. In Appendix Table A6, we impose separately the moment conditions for the six test assets for the early and late sample, therefore imposing that the mean excess returns are matched separately in both subsamples. In Appendix Table A7, we also impose that the equity premium is matched separately in the two subsamples.

We next investigate how the results depend on the variables included in the VAR. We run our estimation using the log price-dividend ratio instead of the log price-earnings ratio in Appendix Table A8. Appendix Table A9 estimates the VAR without the value spread, while Appendix Table A10 drops the default spread. Appendix Table A11 uses the aggregate book-to-market ratio of nonfinancial firms.[6] The results show that the decomposition of the two stock market episodes is not due to financial firms in particular, but reflects economy-wide news. We use the book-to-market ratio of nonfinancial firms instead of the price-earnings ratio because we can then integrate Compustat data with Ken French's Historical Book Equity Data, to obtain a series that extends back to the 1920s.

Appendix Tables A12 and A13 report the unrestricted and restricted estimates obtained adding CAY, the cointegrating residual between log aggregate consumption, log assets and log labor income (Lettau and Ludvigson 2001), to the VAR. CAY is a powerful predictor of returns and as such can help us better identify discount rate news. However, CAY is only available from 1952, so that we lose the information contained in the period of the Great Depression. Appendix Tables A14 and A15 report the estimates from the baseline model obtained in the same subsample in which CAY is available.

In Appendix Tables A16 to A18 we explore different ways to construct the price-earnings ratio. In particular, in Table A16 we construct the ratio by deflating both the price and the earnings series by the CPI before taking their ratio. In Table A17, we construct the price-earnings ratio as in the preceding table, but we add inflation, exponentially smoothed with a three-year half-life, to the set of variables in the VAR. In Table A18, we use the short-term nominal rate to deflate the two series, to address concerns that the CPI is poorly measured in the early sample.

In Appendix Table A19, we run the estimation using as a first variable in the VAR the real return on the market portfolio, as opposed to the excess return. This specification undoes a modeling choice of Campbell and Vuolteenaho (2004), who forecast excess returns out of concern that the

---

[6] We thank an anonymous referee for the suggestion of looking at a valuation ratio that excludes financial firms.

CPI was poorly measured in the first part of the sample period, at the cost of combining cash-flow news with real-interest-rate news.
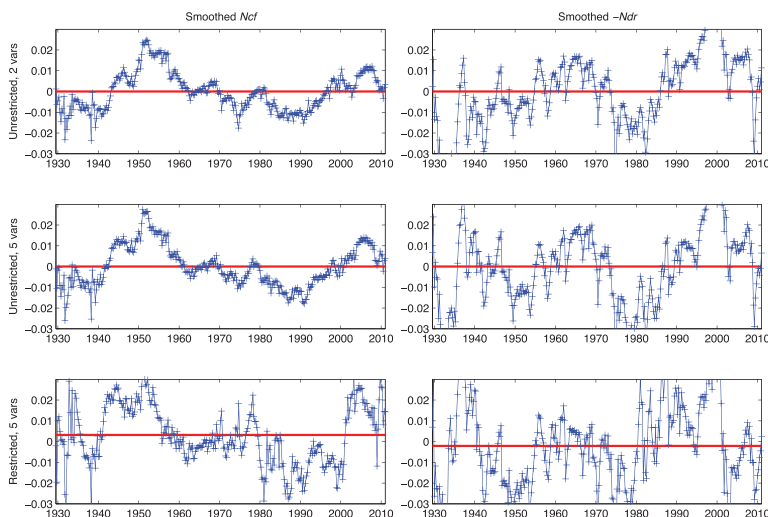
While the estimation results vary across these different specifications, all of the main qualitative patterns turn out to be very robust. First, the restricted news terms always appear more volatile than the ones implied by the unrestricted estimator and, more importantly, they always appear negatively correlated. Therefore, a robust feature of our results is that imposing the cross-sectional restrictions leads us to decompose stock market movements into *either* cash-flow *or* discount-rate news, rather than a mixture of the two. Second, across almost all of the restricted specifications, movements in the price-earnings ratio predict future returns more weakly than in the unrestricted specifications. Third, the default spread becomes in most specifications more important as a predictor of future returns, and appears, as in the baseline estimate, with a negative coefficient. Fourth, the pattern of cash-flow and discount-rate betas for value and growth stocks is robust across specifications for the modern period, just as in our baseline restricted estimates. In almost all robustness tests, value stocks have in the modern sample higher cash-flow beta and lower discount-rate beta. In fact, in some specifications this pattern holds even in the early sample, where the standard CAPM would work reasonably well. Taken together, these patterns are the main features that drive our interpretation of the results, discussed in the next section.

4. Understanding Recent U.S. Financial History

4.1 The VAR approach

What account do these alternative VAR models give of U.S. 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 approximately 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 sum 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 80-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 early 1930s. This was followed by a profit boom in the late 1940s through to the early 1960s, but discount rates remained high in this period (sentiment remained depressed), dampening the effect on stock

Figure 1
Cash-flow and discount-rate news from 1929 to 2010
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 $MA_t(N) = 0.08 N_t + (1\text{-}0.08) MA_{t-1}(N)$. The sample period is 1929:2–2010:4.

prices. Profits were weak in the 1980s and early 1990s, 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 1970s 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 1990s, 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 2007–2008. The restricted model shows a particularly strong hump shape in cash-flow news over this period. Note that since the restricted model predicts a very sharp decomposition between discount rate and cash-flow news, the large negative shock that follows the bust of the tech episode is accompanied by a positive cash flow news that compensates its effect on the return. Therefore, in the restricted model the cash-flow boom of the 2000s appears earlier than in the unrestricted model, but this is mostly due to the large negative discount rate shock of 2001.

There is less consistency across models about cash-flow news in the 1990s, which is estimated to be modestly positive in the two-variable model but not in the five-variable models. The models in Figure 2 also give different accounts of the trough of the financial crisis in the fall of 2008 and the subsequent recovery. According to the unrestricted models, the sharp decline in the price-earnings ratio in late 2008 was a piece of negative discount-rate news;

Figure 2
Cash-flow and discount-rate news from 1995 to 2010
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 $MA_t(N) = 0.08 N_t + (1 - 0.08) MA_{t-1}(N)$. The sample period is 1995:1–2010:4.

these models indicate 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. This interpretation results from the emphasis placed by the restricted model on the default spread, which widened dramatically in the fall of 2008. Between 2008 and 2009, the spike in the default spread reversed, from about 3.16% to just above 1%. Given the importance of the default spread in the restricted model, this implies that the restricted estimates of the model show a much stronger rebound in cash flow news, i.e., much stronger good news for long-term investors, after the sharp decline in 2008.

This result has also strong implications for consumption growth. Simple algebra shows that the consumption growth process in this model follows:

$$\Delta c_{t+1} - E_t \Delta c_{t+1} = (r_{t+1} - E_t r_{t+1}) - (\psi - 1) N_{DR,\, t+1},$$
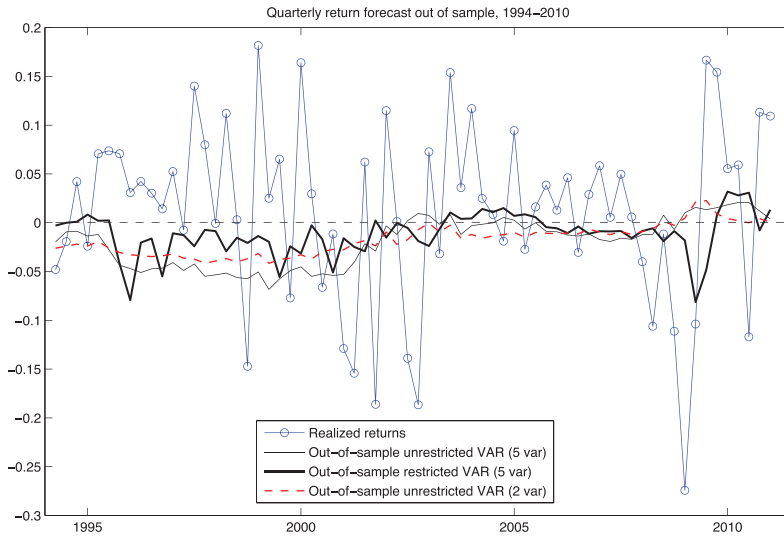
which can be plotted over time for different values of ψ (recall that ψ is not pinned down by asset prices, so we cannot estimate it using our VAR and test asset data). Appendix Figure A1 shows results for $\psi = 0.5, 1, 1.5$. In all cases, a strong consumption recovery follows the financial crisis. Given that the restricted model finds more good news coming from N_{CF} (and consequently relatively less good news from $-N_{DR}$), the rebound in consumption in 2009 is stronger under the restricted model for $\psi < 1$, weaker for $\psi > 1$, and approximately the same as the unrestricted model for $\psi = 1$. Note that the magnitude of the consumption innovations implied by the model is much larger than that of observed aggregate consumption. On the one hand, this points to a limitation of the model in capturing the consumption dynamics. On the other hand, as noted by Campbell (1993), this result could partially be explained by imperfect measurement of the consumption series, or by a divergence between the average consumption in the population and the consumption of asset-market participants, which has been shown to be more volatile (see for example Malloy, Moskowitz, and Vissing-Jorgensen 2009).

Summarizing these results, our VAR models indicate 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.

4.2 Robustness of the return decomposition

To illustrate the robustness of the main results to the specification of the VAR, Appendix Figures A2 and A3 reproduce Figure 2 of the text using different sets of variables for the VAR. Only the five-variable unrestricted and restricted results are plotted. Appendix Figure A2 compares the baseline estimate from Figure 2 with estimates obtained using the price-dividend ratio instead of PE, dropping the value spread, dropping the default spread, and using the aggregate book-to-market ratio of nonfinancial firms. The figure shows that the main features of the return decomposition into cash-flow and discount-rate news for the last two stock market cycles are similar in all cases.

Figure A3 in the Appendix plots the results obtained by adding CAY to the VAR. Once again, the unrestricted estimates confirm the main decomposition of the technology and credit boom-bust cycles. The restricted news terms, however, are very noisy. To further investigate this, we compute and plot in Appendix Figure A2 the baseline VAR (without CAY) using the same sample used to compute the results with CAY, starting in 1952. The resulting news terms, denoted "later sample," are also quite noisy. This shows that what makes the restricted estimated news terms so noisy is using the shorter

Figure 3
Out-of-sample excess return forecasts from 1994 to 2010
This figure plots the 1994–2010 out-of-sample equity premium forecasting performance of the VARs estimated in Tables 2 through 4. The models are estimated on a 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.

sample for which *CAY* is available, not the addition of *CAY* per se. We believe this is a downside of using CUE GMM, which is known to sometimes result in extreme estimates (see Hansen, Heaton, and Yaron 1996), even though it has generally good small-sample properties.

4.3 Out-of-sample return forecasts

Another way to understand the differences between the technology and credit boom-bust cycles 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 and the realized returns.

The striking pattern in Figure 3 is that return forecasts were considerably lower during the boom of the late 1990s than they were during the boom of the mid-2000s, and they increased more strongly and rapidly during the downturn following the year 2000 than they did in 2007–2008. 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 2000s was mitigated, for

long-term investors, by an increase in expected future stock returns. This was much less the case in 2007–2008.

To compare the out-of-sample performance of our restricted and unrestricted models, we employ the methodology of Clark and West (2006, 2007). Our restricted model is nested in the unrestricted model: in population, the predictions of the unrestricted model are the same as those of the restricted model if the ICAPM restrictions are correct. In cases such as these, the commonly used Diebold and Mariano (1995) statistic to compare out-of-sample mean squared prediction errors (MSPE) has a nonstandard limiting distribution. As Clark and West (2006) point out, under the null hypothesis that the restrictions are true, we should expect a better MSPE for the restricted model, since the restrictions imply a gain in estimation efficiency. The MSPE difference statistic is not centered at zero, and the test should take this into account. In addition, the asymptotic distribution of the statistic is in general not normal.
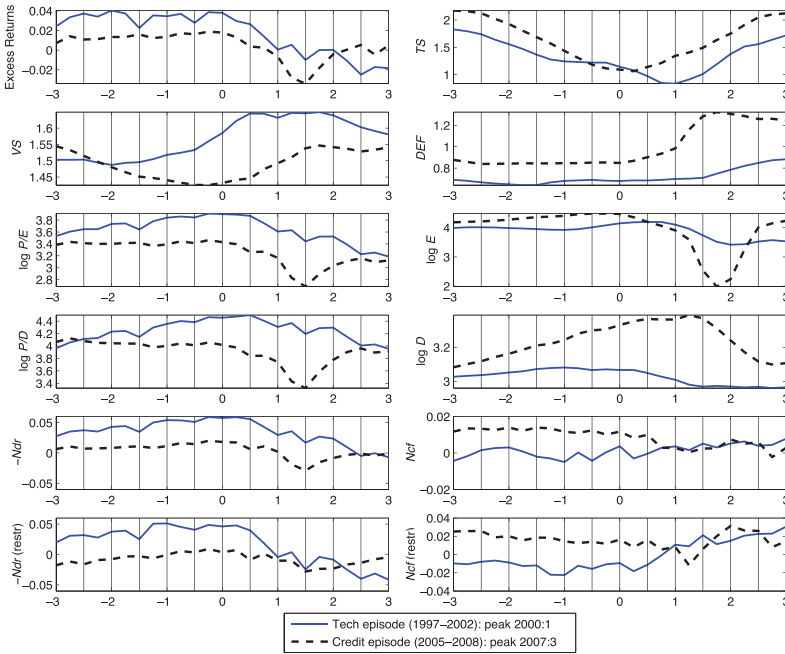
Clark and West (2007) present simulation-based evidence that applying standard critical values to a suitably adjusted MSPE difference statistic produces an approximately correctly sized test. We follow their suggested approach. The difference in the MSPE of the five-variable restricted model relative to the unrestricted model is -0.0008, so the restricted model performs better out-of-sample than the unrestricted model. The adjusted MSPE difference of Clark and West is -0.0001, with a t-statistic of -0.18. The results presented by Clark and West suggest that one should reject the restricted model at a 10% confidence level if the t-statistic of the adjusted MSPE is higher than 1.28, that is, if the restricted model performs sufficiently worse than the unrestricted model after penalizing the restricted model for the expected gain in efficiency under the null hypothesis. In our case, the restricted model improves the MSPE in line with what we would expect if the restrictions were true (the t-statistic is close to zero), and therefore we do not reject our ICAPM model. This confirms that imposing ICAPM restrictions can indeed improve estimation efficiency and out-of-sample predictive power.

4.4 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: 2000:1 for the technology boom, and 2007:3 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 although, 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

Figure 4
Aggregate state variables during the tech and credit episodes
This figure plots the evolution in event time for the key aggregate variables in our analysis for both the tech episode of 1997–2002 and the credit episode of 2005–2010. The event for each period is the market peak (tech: 2000:1, credit: 2007:3). The variables we plot include the excess return on the market, the small stock value spread (*VS*), the default yield spread (*DEF*), the price-earnings ratio (*PE*), the market's smoothed earnings (*E*), the price-to-dividend ratio (*PD*), and the market's smoothed dividends (*D*). Excess returns, *VS* and *DEF* are exponentially smoothed with decay 0.08.

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 in Figure 4 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, and only then started to decline. 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 to 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 (followed by a negative cash flow shock with some delay) 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, although 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.
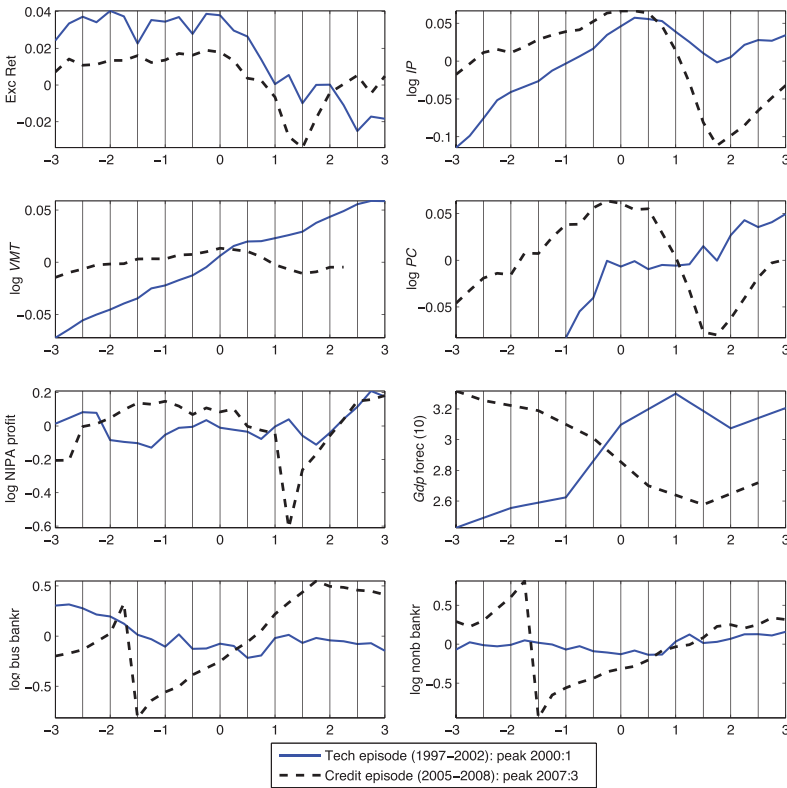
When we do estimate a VAR imposing the restrictions of the ICAPM, the restricted system places greater emphasis on movements in the default spread. Figure 4 shows that the default spread was relatively stable during the tech boom and bust, but spiked upwards in late 2008 and early 2009, between one and 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.

These observations apply as well to the various robustness tests reported in the Appendix. Both the strong negative correlation of cash-flow and

discount-rate news, and the strong negative predictive power of *DEF* for future cash flows (or *VS* in some cases), imply that the recent credit episode is mainly explained by cash-flow news. Similarly, the large movements in *PE* imply a dominant discount-rate news component for the tech boom and bust across all specifications.

4.5 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, reported in event time around the two

Figure 5
Other macroeconomic aggregates during the tech and credit episodes
This figure plots the evolution in event time for several macroeconomic variables for both the tech episode of 1997–2002 and the credit episode of 2005–2010. The event for each period is the market peak (tech: 2000:1, credit: 2007:3). The variables include the smoothed excess return on the market, the seasonally adjusted logs of industrial production (*IP*), vehicle miles travelled in the U.S. (*VMT*), the Ceridian-UCLA Pulse of Commerce Index based on diesel fuel consumption in the U.S. trucking industry (*PC*), the log NIPA profits, the 10-year GDP Forecast from the Survey of Professional Forecasters, and the log of business and nonbusiness bankruptcies.

peaks. The top left panel reports the smoothed excess return on the market. Then, we report several measures of real activity, namely the seasonally-adjusted logs of industrial production (*IP*), vehicle miles travelled in the U.S. (*VMT*), and the Ceridian-UCLA Pulse of Commerce Index based on diesel fuel consumption in the US trucking industry (*PC*). All of these series show much greater declines in the credit bust than in the technology bust. We also show NIPA profits, which do not include capital gains and losses, and highlight the real nature of the credit bust (and successive recovery). Next, the graph of the 10-year GDP forecast from the Survey of Professional Forecasters shows the very different nature of the two episodes not only in terms of current activity, but also in terms of expectations about the future. Finally, we report the logs of business and nonbusiness bankruptcies. Both types of bankruptcy decline sharply in 2006 with a change in U.S. 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 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 2000–2002, while both temporary and permanent shocks were important in 2007–2009. Their finding, while obtained using a very different methodology, is consistent with ours.
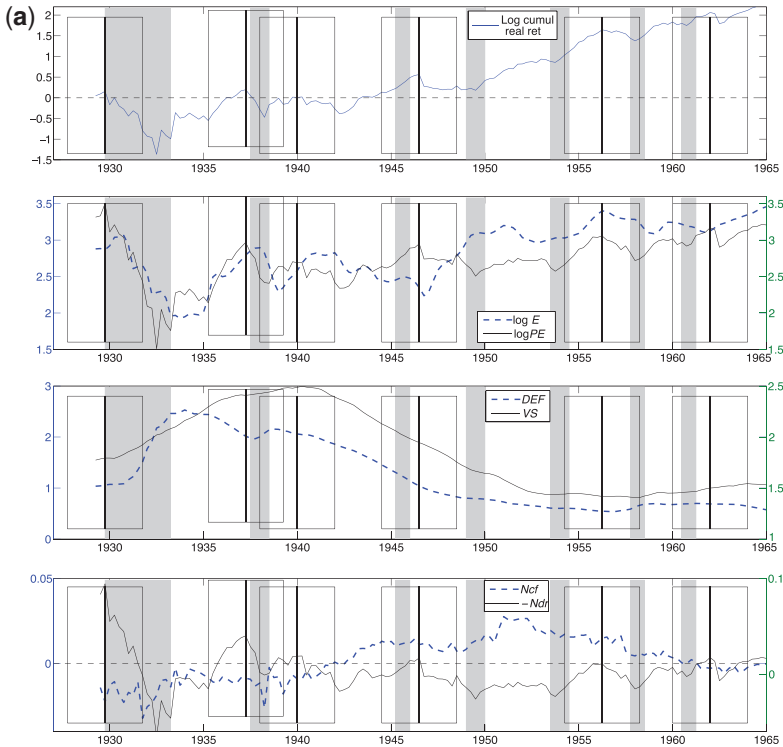
5. Booms and Busts During the Last Century

This analysis can be extended to all booms and busts in the stock market for the last century, in the spirit of Burns and Mitchell (1946) and Blanchard and Watson (1986). We start by designing a simple algorithm to detect boom-bust episodes and their peaks from the time series of log real market returns. After removing a linear trend, we identify quarters that are part of a downturn in the S&P 500 if any of the following conditions is met: the quarter is followed by a cumulative decline of 30% in six quarters, or 25% in four quarters, or 20% in two quarters; this allows us to capture episodes in which the declines occur in different time frames. We then identify "episodes" by grouping together quarters that are deemed to be part of a downturn, as long as they are not more than two quarters apart from each other. The peak of the episode

will then be the quarter of each group with the highest value of the S&P. Finally, we consider only episodes that include at least two quarters of downturn.
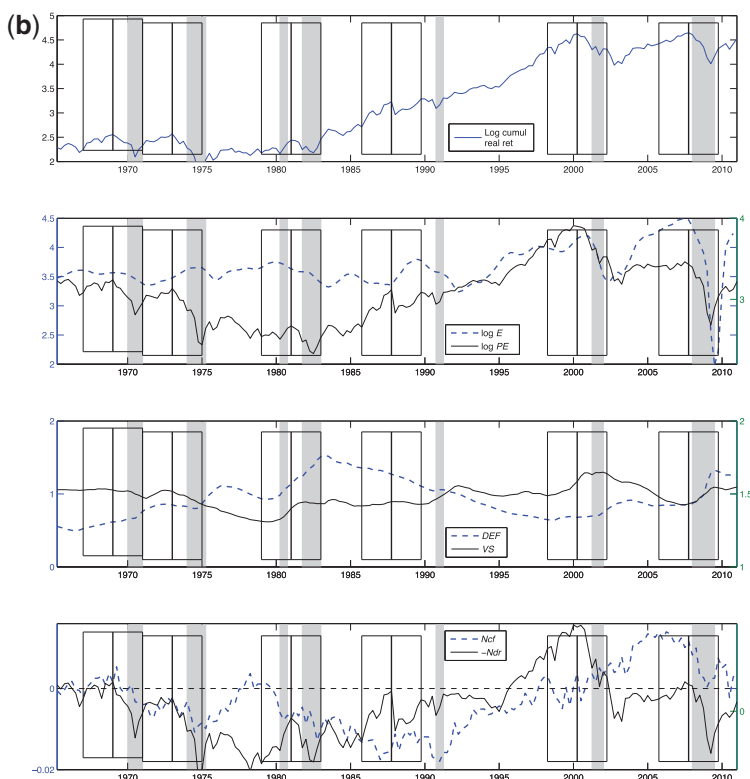
Proceeding in this way, we identify 12 episodes, with peaks in 1929:3, 1937:1, 1940:1, 1946:2, 1956:1, 1961:4, 1968:4, 1972:4, 1980:4, 1987:3, 2000:1, and 2007:3. Figure 6 plots real cumulative log returns (first panel); log E and PE (second panel); log default and value spread (smoothed, third panel); N_{CF} and negative N_{DR} (smoothed, last panel) from 1929:2 to 2010:4 together with two-year intervals centered around each peak. Figure 6a shows the period 1929–1965, while Figure 6b shows the period 1970–2010. Shaded areas correspond to NBER recessions.

Looking at the history of the S&P 500, we can classify boom-bust episodes into three categories. First, we have "hard times": episodes driven mainly by cash-flow news. In these cases, large negative movements in cash-flow news are accompanied by large drops in earnings, a widening of the default spread and of the value spread. NBER recessions start almost immediately after the stock market begins to drop. Besides the recent credit episode, this group includes the downturns of 1929 and 1937.

Figure 6
Booms and busts from 1929 to 2010

34

Figure 6
Continued
This figure plots excess returns, log E, log PE, the log default spread and value spread, together with N_{CF} and N_{DR} from the five-variable unrestricted model between 1929 and 2010. Shaded areas indicate NBER recessions. Boxes are centered around market peaks and include two years before and after each peak. Figure 6a shows the period 1929–1965, while Figure 6b shows the period 1970–2010. Shaded areas correspond to NBER recessions.

Second, we can identify a few "pure sentiment" episodes, in which we observe sharp increases and drops in the market driven almost exclusively by discount-rate news. During these periods, we do not observe large increases in the default spread and value spread (which may instead be decreasing). These episodes are not followed by an NBER recession in the next few years. Pure sentiment episodes are the 1987 market crash, as well as the downturns in 1946, immediately after the end of World War II, and in 1961.

Finally, the other episodes (concentrated between 1955 and 2002, including the technology boom-bust period) are primarily–but not exclusively–driven by sentiment. They feature stable or growing earnings during the boom, and the downturn is driven mostly by discount rate shocks. Cash-flow shocks reinforce the downturn, but with a delay of several quarters. NBER recessions usually follow the sentiment shock with a lag of 9 to 18 months.

Hard times, then, are not new to the recent financial crisis. Sharp drops in the stock market driven primarily by cash-flow news have been observed in particularly bad times: the Great Depression and the period immediately before World War II. Most boom-bust episodes in the stock market in the last century, however, can be attributed to sentiment shocks (alone or accompanied by delayed, and less severe, cash flow shocks) – much more bearable for long-term investors.

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 (e.g., Campbell and Thompson 2008; Cochrane 2008; 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. Market returns 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 using GMM. Although the test of overidentifying restrictions statistically rejects our joint model, we argue that the economic theory we exploit, the ICAPM, contains valuable information that can improve estimation of time variation in the equity premium. Indeed, out-of-sample tests confirm the usefulness of our theory-driven approach. Naturally, other approaches, for example a Bayesian framework, could be used to incorporate theoretical restrictions in the estimation. We do not pursue the Bayesian approach in our paper, but we are confident that in that framework the theoretical restrictions would affect the estimated news terms in a very similar direction to the one we obtain, yielding qualitatively similar results.

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

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron. 2012. Volatility, the Macroeconomy and Asset Prices. Unpublished paper. Duke University and University of Pennsylvania.

Blanchard, O., and M. Watson. 1986. Are Business Cycles All Alike? In Robert J. Gordon (ed.), *The American Business Cycle* Cambridge, MA: NBER.

Burns, A., and W. Mitchell. 1946. *Measuring Business Cycles*. New York, NY: NBER.

Campbell, J. Y. 1991. A Variance Decomposition for Stock Returns. *Economic Journal* 101:157–79.

———. 1993. Intertemporal Asset Pricing Without Consumption Data. *American Economic Review* 83:487–512.

———. 1996. Understanding Risk and Return. *Journal of Political Economy* 104:298–345.

Campbell, J. Y., S. Giglio, and C. Polk. 2012. Hard Times. Appendix. Available online at the authors' websites.

Campbell, J. Y., S. Giglio, C. Polk, and B. Turley. 2011. An Intertemporal CAPM with Stochastic Volatility. Working paper. Harvard University, University of Chicago, and London School of Economics.

Campbell, J. Y., C. Polk, and T. Vuolteenaho. 2010. Growth or Glamour? Fundamentals and Systematic Risk in Stock Returns. *Review of Financial Studies* 23:305–44.

Campbell, J. Y., and R. J. Shiller. 1988a. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies* 1:195–228.

———. 1988b. Stock Prices, Earnings, and Expected Dividends. *Journal of Finance* 43:661–76.

———. 1998. Valuation Ratios and the Long-Run Stock Market Outlook. *Journal of Portfolio Management* 24:11–26.

Campbell, J. Y., and S. Thompson. 2008. Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average? *Review of Financial Studies* 21:1509–31.

Campbell, J. Y., and T. Vuolteenaho. 2004. Bad Beta, Good Beta. *American Economic Review* 94:1249–75.

Chen, L., and X. Zhao. 2009. Return Decomposition. *Review of Financial Studies* 22:5213–49.

Claessens, Stijn, M. Kose, and M. Terrones. 2008. What happens during recessions, crunches and busts? IMF Working Paper 08/274.

Clark, T. E., and K. D. West. 2006. Using Out-of-Sample Mean Squared Prediction Errors to Test the Martingale Difference Hypothesis. *Journal of Econometrics* 135:155–86.

———. 2007. Approximately Normal Tests for Equal Predictive Accuracy in Nested Models. *Journal of Econometrics* 138:291–311.

Cochrane, J. H. 2008. The Dog That Did Not Bark: A Defense of Return Predictability. *Review of Financial Studies* 21:1533–75.

———. 2011. Presidential Address: Discount Rates. *Journal of Finance* 66:1047–1108.

Diebold, F. X., and R. S. Mariano. 1995. Comparing Predictive Accuracy. *Journal of Business and Economic Statistics* 13:253–63.

Engsted, T., T. Q. Pedersen, and C. Tanggaard. 2010. Pitfalls in VAR Based Return Decompositions: A Clarification. Working paper. University of Aarhus.

Epstein, L., and S. Zin. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57:937–69.

———. 1991. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Investigation. *Journal of Political Economy* 99:263–86.

Fama, E. F., and K. R. French. 1989. Business Conditions and Expected Returns on Stocks and Bonds. *Journal of Financial Economics* 25:23–50.

———. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33: 3–56.

Goyal, A., and I. Welch. 2008. A Comprehensive Look at the Empirical Performance of Equity Premium Prediction. *Review of Financial Studies* 21:1455–1508.

Graham, B., and D. L. Dodd. 1934. *Security Analysis*. New York: McGraw Hill.

Hansen, L. P., J. Heaton, and A. Yaron. 1996. Finite-Sample Properties of Some Alternative GMM Estimators. *Journal of Business and Economic Statistics* 14:262–80.

Kelly, B., and S. Pruitt. 2011. Market Expectations in the Cross-Section of Present Values. Working paper. University of Chicago.

Lettau, M., and S. C. Ludvigson. 2001. Consumption, Aggregate Wealth and Expected Stock Returns. *Journal of Finance,* 66:815–49.

———. 2011. Shocks and Crashes. NBER Working Paper No. 16996.

Malloy, C., T. Moskowitz, and A. Vissing-Jorgensen. 2009. Long-run Stockholder Consumption Risk and Asset Returns. *Journal of Finance* 44:2427–79.

Merton, R. C. 1973. An Intertemporal Capital Asset Pricing Model. *Econometrica* 41:867–87.

Merton, R. 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8:1–39.

Newey, W. K., and R. J. Smith. 2004. Higher-Order Properties of GMM and Generalized Empirical Likelihood Estimators. *Econometrica* 72:219–55.

Polk, C., S. Thompson, and T. Vuolteenaho. 2006. Cross-Sectional Forecasts of the Equity Premium. *Journal of Financial Economics* 81:101–41.

Ranish, B. 2009. Was This Stock Market Crash Different? Working paper. Harvard University.

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

Shiller, R. J. 2000. *Irrational Exuberance*. Princeton, New Jersey: Princeton University Press.