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

Carolin Pflueger† Emil Siriwardane‡ Adi Sunderam§

October 31, 2017

Abstract

We document a strong and robust relationship between the one-year real rate and the valuation of high-volatility stocks, which we contend measures precautionary savings motives. Our novel proxy for precautionary savings explains 41% of the variation in the real rate. In addition, the real rate forecasts returns on the low-minus-high volatility portfolio but has little relation to observable measures of the quantity of risk. These results suggest that precautionary savings motives, and thus the real rate, are driven by time-varying attitudes towards risk. Our findings are difficult to rationalize in models with perfect risk sharing and highlight the role that imperfect diversification plays in determining interest rates. We also explore the implications of our findings for monetary policy, arguing that precautionary savings motives should be included in assessing the natural real rate.

*We thank Michael Brennan (discussant), John Campbell, Robert Engle, Xavier Gabaix, Espen Henriksen (discussant), Bryan Kelly, Hanno Lustig (discussant), Thomas Maurer (discussant), Monika Piazzesi, Robert Ready (discussant), Martin Schneider, Andrei Shleifer, Jeremy Stein, Luis Viceira, and seminar participants at the BI-SHoF Conference 2017, CEF 2017, CITE 2017, SITE 2017, FRBSF conference on Advances in Finance Research 2017, London School of Economics, Federal Reserve Board, University of British Columbia, University of Indiana, SFS Cavalcade, and HEC-McGill Winter Finance Workshop for helpful comments. 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

Real interest rates vary considerably over time. The quarterly standard deviation of the one-year real interest rate has been 2.3 percentage points over the last 40 years, large relative to an average level of 1.9 percent. Understanding the forces that drive this variation is critical because the real interest rate is arguably the key asset price in finance and macroeconomics, sitting at the heart of consumption, investment, and savings decisions. Despite its importance in determining macroeconomic outcomes, the question of what moves the real rate is an open one.

Frictionless models of the macroeconomy typically point to two potential drivers. Movements in the real rate can either reflect intertemporal smoothing, with investors borrowing more to smooth consumption when expected growth is high, driving up the real rate. Or they can reflect changes in precautionary savings motives, with investors saving more when uncertainty or aversion to uncertainty is high, driving down the real rate. Large literatures have studied both channels, using both aggregate and micro-level consumption data. But these efforts have not proven conclusive.

In this paper, we show that changes in investors' precautionary savings motives are an important source of real rate variation. Precautionary motives may vary either because of changes in the uncertainty faced by investors or because of changes in investors' aversion to that uncertainty. We show that variation in precautionary motives is primarily driven by variation in aversion to uncertainty. Importantly, we find that investors display time-varying aversion to risks that could in principle be diversified away.

Our results rest on two key empirical innovations. First, we use asset prices, specifically the cross section of stock market valuation ratios, to overcome challenges in measuring time-varying precautionary savings. Asset prices have several advantages over the standard approach of measuring precautionary savings with household consumption and investment data (e.g., Carroll and Samwick (1998); Lusardi (1998); Banks et al. (2001); Parker and Preston (2005)). For one, asset prices aggregate over a much wider range of economic agents. In addition, they are available at a much higher frequency than income or consumption. Moreover, asset prices are unique in that they allow us to estimate investors' willingness to pay to avoid uncertainty, whereas consumption and income data generally only allow us to estimate the quantity of uncertainty investors face.

Second, we start with the idea that investors are imperfectly diversified. If investors are differentially exposed to individual stocks, for instance due to market segmentation among professional investors or undiversifiable household labor income risk, the price of volatile stocks (henceforth "PVS_t") relative to low-volatility stocks provides a gauge of precautionary savings motives in the stock market. Intuitively, time-variation in aggregate precautionary savings should disproportionately reflect those investors who have the strongest precautionary saving motives, i.e., those who are exposed to the greatest portfolio volatility. An increase in high-volatility investors' precautionary savings motives should then make them less willing to hold volatile assets and increase their demand for risk-free bonds, leading to simultaneous declines in the price of volatile stocks and real rates. In contrast, low-volatility investors' bond demand may reflect competing time-varying intertemporal substitution and precautionary savings motives, leading to an ambiguous relation between low-volatility stock valuations and real rates. This intuition suggests that, by taking the difference between the price of volatile stocks minus less volatile stocks, PVS_t isolates precautionary savings demand in the stock market from other market-wide drivers of equity valuations.

We begin by establishing several new 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.[1] 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 idiosyncratic volatility were a major driver of risk-free bond valuations. The relationship is remarkably consistent through very different macroeconomic environments, economically significant, and robust in both levels and first differences. The headline result of the paper is that PVS_t explains 41% of the variation in the real rate from 1970 to 2016.

Our emphasis on the cross section is important, as the valuation of the aggregate stock market

[1] We define PVS_t as the average book-to-market ratio of low-volatility stocks minus the average book-to-market ratio of high-volatility stocks. This definition is consistent with the variable name PVS_t or the "Price of Volatile Stocks," because an increase in the market value of high-volatility stocks leads to a decrease in the book-to-market ratio for high-volatility stocks and an increase in PVS_t.

has little explanatory power for the real rate. This second finding echoes the Campbell and Ammer (1993) result that most variation in the aggregate stock market cannot be attributed to real rate news. We therefore conclude that PVS_t is not just another proxy for risk aversion to aggregate market fluctuations. While PVS_t is related to other variables that capture the pricing of idiosyncratic risk, notably the price of credit risk, its explanatory power for the real rate is new and stronger, relative to these other measures.

We favor the precautionary savings interpretation of our headline result because volatility is critical to our findings. Sorting on other characteristics, such as beta, which should matter if investors are perfectly diversified, cannot explain the real rate nearly as well as sorting on volatility. Using both horse races and double sorts, we show that PVS_t contains information about real rate variation that is independent of valuation-ratio spreads generated from sorting stocks based on size, value, leverage, duration of cash flows, CAPM beta, and cash flow beta, all characteristics that are known to describe the cross section of stock returns.[2] The link between the real rate and PVS_t is robust to whether we sort by stock return volatility over the past two months or past two years, indicating that the results are not driven by stocks quickly rotating in and out of high- and low-volatility portfolios. Furthermore, PVS_t retains its explanatory power for real rate variation in the pre-crisis sample, when excluding recessions, and when controlling for changes in macroeconomic uncertainty (e.g., total factor productivity volatility), the output gap, and inflation. The output gap and inflation are particularly useful for ruling out concerns that PVS_t simply captures monetary policy, as these two control variables enter directly into the Taylor (1993) monetary policy rule.

What drives the strong relationship between PVS_t and the real rate? 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 strongly forecasts future

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

returns on a portfolio that is long low-volatility stocks and short high-volatility stocks, but does not forecast cash flows (*ROE*) for the same low-minus-high volatility portfolio. Using a covariance decomposition based on Vuolteenaho (2002), we find that nearly 90% of the comovement between the real rate and PVS_t arises because the real rate forecasts future returns to volatility-sorted stocks. These results also alleviate concerns that PVS_t might load onto time-varying growth expectations and are again consistent with PVS_t capturing how investors price volatility.

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 relatively little. We show that a wide range of proxies, including the realized return volatility of the low-minus-high volatility stock portfolio, the Herskovic et al. (2016) common idiosyncratic volatility factor, the realized volatility of the aggregate stock market, and TFP volatility explain much less variation in the real rate than PVS_t. Moreover, controlling for these proxies does not affect the economic and statistical explanatory power of PVS_t for the real rate. Furthermore, real rates do not strongly forecast realized volatility of the low-minus-high volatility stock portfolio or the realized volatility of the aggregate stock market. While it is impossible to account for all possible sources of time-varying volatility, these results strongly indicate that time-varying aversion to diversifiable shocks plays an important role in driving precautionary savings motives.

Our new link between bond and stock markets raises significant questions for theories of asset pricing and macroeconomics. The fact that total volatility is important for the empirical findings is challenging for models with perfect risk sharing, where agents should care about beta and not volatility. We show empirically that an important class of equity portfolios are indeed imperfectly diversified across high- and low-volatility stocks by studying mutual funds. In particular, we show that there is substantial variation in the exposure of US equity mutual funds to high-volatility stocks. Moreover, high-volatility funds suffer relatively higher outflows than low-volatility funds during periods of low interest rates, as would be the case if adverse shocks lead high-volatility investors to decrease their equity exposure and increase precautionary savings demand for bonds.

We illustrate the tension between our empirical results and standard models with perfect risk sharing in a stylized asset pricing model. In the model, volatility, not beta, is the proper measure of risk because markets are segmented and investors are imperfectly diversified. We think of this feature as capturing segmented institutional investors with 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)) or under-diversified households (Benartzi (2001)). Investors are borrowing constrained, so the real risk-free rate is typically determined by agents who have strong precautionary savings motives because of their exposure to high-volatility stocks. As a result, the time-varying risk attitudes of these investors move around both risk premia on high-minus-low volatility stocks and the real rate. Importantly, and consistent with the data, market segmentation breaks the link between aggregate market valuations and the risk-free rate. Market segmentation is of course only one form of imperfect diversification and we also discuss several related mechanisms that are consistent our empirical evidence.

The relation between precautionary savings and the real interest rate has important implications for monetary policy. In a standard New Keynesian framework, monetary policy tightness depends on the gap between the actual real rate and the natural real rate of interest: the fundamental interest rate consistent with stable inflation and output at its natural rate (Clarida et al., 1999). Our results imply that estimates of the natural rate should also account for varying precautionary savings. Impulse response functions based on a standard recursive identification scheme (Bernanke and Mihov (1998); Gilchrist and Zakrajšek (2012)) show that precautionary savings shocks lead to very different inflation and output responses than independent monetary policy shocks, further supporting the notion that PVS_t comoves with the natural real rate as opposed to capturing monetary policy.

Our paper is related to several strands of the literature. 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)). We build on this research by showing that the pricing of volatility in the cross-section of stocks sheds light on the fundamental drivers of the real rate, despite the fact that aggregate stock market valuations do not reliably explain the real rate. While the literature on the pricing of idiosyncratic risk in the stock market has focused on the average risk premium (Ang et al. (2006a); Johnson (2004); Ang et al. (2009); Fu (2009); Stambaugh et al. (2015); Hou and Loh (2016)), we contribute by studying the time-variation in the risk premium of low-minus-high volatility stocks. A related paper is Herskovic et al. (2016), who argue that idiosyncratic firm-level shocks matter for households and generate 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 studying how this exposure is priced. On the other hand, our focus is on how the relative valuation of high- and low-volatility stocks connects 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 preferences.

This paper also contributes to a recent literature in macroeconomics that focuses on the trend in the natural rate of interest (Laubach and Williams (2003); Cúrdia et al. (2015); Del Negro et al. (2017)), attributing decade-by-decade changes in real rates primarily to expected growth and Treasury convenience yields. By contrast, our findings emphasize that time-varying precautionary savings play an important role for understanding quarterly real rate variation around long-term trends. 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, by relying on information from the cross section of stocks, we provide new evidence that

[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.

time-varying aversion to volatility is a fundamental determinant of the natural real rate of interest.[4] Our paper is also related to Berger et al. (2016), who argue that heterogeneity in the effects of recessions on individual income and consumption make a dual mandate in monetary policy optimal. Our paper is also related to the literature in macroeconomics arguing that precautionary savings matter for the origins of business cycles, the effectiveness of conventional and unconventional monetary policy, and firms' cash holdings.[5]

The remainder of this paper is organized as follows. Section 2 describes the data and portfolio construction. Section 3 presents the main empirical results. Section 4 explores monetary policy implications. Section 6 describes the model, shows that it can replicate the empirical findings, and discusses alternative interpretations. Finally, Section 7 concludes.

2 Data

We construct a quarterly data set running from 1970q2, when survey data on inflation expectations begins, to 2016q2. 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 the Online 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 mostly 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

[4] In contrast to Hartzmark (2016), we find little relation between macroeconomic volatility and the real rate. Moreover, we show that measures of macroeconomic volatility do not subsume the link between the real rate and PVS_t. We explore the relationship between our results and Hartzmark (2016) in Section 1.4 of the Online Appendix.

[5] See, e.g., Bloom (2009); Riddick and Whited (2009); Bloom et al. (2014); Cochrane (2016); McKay et al. (2016); Duchin et al. (2016); Hall (2016); Caballero and Simsek (2017).

(1993). If book equity is not available in COMPUSTAT Quarterly, we look for it in the annual file and then the book value data of Davis, Fama, and French (2000), in that order. 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.[6]

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:

$$PVS_t = \left(\overline{B/M}\right)_{low\ vol,t} - \left(\overline{B/M}\right)_{high\ vol,t}. \tag{1}$$

Again, PVS_t stands for the "price of volatile stocks." When market valuations are high, book-to-market ratios are low. Thus, PVS_t is high when the price of high-volatility stocks is large relative to low-volatility stocks. Quarterly realized returns in a given quintile are computed in an analogous fashion, aggregated up using monthly data from CRSP.

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

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 use a short maturity interest rate because inflation risk is small at this horizon, meaning inflation risk premia are unlikely to affect our measure of the risk-free rate. In addition, our focus is on understanding cyclical fluctuations in the real rate, as opposed to low-frequency movements that are likely driven by secular changes in growth expectations (Laubach and Williams (2003)). To keep things as simple and transparent as possible, we use a linear trend to extract the cyclical component of the real rate. In the Online Appendix, we show that all of our results are essentially unchanged if we just use the raw real rate or if we employ more sophisticated filtering methods that allow for stochastic trends.[7]

2.2 Summary Statistics

Table 1 contains basic summary statistics on our volatility-sorted portfolios. Panel A of the table reports statistics on book-to-market-ratios, while Panel B reports statistics on excess returns. The first thing to note in Panel A is that low-volatility stocks have lower book-to-market ratios than high-volatility stocks: on average, PVS_t is negative. However, as Fig. 1 shows, this masks considerable variation in PVS_t. Indeed, the standard deviation of PVS_t is about twice the magnitude of its mean. This variation is at the heart of our empirical work.

Panel B shows that returns on the low-minus-high volatility portfolio are themselves quite volatile, with an annualized standard deviation of 29.6%. While high-volatility stocks in our sample have high book-to-market ratios, the highest-volatility quintile of stocks on average has excess returns that are 2.71 percentage points per year 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

[7] In particular, we extract the cyclical component of the real rate using filtering methods that are robust to unspecified and potentially stochastic trends (Hamilton (2017)).

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 Panel B shows that high-volatility portfolios load significantly on 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:

$$\text{Real Rate}_t = a + b \times PVS_t + \varepsilon_t. \tag{2}$$

We report Newey and West (1987) standard errors using five lags. In the Online Appendix, we also consider several other methods for dealing with the persistence of these variables (e.g., parametric corrections to standard errors, generalized least squares, simulated bootstrap p-values, etc.). Our main conclusions are robust to these alternatives.

Column (1) of Table 2 reveals a strong positive correlation between the real rate and PVS_t — the real rate tends to be high when investors favor high-volatility stocks, and is low when investors prefer low-volatility stocks. This simple fact is the first piece of evidence that PVS_t captures variation in precautionary savings motives. 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 41%, indicating that PVS_t explains a large fraction of variation in the real rate. Column (2) of Table 2 separates PVS_t into its constituent parts. The valuations of low-volatility and high-volatility stocks enter with opposite signs, so both components of PVS_t play a role in driving the relation. The magnitude of the effect is economically large and measured precisely.

Figures 2 and 3 present visual evidence of our primary finding. Fig. 2 plots the time series of the real rate against the fitted value from regression in Eq. (2). As the figure shows, PVS_t tracks a remarkable amount of real rate variation since 1970. Additionally, the scatter plot in Panel A of Fig. 3 reinforces our linear regression specification and confirms that outliers are not driving our results. Panel B of Fig. 3 shows that the relationship is equally strong if we remove recession quarters, which are shaded in light grey. Thus, the relationship between PVS_t and the real rate is stable across different macroeconomic environments.

Column (3) of Table 2 indicates that our focus on the cross section of stock valuations is important. There appears to be no relationship between the book-to-market ratio of the aggregate stock market and the real rate. This fact is not just an issue of statistical precision. The economic magnitude of the point estimate on the aggregate book-to-market ratio is also quite small – a one-standard deviation movement in the aggregate book-to-market ratio is associated with only a 17 basis point movement in the real rate.[8] In contrast, column (3) of Table 2 shows that the statistical significance and the magnitude of the coefficient on PVS_t are unchanged when controlling for the aggregate book-to-market ratio.

In column (4), we control for variables thought to influence monetary policy: four-quarter inflation, as measured by the GDP price deflator, 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.

[8] As we discuss further in the Online Appendix, the aggregate book-to-market ratio does enter significantly in a small number of variants on our baseline specification. However, the statistical significance is irregular across various specifications, and the economic significance is always negligible.

The main takeaway is that the relationship between the real rate and PVS_t is stable throughout all of these regression specifications, implying that PVS_t does not just capture monetary policy. We revisit the relationship between monetary policy, the real interest rate, and PVS_t in Section 5.

In columns (5)-(8) of Table 2, we rerun the preceding regression analysis in first differences 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. Running regression (2) in differences yields very similar results to running it in levels. 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 when we run the regression in levels. The differenced regression also reinforces the nonexistent relationship between the real rate and the aggregate book-to-market ratio. Overall, the evidence in Table 2 indicates a strong and 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 below, these results stand up to the inclusion of a battery of additional control variables and various regression specifications.

3.2 Robustness and Alternative Cross-Sectional Sorts

Because the relation between PVS_t and the real rate is at the heart of our empirical results, we now show that this relation is robust to a wide range of additional tests. We first show robustness to alternative variable definitions, to controlling for a set of alternative cross-sectional valuation spreads, and to double-sorting by volatility and alternative characteristics. In short, we find that investors' willingness to hold volatile stocks indeed plays a special role for understanding real rate variation. The tests in this subsection take the following form:

$$\text{Real Rate}_t = a + b \times PVS_t + \theta^{'} X_t + \varepsilon_t, \tag{3}$$

where X_t is a vector of control variables that always includes the aggregate book-to-market ratio. In our horse races, it also contains book-to-market spreads based on alternative cross-sectional sorts.

We run these tests in both levels and changes, using both the full sample and the pre-crisis sample. The economic and statistical significance of PVS_t remains essentially unchanged throughout these robustness checks.

To start, we explore alternative definitions of PVS_t. Row (1) of Table 3 reproduces our baseline result from columns (3) and (7) of Table 2. In row (2) of Table 3, we recompute PVS_t by value-weighting the book-to-market ratio of stocks within each volatility quintile, as opposed to equal-weighting. The coefficients and statistical significance are comparable to the baseline, showing that our results are not exclusively driven by small stocks. In row (3), we construct PVS_t by sorting stocks on volatility measured over a two-year window, rather than a two-month window. As row (3) shows, this variant of PVS_t is still highly correlated with the real rate. Computing volatility over a long period helps ensure that our results are not driven by changing portfolio composition. That is, we are capturing changes in the valuations of stocks with a long history of being volatile, not changes in the volatility of value stocks. This distinction is critical to our interpretation of PVS_t as a measure of investors' willingness to hold volatile stocks.

In row (4), we run a horse race of PVS_t against a measure of liquidity premia in the fixed income market, the spread between 10-year off-the-run and on-the-run Treasury yields (Krishnamurthy (2002)).[9] The explanatory power of PVS_t for the real rate is unchanged, suggesting that PVS_t is not linked to the real rate because it captures variation in the demand for liquid assets like on-the-run Treasuries.

Next, we test whether volatility simply proxies for another characteristic that may drive the relation between the real rate and the cross-section of stocks. We do so by controlling for book-to-market spreads based on alternative characteristics in regression (3). For an alternative characteristic Y, we sort stocks in quintiles based on Y and then compute the difference between the book-to-market ratio of the lowest Y and highest Y quintiles. In other words, we construct book-to-market spreads for other characteristics in the same way we construct PVS_t. Rows (5)-(9) of Table

[9] The off-the-run spread is the difference between the continuously compounded 10-year off-the-run and on-the-run bond yields. On-the-run bond yields are from the monthly CRSP Treasury master file. The off-the-run bond yield is obtained by pricing the on-the-run bond's cash flows with the off the- run bond yield curve of Gürkaynak et al. (2007). For details of the off-the-run spread construction see Kang and Pflueger (2015).

3 shows the coefficient on PVS_t, while controlling for the Y-sorted book-to-market spread and the aggregate book-to-market. As before, we run these horse races for both the full and pre-crisis samples, as well as in levels and in changes.

Row (5) of Table 3 considers cash flow duration as a alternative characteristic. If low-volatility stocks simply have longer duration cash flows than high-volatility stocks, then a decline in real rates would increase their valuations relative to high-volatility stocks, potentially driving our results. To rule out this particular reverse causality story, we follow Weber (2016) and construct the expected duration of cash flows for each firm in our data. The duration-sorted valuation spread does not drive PVS_t out of the regression. This observation cuts against the idea that low-volatility stocks are "bond-like" because of their cash flow duration (e.g., Baker and Wurgler (2012)) and instead supports our point that volatility is the key characteristic determining whether stocks are bond-like.

Row (6) shows that PVS_t is robust to controlling for leverage-sorted valuation ratios. We define leverage as the book value of long-term debt divided by the market value of equity. Highly-levered firms may suffer disproportionately from a decrease in the real rate because they are effectively short bonds. But highly-levered firms also have high volatility, which could confound our results. Row (6) helps alleviate these concerns, as the leverage-sorted valuation ratio does not impact PVS_t in the regression.

In row (7), we show that the economic and statistical significance of PVS_t is unchanged when controlling for spreads based on systematic risk (i.e., beta). This test has important implications for interpreting our results, because perfectly diversified investors should care about beta and not volatility. We use the past two months of daily returns to compute beta, mimicking our construction of volatility.[10] The regression coefficient on PVS_t remains statistically significant at a 5% level in all cases, and the economic magnitudes are very similar. Thus, it does not appear that our measure of volatility is simply picking up on beta. The results in row (7) hence further corroborate the

[10] In the Section 1.1 of the Online Appendix, we try a number of additional constructions of beta. Specifically, we compute beta using (i) the past two years of monthly returns and (ii) the past ten years of semi-annual returns. In addition, we compute a measure of cash-flow beta as opposed to stock market beta, using rolling twelve-quarter regressions of quarter-on-quarter EBITDA growth on quarter-on-quarter national income growth. Our results are essentially unchanged using any of these additional measures.

notion that PVS_t does not simply capture risk aversion in the aggregate stock market, consistent with the weak relation between the real rate and the aggregate book-to-market ratio in Table 2.

In addition, we compare PVS_t to book-to-market spreads based on the popular Fama-French sorting variables, size and value. Consistent with our value-weighted results in row (2), the horse race in row (8) shows that the relationship between the real rate and PVS_t is robust to controlling for the difference in valuation between small and large stocks. Row (9) shows that PVS_t is robust to controlling for the book-to-market spread between value and growth stocks. The robustness to value-sorted book-to-market spreads is reassuring, because if value proxies for growth options, this suggests that the relation between PVS_t and the real rate is robust to controlling for the time-varying value of growth options.

In rows (10)-(16), we use double sorts as a complementary way to rule out alternative explanations for why PVS relates to the real rate. Specifically, we 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. For example, in row (10) we form a duration-neutral version of PVS_t by first sorting stocks into terciles based on their cash flow duration. Within each tercile we then compute the book-to-market spread between low and high volatility firms. The duration-neutral version of PVS is the average low-minus-high volatility valuation spread across the three duration terciles. In rows (10)-(14) of Table 3, we show that these double sorted book-to-market spreads are still strongly correlated with the real rate.

Row (15) computes a dividend-adjusted version of PVS_t. We first divide stocks based on whether they have paid a dividend over the previous twenty-four months.[11] We then compute PVS_t separately within the set of dividend-paying and non-dividend paying firms. The dividend-adjusted PVS_t is just the average across the two. Row (15) indicates that the explanatory power of PVS_t for the real rate is robust to controlling for dividends in this fashion.

Finally, our PVS_t measure might be simply capturing the value of industries that are particularly

[11] We use CRSP total return and ex-dividend adjusted returns to determine each firm's dividend yield in a given month.

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 (16) shows that this industry-adjusted spread still possesses significant explanatory power for the real rate.

The upshot of these robustness tests is that the sorting stocks on volatility is the key to our construction of PVS_t. Sorting on other characteristics does not perform nearly as well in terms of informational content about the real rate. This is a key reason we view PVS_t as measuring the economy's precautionary savings motive.

3.3 Unpacking the Mechanism

3.3.1 Returns on Volatility-Sorted Portfolios and the Real Rate

Why are PVS_t and the real rate related? Standard present value logic (Campbell and Shiller (1988); Vuolteenaho (2002)) suggests that variation in PVS_t itself is driven by changes in future expected returns of a portfolio that is long low-volatility stocks and short high-volatility stocks (i.e., the portfolio underlying PVS_t) or future expected cash flow growth of this portfolio. To explore what drives variation in PVS_t, we begin by forecasting the return on the volatility-sorted portfolio with either PVS_t or the real rate. Formally, we run:

$$R_{t \to t+k} = a + b \times X_t + \xi_{t+k}, \tag{4}$$

where X_t is either PVS_t or the real rate. Table 4 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 Newey and West (1987) with five lags. For regressions with a four-quarter horizon, we use Hodrick (1992) standard errors to be maximally conservative in dealing with overlapping returns.

Column (1) of Table 4 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 PVS_t is associated with a 5.3 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 variation in PVS_t largely reflects variation in expected returns, consistent with much of the empirical asset pricing literature (e.g., Cochrane (2011)).

Column (2) 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. In economic terms, the real rate forecasts returns on the long-short portfolio nearly as well as PVS_t. A one-standard deviation increase in the real rate is associated with a 3.1 percentage point increase in returns on the long-short portfolio. As we discuss in further detail below, 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.

In the remaining columns of Table 4 Panel A, we show that the relationship between real rate and returns on the volatility-sorted portfolio is not driven by other characteristics. Specifically, we explore the relationship between the real rate and the Fama and French (1993) factors. The columns show 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 3, the overall evidence suggests that volatility, not size, is the main driver of our results.

Panel B of Table 4 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. The forecasting R-squared of 0.31 is large. For comparison, the aggregate price-dividend ratio forecasts aggregate annual stock returns with an R-squared of 0.15 (Cochrane (2009)). Taken together, the forecasting evidence in Table 4 further corroborates that the comovement of the real rate with PVS_t is not due to time-varying growth expectations, but to time-varying precautionary savings.

3.3.2 Covariance Decomposition

Since variation in PVS_t is driven by changes in future expected returns or future expected cash flow growth, the real rate must covary with PVS_t because it covaries with one of these two factors. In Section A.3 of the Online Appendix, we use the present value decomposition in Vuolteenaho (2002) to make make this argument more explicit. In particular, we show that the covariance between PVS_t and the real rate can be approximately decomposed as follows:

$$\begin{aligned} Cov\left(\text{Real Rate}_t, PVS_t\right) \approx & \left(1-\rho\phi\right)^{-1} \times \left[Cov\left(\text{Real Rate}_t, Ret_{t+1}\right)\right. \\ & \left. -Cov\left(\text{Real Rate}_t, ROE_{t+1}\right) + Cov\left(\text{Real Rate}_t, \xi_{t+1}\right)\right]. \end{aligned} \tag{5}$$

Here, ρ is a log-linearization constant, $\phi$is the persistence of PVS_t, Ret_{t+1} is the return on the volatility-sorted portfolio, ROE_{t+1} is the return on equity of the same portfolio. We follow Vuolteenaho (2002) in setting $\rho = 0.969$. The parameter $\phi = 0.88$ is estimated using a simple AR(1) model. ξ_{t+1} is an error term that is comprised mainly of future innovations to PVS_t, but also collects the usual approximation errors that arise from these types of present-value decompositions.

Eq. (5) lets us quantitatively decompose the comovement between the real rate and PVS_t estimating each of the terms on the right hand side. The first covariance term on the right hand side can be inferred by forecasting future returns on the volatility-sorted portfolio with the real rate, as we did in Table 4. Similarly, the second term can be estimated by forecasting ROE_{t+1} on the volatility-sorted portfolio with the real rate. In the Internet Appendix, we directly show that neither PVS_t nor the real rate forecast ROE for low- versus high-volatility stocks.[12]

Combining these estimates, we find that nearly 90% of the comovement between the real rate and PVS_t arises because the real rate forecasts future returns to volatility-sorted stocks. The intuition is that most of the predictability in low-minus-high volatility portfolios can be accounted for with mean-reversion in PVS_t, meaning there is little room for changes in expected cash flows. Since most of the variation in PVS_t is driven by changing expected returns, most of its covariation with the real rate must be driven by covariation between the real rate and expected returns.

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. Finding little evidence for a variety of uncertainty measures and timing assumptions suggests that the relationship between the real rate and returns on the long-short portfolio sorted on volatility is likely driven by changing aversion to shocks.

We begin by showing that the relation between the real rate and PVS_t is not explained by contemporaneous volatility. Specifically, we run the regression in Eq. (2) with measures of contemporaneous realized volatility on the right-hand side. In particular, we include the realized return volatility on our low-minus-high volatility portfolio in quarter t, computed with daily data. To proxy for macroeconomic volatility, we include the volatility of TFP growth implied from a

[12] Furthermore, we can show that this is not simply a product of sampling error in the regression. Following Cochrane (2007)'s logic, the Vuolteenaho (2002) decomposition of returns implies that

$$\beta \quad = \quad 1 - \rho\phi + \beta_{ROE},$$

where β is the coefficient from a regression of future returns on log book-to-market and β_{ROE} is the coefficient from a regression of future log ROE on log book-to-market. Our point estimates are $\beta = 0.14$ and $\phi = 0.88$, implying a point estimate of $\beta_{ROE} = \beta - (1 - \rho\phi) = -0.01$. Thus, both direct evidence from cash flow forecasting regressions and indirect evidence from return forecasting regressions show that movements in PVS_t reflect changes in future returns, not future cash flows.

GARCH model, as in Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012).[13] In addition, we include the realized within-quarter volatility of the Fama and French (1993) factors, computed using daily data, and the common factor in idiosyncratic volatility variable of Herskovic et al. (2016).

The results are presented in Table 5. Column (1) finds no relationship between the real rate and the volatility on our low-minus-high volatility portfolio, so the baseline relation between the real rate and PVS_t does not appear to be driven by changes in the volatility 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. However, the five volatility measures in column (2) jointly achieve only an R-squared of 0.15, while the R-squared rises to 0.60 when we include PVS_t in column (3). In columns (4) to (6), we obtain similar results when running the analysis in first differences. Either way, the only variable robustly correlated with the real rate is PVS_t, whereas the volatility variables have little impact.

Next, we examine the possibility that PVS_t is related to expectations of future volatility instead of contemporaneous or lagged volatility. In Table 6, we try to forecast volatility directly using either PVS_t or the real rate. Formally, we run:

$$Vol_{t+1} = a + b \times X_t + \varepsilon_{t+1}, \tag{6}$$

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 the common factor in idiosyncratic volatility variable of Herskovic et al. (2016). In columns (2) and (3) of Panel A, PVS_t has marginally significant forecasting power for aggregate market volatility and the common factor in idiosyncratic volatility, as one might expect if an increase in volatility leads to stronger precautionary savings motives. But again, the R-squareds are small. In Panel B, we find that the real rate does not forecast any of the volatility

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

measures. While the real rate has marginally significant forecasting power for the volatility of the SMB portfolio and the common factor in idiosyncratic volatility, the signs of these coefficients are the opposite of what one would expect if a lower rate reflects concerns about higher volatility.[14]

Overall, the results presented in Tables 5 and 6 indicate that our baseline finding of a strong relation between the real rate and the pricing of volatility in the stock market cannot be explained by variation in volatility alone. While we cannot test for hard-to-observe volatility components, these results strongly suggest that time-varying aversion to shocks, and not time variation in the size of these shocks, is the key driver of precautionary savings during our sample. Consequently, these results suggest that models featuring time-varying pricing of shocks are needed to adequately capture real rate variation,

3.5 Interpretation

Thus far we have shown that the real interest rate is strongly correlated with PVS_t, with the correlation owing to the fact that the real rate covaries with expected returns on volatile stocks. When the real rate is low, investors demand more compensation for holding volatile stocks. Furthermore, the relationship appears to be driven by changes in investor aversion to volatility, rather than changes in the quantity of volatility.

Our interpretation of these results is that the demand of imperfectly diversified investors for precautionary savings is an important driver of real interest rates. If investors are imperfectly diversified, then they will care about the total volatility of the stocks they hold, as opposed to the CAPM betas of those stocks. When imperfectly diversified investors become more averse to portfolio volatility, they drive PVS_t down. Simultaneously, their demand for bonds increases through the precautionary savings channel, resulting in lower real rates. This interpretation underscores the importance of our emphasis on total volatility. If investors were perfectly diversified, their attitudes toward stocks with high CAPM betas would be the best measure of their attitudes to-

[14] In the Online Appendix, we also show that the quantity of risk also has no ability to forecast excess returns on the long-short portfolio of volatility sorted stocks, which is again consistent with our results being driven by time variation in investor aversion to volatility.

wards risk. However, we find that attitudes towards total volatility, as measured by PVS_t, have significantly more explanatory power for the real rate.

A potential alternative explanation for our results is that PVS_t comoves with the real rate because it captures aggregate growth expectations. For instance, more volatile stocks could have cash flows that are more sensitive to aggregate growth expectations. We think this alternative explanation is unlikely for two reasons. First, as discussed above, we find no evidence that PVS_t forecasts the cash flows of volatile stocks, while it strongly forecasts expected returns on volatility-sorted portfolios. Thus, for aggregate growth expectations to explain our results, it would have to be the case that aggregate growth affects the discount rate associated with volatile stocks, yet does not impact the cash flows of those stocks. It is difficult to imagine a reasonable model of asset markets and the macroeconomy that has these features. Second, if aggregate growth expectations were important, one would expect that aggregate stock market valuations as well as duration-sorted stock valuations would explain more variation in the real rate. However, in the data, neither can match the explanatory power of PVS_t for the real rate.

On balance, we therefore believe that the most natural interpretation of our results is that the precautionary savings demand of imperfectly diversified investors is a fundamental determinant of real interest rates.

4 Additional Supporting Evidence

4.1 Evidence from Mutual Fund Flows

So far, we have inferred investor preferences from asset prices, which have the advantage of aggregating over a broad range of investors, including households, institutions, firms, and international investors. In this section, we provide evidence that a specific but important class of investors, namely mutual funds, behaves consistently with the evidence from prices. The mutual fund industry is an ideal laboratory to test predictions of under-diversified investors, not only because the mutual fund industry itself is large,[15] but also because its behavior is likely representative of an even larger class of intermediated investors, such as private and public pensions funds, sovereign wealth funds, etc. Mutual fund flows are useful, because they provide direct evidence for portfolio under-diversification across high-volatility and low-volatility stocks and allow us to separately verify our baseline results in a completely different data set. If real rate variation indeed reflects time variation in precautionary savings motives, we expect investors to leave high-volatility mutual funds when the real rate is low. Specifically, an increase in risk aversion of investors exposed to high volatility should lead to outflows from high-volatility mutual funds, an increase in the demand for bonds, and a drop in the real rate.

Our sample is the CRSP mutual fund data base, from which we have monthly return data from 1973q2 through 2015q3. To test our prediction, we first need to determine whether some mutual funds are more exposed to high-volatility stocks than others. We use two simple measures. First, we estimate the return beta of each fund with respect to the high-volatility portfolio. Second, we simply calculate the volatility of the fund's returns. We use the full sample of monthly return data available for each fund to minimize measurement error. We then compute quarterly fund flows for each fund, winsorizing at the 5th and 95th percentiles, and restrict our data set to fund-quarter observations where the fund has total net assets of over \$100 million to ensure that our results are not driven by small funds.[16]

We find substantial heterogeneity in mutual funds' volatility exposure, reported in the fund-quarter summary statistics in Table 7 Panel A. The average fund appears in our sample for 31 quarters. The annualized volatility of returns, σ_f, over the full sample is about 12% for the average fund with a substantial cross-sectional standard deviation 8%. The beta of fund returns with respect to the high-volatility portfolio is 0.30 for the average fund, with a substantial cross-sectional standard deviation of 0.24. The substantial cross-fund heterogeneity indicates that investors are indeed imperfectly diversified across high-volatility and low-volatility stocks, as would need to be

[15] According to the 2017 ICI factbook, the U.S. mutual fund industry had USD 19.1 trillion assets under management as of the end of 2016.

[16] We obtain similar results if we use the full sample.

the case for total volatility to matter for precautionary savings motives.

In Panel B of Table 7, we explore the relationship between fund flows, the real rate, and fund volatility. Specifically we run

$$Flows_{f,t} = \alpha_f + \theta_1 Real\ Rate_t + \theta_2 Real\ Rate_t \times Vol_f + \varepsilon_{f,t},$$

where Vol_f is a measure of the fund's exposure to high-volatility stocks. In columns (1)-(3) Vol_f is the beta of the fund's returns with respect to the high-volatility portfolio. In columns (4)-(6), it is the volatility σ_f of the fund's returns.

Panel B of Table 7 shows that mutual fund flows indeed tell the same story as our baseline results. The magnitudes are economically meaningful. In column (1), a one percentage point drop in the real rate is associated with a 0.9 percentage point outflow for a fund with zero exposure to the high-volatility portfolio. A one-standard deviation increase in volatility exposure increases the impact of the real rate by over 50%: a one percentage point drop in the real rate is now associated with a 1.4 percentage point outflow. Column (2) shows the results are robust to including time fixed effects. Column (3) shows that they are robust to controlling for the fund's contemporaneous and lagged performance, so we are not simply picking up a performance-flow relationship. Similarly, Columns (4) through (6) show that mutual funds with higher overall volatility tend to experience outflows when the real rate is low.

Overall, the results in this section show that investor behavior, as measured by mutual fund flows, is consistent with our main results and support the interpretation that the real rate varies due to the time-varying precautionary savings motives of heterogeneous and imperfectly diversified investors.

4.2 Alternative Interest Rate Measures

We next explore the correlation between PVS_t and different interest rate measures in Table 8. We again run regressions like Eq. (2), but explore different dependent variables. Row (1) reproduces our baseline result from Table 2, where the real rate is defined as the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters and is linearly detrended.

In row (2), we use the raw real rate and do not detrend. The results are very similar to the baseline. In rows (3) and (4), we decompose the raw real rate into the one-year nominal Treasury bill rate and inflation expectations, so that the difference between the coefficients in row (3) and row (4) equals the coefficient in row (2). This decomposition shows that the correlation between PVS_t and the real rate primarily comes from the nominal rate, as one would expect if precautionary savings were driving demand for government bonds.

In rows (5)-(8), we try to separate movements in the real rate that can be attributed to monetary policy from monetary policy-independent variation. Taylor (1993) proposed that monetary policy can be well-described by a linear rule for the real short-term interest as a function of inflation and the output gap. Specifically, we decompose the real rate into a Taylor (1993) rule component and a residual. We do two versions of this decomposition. First, in rows (5) and (6), we use the original monetary policy coefficients on the output gap and inflation from Taylor (1993). Specifically, we compute

$$Taylor1993_t = 0.5 \times (OutputGap_t) + 0.5 \times (Inflation_t - 2) + 2$$

,where $Taylor1993_t$ is the real rate that obtains if the central bank follows the Taylor rule. We define the residual as the raw real rate minus $Taylor1993_t$. Rows (5) and (6) show that in this construction the explanatory power of PVS_t for the real rate comes from its explanatory power for Taylor (1993) rule residuals. In rows (7) and (8), we do a second version of the decomposition, where we estimate the coefficients on the output gap and inflation. Specifically, we run a regression of the raw real rate on the output gap and inflation and call the fitted value the Taylor rule component. Rows (7) and (8) show that in this construction, the explanatory power of PVS_t for the real rate again comes from its explanatory power for the residuals. These results indicate that PVS_t reflects unobservable demand for bonds due to precautionary savings, rather than the monetary policy reaction function to short-term business cycle fluctuations.

Finally, in rows (9)-(11), we examine the relationship between PVS_t and several interest rate spreads. Row (9) shows that PVS_t is negatively correlated with the term spread; when precautionary savings motives are strong, PVS_t is low and the term spread is high. These results indicate that long-term bond yields respond less with precautionary savings motives than short-term bonds, consistent with long-term bonds having higher return volatility than short-term bonds.

Rows (10) and (11) show the relationship between PVS_t and two measures of corporate credit spreads: the BAA-Treasury spread and the GZ spread of Gilchrist and Zakrajšek (2012). It is natural that PVS_t should be negatively correlated with both measures, because default risk priced into corporate bonds reflects total firm uncertainty, including both idiosyncratic and systematic risk. Thus, as precautionary savings motives increase and investors become more averse to firm-specific shocks, we would expect to see a widening of corporate bond spreads. In untabulated results, we find that PVS_t has more explanatory power for the real rate than credit spreads and drives out the credit spread when both are included in the regression. This likely reflects the fact that, by taking differences, PVS_t isolates attitudes towards taking on idiosyncratic risk from market-wide pricing factors, such aggregate risk aversion priced into the aggregate stock market and time-varying growth expectations. In contrast, the average credit spread also reflects these market-wide factors, as well as short-term spikes in corporate bond illiquidity (Bao et al. (2011)).

5 Implications for Monetary Policy

Our finding that precautionary savings is an important driver of the real interest rate has potentially significant implications for monetary policy. This section provides suggestive evidence that our new stock-market implied measure of precautionary savings is indeed linked to fundamental elements of the macroeconomy. In particular, our intuition suggests that precautionary savings motives drive variation in the natural real rate, so whether monetary policy is tight or loose should be evaluated relative to a natural real rate that accounts for precautionary savings (Clarida et al., 1999; Woodford, 2003). In this section, we show that recursively identified impulse responses support the interpretation that the relation between PVS_t and the real rate captures comovement with the natural real rate rather than monetary policy shocks.

We can illustrate the macroeconomic implications most easily by writing the Euler equation in the style of a New Keynesian model:

$$x_t \quad = \quad E_t x_{t+1} - \psi \left(r_t - r_t^n \right). \tag{7}$$

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. The important term in (7) is the last one, showing that macroeconomic activity today depends negatively on the gap between the real rate and the natural rate. The fact that r_t and r_t^n enter with opposite signs leads to two predictions. First, a typical monetary policy shock, or a positive shock to r_t, has the standard contractionary effect on output. By contrast, a positive shock to PVS_t, indicating that investors are less averse to volatility and weaker precautionary savings motives, should lead to an increase in r_t^n. Since r_t^n enters into (7) with a positive sign, output should increase in response to a positive PVS_t shock. Of course, the exact output response in this second case depends on how the central bank responds to precautionary savings shocks. If inflation is linked to the output gap through a New Keynesian Phillips curve, we expect these output responses to be accompanied by similarly signed, but potentially more sluggish, inflation responses.

5.1 VAR Framework

We estimate a VAR that is as simple and transparent as possible, while following a common set of recursiveness assumptions, similarly to Sims (1980), Bernanke and Mihov (1998) and Gilchrist and Zakrajšek (2012). 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.[17]

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

$$Y_t = \sum_{i=1}^{k} B_i Y_{t-i} + \sum_{i=1}^{k} C_i P_{t-i} + A^y v_{y,t} \tag{8}$$

$$P_t = \sum_{i=0}^{k} D_i Y_{t-i} + \sum_{i=0}^{k} G_i P_{t-i} + \begin{bmatrix} v_{PVS,t} \\ v_{MP,t} \end{bmatrix}. \tag{9}$$

Here, Y_t is a vector of quarterly non-policy variables, consisting of the output gap and detrended inflation. P_t is a vector of policy variables consisting of PVS_t and the detrended real rate. Eq. (8) describes a set of structural relationships in the economy, where macroeconomic variables depend on lagged values of macroeconomic and policy variables. Eq. (9) describes the stance of monetary policy conditional on contemporaneous macroeconomic variables. By analogy to the treatment of money demand shocks in Bernanke and Mihov (1998), the vector of policy shocks contains a precautionary savings demand shock for the risk-free asset $v_{PVS,t}$ and a monetary policy shock.

Similarly to Bernanke and Mihov (1998), we estimate the structural policy shocks under the restriction that $v_{PVS,t}$ does not respond to $v_{MP,t}$ contemporaneously, but $v_{MP,t}$ may respond to $v_{PVS,t}$. 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. 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 use two-step GMM using a Bartlett kernel with two lags and the initial weighting matrix equal to the identity. See the Appendix for estimation

[17] However, our results are not sensitive to this second assumption, and we show that impulse responses are substantially unchanged if instead we assume that PVS_t is faster than the real rate.

details.

Both a Wald test and a Hansen J-test provide clear evidence that the real rate reacts contemporaneously to PVS_t, consistent with our baseline results that PVS_t explains substantial variation in the real rate. For the reaction coefficient of PVS_t onto the monetary policy innovation $v_{MP,t}$, we estimate a point estimate of 2.32 with a t-statistic of 3.27. The over-identifying restriction that the real rate does not react contemporaneously to v_{PVS} is rejected by a Hansen J-test at any conventional significance level with a p-value of 0.0009.

5.2 Impulse Response Functions

Fig. 4 displays impulse responses of output, inflation, PVS_t, and the real rate to one-standard-deviation increases in v_{MP} and v_{PVS}. Dashed lines indicate 95% confidence bands. The left panels display responses to a monetary policy shock, corresponding to an unanticipated tightening by the central bank and an increase in the real rate. Reassuringly, our simple VAR identification scheme produces results that are consistent with the 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 precautionary savings shocks drive the real rate, and not vice versa.

By contrast, the right panel of Fig. 4 shows that a positive PVS_t shock leads to a significant increase in output and a flat inflation response, despite being associated with a similar increase in the real rate as the MP shock. The intuition is that a positive innovation to PVS_t reflects a decrease in investors' aversion to volatility, decreasing the motive to save and effectively acting as a positive demand shock for consumption and investment. The difference in responses across the left and right panels in Fig. 4 is exactly what the Euler equation would suggest if PVS_t directly impacts the natural real rate, with a less than full accommodation by monetary policy.[18]

At first glance, it may seem that Fig. 4 suggests that the relationship between PVS_t and the real rate we find reflects variation in expected growth, not precautionary savings as we argue. After all, Fig. 4 shows that following a shock to PVS_t, aggregate growth is high and the output gap responds positively. We maintain that PVS_t does not reflect growth expectations for several reasons. First, in Section 3, we abstracted from monetary frictions in the economy and argued that if volatility-sorted portfolios were differentially exposed to expected growth, PVS_t should forecast relative portfolio cash flows, but not returns. Finding empirically that PVS_t forecasts returns but not cash flows, we concluded that PVS_t does not appear to proxy for growth. Second, if PVS_t was simply pricing investors' growth expectations, PVS_t and the output gap should have a high overall correlation, which they do not. Third, if PVS_t proxied for growth, it should decline following a monetary policy shock, the way the output gap does. In contrast, the evidence in the left column of Fig. 4 shows that PVS_t does not respond to monetary policy shocks.

The most natural explanation for the positive output gap response to a PVS_t shock is therefore that in the presence of monetary frictions a PVS_t shock acts as a traditional demand or discount rate shock in a New Keynesian model. Intuitively, if a positive demand shock leads to an increase in the natural rate r_t^n and the monetary policy authority does not fully and immediately raise policy rates to offset this shock, $r_t - r_t^n$ will decline, resulting in a temporary boom through the Euler equation (7). Taken together, the impulse responses in this section provide highly suggestive evidence of a link between PVS_t as a measurable proxy for demand shocks and monetary frictions in the economy.

[18] In the appendix, we show that conclusions are unchanged if instead we make the opposite identification assumption that PVS_t responds to the real rate contemporaneously, but the real rate reacts to precautionary savings demand for a lag. This second identification assumption is different from saying that the Fed does not pay attention to the stock market. It merely requires that the Fed historically did not react instantaneously to the cross-sectional valuation spread newly documented in this paper. Impulse responses are also robust to excluding the post-crisis period and to including additional lags.

6 Stylized Segmented Markets Model

This section presents a stylized model to rationalize the empirical findings. Markets are segmented, so total volatility, not just systematic volatility, is priced and generates precautionary savings motives. Time-varying risk aversion drives both prices and expected returns on stocks and precautionary savings demand for bonds. We take segmentation as our starting point, building on the long tradition of segmented market models going back at least to Merton (1987). We think of this assumption as potentially representing short-lived and risk-averse institutional investors with limited arbitrage ability (Lee et al., 1991) or households with home or familiarity bias (see Barberis and Thaler (2003) for an overview) and possibly heterogeneous background risk as in Schmidt (2016) and Berger et al. (2016). Investors are borrowing-constrained, so bonds are priced by those investors with the highest bond valuations, who are typically investors with strong precautionary savings demand due to volatility exposure.

6.1 Endowments and Preferences

The economy consists of a share p_H of high-volatility stocks and a share $p_L = 1 - p_H$ of low-volatility stocks. 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:

$$
\begin{aligned}
c_{H,t} &= \mu t + \varepsilon_{H,t}, && (10)\\
c_{L,t} &= \mu t + \varepsilon_{L,t}, && (11)\\
\begin{bmatrix} \varepsilon_{L,t} \\ \varepsilon_{H,t} \end{bmatrix} &\overset{iid}{\sim} N\left(0, \begin{pmatrix} \sigma_L^2 & 0 \\ 0 & \sigma_H^2 \end{pmatrix}\right). && (12)
\end{aligned}
$$

Here, we use lower case letters to denote logs. 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 $c_{i,t}$.

We introduce time-varying risk aversion to generate time-varying risk premia in stocks. For tractability, we follow the Menzly et al. (2004) version of Campbell and Cochrane (1999) habit formation preferences. While time-varying pricing of risk is clearly important for replicating the empirical results, we consider the specific functional form illustrative.

Agents of type i maximize the expected discounted sum of log consumption utility relative to external habit $X_{i,t}$ as in Campbell and Cochrane (1999):

$$U_{i,t} = E_t\left[\sum_{\tau=t}^{\infty}\beta^{\tau}\log\left(C_{i,\tau}-X_{i,\tau}\right)\right],\ i=H,L. \tag{13}$$

Relative risk aversion $G_{i,t}=S_{i,t}^{-1}$ varies with surplus consumption:

$$S_{i,t} = \frac{C_{i,t}-X_{i,t}}{C_{i,t}}. \tag{14}$$

Habit dynamics are defined implicitly via (14) and by assuming that inverse surplus consumption of agent i follows a process of the form:

$$G_{i,t+1} = \kappa\bar{G}+(1-\kappa)G_{i,t}-\alpha\left(G_{i,t}-\lambda\right)\varepsilon_{i,t+1}. \tag{15}$$

6.2 Asset Markets

6.2.1 Equity

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:

$$\frac{P_{i,t}}{C_{i,t}} = a+bS_{i,t}, \tag{16}$$

for positive constants a and b that are given in Online Appendix A. 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 moving average of past stock prices. The book-to-market ratio then is computed as

$$BM_{i,t} = \frac{B_{i,t}}{P_{i,t}},\ B_{i,t} = (1-\rho)\sum_{\tau=1}^{\infty}\rho^{\tau}P_{i,t-\tau}. \tag{17}$$

Eqs. (16) and (17) show that when agent i suffers an adverse shock, lowering surplus consumption, this drives down the stock price, raises the book-to-market ratio, and raises future expected returns for stock i.

6.2.2 Bonds

Real risk-free one-period bonds 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:

$$r_{f,t} = min\left\{r_{f,H,t}, r_{f,L,t}\right\}, \tag{18}$$

$$r_{f,i,t} = \mu - \log\beta - \frac{\sigma_i^2}{2} - \varepsilon_{i,t} \tag{19}$$

$$-\log\left((1-\kappa) + \alpha\sigma_i^2 + \left(\kappa\bar{G} - \sigma_i^2\alpha\lambda\right)G_{i,t}^{-1}\right),\ i \in \{H, L\}. \tag{20}$$

Here, $r_{f,i,t}$ denotes the risk-free rate at which investor type i is indifferent about investing in the bond market. Eq. (20) shows that for high-volatility investors with $\sigma_H^2 > \frac{\kappa\bar{G}}{\lambda\alpha}$ the indifference real rate $r_{f,H,t}$ decreases with risk aversion $G_{H,t}$. Intuitively, an increase in risk aversion generates an incentive to save for a riskier future, making high-volatility investors willing to invest at a lower risk-free rate, capturing the time-varying precautionary savings motive. By contrast, for a low-volatility investor with σ_L^2, the indifference risk-free rate $r_{L,f,t}$ is driven by competing intertemporal substitution and precautionary savings motives, so the relation between the risk-free rate and low-vol risk premia is ambiguous. If $\sigma_L^2 < \frac{\kappa\bar{G}}{\lambda\alpha}$, an adverse shock generates an intertemporal substitution motive, generating an incentive to borrow and driving up $r_{f,L,t}$. From now on, 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.[19]

Eqs. (18) and (20) make clear that the real rate moves positively with the difference between high-volatility and low-volatility investors' risk aversion. Combined with the link between investor risk aversion and equity prices (16), this generates the main intuition in this paper. A decline in the price of high-volatility stocks relative to low-volatility stocks indicates that high-volatility investors are risk averse, increasing high-volatility investors' precautionary savings motives and driving down the risk-free rate.

6.3 Calibration and Results

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 9. 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-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 the condition that $p_H\sigma_H^2 + p_L\sigma_L^2 = \frac{\kappa\bar{G}}{\lambda\alpha}$.

[19] Formally, we assume that $\sigma_H^2 > \frac{\kappa\bar{G}}{\lambda\alpha}$ and $\sigma_L^2 < \frac{\kappa\bar{G}}{\lambda\alpha}$. 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, or $p_H\sigma_H^2 + p_L\sigma_L^2 = \frac{\kappa\bar{G}}{\lambda\alpha}$. This ensures that the aggregate book-to-market, which reflects a weighted average of high-volatility and low-volatility stock prices, is approximately uncorrelated with the real rate.

Table 10 shows that 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. The risk-free rate in the model is low and comparably volatile to the data.[20]

The third panel in Table 10 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, 7% of our simulations generate low-minus-high volatility excess returns that are as large or larger than in the 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.

Importantly, the bottom panel of Table 10 shows that the model replicates the main empirical facts documented in this paper. A regression of the risk-free rate on the aggregate book-to-market ratio yields a small and insignificant slope coefficient, while a regression onto PVS_t generates a strongly positive coefficient, as in the data. Moreover, the risk-free rate forecasts excess returns on the low-minus-high volatility equity portfolio with empirically reasonable magnitudes.

Intuitively, an increase in high-volatility investors' risk aversion raises required expected returns on high-volatility stocks relative to low-volatility stocks and drives down the price of high-volatility stocks. At the same time, an increase in high-volatility investors' risk aversion increases their precautionary savings motives and drives down the risk-free rate. This generates a positive

[20] All model moments are 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. 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.

relation between low-minus-high volatility book-to-market and the risk-free rate, as in the data. By contrast, the aggregate book-to-market ratio is largely driven by low-volatility investors, whose risk aversion has little correlation with the real rate. 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 that the assumption of segmented markets is 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. In column (3), we assume a representative consumer, who consumes aggregate consumption, and with preferences of the form (13) through (15). Without market segmentation, 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 volatility excess return. In addition, in the representative agent model time-variation in the representative agent's surplus consumption ratio drives both aggregate equity risk premia and the risk-free rate, generating a counterfactually positive relation between the aggregate book-to-market and the risk-free rate.

Column (4) shows that time-varying prices of risk are 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. Column (4) considers the case of segmented investors with log utility and no habit formation. This case is nested by our baseline model with $\alpha = 0$, $\kappa = 1$, 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.

6.4 Discussion

The new empirical results in Sections 3 through 5 provide clear evidence of time-varying demand for precautionary savings as a significant determinant of time-varying real interest rates. Our empirical findings raise significant questions for representative-agent asset pricing models and indicate that future research into the drivers of time-varying demand for precautionary savings is likely to be fruitful. The main ingredients of any model that is consistent with our empirical findings are as follows. First, total volatility must give rise to a precautionary savings motive, so total volatility, not systematic risk, must enter investors' pricing kernel. Second, the pricing of risk must vary over time. 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.

A variety of models could have these ingredients and generate our empirical results. Investor heterogeneity need not arise exclusively from households, as it does in the model in Section 6. Constrained institutional investors or financial intermediaries like mutual funds and broker-dealers are natural candidates for generating heterogeneity, given the evidence that intermediaries take concentrated positions in individual stocks (Kacperczyk et al. (2005); Veldkamp (2006); Cremers and Petajisto (2009); Agarwal et al. (2013); Merton (1987)) and recent research indicating that they appear to price a wide range of assets (He and Krishnamurthy (2013); Adrian et al. (2014); He et al. (2016)). We regard our empirical results on mutual fund flows as providing suggestive evidence for this class of channels. Another potential source of heterogeneity is international investors, who the literature has argued may be partly responsible for the current low level of interest rates (Caballero et al. (2008); Caballero and Krishnamurthy (2009)). Our stock market-based measure automatically aggregates over investors that participate in the US stock market and as such should incorporate shifts in both domestic and international investors' attitudes towards volatility. We regard this aggregation quality as a major advantage of our stock-market based approach compared to household data.

Similarly, we consider the habit formation mechanism in the model in Section 6 as illustrative of a broader range of sources for time-varying prices of risk. Our results are consistent with alternatives, including shifts in the wealth distribution across investors with different degrees of risk aversion (Chan and Kogan (2002); Hall (2016); Barro and Mollerus (2014)) and long-run risks with time-varying quantities of risk that are hard to observe (Bansal et al. (2012)).

7 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 potentially diversifiable shocks, 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 drivers of precautionary savings in the presence of partially segmented markets is likely to be fruitful.

References

Adrian, T., E. Etula, and T. Muir (2014). Financial intermediaries and the cross-section of asset returns. *Journal of Finance 69*, 2557–2596.

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.

Bao, J., J. Pan, and J. Wang (2011). The illiquidity of corporate bonds. *The Journal of Finance 66*(3), 911–946.

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.

Berger, D., I. Dew-Becker, K. Milbradt, and L. D. W. Schmidt (2016, April). Layoff risk, the welfare cost of business cycles, and monetary policy. *Working Paper*.

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.

Caballero, R. J. and A. Simsek (2017, October). A risk-centric model of demand recessions and macroprudential policy. *Working Paper, MIT*.

Campbell, J. Y. and J. Ammer (1993). What moves the stock and bond markets? a variance decomposition for long-term asset returns. *The Journal of Finance 48*(1), 3–37.

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. (2007). The dog that did not bark: A defense of return predictability. *The Review of Financial Studies 21*(4), 1533–1575.

Cochrane, J. H. (2009). *Asset Pricing:(Revised Edition)*. Princeton university press.

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*.

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.

Davis, J. L., E. F. Fama, and K. R. French (2000). Characteristics, covariances, and average returns: 1929 to 1997. *The Journal of Finance 55*(1), 389–406.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity*, 235–316.

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.

Gilchrist, S. and E. Zakrajšek (2012, June). Credit spreads and business cycle fluctuations. *American Economic Review 102*(4), 1692–1720.

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

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007). The us treasury yield curve: 1961 to the present. *Journal of Monetary Economics 54*(8), 2291–2304.

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

Hamilton, J. D. (2017). Why you should never use the hodrick-prescott filter. *Review of Economics and Statistics forthcoming*.

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

He, Z., B. Kelly, and A. Manela (2016). Intermediary asset pricing: New evidence from many asset classes. *Journal of Finance XX*, XX.

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *American Economic Review 103*, 732–770.

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.

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.

Kang, J. and C. E. Pflueger (2015). Inflation risk in corporate bonds. *The Journal of Finance 70*(1), 115–162.

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*.

Krishnamurthy, A. (2002). The bond/old-bond spread. *Journal of Financial Economics 66*(2), 463–506.

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.

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.

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.

Schmidt, L. D. W. (2016, March). Climbing and falling off the ladder: Asset pricing implications of labor market event risk. *SSRN Working Paper*.

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

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica 48*, 1–48.

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.

Veldkamp, L. (2006). Information markets and the comovement of asset prices. *Review of Economic Studies 73*, 823–845.

Vuolteenaho, T. (2002). What drives firm-level stock returns? *The Journal of Finance 57*(1), 233–264.

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

Woodford, M. (2003). Optimal Interest-Rate Smoothing. *The Review of Economic Studies 70*(4), 861–886.

FIGURES

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

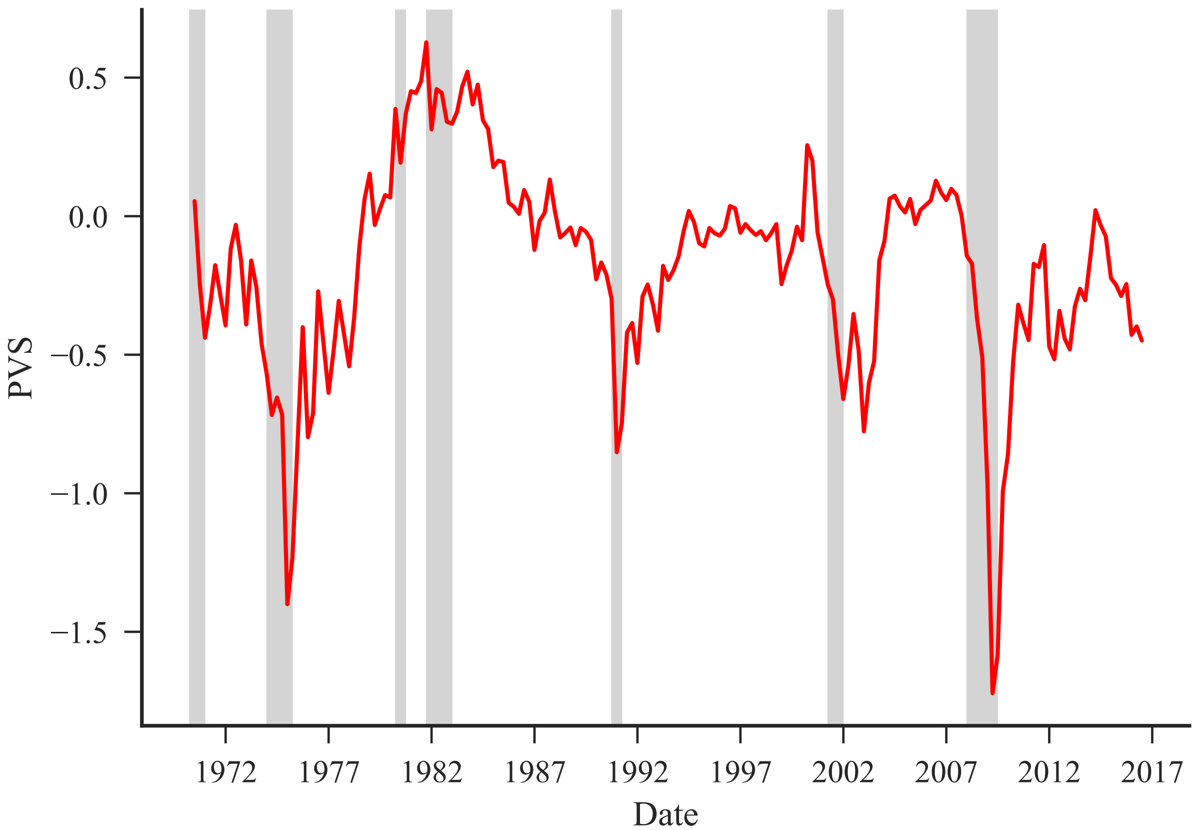

Notes: This figure plots the spread in book-to-market ratios between low and high volatility stocks. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at the end of each quarter using the previous sixty days of daily returns. We then form equal-weighted portfolios based on the quintiles of volatility. Within each quintile, we compute the average book-to-market (BM) ratio, which we call PVS_t. The Online 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 1970Q2-2016Q2.

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

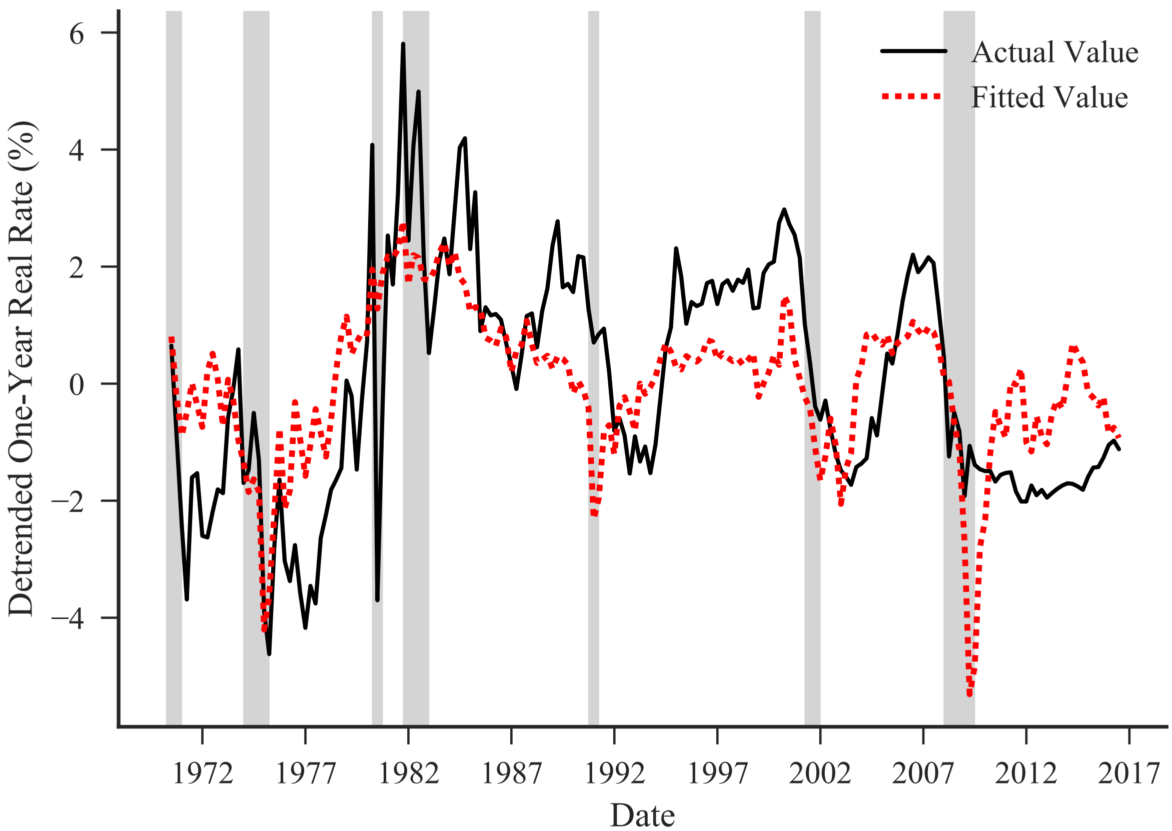

Notes: This figure plots the linearly detrended one-year real rate, as described in Table 3, and the fitted value from a regression of the real rate on the spread in book-to-market ratios between low and high volatility stocks (PVS_t). For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility 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 Online 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 1970Q2-2016Q2.

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

Panel A: Full Sample

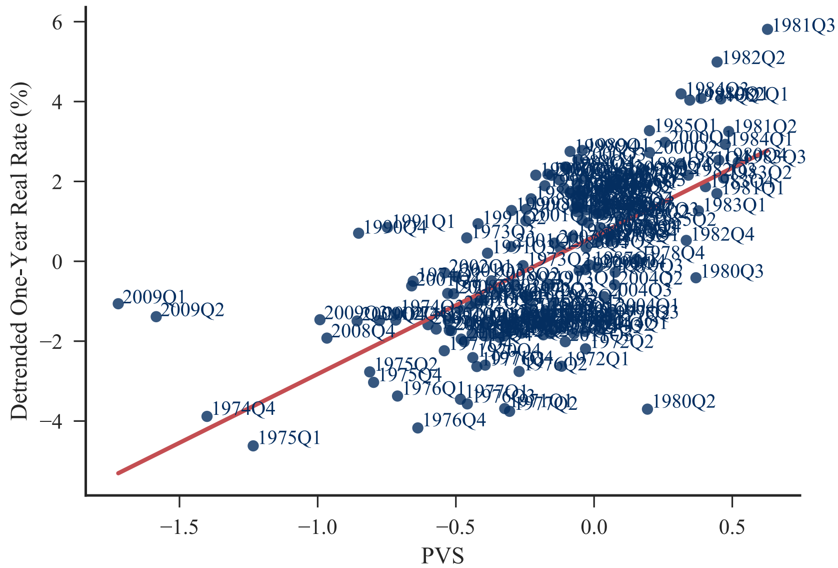

Panel B: Excluding NBER Recessions

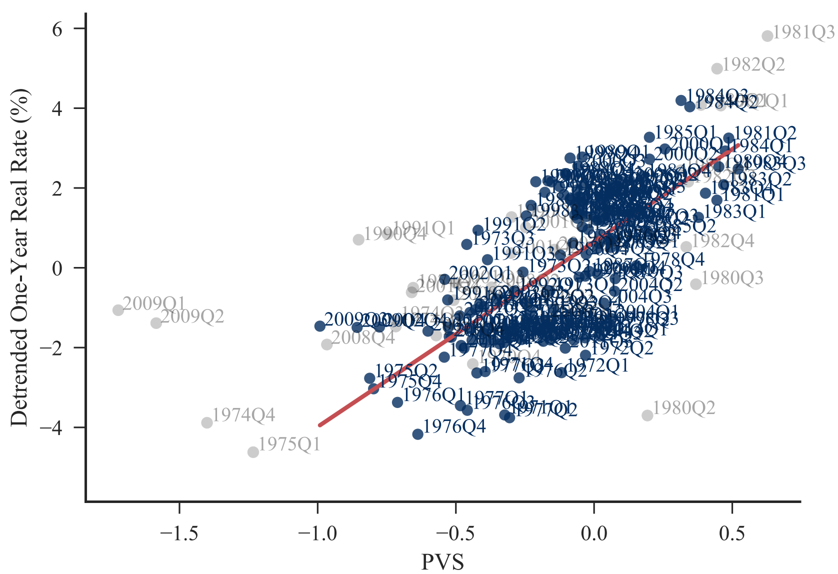

Notes: This figure plots the detrended one-year real interest rate against the spread in book-to-market ratios between low and high volatility stocks (PVS_t). Panel A shows the full sample, while Panel B excludes NBER recessions. For all NYSE, AMEX, and NASDAQ firms in CRSP, we compute volatility at 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 1970Q2-2016Q2.

Figure 4: Impulse Responses to Monetary Policy and PVS Shocks

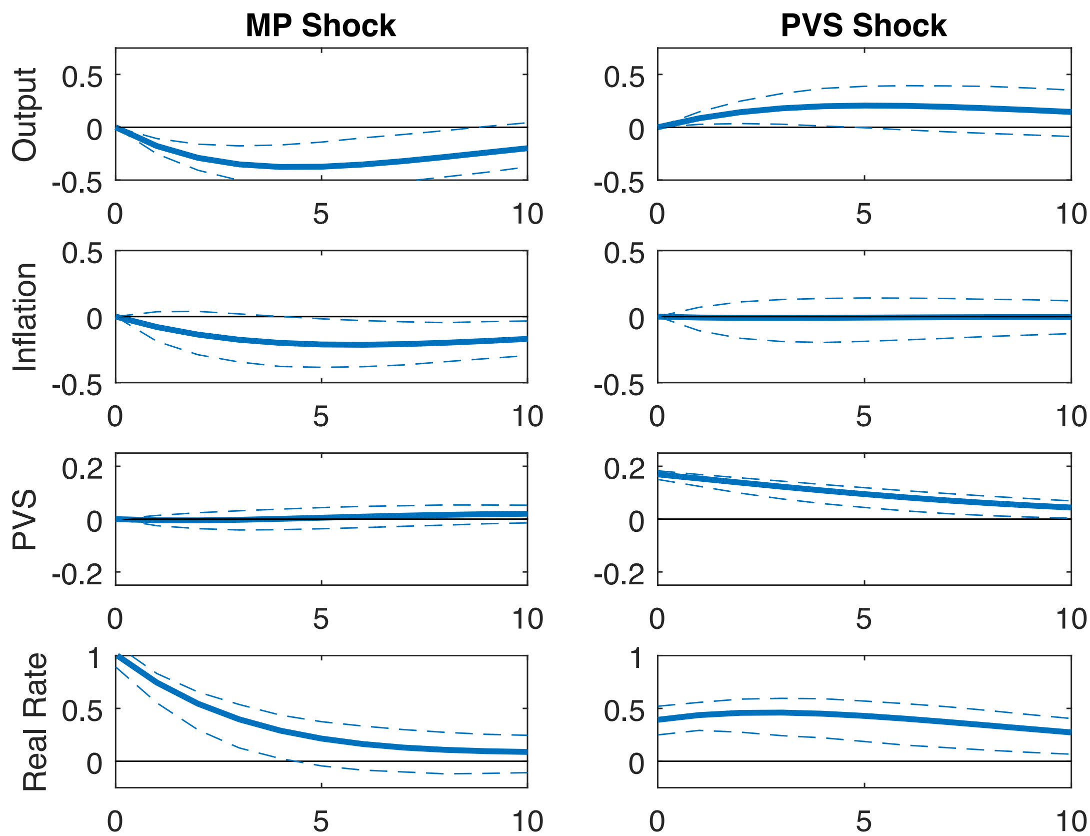

Notes: This figure plots impulse responses to monetary policy and PVS shocks. Impulse responses to one-standard deviation shocks are estimated from a four-variable VAR in the output gap, inflation, PVS, and de-trended real rate with one lag using quarterly data 1973Q1-2015Q4. Following Bernanke and Mihov (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 and the Real Rate

***Panel A**: Book-to-Market Ratios of Volatility Sorted Portfolios*

	High Volatility → Low Volatility					PVS
	5	4	3	2	1	1-5
Mean	1.04	0.87	0.83	0.82	0.86	-0.18
Volatility	0.45	0.31	0.26	0.25	0.28	0.37
Min	0.45	0.48	0.48	0.51	0.54	-1.72
Median	0.92	0.78	0.78	0.74	0.75	-0.12
Max	3.10	2.13	1.80	1.71	1.70	0.63

***Panel B**: Realized Excess Returns of Volatility Sorted Portfolios*

	5	4	3	2	1	1-5
Mean	7.44	9.65	12.04	11.15	10.15	2.71
Volatility	39.17	31.19	25.07	19.99	15.42	29.57
Median	-0.11	6.83	12.07	13.13	12.60	9.47
Min	-44.87	-37.31	-31.72	-29.25	-22.28	-49.51
Max	74.19	55.22	45.14	35.82	27.32	50.48
α	-4.99	-0.96	2.57	2.41	2.94	7.92
$t(\alpha)$	-2.08	-0.84	3.87	3.01	2.52	2.58
CAPM-β	1.25	1.17	1.03	0.92	0.73	-0.51
SMB-β	1.86	1.39	1.04	0.66	0.40	-1.46
HML-β	0.19	0.09	0.18	0.33	0.37	0.18

***Panel C**: Real Rate*

	Mean	Volatility	Median	Min	Max
Raw Real Rate	1.86	2.30	2.18	-1.86	8.72
Detrended Real Rate	0.00	1.96	-0.21	-4.62	5.81

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 1970Q1 to 2016Q2. 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. 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. We detrend the real rate using a linear trend and explore alternative methodologies in the Online Appendix..

Table 2: What Explains Real Rate Variation?

Dep. Variable:	One-Year Real Rate							
	Levels				First-Differences			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PVS	3.44**		3.45**	3.34**	2.22**		2.45**	2.09**
	(5.36)		(5.01)	(4.70)	(2.73)		(2.65)	(2.36)
BM Low-Vol		3.02**				1.99*		
		(3.11)				(1.80)		
BM High-Vol		-3.46**				-2.22**		
		(-5.39)				(-2.70)		
Aggregate BM			-1.10	0.26			2.30	3.59
			(-0.71)	(0.12)			(0.88)	(1.16)
Output Gap				0.09				0.36**
				(0.79)				(2.51)
Inflation				-0.12				0.22
				(-0.95)				(1.16)
Constant	0.62**	1.01	0.62**	0.60**	-0.00	-0.00	-0.01	-0.01
	(2.64)	(1.49)	(2.65)	(2.52)	(-0.07)	(-0.08)	(-0.13)	(-0.31)
Adj. R^2	0.41	0.41	0.42	0.42	0.13	0.12	0.13	0.18
N	185	185	185	185	184	184	184	184

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 (BM Low Vol) and top quintile (BM High Vol) 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 annualized four quarter percentage growth in the 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. We also independently detrend the output gap, inflation, and the aggregate book-to-market ratio. Results using the raw series for all variables is contained in the Online Appendix. Standard errors are computed using Newey-West (1987) with five lags. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. Data is quarterly and spans 1970Q2-2016Q2.

Table 3: Robustness: The Real Rate and PVS_t

		Levels						First-Differences					
		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
(1)	Baseline	**3.45**	5.01	0.42	**4.11**	7.61	0.47	**2.45**	2.65	0.13	**3.58**	3.67	0.20
(2)	VW	**3.16**	4.48	0.32	**4.01**	6.01	0.41	**1.39**	2.45	0.08	**1.81**	2.59	0.10
(3)	2YR Vol	**4.49**	6.27	0.52	**5.13**	8.20	0.54	**2.22**	2.32	0.05	**3.64**	4.21	0.10
Horse-Races:													
(4)	Liquidity	**3.80**	6.54	0.47	**4.30**	7.73	0.51	**2.06**	2.14	0.15	**3.17**	3.02	0.21
(5)	Duration	**3.23**	4.26	0.42	**3.62**	5.24	0.49	**2.46**	3.13	0.12	**3.46**	4.31	0.20
(6)	Leverage	**4.12**	6.15	0.44	**4.53**	7.57	0.48	**3.20**	2.87	0.14	**4.19**	3.34	0.21
(7)	2M-Beta	**3.47**	5.50	0.41	**4.02**	7.73	0.48	**1.77**	2.55	0.15	**2.82**	4.35	0.22
(8)	Size	**3.04**	2.48	0.42	**4.01**	3.80	0.47	**3.44**	2.42	0.13	**4.18**	2.59	0.20
(9)	Value	**4.16**	4.97	0.43	**4.70**	7.03	0.48	**3.88**	3.05	0.16	**4.49**	3.24	0.22
Double-Sorts:													
(10)	Duration	**3.85**	3.96	0.17	**3.98**	3.73	0.14	**3.19**	2.70	0.10	**4.10**	3.23	0.15
(11)	Leverage	**4.76**	5.08	0.35	**5.40**	5.79	0.38	**3.53**	2.81	0.12	**4.96**	3.52	0.18
(12)	2M-Beta	**4.52**	5.54	0.43	**5.25**	8.29	0.50	**1.90**	2.33	0.04	**2.88**	4.10	0.07
(13)	Size	**5.21**	4.79	0.39	**6.30**	7.90	0.46	**3.76**	2.60	0.12	**4.85**	3.12	0.17
(14)	Value	**8.61**	4.96	0.31	**9.49**	4.88	0.32	**6.16**	2.25	0.09	**8.97**	2.73	0.15
(15)	Dividend-Adj	**3.35**	4.20	0.28	**3.97**	4.91	0.28	**2.00**	2.59	0.08	**2.52**	2.93	0.11
(16)	Industry-Adj	**3.56**	5.18	0.34	**4.03**	5.65	0.35	**1.66**	2.39	0.06	**2.74**	3.96	0.11

Notes: This table reports a battery of robustness exercises for our main results. Specifically, we report time-series regression results of the following form: $\text{Real Rate}_t = a + b \times PVS_t + \theta X_t + \varepsilon_t$. We run this regression directly in levels and in first differences. 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 (linearly detrended), 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. Row (4) controls for the spread between off-the-run and on-the-run Treasury yields (CITATION). Rows (5)-(9) run bivariate horse races by adding book-to-market spreads based on other characteristic sorts to our control variables X_t. See the Table 2 of the Online Appendix for a description of each characteristic. In rows (10)-(16), 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 Newey-West (1987) 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 1970Q2-2016Q2 (pre-crisis ends in 2008Q4).

Table 4: Forecasting Returns with PVS and the Real Rate

***Panel A** - Quarterly Forecasting*

	Vol-Sorted Ret$_{t,t+1}$		Mkt-Rf$_{t,t+1}$		SMB$_{t,t+1}$		HML$_{t,t+1}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PVS_t	14.42**		-2.22		-2.84**		-0.05	
	(5.07)		(-1.25)		(-2.83)		(-0.05)	
Real Rate$_t$		1.57**		-0.26		-0.52**		0.31
		(2.81)		(-0.91)		(-2.65)		(1.19)
Constant	3.15**	0.58	1.39**	1.79**	-0.17	0.34	1.14**	1.15**
	(2.83)	(0.53)	(2.04)	(2.82)	(-0.35)	(0.84)	(2.02)	(2.34)
Adj. R^2	0.12	0.04	0.00	-0.00	0.03	0.03	-0.01	0.01
N	184	184	184	184	184	184	184	184

***Panel B** - Annual Forecasting*

	Vol-Sorted Ret$_{t,t+4}$		Mkt-Rf$_{t,t+4}$		SMB$_{t,t+4}$		HML$_{t,t+4}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PVS_t	40.99**		-6.27		-7.39**		2.51	
	(4.11)		(-0.90)		(-2.16)		(0.50)	
Real Rate$_t$		4.13**		0.03		-1.21*		0.45
		(2.13)		(0.03)		(-1.76)		(0.54)
Constant	9.77**	2.49	5.85*	6.95**	0.65	1.97*	4.95**	4.50**
	(2.46)	(0.59)	(2.24)	(2.73)	(0.41)	(1.86)	(2.53)	(2.59)
Adj. R^2	0.26	0.07	0.01	-0.01	0.06	0.04	-0.00	-0.00
N	181	181	181	181	181	181	181	181

Notes: This table reports several return forecasting regressions. 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 Online 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. Vol-Sorted Ret in the forecasting regression corresponds to returns on 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 (3)-(8) in both panels forecast returns on the Fama and French (1993) factors. For quarterly regressions, standard errors are computed using Newey-West (1987) with two lags. 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 1970Q2-2016Q2. Returns are in percentage points.

Table 5: The Real Rate and Contemporaneous Volatility

Dependent Variable:	Real Rate (Level)			Real Rate (Changes)		
	(1)	(2)	(3)	(4)	(5)	(6)
σ (Vol-Sorted Portfolio)	0.00		-0.04	0.02		0.03
	(0.05)		(-1.63)	(0.84)		(1.19)
CIV_t		0.03	0.06**		-0.04**	0.00
		(1.19)	(3.28)		(-2.29)	(0.12)
σ (Mkt-Rf)		-0.18**	-0.05		-0.02	-0.05*
		(-3.94)	(-1.44)		(-0.66)	(-1.70)
σ (SMB)		0.31**	0.07		0.09	0.07
		(3.88)	(1.36)		(1.28)	(1.22)
σ (HML)		0.07	0.10**		0.03	0.02
		(0.94)	(2.22)		(1.40)	(0.88)
σ (TFP Growth)		0.09	0.06		0.71**	0.72**
		(0.19)	(0.31)		(2.54)	(2.81)
PVS_t			4.01**			2.50**
			(8.66)			(3.10)
Adj R^2	-0.01	0.15	0.60	0.00	0.08	0.20
N	185	185	185	184	184	184

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 first differences. Standard errors are computed using Newey-West (1987) with five lags. * 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 1970Q2-2016Q2.

Table 6: Forecasting Realized Volatility

***Panel A**: Forecasting Volatility Using PVS_t*

Dep Variable	Realized Volatility$_{t\rightarrow t+1}$					
	Vol-Sorted	CIV	MktRf	SMB	HML	TFP
	(1)	(2)	(3)	(4)	(5)	(6)
PVS_t	-2.46	-7.27*	-3.95*	-0.03	-2.34	-0.18
	(-1.21)	(-1.89)	(-1.88)	(-0.03)	(-1.08)	(-0.34)
Constant	10.95**	49.02**	13.88**	7.88**	6.49**	3.14**
	(11.89)	(22.27)	(17.03)	(17.61)	(11.69)	(16.57)
Adj. R^2	0.01	0.03	0.03	-0.01	0.03	0.00
N	184	184	184	184	184	184

***Panel B**: Forecasting Volatility Using the Real Rate*

Dep. Variable	Realized Volatility$_{t\rightarrow t+1}$					
	Vol-Sorted	CIV	MktRf	SMB	HML	TFP
	(1)	(2)	(3)	(4)	(5)	(6)
Real Rate$_t$	0.16	1.90*	0.15	0.38*	0.26	-0.05
	(0.33)	(1.84)	(0.39)	(1.72)	(0.97)	(-0.56)
Constant	11.39**	50.30**	14.59**	7.89**	6.91**	3.17**
	(13.15)	(23.74)	(15.06)	(19.20)	(11.04)	(22.89)
Adj. R^2	-0.00	0.07	-0.00	0.04	0.01	0.01
N	184	184	184	184	184	184

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. 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. 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 1970Q2-2016Q2.

Table 7: The Real Rate and Mutual Fund Flows

***Panel A**: Summary Statistics*

	Mean	Std. Dev.	p25	p50	p75	Min	Max	# Funds
# Year-Quarter Obs.	31	28	11	24	43	2	170	20,253
AUM ($ mm)	754	2,049	155	266	597	100	65,339	20,253
Net Inflows (%)	5.55	8.53	0.49	3.21	7.80	-19.70	66.54	20,253
Quarterly Return (%)	1.47	2.32	0.65	1.38	2.34	-38.77	58.76	20,253
Annual Volatility (%)	11.84	7.92	4.57	12.61	17.29	0.31	36.62	20,253
$\beta_{f,HVOL}$	0.30	0.24	0.02	0.32	0.49	-0.05	0.83	20,253

***Panel B**: High Volatility Funds and the Real Rate*

Dependent Variable			$Flows_{f,t}$			
	(1)	(2)	(3)	(4)	(5)	(6)
Real Rate$_t$	0.92**			0.94**		
	(4.56)			(4.27)		
Real Rate$_t \times \beta_{f,HVOL}$	2.08**	2.09**	1.52**			
	(4.11)	(4.17)	(4.30)			
Real Rate$_t \times \sigma_f$				0.04**	0.04**	0.03**
				(3.09)	(3.12)	(2.87)
$Ret_{f,t}$			0.22**			0.23**
			(6.47)			(6.46)
$Ret_{f,t-1}$			0.22**			0.22**
			(6.85)			(6.84)
FE	f	(f,t)	(f,t)	f	(f,t)	(f,t)
Adj. R^2	0.11	0.15	0.16	0.11	0.14	0.16
N	630,592	630,592	630,592	630,592	630,592	630,592

Notes: This table studies whether high-volatility mutual fund flows are more sensitive to real rate movements, relative to low-volatility mutual funds. In Panel B, our baseline regression is $Flow_{f,t} = FE(f) + b_1$Real Rate$_t$ $+ b_2$Real Rate$_t \times \beta_{f,HVOL} + \varepsilon_{f,t}$. $Flow_{f,t}$ is the net percentage inflow into fund f at time t, computed as the dollar inflow divided by assets under management. Flows are winsorized at the 5% tails. $\beta_{f,HVOL}$ is the beta of fund f's return with respect to a portfolio of high-minus-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. Betas of each fund are computed using the high-minus-low volatility portfolio return over the life of the fund. σ_f is the return volatility of the fund, computed using the full sample of year-quarter observations. We drop fund with assets under management of under $100 mm. Panel A presents summary statistics for the funds in our sample. We first compute statistics for each fund (across time), and then report summary stats across funds. In Panel B, standard errors are computed by using Driscoll-Kraay (1998) with five lags within each fund cluster. Quarterly mutual fund data derives from CRSP and spans 1973Q2-2015Q3. Returns are in percentage terms.

Table 8: Alternative Interest Rates and PVS_t

		Regression on PVS		
		b	$t(b)$	R^2
(1)	Baseline Detrended	3.44	5.36	0.41
(2)	Baseline Raw	3.87	5.65	0.38
(3)	Nominal 1-Year Rate	5.20	3.22	0.27
(4)	Expected Inflation	1.33	1.19	0.06
(5)	Fixed Taylor Rule Implied Rate (Taylor, 1993)	0.97	1.06	0.04
(6)	Residual	2.90	2.74	0.24
(7)	Fitted Taylor Rule Implied Rate	0.58	1.05	0.04
(8)	Residual	3.29	4.06	0.33
(9)	10Y-1Y Term Spread	-1.11	-2.36	0.11
(10)	BAA-10Y Spread	-0.84	-2.52	0.18
(11)	GZ Spread	-1.26	-2.24	0.23

Notes: This table reports univariate regressions of several variables on PVS. The Data Appendix contains full details on how we compute PVS_t , defined as the difference in book-to-market ratios between low and high volatility stocks. In Row (1), the dependent variable in the regression is the linearly detrended one-year real rate. The dependent variable in Row (2) is the raw one-year real rate. 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. Rows (3) and (4) decompose the raw one-year real rate into the one-year nominal rate and expected inflation. Rows (5) and (6) decompose the raw one-year real rate into a Taylor (1993) rule component and a residual component. The Taylor (1993) rule component is defined as $Taylor1993_t = 0.5 \times (OutputGap) + 0.5 \times (Inflation - 2) + 2$. The output gap is the percentage deviation of real GDP from the CBO's estimate of potential real GDP. Inflation is the annualized percentage four-quarter growth in the GDP price deflator from the St. Louis Fed (GDPDEF). The Taylor (1993) rule residual used in Row (5) is then Raw Real Rate$_t - Taylor1993_t$. Rows (7) and (8) use the same decomposition, where the fitted Taylor rule is defined as the fitted value from a regression of the raw real rate on the output gap and inflation. The Fitted Taylor Rule residual in Row (8) is the residual from the aforementioned regression. In Rows (9)-(11), the dependent variables are the 10-year minus 1-year nominal Treasury rate, the spread between Moody's BAA credit yields and the 10-year Treasury rate, and the credit spread index from Gilchrist and Zakrajšek (2012), respectively. Standard errors are computed using Newey-West (1987) with five lags. Data is quarterly and the full sample spans 1970Q2-2016Q2. The Gilchrist and Zakrajšek (2012) credit spread index runs from 1973Q1-2010Q3.

Table 9: 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 10: Model Moments

		(1) Data	(2) Model	(3) Rep. Agent	(4) Log Utility
Aggregate Stocks and Bonds					
Equity Premium	$E\left(R^e_{t+1} - R_{f,t}\right)$	6.54	**6.34**	**4.85**	0.98
Equity Volatility	$Std\left(R^e_{t+1}\right)$	17.69	**19.88**	**26.53**	2.35
AR(1) Agg. Book/Market	$AR\left(B_t/P_t\right)$	0.81	0.41	0.41	0.63
Std. Risk-Free Rate	$Std(R_{f,t})$	1.96	**1.75**	1.38	1.35
Low-Minus-High Volatility Portfolio					
PVS_t	$E\left(B_{L,t}/P_{L,t} - B_{H,t}/P_{H,t}\right)$	-0.18	-0.01	0.00	0.00
Std. PVS_t	$Std\left(B_{L,t}/P_{L,t} - B_{H,t}/P_{H,t}\right)$	0.37	0.25	0.00	0.01
Low-Minus-High Return	$E\left(R_{L,t} - R_{H,t}\right)$	2.71	**-1.63**	0.00	-0.03
The Risk-Free Rate and Equity Risk Premia					
Risk-Free Rate on Agg. Book-Market	$slope\left(R_{f,t}, \frac{B_t}{P_t}\right)$	-1.01	**-1.10**	8.73	46.67
Risk-Free Rate on PVS_t	$slope\left(R_{f,t}, \frac{B_{L,t}}{P_{L,t}} - \frac{B_{H,t}}{P_{H,t}}\right)$	3.44	**3.31**	NaN	-47.05
Return Spread on Lag Risk-Free Rate	$slope\left(R_{L,t+1} - R_{H,t+1}, R_{f,t}\right)$	4.13	**8.49**	0.00	-0.95

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. The risk-free rate and the aggregate book-to-market are detrended, as in the data. 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.