A Measure of Risk Appetite for the Macroeconomy

Internet Appendix

Carolin Pflueger* Emil Siriwardane† Adi Sunderam‡

Abstract

In this appendix, we provide details about the data construction for all variables used in the main text. We then present a battery of tests and additional analysis demonstrating the robustness of the relationship between the real rate and PVS_t. In addition, we show that roughly 90% of the covariation between the real rate and PVS_t stems from the fact that the real rate forecasts future returns on the vol-sorted portfolio. We also show that PVS_t is effectively uncorrelated with objective measures of aggregate uncertainty that come from statistical forecasting models. Moreover, we offer complementary VAR and local projection evidence that shocks to risk appetite, as measured by PVS_t, lead to a boom in the real economy. We then document that periods of low risk appetite coincide with investor outflows from high-volatility mutual funds. Finally, we provide proofs for the propositions contained in the model section of the main text.

*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.

A1 Data Construction

In this section we provide details on how we construct our main variables. We then provide details on the variables used in each table of the main text.

Construction of PVS_t

Valuation Ratios

Our valuation ratios (book-to-market) derive from the CRSP-COMPUSTAT merged databases. We augment CRSP-COMPUSTAT with the book value data used in Davis, Fama, and French (2000). We provide additional details of our variable construction below, but at a high level our procedure is as follows: for a given firm f on date t, we look for a valid value of book equity in COMPUSTAT Quarterly, then COMPUSTAT Annual, and finally the book values contained in Davis, Fama, and French (2000). We assume balance sheet information is known with a one-quarter lag. Finally, we combine the aforementioned book value with the trailing 6-month average of equity market capitalization to form a book-to-market ratio for firm f. We have confirmed that our results are not sensitive to these variable definition choices.

COMPUSTAT Quarterly: From COMPUSTAT Quarterly (COMPQ). Specifically, we obtain information on all firms (INDFMT = INDL) with a standardized data format (DATAFMT = STD) that report financial information at a consolidated level (CONSOL = C). In order to avoid the well-known survival bias in COMPUSTAT, we only include firms once they have at least 2 years of data.

We define book common equity (BE) according to the standard Fama and French (1993) definition. Specifically, BE is the COMPUSTAT book value of shareholder equity, plus balance-sheet deferred taxes and investment tax credit, minus the book value of preferred stock. We use the par value of preferred stock in COMPQ to estimate the value of preferred stock.

COMPUSTAT Annual: When using COMPUSTAT Annual (COMPA) for balance sheet information, we obtain information on all firms (INDFMT = INDL) with a standardized data format (DATAFMT = STD) that report financial information at a consolidated level (CONSOL = C). In order to avoid the well-known survival bias in COMPUSTAT, we only include firms once they have at least 2 years of data. For firms that change fiscal year within a calendar year, we take the last reported date when extracting financial data. This leaves us with one set of observations for each firm (gvkey) in each year.

We define book common equity (BE) according to the standard Fama and French (1993) definition. Specifically, BE is the COMPUSTAT book value of shareholder equity, plus balance-sheet deferred taxes and investment tax credit, minus the book value of preferred stock. Following Fama and French (1993), we use the redemption, liquidation, or par value (in that order) to estimate the value of preferred stock.

Defining Valuation Ratios: We then build book-to-market ratios at end of each quarter t as follows:

- The book equity comes from COMPQ, and we assume this data is known with a 3-month lag. This means we add three months to the DATADATE field in COMPQ to define the

"KNOWNDATE". Then at the end of each quarter, we take the book equity on the last available KNOWNDATE. For instance, this means that in June of a given year, we are using the book value of equity from COMPQ as of March in that same year. We prefer this definition because it uses up-to-date balance sheet information, while still allowing for reasonable lags to ensure the information was actually known by market participants at each date.

- If COMPQ does not have a valid book value, we obtain book equity from COMPA, again assuming a one-quarter information lag for balance sheet information. If COMPA also does not have a valid book value for a firm, we check the book equity values from Davis, Fama, and French (2000), which we downloaded from Ken French's website. For the book equity in Davis, Fama, and French (2000), we use their assumption that book values are known as of June 30 of the "Last_Moody_Year" variable.

- For the purposes of computing book-to-market ratios, we use the trailing 6-month average of market capitalization using CRSP Monthly. For instance, in June of a given year we take the average end-of-month market capitalization from January through June of that year. We prefer this definition because it smoothes out any high-frequency movements in equity valuations.

Book-to-market ratios for a given firm then follow naturally. We have also used the Fama and French (1993) definition of book-to-market ratios and obtain very similar results. Fama and French (1993) assume a more conservative lag in terms of when balance sheet is known and also use lagged market capitalization (e.g. in June of a year, use the previous December's market capitalization).

Volatility Used for Portfolio Sorts

At the end of each quarter, we compute each firm's stock return volatility as the standard deviation of ex-dividend returns (variable RETX) using daily data from the previous two months. 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.[1]

Computing PVS

At the end of each quarter t, we sort all stocks in the NYSE, AMEX, and NASDAQ into quintiles based on their total volatility. Total volatility is computed as described above. We then form equal-weighted portfolios based on the quintiles of volatility. Our measure of risk appetite is defined as:

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

In words, PVS_t is the average book-to-market ratio of firms in the low-volatility quintile minus the average book-to-market ratio of firms in the high-volatility quintile. Thus, PVS_t is high when the market valuations of high-volatility firms is large relative to low-volatility firms.

Finally, we define the aggregate book-to-market ratio for our universe of firms as the their total book value divided by their total market capitalization at time t.

[1] Our long-short portfolio effective replicates the one on Ken French's website. If we regress our portfolio on his, the point estimate is 0.84, the constant in the regression is statistically indistinguishable from zero, and the R^2 is 96%.

A1.1 Table 1 - Summary Statistics for Volatility-Sorted Portfolios and the Real Rate

The one-year real interest rate is the one-year constant maturity nominal Treasury rate from the U.S. Federal Reserve minus the one-year expectation of inflation (GDP deflator) from the Survey of Professional Forecasters. We linearly detrend the one-year real rate for all of our analysis in the main text, though we show that our analysis is robust to no detrending and alternative detrending methods in Section A2 of this appendix.

A1.2 Table 2 - What Explains Real Rate Variation?

The caption contains complete details on the variables used in the table.

A1.3 Table 3 - Robustness: The Real Rate and PVS

Panel A

Five-Year and Ten-Year Real Rate The five-year real interest rate is the five-year constant maturity nominal Treasury rate from the U.S. Federal Reserve minus the one-year expectation of inflation (GDP deflator) from the Survey of Professional Forecasters. The ten-year rate is defined analogously. We use one-year expectations of inflation because of data availability. For the analysis in the table, we linearly detrend both the five and ten-year real interest rates.

Value-Weighted Version of PVS_t The value-weighted version of PVS_t is the value-weighted average book-to-market ratio of low-volatility stocks at time t minus the value-weighted average book-to-market ratio of high-volatility stocks at time t. The value weights are determined by market capitalizations at the end of quarter t.

PVS_t Based on Two-Year Volatility The variable "2-Year Volatility" listed under "Alternative Constructions" in the table uses each firm's trailing 2-year volatility to form our volatility-sorted portfolios. We use monthly return data from CRSP to compute this measure of volatility.

Off-the-Run Minus On-the-Run Treasury Yields The off-the-run minus on-the-run Treasury yield 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).

Other Variables Used in Horse Races For a description of the other variables used in the horse races contained in Table 3, see Section A2.6 of this appendix.

Panel B

Panel B of Table 3 in the main text compares PVS_t to other measures of financial market conditions. Most of the variables are described in the caption of the table. Here, we focus on our measure of the time-t expectation of excess aggregate stock market returns from t to $t+4$, denoted $\mathbb{E}_t\left[\text{Mkt-Rf}_{t,t+4}\right]$. We obtain a statistically optimal measure of expected excess returns following the methodology developed in Kelly and Pruitt (2013). Specifically, we use the three-pass regression filter (3PRF) in Kelly and Pruitt (2013). In particular, we use the entire sample to estimate the 3PRF and assume two latent factors. In our experimentation with the procedure, using two factors balances the desire to have a good in-sample predictor of market returns against overfitting.[2]

The variables that we use as the predictors in the Kelly and Pruitt (2013) procedure are five BM ratios from sorting on each of the following variables: size, BM ratios, cash-flow duration (Weber (2016)), leverage, cash-flow beta with respect to aggregate cash flows, leverage, beta with respect to the aggregate market (using a 5-year window and a 10-year window), and total volatility. We construct BM ratios based on these sorts in the same way we do for PVS_t. In addition, we include the aggregate BM ratio, aggregate dividend yield, and CAY from Lettau and Ludvigson (2004). This gives us a total of 43 predictors that we feed into the 3PRF to forecast annual excess market returns. The R-squared in the forecasting regression is 14.2%. As a point of comparison, we are able to nearly double the forecasting power (in-sample) of CAY alone, which gives a forecasting R-squared of 7.5%. The sample size for this analysis is 180, and is lower than our main sample ($N = 185$) because the duration sorted portfolios that we include as predictor variables have a shorter sample.

A1.4 Table 4 - PVS, the Real Rate, and Future Returns to Volatile Assets

Panel A

Panel A of the table uses PVS_t and the one-year real rate to forecast returns and earnings surprises. Columns (1) and (2) forecast stock returns on the low-minus-high volatility portfolio. Columns (3) and (4) forecast this portfolio's accounting return on equity (ROE), which is defined according the clean-surplus accounting formula from Cohen et al. (2003). For each firm i and date t, we compute future annual ROE based on the next four quarters of financial statements after date t. Financial statement information is from the COMPUSTAT quarterly file. Because firms have different reporting periods, the calendar time over which we compute annual ROE differs across firms. Once we compute the future annual ROE for each firm i in quarter t, we aggregate to the portfolio level by taking the equal-weighted averages within each volatility quintile. Columns (5) and (6) uses PVS_t and the real rate to forecast excess returns on the CRSP Value-Weighted Index, which we obtained from Ken French's website.

Panel B

In Panel B of the table, we use both PVS_t and the one-year real rate to forecast returns on the low-minus-high volatility trade in other asset classes. To do so, we use the test assets from He

[2] We have tried a truly out-of-sample version of the 3PRF and obtain similar conclusions regarding the correlation with the real rate.

et al. (2017), henceforth HKM. We focus on the following asset classes from HKM: equities, U.S. corporate bonds, sovereign bonds, options, credit default swaps (CDS), commodities, and foreign exchange (FX). We refer the reader to HKM for more detail on each of these portfolios.

Within each asset class, we form a portfolio that is long the low volatility portfolio in that asset class, and short the high-volatility portfolio. For each portfolio in each asset class, we compute the volatility at each quarter using the trailing 5-year history of monthly portfolio returns, requiring a minimum of four years of data. We are constrained to use monthly data because HKM do not have daily asset class data. For example, suppose we want to form the low-minus-high volatility portfolio for U.S. corporate bonds in quarter t.[3] We then compute the volatility of each of the 10 HKM corporate bond portfolios over the previous 5 years. We then go long the portfolio with the lowest trailing volatility and short the portfolio with the highest volatility. We hold this long-short portfolio for one quarter, and then repeat the process. Denote the returns to this long-short strategy as $LMHV_t^c$, where the superscript c denotes the asset class we are studying and the subscript denotes time of the return. The forecasting regressions in Panel B of the table use PVS_t or Real Rate$_t$ to forecast $LMHV_{t+1}^c$ for several different c.

A1.5 Table 5 - Volatility-Sorted Returns and Monetary Policy Surprises

See the main text and the table caption.

A1.6 Table 6 - PVS and Real Outcomes

For both panels, see the main text and the table captions.

A1.7 Table 7 - PVS and Investor Expectations

This table looks at the contemporaneous correlation between PVS_t and several measures of expected cash-flow growth and risk. Most of the measures are defined in the main text and in the caption in the table. Here, we elaborate on our risk expectation measures (Panel B) that come from analyst forecasts and option prices. Note that our portfolio-level risk measures take the difference between high and low volatility firms. In other parts of the paper, such as return forecasting, we instead take the difference between low and high volatility firms. We reversed the direction of the long-short portfolio for our risk expectation analysis because we found it more intuitive to compare the price of volatile stocks (PVS_t) to the risk of high-volatility stocks (relative to low-volatility stocks).

IBES Based Measures Two of the measures we use derive from analyst forecasts of earnings-per-share (EPS) that come from Thompson Reuters IBES data. More specifically, we use the unadjusted summary file from WRDS. Data in IBES is organized by firm i, estimation date d, earnings announcement date u, and earnings type t. The two earnings types that we consider are annual and quarterly. We require at least two analyst forecasts for each (i,d,u,t). For this particular cut of the IBES data, we start the sample in 1989. Prior to 1989, the number of high-volatility firms that have a match in IBES fluctuates wildly, but steadily increases from 1989 onward.

[3]This corresponds to US_bond11 through US_bond20 in HKM's data.

For each firm i, quarter t, earnings announcement date u, and earnings t, we first select the last estimation date d that occurs prior to t. We then define the quarter t dispersion of firm i's earnings at time $u > t$:

$$\sigma^s_{i,t}(EPS_u) = \frac{\text{Range EPS Forecasts}_{i,d}(u)}{\text{Median EPS Forecasts}_{i,d}(u)}.$$

This is our proxy for analyst time-t expectations of earnings volatility at time u. We exclude firms where the median EPS forecasts is zero. In addition, because $\sigma_{i,t}(u)$ can be large for low median EPS forecasts, we winsorize it at its 5% and 95% tails.

In the table, we consider two different forecast horizons u. First, for quarterly earnings, we select u for each firm such that the earnings announcement corresponds to the next fiscal quarter (fpi = 6). We denote this case by $\sigma^s_{i,t}(EPS_{u=t+1})$. For annual earnings, we choose u such that the earnings announcement corresponds to two fiscal periods from t (fpi = 2). For our annual IBES measure, the average difference between u and t is five quarters, but it can vary depending on fiscal reporting periods and the availability of analyst forecasts. We denote this firm-level measure by $\sigma^s_{i,t}(EPS_{u=t+5})$.

Finally, $\sigma_t(\text{EPS}_{t+1})$ is the median $\sigma^s_{i,t}(EPS_{t+1})$ for high-volatility firms minus the median for low-volatility firms. $\sigma_t(\text{EPS}_{t+5})$ is assembled the same way from $\sigma^s_{i,t}(EPS_{t+5})$ at the firm level. In all cases, our classification of whether a firm is high or low-volatility at time t matches the portfolios used to compute PVS_t.

Option-Based Measure In Table 7 Panel B, we also build a measure of expected return volatility based on options data. Our options data comes derives from the Standardized Volatility Surface from OptionsMetrics on WRDS. For each firm, date, and horizon, the volatility surface contains at-the-money (ATM) put and call options. We define the expected return volatility on date d for horizon h as the average of the put and call implied volatilities. We denote this quantity by $\sigma^{IV}_{i,d}(Ret_{t,t+h})$. For this particularly analysis, we use h = 4 quarters. Option-based measures of expected volatility typically use the entire spectrum of option strike prices (i.e. the VIX). Due to the relatively scarcity of out-of-the-money options, especially for the high-volatility firms in our sample, we instead use ATM implied volatilities.

For each firm i and quarter t, we find the last available $\sigma^{IV}_{i,d}(Ret_{t,t+4})$ from the OptionMetrics database, requiring that the implied volatility was computed no more than 21 days prior to t. To aggregate to the portfolio level, we take median $\sigma^{IV}_{i,t}(Ret_{t,t+4})$ for high-volatility firms minus the median of low-volatility firms. The resulting variable is what we define in the table as $\sigma_t(Ret_{t,t+4})$. This measure begins in 1996Q3 because, prior to this date, we do not have any matches in OptionsMetrics for our high-volatility firms.

Model-Based Measure The variable "Model-Based $\sigma_t(Ret_{t,t+1})$" comes from a simple statistical forecasting model. Define the average realized volatility of high-volatility stocks in the portfolio at time t as $rv_{H,t}$, where each firm's volatility is computed as the daily standard deviation of returns in quarter t. $rv_{,t}$ is the same object for low-volatility firms and $rv_t \equiv rv_{H,t} - rv_{L,t}$. We fit an AR(1) process to rv_t using the full sample of returns. The estimated AR(1) coefficient for this series is 0.92, so rv_t is relatively persistent. The AR(1) model also fits the data well in terms of forecasting, as a simple regression of rv_{t+1} on rv_t yields an R^2 of 85%. Finally, the variable Model-Based $\sigma_t(Ret_{t,t+1})$ is defined as the $\mathbb{E}_t[rv_{t+1}]$ that emerges from the AR(1) model.

A1.8 Table 8 - What occurs in the rest of the economy during the build up of PVS?

We compute the trailing annual ROE (LMH-Vol $\text{ROE}_{t-4\rightarrow t}$) of the low-minus-high volatility portfolio in the same manner as described in Section A1.4.

A1.9 Table 9 - PVS and Revisions in Expectations

Our analysis of revisions in expected risk builds off the variable construction described in Section A1.7. In row (1) of the table, we forecast ROE surprises. Our measure of ROE surprises derives from the IBES Surprise History (unadjusted) file, downloaded from WRDS. The IBES Surprise History. We use the IBES-CRSP linking table on WRDS to match IBES data to CRSP-COMPUSTAT. We define each firm's quarterly ROE surprise as its quarterly earnings-per-share (EPS) surprise scaled by the previous quarter's book value of equity. We define the future annual ROE surprise for each firm i between dates $t+1$ and $t+4$ as the average of the four quarterly ROE surprises over this period. The future ROE surprise for the low-minus-high volatility portfolio is the median of future ROE surprises across high-volatility firms minus the median across low-volatility firms. We use medians to aggregate to the portfolio level in order to ensure that outliers in the IBES surprise do not exert too much influence on our results.

In row (2), we forecast revisions in ROE expectations for high-volatility firms relative to low-volatility firms. For each firm i and date t, we find the set of EPS forecasts in IBES corresponding to earnings that are three fiscal quarters away (fpi = 9). In calendar time, this ends up corresponding to quarterly earnings realized at $t+3$ for most firms, hence our notation. We translate the average EPS forecast to ROE by dividing by book equity at time t, where book equity comes from COMPUSTAT and follows the accounting timing that we use for PVS_t itself. At the firm level, we define the expected ROE as $\mathbb{E}_t[\text{ROE}_{i,t+3}]$. To build $\mathbb{E}_{t+2}[\text{ROE}_{i,t+3}]$, we hold the forecast date constant and recompute the expectation at time $t+2$. The revision in ROE expectation is then simply $\mathbb{E}_{t+2}[\text{ROE}_{i,t+3}]$ - $\mathbb{E}_t[\text{ROE}_{i,t+3}]$. To aggregate to the portfolio level, we compute the median revision in ROE expectation for high-volatility firms, less the median for low-volatility firms. The volatility quintile of firms is determined at time t.

In row(3), the variable $\sigma_{t+2}(\text{EPS}_{t+3}) - \sigma_t(\text{EPS}_{t+3})$ proxies for the revision in expected earnings volatility for earnings at time $t+3$. Let's start with how we construct $\sigma_t(\text{EPS}_{t+3})$. At time t and for each firm i, we find the set of IBES forecasts corresponding to earnings that are three fiscal quarters away (fpi = 9). Again, in calendar time, this ends up corresponding to quarterly earnings realized at $t+3$ for most firms. For this forecast horizon, we then build our dispersion measure, denoted by $\sigma^s_{i,t}(EPS_{u=t+3})$, as in Section A1.7. We also apply the same filters and methodology as described in that section. For each firm, we then hold fixed the date on which earnings will be realized and recompute our dispersion measure at time t. This delivers us $\sigma^s_{i,t+2}(EPS_{t+3}) - \sigma^s_{i,t+2}(EPS_{t+3})$, which is the news between time t and $t+2$ about expected earnings volatility at $t+3$. Again, our working assumption here is that our dispersion measure is a good proxy for expected earnings volatility. To aggregate this the portfolio level, we take the median $\sigma^s_{i,t+2}(EPS_{t+3}) - \sigma^s_{i,t+2}(EPS_{t+3})$ for high-volatility firms in the portfolio at time t minus the median for low-volatility firms. The resulting variable is defined as $\sigma_{t+2}(\text{EPS}_{t+3}) - \sigma_t(\text{EPS}_{t+3})$.

The measure $\sigma^{IV}_{t+3}(Ret_{t+4}) - \sigma^{IV}_t(Ret_{t+4})$ comes from options data. At time t and for each firm i, we define $\sigma^{IV}_{i,t}(Ret_{t+4})$ as the implied volatility of stock returns in quarter $t+4$. We use the term structure of option-implied volatilities to compute this measure. Specifically, we take the implied variance of 365-day options at time t and subtract off the implied variance of 273-day options at time t, which we then convert to an implied volatility measure[4] This is a valid approach to estimating the option-implied expected volatility between in quarter $t+4$ so long as there is negligible return autocorrelation at the quarterly frequency. For each firm, $\sigma_{i,t+3}^{IV}(Ret_{t+4})$ is the 90-day implied volatility at time $t+3$. This allows us to construct $\sigma_{i,t+3}^{IV}(Ret_{t+4}) - \sigma_{i,t}^{IV}(Ret_{t+4})$ for each firm, which we then aggregate to the portfolio level by taking the median across high-volatility firms minus the median for low-volatility firms.

Finally, for the realized risk measure $\Delta_4\sigma_{t+4}$(HML-Vol) is constructed as follows. For each firm in the high-volatility quintile at time t, we take its realized quarterly volatility at time $t+4$ and subtract its realized quarterly volatility at time t. We then average this difference across high-volatility firms and then repeat the entire process for low-volatility firms. $\Delta_4\sigma_{t+4}$(HML-Vol) is the resulting spread between the two groups and it measures the average change in volatility for high-volatility firms minus the average change for low-volatility firms.

A1.10 Table 10 - PVS and Implied Volatility Forecast Errors

See the main text and the table caption. For the time-t implied volatility of returns between $t+3$ and $t+4$, we use the same firm-level measure $\sigma_{i,t}^{IV}(Ret_{t+4})$ defined in Section A1.9 of this appendix.

A2 Robustness: PVS and the Real Rate

The purpose of this section is to conduct a several of robustness tests to ensure that our statistical inference regarding the relationship between the real rate and PVS is not driven by specific choices in defining our main variables. We begin by discussing alternative methods of filtering the real rate (e.g. using a deterministic versus stochastic trend). We then show that our results are largely unchanged with these alternative filters or if we simply study the raw real rate. We conclude the section by exploring several ways of adjusting the standard errors in our regressions of the real rate on PVS that account for the fact that these variables are persistence. The main takeaway of the section is that there is a robust relationship - both in economic and statistical terms - between the real rate and PVS_t. For the remainder of this appendix, we use R_t to denote the raw real rate.

A2.1 Filtering the Raw Real Rate

The top panel of Figure A.1 plots the raw real rate R_t from 1970Q2 to 2016Q2. The downward trend in R_t has received recent attention from many macroeconomists who argue that it reflects a form of economic secular stagnation (e.g. Summers (2015)). In this paper, we do not focus on the longer-run trend in R_t, but rather the large cyclical variation around this trend. Our goal is to better understand the determinants of cyclical (i.e., quarterly) movements in the real rate.

[4]Implied volatilities from OptionMetrics standardized volatility surfaces are annualized, so we first translate them to annualized implied variances. We take the 365-day implied variance minus 0.75 times the 273-day variance. This provides an unannualized estimate of return variance between $t+3$ and $t+4$, which we then annualize by multiplying by 4 and then take the square root to arrive at an implied volatility measure.

To achieve this goal, we need to empirically extract the cyclical component of the real rate. In the main text, we use a simple linear deterministic trend to do so:

$$R_t = \beta_0 + \beta_1 t + r_t \tag{1}$$

Here, the detrended real rate r_t is just the sequence of residuals from the regression. We chose this approach because it is simple and transparent. Still, it is fair to wonder whether a deterministic (downward) linear trend is a plausible model of the economy's real interest rate. No economic theory would predict the real rate to tend towards negative infinity over the next fifty years. A natural alternative that we explore now is to allow for a stochastic drift in the real interest rate. In short, real rates look extremely similar whether we remove a linear or stochastic trend, consistent with the finding that it is extremely difficult to distinguish between deterministic and stochastic trends in finite samples (Campbell and Perron (1991)).[5]

Specifically, we follow Hamilton (2017) to extract the cyclical component of R_t in the presence of a potentially stochastic drift. For quarterly data, Hamilton (2017) recommends the following regression to achieve the filter:

$$R_t = k_0 + k_1 \times R_{t-8} + k_2 \times R_{t-9} + k_3 \times R_{t-10} + k_4 \times R_{t-11} + \widetilde{r}_t \tag{2}$$

where the cyclical component of R_t is captured by the regression residuals, denoted here by $\widetilde{r}_t$. Importantly, this filtering methodology is relatively agnostic about the underlying trend driving the series.[6] This is particularly useful in our context because, again, we are not interested in understanding longer-run trends in R_t. Hamilton (2017) also provides an extensive argument for why regression (2) is superior to the more standard Hodrick-Prescott filter.

The bottom panel of Figure A.1 plots the linearly detrended real rate (r_t) and what we call the Hamilton-filtered real rate ($\widetilde{r}_t$). A visual inspection shows that r_t and $\widetilde{r}_t$ are quite similar. That is, linearly detrending and using the Hamilton-filter appear to give similar estimates for the cyclical component of the real rate. A regression of one on the other, run in both levels and first-differences, confirms this intuition:

$$\begin{aligned} \widetilde{r}_t &= \underset{(0.01)}{0.002} + \underset{(8.56)}{0.71} \times r_t, \qquad R^2 = 0.56 \\ \Delta\widetilde{r}_t &= \underset{(-0.29)}{-0.01} + \underset{(40.44)}{0.99} \times \Delta r_t, \qquad R^2 = 0.85 \end{aligned}$$

where Newey-West t-statistics with five lags are listed below point estimates. Both specifications indicate that the linearly detrended real rate is fairly close to the Hamilton-filtered rate. The constant in both regressions is near zero, the point estimate on r_t is near one, and the R-squared's are pretty large. As a result, we focus on the simpler, linearly detrended real rate in the main text and repeat our core analysis on the Hamilton-filtered rate now. To be certain that detrending (in any fashion) is not driving our conclusions, we also show our results using the raw real rate R_t.

[5] We think of the stochastic or non-stochastic drift as a simple way of controlling for long-run output growth. For example, Holston et al. (2016) embed this type of thinking in their statistical model of the natural rate of interest. They model the natural rate of interest as the sum of two random walks, one of which also drives the stochastic drift of potential output growth.

[6] In fact, Hamilton (2017) argues that it is still a useful method for extracting the cyclical component of a series that has a deterministic time trend.

A2.2 Results Using $\tilde{r}_t$ and the Raw Real Rate

A2.2.1 The Real Rate and PVS

Table A.1 shows regressions of the form:

$$Y_t = a + b \times PVS_t + \theta^{'}\mathbf{X_t} + \xi_t$$

where $\mathbf{X}_t$ is a vector of control variables and Y_t is either the Hamilton-filtered rate $\tilde{r}_t$ or the raw rate R_t. In all cases, we standardize PVS_t to have a mean of zero and variance one. We do the same to the aggregate book-to-market ratio when it is included as a control variable.

Results with $\tilde{r}_t$ Columns (1)-(6) run the regression for the Hamilton-filtered real rate, $\widetilde{r}_t$. The control variables that we use are the aggregate book-to-market ratio, the output gap, and the inflation rate. For consistency, we extract the cyclical components of these variables using Hamilton (2017) before including them in the regression. Echoing our results in the main text, the relationship between PVS and $\widetilde{r}_t$ is robust across level and first-difference specifications, and is not altered much by the addition of our control variables. Column (2) adds the aggregate book-to-market ratio as a control to the regression, which has very little effect on both the point estimate on PVS, as well as the R^2 in the regression. Indeed, a univariate regression of $\widetilde{r}_t$ on the aggregate book-to-market yields an R^2 of less than 1%. In terms of economic significance, a one-standard deviation move in PVS_t impacts $\widetilde{r}_t$ nearly three times as much as a one-standard deviation in the aggregate book-to-market. Column (3) of Table A.1 adds the output gap and inflation to the level regression of $\widetilde{r}_t$ on PVS. Again, we include these variables to check whether PVS is just picking up on Taylor (1993) rule variables. The Hamilton-filtered rate does load positively and significantly on the output gap, which is what we would expect if the central bank follows some version of a Taylor (1993) rule. The important thing though is that the inclusion of these variables does not impact the point estimate or statistical significance of PVS in the regression. The results using the Hamilton-filtered rate also compared favorably to those using the simple linear detrending in the main text.

Results with R_t Columns (7)-(9) repeat the analysis for the raw real rate R_t. Importantly, in this case, we do also not do any filtering to the control variables – these regressions only use raw variables. Column (7) runs a univariate regression of the raw real rate on PVS. The regression coefficient of 1.42 is economically comparable to the point estimate of we get when using the detrended real rate (see Table 2 in the main text). The R-squared is also comparable to our main results at 0.38. Columns (8) and (9) add the aggregate book-to-market and the output gap and inflation as control variables. While the aggregate book-to-market enters significantly, the R-squared in columns (7) and (8) is almost the same, indicating that the explanatory power of PVS_t for the real rate is much stronger than that of the aggregate book-to-market. A univariate regression of the raw real rate on the raw aggregate book-to-market ratio delivers an R^2 of less than 10%, much less than when using PVS_t. More importantly, none of our conclusions regarding the relationship between the real rate and PVS are impacted.

A2.2.2 The Real Rate and the Aggregate Stock Market

In the previous section, there were some specifications where the point estimate on the aggregate book-to-market ratio was estimated with some measure of statistical precision. Overall though, there is very little evidence suggesting that the valuation of the aggregate stock market contains meaningful information about the dynamics of the real interest rate. For one, the aggregate book-to-market ratio explains a very small amount of variation in the real rate. This is true regardless of how or whether we detrend these variables. Moreover, the relationship is nonexistent when we difference the data and when we linearly detrend the real rate and the aggregate BM ratio, as we do in the main text. In addition, there is ample empirical evidence that variation in the aggregate value of the stock market is largely disconnected from real rate variation (e.g. Campbell and Ammer (1993)). In sum, we do not view the evidence in Section A2.2.1 to reveal a robust link between the real rate and the aggregate BM ratio. In unreported results, we draw similar conclusions if we instead use Shiller's CAPE ratio or CAY from Lettau and Ludvigson (2004).

Even if there is a weak relationship between the real rate and the aggregate value of the stock market in our sample, it is likely unrelated to risk premia. Standard Gordon growth model logic suggests that the aggregate dividend-yield is driven by the risk-free rate r_f, the market risk premium $\mathbb{E}\left[r_m - r_f\right]$, and the growth rate of aggregate dividends g:[7]

$$D/P = r_f + \mathbb{E}\left[r_m - r_f\right] - g$$

The simple formula immediately illustrates the mechanical relationship between the risk-free rate and the dividend-yield. Of course, D/P and r_f may also correlate if the risk-free rate is also related to the market risk premium or aggregate dividend growth. However, Table 4 in the main text and Panel A of Table A.2 demonstrate that the real rate contains no forecasting power for excess market returns. Moreover, in Table A.1, in the cases where the aggregate book-to-market enters significantly for the real rate, the point estimate is positive. This is the opposite of what we would expect if risk appetite drives both the aggregate market and the real risk-free rate. On the contrary, the positive point estimates are consistent with a simple Gordon growth formula above.

Furthermore, Panel A of Table A.2 also demonstrates that both the Hamilton-filtered and raw real rate still forecast returns on the low-minus-high volatility portfolio. In Panel B of Table A.2, we show that the real rate – both the Hamilton-filtered and raw series – has no forecasting power for aggregate real earnings growth or aggregate real dividend growth. In conclusion, the link between the aggregate BM ratio and the real rate appears unrelated to risk appetite, or more likely, nonexistent.

A2.3 Time-Series Inference

The AR(1) coefficients of the Hamilton-filtered rate $\widetilde{r}_t$, the linearly detrended rate r_t, and PVS_t are 0.81, 0.85, and 0.88, respectively. While the persistence of PVS_t may appear high, it is useful to keep in mind that it is much less persistent than the aggregate valuation ratios, where persistent regressor biases have found the most attention in asset pricing (Stambaugh (1999)). While PVS_t has a quarterly AR(1) coefficient of 0.88, corresponding to a half-life of about 1.5 years, the aggregate book-to-market has an AR(1) coefficient of 0.98, corresponding to a much longer half-life of around 10 years. This simple comparison already suggests that inference problems from persistent regressors are likely to be much less severe in our setting than for aggregate valuation ratios.

[7] A similar argument holds for the aggregate book-to-market ratio, but the dividend-price ratio is easier for the purposes of this illustration. As an empirical matter, the two are 98% (60%) correlated in levels (first-differences) for our sample.

Nonetheless, we use a battery of approaches to formally establish that the relationship between the real rate and PVS_t is not driven by serially correlated regressors. First, we run all our main results in differences, as shown throughout the main text and the appendix. In this section, we explore several ways of adjusting standard errors, GLS, and a bootstrap simulation exercise.

Note: For this particular analysis, we leave PVS_t in its natural units, though for most of the analysis in the main text and in the rest of the appendix we standardize it to have mean zero and variance one.

A2.3.1 Standard Error Corrections

Our baseline univariate regression of the linearly detrended real rate (r_t) on PVS yields the following estimates:

$$\begin{array}{ccccccc} r_t & = & 0.62 & + & 3.44 & \times & PVS_t \\ & & (5.02) & & (11.41) & & \\ & & [2.64] & & [5.36] & & \end{array}$$

where the parenthesis below the point estimates contain OLS t-statistics and the square brackets contain Newey-West t-statistics with five lags. The first thing to note from this simple regression is that Newey-West correction still indicates the point estimate on PVS is statistically significant. The second thing to note is that the nonparametric Newey-West correction shrinks the OLS t-statistic by a factor of nearly two. This owes in part to the fact that the regression residuals have a first-order autocorrelation of 0.76. We address this persistence directly by using a standard parametric correction based on on the estimated residual autocorrelation. Specifically, we multiply the standard errors in the regression by a factor of $C = (1+\rho)/(1-\rho)$, where ρ is the autocorrelation of the regression residuals. $\rho = 0.76$ means that $C \approx 7.3$, thereby implying that the OLS t-statistics need to be divided by a factor of $\sqrt{C} = 2.71$. The parametric correction therefore shrinks the t-statistic on PVS from 11.41 to 4.21, so the point estimate is still statistically significant.

For completeness, we repeat the analysis using the Hamilton-filtered real rate $\widetilde{r}_t$. In this case, a univariate regression of $\widetilde{r}_t$ on PVS_t gives:

$$\begin{array}{ccccccc} \widetilde{r}_t & = & 0.59 & + & 3.28 & \times & PVS_t \\ & & (5.11) & & (11.57) & & \\ & & [2.73] & & [6.53] & & \end{array}$$

The first-order autocorrelation of the residuals for this specification is 0.69, implying that the OLS t-statistic of 11.57 should be adjusted to 4.96.

The broader takeaway here is that no matter how we adjust our standard errors, we are still able to comfortably reject the null that the point estimate on PVS is equal to zero.

A2.3.2 Generalized Least Squares (GLS)

For statistically efficiency and to account for the role of outliers, we also estimate the relationship between the linearly detrended real rate and PVS using generalized least squares. This is just a Prais-Winsten regression, which amounts to quasi-differencing the data before running the regression. GLS gives the following estimates:

$$\begin{array}{ccccccc} r_t & = & 0.44 & + & 2.47 & \times & PVS_t \\ & & (1.32) & & (6.15) & & \end{array}$$

where the GLS t-statistics are listed below point estimates. We also estimate the same system using the Hamilton-filtered real rate $\widetilde{r}_t$:

$$\underset{}{\widetilde{r}_t} = \underset{(1.90)}{0.49} + \underset{(6.35)}{2.59} \times PVS_t$$

Regardless of the detrending method, the relationship between the real rate and PVS remains economically and statistically significant when using GLS.

Moreover, if we run the regression using data up until the financial crisis (pre-2009), we get fairly similar point estimates on PVS across simple OLS and GLS estimation methods. For example, when using the Hamilton-filtered real rate, OLS gives a point estimate on PVS of 3.44 and GLS gives a point estimate of 3.28.

A2.3.3 Simulation Evidence

Finally, one might be concerned that our results are biased in a Granger-Newbold sense. We use simulations to show that the standard error and R^2 from our baseline regression are not just a result of regressor persistence. Specifically, we fit independent AR(1)-GARCH(1,1) models to r_t and PVS and simulate these processes mimicking the persistence properties of r_t and PVS_t and with identical sample length as in the data. In the simulated data, where by construction r_t and PVS_t are unrelated, we regress r_t on PVS_t, retaining the Newey-West corrected t-statistic (five lags) for PVS and the R^2 in the simulated regression. Figure A.2 presents histograms of the simulated t-statistics and R^2 from this exercise for 10,000 independent simulations. The plot also shows the actual t-statistic on PVS and the R^2 that we estimate in the data. The p-values listed in the plot are just the proportion of simulations where the t-statistic (or R^2) exceed the actual t-statistic we estimate in the data. For both the t-statistic and R^2, less than 0.5% of simulations can match the regression of the real rate on PVS that we estimate using actual data. Combined with the other analysis in the paper, this tells us that under the null of no relation between PVS_t and r_t it would be highly unlikely to observe the t-statistics and R^2s that we see in the data. This simulation once again adds to our evidence that the relation between PVS_t and r_t is a real feature of the data and not just an erroneous statistical artifact.

A2.4 Subsample Stability

Our main sample runs from 1970Q2 through 2016Q2. In this subsection, we study the sub-sample stability of the relationship between the real rate and PVS_t. We start by showing that our results are not dependent on the period from 1977 to 1987, a time when the U.S. suffered unusually high inflation and the Federal Reserve – led by Paul Volcker – tightened monetary policy to regain control over inflation. In addition, we expand our sample back to 1953Q2 and show the relationship between PVS_t and the real rate is equally strong in this longer sample. The beginning of this extended sample coincides with the beginning of the series for the constant maturity nominal one-year rate that is available from the St. Louis FRED database. The Survey of Professional Forecasters inflation forecasts are not available prior to 1970Q2, so to construct a one-year real rate series from 1953Q2 to 1970Q2, we use the four-quarter moving average of realized inflation as our measure of expected one-year inflation. This approach for forming expected inflation forecasts is motivated by the findings of Atkeson and Ohanian (2001).[8] To extend PVS_t back to 1953Q2, we use the accounting data from Davis et al. (2000). Specifically, we look for book values from COMPUSTAT quarterly, then COMPUSTAT annual, then Davis et al. (2000), in that order and depending on data availability.

In all cases, we run regressions of the following form, in both levels and first differences:

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

Table A.3 collects the results of our subsample analysis. For reference, column (1) of the table presents the level-regression results using the baseline sample in the main text. In column (2), we find similar results when we exclude the period from 1977 to 1987, providing some comfort that our results are not dependent on the so-called "Volcker period". Column (3) runs the regression over 1953Q2-2016Q2. For this sample, we use the raw real rate because it appears stationary during this time period.[9] Here, we once again see that the correlation between the real rate and PVS_t is present in the longer sample. In column (4), we focus on the portion of the longer sample that precedes the Volcker period, again confirming a strong link between the real rate and PVS_t. Columns (5)-(8) indicate that we obtain similar conclusions when running these regressions in first-differences. Overall, the main takeaway from Table A.3 is that the relationship between PVS_t and the real rate is robust across subsamples.

A2.5 Inflation Expectations and the Taylor Rule

Our construction of the one-year real interest rate is simply the nominal one-year Treasury rate minus expected one-year inflation from the Survey of Professional Forecasters. Thus, PVS_t can correlate with our real rate variable because it correlates with one of these components. To explore this potential further, we decompose the real rate into its constituent parts and regress both on PVS_t. Table A.4 contains the results, and in all regressions, PVS_t is standardized to have zero mean and unit variance. For sake of comparison, we present the results of regressing the detrended real rate and the raw real rate on PVS_t in rows (1)-(2) of the table, respectively. 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, not inflation expectations.

In rows (5)-(8), we try to separate movements in the real rate that can be attributed to the Taylor (1993) rule, which sets the real short-term interest rate 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 explore 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

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

[8] There is a large body of research that studies optimal inflation forecasts, with varying conclusions depending on the subsample of interest. The four-quarter random walk benchmark studied in Atkeson and Ohanian (2001) is surprisingly successful and we use it here due to its simplicity.

[9] For this sample, the augmented Dickey-Fuller test rejects the null of a random walk with no drift at conventional significance levels.

where $Taylor1993_t$ is the real rate that obtains if the central bank follows the Taylor rule exactly. 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 does not simply capture the reaction of monetary policy along a standard Taylor (1993) rule.

A2.6 The Real Rate and Other Valuation Spreads

A2.6.1 Univariate Analysis

We now explore alternative explanations for the empirical relationship between the real rate and PVS_t. This analysis complements our findings in Section 3.1.2 of the main text. Specifically, we examine the possibility that volatility is simply correlated with another characteristic that is more important for explaining the real rate. We sort stocks along a variety of dimensions and form book-to-market spreads based on the sorting variable. For instance, when examining size as a characteristic, we sort stocks in quintiles based on their market capitalization, then compute the difference between the book-to-market ratio of the smallest (i.e., the lowest quintile of market capitalization) and the largest stocks. Recall that PVS_t is the book-to-market spread that emerges when the characteristic Y is trailing 60-day volatility. We then run the following regression relating the real rate to the spread in book-to-market based on each sort:

$$\text{Real Rate}_t = a + b \times Y_t + \varepsilon_t \tag{3}$$

where Y_t is the book-to-market spread based on sorting on characteristic Y. In all cases, we standardize Y_t to have a mean of zero and a variance of one. For reference, column (1) of Table 1 of the main text runs regression (3) with PVS_t as the explanatory variable. There, we find that a one standard deviation increase in PVS_t is associated with a 1.27 percentage point increase in the one-year real rate and PVS_t alone explains 41% of real rate variation over our main sample.

The results are displayed in Table A.5. In row (1), we relate the real rate to the spread in book-to-market sorting stocks based on the expected duration of their cash flows. If high volatility stocks simply have higher duration cash flows than low duration cash flows, then their valuations should fall more when real rates rise. This is one sense in which low volatility stocks may be more "bond-like" than high volatility stocks (e.g., Baker and Wurgler (2012)). In this case, a mechanical duration effect could explain the relationship between the real rate and PVS. To examine this possibility, we follow Weber (2016) and construct the expected duration of cash flows for each firm in our data. We then sort stocks based on this duration measure and calculate the spread in book-to-market between high and low duration stocks. As row (1) shows, the relationship between this duration spread and the real rate is negative. However, it is not consistently statistically significant across specifications and is in general much smaller in magnitude than PVS.

Row (2) displays the same exercise when looking at the relative valuations of low-leverage versus high-leverage stocks. We define leverage as the book value of long-term debt divided by the market value of equity. It seems natural to think that high-leverage firms have high volatility, and since these firms effectively are short bonds, their equity may suffer disproportionately from a decrease in the real rate. The positive coefficient in row (2) indicates that this intuition bears out in the data. When the real rate falls, the book-to-market spread between low-and-high leverage firms also falls. In other words, high-leverage firms become cheaper when the real rate falls.

In rows (3)-(5), we sort stocks based on three types of market (CAPM) betas:

1. The first CAPM beta we compute is a two-year rolling beta. In a given quarter, we use the previous twenty-four months worth of monthly return data to compute a CAPM beta. In order to have a valid two-year beta, a firm must have at least 80% of its observations over the previous two years.

2. The second CAPM beta we compute is a "long-run" beta. We first aggregate monthly returns into six-month returns. Then at the end of each quarter we use the previous ten years worth of data to compute betas from our six-month return series (e.g. 20 observations per regression). Once again, firms must have 80% of their observations in order to have a valid long-run beta.

3. The third CAPM beta we compute uses a two-month window. For each firm, we use daily stock data from the previous two months to compute a high-frequency measure of CAPM beta. We exclude firms that do not have at least 20 observations over this time frame.

In all cases, our benchmark index is the CRSP Value-Weighted Index. For the first two measures of CAPM Beta, all of our individual firm data derives from the CRSP Monthly dataset. We deal with delisted returns as in Shumway (1997) by setting missing delisted returns with codes 400-591 to a value of -30%.

Row (3) indicates that the book-to-market spread based on a two-year CAPM beta is correlated with the real rate. Row (4) sort stocks based on CAPM betas that we compute using long-horizon returns. The motivation for studying longer-run CAPM betas is that long-horizon returns are more plausibly driven by cash flow news rather than discount rate news. Thus, long horizon CAPM betas can be viewed as a measure of aggregate cash flow beta. Row (4) indicates a positive relationship between the book-to-market spread based on long-run CAPM beta in levels, but the relationship is not particularly strong in a statistical sense when moving to first-differences. Row (5) uses our measure of CAPM beta that is computed using daily data over rolling 60-day windows. This construction mimics how we compute volatility (and hence PVS). There is again a positive relationship between 2-month beta and the real rate, but not one that is robust across specifications.

In row (6), we sort stocks on the estimated beta of their cash flows with respect to aggregate cash flows. Specifically, cash flow betas are computed via rolling twelve quarter regressions of quarter-on-quarter EBITDA growth on quarter-on-quarter national income growth. EBITDA is defined as the cumulative sum of operating income before depreciation (series oibdpq from COMPUSTAT quarterly). We require a minimum of 80% of observations in a window to compute a cash flow beta. If high volatility stocks have higher cash flow betas than low volatility stocks, then their valuations should fall more when aggregate growth expectations are low. In this case, our results using PVS could be explained by changes in aggregate growth expectations rather than change in the precautionary savings motive. Row (6) shows that the book-to-market based on cash flow betas is not significantly correlated with the real rate.

Keep in mind that the preceding regressions are all univariate. The relevant question for us is whether PVS is just picking up on the information carried in these various book-to-market spreads.

Two pieces of evidence strongly suggest that PVS carries independent information about the real rate. For one, in Table 3 of the main text, we run bivariate horse races of PVS against each of these alternative sorting variables. None of these alternative sorting variables drive out PVS from the regression. This is true when running the horse races in levels, first differences, and across different subsamples.

As a second piece of evidence, in row (9) of Table A.5 we run a "kitchen-sink" regression of the following form:

$$\text{Real Rate}_t = a + b_{PVS} \times PVS + \theta^{'} X_t + \varepsilon_t$$

where X_t contains all of the valuation spreads discussed above. Row (9) of the table reports the estimated b_{PVS}, its associated t-statistic, and the adjusted R^2 from the regression. The simple takeaway from the kitchen-sink regression is that none of the control variables drive out the explanatory power of PVS for the real rate. The coefficient on PVS remains statistically significant in both the levels and first-differenced specifications, and the point estimate compares favorably to those found in the main text. If anything, including the other control variables increases the economic relationship between PVS and the real rate. These results suggest that the relative valuation of high and low-volatility stocks contains unique information about the real rate.

A2.6.2 Double-Sorted Versions of PVS_t

In this subsection, we create double-sorted versions of PVS_t as an alternative way to address the possibility that volatility just proxies for another characteristic whose price is correlated with the real rate. More precisely, consider characteristic Y. We construct a Y-neutral version of PVS_t by first grouping stocks at time t based on whether they have above or below median values of characteristic Y. We define "low Y" firms as those firms with below-median values of Y and "high Y" firms are defined analogously. Next, within low-Y firms, we further sort firms into terciles based on volatility. $PVS_t^{L,Y}$ is defined as the average book-to-market ratio of low-volatility and low-Y firms minus the average book-to-market value of high-volatility and low-Y firms. $PVS_t^{H,Y}$ is defined in the same manner, except for high-Y firms. Finally, the Y-neutral version of PVS_t is defined as $(PVS_t^{L,Y} + PVS_t^{H,Y})/2$. This spread measures the difference in valuations of low volatility and high volatility stocks that have similar values of characteristic Y.

For example, suppose the characteristic that we are interested in is CAPM-Beta. We then split stocks into low and high beta firms based on the median CAPM-Beta at time t. Then within each CAPM-Beta bucket, we compute the difference in book-to-market ratios of low and high volatility stocks. Finally, we average the spread between low- and high-volatility stocks across low and high CAPM-Beta firms. This procedure delivers us a version of PVS that is immunized to CAPM-Beta but differentially exposed to volatility. The sorting variables we use are described in Section A2.6 of this appendix. In addition, we construct an industry-neutral version of PVS_t in the same way by first grouping stocks into industries based on their SIC codes and the 48 industry definitions on Ken French's website. We also form PVS_t in the subset of dividend paying and non-dividend paying stocks, where we define a dividend-paying stock at time t as one that has paid a divided any time in the previous two years.[10]

After we build double-sorted versions of PVS_t, we run the following regression in both levels

[10] We determine dividend yields by looking at the total return and the ex-dividend return in CRSP.

and first differences:

$$\text{Real Rate}_t = a + b \times \text{Y-Neutral PVS}_t + \varepsilon_t \tag{4}$$

In all cases, we standardize the double-sorted version of PVS_t (or its first difference) to have a mean of zero and a variance of one. Table A.6 contains the point estimates of b, their associated t-statistics, and the adjusted R^2 from these regressions. Echoing our analysis from Section A2.6.1, we find that all of the double-sorted versions of PVS_t exhibit an economically and statistically significant positive correlation with the real interest rate. By and large, the point estimate on the Y-neutral version of PVS_t is comparable to what we obtain in the main text when using the raw version of PVS_t. The fact that the industry-adjusted version of PVS_t continues to explain a large fraction of real rate variation indicates that PVS_t does not just load up on industries that are more exposed to interest rate movements. A similar conclusion holds when looking at dividend versus non-dividend paying stocks. Overall, these facts lend further support of the idea that volatility is the key characteristic underlying the construction of PVS_t.

A2.6.3 Total Volatility vs. Alternative Measures of Risk

As discussed in the main text, we use total volatility of stock returns because it is a robust measure of risk. Intuitively, volatility increases with stocks' exposure to any risk factors that investors care about, and PVS_t captures how much of a price discount investors require for holding risky stocks. To confirm that our results are robust to variations in how we measure risk, we verify that the spread in book-to-market ratios is similar when we sort stocks by their CAPM betas instead of total stock return volatility. The CAPM beta captures systematic risk provided that investors are well-diversified and that investors' aggregate wealth portfolio equals the aggregate stock market. Indeed, we find that PVS_t and the beta-sorted book-to-market spread are 82% correlated in levels and 51% correlated in first-differences. As our preceding analysis shows, the link between the real risk-free rate and the spread in book-to-market ratios when sorting on two-year CAPM betas is similar, albeit weaker, than our baseline results for PVS_t.

Of course, investors may care about risk factors other than the aggregate stock market. To allow for a broader set of factors, we sort firms into quintiles based on the volatility of the fitted value from a regression of daily stock returns on the Fama and French (1993) factors. To match the construction of our benchmark sorting variable, i.e. total volatility, we use trailing 60-day returns at the end of each quarter t. As expected, the resulting book-to-market spread is even more closely correlated with PVS_t, with correlations of 87% in levels and 84% in first-differences. Overall, we find that PVS_t is not sensitive to small variations in our measure of risk. We use total volatility as our benchmark sorting variable because it does not require us to take a stand on the underlying risk-factors that investors care about.

A3 Additional Results

A3.1 Additional Return Forecasting Results

A3.1.1 Covariance Decomposition

Vuolteenaho (2002) derives the following relation tying a firm i's log book-to-market ratio to its future log return and log accounting return (ROE):

$$\theta_{i,t} = r_{i,t+1} - e_{i,t+1} + \rho \theta_{i,t+1} + \nu_{it}$$

where θ_i is the log book-to-market of firm i, $r_{i,t+1}$ is its log stock return, and $e_{i,t+1}$ is the log ROE. ρ is a log-linearization constant and $\nu_{i,t}$ is an approximation error, such that $\theta_{i,t} \approx r_{i,t+1} - e_{i,t+1} + \rho \theta_{i,t+1}$. To map this expression to the current setting, we define the log version of PVS_t, denoted by pvs_t, as follows:

$$pvs_t \equiv \left[\frac{1}{N_{L,t}} \sum_{i \in \text{Low Vol}_t} \theta_{i,t} \right] - \left[\frac{1}{N_{H,t}} \sum_{i \in \text{High Vol}_t} \theta_{i,t} \right]$$

where, for example, $N_{L,t}$ is the number of firms in the low vol portfolio at time t. The Vuolteenaho (2002) decomposition then implies that:

$$\begin{aligned} pvs_t &\approx r^{PVS}_{t+1} - e^{PVS}_{t+1} + \rho \times pvs_{t+1} \\ r^{PVS}_{t+1} &\equiv \left[\frac{1}{N_{L,t}} \sum_{i \in \text{Low Vol}_t} r_{i,t+1} \right] - \left[\frac{1}{N_{H,t}} \sum_{i \in \text{High Vol}_t} r_{i,t+1} \right] \\ e^{PVS}_{t+1} &\equiv \left[\frac{1}{N_{L,t}} \sum_{i \in \text{Low Vol}_t} e_{i,t+1} \right] - \left[\frac{1}{N_{H,t}} \sum_{i \in \text{High Vol}_t} e_{i,t+1} \right] \end{aligned} \tag{5}$$

In addition, we assume that pvs_t follows an AR(1) process, $pvs_{t+1} = a + \phi pvs_t + \xi_{t+1}$. Next, combining the AR-process with Equation (5), plus some rearranging yields:

$$\begin{aligned} Cov(\text{Real Rate}_t, pvs_t) \approx (1-\rho\phi)^{-1} \times [&Cov\left(\text{Real Rate}_t, r^{PVS}_{t+1}\right) - Cov\left(\text{Real Rate}_t, e^{PVS}_{t+1}\right) \\ &+ \rho Cov(\text{Real Rate}_t, \xi_{t+1})] \end{aligned}$$

Dividing both sides by $Cov(\text{Real Rate}_t, pvs_t)$ delivers a simple covariance decomposition:

$$1 = \Psi_r - \Psi_e + \Psi_\xi \tag{6}$$

where $\Psi_r \equiv (1-\rho\phi)^{-1} \times Cov\left(\text{Real Rate}_t, r^{PVS}_{t+1}\right) / Cov(\text{Real Rate}_t, pvs_t)$, and so forth.

Equation (6) states that covariation between today's real rate and pvs_t can arise for three reasons: (i) today's real rate forecasts future returns to the volatility-sorted portfolio, r^{PVS}; (ii) today's real rate forecasts future cash flows on the same portfolio, e^{PVS}; or (iii) today's real rate forecasts future innovations in tomorrow's pvs.

To operationalize the decomposition, we need to first estimate ϕ and ρ. We fit a simple AR(1) for pvs and find that $\phi = 0.88$ for quarterly data. With regards to ρ, we consider a range of values from 0.9 to 0.97.[11] All of the other components needed for the covariance decomposition are estimated from simple covariances in the data, namely one-quarter ahead forecasting regressions of returns and ROEs on PVS_t.[12]

[11] Vuolteenaho (2002) sets $\rho = 0.967$ for annual data. We use a range of values to get a sense of how sensitive our decomposition is to the approximation constant.

[12] Note that in estimating $Cov\left(\text{Real Rate}_t, e^{PVS}_{t+1}\right)$ by forecasting future ROE with PVS_t, we are imposing that investors have rational expectations of the cash flows of high-volatility versus low-volatility firms. As discussed in the main text, this assumption is justified by the fact that PVS_t does not forecast surprises in ROE based on analyst forecasts, nor does it correlate with analyst expectations of cash flows. To be clear, it could be that movements in the expected return of high-versus-low volatility stocks are still driven by behavioral forces and irrational expectations of risk. We explore this issue in the main text as well.

For all of the ranges of ρ that we consider, Ψ_r is never less than 70% and approaches 100% for larger values of ρ. Moreover, for all of the ranges of ρ considered in Vuolteenaho (2002), Ψ_r is never below 90%. This is rather unsurprising given that the real rate does not forecast future ROE for the low-minus-high volatility portfolio. We therefore conclude that a large majority of the covariation (around 90%) between PVS_t and the real rate comes from the real rate forecasts future returns on the volatility-sorted portfolio. Put differently, PVS_t and the real rate correlate because discount rate shocks to high-volatility stocks coincide with shocks to the real rate. This simple fact is a large reason we interpret PVS_t as a measure of risk appetite.

A3.2 Monetary Policy Shocks - All announcements

In Section 3.3.1 of the main text, we show that monetary policy shocks do not differentially affect high-volatility stocks. In the main text, we exclude unscheduled FOMC meetings because surprise policy changes made outside of regularly scheduled meetings may be driven by financial market conditions. Here we examine the full sample of FOMC meetings. The results are in Table A.9. If reverse causality was responsible for our baseline result, high-volatility stocks should increase in response to a positive shock to interest rates. Since the independent variable is the low-minus-high volatility return, reverse causality should therefore show up as negative coefficients. Generally the coefficients are statistically insignificant with inconsistent signs. Using daily returns, we find a positive correlation that is borderline statistically significant for some specifications. However, this is the opposite of what we would expect if there was reverse causality. Instead, a positive correlation is consistent with the Fed cutting interest rates and stabilizing high-volatility stocks in times of market turmoil. Consistent with this interpretation, in untabulated results we find that the positive correlation is entirely driven by surprise changes in 2001. In that year, the Fed cut rates aggressively outside of regularly scheduled meetings in the aftermath of the technology bubble.

A3.3 PVS and Real Outcomes

A3.3.1 Evidence from VARs

In the main text, we used local projections from the main text that show how shocks PVS_t impact real investment, GDP growth, and unemployment. In this section, we complement that analysis with standard vector autoregression (VAR) evidence. This evidence shows that monetary policy shocks and shocks to PVS_t have opposite effects on economic activity, as predicted by the standard New Keynesian model (Clarida et al., 1999; Woodford, 2003).

We estimate a VAR that is as simple and transparent as possible, while following a common set of recursiveness assumptions, similar to Sims (1980), Bernanke and Mihov (1998) and Gilchrist and Zakrajšek (2012). We use the following strategy for measuring dynamic effects:

$$Y_t = \sum_{i=1}^{k} B_i Y_{t-i} + \sum_{i=1}^{k} C_i P_{t-i} + A^y v_{y,t} \tag{7}$$

$$P_t = \sum_{i=0}^{k} D_i Y_{t-i} + \sum_{i=0}^{k} G_i P_{t-i} + \begin{bmatrix} v_{PVS,t} \\ v_{MP,t} \end{bmatrix}. \tag{8}$$

Here, Y_t is a vector of quarterly non-policy variables, consisting of unemployment, the investment-to-capital ratio, and detrended inflation. P_t is a vector of policy variables consisting of PVS_t and the detrended real rate. Eq. (7) describes a set of structural relationships in the economy, where macroeconomic variables depend on lagged values of macroeconomic and policy variables. Eq. (8) describes the stance of monetary policy conditional on contemporaneous macroeconomic variables. Our baseline estimation uses $k = 1$ lag.

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 $v_{PVS,t}$, consistent with the Federal Reserve actively monitoring macroeconomic and financial variables. It is plausible that investors' risk preferences shift gradually over time and do not jump in response to monetary policy actions. Indeed, this identification restriction is supported by our analysis of monetary policy shocks in the main text.[13]

More formally, we can rewrite the system (7)-(8) in VAR form with only lagged variables on the right-hand-side and estimate by OLS. Let $u_p = A^p v_p$ be the VAR residuals in the policy block, which are assumed to be orthogonal to the VAR residuals in the non-policy block as in Bernanke and Mihov (1998). We further assume that the market for the risk-free bond is described by the following set of equations:

$$u_{PVS} = \alpha v_{MP} + v_{PVS}, \tag{9}$$

$$u_{rr} = \phi v_{PVS} + v_{MP}. \tag{10}$$

Eq. (9) is the innovation in investors' demand for bonds. It states that the demand for low-volatility assets depends on monetary policy shocks v_{MP} and the structural innovation v_{PVS}. Eq. (10) describes central bank behavior. We assume that the Fed observes and responds to risk appetite shocks. The model described by Eqs. (9) and (10) has four unknown parameters: α, ϕ, and the two structural shock variances, σ^2_{PVS}, and σ^2_{MP}.

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

$$E\left[u^2_{PVS} - \sigma^2_{PVS}\right] = 0, \tag{11}$$

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

$$E\left[u^2_{rr} - \phi\sigma^2_{PVS} - \sigma^2_{MP}\right] = 0. \tag{13}$$

We estimate the parameters ϕ, σ_{PVS}, and σ_{MP} by two-step GMM using a Bartlett kernel with two lags and the initial weighting matrix equal to the identity. Our baseline identification is that $\alpha = 0$, so we have a just-identified setup with three moments identifying three parameters (ϕ, σ_{PVS}, σ^2_{MP}).

Both Wald and Hansen J-tests provide clear evidence that the real rate reacts contemporaneously to PVS_t, consistent with the Federal Reserve reacting to risk appetite shocks. For the reaction coefficient $v_{MP,t}$ onto $v_{PVS,t}$, we obtain a point estimate of 2.33 with a t-statistic of 3.71. 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 0.0008.

[13] This identification restriction is not crucial to our findings. As we show below, our 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 risk appetite demand with 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.

As a baseline, the left panel of Figure A.3 shows responses to an unexpected tightening by the Federal Reserve. Consistent with the long literature on monetary policy shocks, summarized in Christiano et al. (1999), unemployment increases and and inflation decreases after a one-standard-deviation shock to the real interest rate. The effect on the investment-to-capital ratio is not statistically different from zero. Interestingly, PVS_t does not respond to monetary policy shocks with tight 95% confidence intervals, consistent with our interpretation that risk appetite shocks drive the real rate, and not vice versa.

The right panel of Figure A.3 shows that a positive PVS_t shock significantly decreases unemployment and increases real investment, despite being associated with a similar increase in the real rate as the MP shock. The difference in responses across the left and right panels in Figure A.3 is exactly what the Euler equation would suggest if an increase in PVS_t acts as a demand shock increasing the natural real rate.

Risk appetite shocks are both statistically significant and quantitatively important for unemployment and investment, as shown by forecast error variance decompositions. Ten quarters after the shock, PVS_t shocks explain 12% of variation in the unemployment rate and 21% of the variation in investment-to-capital ratios. It is intuitive that risk appetite shocks should matter most for real investment, since it is the interface between financial market attitudes towards risk and the real economy. For comparison, the monetary policy shocks explain 24% of variation in unemployment and only 2.5% of variation in the investment-to-capital ratio.

Robustness Figure A.4 shows that impulse responses look similar to Figure A.3 if we use the pre-crisis sample. Figure A.5 shows that our findings are not dependent on the specific identification assumption. We see that unemployment and investment responses are similar if we make the alternative identification assumption that PVS_t is faster than the real rate. Under this alternative identification restriction, we estimate α freely while assuming that $\phi = 0$. With this alternative ordering, the forecast error variance decomposition at 10 quarters out attributes an even greater role to risk appetite shocks than our baseline specification. With this alternative ordering, PVS_t shocks explain 22% of variation in the unemployment rate and 23% of variation in the investment-to-capital ratio 10 quarters out.

Figure A.6 shows that again the conclusion is similar if instead of estimating a VAR(1) we include additional lags in our estimation and base the impulse responses on a VAR(4). Finally, Figure A.7 shows that we obtain similar results if we replace the unemployment rate by the output gap. Of course, the output gap responses have the opposite signs of the unemployment responses in our baseline specification, because the output gap decreases in recessions, whereas unemployment increases. So, our results linking shocks to risk appetite and the real economy are not specific to a particular measure of economic activity.

A3.3.2 Evidence from Jordà (2005) Local Projections

In Section 3.4.2 of the main text, we use Jordà (2005) local projections to trace out how a shock to risk appetite (PVS) propagates through the macroeconomy, finding that risk appetite shocks lead to a boom in investment, an expansion of output, and a decline in unemployment. In that analysis, we control for lagged outcome variables and the real interest rate. Here, we explore the robustness of those local projections by running the following sequence of regressions:

$$y_{t+h} = a + b^h_{PVS} \times PVS_t + b^h_{RR} \times RealRate_t + b^h_y \times y_t + b^h_{mkt} \times \text{Agg BM}_t + b^h_{cp} \times CP_t + \varepsilon_{t+h}$$

where h is the forecast horizon. y_{t+h} is either the investment-to-capital at time $t+h$, the real output gap at $t+h$, or the change in the unemployment rate between t and $t+h$. In this regression, we also control for the aggregate book-to-market ratio (Agg BM_t) and the Cochrane and Piazzesi (2005) bond risk-factor, the latter of which we construct using quarterly data and forward rates from Gürkaynak et al. (2007). We include the aggregate book-to-market ratio and the Cochrane and Piazzesi (2005) factor to test whether PVS_t reflects redundant information embedded in measures of financial market activity from aggregate stock and bond markets.

Figure A.8 displays the results of these local projections. The main thing to notice is the magnitude of the response of the macroeconomy to a risk appetite shock is very similar in these specifications compared to those shown in Figure 3 of the main text. Following a risk appetite shock, investment and the output gap both rise and unemployment falls, even when controlling for the value of the aggregate stock market and the Cochrane and Piazzesi (2005). These results therefore suggest that PVS_t contains information about the real side of the economy that are not contained in these alternative financial market indicators.

A3.4 Additional Proxies for Objective Expectations of Risk

In Section 3.3 of the main text, we studied the contemporaneous relationship between PVS_t and the following measures of expected risk: (i) subjective expectations of earnings volatility of high-volatility stocks (relative to low-volatility) from analyst forecasts; (ii) expected return volatility of high-volatility stocks based on option prices; (iii) objective expectations of return volatility for high-volatility stocks, where objective expectations are defined from the perspective of a statistical forecasting model; and (iv) the percent of loan officers loosening lending standards, which is plausibly related to their subjective expectations of risk. Our main finding is that PVS_t is highly correlated with subjective measures of risk and weakly correlated with objective measures of risk from statistical forecasting models. In this section, we extend our analysis to include objective measures of risk for several macroeconomic variables and also consider an alternative technique for forming objective expectations of the risk of high-volatility stocks, relative to low-volatility stocks.

To start, we construct empirical forecasts for the volatility of the following aggregate quantities: real GDP growth, real consumption growth, and industrial production growth. For each of these variables, we fit ARMA(1,1)-GARCH(1,1) models and then define the objective expectation of risk as the one-period forecast of volatility from the GARCH component of the model, denoted by $\sigma_{t+1}(Y)$ for variable Y. The recursive structure of the GARCH model means that σ_{t+1} is known at time t, so we then compare $\sigma_{t+1}(\cdot)$ with both PVS_t and the real rate at time t.

Rows (1)-(3) of Table A.7 collect the results of this exercise. We standardize the expected risk measures variables to facilitate comparison of magnitudes. The main takeaway is a null result — in our sample there is almost no evidence that the real rate or PVS_t correlate with econometric forecasts of macroeconomic uncertainty. The first set of columns regress PVS_t on the econometric forecasts of macroeconomic uncertainty and finds small t-statistics and low R^2s, indicating little correlation. The second set of columns regress the real rate on the same right-hand-side variables and again find very little correlation. The latter finding is related to Hartzmark (2016), who studies the relationship between the real interest rate and a measure of macroeconomic uncertainty. Hartzmark (2016) defines macroeconomic uncertainty is as the variance forecast that comes from fitting an ARMA(1,1)-GARCH(1,1) model to several different macroeconomic series, so very much in the same manner that we do. Using annual data from 1890-2010, Hartzmark (2016) finds that there is a negative relationship between the real interest rate and macroeconomic uncertainty. In contrast, we find that for our subsample of quarterly data, macroeconomic uncertainty from econometric models shows a relatively weak correlation with the real rate, especially compared to PVS_t. In row (4), we find similar results using the macroeconomic uncertainty index from Jurado et al. (2015), which may better filter noise from volatility forecasts. Jurado et al. (2015) define the uncertainty of a macroeconomic series as the conditional volatility of the purely unforecastable component of that series. They employ sophisticated econometric techniques to compute uncertainty measures for a wide range of macroeconomic and financial series, and then combine them into a single aggregate index of macroeconomic uncertainty. As row (4) shows, the Jurado et al. (2015) macroeconomic uncertainty index has little correlation with both PVS_t and the real rate.

In row (5), we build a measure of the expected volatility of the aggregate stock market. From 1986 onward, we follow Bloom (2009) and use the VXO implied volatility index of the S&P 100. Options data is not available prior to 1986, so we use the one-step ahead forecast from fitting an AR(1) model to the within-quarter realized volatility of the aggregate stock market, which we scale to create a smooth series when the VXO becomes available. The results show that both PVS_t and the real rate tend to be lower when expected market volatility is high, but the strength of the relationship is statistically weak and the R^2s from these regressions are quite low. Overall, rows (1)-(5) of the table indicate a weak relationship between PVS_t and these objective measures of uncertainty based on aggregate variables, which is consistent with our broader view that the cross-section of equities that is embedded in PVS_t provides valuable information about the state of the economy.

To that end, in rows (6) and (7), we examine measures of risk specific to the portfolio that underlies PVS_t. In row (6), we forecast the volatility of the average high-volatility stock and subtract out the forecasted volatility of the average low-volatility stock. We denote the time-t value of this difference by $\overline{\sigma}_{H,t} - \overline{\sigma}_{L,t}$. To form our forecast $\mathbb{E}_t\left[\overline{\sigma}_{H,t+1} - \overline{\sigma}_{L,t+1}\right]$, we fit an AR(1) model to each component and then use the forecast that comes from fitting the model. This approach exactly replicates column (4) of Table 5 in the main text. Relative to aggregate quantities of risk, the first set of regressions in row (6) shows an increase in the correlation between PVS_t and this portfolio-based measure of risk relative, though the R^2 of 9% in the regression indicates a somewhat modest relationship between the two. In row (7), we follow Hamilton (2017) to extract the cyclical component of $\overline{\sigma}_{H,t} - \overline{\sigma}_{L,t}$ before forming forecasts with an AR(1) model. We do so because it is well known that the individual stock volatility has increased during our sample period Campbell et al. (2001). As row (7) shows, this measure of objective expected risk explains 16% of the variation in PVS_t. Rows (6) and (7) both indicate that PVS_t is indeed high when objective expectations of the risk of high-volatility stocks is low.

As discussed in the main text, the fact that PVS_t is largely uncorrelated with objective forecasts of macroeconomic risk – yet is quite correlated with subjective expectations of risk – suggests that movements in risk appetite may reflect a departure from fully rational expectations. The possibility that investors have less-than-rational expectations dates back to at least Keynes (1937) and Minsky (1977). At the same time, our results do not rule out complementary channels of time-varying risk appetite, such as time-varying risk aversion as in Campbell and Cochrane (1999). For example, in the extreme, if PVS_t was driven entirely by time-varying risk aversion then we would expect to see a weak relationship between PVS_t and objective expectations of risk. Still, the fact that we observe a strong correlation between PVS_t and subjective expectations of risk suggests that time-varying risk aversion is unlikely to be the primary reason that PVS_t fluctuates. Our paper highlight expectations of risk, not risk aversion, in part due to data limitations – there are more direct measures of expectations of risk than of risk aversion.

A3.5 Revisions in Expected Risk

In the main text, we use options to study revisions from quarter t to $t+3$ in the expected volatility of stock returns that will be realized between $t+3$ and $t+4$. In particular, we test whether PVS_t can forecast these revisions, which we infer from the implied volatility embedded in option prices.

To formalize our approach, first define the time-t conditional variance of returns between $t+k$ and $t+h$, denoted by $R_{t+k,t+h}$, as:

$$\begin{aligned} \mathbb{V}_t\left(R_{t+k,t+h}\right) &\equiv \mathbb{E}_{t+k}\left[R^2_{t+k,t+h}\right] - \mathbb{E}^2_{t+k}\left[R_{t+k,t+h}\right] \\ &= \mathbb{E}_t\left[\mathbb{V}_{t+k}\left(R_{t+k,t+h}\right)\right] + \mathbb{V}_t\left(\mathbb{E}_{t+k}\left[R_{t+k,t+h}\right]\right) \end{aligned} \tag{14}$$

where the second equality follows from the law of total variance.

Next, define the news about variance between t and $t+k$ as:

$$\begin{aligned} \eta_{t+k} &\equiv \mathbb{E}_{t+k}\left[\mathbb{V}_{t+k}\left(R_{t+k,t+h}\right)\right] - \mathbb{E}_t\left[\mathbb{V}_{t+k}\left(R_{t+k,t+h}\right)\right] \\ &= \mathbb{V}_{t+k}\left(R_{t+k,t+h}\right) - \mathbb{E}_t\left[\mathbb{V}_{t+k}\left(R_{t+k,t+h}\right)\right] \end{aligned} \tag{15}$$

Our approach in the main text effectively focuses on the following object:

$$\theta_{t+k} \equiv \mathbb{V}_{t+k}\left(R_{t+k,t+h}\right) - \mathbb{V}_t\left(R_{t+k,t+h}\right),$$

which we can easily construct using option prices at time t and $t+k$. θ_{t+k} is not exactly the same as the news about expected variance, but is close. To concretely relate the two, substitute Eq. (14) into Eq. (15) and rearrange to get:

$$\theta_{t+k} = \eta_{t+k} - \mathbb{V}_t\left(\mathbb{E}_{t+k}\left[R_{t+k,t+h}\right]\right) \tag{16}$$

In the data, we use PVS_t to forecast θ_{t+k}. However, to ensure our point estimates on PVS_t are not biased in this regression, we should also control for $\mathbb{V}_t\left(\mathbb{E}_{t+k}\left[R_{t+k,t+h}\right]\right)$. Thus, the remaining task is to construct $\mathbb{V}_t\left(\mathbb{E}_{t+k}\left[R_{t+k,t+h}\right]\right)$ in the data. To do so, let's focus on the case where $k=3$ and $h=4$, as we do in the main text. Next, notice that we can write:

$$\mathbb{V}_t(\mathbb{E}_{t+3}[R_{t+3,t+4}]) = \mathbb{E}_t\left\{\mathbb{E}^2_{t+3}[R_{t+3,t+4}]\right\} - \mathbb{E}^2_t[R_{t+3,t+4}] \tag{17}$$

We can form an estimate of $\mathbb{E}_{t+3}[R_{t+3,t+4}]$ by regressing $R_{t+3,t+4}$ on PVS_{t+3}. In turn, the square of the fitted value from this forecasting regression provides an estimate of $\mathbb{E}_t\left\{\mathbb{E}^2_{t+3}[R_{t+3,t+4}]\right\}$. Similarly, we can construct an estimate $\mathbb{E}^2_t[R_{t+3,t+4}]$ based on the square of the fitted value from a regression of $R_{t+3,t+4}$ on PVS_t. Combining the two yields an proxy for $\mathbb{V}_t\left(\mathbb{E}_{t+k}\left[R_{t+k,t+h}\right]\right)$, which we then add as a control to our forecasting regression. The results are presented below:

$$\theta_{t+3} = \underset{(0.28)}{\underset{0.05}{a}} + \underset{(2.99)}{\underset{0.47}{b_1}} \times PVS_t + \underset{(5.05)}{\underset{0.004}{b_2}} \times \mathbb{V}_t\left(\mathbb{E}_{t+3}\left[R_{t+3,t+4}\right]\right)$$

where point estimates are listed below the coefficients and t-statistics based on Newey-West standard errors with five lags are in parenthesis. As is clear from the regression, controlling for the time-t variance of expected returns at $t+3$ does not change the main conclusion that PVS_t forecasts revisions in risk. Moreover, if we just regress θ_{t+3} onto PVS_t, the point estimate is basically unchanged at 0.48. With this in mind, and to keep the exposition as simple as possible, in the main text we focus on predicting revisions in volatility as opposed to variance. In addition, we do not control for the time-t variance of expected returns at $t+3$.

A3.6 PVS, the Real Rate, and Mutual Fund Flows

Throughout the main text, we 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 subsection, we provide evidence that a specific but important class of investors, namely mutual funds investors, behaves consistently with the evidence from prices. If real rate variation indeed reflects time variation in risk appetite, we expect investors to leave high-volatility mutual funds when the real rate is low. Specifically, a decrease in risk appetite should lead to outflows from high-volatility mutual funds, an increase in the demand for bonds, and a drop in the real rate. Mutual fund flows are also useful because they allow us to separately verify our baseline results in a completely different data set.

Our sample is the CRSP mutual fund data base, from which we have monthly return data from 1973q2 through 2015q3. 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.[14]

Panel A of Table A.8 contains summary statistics for our sample of mutual funds. The average fund appears in our sample for 31 quarters and has around $750 million in assets under management. We find substantial heterogeneity in mutual funds' exposure to volatile stocks, regardless of how we measure exposure. The average fund has an annualized return volatility of about 12%, though this ranges from 4.6% to 17.3% when moving from the 25th to 75th percentile of fund volatility. Similarly, the beta of fund returns with respect to the high-volatility portfolio is 0.30 for the average fund, with a cross-sectional standard deviation of 0.24. This cross-sectional dispersion in volatility exposure allows us to study how movements in risk appetite differentially impacts our sample of mutual funds.

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

In Panel B of Table A.8, we explore the relationship between fund flows, the real rate, and fund volatility. Specifically we run

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

where $VolExp_f$ is a measure of the fund's exposure to high-volatility stocks. In columns (1)-(3) $VolExp_f$ 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. For all regressions, we use Driscoll-Kraay standard errors, clustered by fund and time with five lags.[15]

Panel B of Table A.8 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 the fund's 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.[16] 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 PVS_t is a good measure of risk appetite. An obvious caveat to our mutual fund analysis is that there must be a buyer for every seller, meaning outflows do not necessarily have to lead to a change in the price of high-volatility securities. Nonetheless, there is ample evidence that mutual fund outflows cause price pressure in equity markets (e.g., Coval and Stafford (2007)). Furthermore, we are not claiming that flows out of high-volatility equity mutual funds are solely responsible for the contemporaneous fall in real rates. These results simply provide a glimpse into the behavior of investors that we think is representative of the broader economy. Indeed, the fact that the real rate forecasts returns on the low-minus-high volatility trade in other asset classes suggests that investors in those asset classes likely behave similarly.

A4 Model Appendix

In this appendix, we provide proofs for the model propositions.

A4.1 Subjective Volatility Process

Consumption growth is described by the following process:

$$\begin{aligned} \Delta c_{t+1} &= \varepsilon_{t+1}, && (18)\\ \varepsilon_{t+1} &= \sigma_t \eta_{t+1}, && (19)\\ \sigma_t^2 &= exp(a - b\varepsilon_t), && (20) \end{aligned}$$

[15] We have also tried double clustered errors by fund and time and are reporting the more conservative standard errors.

[16] Because we include fund fixed effects, the base effect of the fund's volatility is absorbed.

where η_{t+1} is iid standard normal.

Our derivation of the subjective distribution follows Gennaioli and Shleifer (2018), Chapters 5 and 6. Their Proposition 1 in Chapter 5 states the following:

> Suppose that $ln\tilde{X}|I_0 \sim N\left(\mu_0, \sigma_0^2\right)$ and $ln\tilde{X}|I_{-1} \sim N\left(\mu_{-1}, \sigma_{-1}^2\right)$. Then, provided $(1+\theta)\sigma_{-1}^2 - \theta\sigma_0^2 > 0$, the distorted density $f^{\theta}(\tilde{X}|I_0)$ is also lognormal with mean $\mu_0(\theta)$ and variance $\sigma_0^2(\theta)$ given by:

$$\mu_0(\theta_0) = \mu_0 + \frac{\theta\sigma_0^2}{\sigma_{-1}^2 + \theta\left(\sigma_{-1}^2 - \sigma_0^2\right)}(\mu_0 - \mu_{-1}), \tag{21}$$

$$\sigma_0^2(\theta) = \sigma_0^2 \frac{\sigma_{-1}^2}{\sigma_{-1}^2 + \theta\left(\sigma_{-1}^2 - \sigma_0^2\right)} \tag{22}$$

As in (Gennaioli and Shleifer, 2018), Chapter 6, we assume that agents' reference distribution is the distribution at the state vector in the absence of news. That is, the reference distribution for ε_{t+1} before learning ε_t is the distribution at the conditional average of ε_t, i.e. at $\mathbb{E}_{t-1}(\varepsilon_t) = 0$. This gives $\mu_{-1} = 0$, $\sigma_{-1}^2 = exp(a)$, $\mu_0 = 0$, and $\sigma_0^2 = exp(a - b\varepsilon_t)$.[17] Substituting into the Proposition gives the subjective mean and variance for ε_{t+1} after having learned ε_t:

$$\mathbb{E}_t^{\theta}(\varepsilon_{t+1}) = 0, \tag{23}$$

$$\mathbb{V}_t^{\theta}(\varepsilon_{t+1}) = exp(a - b\varepsilon_t)\frac{1}{1+\theta(1 - exp(-b\varepsilon_t))}, \tag{24}$$

$$= \frac{1}{1+\theta(1 - exp(-b\varepsilon_t))}\sigma_t^2. \tag{25}$$

The subjective variance is therefore distorted relative to the objective variance by a factor of $\frac{1}{1+\theta(1-exp(-b\varepsilon_t))}$.[18]

Taking the comparative static in the vicinity of $\varepsilon_t = 0$ then gives Proposition 2 a):

$$\left.\frac{d\mathbb{V}_t^{\theta}(\varepsilon_{t+1})}{d\varepsilon_t}\right|_{\varepsilon_t=0} = -exp(a)b(1+\theta) < 0. \tag{26}$$

A4.2 Stochastic Discount Factor and Real Rate

The representative agent maximizes expected discounted power utility over consumption:

$$U(C_t, C_{t+1}, ...) = \sum_{i=0}^{\infty}\beta^i \frac{C_{t+i}^{1-\lambda}}{1-\lambda}. \tag{27}$$

[17] We follow Bordalo et al. (2018) in considering the distribution at the conditional average of ε_t as the reference distribution for simplicity and tractability. Alternatively, we could also consider the case where the reference distribution equals the conditional distribution of ε_{t+1} conditional on knowing ε_{t-1}. This would give $\mu_{-1} = 0$ and $\sigma_{-1} = \sqrt{exp(a + \frac{1}{2}b^2\sigma_{t-1}^2)}$. This would make the solution more complicated but preserve the main qualitative feature that the subjective variance $\mathbb{V}_t^{\theta}(\varepsilon_{t+1})$ reacts more to ε_t than the objective variance $\mathbb{V}_t(\varepsilon_t)$.

[18] Technically the proposition only applies if $1+\theta(1-exp(-b\varepsilon_t)) > 0$, or if the variance does not increase excessively. We follow (Bordalo et al., 2018) in imposing this condition, which holds with probability one in the the perfectly rational limit with $\theta \to 0$.

The representative agent's optimal portfolio choice gives that the time-t price of a time $t+1$ payoff is priced by the stochastic discount factor:

$$M_{t+1} = \delta\left(\frac{C_{t+1}}{C_t}\right)^{-\lambda}, \tag{28}$$

$$= \delta exp\left(-\lambda\varepsilon_{t+1}\right). \tag{29}$$

The time-t log real risk-free rate is then given by the asset pricing Euler equation under the subjective distribution:

$$1 = \mathbb{E}_t^{\theta}\left[exp(r_{ft})M_{t+1}\right], \tag{30}$$

$$= exp(r_{ft})\delta exp\left(\frac{1}{2}\lambda^2\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)\right), \tag{31}$$

implying that

$$r_{ft} = -\ln(\delta) - \frac{1}{2}\lambda^2\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right). \tag{32}$$

Taking the comparative static of r_{ft} in the vicinity of $\varepsilon_t = 0$ gives:

$$\left.\frac{dr_{ft}}{d\varepsilon_t}\right| = -\frac{1}{2}\lambda^2 \left.\frac{d\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)}{d\varepsilon_t}\right|_{\varepsilon_t=0}. \tag{33}$$

Substituting in for $\left.\frac{d\mathbb{V}_t^{\theta}(\varepsilon_{t+1})}{d\varepsilon_t}\right|_{\varepsilon_t=0}$ from equation (26) gives Proposition 2 c):

$$\left.\frac{dr_{ft}}{d\varepsilon_t}\right| = \frac{1}{2}\lambda^2 exp(a)b(1+\theta) > 0. \tag{34}$$

A4.3 Firms and Investment

Investors expect that the output of firm i follows:

$$Y_{i,t+1} = \exp\left(s_i\varepsilon_{t+1} - \frac{1}{2}s_i^2\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)\right)K_{i,t+1}^{\alpha}$$

where s_i captures how exposed firm i is to the fundamental shock and $K_{i,t+1}$ is the capital stock of the firm.

We assume that capital depreciates fully each period so that firm i's period $t+1$ capital, $K_{i,t+1}$, equals period t investment in firm i, $I_{i,t}$:

$$K_{i,t+1} = I_{i,t}.$$

The marginal return to capital in firm i is

$$R_{i,t+1}^K = \frac{dY_{i,t+1}}{dK_{i,t+1}}, \tag{35}$$

$$= \alpha\exp\left(s_i\varepsilon_{t+1} - \frac{1}{2}s_i^2\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)\right)K_{i,t+1}^{\alpha-1}. \tag{36}$$

The subjective log expected return on firm i capital is

$$\ln\left(\mathbb{E}_t^{\theta}\left[R_{i,t+1}^K\right]\right) = \ln\alpha + (\alpha-1)k_{i,t+1}, \tag{37}$$

where $k_{i,t+1} = \ln(K_{i,t+1})$ is defined as log capital. For an objective observer with $\theta = 0$, the log expected return on firm i capital equals:

$$\ln\left(\mathbb{E}_t\left[R_{i,t+1}^K\right]\right) = \ln\alpha + (\alpha-1)k_{i,t+1}. \tag{38}$$

As in (Campbell, 2017) Chapter 7, firm i's capital stock then satisfies the investment Euler equation under the subjective probability distribution:

$$\begin{aligned}
1 &= \alpha K_{i,t+1}^{\alpha-1}\mathbb{E}_t^{\theta}\left[M_{t+1}\left(s_i\varepsilon_{t+1} - \frac{1}{2}s_i^2\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)\right)\right], && (39)\\
&= \alpha\delta K_{i,t+1}^{1-\alpha}\mathbb{E}_t^{\theta}\left[exp\left(-\left(\lambda-s_i\right)\varepsilon_{t+1} - \frac{1}{2}s_i^2\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)\right)\right], && (40)\\
&= \alpha\delta K_{i,t+1}^{\alpha-1}exp\left(\frac{1}{2}\left(\left(\lambda-s_i\right)^2 - s_i^2\right)\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)\right). && (41)
\end{aligned}$$

Solving for $k_{i,t+1}$ then gives:

$$k_{i,t+1} = \frac{\ln(\alpha\delta)}{1-\alpha} - \frac{\lambda s_i - \frac{1}{2}\lambda^2}{1-\alpha}\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right). \tag{42}$$

Taking the comparative static of $k_{i,t}$ with respect to ε_t in the vicinity of $\varepsilon_t = 0$ and applying the chain rule with (26) gives:

$$\left.\frac{dk_{i,t}}{d\varepsilon_t}\right|_{\varepsilon_t=0} = \frac{\lambda s_i - \frac{1}{2}\lambda^2}{1-\alpha}exp(a)b(1+\theta) \tag{43}$$

Defining aggregate log investment as $k_{H,t} + k_{L,t}$, Proposition 2 parts d) and e) then follow from (43).

Substituting the solution for $k_{i,t+1}$ back into (37) gives the log subjective expected return on capital as:

$$\ln\left(\mathbb{E}_t^{\theta}\left[R_{i,t+1}^K\right]\right) = -\ln\delta + \left(\lambda s_i - \frac{1}{2}\lambda^2\right)\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right), \tag{44}$$

which with equation (32) for the real risk-free rate becomes equation (14) in the main paper:

$$\ln\left(\mathbb{E}_t^{\theta}\left[R_{i,t+1}^K\right]\right) = r_{ft} + \lambda s_i\mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right). \tag{45}$$

We define model PVS_t as the objective expected return differential between low- and high-risk firms

$$PVS_t^{model} = \ln\left(\mathbb{E}_t\left[R_{L,t+1}^K\right]\right) - \ln\left(\mathbb{E}_t\left[R_{H,t+1}^K\right]\right). \tag{46}$$

Substituting the solution for log capital (42) into the objective log expected return (38) then gives equation (17) in the main paper:

$$PVS_t^{model} = \lambda \left(s_L - s_H\right) \mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right). \tag{47}$$

Taking the comparative static of (47) in the vicinity of $\varepsilon_t = 0$ and applying the chain rule with (26) gives Proposition 2 b):

$$\left.\frac{dPVS_t^{model}}{d\varepsilon_t}\right|_{\varepsilon_t=0} = \left.\frac{d\left(\ln\left(\mathbb{E}_t\left[R_{L,t+1}^K\right]\right) - \ln\left(\mathbb{E}_t\left[R_{H,t+1}^K\right]\right)\right)}{d\varepsilon_t}\right|_{\varepsilon_t=0}, \tag{48}$$

$$= \lambda(s_H - s_L)exp(a)b(1+\theta) > 0. \tag{49}$$

From the expressions in Proposition 2, it is clear that all comparative statics increase in magnitude with θ, showing that all effects are amplified if $\theta > 0$.

To show Proposition 3, we note that the difference between the objective and the subjective variance equals:

$$\mathbb{V}_t\left(\varepsilon_{t+1}\right) - \mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right) = \left(1 - \frac{1}{1+\theta\left(1-exp(-b\varepsilon_t)\right)}\right) exp\left(a - b\varepsilon_t\right). \tag{50}$$

We then show Proposition 3 by taking the comparative static of (50) with respect to PVS_t^{model} in the vicinity of $\varepsilon_t = 0$ and applying the chain rule:

$$\left.\frac{d\left(\mathbb{V}_t\left(\varepsilon_{t+1}\right) - \mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)\right)}{dPVS_t^{model}}\right|_{\varepsilon_t=0} = \left(\left.\frac{d\left(\mathbb{V}_t\left(\varepsilon_{t+1}\right) - \mathbb{V}_t^{\theta}\left(\varepsilon_{t+1}\right)\right)}{d\varepsilon_t}\right|_{\varepsilon_t=0}\right) / \left(\left.\frac{dPVS_t^{model}}{d\varepsilon_t}\right|_{\varepsilon_t=0}\right),$$

$$= exp(a)b\theta / \left(\lambda\left(s_H - s_L\right)exp(a)b(1+\theta)\right), \tag{51}$$

$$= \frac{\theta}{1+\theta}\frac{1}{\lambda(s_H - s_L)}. \tag{52}$$

References

Atkeson, A. and L. E. Ohanian (2001). Are Phillips curves useful for forecasting inflation? *Federal Reserve Bank of Minneapolis Quarterly Review 25*(1), 2–11.

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.

Bernanke, B. S. and K. N. Kuttner (2005). What explains the stock market's reaction to federal reserve policy? *Journal of Finance 60*(3), 1221–1257.

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

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica 77*(3), 623–685.

Bordalo, P., N. Gennaioli, and A. Shleifer (2018). Diagnostic expectations and credit cycles. *Journal of Finance 73*(1), 199–227.

Campbell, J. Y. (2017). *Financial decisions and markets: A course in asset pricing*. Princeton University Press.

Campbell, J. Y. and J. Ammer (1993). What moves the stock and bond markets? A variance decomposition for long-term asset returns. *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., M. Lettau, B. G. Malkiel, and Y. Xu (2001). Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *Journal of Finance 56*(1), 1–43.

Campbell, J. Y. and P. Perron (1991). Pitfalls and opportunities: What macroeconomists should know about unit roots. *NBER Macroeconomics Annual 6*(Chapter 3), 141 – 220.

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

Clarida, R., J. Galí, and M. Gertler (1999). The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature 37*(4), 1661–1707.

Cochrane, J. H. and M. Piazzesi (2005). Bond risk premia. *American Economic Review 95*(1), 138–160.

Cohen, R. B., C. Polk, and T. Vuolteenaho (2003). The value spread. *Journal of Finance 58*(2), 609–641.

Coval, J. and E. Stafford (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics 86*(2), 479–512.

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

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.

Gennaioli, N. and A. Shleifer (2018). *A Crisis of Beliefs: Investor Psychology and Financial Fragility*. Princeton University Press.

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

Gorodnichenko, Y. and M. Weber (2016). Are sticky prices costly? Evidence from the stock market. *American Economic Review 106*(1), 165–199.

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.

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 (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics 126*(1), 1–35.

Holston, K., T. Laubach, and J. Williams (2016). Measuring the natural rate of interest: International trends and determinants. Working paper, Federal Reserve Bank of San Francisco.

Jordà, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review 95*(1), 161–182.

Jurado, K., S. C. Ludvigson, and S. Ng (2015, March). Measuring uncertainty. *American Economic Review 105*(3), 1177–1216.

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

Kelly, B. and S. Pruitt (2013). Market expectations in the cross-section of present values. *Journal of Finance 68*(5), 1721–1756.

Keynes, J. M. (1937). The general theory of employment. *Quarterly Journal of Economics 51*(2), 209–223.

Lettau, M. and S. C. Ludvigson (2004). Understanding trend and cycle in asset values: Reevaluating the wealth effect on consumption. *American Economic Review 94*(1), 276–299.

Minsky, H. P. (1977). The financial instability hypothesis: An interpretation of keynes and an alternative to "standard" theory. *Nebraska Journal of Economics and Business 16*(1), 5–16.

Nakamura, E. and J. Steinsson (2018). High frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics 133*(3), 1283–1330.

Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review 94*(4), 1055–1084.

Shumway, T. (1997). The delisting bias in CRSP data. *Journal of Finance 52*(1), 327–340.

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

Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics 54*, 375–421.

Summers, L. H. (2015). Demand side secular stagnation. *American Economic Review 105*(5), 60–65.

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

Vuolteenaho, T. (2002). What drives firm-level stock returns? *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. *Review of Economic Studies 70*(4), 861–886.

APPENDIX FIGURES

Figure A.1: Comparing Filtering Methods for the Real Rate

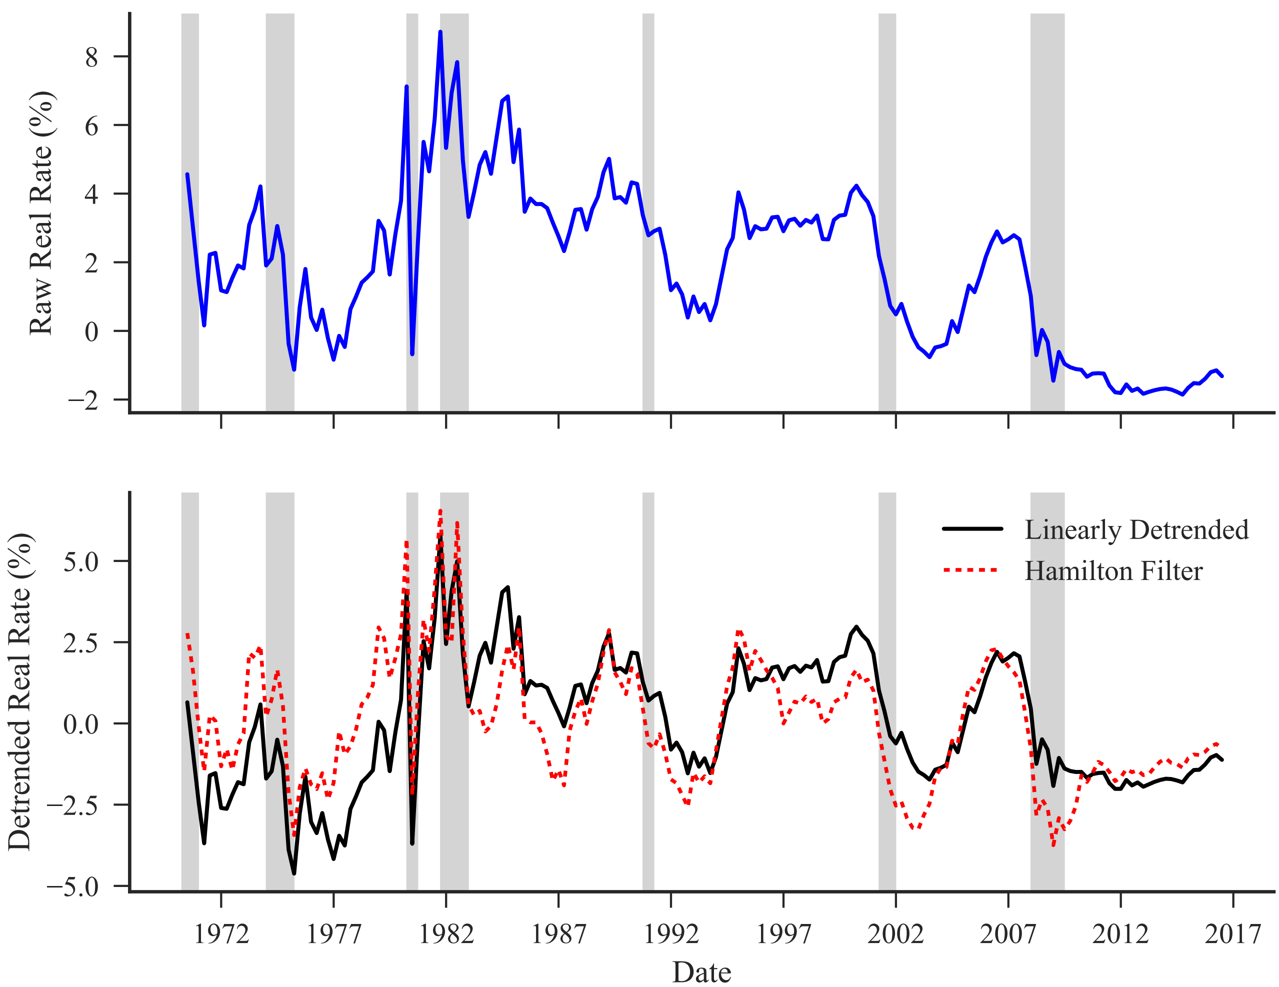

Notes: The top panel of the figure plots the raw one-year real rate. The raw real rate is the one-year Treasury bill rate net of one-year survey expectations of the inflation (the GDP deflator) from the Survey of Professional Forecasters, expressed in percentage terms. The bottom panel of the figure compares two different methods for extracting the cyclical component of the real rate. The first just uses a deterministic time trend. The second uses the methodology of Hamilton (2017), with full details in Section A2.1. Data is quarterly and spans 1970Q2-2016Q2. Shaded bars indicate NBER recessions.

Figure A.2: Simulated t-statistics and R^2

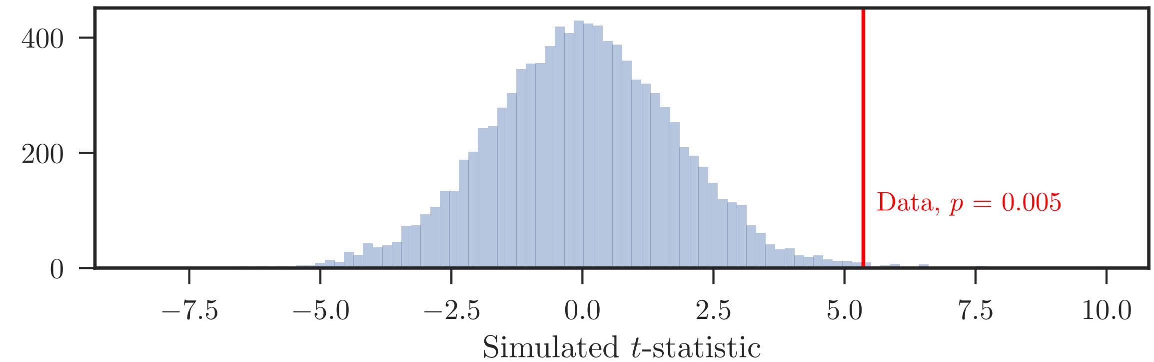

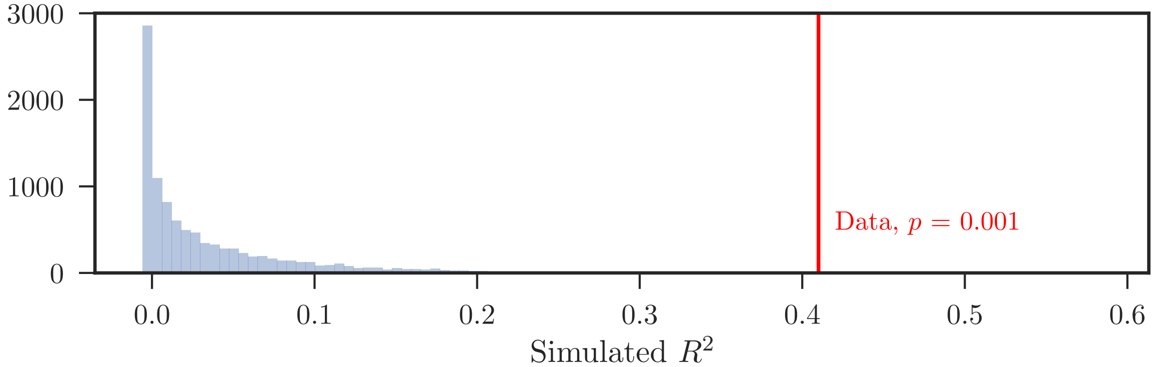

Notes: This figure plots simulated t-statistics and R^2 for a univariate regression of the real rate on PVS. We independently fit AR(1)-GARCH(1,1) processes to each series and simulate each 10,000 times. Within each simulation, we regress the real rate on PVS, saving the Newey-West t-statistic (with five lags) and the R^2. The top panel of the figure shows the distribution of the t-statistics from this procedure and the bottom panel shows the R^2. The red bar shows the actual estimate of each statistic in the data. The p-values listed in the plot are computed as the proportion of simulations that have a t-statistic (or R^2) that exceeds the actual value in the data. 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. See the Section A1 for details on how we construct PVS.

Figure A.3: Impulse Responses to Monetary Policy and PVS Shocks (Traditional VAR)

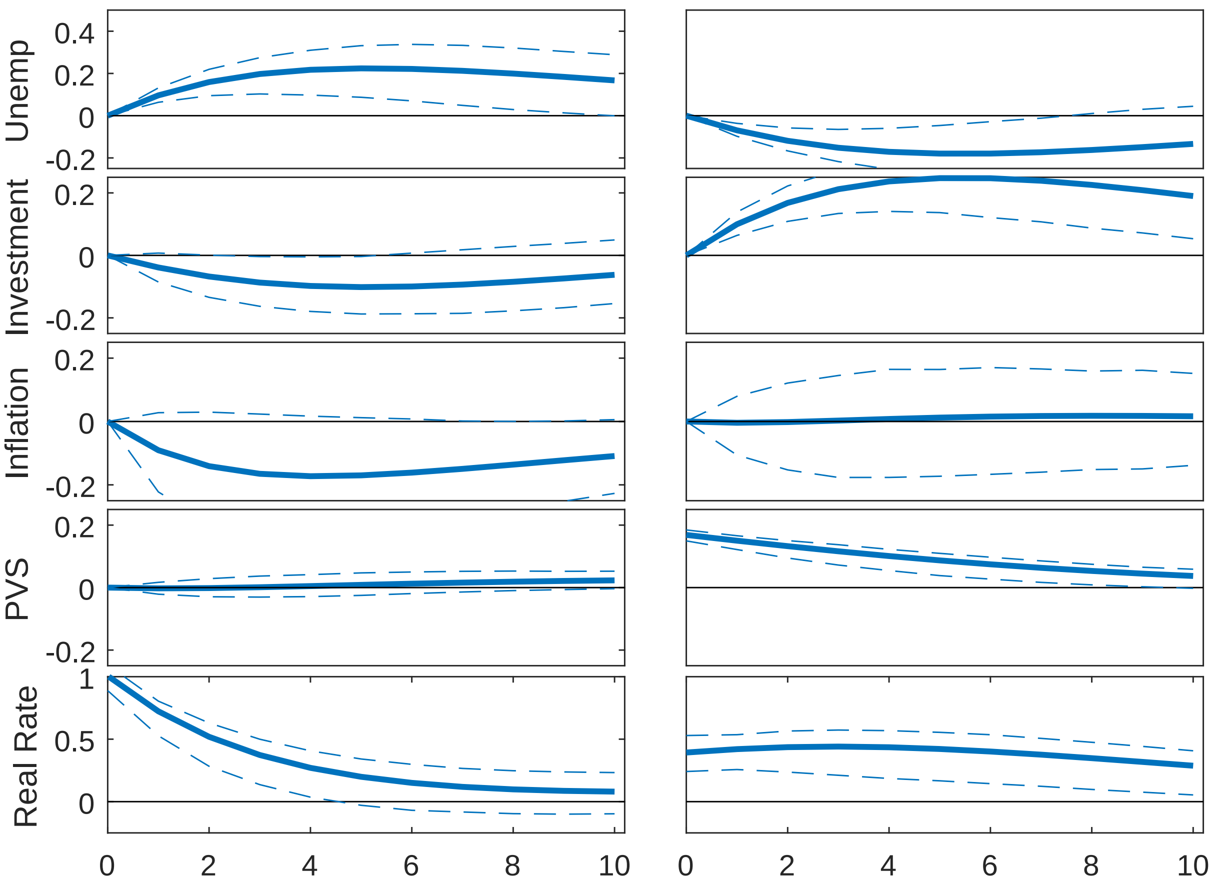

Notes: This figure plots impulse responses to monetary policy shocks (left panel) and PVS shocks (right panel). Impulse responses to one-standard deviation shocks are estimated from a five-variable VAR(1) in unemployment, the investment-capital ratio, inflation, PVS, and the linearly detrended real rate with one lag using quarterly data 1970Q-2016Q2. Unemployment is the civilian unemployment rate (UNRATE). The investment-capital ratio is computed as private nonresidential fixed investment (PNFI) divided by the previous year's current-cost net stock of fixed private nonresidential assets (K1NTOTL1ES000). Following Bernanke and Mihov (1998), structural innovations in the real rate are assumed to affect output, inflation, and precautionary savings demand with a lag. Risk appetite (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.

Figure A.4: Impulse Responses Pre-Crisis

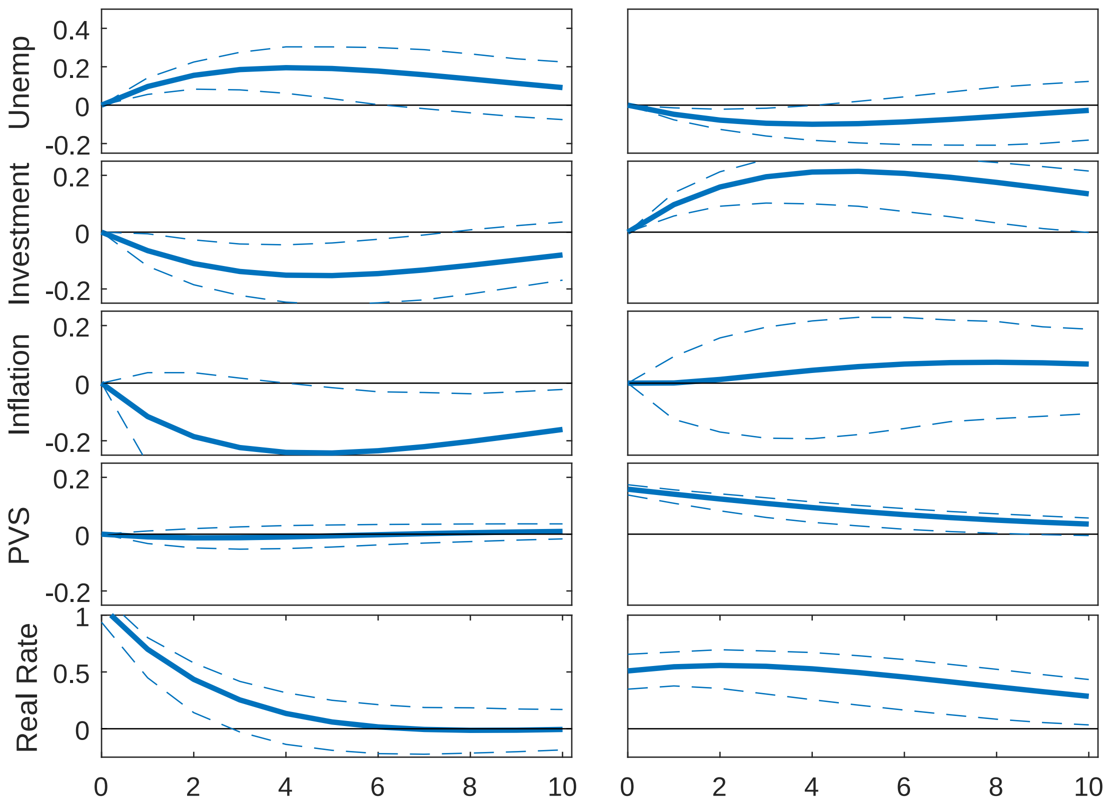

Notes: This figure plots impulse responses to monetary policy shocks (left panel) and *PVS* shocks (right panel). It corresponds to Figure A.3 of this appendix, but uses the pre-crisis sample that ends in 2008Q4.

Figure A.5: Impulse Responses Alternative Ordering

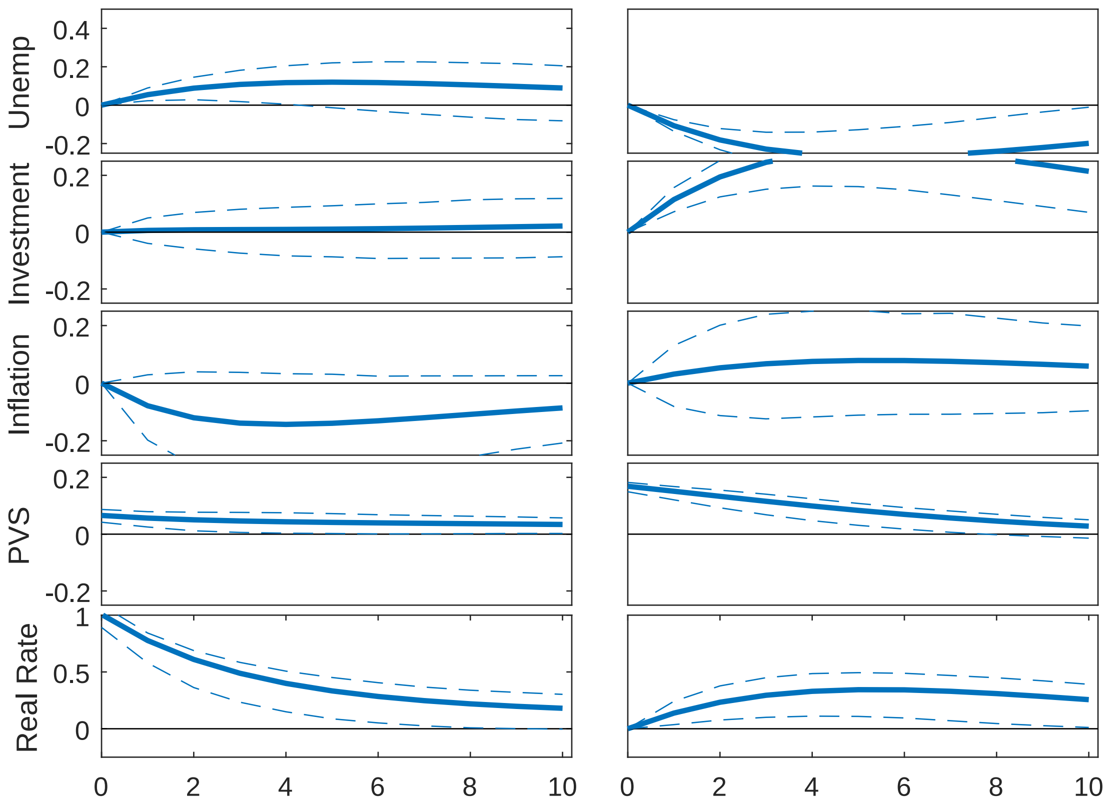

Notes: This figure plots impulse responses to monetary policy shocks (left panel) and PVS shocks (right panel). It differs from Figure A.3 of this appendix in that here we construct impulse responses under the assumption that PVS_t reacts to the real rate immediately, but the real rate reacts to PVS_t with a lag.

Figure A.6: Impulse Responses for VAR(4)

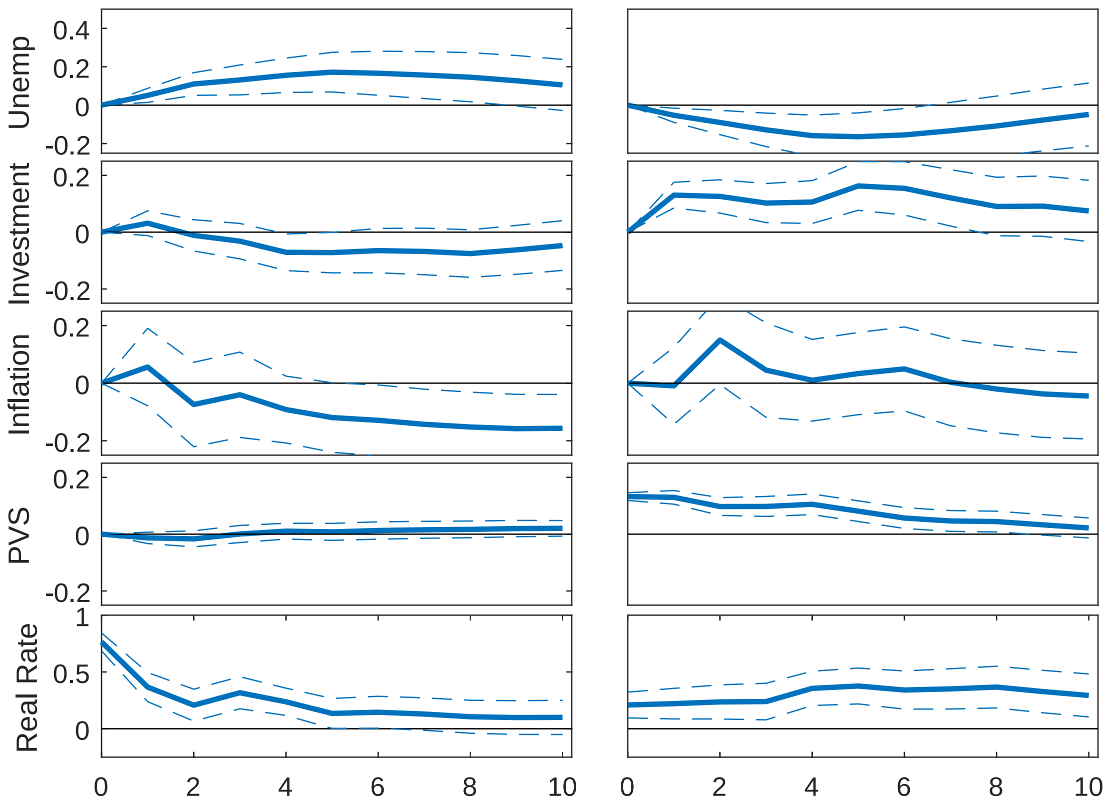

Notes: This figure plots impulse responses to monetary policy shocks (left panel) and *PVS* shocks (right panel). It differs from Figure A.3 of this appendix in that impulse responses are based on a VAR(4) instead of a VAR(1).

Figure A.7: Impulse Responses with Output Gap

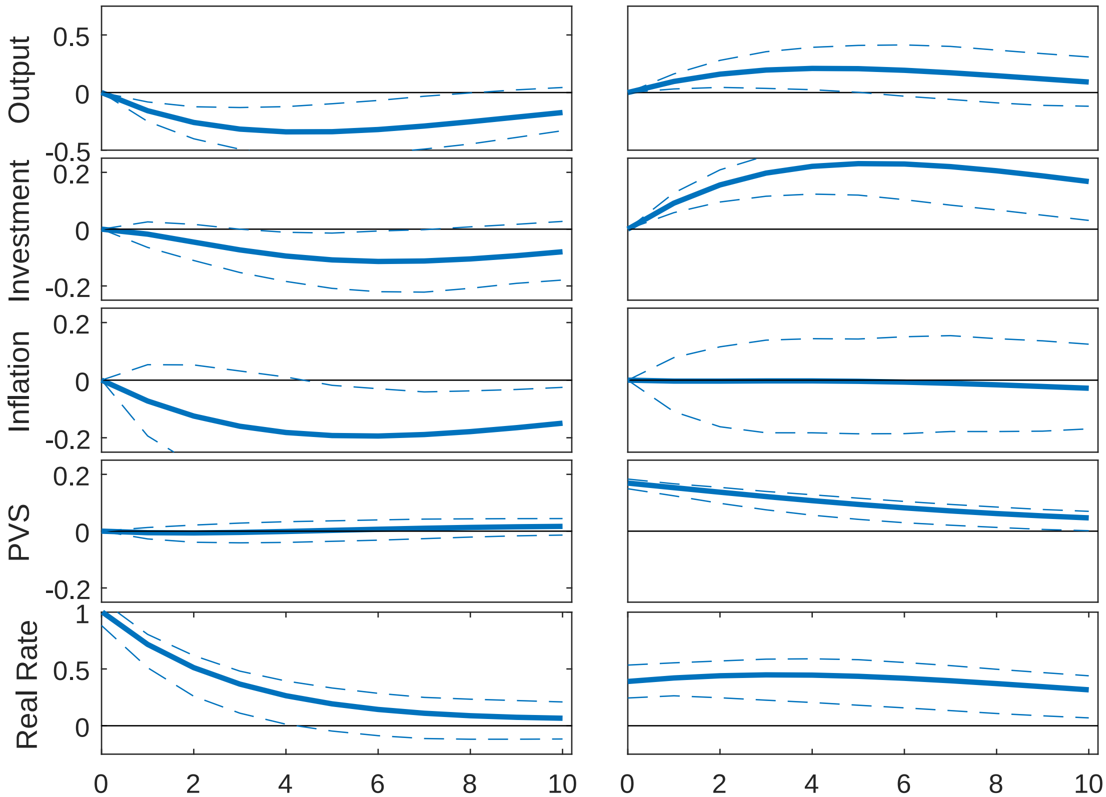

Notes: This figure plots impulse responses to monetary policy shocks (left panel) and *PVS* shocks (right panel). It differs from Figure A.3 of this appendix in that it uses the output gap instead of the unemployment rate.

Figure A.8: Impulse Responses of the Macroeconomy to PVS Shocks (Local Projections)

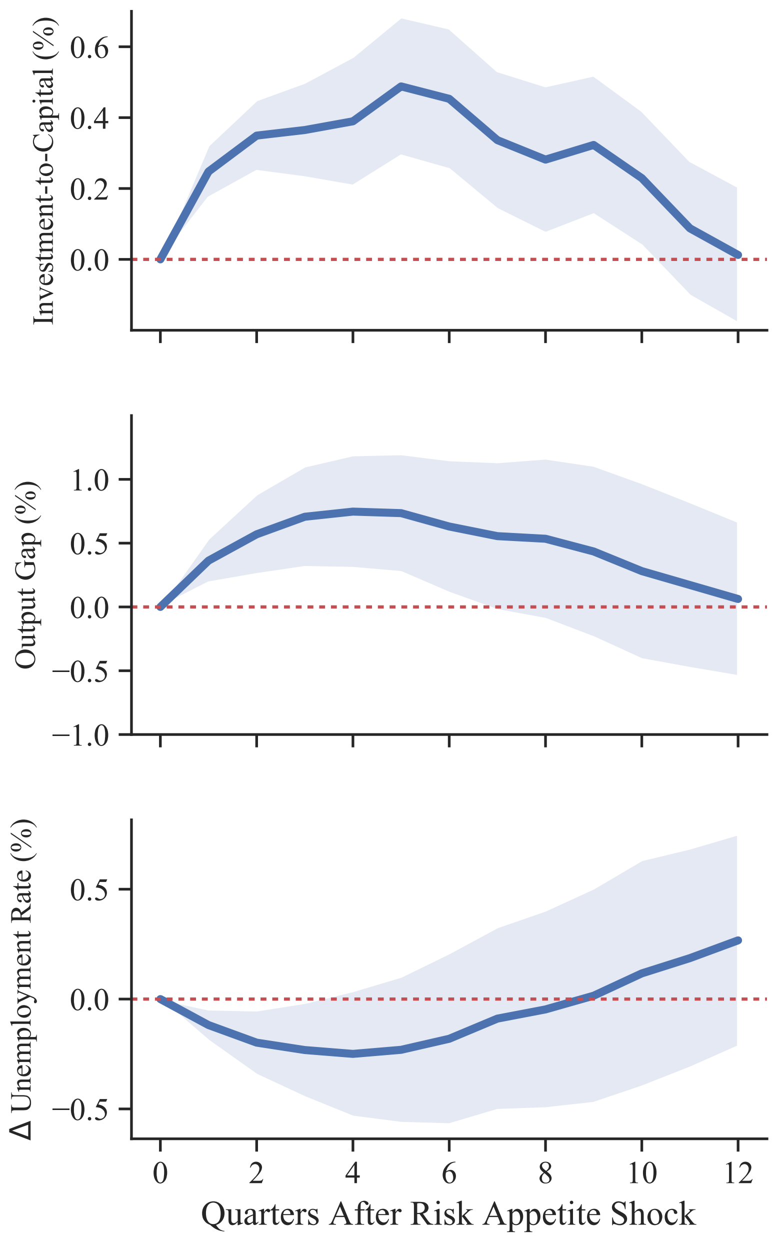

Notes: This figure plots the estimated impulse response (and its associated 95% confidence band) of several macroeconomic variables to a one-standard deviation shock to PVS_t using local projections. We compute impulse responses using Jordà (2005) local projections of each macroeconomic outcomes onto PVS$_t$. In all cases, we run regressions of the following form: $y_{t+h} = a + b^h_{PVS} \times PVS_t + b^h_{RR} \times RealRate_t + b^h_y \times y_t + b^h_{mkt} \times \text{Agg BM}_t + b^h_{cp} \times CP_t + \varepsilon_{t+h}$, where Agg BM$_t$ is the aggregate book-to-market ratio and CP_t is the Cochrane and Piazzesi (2005) bond risk-factor. We consider three different macroeconomic outcomes for the y-variable. The first is the investment-to-capital ratio, defined as the level of real private nonresidential fixed investment (PNFI) divided by the previous year's current-cost net stock of fixed private nonresidential assets (K1NTOTL1ES000). The second is the real output gap, defined as the percent deviation of real GDP from real potential output. The third is the change in the U.S. civilian unemployment rate. When forecasting the investment-capital ratio, y_{t+h} is the level of the investment-capital ratio at time $t+h$. For the output gap, y_{t+h} is the level of the output gap at time $t+h$. Finally, for the unemployment rate, y_{t+h} is the change in the unemployment rate between t and $t+h$, and y_t is the change between $t-1$ and t. All macroeconomic variables come from the St. Louis FRED database and are expressed in percentage points. PVS_t is defined as in the main text. 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. For all regressions, we use Newey-West standard errors with five lags. Data is quarterly and spans 1970Q2-2016Q2.

APPENDIX TABLES

Table A.1: Real Rate Variation (Alternative Filters and Raw Series)

Dependent. Variable:	Hamilton-Filtered Real Rate ($\tilde{r}_t$)						Raw Real Rate		
	Levels			First-Differences			Levels		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PVS_t	1.21**	1.44**	1.36**	0.39**	0.39**	0.32**	1.42**	1.37**	1.30**
	(6.53)	(8.29)	(6.63)	(3.00)	(3.07)	(2.66)	(5.65)	(7.02)	(5.71)
Aggregate BM		0.51**	0.47**		0.02	0.13		0.57**	0.38
		(3.76)	(2.40)		(0.31)	(1.48)		(2.37)	(0.87)
Output Gap			0.25**			0.42**			0.13
			(3.66)			(2.89)			(1.07)
Inflation			0.15*			0.01			0.14
			(1.90)			(0.06)			(0.94)
Constant	0.00	0.00	0.01	-0.02	-0.02	-0.02	1.86**	1.86**	0.93
	(0.01)	(0.01)	(0.03)	(-0.32)	(-0.32)	(-0.39)	(6.56)	(6.90)	(1.48)
Adj. R^2	0.42	0.48	0.58	0.10	0.10	0.18	0.38	0.44	0.46
N	185	185	185	184	184	184	185	185	185

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. Section A1 contains full details on how we compute BM ratios. PVS_t is defined as the difference in BM ratios between the bottom and top quintile portfolios. Aggregate BM is computed by summing book equity values across all firms and divided by the corresponding sum of market equity values. The output gap is the percentage deviation of real GDP from the CBO's estimate of potential real GDP. Inflation is the annualized percentage four-quarter 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. In columns (1)-(6), we follow Hamilton (2017) in extracting the cyclical component of the real rate and use it in the regression. We do the same for inflation, the aggregate BM ratio, and the output gap. See Section A2.1 for more details on the procedure. Columns (7)-(9) use the raw real rate. We do not present first-difference analysis of the raw real rate because this maps directly to the first-difference analysis of the linearly detrended real rate that we present in the main text. Standard errors are computed using Newey-West (1987) with five lags. In each set of regressions, we normalize PVS (or its first difference) to have a mean of zero and a variance of one. * 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 A.2: The Real Rate and the Aggregate Stock Market

Panel A: *Return Forecasting*

	Vol-Sorted $\text{Ret}_{t \to t+1}$		$\text{Mkt-Rf}_{t \to t+1}$	
	(1)	(2)	(3)	(4)
Hamilton-Filtered Real Rate ($\tilde{r}_t$)	1.49**		-0.19	
	(2.56)		(-0.49)	
Raw Real Rate		1.17**		-0.24
		(2.55)		(-0.88)
Adj. R^2	0.03	0.03	-0.00	-0.00
N	184	184	184	184

Panel B*: Aggregate Earnings and Dividend Growth Forecasting*

Dep. Variable:	$g^E_{t,t+1}$		$g^E_{t,t+4}$		$g^D_{t,t+1}$		$g^D_{t,t+4}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\tilde{r}_t$	-4.32		-11.45		0.10		-0.06	
	(-0.68)		(-1.51)		(0.19)		(-0.10)	
R_t		-3.68		-7.45		-0.64		-0.75
		(-0.90)		(-1.47)		(-1.64)		(-1.56)
Adj R^2	0.00	0.00	0.06	0.04	-0.00	0.04	-0.01	0.07
N	184	184	181	181	184	184	181	181

Notes: **Panel A** of this table uses the one-year real interest rate to forecast returns on either the low-minus-high volatility equity portfolio or the excess returns on the aggregate stock market. 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. Volatility-sorted returns are returns on the lowest minus highest volatility quintile portfolios. Vol-Sorted Ret in the forecasting regression corresponds to returns on this low-minus-high volatility portfolio. When forecasting the aggregate stock market, we use the excess return of the CRSP Value-Weighted index obtained from Ken French's website. For quarterly regressions, standard errors are computed using Newey-West (1987) with two lags. **Panel B** of the table reports forecasting regressions of real aggregate earnings growth (g^E) or real aggregate dividend growth (g^D) using the the one-year real rate. Real earnings and real dividends come from Robert Shiller'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. In the table, $\tilde{r}_t$ is the cyclical component of the real rate, extracted using Hamilton (2017). See Section A2.1 for more details on the procedure. R_t is simply the raw real rate. Standard errors are computed using Newey-West (1987) with two lags for quarterly regressions and five lags for annual. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. For both panels, all regressions have a constant, but we omit the estimates to save space. Data is quarterly and spans 1970Q2-2016Q2. Growth rates and returns and expressed in percentage terms.

Table A.3: Subsample Analysis of PVS_t and the Real Rate

Dep. Variable:	One-Year Real Rate							
	Levels				First-Differences			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PVS	1.27**	1.09**	1.05**	0.53**	0.39**	0.24**	0.29**	0.24**
	(5.36)	(3.85)	(5.23)	(4.76)	(2.73)	(2.39)	(2.92)	(2.66)
Subsample	Main	Main, Ex. Volcker	Long	Pre-1977	Main Sample	Main, Ex. Volcker	Long	Pre-1977
N	185	145	253	95	184	143	252	94
Adj. R^2	0.41	0.31	0.26	0.21	0.13	0.13	0.08	0.11

Notes: This table reports regression estimates of the one-year real rate on the contemporaneous spread in book-to-market (BM) 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 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. 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. For the columns listed as "Main Sample" and "Ex. Volcker", we linearly detrend the one-year real rate using the sample 1970Q2-2016Q2. t-statistics are listed below each point estimate in parentheses and are computed using Newey-West (1987) standard errors with five lags. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. In the table, PVS_t is standardized to have mean zero and variance one. This is true in both the levels regression and the first-differenced regressions. Data is quarterly and the subsamples are as follows: (i) "Main" is from 1970Q2-2016Q2 and corresponds to the sample used in the main text; (ii) "Main, Ex. Volcker" is 1970Q2-1976Q4 and 1987Q4-2016Q2. It is the main example but excludes the period 1977Q1-1986Q4; (iii) "Long" is the period 1953Q2-2016Q2. To extend PVS back to 1953, we use the book equity data from Davis, Fama, and French (2000). In addition, to compute the one-year real rate prior to 1970Q2, we take the 1-year nominal rate minus the four-quarter moving average of inflation; and (iv) "Pre-1977" is 1953Q2-1976Q4.

Table A.4: Decomposition of the Real Interest Rate and PVS_t

	One-Year Real Rate Decomposition	Regression on PVS b	$t(b)$	Adj. R^2
(1)	Baseline Detrended Real Rate	1.27	5.36	0.41
(2)	Baseline Raw Real Rate	1.42	5.65	0.38
(3)	Nominal 1-Year Rate	1.91	3.22	0.27
(4)	Expected Inflation	0.49	1.19	0.06
(5)	Fixed Taylor Rule Implied Rate (Taylor, 1993)	0.32	0.97	0.03
(6)	Residual	1.10	2.84	0.23
(7)	Fitted Taylor Rule Implied Rate	0.16	0.82	0.03
(8)	Residual	1.27	4.66	0.35

Notes: This table reports univariate regressions of several variables on PVS. Section A1 of the internet 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. Standard errors are computed using Newey-West (1987) with five lags. Data is quarterly and the full sample spans 1970Q2-2016Q2. In all cases, PVS_t is standardized to have a mean of zero and a variance of one.

Table A.5: The Real Rate and Valuation of Other Characteristic-Sorted Portfolios

		Real Rate$_t = a + b \times X_t + \varepsilon_t$					
		Levels			First-Differences		
		b	$t(b)$	Adj. R^2	b	$t(b)$	Adj. R^2
Univariate:							
(1)	Duration	**-0.69**	-2.65	0.12	*-0.22*	-1.64	-0.01
(2)	Leverage	**0.54**	2.20	0.07	**0.17**	2.00	0.02
(3)	Beta	**1.20**	5.81	0.37	0.13	1.56	0.01
(4)	LR Beta	**1.12**	6.18	0.32	*0.13*	1.83	0.01
(5)	2M-Beta	0.35	1.40	0.03	**0.36**	2.56	0.11
(6)	CF Beta	-0.02	-0.08	-0.01	-0.04	-0.64	-0.00
(7)	Size	**-1.12**	-5.75	0.32	**-0.25**	-2.25	0.05
(8)	Value	**0.69**	3.33	0.12	**0.18**	2.02	0.02
Kitchen-Sink:							
(9)	PVS	**2.18**	4.09	0.60	**0.45**	2.80	0.17

Notes: This table reports regression estimates of the one-year real rate on the book-to-market spreads of portfolios formed on various sorting characteristics. Rows (1)-(8) run the following regression, in both levels and first-differences: Real Rate$_t = a + b \times$ Y-Sorted BM Spread$_t + \varepsilon_t$, where Y-Sorted BM Spread$_t$ is the spread in book-to-market ratios between stocks sorted on characteristic Y. Our main variable of interest in the study is the spread in book-to-market 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. PVS_t is defined as the difference in BM ratios between the bottom and top quintile portfolios. We form book-to-market spreads in the same fashion for other sorting variables. The sorting variables we use are: (1) Duration (Weber (2016)); (2) Leverage, measured as long-term debt from COMPUSTAT divided by market equity; (3) CAPM Beta, measured using monthly data over rolling 5 year windows; (4) Long-Run (LR) CAPM Beta, measured using semi-annual data over a rolling ten year window; (5) 2M-Beta, computed at the end of each quarter using the previous sixty days of daily returns; (6) Cashflow (CF) Beta, which is measured by regressing EBITDA growth on national income growth; (7) market capitalization; and (8) book-to-market ratios themselves (value). Spreads are always between the high quintile and the low quintile of the sorting variable. In Row (9), we run a kitchen-sink regression of the real rate on PVS_t plus all of the book-to-market spreads in rows (1)-(8) and report the estimated coefficient on PVS_t. 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. In all cases, we normalize all book-to-market spreads (or their first-difference) to have mean zero and variance one. Standard errors are computed using both Newey-West (1987) with five lags. Italicized point estimates indicates a p-value of less than 0.1 and bold point estimates indicate a p-value of less than 0.05. Data is quarterly and spans 1970Q2-2016Q2.

Table A.6: The Real Rate and Double-Sorted Versions of PVS

		Real Rate$_t = a + b \times$ Y-Neutral PVS$_t + \varepsilon_t$					
		Levels			First-Differences		
	Characteristic Y	b	$t(b)$	Adj. R^2	b	$t(b)$	Adj. R^2
(1)	Duration	**0.81**	4.02	0.16	**0.36**	2.91	0.10
(2)	Leverage	**1.16**	5.27	0.35	**0.38**	2.95	0.12
(3)	2M-Beta	**1.30**	5.79	0.43	**0.24**	2.52	0.04
(4)	Size	**1.23**	5.04	0.39	**0.38**	2.76	0.12
(5)	Value	**1.10**	5.12	0.31	**0.34**	2.50	0.09
(6)	Industry-Adjusted	**1.15**	5.42	0.34	**0.29**	2.58	0.06
(7)	Div. Payers	**1.22**	6.86	0.38	**0.27**	2.93	0.06
(8)	Non-Div. Payers	**0.63**	2.46	0.10	**0.31**	2.42	0.08

Notes: This table reports a battery of robustness exercises for our main results. Specifically, we report time-series regression results of the following form, in both levels and first-differences: Real Rate$_t = a + b \times$ Y-Neutral PVS$_t + \varepsilon_t$. For rows (1)-(5), the variableY-Neutral PVS$_t$ is constructed by sorting all NYSE, AMEX, and NASDAQ firms in CRSP into two bins based on the median value of characteristic Y at time t. Within the high-Y (above median) firms, we further sort firms into terciles based on their volatility over the previous sixty days. Within each tercile, we compute the average book-to-market (BM) ratio between the low and high-volatility firms. We repeat this procedure for firms in the low-Y bucket. Y-Neutral PVS$_t$ is then defined as (BM$_t$ of Low-Volatility − BM$_t$ of High-Volatility with High Y)/2 + (BM$_t$ of Low-Volatility − BM$_t$ of High-Volatility with Low Y)/2. In row (6), we compute an industry-adjusted version of PVS_t by first sorting stocks into industries based on their SIC code and the 48 industry definitions on Ken French's website. Within each industry i we sort firms into quintiles based on their trailing 60-day volatility and then define $PVS_{i,t}$ as the average BM ratio of low-volatility firms in industry i minus the average BM ratio of high-volatility firms in industry i. The industry-adjusted PVS_t is defined as the equal weighted $PVS_{i,t}$ across all 48 industries. In row (7), we construct PVS_t only for the set of firms who have paid a dividend over the past twenty-four months. Row (8) repeats the exercise for the set of firms that have not paid a dividend over the past twenty-four months. See Section A2.6.2 of this appendix for more details on how we construct each of these versions of PVS_t. In all cases, we standardize PVS_t (or its first difference) to have a mean of zero and a variance of one. 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. Italicized point estimates indicates a p-value of less than 0.1 and bold point estimates indicate a p-value of less than 0.05. Data is quarterly and the full sample spans 1970Q2-2016Q2.

Table A.7: Objective Measures of Risk

	X-variable	N	$PVS_t = a + b \times X_t$			$\text{RealRate}_t = a + c \times X_t$		
			b	$t(b)$	R^2	c	$t(c)$	R^2
(1)	σ_t (Real GDP Growth$_{t+1}$)	184	0.01	0.07	-0.01	-0.27	-0.72	0.01
(2)	σ_t (Real Consumption Growth$_{t+1}$)	184	0.20	1.07	0.03	0.12	0.32	-0.00
(3)	σ_t (Industrial Production Growth$_{t+1}$)	184	-0.10	-0.43	0.00	-0.46	-1.20	0.05
(4)	Macro Uncertainty Index$_t$	185	0.07	0.31	-0.00	0.35	1.15	0.03
(5)	σ_t (Aggregate Stock Market Ret$_{t+1}$)	185	-0.23	-1.48	0.05	-0.04	-0.20	-0.00
(6)	Raw $\mathbb{E}_t\left[\overline{\sigma}_{H,t+1}\right] - \mathbb{E}_t\left[\overline{\sigma}_{L,t+1}\right]$	184	-0.31	-2.13	0.09	0.32	1.15	0.02
(7)	Detrended $\mathbb{E}_t\left[\overline{\sigma}_{H,t+1}\right] - \mathbb{E}_t\left[\overline{\sigma}_{L,t+1}\right]$	184	-0.41	-2.60	0.16	0.18	0.76	0.00

Notes: This table compares objective expectations of macroeconomic risk with PVS$_t$ and the one-year real interest rate. We define the objective expectation of variable X as the expectation that comes from a statistical forecasting model. Specifically, we fit ARMA(1,1)-GARCH(1,1) models to real GDP growth (Row 1), real consumption growth (Row 2), and industrial production growth (Row 3). σ_{t+1} in the table then represents the GARCH forecast for the volatility of each growth series, which is in the information set at time t. Row (4) uses the macroeconomic uncertainty index from Jurado et al. (2015). σ_{t+1} (Agg. Stock Market) in Row (5) is the expected volatility of the aggregate stock market. From 1986 onward, it is the time t value of the VXO option implied volatility index from the CBOE. To fill in the data prior to 1986, we fit an AR(1) model to within-quarter realized volatility of the CRSP Value-Weighted index. We then use the one-step ahead forecast made at time t from the AR(1) model, reindexed to create a smooth series when appending the VXO after 1986. In row (6), $\overline{\sigma}_H - \overline{\sigma}_L$ is the average volatility of high-volatility stocks minus the average volatility of low-volatility stocks. At time t,we compute our forecast of $\overline{\sigma}_{H,t+1} - \overline{\sigma}_{L,t+1}$ using an AR(1) model fit to the whole sample. This row corresponds to column (4) in Table 5 of the main text. In row (7), we repeat the construction in row (6), but first use the methods in Hamilton (2017) to extract the cyclical component of $\overline{\sigma}_H - \overline{\sigma}_L$ before fitting our forecasting model. The first set of regressions in the table shows the results of a univariate regression of PVS$_t$ on each of these expected risk measures. The second set of regressions in the table shows the results of a univariate regression of the real rate on contemporaneous values of each variable. 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. In all regressions, we standardize both PVS and the expected risk measures to have mean zero and variance one in order to facilitate comparison of magnitudes. The listed t-statistics are computed using Newey-West (1987) standard errors with five lags. Data is quarterly and the full sample spans 1970Q2-2016Q2.

Table A.8: The Real Rate and Mutual Fund Flows

***Panel A**: Summary Statistics*

	Mean	Std. Dev.	p25	p50	p75	Min	Max	# Funds
Quarterly Obs./Fund	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, t-statistics are listed below point estimates in parentheses and are computed using Driscoll-Kraay (1998) standard errors with five lags within each fund cluster. * indicates a p-value of less than 0.1 and ** indicates a p-value of less than 0.05. Quarterly mutual fund data derives from CRSP and spans 1973Q2-2015Q3. Returns are in percentage terms.

Table A.9: Volatility-Sorted Returns and Monetary Policy Surprises

$$\text{Vol-Sorted Ret}_{t\to t+1} = a + b \times \text{MP Shock}_{t\to t+1} + \varepsilon_{t\to t+1}$$

	Quarterly Data				Daily Data				Sample	
	All		Scheduled		All		Scheduled			
MP Shock	b	$t(b)$	b	$t(b)$	b	$t(b)$	b	$t(b)$	Start	End
Romer and Romer (2004)	0.75	0.47	0.71	0.44	-0.04	-0.30	-0.06	-0.43	1970.Q1	1996.Q4
Bernanke and Kuttner (2005)	-2.94	-0.19	-1.65	-0.07	5.55	1.32	-1.08	-0.49	1989.Q2	2008.Q2
Gorodnichenko and Weber (2016)	-1.14	-0.06	1.60	0.03	13.34	1.89	3.67	0.94	1994.Q1	2009.Q4
Nakamura and Steinsson (2018)	1.46	0.06	12.83	0.20	18.74	1.98	5.29	1.03	1995.Q1	2014.Q1

Notes: This table reports regressions of volatility-sorted returns onto monetary policy shocks. 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. Volatility-sorted returns are returns on the lowest minus highest volatility quintile portfolios. Quarterly return regressions aggregate daily monetary policy shocks by summing over all shocks within a quarter. The Romer and Romer (2004) shock is the change in the intended Federal Funds rate inferred from narrative records around monetary policy meetings, after controlling for changes in the Federal Reserve's information. The Bernanke and Kuttner (2005) shock is derived from the price change in Federal Funds future contracts relative to the day before the policy action. The Gorodnichenko and Weber (2016) shock is derived from the price change in Federal Funds futures from 10 minutes before to 20 minutes after an FOMC press release. The Nakamura and Steinsson (2018) shock is the unanticipated change in the first principal component of interest rates with maturity up to one year from 10 minutes before to 20 minutes after an FOMC news announcement. Columns listed as "All" include all policy changes and "Scheduled" includes only changes that occurred at regularly scheduled policy meetings. In restricting the analysis to regularly scheduled meetings, we exclude quarters after 1993Q4 where the Federal Reserve made policy changes outside of scheduled meetings. Prior to 1994, policy changes were not announced after meetings so the distinction between scheduled and unscheduled meetings is not material. The listed t-statistics are computed using Davidson and MacKinnon (1993) standard errors for heteroskedasticity in small samples.