Does Precautionary Savings Drive the Real Interest Rate? Evidence from the Stock Market*

Carolin Pflueger† Emil Siriwardane‡ Adi Sunderam§

August 12, 2017

Abstract

We document a strong and robust relation between the one-year real rate and the valuation of high-volatility stocks, which we argue measures precautionary savings motives. Our novel proxy for precautionary savings explains 44% of variation in the real rate. In addition, the real rate forecasts returns on the low-minus-high volatility portfolio but appears unrelated to observable measures of the quantity of risk. Our results suggest that precautionary savings motives, and thus the real rate, are driven by time-varying attitudes towards risk. These findings are difficult to rationalize in models with perfect risk sharing and highlight the role that imperfect diversification plays in determining interest rates.

*We thank Michael Brennan (discussant), John Campbell, Xavier Gabaix, Bryan Kelly, Hanno Lustig, Thomas Maurer (discussant), Monika Piazzesi, Robert Ready (discussant), Martin Schneider, Andrei Shleifer, Jeremy Stein, and Luis Viceira and seminar participants at the CEF 2017, CITE 2017, London School of Economics, Federal Reserve Board, University of British Columbia, SFS Cavalcade, and HEC-McGill Winter Finance Workshop for helpful comments. Pflueger gratefully acknowledges funding from the Social Sciences and Humanities Research Council of Canada (grant number 430-2014-00796). The Online Appendix to the paper can be found here and the Data Appendix can be found here.

†Pflueger: University of British Columbia. E-mail: carolin.pflueger@sauder.ubc.ca

‡Siriwardane: Harvard Business School. E-mail: esiriwardane@hbs.edu.

§Sunderam: Harvard Business School and NBER. E-mail: asunderam@hbs.edu.

1 Introduction

What drives real interest rates? In recent years, this question has received renewed attention because of unusually low interest rates across the developed world. A potentially important source of variation in the real interest rate is the precautionary savings motive, which may vary due to changes in uncertainty faced by investors or due to changes in investors' aversion to uncertainty. Variation in the precautionary savings motive has significant implications for both financial markets and real investment (e.g., Hall (2016), Cochrane (2016)). In this paper, we provide new evidence from the cross-section of stocks that the precautionary savings channel has historically played a major role in driving real interest rates. Moreover, we provide evidence that variation in aversion to uncertainty is a central reason that the economy's desire for precautionary savings is itself moving around.

Understanding what drives variation in the real rate - a key asset price for consumption, investment, and savings decisions - is fundamental to finance and macroeconomics. The precautionary savings motive, in turn, is important for understanding the origins of business cycles, the effectiveness of conventional and unconventional monetary policy, and firms' cash holdings.[1] Measuring variation in the precautionary savings motive is a challenge. The standard approach relies on estimating volatilities of income or consumption and relating them to investment and savings decisions (e.g., Carroll and Samwick (1998); Lusardi (1998); Banks et al. (2001); Parker and Preston (2005)).

Our key empirical innovation is to use asset prices – specifically the cross section of stock market valuation ratios – to shed light on the strength of the precautionary savings motive over time. Relying on asset prices is advantageous because they automatically aggregate over different agents in the economy and are available at a much higher frequency than income or consumption. Asset prices are also unique in that they allow us to estimate investors' willingness to pay to avoid uncertainty at a given point in time.

[1] See, e.g., Bloom (2009); Bloom et al. (2014); Cochrane (2016); Laubach and Williams (2003); McKay, Nakamura, and Steinsson (2016); Holston, Laubach, and Williams (2016); Riddick and Whited (2009); Duchin et al. (2016).

We start from the intuition that if investors are differentially exposed to idiosyncratic shocks, for instance due to market segmentation among professional investors or households' undiversifiable labor income risk, high aversion to uncertainty and strong precautionary savings motives should drive down valuations for high-volatility stocks relative to low-volatility stocks. Building on this intuition, we use the price of volatile stocks (henceforth "PVS_t") relative to low-volatility stocks, defined as the book-to-market ratio of low-volatility stocks minus the book-to-market ratio of high-volatility stocks, as our key proxy for precautionary savings. Intuitively, an increase in precautionary motives means that investors should be less willing to hold volatile assets and should increase their demand for real risk-free bonds. This intuition suggests that if investors' demand for precautionary savings is an important driver of the real rate, we should expect the real rate to move in the same direction as PVS_t, and PVS_t should explain substantial time-variation in the real rate.

We begin by establishing several novel empirical facts about the relationship between real rates and the cross section of stocks. First, we show that PVS_t is strongly correlated with the real rate, measured as the 1-year Treasury bill rate net of survey expectations of 1-year inflation. Put differently, a low risk-free rate typically coincides with low prices for high-volatility stocks compared to low-volatility stocks, as would be the case if aversion to uncertainty were a major driver of risk-free bond valuations. The relationship is robust in both levels and changes and is strongly economically significant. The headline result of the paper is that PVS_t explains 44% of the variation in the real rate from 1973 to 2015.

Our emphasis on the cross section is important, as the valuation of the aggregate stock market has little explanatory power for the real rate. This indicates that PVS_t is not simply another proxy for risk aversion to aggregate market fluctuations. Our particular focus on equity volatility is also critical. Real rate variation is not explained by valuation-ratio spreads generated from sorting stocks based on size, value, leverage, duration of cash flows, cash flow beta or CAPM beta - all characteristics that are known to describe the cross section of stock returns.[2] The relation between the real rate and PVS_t is robust to whether we sort by stock return volatility over the past two

[2] The relative valuation of small and big stocks does seem to possess some explanatory power but is subsumed by PVS_t.

months or past two years, indicating that results are not driven by stocks quickly rotating in and out of high- and low-volatility portfolios. Furthermore, the ability of the PVS_t to explain real rate variation remains after we account for changes in macroeconomic uncertainty (e.g., total factor productivity volatility), the business cycle, and inflation.

We then delve deeper into what drives the relationship between the real rate and PVS_t. Standard present value identities point to two possible explanations. Because it is a valuation ratio, changes in PVS_t must reflect either differential changes in expected cash flow growth or differential changes in expected returns between low- and high-volatility stocks. In other words, the real rate may correlate with PVS_t because it loads on factors that drive expected cash flow growth or factors that determine expected returns. The data points to expected returns, as the real rate forecasts future returns on a portfolio that is long low-volatility stocks and short high-volatility stocks, but does not forecast *ROE* for the same low-minus-high volatility portfolio. These findings imply that the factors driving expected returns on volatility-sorted portfolios also drive real rate variation.

Taken together, these pieces of evidence paint a clear picture. The book-to-market spread between low- and high-volatility stocks captures the compensation investors demand for bearing uncertainty, and thus their demand for precautionary savings. In turn, the relationship between PVS_t and the real interest rate implies that variation in precautionary savings is a significant driver of movements in the real rate.

We next explore why investor compensation for bearing uncertainty varies over time. Changes in expected returns must reflect either changing investor aversion to volatility or changing quantities of volatility. We look for evidence that the real rate is correlated with observable quantities of risk and find little. Real rates are not contemporaneously correlated with the realized return volatility of the low-minus-high volatility stock portfolio, nor the realized volatility of the aggregate stock market. Furthermore, real rates do not forecast realized volatility of the low-minus-high volatility stock portfolio or the realized volatility of the aggregate stock market. Finally, the forecasting power of the real rate for returns on the long-short portfolio sorted on volatility is robust to controlling for volatility itself. It is hard to rule out comovement between real rates and hard-to-observe components of volatility. However, these results suggest that variation in the precautionary savings motive, and hence variation in the real rate, is driven by changing investor aversion to volatility rather than changing quantities of volatility. To be clear, the quantity of volatility certainly displays significant time variation. However, the data does not provide a strong indication that this variation drives PVS_t or the real interest rate.

The relation between precautionary savings and the real interest rate has important implications for monetary policy. In a standard New Keynesian framework, the central bank optimally adjusts interest rates to fully accommodate shocks to the natural real rate – or the interest rate consistent with output at its natural rate and stable inflation – and monetary policy tightness should be assessed relative to the natural real rate (Clarida et al., 1999). The link between precautionary savings motives and the real interest rate depends only on the investor's Euler equation and is hence independent of any price-setting frictions. If the relation between the real rate and time-varying precautionary savings motives indeed reflects time-variation in the natural real rate, as this intuition would suggest, output and inflation should respond to precautionary savings shocks very differently than to independent real rate shocks due to monetary policy. Impulse response functions following the recursive identification scheme of Bernanke and Mihov (1998) corroborate this prediction in the data, supporting the notion that it is important to account for precautionary savings demand in assessing monetary policy.

Finally, we provide a highly stylized model consistent with our empirical results. In the model, portfolio volatility – not beta — is the proper measure of risk because markets are segmented and investors are imperfectly diversified. We think of this assumption as either representing under-diversified households, consistent with employees' bias towards their own employer's stock in 401(k) plans (Benartzi (2001)), or as capturing segmented institutional investors who take concentrated positions in individual stocks (Shleifer and Vishny (1997); Gromb and Vayanos (2010); Cremers and Petajisto (2009); Kacperczyk et al. (2005); Agarwal et al. (2013)). Investors require time-varying risk premia, which we model as arising from slowly-moving habit (Campbell and Cochrane (1999); Menzly et al. (2004)), generating volatile and predictable stock returns, as in the data. Investors are borrowing-constrained, so the real risk-free rate is determined by whoever values the risk-free asset most highly (Miller (1977)). Marginal bond investors are typically investors with highly uncertain consumption streams and strong time-varying precautionary savings motives. A shock to high-volatility investors' risk aversion raises risk premia and drives down prices of high-volatility stocks relative to low-volatility stocks. Simultaneously, this increase in risk aversion increases the precautionary savings motive of marginal savers, driving down the risk-free rate. Book-to-market ratios and expected excess returns for low-minus-high-volatility stocks hence fall at the same time as the risk-free rate. Market segmentation between high- and low-volatility stock investors implies that only a small fraction of stock market investors are marginal in the bond market. As a result, the risk-free rate is close to uncorrelated with the aggregate book-to-market ratio, as in the data. In a calibrated version of the model, the relationships between the real risk-free rate, PVS_t, and future low-minus-high-volatility equity excess returns are quantitatively consistent with the data. We consider the model illustrative and conclude by discussing several alternative models that are consistent with the channel favored by our empirical evidence.

Our paper is related to several strands of the literature. On the asset pricing side, it contributes to the literature on the pricing of idiosyncratic risk in the stock market (Ang et al. (2006a, 2009); Johnson (2004); Fu (2009); Stambaugh et al. (2015); Hou and Loh (2016); Herskovic et al. (2016)). While this literature has focused on the average returns on low-volatility stocks over high-volatility stocks, we contribute by studying how the valuation of low-minus-high volatility stocks varies over time. The relation between risk premia in bonds and stocks has been a long-standing question in financial economics (Fama and French, 1993; Koijen et al., 2010; Baker and Wurgler, 2012) and we contribute by showing that the pricing of volatility in the cross-section of stocks can help understand fundamental drivers of the real risk-free rate. The model most closely related to ours is Herskovic et al. (2016), where idiosyncratic firm-level shocks matter for households with cross-sectional asset pricing implications. However, Herskovic et al. (2016) focus on a different cross-section of stocks, sorting stocks by their exposure to the common factor driving idiosyncratic volatility, and study how this exposure is priced in the cross section of stocks. On the other hand, our focus is on how the relative valuation of high- and low-volatility stocks connect to real interest rates. Indeed, in their model, the correlation between the risk-free rate and the model equivalent of PVS_t takes the opposite sign of what we find.[3] Rationalizing our findings therefore requires a different pricing mechanism, which we argue can be accomplished with market segmentation and time-varying attitudes towards volatility.

This paper also contributes to a recent literature in macroeconomics that seeks to estimate the time-varying natural rate of interest (Laubach and Williams (2003); Cúrdia et al. (2015)), which uses either long-term historical data or dynamic stochastic equilibrium models. Our findings emphasize that time-varying precautionary savings play an important role in driving investors' demand for savings and are consistent with McKay et al. (2016), who argue that consumers' precautionary savings motive helps explain why forward guidance by central banks has been less effective in stimulating consumption and spending than standard New Keynesian models might suggest, and with a recent corporate finance literature that attributes high recent corporate cash holdings to a precautionary savings motive (Riddick and Whited (2009); Duchin et al. (2016)). A closely related paper is Hartzmark (2016), who estimates changes in expected macroeconomic volatility to argue that precautionary savings is an important driver of real interest rates. In contrast, our approach pins down variation in the precautionary savings motive by using information from the cross section of stocks. Using a stock market based measure of precautionary savings, we contribute over previous findings by showing that time-varying demand for precautionary savings is not just a result of time-varying volatility, but that the time-varying price of volatility is important for understanding the natural real rate of interest.[4]

The remainder of this paper is organized as follows. Section 2 describes the data and portfolio

[3] In their model, a positive shock to idiosyncratic volatility drives down the risk-free rate but drives up the price of high-volatility stocks relative to low-volatility stocks due to a convexity effect. Empirically, we also find little evidence that their common idiosyncratic volatility factor is correlated with the real rate.

[4] In contrast to Hartzmark (2016), we do not find a significant relation between variation in volatility itself and the real rate but instead provide evidence that the pricing of volatility has changed over time. While this might at first appear in contrast with Hartzmark (2016), we note that our empirical sample is substantively different. We estimate precautionary savings during normal business-cycle fluctuations, while Hartzmark (2016) includes data from the 1930s, when both interest rates and uncertainty experienced very large swings. We can therefore reconcile the results in this paper with Hartzmark (2016) if precautionary savings and the quantity of volatility are unrelated during normal times, but move together during rare episodes of extreme economic fluctuations.

construction. Section 3 presents the main empirical results. Section 4 explores monetary policy implications. Section 5 describes the model, shows that it can replicate the empirical findings, and discusses alternative interpretations. Finally, Section 6 concludes.

2 Data

We construct a quarterly data set running from 1973 to 2015. We include all U.S. common equity in the CRSP-COMPUSTAT merged data set that is traded on the NYSE, AMEX, or NASDAQ exchanges. We provide full details of all of the data used in the paper in a separate Data Appendix. Here, we briefly describe the construction of some of our key variables.

2.1 Construction of Key Variables

Valuation Ratios

The valuation ratios used in the paper derive from the CRSP-COMPUSTAT merged database. At the end of each quarter and for each individual stock, we form book-to-market ratios. The value of book equity comes from COMPUSTAT Quarterly and is defined following Fama and French (1993). We assume that accounting information for each firm is known with a one-quarter lag. At the end of each quarter, we use the trailing six-month average of market capitalization when computing the book-to-market ratio of a given firm. This smooths out any short-term fluctuations in market value. We have experimented with many variants on the construction of book-to-market, and our results are not sensitive to these choices.

Volatility-Sorted Portfolio Construction

At the end of each quarter, we use daily CRSP stock data from from the previous two months to compute equity volatility. We exclude firms that do not have at least 20 observations over this time frame. This approach mirrors the construction of variance-sorted portfolios on Ken French's website. We compute each firm's volatility using ex-dividend firm returns.[5]

At the end of each quarter, we sort firms into quintiles based on their volatility. At any given point in time, the valuation ratio for a quintile is simply the equal-weighted average of the valuation ratios of stocks in that quintile. One of the key variables in our empirical analysis is PVS_t, the difference between the average book-to-market ratio of stocks in the lowest quintile of volatility and the average book-to-market ratio of stocks in the highest quintile of volatility. Again, PVS_t stands for the "price of volatile stocks," as PVS_t is high when high-volatility stocks have high market valuations. Quarterly realized returns in a given quintile are computed in an analogous fashion, aggregated up using monthly data from CRSP.

The Real Rate

The real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters. We choose a short maturity interest rate, because at this horizon, inflation risk is small, and inflation risk premia are unlikely to affect our measure of the risk-free rate. In the Online Appendix, we conduct formal unit root tests of the real rate and find that it is trend stationary. Thus, for our main analysis, we use a detrended version of the real rate to ensure our statistical analysis is well behaved. Detrending the real rate throughout our analysis ensures that our results are not driven by secular changes in growth expectations, which may be significant in explaining low-frequency movements in the natural real rate (Laubach and Williams, 2003).[6]

2.2 Summary Statistics

Table 1 contains basic summary statistics on our volatility-sorted portfolios. The first thing to notice is that, on average, PVS_t is negative; that is, low-volatility stocks have lower book-to-market

[5] In earlier versions of the paper, we instead sorted stocks on idiosyncratic volatility as in Ang, Hodrick, Xing, and Zhang (2006b). Our results are essentially unchanged when using idiosyncratic volatility, mainly because the total volatility of an individual stock is dominated by idiosyncratic volatility (Herskovic et al. (2016))

[6] In a previous version of the paper, we used the real rate without detrending. All of our results are qualitiatively and quantitatively similar.

ratios than high-volatility stocks. However, as Fig. 1 shows, this masks considerable variation in PVS_t. Indeed, the standard deviation of PVS_t is bigger in absolute value than its mean. This variation is at the heart of our empirical work.

Returns on the low-minus-high volatility portfolio are themselves quite volatile, with an annualized standard deviation of 29.95%. While high-volatility stocks in our sample have high book-to-market ratios, the quintile of the most volatile stocks on average has excess returns that are 0.66 percentage points lower than for the lowest-volatility quintile. This is related to the well-known idiosyncratic volatility puzzle of Ang et al. (2006a) and Ang et al. (2009). A number of explanations have been proposed in the literature, ranging from shorting constraints (Stambaugh et al. (2015)) to the convexity of equity payoffs (Johnson (2004)). Those papers focus on the unconditional average level of returns, whereas we focus on time-variation in low-minus-high volatility stock returns and valuations.

The second-to-last row of Table 1 shows that high-volatility portfolios load onto the SMB factor, consistent with highly volatile stocks being smaller on average. Small stocks are more likely to be traded by individuals and specialized institutions (Lee et al. (1991)), so this finding supports the notion that markets for these stocks are segmented, exposing specialized investors to both systematic and idiosyncratic shocks. In turn, market segmentation raises the possibility of a link between volatility and investors' desire for precautionary savings. This logic underlies our interest in how the valuation of high-volatility stocks varies through time.

3 Empirical Results

3.1 Valuation Ratios and the Real Rate

We begin by documenting the strong empirical relationship between the real rate and the book-to-market spread between low- and high-volatility stocks. Specifically, we run regressions of the form:

$$\text{Real Rate}_t = a + b \times PVS_t + \varepsilon_t, \tag{1}$$

where PVS_t is the difference in book-to-market valuations between low- and high-volatility stocks. Because both the real rate and PVS_t spread are persistent, we compute standard errors in multiple ways. Specifically, we compute both Hansen and Hodrick (1980) and Newey and West (1987) standard errors using 12 lags and report the more conservative t-statistic. In the Online Appendix, we also consider several other methods for dealing with the persistence of these variables (e.g. maximum likelihood regressions with AR-GARCH errors). Our main conclusions are robust to these alternative estimation techniques.

Column (1) of Table 2 shows a strong positive correlation between the real rate and PVS_t. When market valuations are high, book-to-market ratios are low. Thus, PVS_t is high when the price of high-volatility stocks are large relative to low-volatility stocks. Column (1) of Table 2 therefore indicates that the real rate tends to be high when investors favor high-volatility stocks. Conversely, the real rate tends to be low when investors are averse to high-volatility stocks. This is the first piece of suggestive evidence that PVS_t spread captures variation in precautionary savings motives.

The magnitude of the effect is large in both economic and statistical terms. A one-standard deviation increase in PVS_t is associated with about a 1.3 percentage point increase in the real rate. As a point of reference, the standard deviation of the real rate is 1.9 percentage points. The R^2 of the univariate regression is 44%, indicating that PVS_t explains a large fraction of variation in the real rate. Fig. 2 makes this point visually, plotting the time series of the real rate against the fitted value from regression in Eq. (1). The figure also shows that the regression is not driven by outliers — PVS_t tracks all of the major variation in the real rate since 1970. Fig. 3 displays the same evidence in a scatter plot. The relationship between the real rate and PVS_t is robust and approximately linear throughout the distribution.

Column (2) of Table 2 shows that our focus on the cross section of stock valuations is important. There is no relationship between the book-to-market ratio of the aggregate stock market and the real rate. This is not just an issue of statistical precision; the economic magnitude of the point estimate is much smaller as well: a one-standard deviation increase in the aggregate book-to-market ratio is associated with a 49 basis point increase in the real rate. In column (3) of Table 2, we show that the statistical significance and even the magnitude of the coefficient on PVS_t are unchanged when controlling for the aggregate book-to-market ratio. We also control for variables that are often thought to determine a monetary policy rule, namely GDP price deflator inflation and the output gap from the Congressional Budget Office (Clarida et al. (1999); Taylor (1993)_. While the output gap enters with a positive coefficient, inflation enters with a slightly negative coefficient. However, both coefficients on the output gap and inflation are statistically indistinguishable from the traditional Taylor (1993) values of 0.5. The main takeaway is that the relationship between the real rate and PVS_t is stable throughout all of these regression specifications.

In Table 3, we rerun the same analysis in changes rather than levels. This helps to ensure that our statistical inference is not distorted by the persistence of either the real rate or PVS_t. Because regression residuals may still be autocorrelated, we again compute both Hansen and Hodrick (1980) and Newey and West (1987) standard errors using six lags and report the more conservative t-statistic. Running regression (1) in differences yields very similar results to running it in levels. As is clear from Table 3, changes in the real rate are strongly correlated with changes in PVS_t. Moreover, the magnitudes and statistical significance of the point estimate on PVS_t are close to what we observe in Table 2. In contrast, there is little relation between changes in the real rate and changes in the aggregate book-to-market ratio. Overall, the evidence in Tables 2 and 3 indicate a robust relationship - both in economic and statistical terms - between the real rate and PVS_t. This is the central empirical finding of the paper, and as we show in later sections, these results stand up to the inclusion of a battery of additional control variables and different regression specifications.

3.2 Alternative Cross-Sectional Sorts

We now explore alternative explanations for the empirical relationship between the real rate and stock portfolios sorted on volatility. Specifically, we examine the possibility that volatility is simply correlated with another characteristic that is more important for explaining the real rate. We sort stocks along a variety of dimensions and form book-to-market spreads based on the sorting variable. For instance, when examining size as a characteristic, we sort stocks in quintiles based on their market capitalization and then compute the difference between the book-to-market ratio of the smallest and the largest stocks. We then augment the regression in Eq. (1) by adding the spread in book-to-market based on each sort. For additional robustness, we also run this analysis in first-differences and 4-quarter differences.

The results are displayed in Table 4. To start, we recompute PVS_t using a two-year window of volatility, as opposed to a two-month window. As row (2) shows, this variant of PVS_t is highly correlated with the real rate. One might be concerned that our findings are driven by value stocks rotating in and out of high-volatility and low-volatility portfolios. By computing volatility over a long period, we ensure that our results are not driven by quickly changing portfolio compositions, but instead by changes in valuations of stocks with a long history of being volatile. This distinction will be relevant later when we argue that PVS_t moves around because of time-varying attitudes towards risk, not time-varying quantities of risk.

In row (4), we relate the real rate to the spread in book-to-market sorting stocks based on the expected duration of their cash flows. If low-volatility stocks simply have higher duration cash flows than high-volatility stocks, then their valuations should rise relative to high-volatility stocks when real rates rise.[7] This is one sense in which low-volatility stocks may be more "bond-like" than high-volatility stocks (e.g., Baker and Wurgler (2012)).[8] In this case, a mechanical duration effect could explain our results in Table 2. To examine this possibility, we follow Weber (2016) and construct the expected duration of cash flows for each firm in our data. We then sort stocks based on this duration measure and calculate the spread in book-to-market between high and low duration stocks. As row (1) shows, the relationship between r_t and PVS_t appears robust to controlling for

[7] This is a particular version of the broader possibility that our results are driven by reverse causality. Our interpretation is that both real rates and the relative valuations of low- and high-volatility stocks are responding to the same factor, precautionary saving. Alternatively, it could be the case that changes in real rates are driving changes in valuations. In addition to examining alternative cross-sectional sorts, we have examined this possibility by examining monetary policy shocks. In untabulated results, we verify that the relationship between the real rate and PVS_t is unaffected by controlling for monetary policy shocks, as identified by Romer and Romer (2004), Bernanke and Kuttner (2005)), and McKay et al. (2016). This gives us some comfort that reverse causality is not driving our results.

[8] The alternative sense that low-volatility stocks are more bond-like because they are less volatile and idiosyncratic risk matters is exactly what we are trying to capture.

the duration-based value spread .

Row (5) displays the same exercise after controlling for the valuations of high-leverage versus low-leverage stocks. We define leverage as the book value of long-term debt divided by the market value of equity. It seems natural to think that high-leverage firms have high volatility, and since these firms effectively are short bonds, their equity may suffer disproportionately from a decrease in the real rate. However, as row (5) shows, PVS_t is not driven out by the leveraged-based value spread in any of the specifications.

In rows (6)-(9), we run horse races of PVS_t against spreads based on various measures of systematic risk (i.e., beta). Row (6) constructs a value spread based on beta from the past two years of monthly returns. Row (7) computes beta using the past ten years of semi-annual returns. Row (8) uses the past two months of daily returns to compute beta, mimicking our construction of volatility. The regression coefficient on PVS_t remains statistically significant at a 5% level in nearly all cases, and is significant at a 10% level for all cases. Thus, it does not appear that our measure of volatility is simply picking up on beta. Finally, row (9) runs a horse race against a spread based on the estimated beta of each firm's cash flows with respect to aggregate cash flows. Specifically, cash flow betas are computed via rolling twelve quarter regressions of quarter-on-quarter EBITDA growth on quarter-on-quarter national income growth. EBITDA is defined as the cumulative sum of operating income before depreciation. We require a minimum of 80% of observations in a window to compute a cash flow beta. If high-volatility stocks have higher cash flow betas than low-volatility stocks, then their valuations should fall more when aggregate growth expectations are low. In this case, our results in Table 2 could be explained by changes in aggregate growth expectations rather than changes in the precautionary savings motive. Contrary to this hypothesis, Row (9) shows that the book-to-market based on cash flow betas does not drive out PVS_t.

In addition, we compare PVS_t to book-to-market spreads based on the popular Fama-French sorting variables, size and value. The book-to-market spread between small and large stocks does correlate with PVS_t, as indicated by the fact that the statistical significance of PVS_t is weakened in some specifications. This is apparent in row (2) for the value-weighted version of PVS_t, as well as in the horse races contained in row (10). Still, the significance of PVS_t never drops below 10% when including the size-based spread, and it is not surprising that the two correlate, as it is well known that smaller stocks have more volatility. Below, we also conduct double sorts that demonstrate that the explanatory power of PVS_t is robust to controlling for size. Row (11) repeats the horse races of PVS_t against the book-to-market spread between value and growth stocks. Once again, the effect of PVS_t is robust in these horse races.

We also explore a complementary method of ruling out alternative explanations based on double sorts. Specifically, we construct double sorts based on volatility and another characteristic Y. We then assemble a Y-neutral version of PVS_t: the book-to-market spread from sorting stocks on volatility within each tercile of characteristic Y. This spread measures the difference in valuations of low-volatility and high-volatility stocks that have similar values of characteristic Y. In rows (12)-(16) of Table 4, we show that these double sorted book-to-market spreads are still strongly correlated with the real rate. Finally, our PVS_t measure might be simply capturing the value of industries that are particularly exposed to interest rate changes, like finance. To alleviate this concern, we construct an industry-adjusted version of PVS_t. We first sort stocks into one of the 48 Fama-French industries. Within each industry, we compute the book-to-market spread between low and high-volatility stocks. The industry-adjusted PVS_t is then the average of these spreads across all of the industry. Row (18) shows that this industry-adjusted spread still possesses significant explanatory power for the real rate.

3.3 Returns on Volatility-Sorted Portfolios and the Real Rate

We next seek to understand what drives the correlation between the real rate and PVS_t. At any point in time, PVS_t simply reflects differences in the valuation of high- and low-volatility stocks. It is well known that valuation ratios must reflect either expected cash flow growth or expected returns (Campbell and Shiller, 1988). Thus, the results in Tables 2 and 3 could be driven by growth expectations if the cash flows of high-volatility stocks are more sensitive to aggregate growth than the cash flows of low-volatility stocks. In this case, PVS_t may line up with the real rate because it is a good proxy for variation in expected aggregate growth. Alternatively, PVS_t may be driven by changes in the expected returns of low-volatility stocks, relative to high-volatility stocks. In this case, changes in the compensation investors' demand for bearing uncertainty, and thus their demand for precautionary savings, is one natural explanation for the observed correlation between the real rate and PVS_t.

To disentangle these two possibilities, we run simple return forecasting regressions. Specifically, we forecast the return on a portfolio that is long low-volatility stocks and short high-volatility stocks with either PVS_t or the real rate. Formally, we run:

$$R_{t\to t+k} = a + b \times X_t + \xi_{t+k}, \tag{2}$$

where X_t is either PVS_t or the real rate. Table 5 contains the results of this exercise. In Panel A, we set $k = 1$ and forecast one-quarter ahead returns, while in Panel B we set $k = 4$ and forecast four-quarter returns. For regressions with a one-quarter horizon, standard errors are computed using both Newey and West (1987) and Hansen and Hodrick (1980) with five lags, and we report the more conservative t-statistic of the two. For regression with four-quarter horizons, we use Hodrick (1992) standard errors to be maximally conservative in dealing with overlapping returns.

Column (1) of Table 5 Panel A shows that PVS_t has strong forecasting power for returns on the long-short portfolio. The economic magnitude of the relationship is also strong. A one-standard deviation increase in the spread is associated with a 5.9 percentage point increase in returns on the long-short portfolio. To put this in perspective, the quarterly standard deviation of the long-short portfolio is 15%. Thus, it appears that much of the variation in PVS_t reflects variation in expected returns, consistent with much of the empirical asset pricing literature (e.g., Cochrane (2011)).

Column (2) indicates that this forecasting power remains once we control for the contemporaneous realizations of the Fama and French (1993) risk factors. That is, in regression (2), we control for the realized values of the excess return on the aggregate stock market, HML, and SMB at time $t + 1$. The forecasting power of PVS_t survives the inclusion of these controls, suggesting that we are not just picking up the power of PVS_t to forecast the Fama-French factors - our focus on volatility sorted portfolios is important. However, the magnitude of the coefficient in column (2) is smaller than that in column (1). This reflects the fact that both PVS_t and the real rate have some forecasting power for excess returns of small stocks (SMB).

Column (3) of Table 5 Panel A makes the connection between the real rate and time-varying expected returns on the volatility-sorted portfolio directly. It demonstrates that the real rate also strongly forecasts returns on the long-short portfolio. When the real rate is high, low-volatility stocks tend to do well relative to high-volatility stocks going forward. In contrast, a low real rate means investors require a premium to hold high-volatility stocks, as evidenced by the fact that these stocks tend to do relatively well in the future. A one-standard deviation increase in the real rate is associated with a 3.7 percentage point increase in returns on the long-short portfolio. Thus, movements in the real rate forecast returns on the long-short portfolio nearly as well as movements in PVS_t. This implies that the correlation between the real rate and PVS_t documented in Section 3.1 is largely driven by changes in expected returns, not changes in expected cash flow growth.

Column (4) shows that the relationship between the real rate and returns on the long-short portfolio is weakened when we control for the Fama and French (1993) factors. This again reflects the fact that returns on the long-short portfolio, the real rate, and returns on small stocks (SMB) are correlated.

Panel B of Table 5 shows that we obtain similar results once we move to an annual horizon. The magnitude of the forecasting power of the real rate is again comparable to the forecasting power of PVS_t. Taken together, we interpret the forecasting evidence in Table 5 to mean that variation in the expected return spread between high- and low-volatility stocks captures precautionary savings, and in turn, is strongly correlated with the real interest rate.

In Table 6, we explore in more depth the relationship between the real rate and the Fama and French (1993) factors. The table shows that the real rate and PVS_t have little forecasting power for either the aggregate market excess return or value stocks (HML). Again, this highlights the importance of our focus on volatility sorts as a proxy for the strength of the precautionary savings motive. Neither the market excess return nor cross sectional sorts based on valuations (HML) are strongly related to the real rate. In contrast, there is some evidence that the real rate is related to the return spread between small and large market capitalization stocks (SMB). Intuitively, small stocks tend to have high-volatility, so the two sorts are somewhat correlated. However, based on the horse races and double sorts in Table 4, the overall evidence suggests that volatility, not size, is the main driver of our results.

Table 7 further supports the evidence that the relation between the real rate and PVS_t is driven by discount rates and not cash flows. Table 7 shows that neither PVS_t nor the real rate forecast *ROE* for low- versus high-volatility stocks.

3.4 Prices versus Quantities of Risk

We next dig deeper into the relationship between the real rate and returns on the long-short portfolio sorted on volatility. Changes in expected returns must reflect either changing prices of risk or changing quantities of risk. In other words, aversion to volatility can be moving around over time or the amount of volatility can be moving around over time. In this section, we look for evidence that the real rate is correlated with observable quantities of volatility. Finding no such evidence in a variety of different tests, our evidence supports the view that the relationship between the real rate and returns on the long-short portfolio sorted on volatility is likely driven by changing aversion to volatility.

We begin by showing that the relationship between the real rate and the book-to-market spread is unaffected by controlling for various measures of contemporaneous volatility. Specifically, we run the regression in Eq. (1) and add controls for contemporaneous realized volatility. Our first volatility control is the spread in average realized return volatilities between our low-volatility portfolio and our high-volatility portfolio in quarter t. We compute this variable using daily data. To control for macroeconomic volatility, we include the volatility of TFP growth implied from a GARCH model, as in Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012).[9] In addi-

[9]See Table A.1 of the Online Appendix for further discussion of the estimation of TFP volatility.

tion, we control for the realized within-quarter volatility of the Fama and French (1993) factors, computed using daily data.

The results are presented in Table 8. Columns (1) to (3) contain the results in levels, while columns (4) to (6) use four-quarter changes. Column (1) shows that there is no relationship between the real rate and the relative realized volatility of high and low-volatility stocks. This suggests that it is unlikely that the relationship we document between the real rate and PVS_t is driven by changes in the volatilities of our portfolios. Column (2) shows that there is some evidence that the real rate is related to volatility of the aggregate market and volatility of the SMB portfolio.[10] However, this relationship disappears in column (3) when we include PVS_t. In columns (4) to (6), we obtain similar results when running the analysis in four-quarter changes. The only variable robustly correlated with the real rate is PVS_t, whereas the volatility variables have little impact.

The quantity of risk also has no ability to forecast excess returns on the long-short portfolio of volatility sorted stocks. In Table 9, we re-run the forecasting regression from Eq. (2) and add controls for realized volatility in quarter t. That is, we forecast returns from one-year ahead returns using the current real rate and the current level of volatility.[11]

Column (1) shows that the spread in average realized return volatilities between our low-volatility portfolio and our high-volatility portfolio in quarter t has no forecasting power for returns. Column (2) shows the forecasting power of the real rate remains unchanged when we add this spread in average volatility as a control. In the remaining columns, we run horse races between the real rate and other measures of the quantity of risk: the volatility of TFP growth, the volatility of the market excess return, and the volatilities of the Fama-French factors. None of these measures impacts the forecasting power of the real rate for excess returns on the long-short portfolio of volatility sorted stocks. Column (7) shows a kitchen sink regression in which we include all of the quantity of risk measures simultaneously. There is some reduction in the magnitude and

[10] The opposite signs of aggregate market volatility and SMB volatility are due to the fact that the two variables have an 82% correlation. In untabulated results, where we run univariate regressions of the real rate on either aggregate market volatility or SMB volatility, we find no statistically or economically significant relationship.

[11] In untabulated results, we look at the level of realized volatility over the forecast period, $t+1$ to $t+4$, as well as increases in realized volatility from t to $t+4$. None of these permutations affect the forecasting power of the real rate for returns on the long-short portfolio.

statistical significance of the real rate's forecasting power for returns. However, given the results of the univariate horse races in Columns (2) though (6), this likely simply reflects the limited size of the sample relative to the number of covariates in the regression.

Lastly, one might think that PVS_t is related to expectations of future volatility, but not necessarily to contemporaneous or lagged volatility. In Table 10, we try to forecast volatility directly using either PVS_t or the real rate. Formally, we run:

$$Vol_{t+1} = a + b \times X_t + \varepsilon_{t+1}, \tag{3}$$

where X_t is either PVS_t (Panel A) or the real rate (Panel B). Each column examines a different volatility measure, as specified by the column header. For instance, column (1) examines the spread in average realized return volatilities between our low-volatility portfolio and our high-volatility portfolio, while column (2) examines realized TFP volatility. PVS_t does not forecast any of the volatility variables we examine in Panel A. Similarly, in Panel B, we find that the real rate does not forecast any of the volatility measures. There is some limited evidence that PVS_t forecasts market volatility in Panel A, but this evidence is not robust across to using the real rate instead.

Overall, the results presented in Tables 8, 9, and 10 suggest that our results are not driven by changes in the quantity of risk. We cannot directly test for time variation in the price of risk. However, our results are most consistent with the idea that the real rate is strongly correlated with time variation in investor aversion to volatility, not time variation in the quantity of volatility.

4 Implications for Monetary Policy

The natural rate of interest – or the real interest rate consistent with output at its natural rate and stable inflation – plays a key role in the design of optimal monetary policy (Woodford, 2003). In a standard New Keynesian framework, it is optimal for the central bank to adjust interest rates to fully accommodate shocks to the natural real rate, but to partly counteract fluctuations due to cost-push shocks that drive up inflation and drive down output, such as wage-markup shocks (Clarida et al., 1999). The link between precautionary savings motives and the real interest rate depends only on the investor's Euler equation and is hence independent of any price-setting frictions. We therefore expect the relation between the real rate and time-varying precautionary savings motives to reflect time-variation in the natural real rate. This logic implies that whether monetary policy is tight or loose should be evaluated relative to a natural real rate that accounts for precautionary savings. If precautionary savings shocks drive the natural real rate, they should have very different implications for output and inflation than independent real rate shocks. This section uses impulse responses to document such differences, thereby providing corroborating evidence that precautionary savings motives are an important component for assessing the stance of monetary policy.

In the simplest New Keynesian model, such as Clarida et al. (1999), output equals consumption, so the Euler equation can be written as (up to a constant)

$$x_t = E_t x_{t+1} - \psi\left(r_t - r_t^n\right). \tag{4}$$

Here, x_t is the output gap between current output and its natural rate, r_t is the actual real rate, r_t^n is the natural real rate, and ψ is the elasticity of intertemporal substitution.

Moreover, in a New Keynesian model, the output gap is linked to inflation through the Phillips curve, such as a forward-looking Phillips curve that arises from staggered price setting, as in Calvo (1983):

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t. \tag{5}$$

Here, $\kappa > 0$ depends on the frequency of firms' price setting and the degree of firms' complementarity in setting product prices. In such a framework, monetary policy can affect the real interest rate in the short term, because it can control the nominal rate and prices are sticky. This simple logic suggests that we would expect very different effects from shocks to r_t versus shocks to the natural real rate r_t^n. The Euler equation (4) and Phillips Curve (5) suggest that an unanticipated increase in the real rate r_t while holding r_t^n constant, should lead to decreases in both output and inflation. On the other hand, the effect of a shock to the natural real rate depends crucially on the monetary policy response. If the central bank fully adjusts interest rates in response to a natural rate shock, as optimal policy would require, output and inflation should not respond at all to such a shock. However, if the central bank does not adjust interest rates fully and contemporaneously, a positive shock to r_t^n may even lead to increases in output and inflation.

4.1 VAR Framework

We estimate impulse responses and test whether the real rate responds contemporaneously to precautionary savings motives, following the identification strategy in Bernanke and Mihov (1998). The question of how to best identify macroeconomic effects of monetary policy shocks is by no means settled. We therefore estimate a VAR that is as simple and transparent as possible, while following a common set of recursiveness assumptions in the spirit of Bernanke and Mihov (1998) and Christiano et al. (1999). The key requirements for our identification strategy are that the output gap and inflation respond to the monetary policy variables - precautionary savings demand for bonds and the real rate – with a lag. While precautionary savings demand is assumed to respond to real rate innovations only with a lag, the real rate is permitted to respond to macroeconomic variables and precautionary savings motives contemporaneously, consistent with the Federal Reserve actively monitoring macroeconomic and financial variables.

We use the following strategy for measuring dynamic effects of monetary policy shocks

$$Y_t = B_1 Y_{t-1} + C_1 P_{t-1} + A^y v_{y,t} \tag{6}$$

$$P_t = \sum_{i=0}^{1} D_i Y_{t-i} + \sum_{i=0}^{1} G_i P_{t-i} + A^p v_{p,t}. \tag{7}$$

Here, Y_t is a vector of quarterly non-policy variables, consisting of the output gap and inflation. P_t is a vector of policy variables consisting of PVS_t and the detrended real rate. Eq. (6) describes a set of structural relationships in the economy, where macroeconomic variables depend on lagged values of macroeconomic and policy variables. Eq. (7) describes the stance of monetary policy conditional on contemporaneous macroeconomic variables. The vector of shocks v_p contains the monetary policy disturbance v_{MP}. By analogy to the treatment of the money demand shock in Bernanke and Mihov (1998), we allow v_p to contain a precautionary savings demand shock for the risk-free asset, proxied for by our measure of precautionary savings demand from the stock market.

We can rewrite the system (6)-(7) in VAR form with only lagged variables on the right-hand-side and estimate by OLS. Let $u_p = A^p v_p$ be the VAR residuals in the policy block that are orthogonal to the VAR residuals in the non-policy block. To recover the structural shocks, including v_{MP}, Bernanke and Mihov (1998) require a specific model relating the VAR residuals and the structural shocks in the policy block. Taking the real rate as the policy instrument and the PVS_t as an indicator of demand for the risk-free bond, we assume that the market for the risk-free bond is described by the following set of equations:

$$u_{PVS} = \alpha v_{MP} + v_{PVS}, \tag{8}$$

$$u_{rr} = \phi v_{PVS} + v_{MP}. \tag{9}$$

Eq. (8) is the innovation in investors' precautionary savings demand for bonds. It states that the demand for low-voatility assets depends on monetary policy shocks v_{MP} and the structural innovation v_{PVS}. Eq. (9) describes central bank behavior. We assume that the Fed observes and responds to precautionary savings shocks. The model described by Eqs. (8) and (9) has four unknown parameters: α, ϕ, and the two structural shock variances, σ^2_{PVS}, and σ^2_{MP}.These need to be estimated from three variances and covariances. We therefore need to impose additional restrictions and estimate two versions of (8) through (9), one of which is a just-identified model, and the other is an over-identified model.

Just-Identified Model ($\alpha = 0$). Similarly to Bernanke and Mihov (1998), we consider the restriction that precautionary savings demand does not respond to monetary policy in the short run

($\alpha = 0$). This identification assumption is plausible if investors' risk preferences shift gradually over time and do not jump in response to monetary policy actions. We make this assumption primarily for identification purposes. However, the plausibility of this assumption is corroborated by the fact that PVS_t changes are not correlated with Romer and Romer (1989)'s monetary policy innovations extracted from the Fed's records.

Constant Intercept. The assumption that the central bank follows a Taylor (1993)-type rule with constant intercept corresponds to the parametric assumption $\phi = 0$, that is the central bank does not allow precautionary savings shocks to enter into the real rate contemporaneously. The monetary policy shock implied by this restriction is $v_{MP} = u_{rr}$, i.e., the policy shock is simply the innovation to the real rate. To ensure that this model is nested by the just-identified model and test over-identifying restrictions, we continue to assume that $\alpha = 0$.

4.2 Estimation

We estimate the model using a two-step efficient GMM procedure, as in Bernanke and Mihov (1998). The first step is an equation-by-equation OLS estimation of the VAR coefficients. The second step consists of matching the second moments to the covariance matrix of the policy block VAR residuals. We apply GMM with the following three moments:

$$E\left[u_{PVS}^2 - \sigma_{PVS}^2\right] = 0, \tag{10}$$

$$E\left[u_{PVS}u_{rr} - \phi\sigma_{PVS}\sigma_{MP}\right] = 0, \tag{11}$$

$$E\left[u_{rr}^2 - \phi\sigma_{PVS}^2 - \sigma_{MP}^2\right] = 0. \tag{12}$$

We estimate the parameters ϕ, σ_{PVS}, and σ_{MP} by two-step GMM using a Bartlett kernel with two lags and the initial weighting matrix equal to the identity.

The hypothesis that the real rate does not react to PVS_t contemporaneously ($\phi = 0$) is clearly rejected both by a Wald test and a Hansen J-test. For the unrestricted model, we estimate a point estimate for $\phi = 1.756$ with a standard error of 0.611. The over-identifying restriction of the *Constant*

Intercept model is rejected at any conventional significance level with a p-value of 0.006. Combined, these tests further corroborate the evidence that demand for precautionary savings has an important effect on the real rate, even on a quarterly innovation basis. The estimated instantaneous reaction coefficient of the real rate to PVS_t is $\phi = 1.756$ is about half the estimated coefficient in our baseline levels regression in Table 2, but similar to the estimated coefficient in changes in Table 3, which is potentially consistent with slow monetary policy adjustment to precautionary savings shocks. Parameter estimates and p-values are similar if we restrict the sample to the pre-crisis period.

4.3 Impulse Response Functions

Fig. 4 displays impulse responses of output and inflation to one-standard-deviation increases in v_{MP} and v_{PVS} for the just-identified model. Dashed lines indicate 95% confidence bands. The first column displays responses to a monetary policy shock. A positive shock v_{MP} corresponds to an unanticipated monetary policy tightening by the central bank, increasing the real rate. Reassuringly, our simple VAR identification scheme produces results that are consistent with a long literature on monetary policy shocks, summarized in Christiano et al. (1999). Following a contractionary monetary policy shock, output decreases immediately, reaching its trough four quarters after the shock period. Inflation also declines in response to a contractionary monetary policy shock, but the response is significantly slower and reaches its trough around seven quarters after the shock. Interestingly, the response of PVS_t to monetary policy shocks does not differ significantly from zero, consistent with the interpretation that shocks to preferences for precautionary savings drive the real rate, and not vice versa.

Contrasting columns (1) and (2) of Fig. 4 shows that a positive shock to v_{PVS} leads to significantly different responses in output and inflation, despite being associated with a similar increase in the real rate as the monetary policy shock. The underlying shock in column (2) is a positive shock to the price of volatile stocks, which corresponds to reduction in the precautionary savings motive and a negative bond demand shock, so the real rate increases. We would expect that such a shock has very different effects on output and inflation, if precautionary savings shocks capture a component of the real rate that is very different from monetary policy innovations. The temporary boom in output and inflation following a PVS_t shock is consistent with the central bank reacting to PVS_t shocks slowly over time, with only a partial adjustment within the quarter. Intuitively, the evidence is consistent with a positive v_{PVS} shock decreasing demand for precautionary savings, that is only partially reflected in bond prices, inducing investors to increase contemporaneous consumption and leading to a temporary output boom.

5 Modeling Framework

This section presents a deliberately stylized model to rationalize the empirical findings. Idiosyncratic risk is priced in the model, because markets are segmented and investors are borrowing-constrained. Broadly interpreted, we think of the segmented markets assumption as representing households or professional investors who take concentrated bets in labor and financial markets. While we do not explicitly model why arbitrageurs do not arbitrage away pricing differentials, it is plausible that short-lived and risk-averse arbitrageurs would have limited arbitrage ability, especially for high-volatility stocks that typically have smaller market capitalization (Lee et al., 1991). Alternatively, Merton (1987) argues that agents' informational differences can drive segmentation in the stock market. Furthermore, segmented stock markets are also consistent with long-standing evidence of home or familiarity bias in individual stock holdings (see Barberis and Thaler (2003) for an overview), which may be either due to superior information or a preference for familiarity.

5.1 Endowments and Preferences

To capture the difference between high-volatility and low-volatility stocks, we assume that the economy consists of two Lucas trees, which are uncorrelated and differ only in terms of volatilities of the endowments. A share p_H of stocks is of the high-volatility type and a share $p_L = 1 - p_H$ is low-volatility. We use lower case letters to denote logs. Both stocks' log dividends are distributed i.i.d. around a common trend to ensure that the shares of both stocks in the economy are stationary:

$$c_{H,t} = \mu t + \varepsilon_{H,t}, \tag{13}$$

$$c_{L,t} = \mu t + \varepsilon_{L,t}, \tag{14}$$

$$\begin{bmatrix} \varepsilon_{L,t} \\ \varepsilon_{H,t} \end{bmatrix} \overset{iid}{\sim} N\left(0, \begin{pmatrix} \sigma_L^2 & 0 \\ 0 & \sigma_H^2 \end{pmatrix}\right). \tag{15}$$

We assume that stock markets are segmented. Stock i is priced by agent of type $i \in \{H, L\}$, who receives an endowment consumption stream equal to $c_{i,t}$. Assuming market segmentation helps us explain the empirical results, because it implies that investor-employees cannot diversify away idiosyncratic consumption risk, which hence generates a precautionary savings motive and is priced into the corresponding stock. In addition, we assume external habit formation preferences (Abel, 1990; Constantinides, 1990; Campbell and Cochrane, 1999; Menzly et al., 2004; Santos and Veronesi, 2010; Wachter, 2006; Lettau and Wachter, 2011). This model element is key to generating predictable stock returns, as documented in Section 3. Simultaneously, by generating time-varying curvature of the utility function, habit formation preferences also imply that the desire to hold the risk-free asset varies over time. This is the source of time-variation in the risk-free rate in our model. Agent i maximizes the expected discounted sum of log consumption utility relative to habit $X_{i,t}$:

$$U_{i,t} = E_t\left[\sum_{\tau=t}^{\infty} \beta^{\tau} \log\left(C_{i,\tau} - X_{i,\tau}\right)\right], \; i = H, L. \tag{16}$$

We define type i surplus consumption and inverse surplus consumption ratios:

$$S_{i,t} = \frac{C_{i,t} - X_{i,t}}{C_{i,t}}, \tag{17}$$

$$G_{i,t} = S_{i,t}^{-1}. \tag{18}$$

We build on the tractable habit dynamics of Menzly et al. (2004), which generates closed-form solutions for both the risk-free rate and equity premia, by assuming that inverse surplus consumption of agent i follows a process of the form:

$$G_{i,t+1} = \kappa\bar{G} + (1-\kappa)G_{i,t} - \alpha\left(G_{i,t} - \lambda\right)\varepsilon_{i,t+1}. \tag{19}$$

Habit dynamics are defined implicitly via (17) through (19). The advantage of specification (19) is that it generates closed-form solutions for asset prices, while capturing the essence of Campbell and Cochrane (1999) habit formation.[12] The key role of habit formation in our model is to generate time-varying risk premia, while holding the quantity of risk constant. We view our empirical findings as consistent with alternative models of time-varying risk premia, such as Bansal et al. (2012), if sources of time-varying quantities of risk are hard to observe.

5.2 Equity Markets

Equities are priced by segmented investor clienteles, with investors of type H trading the high-volatility stock and investors of type L trading the low-volatility stock. We obtain closed-form solutions for the price-dividend ratio of stock i:

$$\frac{P_{i,t}}{C_{i,t}} = a + bS_{i,t}, \tag{20}$$

for positive constants a and b that are given in Online Appendix A. Analogously to Campbell and Cochrane (1999), when agent i's consumption receives a negative shock that lowers consumption close to habit, surplus consumption $S_{i,t}$ is low. This raises the local coefficient of relative risk aversion $G_{i,t}$, driving up risk premia on the risky asset traded by agents of type i.

Finally, we define book-to-market ratios as simply as possible. We assume that a fixed fraction of assets are marked to market each year, so the book value represents an exponentially-weighted

[12] Santos and Veronesi (2010) show that if habit equals an exponentially-weighted moving average of past consumption as in Constantinides (1990) and Detemple and Zapatero (1991), this gives rise to dynamics of the form (19) with more complicated coefficients.

moving average of past stock prices. The book-to-market ratio then is computed as

$$BM_{i,t} = \frac{B_{i,t}}{P_{i,t}},\ B_{i,t} = (1-\rho)\sum_{\tau=1}^{\infty}\rho^{\tau}P_{i,t-\tau}. \tag{21}$$

5.3 Bond Market

Different investors potentially have different valuations for real risk-free one-period bonds, which are available in zero net supply. The real risk-free rate in the model is pinned down by assuming that agents are borrowing-constrained, so the risk-free rate is bid down to the minimum of investors' indifference points. The risk-free rate takes the tractable form:

$$r_{f,t} = min\left\{r_{f,H,t}, r_{f,L,t}\right\}, \tag{22}$$

$$r_{f,i,t} = \mu - \log\beta - \frac{\sigma_i^2}{2} - \varepsilon_{i,t} \tag{23}$$

$$-\log\left((1-\kappa) + \alpha\sigma_i^2 + \left(\kappa\bar{G} - \sigma_i^2\alpha\lambda\right)G_{i,t}^{-1}\right),\ i \in \{H, L\}. \tag{24}$$

The risk-free rate at which investor type i is indifferent about investing in the bond market (24) can either increase or decrease with surplus consumption, depending on whether σ_i^2 is greater or smaller than the threshold $\frac{\kappa\bar{G}}{\lambda\alpha}$. If $\sigma_i^2 > \frac{\kappa\bar{G}}{\lambda\alpha}$, the risk-free rate increases with surplus consumption. Intuitively, when surplus consumption is low, investors become more risk-averse over future endowment shocks, inducing them to save for a riskier future. As a result, the risk-free rate declines at the same time as surplus consumption. Alternatively, if $\sigma_i^2 < \frac{\kappa\bar{G}}{\lambda\alpha}$, the risk-free rate is inversely related to surplus consumption. This case captures a consumer who faces only little uncertainty about his future consumption stream. Therefore, a decline in surplus consumption increases the marginal utility of consumption and the desire to borrow. Since volatility for this consumer is low, this effect is not offset by the consumer wanting to save for a riskier future. As a result, the investor wants to borrow, driving up the risk-free rate when surplus consumption is low.[13]

[13]The intuition is similar to a standard endowment economy with i.i.d. log consumption growth and power utility with risk aversion γ, where a decrease in surplus consumption S_t acts similarly to an increase in risk aversion. In the standard endowment economy, the consumption Euler equation is $r_{f,t} = -\log\beta + \gamma\mu - \frac{\gamma^2}{2}\sigma^2$. If the endowment

We assume that for high-volatility investors, the time-varying precautionary savings effect dominates, while for low-volatility investors, the time-varying intertemporal substitution effect dominates. In addition, we assume that on average, the low-volatility type's time-varying intertemporal substitution effect balances the high-volatility type's desire for precautionary savings:

$$p_H \sigma_H^2 + p_L \sigma_L^2 = \frac{\kappa \bar{G}}{\lambda \alpha}. \tag{25}$$

Since high-volatility investors have a stronger precautionary savings motive, they tend to value the risk-free bond more highly. Thus, they are typically the marginal investor in the bond market. In fact, unless low-volatility investors expect exceptionally low consumption growth from this period to the next, the risk-free asset is priced by high-volatility investors.[14]

5.4 Calibration

We calibrate the model to illustrate that the magnitudes of our empirical findings are within the range of reasonable values. Calibration parameters are reported in Table 11. Most parameters are set to standard values in the literature. We set the discount rate to 0.96, as in Menzly et al. (2004) and the consumption growth rate to 0.03. We set $\lambda = 10$, corresponding to an upper bound for the surplus consumption ratio of 0.1 as in Campbell and Cochrane (1999). The share of high-volatility stocks is $p_H = 0.2$, corresponding to the top quintile of stocks by volatility in the empirical analysis. We set the standard deviations of consumption volatility to $\sigma_H = 0.02$ and $\sigma_L = 0.01$, so high-volatility stocks are subject to twice as much volatility as low-volatility stocks, matching the empirical ratio of return standard deviations of high-volatility and low-volatility portfolios. Conditional on these values, we pick the parameter α, which determines the volatility of marginal utility, to match the empirical equity volatility of the aggregate stock market. Finally, we set the mean-

volatility σ^2 is sufficiently small, the risk-free rate moves positively with utility curvature γ. However, if the endowment volatility σ^2 is large, an increase in utility curvature increases the desire for precautionary savings and leads to a decrease in the risk-free rate.

[14]This follows from the observation that at $\varepsilon_{H,t} = \varepsilon_{L,t} = 0$ we have that $r_{f,t} = r_{f,H,t}$ for any values $G_{H,t}, G_{L,t} \in [0, \lambda^{-1}]$.

reversion parameter κ to a small value 0.01 to maximize the persistence of the log price-dividend ratio. We set the decay parameter for mark-to-market to 0.933, corresponding to a half-life of book assets of 10 years, or a depreciation rate of 7%. Finally, we obtain the average inverse surplus consumption ratio from condition (25).

Table 12 shows simulated model moments from 1000 simulations, each 36 years in length, corresponding to our empirical sample size. Model moments are shown in bold if we cannot reject the null hypothesis that both are equal at the 95% level. The model matches the equity premium, equity volatility, and book-to-market ratio for the aggregate stock market. The aggregate book-to-market ratio is persistent, but less persistent than in the data, despite the low value for κ, a common problem in these types of habit formation models.[15] The risk-free rate in the model is low and comparably volatile to the data.

The third panel in Table 12 shows that in the model the book-to-market ratios and excess returns of low-volatility stocks are lower than for high-volatility stocks. This might at first seem to contrast with the well-known idiosyncratic volatility puzzle (Ang et al., 2006a, 2009), which finds that low-volatility stocks have historically earned higher returns than high-volatility stocks. In our calibrated model, however, it would not be unusual to observe a positive return comparable to that in the data. In fact, 14.1% of our simulations generate low-minus-high volatility excess returns that are as large as the observed data. In addition, a wide range of additional explanations offered in the literature (Johnson, 2004; Fu, 2009; Stambaugh et al., 2015; Hou and Loh, 2016) may further contribute to the high average excess returns on low-volatility stocks.

The bottom panel of Table 12 shows that the model can generate the empirical relation between the risk-free rate and the cross section of equity valuations that we find in the data. Regressing the risk-free rate on the aggregate book-to-market ratio yields a small slope coefficient. But regressing the risk-free rate on the book-to-market spread between low- and high-volatility stocks yields a

[15] Due to the convexity inherent in the analytically convenient Menzly et al. (2004) specification of inverse surplus consumption as a mean-reverting process, while stock market valuations are a function of surplus consumption, further decreases in the mean-reversion parameter κ do not increase persistence of book-to-market ratios. Instead, we face a tension in choosing the volatility of innovations to $G_{i,t}$, because more volatile innovations allow us to match the high volatility of equity returns, but also exacerbate the convex relation between equity valuations and $G_{i,t}$, thereby driving down the persistence of book-to-market ratios.

strong positive coefficient. In addition, the risk-free rate forecasts excess returns on the low-minus-high volatility equity portfolio with empirically reasonable magnitudes.

Intuitively, a decrease in high-volatility investors' surplus consumption makes these investors more risk averse, raising risk premia on high-volatility stocks relative to low-volatility stocks and driving down the low-minus-high volatility book-to-market ratio. At the same time, an increase in high-volatility investors' risk aversion increases their demand for precautionary savings and drives down the risk-free rate. This generates a positive relation between low-minus-high volatility book-to-market and the risk-free rate, as in the data. In the model, time-varying discount rates due to habit formation drive most of the variation in equity valuations, analogously to Campbell and Cochrane (1999), so book-to-market forecasts stock returns. Since the risk-free rate is related to the low-minus-high volatility book-to-market, it then also forecasts excess returns on the low-minus-high volatility portfolio. Finally, the aggregate book-to-market ratio is largely driven by low-volatility investors, who represent the largest share of the market, while the risk-free rate is driven by the surplus consumption ratio of high-volatility investors, who tend to be the marginal risk-free bond investors. Thus, the model generates a low correlation between the aggregate book-to-market and the risk-free rate.

To illustrate the role of segmented markets and time-varying risk premia for generating our main results, column (3) and (4) switch these features off one at a time, while holding all other parameter values constant. Column (3) shows model moments when assets are priced by a representative consumer, who consumes aggregate consumption, and with preferences of the form (16) through (19). With a representative consumer, risk premia for high- and low-volatility stocks move in lockstep, both being determined by the representative agent's surplus consumption ratio. As a result, the representative agent model generates no variation in PVS_t and no predictability in the low-minus-high vol excess return. In addition, in the representative agent model time-variation in the representative agent's surplus consumption ratio drives both equity risk premia and the risk-free rate, generating a counterfactually positive relation between the aggregate book-to-market and the risk-free rate. The assumption of segmented markets is hence essential for generating time-variation in PVS_t, return predictability in low-minus-high volatility excess returns, and for replicating the empirically weak relation between aggregate book-to-market and the risk-free rate.

Column (4) considers the case of segmented investors, who have log utility. This case is nested in our baseline model if we set $\alpha = 0$, $\kappa = 1$, and $\bar{G} = 1$. Consistent with the equity premium puzzle of Mehra and Prescott (1985), log utility generates equity volatility and an equity premium much smaller than observed in the data. Variation in PVS_t and low-minus-high volatility return predictability is extremely small in magnitude, arising only from temporary fluctuations in log dividend growth. As a result, the log utility model implies a very large and negative slope of the risk-free rate onto PVS_t and the risk-free rate forecasts low-minus-high volatility excess returns with a small and negative coefficient, contrary to the data. Time-varying prices of risk are therefore essential for generating the relation between the risk-free rate and risk premia on low-minus-high volatility stocks, that we document in the data.

5.5 Alternative Explanations

The novel empirical results in Section 3 provide clear evidence of time-varying demand for precautionary savings as a significant determinant of time-varying real interest rates. We view the contribution of this paper as distinguishing between broad classes of real interest rate drivers without taking a stand on the exact channel that generates time-varying demand for precautionary savings. The main ingredients of any model that is consistent with our empirical findings are the following. First, idiosyncratic risk must give rise to a precautionary savings motive, so idiosyncratic risk must enter into investors' pricing kernel. Second, the pricing of idiosyncratic risk must vary over time. This can be achieved through habit formation as in Section 5, or through a time-varying distribution of heterogeneous investors as we discuss below. Third, the model must have at least two state variables. This is necessary to match the close relationship between the real rate and PVS_t, while maintaining no relationship between the real rate and aggregate book-to-market.

To give a sense of the range of explanations that could generate our empirical results, this section discusses alternative explanations for the empirical findings in addition to the model described in Section 5. First, we discuss how shifts in wealth between investors with different levels of risk aversion could generate time-varying precautionary savings motives. Second, we discuss a model, in which consumers care about the volatility of individual goods in their consumption basket, thereby linking idiosyncratic stock returns to precautionary savings even in the absence of any market segmentation.

While the model in Section 5 generates time-varying real rates from habit formation preferences, shifts in the wealth distribution towards agents who require a higher price of risk may act similarly on stocks and bonds (Chan and Kogan (2002); Hall (2016); Barro and Mollerus (2014)). We view this channel as a complementary way of generating time-varying attitudes towards precautionary savings and within the same broad class of drivers for real rate variation as our model. Heterogeneity in risk aversion may arise either because agents have different preferences, or because agents' labor income is idiosyncratic, linking idiosyncratic stock return volatility with the quantity of background risk that individual agents are exposed to. Such a model would be able to match our main empirical result if a relative increase of wealth of investors with high risk aversion or high background risk increases the precautionary savings demand priced into the risk-free rate and the cross-section of stocks. At the same time, if different agents are marginal for the majority of stocks than for the risk-free asset, aggregate market valuations could be relatively unrelated to the risk-free rate, as in the data.

The literature on heterogeneous agents has argued that risk-averse international investors are partly responsible for low interest rates (Caballero et al. (2008); Caballero and Krishnamurthy (2009)). To this end, we note that foreign ownership of Federal Treasury debt has increased fairly steadily since the 1970s. PVS_t, on the other hand, isolates a business-cycle frequency component, indicating that we measure a component of real rates that is different from increasing demand from abroad. Similarly, demographics tend to change more steadily than the component of precautionary savings that we isolate from the stock market.

Idiosyncratic stock return volatility may be priced and informative about precautionary savings even in the absence of segmented markets, provided that consumers care about the volatility of individual goods in their consumption basket. If investors do not form habits over an aggregate consumption bundle, but instead get used to individual goods, such as cars, refrigerators, coffee, etc, risk aversion over shocks to a volatile good may increase precautionary savings motives. Such shocks could move the risk-free rate and the risk premium wedge between high- and low-volatility stocks in a manner that is consistent with the data.

6 Conclusion

This paper uses the cross-section of equity valuations to provide new empirical evidence for one broad driver of real interest rates: investors' time-varying demand for precautionary savings. Decomposing time-varying demand for precautionary savings into price of risk and quantity of risk, we find evidence of time-varying attitudes towards uncertainty, but little evidence of a link between the motive for precautionary savings and time-variation in uncertainty itself. We explore the implications of our findings for monetary policy and present a stylized model of segmented equity markets to rationalize these empirical findings. The goal of this paper is to distinguish between broad drivers of real interest rates. Our results indicate that future research on models and drivers of precautionary savings is likely to be fruitful.

References

Abel, A. B. (1990). Asset Prices under Habit Formation and Catching up with the Joneses. *The American Economic Review*, 38–42.

Agarwal, V., W. Jiang, Y. Tang, and B. Yang (2013). Uncovering hedge fund skill from the portfolio holdings they hide. *The Journal of Finance 68*(2), 739–783.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006a). The Cross-Section of Volatility and Expected Returns. *The Journal of Finance 61*(1), 259–299.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006b). The Cross-Section of Volatility and Expected Returns. *The Journal of Finance 61*(1), 259–299.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2009). High Idiosyncratic Volatility and Low Returns: International and Further US Evidence. *Journal of Financial Economics 91*(1), 1–23.

Baker, M. and J. Wurgler (2012). Comovement and Predictability Relationships Between Bonds and the Cross-Section of Stocks. *Review of Asset Pricing Studies 2*(1), 57–87.

Banks, J., R. Blundell, and A. Brugiavini (2001). Risk pooling, Precautionary Saving and Consumption Growth. *Review of Economic Studies 68*(4), 757–779.

Bansal, R., D. Kiku, A. Yaron, et al. (2012). An empirical evaluation of the long-run risks model for asset prices. *Critical Finance Review 1*(1), 183–221.

Barberis, N. and R. Thaler (2003). A survey of behavioral finance. *Handbook of the Economics of Finance 1*, 1053–1128.

Barro, R. J. and A. Mollerus (2014). Safe Assets. *NBER Working Paper wp20652*.

Benartzi, S. (2001). Excessive Extrapolation and the Allocation of 401 (k) Accounts to Company Stock. *The Journal of Finance 56*(5), 1747–1764.

Bernanke, B. S. and K. N. Kuttner (2005). What Explains the Stock Market's Reaction to Federal Reserve Policy? *Journal of Finance 60*(3), 1221–1257.

Bernanke, B. S. and I. Mihov (1998). Measuring monetary policy. *The Quarterly Journal of Economics 113*(3), 869–902.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica 77*, 623–85.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta, and S. Terry (2014). Really Uncertain Business Cycles. *Working Paper, Stanford University, University College London, and Duke University*.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2012). Really uncertain business cycles.

Caballero, R. and A. Krishnamurthy (2009). Global Imbalances and Financial Fragility. *American Economic Review 99*(2), 584–588.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2008). An Equilibrium Model of Global Imbalances and Low Interest Rates. *American Economic Review 98*(1), 358–93.

Calvo, G. A. (1983). Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics 12*(3), 383–398.

Campbell, J. Y. and J. H. Cochrane (1999). By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy 107*(2), 205–251.

Campbell, J. Y. and R. J. Shiller (1988). The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies 1*(3), 195–228.

Carroll, C. D. and A. A. Samwick (1998). How Important is Precautionary Saving? *Review of Economics and Statistics 80*(3), 410–419.

Chan, Y. L. and L. Kogan (2002). Catching Up with the Joneses: Heterogeneous Preferences and the Dynamics of Asset Prices. *Journal of Political Economy 110*(6), 1255–1285.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999). Monetary policy shocks: What have we learned and to what end? *Handbook of Macroeconomics 1*, 65–148.

Clarida, R., J. Galí, and M. Gertler (1999). The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature 37*, 1661–1707.

Cochrane, J. H. (2011). Presidential Address: Discount Rates. *The Journal of Finance 66*(4), 1047–1108.

Cochrane, J. H. (2016, July). Macro-finance. *Working Paper, Stanford University*.

Constantinides, G. M. (1990). Habit Formation: A Resolution of the Equity Premium Puzzle. *Journal of Political Economy*, 519–543.

Cremers, K. and A. Petajisto (2009). How active is your fund manager? a new measure that predicts performance. *The Review of Financial Studies 22*(9), 3329–3365.

Cúrdia, V., A. Ferrero, G. C. Ng, and A. Tambalotti (2015). Has US Monetary Policy Tracked the Efficient Interest Rate? *Journal of Monetary Economics 70*, 72–83.

Detemple, J. B. and F. Zapatero (1991). Asset Prices in an Exchange Economy. *Econometrica 59*(6), 1633–1657.

Duchin, R., T. Gilbert, J. Harford, and C. Hrdlicka (2016). Precautionary Savings with Risky Assets: When Cash is not Cash. *Journal of Finance forthcoming*.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Fu, F. (2009). Idiosyncratic Risk and the Cross-Section of Expected Stock Returns. *Journal of Financial Economics 91*(1), 24–37.

Gromb, D. and D. Vayanos (2010). Limits of arbitrage. *Annual Review of Financial Economics 2*(1), 251–275.

Hall, R. E. (2016, April). Understanding the Decline in the Safe Real Interest Rate. *NBER Working Paper* (22196).

Hansen, L. P. and R. J. Hodrick (1980). Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis. *The Journal of Political Economy*, 829–853.

Hartzmark, S. M. (2016). Economic Uncertainty and Interest Rates. *Review of Asset Pricing Studies 6*(2), 179–220.

Herskovic, B., B. Kelly, H. Lustig, and S. Van Nieuwerburgh (2016). The Common Factor in Idiosyncratic Volatility: Quantitative Asset Pricing Implications. *Journal of Financial Economics 119*(2), 249–283.

Holston, K., T. Laubach, and J. Williams (2016). Measuring the Natural Rate of Interest: International Trends and Determinants.

Hou, K. and R. K. Loh (2016). Have we Solved the Idiosyncratic Volatility Puzzle? *Journal of Financial Economics 121*(1), 167–194.

Johnson, T. C. (2004). Forecast Dispersion and the Cross Section of Expected Returns. *Journal of Finance 59*(5), 1957–1978.

Kacperczyk, M., C. Sialm, and L. Zheng (2005). On the industry concentration of actively managed equity mutual funds. *The Journal of Finance 60*(4), 1983–2011.

Koijen, R. S., H. Lustig, and S. Van Nieuwerburgh (2010). The cross-section and time-series of stock and bond returns. *NBER Working Paper wp15688.*

Laubach, T. and J. C. Williams (2003). Measuring the Natural Rate of Interest. *Review of Economics and Statistics 85*(4), 1063–1070.

Lee, C. M. C., A. Shleifer, and R. H. Thaler (1991). Investor Sentiment and the Closed-End Fund Puzzle. *Journal of Finance 46*(1), 75–109.

Lettau, M. and J. A. Wachter (2011). The Term Structures of Equity and Interest Rates. *Journal of Financial Economics 101*(1), 90–113.

Lusardi, A. (1998). On the Importance of the Precautionary Saving Motive. *American Economic Review 88*(2), 449–453.

McKay, A., E. Nakamura, and J. Steinsson (2016). The Power of Forward Guidance Revisited. *American Economic Review 106*(10), 3133–3358.

Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics 15*(2), 145–161.

Menzly, L., T. Santos, and P. Veronesi (2004). Understanding Predictability. *Journal of Political Economy 112*(1), 1–47.

Merton, R. C. (1987). A simple model of capital market equilibrium with incomplete information. *The Journal of Finance 42*(3), 483–510.

Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *The Journal of finance 32*(4), 1151–1168.

Newey, W. K. and K. D. West (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica 55*(3), 703–708.

Parker, J. A. and B. Preston (2005). Precautionary Saving and Consumption Fluctuations. *American Economic Review 95*(4), 1119–1143.

Riddick, L. A. and T. M. Whited (2009). The corporate propensity to save. *Journal of Finance 64*(4), 1729–1766.

Romer, C. D. and D. H. Romer (1989). Does monetary policy matter? a new test in the spirit of friedman and schwartz. *NBER Macroeconomics Annual 4*, 121–170.

Romer, C. D. and D. H. Romer (2004). A New Measure of Monetary Shocks: Derivation and Implications. *American Economic Review 94*(4), 1055–1084.

Santos, T. and P. Veronesi (2010). Habit formation, the Cross Section of Stock Returns and the Cash-Flow Risk Puzzle. *Journal of Financial Economics 98*(2), 385–413.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance 52*(1), 35–55.

Stambaugh, R. F., J. Yu, and Y. Yuan (2015). Arbitrage Asymmetry and the Idiosyncratic Volatility Puzzle. *Journal of Finance 70*(5), 1903–1948.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester Conference Series on Public Policy*, Volume 39, pp. 195–214. Elsevier.

Wachter, J. A. (2006). A Consumption-Based Model of the Term Structure of Interest Rates. *Journal of Financial Economics 79*(2), 365–399.

Weber, M. (2016). Cash Flow Duration and the Term Structure of Equity Returns. *Working Paper, University of Chicago*.

Woodford, M. (2003). *Interest and Prices*. Princeton University Press.

FIGURES

Figure 1: Book-to-Market Spread Between Low- and High-Volatility Stocks (PVS)

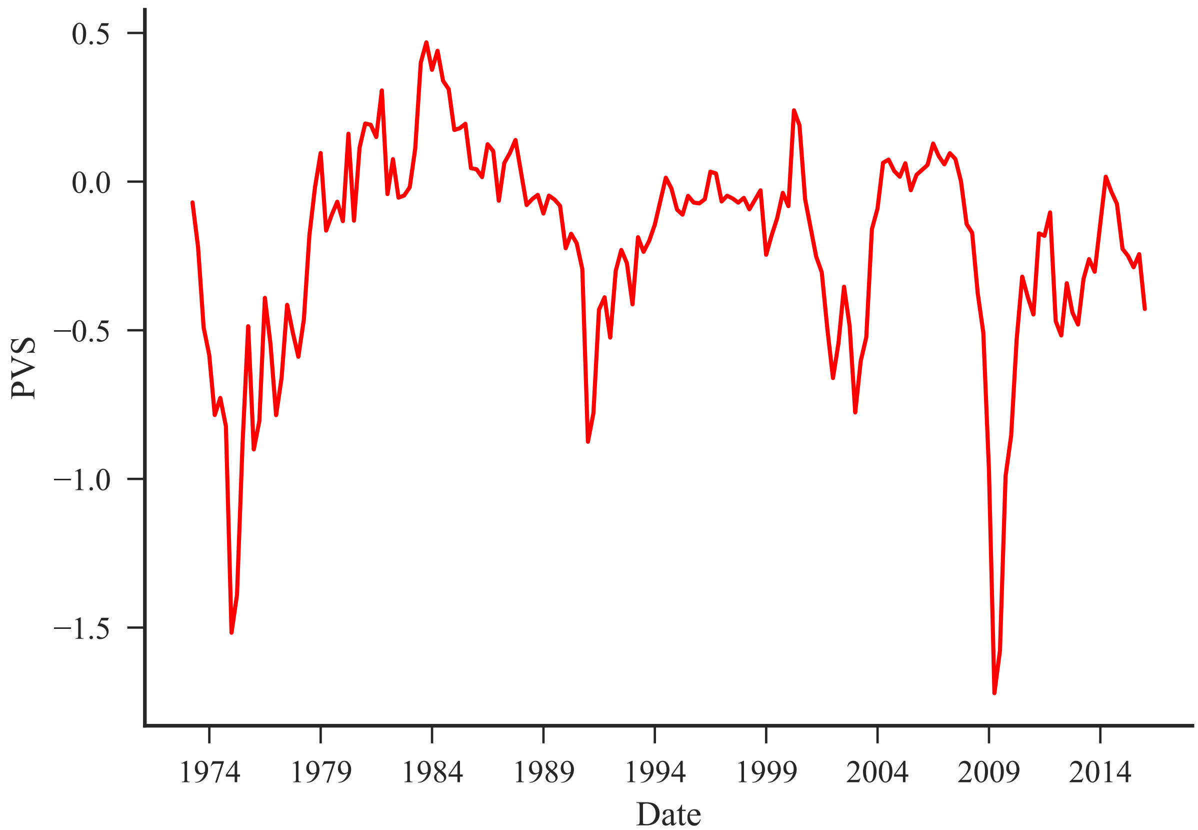

Notes: This figure plots the spread in book-to-market ratios between low and high volatility stocks. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio, which we call PVS_t. The Data Appendix contains full details on how we compute BM ratios. The plotted series is the difference in average book-to-market ratios between the low volatility and high volatility portfolios. Data is quarterly and spans 1973Q1-2015Q4.

Figure 2: One-Year Real Rate: Actual and Fitted Value

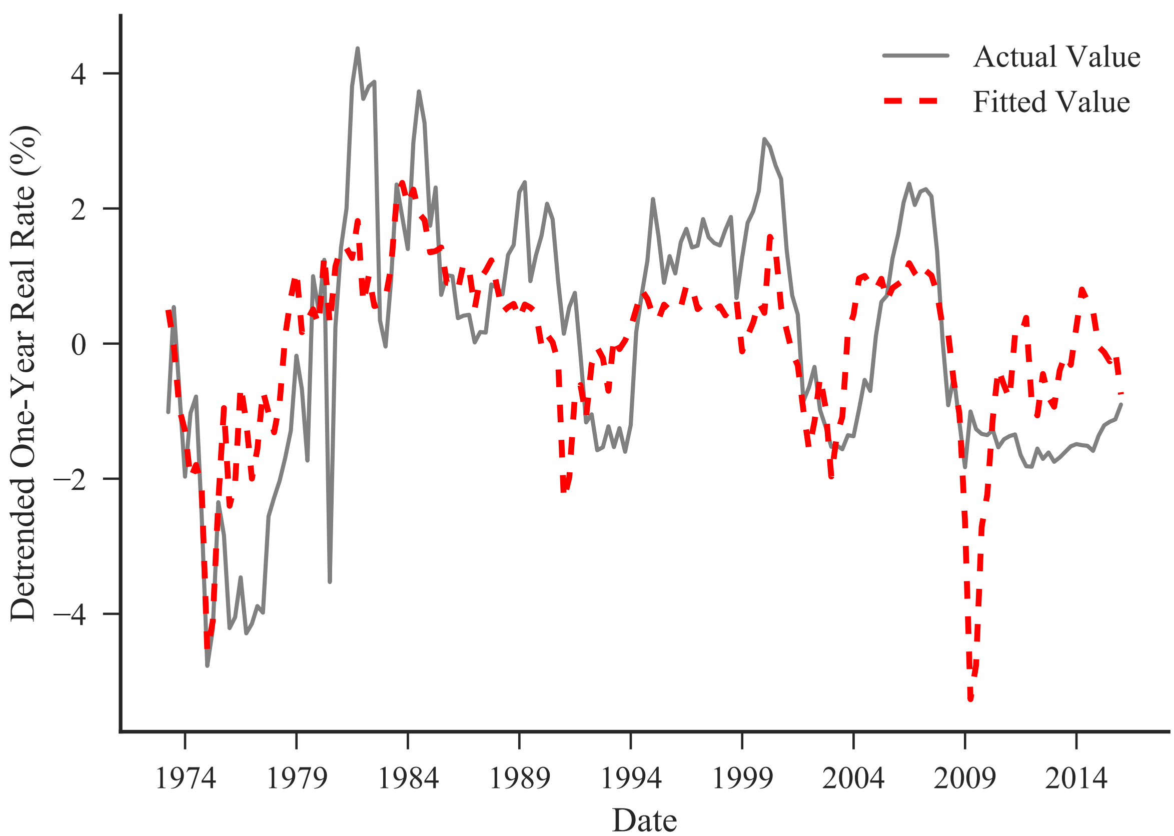

Notes: This figure plots the one-year real rate and the fitted value from a regression of the real rate on the spread in book-to-market ratios between low and high volatility stocks (PVS_t). For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. The one-year real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percentage terms and linearly detrended. Data is quarterly and spans 1973Q1-2015Q4.

Figure 3: Scatterplot of One-Year Real Rate against PVS

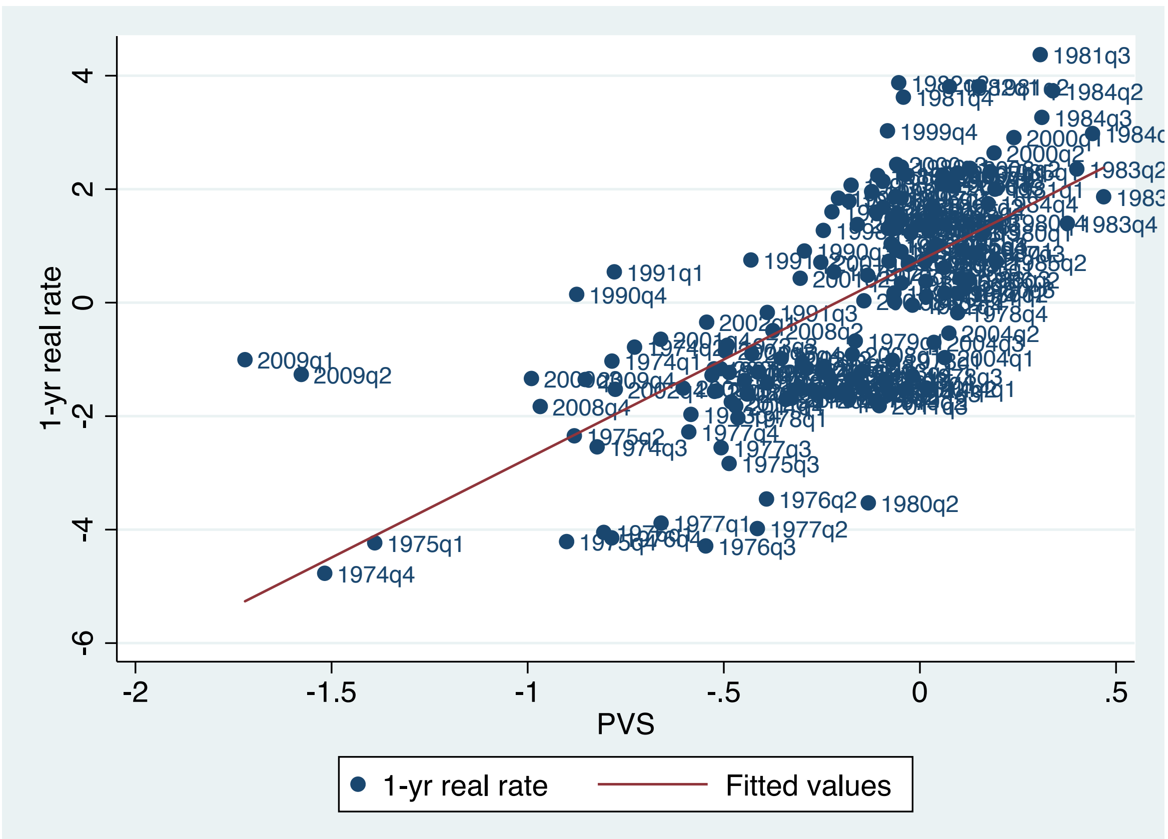

Notes: This figure plots the one-year real interest rate against the spread in book-to-market ratios between low and high volatility stocks (PVS_t). For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. The one-year real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percentage terms and linearly detrended. Data is quarterly and spans 1973Q1-2015Q4.

Figure 4: Impulse Responses to Monetary Policy and PVS Shocks

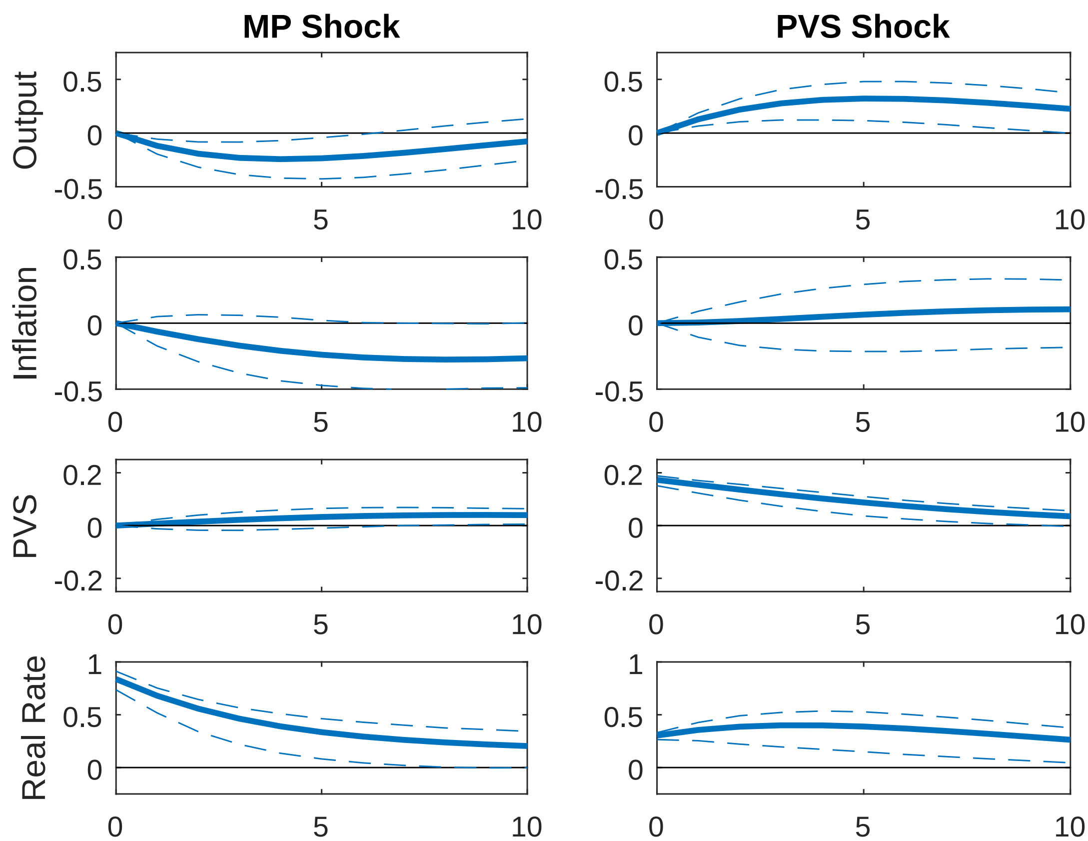

Notes: This figure plots impulse responses to monetary policy and PVS shocks. Impulse responses to one-standard deviation shocks are estimated from a four-variable VAR in the output gap, inflation, PVS, and de-trended real rate with one lag using quarterly data 1973Q1-2015Q4. Following Bernanke and Mihov (1998), structural innovations in the real rate are assumed to affect output, inflation, and precautionary savings demand with a lag. Precautionary savings (PVS) shocks are assumed to affect output and inflation with a lag, but have a contemporaneous effect on the real rate. Dashed lines denote 95% confidence bands, generated by simulating 1000 data processes with identical sample length as in the data from the estimated VAR dynamics.

TABLES

Table 1: Summary Statistics for Volatility-Sorted Portfolios

Panel A: Book-to-Market Ratios

	High-Vol	4	3	2	Low-Vol	Low-High
Mean	1.07	0.88	0.84	0.82	0.86	-0.21
Volatility	0.47	0.34	0.28	0.25	0.27	0.36
Min	0.45	0.48	0.48	0.51	0.54	-1.72
Median	0.93	0.77	0.77	0.73	0.76	-0.11
Max	3.16	2.18	1.85	1.69	1.64	0.47

Panel B: Realized Excess Returns

	High-Vol	4	3	2	Low-Vol	Low-High
Mean	8.89	9.97	12.05	10.77	9.55	0.66
Volatility	39.10	30.70	24.59	19.59	15.15	29.95
Median	3.00	7.04	12.39	13.81	11.73	5.16
Min	-44.87	-37.32	-31.57	-29.27	-22.19	-51.02
Max	76.69	56.20	43.42	34.54	25.67	50.52
α	-4.58	-1.20	2.27	1.90	2.34	6.91
$t(\alpha)$	-1.77	-0.98	3.07	2.13	1.82	2.13
CAPM-β	1.29	1.18	1.03	0.92	0.74	-0.55
SMB-β	1.84	1.37	1.01	0.60	0.34	-1.50
HML-β	0.23	0.11	0.19	0.34	0.37	0.14

Notes: This table presents summary statistics for portfolios formed on volatility. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Panel A shows summary statistics on the average book-to-market (BM) ratio within each quintile. The Data Appendix contains full details on how we form portfolios and compute book-to-market ratios. Panel B displays summary statistics on the realized excess returns of each quintile (in percentage terms). The α is the (annualized) intercept from a regression of excess returns on the Fama and French (1993) factors. Standard errors are computed via GMM by pooling all portfolios. We allow for within-portfolio heteroskedasticity and cross-portfolio correlations. The mean, volatility, and median returns are all annualized. Data is quarterly and runs from 1973Q1 to 2015Q4. The riskless rate for computing excess returns and quarterly returns on the Fama and French (1993) factors are aggregated using monthly data from Ken French's website.

Table 2: What Explains Real Rate Variation (Level Regression)?

Dependent Variable:		One-Year Real Rate	
	(1)	(2)	(3)
PVS_t	3.49**		3.25**
	(4.59)		(3.73)
Aggregate BM		-1.24	0.32
		(-0.56)	(0.17)
Output Gap			0.11
			(0.77)
Inflation			-0.11
			(-0.93)
Constant	0.74**	0.75	1.04
	(2.64)	(0.63)	(1.19)
Adj. R^2	0.44	0.02	0.45
N	172	172	172

Notes: This table reports regression estimates of the one-year real rate on the spread in book-to-market (BM) ratios between high volatility and low volatility stocks (PVS_t). For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. PVS_t is defined as the difference in BM ratios between the bottom and top quintile portfolios. Aggregate BM is computed by summing book equity values across all firms and divided by the corresponding sum of market equity values. The output gap is the percentage deviation of real GDP from the CBO's estimate of potential real GDP. Inflation is the percentage growth in implicit GDP price deflator from the St. Louis Fed (GDPDEF). The one-year real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Standard errors are computed using both Newey-West (1987) and Hansen-Hodrick (1980) with five lags, and we report the more conservative t-statistic of the two. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. Data is quarterly and spans 1973Q1-2015Q4.

Table 3: What Explains Real Rate Variation (First-Differenced Regression)?

Dependent Variable:	1-Qtr ΔReal Rate		
	(1)	(2)	(3)
ΔVol BM Spread	1.52**		1.15*
	(2.26)		(1.81)
ΔAgg. BM Ratio		-4.82*	0.89
		(-1.94)	(0.39)
ΔOutput Gap			0.31**
			(2.84)
ΔInflation			-0.07
			(-1.01)
Constant			0.01
			(0.23)
Adj. R^2	0.09	0.04	0.14
N	171	171	171

Notes: This table reports regression estimates of the one-year real rate on the spread in book-to-market (BM) ratios between high volatility and low volatility stocks (PVS_t). For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. PVS_t is defined as the difference in BM ratios between the bottom and top quintile portfolios. Aggregate BM is computed by summing book equity values across all firms and divided by the corresponding sum of market equity values. The output gap is the percentage deviation of real GDP from the CBO's estimate of potential real GDP. Inflation is the percentage growth in implicit GDP price deflator from the St. Louis Fed (GDPDEF). The one-year real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Columns (1)-(3) run the regression in first differences (i.e. one-quarter changes). Columns (4)-(6) run the regression in quarter-on-quarter differences (i.e. four-quarter changes). Standard errors are computed using both Newey-West (1987) and Hansen-Hodrick (1980) with five lags, and we report the results from the specification that delivers the most conservative standard error for the PVS_t. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. Data is quarterly and spans 1973Q2-2015Q4.

Table 4: Robustness: The Real Rate and PVS_t

		Levels						4Q Changes						Q-Q Changes					
		Full			Pre-Crisis			Full			Pre-Crisis			Full			Pre-Crisis		
		b	$t(b)$	R^2	b	$t(b)$	R^2	b	$t(b)$	R^2	b	$t(b)$	R^2	b	$t(b)$	R^2	b	$t(b)$	R^2
(1)	Baseline	**3.45**	4.63	0.45	**3.95**	6.80	0.51	**1.77**	2.22	0.12	**3.32**	4.80	0.24	*1.40*	1.85	0.09	**2.25**	2.97	0.15
(2)	VW	**2.99**	3.92	0.34	**3.63**	5.06	0.42	*0.97*	1.92	0.06	**1.67**	2.46	0.09	0.70	1.31	0.06	0.94	1.42	0.08
(3)	2YR Vol	**4.81**	5.39	0.57	**5.09**	6.04	0.56	**2.61**	2.64	0.17	**4.14**	4.68	0.28	1.16	1.22	0.05	**2.35**	2.27	0.09
Horse-Races																			
(4)	Duration	**3.21**	3.25	0.45	**4.38**	6.20	0.51	*1.77*	1.90	0.12	**3.35**	4.49	0.23	**1.26**	2.16	0.09	**2.00**	3.65	0.16
(5)	Leverage	**4.81**	6.63	0.49	**5.25**	9.25	0.55	**2.24**	3.01	0.13	**3.52**	5.42	0.23	1.33	1.51	0.09	**1.98**	2.01	0.15
(6)	Beta	**2.53**	3.14	0.47	**3.08**	5.86	0.54	**1.97**	2.29	0.12	**3.24**	4.47	0.23	**1.85**	2.45	0.11	**2.58**	3.21	0.17
(7)	LR Beta	**2.69**	2.73	0.46	**3.74**	3.85	0.50	*1.94*	1.90	0.12	**4.05**	4.87	0.24	*1.37*	1.87	0.09	**2.36**	3.42	0.15
(8)	2M-Beta	**3.51**	5.31	0.45	**3.90**	7.31	0.51	**1.92**	2.56	0.12	**3.33**	4.47	0.23	**1.66**	2.00	0.09	**2.58**	3.07	0.15
(9)	CF Beta	**3.65**	5.25	0.47	**4.00**	6.54	0.50	**1.75**	2.20	0.12	**3.26**	4.89	0.23	*1.43*	1.95	0.09	**2.28**	2.87	0.15
(10)	Size	*3.02*	1.88	0.45	**4.88**	4.24	0.51	*1.84*	1.66	0.12	**3.52**	3.92	0.23	1.65	1.43	0.09	*2.22*	1.70	0.15
(11)	Value	**4.84**	4.76	0.48	**5.73**	8.51	0.55	**2.51**	3.01	0.13	**3.89**	6.78	0.24	*1.72*	1.85	0.09	**2.10**	2.03	0.15
Double-Sorts																			
(12)	Duration	**4.29**	4.26	0.24	**4.52**	4.64	0.22	*1.70*	1.79	0.06	**2.77**	2.20	0.10	**2.00**	2.00	0.09	**2.81**	2.69	0.13
(13)	Leverage	**5.03**	4.98	0.45	**5.72**	7.47	0.52	**2.93**	2.58	0.15	**4.99**	4.78	0.26	*1.63*	1.94	0.07	**2.65**	3.05	0.11
(14)	2M-Beta	**4.36**	5.18	0.45	**4.76**	7.03	0.48	**2.46**	2.29	0.13	**4.48**	4.72	0.24	0.69	0.76	0.04	1.39	1.47	0.06
(15)	Size	**5.18**	3.71	0.38	**6.52**	6.84	0.48	**2.65**	2.04	0.11	**5.08**	3.80	0.21	*2.06*	1.91	0.08	**2.99**	2.61	0.12
(16)	Value	**9.12**	4.80	0.42	**9.96**	5.57	0.43	**5.29**	2.50	0.15	**9.26**	4.22	0.26	*4.00*	1.89	0.09	**6.39**	2.65	0.14
(17)	Industry-Adj	**3.71**	4.89	0.40	**4.04**	5.99	0.42	**1.47**	1.98	0.08	**2.88**	3.45	0.16	0.91	1.48	0.05	**1.69**	2.76	0.10

Notes: This table reports a batter of robustness exercises for our main results. Specifically, we report time-series regression results of the following form: Real Rate$_t = a + b \times \text{PVS}_{t\,t} + \theta X_t + \varepsilon_t$. We run this regression directly in levels and in Quarter-on-Quarter changes (Q-Q Changes). For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. PVS_t is defined as the difference in BM ratios between the bottom and top quintile portfolios. X_t is a vector of control variables, which always includes the Aggregate BM, computed by summing book equity values across all firms and divided by the corresponding sum of market equity values. Row (1) uses our baseline PVS_t measure and the full sample. Row (2) uses value weights instead of equal weights when forming our PVS_t. Row (3) constructs our PVS_t using the past two years of return volatility, as opposed to the past two months. Columns (4)-(11) run bivariate horse races by adding book-to-market spreads based on other characteristic sorts to our control variables X_t. See the Online Appendix for a description of each characteristic. In rows (12)-(17), we instead use a double-sorted PVS_t, with complete details also contained in the Online Appendix. The one-year real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Standard errors are computed using Hansen-Hodrick (1980) with five lags. Italic point estimates indicates a p-value of less than 0.1 and bold indicates a p-value of less than 0.05. Data is quarterly and the full sample spans 1973Q1-2015Q4 (pre-crisis ends in 2008Q4).

Table 5: Forecasting Returns of Portfolios Sorted on Volatility

Panel A - Quarterly Forecasting		Returns$_{t \to t+1}$		
	(1)	(2)	(3)	(4)
PVS_t	16.32**	9.45**		
	(5.98)	(4.17)		
Real Rate$_t$			1.92**	0.56
			(2.85)	(1.17)
Constant	3.58**	3.69**	0.11	1.81**
	(3.13)	(4.37)	(0.10)	(2.31)
Fama-French$_{t \to t+1}$	N	Y	N	Y
Adj. R^2	0.15	0.59	0.05	0.55
N	171	171	171	171

Panel B - Annual Forecasting		Returns$_{t \to t+4}$		
	(1)	(2)	(3)	(4)
PVS_t	45.92**	30.03**		
	(4.09)	(2.46)		
Real Rate$_t$			5.77**	2.58
			(2.52)	(0.91)
Constant	9.51**	9.01*	-0.33	3.91
	(2.25)	(1.91)	(-0.07)	(0.92)
Fama-French$_{t \to t+4}$	N	Y	N	Y
Adj. R^2	0.31	0.61	0.13	0.52
N	168	168	168	168

Notes: This table reports forecasting regressions of portfolios formed on volatility. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. PVS_t is defined as the difference in BM ratios between the bottom and top quintile portfolios. Returns in the forecasting regression correspond to the low-minus-high volatility portfolio. The real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Columns (2) and (4) include the three Fama-French factors as controls. For quarterly regressions, standard errors are computed using both Newey-West (1987) and Hansen-Hodrick (1980) with five lags, aand we report the results from the specification that delivers the most conservative standard error for the PVS_t or the real rate. For annual horizons we use Hodrick (1992) standard errors. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. Data is quarterly and spans 1973Q2-2015Q4. Returns are in percentage points.

Table 6: Forecasting Annual Returns of the Fama-French Factors

	$\text{Mkt-Rf}_{t\to t+4}$		$\text{SMB}_{t\to t+4}$		$\text{HML}_{t\to t+4}$	
	(1)	(2)	(3)	(4)	(5)	(6)
PVS_t	-8.40		-11.16**		1.44	
	(-1.16)		(-3.25)		(0.29)	
Real Rate_t		-0.49		-2.13**		0.38
		(-0.39)		(-3.06)		(0.41)
Constant	5.38*	7.17**	0.31	2.69*	4.95**	4.64**
	(1.95)	(2.75)	(0.21)	(1.86)	(2.39)	(2.58)
Adj. R^2	0.02	0.00	0.15	0.15	-0.00	-0.00
N	168	168	168	168	168	168

Notes: This table reports forecasting regressions of annual returns on the Fama and French (1993) factors. PVS_t is the spread in book-to-market ratios between stocks with low and high volatility. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. The real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Standard errors are computed according to Hodrick (1992). * indicates a p-value of 0.1 and ** indicates a p-value of 0.05. Data is quarterly and spans 1973Q2-2015Q4. Returns are in percentage points.

Table 7: Forecasting ROE of Volatility-Sorted Portfolios

Dep. Variable	$ROE_{t \to t+4}$	
	(1)	(2)
PVS_t	-2.72	
	(-0.82)	
Real Rate$_t$		0.51
		(0.86)
Constant	9.99**	10.55**
	(6.30)	(6.31)
Adj. R^2	0.01	0.00
N	168	168

Notes: This table reports ROE forecasting regressions of the portfolios formed on volatility. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. PVS_t is defined as the difference in BM ratios between the bottom and top quintile portfolios. ROE in the forecasting regression correspond to the low-minus-high volatility portfolio, which we compute following Cohen, Polk, and Vuolteenaho (2003). The real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Standard errors are computed using both Newey-West (1987) and Hansen-Hodrick (1980) with five lags, and we report the results from the specification that delivers the most conservative standard error for the PVS_t or the real rate. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. Data is quarterly and spans 1973Q2-2015Q4. ROEs are in percentage points.

Table 8: Contemporaneous Real Rate Variation and the Quantity of Risk

Dependent Variable:	Real Rate (Level)			4-Q ΔReal Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
σ (LMH-Vol Portfolio)	-0.00		-0.02	-0.00		0.01
	(-0.08)		(-0.67)	(-0.08)		(0.42)
σ (TFP Growth)		-0.09	0.21		-0.37*	-0.09
		(-0.18)	(0.93)		(-1.85)	(-0.48)
σ (Mkt-Rf)		-0.19**	-0.06		-0.06*	-0.06**
		(-3.20)	(-1.63)		(-1.79)	(-2.07)
σ (SMB)		0.28**	0.05		0.09*	0.04
		(3.41)	(1.37)		(1.92)	(0.99)
σ (HML)		0.10	0.12**		0.11	0.10
		(1.03)	(2.68)		(1.50)	(1.33)
CIV_t		0.01	0.04*		-0.04*	-0.01
		(0.49)	(1.83)		(-1.93)	(-0.59)
PVS_t			4.00**			1.42**
			(7.62)			(2.54)
Adj R^2	-0.01	0.14	0.57	-0.01	0.10	0.16
N	172	172	172	168	168	168

Notes: This table reports regression estimates of the one-year real rate on various measures of risk. σ (TFP Growth) is the volatility of TFP growth that is implied by a GARCH model (see Table A1 of the Online Appendix). σ (Mkt-Rf), σ (SMB), and σ (HML) are the within-quarter annualized volatility (percentage terms) of the three Fama and French (1993) factors, which we compute using daily data. CIV_t is the average idiosyncratic volatility factor of Herskovic et al. (2016). PVS_t is the difference in book-to-market ratios between high volatility and low volatility stocks. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. σ (LMH-Vol Portfolio) is the annualized percentage volatility of the low-minus-high volatility portfolio, which we compute using daily data. The real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Columns (1)-(3) run the regression in levels. Columns (4)-(6) run the regression in quarter-on-quarter differences (i.e. four-quarter changes). Standard errors are computed using both Newey-West (1987) and Hansen-Hodrick (1980) with five lags. We report the more conservative t-statistic of the two. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. All regressions have a constant, but we omit the estimates to save space. Data is quarterly and spans 1973Q1-2015Q4.

Table 9: Forecasting Returns of Volatility-Sorted Portfolios with the Quantity of Risk

Dependent Variable:	Returns$_{t\to t+4}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Real Rate$_t$		5.77**	5.58**	5.41**	5.84**	5.92**	3.89**
		(2.50)	(2.62)	(2.38)	(2.64)	(2.61)	(2.19)
σ (LMH-Vol Portfolio)	-0.18	-0.17					0.72
	(-0.20)	(-0.22)					(0.74)
σ_t (TFP Growth)			-3.03				-0.20
			(-0.60)				(-0.04)
σ_t (Mkt-Rf)				-0.94			-2.59**
				(-1.58)			(-2.29)
σ_t (SMB)					-0.30		4.00*
					(-0.25)		(1.94)
σ_t (HML)						-1.30	-0.63
						(-1.05)	(-0.36)
Constant	1.79	1.61	9.30	13.56*	2.12	8.73	2.87
	(0.21)	(0.20)	(0.54)	(1.77)	(0.25)	(1.14)	(0.16)
Adj. R^2	-0.01	0.13	0.13	0.18	0.13	0.16	0.27
N	168	168	168	168	168	168	168

Notes: This table reports annual return forecasting regressions of portfolios formed on volatility. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Returns in the forecasting regression correspond to the low-minus-high volatility portfolio. σ (LMH-Vol Portfolio) is the realized return volatility for the low-minus-high volatility portfolio, which we compute using daily data. σ (TFP Growth) is the volatility of TFP growth that is implied by a GARCH model (see Table A1 of the Online Appendix). σ (Mkt-Rf), σ (SMB), and σ (HML) are the within-quarter volatility of the three Fama and French (1993) factors, which we compute using daily data. All volatility measures are annualized and expressed in percentage terms. The real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Column (2) includes the three Fama-French factors as controls. Standard errors are computed according to Hodrick (1992). * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. Data is quarterly and spans 1973Q2-2015Q4. Returns are in percentage points.

Table 10: Forecasting Realized Volatility

Panel A: Forecasting Volatility Using the PVS_t

Dependent Variable:	Realized Volatility$_{t\rightarrow t+1}$				
	LMH-Vol	TFP	MktRf	SMB	HML
	(1)	(2)	(3)	(4)	(5)
PVS_t	-2.51	-0.64	-4.20*	0.21	-2.54
	(-0.96)	(-1.10)	(-1.68)	(0.14)	(-0.87)
Constant	10.77**	3.03**	13.98**	8.14**	6.48**
	(7.45)	(9.18)	(12.73)	(13.15)	(8.14)
Adj. R^2	0.01	0.07	0.03	-0.00	0.04
N	171	171	171	171	171

Panel B: Forecasting Volatility Using the Real Rate

Dependent Variable:	Realized Volatility$_{t\rightarrow t+1}$				
	LMH-Vol	TFP	MktRf	SMB	HML
	(1)	(2)	(3)	(4)	(5)
Real Rate$_t$	0.09	-0.07	-0.05	0.32	0.23
	(0.16)	(-0.61)	(-0.09)	(1.12)	(0.61)
Constant	11.30**	3.17**	14.87**	8.09**	7.02**
	(10.16)	(16.39)	(12.00)	(16.53)	(8.64)
Adj. R^2	-0.01	0.02	-0.01	0.02	0.00
N	171	171	171	171	171

Notes: This table reports forecasting regressions of realized volatility. TFP volatility is the volatility of TFP growth that is implied by a GARCH model (see Table A1 of the Online Appendix). MktRf Vol, SMB Vol, and HML Vol are the within-quarter volatility of the three Fama and French (1993) factors, which we compute using daily data. PVS_t is the difference in book-to-market ratios between high volatility and low volatility stocks. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio. The Data Appendix contains full details on how we compute BM ratios. LMH-Vol is the realized return volatility of the low-minus-high volatility portfolio, which we compute using daily data. All volatility measures are expressed in annualized percentage terms. The real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percent and linearly detrended. Standard errors are computed using both Newey-West (1987) and Hansen-Hodrick (1980) with five lags, and we report the more conservative t-statistic of the two. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. Data is quarterly and spans 1973Q2-2015Q4.

Table 11: Model Parameters

Variable Name	Parameter	Value
Share High-Volatility Stocks	p_H	0.20
Discount Rate	β	0.96
Consumption Growth	μ	0.03
Lower Bound G	λ	10
Heteroskedasticity Parameter	α	350
Average G	$\bar{G}$	56
Mean-Reversion G	κ	0.01
High Consumption Vol.	σ_L	0.01
Low Consumption Vol.	σ_H	0.02
Decay Parameter Mark-to-Market	ρ	0.933

Notes: This table displays the parameter values for the calibrated version of the model in Section 4.

Table 12: Model Moments

		(1) Data	(2) Model	(3) Rep. Agent	(4) Log Utility
Equity Market					
Equity Premium	$E\left(r^{e}_{t+1} - r_{f,t}\right)$	5.52	**6.03**	**4.31**	0.91
Equity Volatility	$Std\left(r^{e}_{L,t+1}\right)$	17.90	**15.64**	**19.78**	1.27
Agg. Book/Market Ratio	$E\left(B_t/P_t\right)$	0.62	**0.63**	**0.64**	0.62
AR(1) Agg. Book/Market	$AR\left(B_t/P_t\right)$	0.91	0.48	0.47	**0.92**
Risk-Free Rate					
Risk-Free Rate	$E(r_{f,t})$	2.40	0.72	3.08	6.18
Std. Risk-Free Rate	$Std(r_{f,t})$	1.90	**1.76**	1.37	1.30
Low-Minus-High Volatility Portfolio					
PVS_t	$E\left(B_{L,t}/P_{L,t} - B_{H,t}/P_{H,t}\right)$	-0.23	-0.01	0.00	0.00
Std. PVS_t	$Std\left(B_{L,t}/P_{L,t} - B_{H,t}/P_{H,t}\right)$	0.35	0.25	0.00	0.01
Low-Minus-High Return	$E\left(r_{L,t} - r_{H,t}\right)$	1.32	**-1.24**	0.00	-0.02
The Risk-Free Rate and Equity Risk Premia					
Risk-Free Rate on Agg. Book-Market	$slope\left(r_{f,t}, \frac{B_t}{P_t}\right)$	2.05	**-1.11**	7.77	**9.87**
Risk-Free Rate on PVS_t	$slope\left(r_{f,t}, \frac{B_{L,t}}{P_{L,t}} - \frac{B_{H,t}}{P_{H,t}}\right)$	3.36	**3.35**	NaN	-45
Return Spread on Lag Risk-Free Rate	$slope\left(r_{L,t+1} - r_{H,t+1}, r_{f,t}\right)$	6.02	**5.22**	0.00	-0.88

Notes: Model moments are averaged over 1000 simulations of length 36 years. Simulations use a burn-in period of 20 years. Bold indicates a one-sided p-value > 0.05. One-sided p-values are computed as the percentage of simulations where the model moments is less than the data moment. The entries corresponding to the risk-free rate and equity risk premia report regression results. For instance, "Slope Risk-Free Rate on Book-Market" reports the estimated coefficient from a regression of the risk-free rate on the aggregate book-to-market ratio. Average equity premium and low-minus-high return include Jensen's inequality adjustments. Column (3) reports model moments when assets are priced by a representative agent with habit formation preferences. Column (4) reports model moments when assets are priced by segmented investor clienteles with log utility, switching off the habit formation channel in the baseline model.