

Journal Pre-proof

Time-Varying Inflation Risk and Stock Returns

Martijn Boons, Fernando Duarte, Frans de Roon,
Marta Szymanowska

PII: S0304-405X(19)30242-9
DOI: <https://doi.org/10.1016/j.jfineco.2019.09.012>
Reference: FINEC 3125

To appear in: *Journal of Financial Economics*

Received date: 14 November 2018
Revised date: 20 March 2019
Accepted date: 29 March 2019

Please cite this article as: Martijn Boons, Fernando Duarte, Frans de Roon, Marta Szymanowska, Time-Varying Inflation Risk and Stock Returns, *Journal of Financial Economics* (2019), doi: <https://doi.org/10.1016/j.jfineco.2019.09.012>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2019 Published by Elsevier B.V.



Time-Varying Inflation Risk and Stock Returns [☆]Martijn Boons^{a,*}, Fernando Duarte^b, Frans de Roon^c, Marta Szymanowska^d^a*Nova School of Business and Economics, Rua Holanda 1, 2775-405 Carcavelos, Portugal*^b*Federal Reserve Bank of New York, 33 Liberty Street, New York, NY 10045, USA*^c*Department of Finance, Tilburg School of Economics and Management, 5000 LE Tilburg, the Netherlands*^d*Department of Finance, Rotterdam School of Management, Erasmus University, 3000 DR Rotterdam, the Netherlands*

Abstract

We show that inflation risk is priced in stock returns and that inflation risk premia in the cross-section and the aggregate market vary over time, even changing sign as in the early 2000s. This time variation is due to both price and quantities of inflation risk changing over time. Using a consumption-based asset pricing model, we argue that inflation risk is priced because inflation predicts real consumption growth. The historical changes in this predictability and in stocks' inflation betas can account for the size, variability, predictability, and sign reversals in inflation risk premia.

JEL classification: G11, G12*Keywords:* Inflation, Time-varying inflation risk premium, Inflation hedging, Individual stock returns, Cross-sectional asset-pricing, Nominal-real covariance

[☆]We thank an anonymous referee, Tobias Adrian, Hengjie Ai, Christine Breiner, Ricardo Caballero, Hui Chen, Richard Crump, Frank de Jong, Joost Driessen, Maya Eden, Rik Frehen, Xavier Gabaix, Jonathan Goldberg, Leonid Kogan, Jennifer La'O, Guido Lorenzoni, Gustavo Manso, Matt Notowidigdo, Sahar Parsa, Pedro Santa-Clara, Piet Sercu, Alp Simsek, Ken Singleton, Dongho Song, Laura Spierdijk, Adrien Verdelhan, Luis Viceira, and Ivo Welch; conference participants at the EFA annual meeting 2018, Auckland Finance Meeting 2013, Dauphine-Amundi Annual Workshop 2013, ESSFM Gerzensee 2012, LUBRAFIN 2014, and Netspar International Pension Workshop 2012; and seminar participants at Rotterdam School of Management, Tilburg University, University of Auckland, MIT, MIT Sloan School of Management, Duke University's Fuqua School of Business, Tepper School of Business, Smith School of Business, Federal Reserve Board, and the Federal Reserve Bank of New York. This paper has received financial support from the Dauphine Chair in Asset Management, an initiative of Amundi and the University Paris-Dauphine, under the aegis of the Dauphine Foundation. The views expressed in this paper are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. This paper replaces and supersedes two papers formerly circulated as "The Stock Market Price of Inflation Risk and Its Variation over Time" by Martijn Boons, Frans de Roon, and Marta Szymanowska; and, "Inflation Risk and the Cross-Section of Stock Returns" by Fernando Duarte.

*Corresponding author. Tel.: +351 213821659.

Email addresses: martijn.boons@novasbe.pt (Martijn Boons), fernando.duarte@ny.frb.org (Fernando Duarte), f.a.deroon@uvt.nl (Frans de Roon), mszymanowska@rsm.nl (Marta Szymanowska)

1. Introduction

We show that inflation risk is priced in stock returns and that both the price and quantities of inflation risk are strongly time-varying, even changing sign over our sample from 1962 to 2014. We argue that inflation is risky because it predicts real consumption growth in a time-varying way. Positive shocks to inflation sometimes contain bad news for future consumption, whereas at other times they contain good news. We provide new empirical evidence, for both the cross-section of stock returns and the aggregate stock market, and develop an equilibrium model to argue that it is the time-variation in the predictive content of inflation that drives the observed time-variation in the price and quantities of inflation risk.

To quantify the predictive content of inflation at a given point in time, we define the *nominal-real covariance* as the slope coefficient in a conditional (that is, rolling) regression of real consumption growth on lagged inflation. We find that this nominal-real covariance is an economically strong and statistically significant predictor of the inflation risk premium in the cross-section of stocks. We measure the quantity of inflation risk of stocks by their conditional inflation beta, that is, the slope coefficient in a rolling regression of returns on inflation shocks. A one standard deviation increase in the nominal-real covariance predicts an increase in annualized expected return of 5.3% for a high-minus-low portfolio constructed by sorting stocks in deciles according to their inflation beta. Given an unconditional average return for the high-minus-low inflation beta portfolio of -4.2%, this increase implies that the inflation risk premium can switch sign.

To understand why, consider periods when the nominal-real covariance is negative, such as during the 1970s. In such periods, higher inflation is a bad state of nature as it predicts lower future consumption growth. As a result, high inflation beta stocks are attractive to hedge against bad states and therefore command lower premia. In contrast, in periods when the nominal-real covariance is positive, such as since the early 2000s, a positive shock to inflation contains good news for future consumption. Now, stocks with lower inflation betas are more attractive to hedge and the inflation risk premium switches sign.

We provide further evidence for this explanation by running pooled predictive regressions for the returns of ten inflation beta-sorted portfolios on their conditional inflation betas, the nominal-real covariance, and an interaction term. We find that the interaction term predicts returns well, whereas the other two (non-interaction) terms do not. Thus, understanding the cross-section of inflation risk premia requires knowledge not only of the relative magnitude of inflation betas across stocks, but also of the nominal-real covariance at each point in time.

To understand how variation in quantities and price of inflation risk interact to determine inflation risk premia, we decompose conditional inflation betas into portfolio-specific and time-specific components in our pooled regressions. We present three results that are new to the literature. First, the cross-sectional distribution of inflation risk premia widens and compresses over time due to variation in the price of inflation risk. This result follows from the fact that the interaction between the portfolio-specific component of inflation

betas and the nominal-real covariance fully captures the time-series predictability of the high-minus-low inflation beta portfolio return. Time-variation in the relative quantities of inflation risk (inflation betas) in the cross-section does not add any information.

Second, quantities of inflation risk vary over time with the nominal-real covariance. This time-variation is largely common across stocks. Pairwise correlations between the inflation betas of any two decile portfolios and between any decile portfolio and the aggregate stock market are well above 90%.

Third, the aggregate stock market risk premium contains an inflation component that is time-varying, due to time-variation in both the price and quantity of inflation risk in the aggregate market. This conclusion follows from the fact that the interaction between the inflation beta of the market portfolio and the nominal-real covariance predicts market returns, whereas neither the nominal-real covariance nor the market's inflation beta predict returns in isolation. The joint effect of the price and quantity of inflation risk is large: when the inflation beta of the market portfolio and the nominal-real covariance are both one standard deviation below their mean, the market risk premium equals 13.6%, more than double its unconditional mean. The intuition is that the aggregate stock market is riskier when it comoves negatively with inflation shocks that contain bad news for future consumption growth.

To provide a theoretical foundation for this empirical evidence, we develop a consumption-based equilibrium model. Our model takes the joint stochastic processes for consumption, inflation, and the nominal-real covariance as given, while asset prices are determined endogenously through the Euler equation of a representative agent. There are three key ingredients in our model. The first ingredient is that the representative agent has Epstein-Zin-Weil (EZ) utility. The second ingredient is that inflation is persistent.¹ The third ingredient is that, as in the data, inflation shocks predict future real consumption growth in a time-varying way that allows for sign reversals.² It is this third ingredient that is new to the literature.

Because EZ preferences are not time-separable, marginal utility depends not only on current consumption growth, as in the classic consumption-CAPM, but also on the entire future path of consumption. Since inflation shocks predict future consumption growth, inflation shocks represent shocks to marginal utility and are therefore priced. When the nominal-real covariance is positive, the price of inflation risk is also positive, because higher inflation today represents a good state of nature as future consumption growth is expected to be higher. As the nominal-real covariance changes sign over time, so does the price of inflation risk. To study the cross-section of stock returns, we model dividends as being exogenous and exposed to the same sources of risk as consumption. Because inflation shocks predict the future path of dividends with sign and magnitude depending on the nominal-real covariance, the endogenous quantities of inflation risk

¹Inflation persistence is widely documented in the literature (see, e.g., Fuhrer and Moore, 1995; Stock and Watson, 2005; Campbell and Viceira, 2002; Ang, Bekaert, and Wei, 2007).

²Although interesting, our aim is not to explain the source of this predictability. Rather, our aim is to explain the asset pricing implications from exogenous variation in the predictive relation between inflation and consumption growth. See Duffee (2018) for a survey and the challenges in explaining the time-variation in the nominal-real covariance.

of stocks are also functions of the nominal-real covariance. Our proposed calibration shows that the model can quantitatively reproduce the empirical results discussed above, including the size and time-variation of the price and quantities of inflation risk in the cross-section of stocks, and their relation to the nominal-real covariance. Inflation persistence is key in allowing us to quantitatively match the relatively high inflation risk premium observed in the data, as a more persistent inflation has a longer-lasting effect on the future path of expected consumption and hence on marginal utility.

Empirically, we find that the cross-section of inflation betas is wide, which is instrumental in identifying the inflation risk premium. Our evidence suggests that the cross-sectional variation in inflation betas is mainly driven by variation in inflation exposure in cash flows rather than in discount rates. There is heterogeneity in inflation betas both across and within industries. Industry affiliation is an important determinant of inflation risk unconditionally. On average, the best inflation hedgers are oil, gold, and mining stocks. However, inflation risk is also priced in a time-varying way within industries, even outside the best inflation-hedging industries. This finding is new to the literature and motivates us to abstract from industry effects in our model.

Our empirical results are robust. We consider alternative measures of the nominal-real covariance, using either the covariance between inflation and industrial production or the stock market beta of a long-term bond. We analyze alternative ways to measure inflation risk premia, estimating a maximum correlation mimicking portfolio (Breedon, Gibbons, and Litzenberger, 1989) and a Fama and MacBeth (1973) cross-sectional regression for individual stocks. We use different methods to estimate inflation risk, for instance, by filtering inflation shocks using first-differences or an AR(1)-process instead of an ARMA(1,1)-process, which is our main specification. We also consider a truly out-of-sample exercise using real-time vintage inflation, similar to Ang, Brière, and Signori (2012). We find that our inflation beta-sorted portfolios are not only exposed to shocks in realized inflation but also to changes in breakeven inflation from Treasury Inflation Protected Securities (TIPS). Finally, our results carry through when we control for benchmark predictors or exposure to benchmark asset pricing factors.³

We contribute to the empirical literature on inflation by providing a rich set of new facts concerning inflation risk premia in the stock market. Measures of the inflation risk premium have had a natural starting point in the yield curve. With the development of sophisticated no-arbitrage term structure models and the emergence of inflation-linked bonds, estimates of the inflation risk premium in the bond market have become more reliable and widely available.⁴ Another conventional way to estimate the inflation risk premium is to

³The benchmark predictors are the dividend yield, term spread, default spread, and the consumption-wealth ratio of Lettau and Ludvigson (2001a,b). The benchmark factors are MKT, SMB, HML, WML, RMW, and CMA from the CAPM (Sharpe, 1964; Lintner, 1965; Mossin, 1966) and the multi-factor models of Fama and French (1993, 2015) and Carhart (1997).

⁴See, e.g., Ang and Piazzesi (2003); Ang, Piazzesi, and Wei (2006); Ang, Bekaert, and Wei (2007, 2008); Dai, Singleton, and Yang (2007); Le, Singleton, and Dai (2010); Le and Singleton (2010); Haubrich, Pennacchi, and Ritchken (2012); Gürkaynak, Sack, and Wright (2010); Chen, Liu, and Cheng (2010); Campbell, Sunderam, and Viceira (2017).

study the joint time-series behavior of inflation and aggregate market returns.⁵

Similar to us, Chen, Roll, and Ross (1986), Ferson and Harvey (1991), and Ang, Brière, and Signori (2012) estimate the inflation risk premium in the cross-section of stocks. However, these papers focus on the unconditional asset pricing implications of inflation risk. The typical point estimate is a negative inflation risk premium that is not statistically significant and varies considerably across sample periods. We provide a natural explanation for this evidence, which is a time-varying inflation risk premium that can change sign. We also contribute to this literature by studying the macroeconomic mechanism behind this time-variation, both empirically and in an equilibrium model.

Our model builds on the mathematics and economic intuition of the long-run risk model of Bansal and Yaron (2004), but the economic sources of risk are different. We contribute to the long-run risk literature in several ways. First, we use directly observable inflation as a predictor of consumption instead of an unobserved long-run risk component that must be inferred from asset prices using the model's assumptions. The observability of inflation thus resolves one of the main limitations of the long-run risk model. Our reliance on observables is advantageous also relative to other consumption-based models that contain varying degrees of "dark matter" (Chen, Dou, and Kogan, 2015; Cochrane, 2017). Second, predictability of consumption using inflation, and inflation persistence itself, both operate at business cycle frequency. In contrast, long-run risk operates at substantially lower frequencies. In addition, inflation shocks have a higher variance than long-run risk shocks. The combination of higher volatility and lower persistence of inflation shocks makes the inflation premium comparable in magnitude to the premium earned by the low volatility, but high persistence, long-run risk shocks. Third, because inflation is less persistent than long-run risk, the timing premium that agents are willing to pay for the early resolution of uncertainty in our model is considerably lower than in the long-run risk model by about 50%, which alleviates the criticism of Epstein, Farhi, and Strzalecki (2014). Finally, inflation predicts consumption growth in a time-varying way as determined by the nominal-real covariance, while the predictability of consumption with long-run risk is constant over time.

Recent articles in the bond market literature argue that the nominal-real covariance is an important driver of the time-variation in bond prices and in the comovement of stocks and bonds. Campbell, Sunderam, and Viceira (2017) document a simultaneous change in the sign of the term premium for US Treasuries and the nominal-real covariance at the turn of the century. Campbell, Pflueger, and Viceira (2015) and Song (2017) interpret the time-variation in bond risk through changes in monetary policy. We argue that monetary policy is unlikely to be the main driver of our results. The high and low inflation portfolio are not differentially exposed to monetary policy shocks. Moreover, the nominal-real covariance predicts returns also when controlling for different monetary policy regimes. David and Veronesi (2013) contend that it is the time-varying signaling role of inflation (plus learning) that drives the joint dynamics of stock and bond

⁵See, e.g., Modigliani and Cohn (1979); Wachter (2006); Gabaix (2008); Bansal and Shaliastovich (2012).

markets. Kang and Pflueger (2015) find that the nominal-real covariance affects corporate bond yield spreads through a credit channel. The nominal-real covariance also features prominently in recent term structure literature that analyzes the implications of the zero lower bound (see, e.g., Gourio and Ngo, 2016; Bilal, 2017; Bretscher, Hsu, and Tamoni, 2017).

Although we estimate an unconditional inflation risk premium in the bond market that is close to what we estimate in the stock market, our evidence suggests that the nominal-real covariance is not a strong time-series predictor of excess bond returns, consistent with Campbell, Sunderam, and Viceira (2017).⁶ Furthermore, previous papers study the nominal-real covariance as measured by the correlation between the returns of long-term government bonds and the aggregate stock market. To the best of our knowledge, we are the first to show that inflation predicts real consumption growth in a time-varying way, and to use this conditional (and exclusively macroeconomic) relation to measure the nominal-real covariance. We also study the pricing implications of time-variation in the nominal-real covariance in an asset class not studied before (the cross-section of stocks) and present new evidence on the relative importance of price and quantities of risk.

Finally, Gorodnichenko and Weber (2016) show that the conditional volatility of stock returns increases more for firms with stickier prices after monetary policy shocks. Weber (2015) finds that firms with stickier prices are more exposed to monetary policy shocks and earn a premium that is fully explained by market beta, because sticky price firms have more cyclical cash flows and both higher discount rate and cash flow beta. Although we also find that exposure in cash flows is important, and despite inflation risk and price stickiness being related, we find that neither exposure to the market nor to monetary policy shocks explains the variation in inflation premia that we document. Together, these differences suggest that price stickiness and our inflation-sorted portfolios are subject to different sources of risk.

2. Data

We use monthly inflation (Π_t) calculated as the percentage change in the seasonally adjusted Consumer Price Index for All Urban Consumers (CPI) available from the U.S. Bureau of Labor Statistics. Monthly per capita real consumption growth (ΔC_t) is calculated using seasonally adjusted aggregate nominal consumption expenditures on nondurables and services (from National Income and Product Accounts (NIPA) Table 2.8.5) combined with population numbers (from NIPA Table 2.6) and price deflator series (from NIPA Table 2.8.4).⁷ In our asset pricing tests, we use all ordinary common stocks traded on the NYSE, Amex, and Nasdaq from CRSP (excluding firms with negative book equity from Compustat). The CRSP value-weighted market portfolio, the one-month T-bill return, benchmark asset pricing factors, and industry portfolios are from

⁶Even though in their model, just like in ours, the nominal-real covariance is a key state variable driving risk premia, Campbell, Sunderam, and Viceira (2017) conclude in Section 5.3 that “bond returns have very limited predictability.”

⁷In Section 8, we replicate our main result using quarterly consumption data.

Kenneth French's website. Table 1 presents descriptive statistics for the sample period we use in our tests: July 1962 to December 2014. The start of the sample period coincides with the introduction of Amex stocks in the CRSP file and is common to many empirical studies of the cross-section.

3. The time-varying relation between inflation and consumption

To analyze the conditional relation between current inflation and future real consumption growth, we present in Panel A of Table 2 results from simple OLS regressions of future consumption growth from month $t + 1$ to $t + K$ on inflation over month t :

$$\Delta C_{t+1:t+K} = d_0^K + d_1^K \Pi_t + e_{t+1:t+K}, \quad K = 1, 3, 6, 12. \quad (1)$$

We see that the unconditional relation between current inflation and future consumption is negative, as measured by the coefficient estimate for d_1^K .⁸ Thus, an increase in inflation is bad news historically, which is consistent with evidence in, for instance, Piazzesi and Schneider (2006). However, over our relatively long sample period, the coefficient on inflation is not significantly different from zero for any horizon K . A potential explanation for this seeming inconsistency with the existing literature is that the unconditional regression masks variation over time in the relation between consumption and inflation.

To quantify the magnitude of this time-variation, we run the following two-stage test:

$$\Delta C_{s+1:s+K} = a_t^K + b_t^K \Pi_s + e_{s+1:s+K}, \quad \text{for } s = 1, \dots, t - K, \quad (2)$$

$$\Delta C_{t+1:t+K} = d_0^{c,K} + d_1^{c,K} (\widehat{a}_t^K + \widehat{b}_t^K \Pi_t) + e_{t+1:t+K}. \quad (3)$$

In the first stage, Eq. (2) regresses consumption growth on lagged inflation over a backward-looking window using all data available up to month t . Hence, the expanding window s runs from 1 to $t - K$. This regression is estimated using weighted least squares (WLS) where weights are given by an exponential weighting scheme that has a half-life converging to 60 months for large t . The expanding window ensures that we use as much information as possible, whereas an exponential decay in weights ascertains that the largest weight is given to the most recent observations. In the second stage, Eq. (3) uses the estimated coefficients and inflation observed at time t , that is, $\widehat{a}_t^K + \widehat{b}_t^K \Pi_t$, to predict consumption growth from $t + 1$ to $t + K$. This setup ensures that we use no forward-looking information when we predict consumption growth in the second stage. The absence of a look-ahead bias is important, because we will later use the estimated first-stage coefficient b_t^K to predict various returns in the time-series. Hence, we prefer our simple setup over more highly parametrized models that require the full sample for estimation, such as a state-space model with time-varying parameters.

If our two-stage structure correctly models the conditional expectation of consumption growth, we should find that $d_0^{c,K} = 0$ and $d_1^{c,K} = 1$. Panel B of Table 2 presents the results. To test significance, we report

⁸In the data, we lag inflation by an additional month, to account for the reporting delay of CPI numbers.

asymptotic Newey-West (with K lags) and bootstrapped t -statistics.⁹ The estimate of $d_1^{c,K}$ is significantly larger than zero at all horizons. For no horizon K can we reject the hypothesis that $d_0^{c,K} = 0$ nor that $d_1^{c,K} = 1$, suggesting that the two-stage structure provides an unbiased estimate of the conditional expectation of consumption growth. Moreover, in the second stage, we capture an economically large share of the variation in consumption growth, with an R^2 that increases from 3% at the one-month horizon to 15% at the one-year horizon.

Following the literature on predictive regressions (see, e.g., Welch and Goyal, 2007), we also calculate an out-of-sample R^2 ($R^2 - OOS$). The $R^2 - OOS$ measures the improvement from conditioning on inflation relative to a conditional model that includes only a time-varying constant, a_t^K , which measures lagged average consumption growth. We find that the $R^2 - OOS$ is similarly increasing from 2% at the one-month horizon to 10% at the one-year horizon. We conclude that inflation strongly predicts consumption growth when we account for the time-varying relation between these macro-variables.

Our main proxy for the nominal-real covariance, which we henceforth denote NRC_t^C , is the time-varying coefficient from the first-stage regression in Eq. (2) of 12-month consumption growth on lagged inflation, \widehat{b}_t^{12} . In robustness checks, we consider two alternative proxies for the nominal-real covariance. First, the time-varying relation between inflation and industrial production growth, denoted NRC_t^{IP} , is estimated by substituting industrial production for consumption in the left-hand side of Eq. (2).¹⁰ This alternative measure effectively treats inflation as a recession state variable of the type advocated in Cochrane (2005) and Kojien, Lustig, and Van Nieuwerburgh (2017). Second, we consider the measure based on asset prices that is used in Campbell, Pflueger, and Viceira (2015): the negative of the stock market beta of a long-term bond, denoted NRC_t^{-BB} . We estimate this time-varying beta using a 60-month rolling window regression of the 10-year constant maturity Treasury bond return on the aggregate stock market return.

Fig. 1 plots the three proxies of the nominal-real covariance and shows strong comovement, with correlations around 0.75. Consistent with Campbell, Pflueger, and Viceira (2015), NRC_t^{-BB} has changed sign from negative to positive in the early 2000s. In fact, all three measures have increased markedly since the turn of the century, as noted also in Bekaert and Wang (2010), Campbell, Sunderam, and Viceira (2017), and David and Veronesi (2013). In the remainder of this paper, we analyze empirically and theoretically the asset pricing implications of the time-variation in inflation risk driven by this nominal-real covariance.

4. The time-varying inflation risk premium

This section analyzes whether the inflation risk premium varies over time in a manner consistent with the time-variation in the nominal-real covariance.

⁹As detailed in Section I of the Online Appendix, these standard errors are derived from 1,000 block-bootstrap replications of the coefficient estimates to address concerns about errors-in-variables (EIV) bias.

¹⁰Seasonally adjusted industrial production growth is from the *FRED*[®] database of the St. Louis FED.

4.1. Estimating the inflation risk premium

We first describe our method to estimate the inflation risk premium.

4.1.1. Inflation betas

At the end of each sample month t , we measure the inflation exposure of firm i by estimating its historical “beta” of excess returns, $R_{i,t}$, with respect to monthly innovations in inflation. Following Fama and Gibbons (1984), Vassalou (2000), and Campbell and Viceira (2002), we filter out these innovations, denoted $u_{\Pi,t}$, using an ARMA(1,1)-model. This model empirically captures many of the important dynamics of inflation.¹¹

We estimate historical inflation betas in the same way that we estimate the nominal-real covariance in Eq. (2): using a WLS regression with exponential weights over an expanding window that uses all observations from the first month the stock is included in the sample up to month t . We require that stocks have at least 24 out of the last 60 months of returns available. The estimator of a stock’s inflation exposure, $\beta_{\Pi,i,t}$, is then given by

$$\left(\widehat{\alpha}_{i,t}, \widehat{\beta}_{\Pi,i,t}\right) = \arg \min_{\alpha_{i,t}, \beta_{\Pi,i,t}} \sum_{\tau=1}^t K(\tau) (R_{i,\tau} - \alpha_{i,t} - \beta_{\Pi,i,t} u_{\Pi,\tau})^2 \quad (4)$$

$$\text{with weights } K(\tau) = \frac{\exp(-|t - \tau|/h)}{\sum_{\tau=1}^{t-1} \exp(-|t - \tau|/h)}. \quad (5)$$

Using $h = \log(2)/60$, the half-life of the weights $K(\tau)$ converges to 60 months for large t , which is consistent with the standard five-year rolling window used in empirical asset pricing tests. Following Elton, Gruber, and Urich (1978), Cosemans, Frehen, Schotman, and Bauer (2016), and Levi and Welch (2017), we transform the estimated $\widehat{\beta}_{\Pi,i,t}$ using a Vasicek (1973) adjustment

$$\widehat{\beta}_{\Pi,i,t}^v = \widehat{\beta}_{\Pi,i,t} + \frac{\text{var}_{TS}(\widehat{\beta}_{\Pi,i,t})}{\left[\text{var}_{TS}(\widehat{\beta}_{\Pi,i,t}) + \text{var}_{CS}(\widehat{\beta}_{\Pi,i,t})\right]} \left[\text{mean}_{CS}(\widehat{\beta}_{\Pi,i,t}) - \widehat{\beta}_{\Pi,i,t}\right]. \quad (6)$$

Thus, each $\widehat{\beta}_{\Pi,i,t}^v$ is a weighted average of the stock’s beta estimated in the time-series and the average of all betas in the cross-section of month t , $\text{mean}_{CS}(\widehat{\beta}_{\Pi,i,t})$. The former receives a larger weight when it is estimated more precisely, that is, when the square of the standard error of the estimated beta, var_{TS} , is small relative to the cross-sectional variance of betas, var_{CS} . From this point forward, inflation betas refer to the WLS and Vasicek-adjusted betas and we drop the hat and superscript v . Section 8 presents results for a variety of alternative specifications that test the robustness of these methodological choices.

¹¹For example, Ang et al. (2007) write: “The motivation for the [ARMA(1,1)] model derives from a long tradition in rational expectations macroeconomics (see Hamilton, 1985) and finance (see Fama, 1975) that models inflation as the sum of expected inflation and noise. If expected inflation follows an AR(1) process, then the reduced-form model for inflation is given by an ARMA(1,1) model. The model also nicely fits the slowly decaying autocorrelogram of inflation.”

4.1.2. Inflation portfolios

Each month in our sample we create 30 value-weighted portfolios by two-way sorting all stocks at the intersection of ten inflation beta deciles and three size groups (Micro, Small, and Big). Following Fama and French (2008), these size groups are defined by the 20th and 50th percentiles of NYSE market capitalization at the end of the previous month. We then collapse the 30 portfolios into ten size-controlled inflation beta-sorted portfolios by averaging over the three size groups in each inflation beta decile. On the one hand, the smallest of stocks are illiquid, not in the set of stocks typically held by institutions that care most about inflation (such as pension funds), and their betas are harder to estimate. On the other hand, Ang, Brière, and Signori (2012) find that the best inflation hedgers in the CRSP file are the smallest stocks. To avoid favoring either hypothesis, we give equal weight to each size group in the inflation risk premium. Also, controlling for size in this way makes exposures to the market and other risk factors more homogeneous across portfolios.¹² With a burn-in period of 60 months, this leaves us with a sample of post-ranking returns from July 1967 to December 2014.

4.1.3. The inflation risk premium

We use three standard approaches to measure the inflation risk premium (IP_t). The first estimate of the inflation risk premium is the difference in return between the High and Low inflation beta decile portfolios (denoted IP_t^{HL}).

The second estimate of the inflation risk premium is the return of a Breeden, Gibbons, and Litzenberger (1989) maximum correlation mimicking portfolio (denoted IP_t^{MC}). We construct this portfolio by projecting inflation innovations on the returns of the ten inflation beta-sorted portfolios, as they should contain a large share of the information in $u_{\Pi,t}$ that is relevant for the cross-section of stock returns. Table OA.1 of the Online Appendix presents the weights and confirms that the portfolios contain relevant inflation information.

Finally, each month we run a cross-sectional regression of individual stock returns on lagged inflation beta, controlling for market capitalization, book-to-market, and momentum (see, e.g., Fama and French, 2008; Chordia, Goyal, and Shanken, 2015). The time-series of estimated slope coefficients on lagged inflation beta represents our third estimate of the inflation risk premium (denoted IP_t^{CS}).

4.1.4. Ex post inflation exposures

We present in Table 3 descriptive statistics for the set of ten inflation beta-sorted portfolios, as well as the three estimates of the inflation risk premium: IP_t^{HL} , IP_t^{MC} , and IP_t^{CS} . Panel A reports inflation exposures estimated with a simple OLS regression of portfolio returns on inflation innovations over the full sample. Analyzing whether these ex post exposures are large, economically and statistically, is important to assess whether inflation is a useless factor in the sense of Kan and Zhang (1999) and also serves as a reality check of

¹²In Section 8, we show that our conclusions on the time-variation in inflation risk premia extend when we perform a single sort of stocks on inflation beta, without controlling for size.

the estimation procedure. We see that the ex post exposures line up almost monotonically from High to Low. The dispersion is wide, with a large post-ranking beta of 3.00 for IP_t^{HL} . A beta of three translates to an incremental monthly return of 76 basis points on average when $u_{\Pi,t}$ increases by one standard deviation. For comparison, the CRSP value-weighted market portfolio, with an inflation beta of -1.96, loses 49 basis points on average for the same increase in $u_{\Pi,t}$. To accommodate comparison and interpretation, we have scaled the alternative measures of the inflation risk premium, IP_t^{MC} and IP_t^{CS} , to have identical ex post inflation exposure to IP_t^{HL} . For these alternative estimates of the inflation risk premium, the ex post exposure is also strongly significant.

Fig. 2 plots the post-ranking inflation beta one month, one year, two years, five years, and ten years after sorting. To do so, we fix the portfolio composition at the sorting date t and calculate monthly value-weighted returns up to ten years after.¹³ We then run a regression of monthly returns in $t + 1$, $t + 12$, $t + 24$, $t + 60$, and $t + 120$ on contemporaneous (with the returns) innovations in inflation. In this way, we mimic the ex post monthly inflation exposure for an investor that rebalances infrequently with respect to inflation beta. We see that post-ranking inflation beta is increasing in pre-ranking beta up to ten years after the sort. The post-ranking beta for IP_t^{HL} falls from 3.00 one month after sorting to a still large and significant 1.68 ten years after sorting.¹⁴ We conclude that our sort on ex ante inflation betas is powerful, because inflation betas are predictable many years ahead. Moreover, there is large cross-sectional dispersion in inflation risk, which is instrumental in identifying the inflation risk premium in the stock market.

We find that our ARMA(1,1)-inflation shocks are strongly correlated at 0.59 with first-differences in breakeven inflation, which is the difference between the yields of nominal and real ten-year constant maturity Treasury bonds (available from *FRED*[®] since 2003). Furthermore, Panel A of Table 4 shows that our sort (on exposures to ARMA(1,1)-innovations in inflation) also creates large ex post exposures to first-differences in breakeven inflation, which we denote ΔBE_t . The coefficient in a regression of returns of portfolios sorted on inflation beta on ΔBE_t increases monotonically in the inflation beta of the portfolios. The High-minus-Low difference in ΔBE_t -exposure is 9.79 with a t -statistic of 4.12. In Panel B, we sort individual stocks into deciles based on their exposure to ΔBE_t . We see that ex post exposure to ΔBE_t increases monotonically in ex ante exposure. More importantly, the High-minus-Low difference in ΔBE_t -exposure is not much larger than in Panel A at 12.47 ($t = 4.82$). Moreover, the correlation coefficient between the returns of the two High-minus-Low portfolios is 0.86, which suggests strong comovement between the two series.

¹³When a stock leaves the sample, we reallocate its market cap to the remaining stocks in a given portfolio.

¹⁴This finding is seemingly inconsistent with the conclusion in Ang, Brière, and Signori (2012) that inflation betas are hard to estimate out-of-sample. However, their conclusion is based on a smaller sample of stocks in the S&P500 from 1990 to 2009. In the Appendix to their paper, the authors report results that are consistent with ours for a sort using all stocks in the CRSP universe from 1967 to 2009.

4.2. The inflation risk premium over subsamples

We present in Panels B and C of Table 3, respectively, the inflation risk premium over the full sample as well as in two subsamples split around 2002. This split is motivated by the fact that the nominal-real covariance, NRC_t^C , increased above its historical mean during 2002, without falling below its mean again until the end of the sample (see Fig. 1). We present annualized average excess returns, Sharpe ratio, and CAPM alpha. We show in Section 8 that our main conclusions are robust to controlling for a broad set of benchmark asset pricing factors.

Over the full sample, average returns for the inflation beta-sorted portfolios are decreasing in inflation beta from 9.49% for Low to 5.26% for High. This dispersion translates to a marginally significant return for IP_t^{HL} equal to -4.23% ($t = -2.08$). This effect is similar in both magnitude and significance for IP_t^{MC} and IP_t^{CS} as well as in alpha. This unconditional inflation risk premium translates to a Sharpe ratio around -0.35, which is in the same order of absolute magnitude as the Sharpe ratio of the aggregate stock market. The last column of the table presents the p -value from the monotonicity test of Patton and Timmermann (2010), where the null hypothesis is a non-decreasing relation of returns in inflation beta. For average returns, Sharpe ratios, and alphas, this null is marginally rejected.

Note that this relatively modest and only marginally significant unconditional estimate of the inflation risk premium does not imply that the inflation risk premium is economically unimportant. Rather, we argue that this unconditional estimate masks important variation over time that is driven by the nominal-real covariance and, in particular, its recent reversal. To see this, we show in Panel C that the inflation risk premium is considerably larger in the pre-2002 subsample, and strongly statistically significant. Average returns for IP_t^{HL} , IP_t^{MC} , and IP_t^{CS} are -7.49%, -5.85%, and -6.81%, respectively, with t -statistics that all pass the data-mining corrected cutoff of three proposed in Harvey, Liu, and Zhu (2016). These average returns translate to a price of risk below -0.50 as measured by the Sharpe ratio, well above the aggregate stock market (in absolute value). Average returns are almost monotonically decreasing in inflation beta from 9.33% for the Low portfolio to 1.84% for the High portfolio. Consequently, the null of a non-decreasing relation is comfortably rejected at a p -value of 0.005. We see similar results for CAPM alphas, suggesting that the one-factor CAPM does not price the cross-section of inflation portfolios in this subsample.

Post-2002, the inflation risk premium changes dramatically. Average returns are now increasing in inflation beta, resulting in an average return of 5.42% for the High-minus-Low portfolio and 3.49% for the cross-sectional regression portfolio. Although these estimates are non-negligible economically, they are not statistically significant over the short sample of 12 years. Moreover, the alternative estimate of the inflation risk premium, IP_t^{MC} , is virtually zero at -0.06% and all three CAPM alphas are small and insignificant. Thus, the evidence suggests that the inflation risk premium is relatively small post-2002 and, if anything, positive, which is markedly different from the period pre-2002.

The last panel of the table presents the post-2002 minus pre-2002 differences. Average returns and alphas

are monotonically decreasing from High to Low from 13.53% to 0.62% and 4.28% to -5.13%, respectively. Although the effects for IP_t^{HL} , IP_t^{MC} , and IP_t^{CS} are not always significant, these differences are economically large and the monotonicity tests comfortably reject (at p -values below 0.01) the null of a non-increasing relation between the average or CAPM-adjusted (post-2002 minus pre-2002 difference in) return and inflation beta. Table OA.2 of the Online Appendix shows a slightly larger increase in the inflation risk premium when we split the sample in months where NRC_t^C is above versus below its mean.

In short, we find that the inflation risk premium has increased markedly over time, consistent with the reversal in the nominal-real covariance from negative (on average since the 1960s) to positive (since the early 2000s).

4.3. The inflation risk premium conditional on the nominal-real covariance

Because a sample split provides only a crude measure of the amount of time-variation in the inflation risk premium, we now turn to predictive regressions. We regress excess returns of the inflation portfolios (compounded over horizons K of one, three, and 12 months) on NRC_t^C using:¹⁵

$$R_{p,t+1:t+K} = L_0 + L_{NRC}NRC_t^C + \varepsilon_{t+1:t+K}. \quad (7)$$

For each horizon, Panel A of Table 5 presents the estimated coefficients (annualized), both asymptotic Newey-West (with K lags) and bootstrapped (see Section 1 of the Online Appendix) t -statistics, and the adjusted R^2 for the individual decile portfolios and the three estimates of the inflation risk premium (IP_t^{HL} , IP_t^{MC} , and IP_t^{CS}). We also present the p -values from Patton and Timmermann (2010) monotonicity tests. NRC_t^C is standardized to have mean equal to zero and standard deviation equal to one to accommodate interpretation. Thus, L_0 measures the average excess return of the respective portfolio, that is, the unconditional inflation risk premium, whereas L_{NRC} measures the increase in annualized portfolio return for a one standard deviation increase in the nominal-real covariance.

We find that the effect of NRC_t^C on the inflation risk premium is positive, economically large, and statistically significant (based on both asymptotic and bootstrap inference). For a one standard deviation increase in the nominal-real covariance, the inflation risk premium as estimated using IP_t^{HL} , IP_t^{MC} , and IP_t^{CS} increases by about 3% to 4% for $K = 1$. This increase is relative to an unconditional inflation risk premium of about -4%. The effect strengthens at longer horizons, growing to about 4% to 5% for $K = 12$. The R^2 are also increasing in horizon, from around 0.5% for $K = 1$ to over 6.5% for $K = 12$. Finally, looking at the decile portfolios, we see that L_{NRC} is increasing monotonically from Low to High inflation beta, such that the null of non-monotonicity is comfortably rejected.

¹⁵To see that these regressions use no forward-looking information, recall that NRC_t^C is the slope coefficient b_t^{12} in the regression of Eq. (2): $\Delta C_{s+1:s+12} = a_t^{12} + b_t^{12}\Pi_s + e_{s+1:s+12}$, with $s = 1, \dots, t - 12$. Consistent with the fact that b_t^{12} moves slowly over time, our results are insensitive to lagging the predictor by an additional month to control for a reporting delay in consumption data.

We conclude that the inflation risk premium is strongly time-varying with the nominal-real covariance. The economic intuition for this result is that high inflation beta stocks are attractive to hedge consumption risk when inflation predicts consumption with a negative sign as it did historically. However, these same stocks are not attractive as a hedge anymore, and could even expose investors to additional consumption risk, once the nominal-real covariance starts increasing, as it does towards the end of our sample, for instance. As a result, the returns of high versus low inflation beta stocks should be increasing in the nominal-real covariance, which is what we find.

Panel A of Table OA.3 of the Online Appendix shows that this conclusion is robust for two alternative measures of the nominal-real covariance, namely, the time-varying relation between inflation and industrial production growth (NRC_t^{IP}) and the negative of the stock market beta of a long-term Treasury bond (NRC_t^{-BB}). Each alternative measure predicts the inflation risk premium with a positive coefficient that is comparable in magnitude, significance, and horizon pattern to what we have seen before. Panel B of Table OA.3 shows that the nominal-real covariance predicts the inflation risk premium with a positive sign in both the first and second half of our sample period, although the predictive relation is stronger in the latter subsample. This result indicates that the time-variation in the inflation risk premium that we uncover is not exclusively driven by the recent reversal in the nominal-real covariance. Moreover, when we end our sample in 2007, before interest rates reached the zero lower bound (ZLB), we find virtually identical effects to what we find in the full sample up to 2014. Thus, the time-variation that we uncover cannot be fully driven by asset pricing mechanisms that rely on interest rates at or close to the ZLB, such as those proposed in, for instance, Gourio and Ngo (2016), Bilal (2017), and Bretscher, Hsu, and Tamoni (2017).

5. Origins of the time-varying inflation risk premium

To understand the drivers of the time-variation in the inflation risk premium, we next analyze (i) the relative contribution of the price versus quantities of inflation risk in our inflation beta-sorted decile portfolios, (ii) the importance of industry affiliation, and (iii) the relative contribution of cash flows versus discount rates to inflation exposure.

5.1. Price and quantities of inflation risk

Fig. 3 shows the time-variation in inflation beta for the High and Low decile portfolios, the difference between these two betas, and the average inflation beta of all ten portfolios.¹⁶ These inflation betas are estimated by applying Eqs. (4) and (5) to the post-ranking returns of each portfolio.

We see that there is considerable variation in inflation betas over time. For instance, the sudden increase in inflation betas around the turn of the century, consistent with the reversal in the nominal-real covariance,

¹⁶The average beta is almost perfectly correlated (at 0.99) with the inflation beta of the CRSP value-weighted market portfolio.

stands out. The correlation between the inflation beta of the High and Low portfolios is 0.93, and both share a correlation of 0.98 with the cross-sectional average of all ten inflation betas. Moreover, although the High-minus-Low inflation beta is large and positive on average (and for most of the sample), the standard deviation of the High-minus-Low beta is much smaller than the standard deviation of either portfolio's beta in isolation (1.36 versus around 3.50). These results imply that time-variation in the quantity of inflation risk is mostly common across portfolios (and in the stock market, more generally). This result is new to the literature and provides important information for our understanding of the inflation risk premium in the stock market.

To see this and quantify the relative contribution of quantities versus price of risk, we run pooled regressions of returns of the inflation beta-sorted decile portfolios, $p = High, 2, \dots, Low$ on inflation beta ($\beta_{\Pi,p,t}$), the nominal-real covariance (NRC_t^C), and an interaction term:

$$R_{p,t+1:t+K} = L_0 + L_{\beta_{\Pi}}\beta_{\Pi,p,t} + L_{NRC}NRC_t^C + L_{\beta_{\Pi} \times NRC}(\beta_{\Pi,p,t} \times NRC_t^C) + \varepsilon_{p,t+1:t+K}. \quad (8)$$

Panel A of Table 6 presents the estimated coefficients (annualized) and t -statistics using asymptotic standard errors calculated following Driscoll and Kraay (1998), which are robust to heteroskedasticity and rather general forms of cross-sectional and temporal dependence when the time dimension becomes large. We focus on the annual horizon ($K = 12$) and present consistent evidence for shorter horizons in Table OA.4 of the Online Appendix. To accommodate interpretation, NRC_t^C is standardized in the time-series (as in Table 5) and $\beta_{\Pi,p,t}$ is demeaned in the pool, that is, its mean across portfolios and time is zero.

In Column 1, we do not include the interaction term (set $L_{\beta_{\Pi} \times NRC} = 0$) and find that neither inflation beta nor the nominal-real covariance is a significant predictor of returns in the pool of inflation beta-sorted portfolios. The R^2 is also low at 0.71%. In Column 2, we include the interaction term and see a large increase in the R^2 to 8.80%. The estimated coefficient on the interaction term is large and significant at 2.36 ($t = 3.63$). This estimate indicates that expected returns are decreasing in inflation beta when the nominal-real covariance is negative, but increasing when it is positive. This result is in line with the large negative inflation risk premium we find in the subsample pre-2002, and the small positive inflation risk premium post-2002.

To further quantify the interaction effect between inflation beta and the nominal-real covariance, Panel B presents predicted risk premia in four distinct cases when: i) inflation beta is at plus or minus one (pooled) standard deviation from its mean and ii) the nominal-real covariance is at plus or minus one (time-series) standard deviation from its mean. These predictions effectively evaluate the joint significance of the coefficients, thus alleviating concerns about correlation in the independent variables. We present both asymptotic and bootstrapped t -statistics.¹⁷

¹⁷For details on the bootstrap procedure, see Section 1 of the Online Appendix.

We find that the predicted risk premium is largest at 16.56% (with a bootstrapped t -statistic > 3) for the case in which a portfolio has a low inflation beta and the nominal-real covariance is also low. This finding is intuitive: a portfolio of stocks is most risky when it comoves negatively with inflation shocks that contain bad news about future consumption growth. The predicted risk premium is also positive, but smaller at 7.65% ($t_{boot} = 1.65$), when both inflation beta and the nominal-real covariance are high. When inflation beta is high (low) and the nominal-real covariance is low (high), predicted premia are close to zero. To interpret these predictions, note that the difference in expected returns between high and low inflation beta is positive when the nominal-real covariance is high and negative when the nominal-real covariance is low. This finding is consistent with the idea that the stocks that are most attractive to hedge high marginal utility states are those that have low (high) returns when inflation increases and the nominal-real covariance is positive (negative). This difference is largest in magnitude and significance when the nominal-real covariance is low, because in this case the covariance between inflation and future consumption growth is stronger in absolute terms (the nominal-real covariance at one standard deviation below the mean equals -2.24, which is relative to -0.51 one standard deviation above the mean).¹⁸

We conclude that variation in inflation beta, that is, the quantity of inflation risk, contains important information about returns when combined with the nominal-real covariance. We further decompose inflation beta into two components: (i) a portfolio-specific component that captures cross-sectional variation in unconditional average inflation betas, $\widehat{\beta}_{\Pi,p} = T^{-1} \sum_{t=1}^T \beta_{\Pi,p,t}$; and, (ii) a time-specific component that captures time-series variation in the cross-sectional average inflation beta, $\widehat{\beta}_{\Pi,t} = 10^{-1} \sum_{p=1}^{10} \beta_{\Pi,p,t}$.

In Column 3, we see that the interaction between the portfolio-specific component of inflation beta and the nominal-real covariance ($\widehat{\beta}_{\Pi,p} \times NRC_t^C$) contains independent information about returns in the pool (relative to the individual effects of inflation beta and the nominal-real covariance). The coefficient estimate for the interaction term is 2.45 ($t = 2.80$). Setting the quantity of inflation risk equal to the unconditional spread in inflation beta between the High and Low inflation beta portfolios, a one standard deviation increase in the nominal-real covariance translates to a large and significant increase of 5.53% in the spread in expected returns between the High and Low portfolios ($5.53 = L_{\widehat{\beta}_{\Pi,p} \times NRC} \times \sigma_{NRC} \times (\widehat{\beta}_{\Pi,High} - \widehat{\beta}_{\Pi,Low}) = 2.45 \times 1 \times 2.26$, $t_{boot} = 2.17$).¹⁹ The R^2 in Column 3 is only 1.40%, which implies that this specification still misses variation in inflation betas that is relevant for the pool of inflation risk premia.

The missing variation is time-variation in the cross-sectional average inflation beta. To see this, Column

¹⁸Although the magnitude of the difference between high and low inflation beta may seem large, at -17.05% when the nominal-real covariance is low versus 10.98% when it is high, these effects are foreshadowed by the sorts we report in Table 3. The difference between the High and Low inflation beta portfolio is -7.49 pre-2002 and 5.42% post-2002, and the 2002 cutoff only approximately splits the sample in low versus high nominal-real covariance months.

¹⁹This marginal effect is consistent with the regressions in Table 5, which ask whether variation in the nominal-real covariance predicts High-minus-Low returns over time. The pooled specification in Column 3 similarly asks whether the nominal-real covariance explains variation over time in the relative return of high versus low inflation beta portfolios, because the product $\widehat{\beta}_{\Pi,p} \times NRC_t^C$ varies monotonically from High to Low at each point in time due to the unconditional variation in inflation betas.

4 shows that the interaction between the time-specific component of inflation betas and the nominal-real covariance ($\widehat{\beta}_{\Pi,t} \times NRC_t^C$) contains relevant information about returns, given the large R^2 of 8.26% in this specification. The coefficient estimate is 2.45 ($t = 3.45$), which translates to a large and significant increase of 6.92% in the difference in expected return between periods of low (that is, the nominal-real covariance is one standard deviation below its mean) and average nominal-real covariance, when the cross-sectional average inflation beta decreases by one standard deviation ($6.92 = L_{\widehat{\beta}_{\Pi,t} \times NRC} \times -\sigma_{\widehat{\beta}_{\Pi,t}} \times -\sigma_{NRC} = 2.45 \times -2.83 \times -1$, $t_{boot} = 2.14$). The intuition for this result is that at a time when the average inflation beta in the cross-section is low, the average stock is especially risky when the nominal-real covariance is negative (and a shock to inflation signals lower future growth).

In the last column of Panel A in Table 6, we include the two components of inflation beta (interacted with NRC_t^C) simultaneously. The R^2 equals 8.80%, just like the model in Column 2, which includes the interaction term between raw inflation betas and the nominal-real covariance, $\beta_{\Pi,p,t} \times NRC_t^C$. This result suggests that the portfolio- and time-specific components of inflation betas together capture all the relevant variation in inflation betas for risk premia.²⁰ Unconditional differences in the quantity of inflation risk, combined with the nominal-real covariance, drive the widening and compressing of the cross-section of inflation risk premia over time. With inflation betas held constant, this time-variation is thus driven by variation in the price of inflation risk.²¹ Time-variation in the average inflation beta, combined with the nominal-real covariance, determines the level of inflation risk premia in the cross-section. This last result suggests that the aggregate stock market risk premium contains an inflation component that is time-varying due to both the price and quantity of inflation risk.

To test this suggestion, we run time-series regressions of CRSP value-weighted market (excess) returns on the market's inflation beta, the nominal-real covariance, and an interaction term:

$$R_{m,t+1:t+12} = L_{0,m} + L_{\beta_{\Pi,m}}\beta_{\Pi,m,t} + L_{NRC,m}NRC_t^C + L_{\beta_{\Pi} \times NRC,m}(\beta_{\Pi,m,t} \times NRC_t^C) + \varepsilon_{t+1:t+12}, \quad (9)$$

where $\beta_{\Pi,m,t}$ is estimated by regressing market returns on inflation innovations following Eq. (4).²² We present the results in Table 7. To accommodate interpretation, NRC_t^C is standardized and $\beta_{\Pi,m,t}$ is demeaned in the time-series (as in Table 6).

In Panel A, we see that the market's inflation beta and the nominal-real covariance do not predict market returns in isolation. However, including their interaction term ($\beta_{\Pi,m,t} \times NRC_t^C$) increases the R^2 from close

²⁰By implication, variation in inflation betas that is portfolio×time-specific, that is, variation that is not captured by a portfolio or a time fixed effect, contains no information for risk premia in the pool of inflation beta-sorted portfolios.

²¹Previous evidence in Table 5 is consistent with this conclusion. To see why, recall that the inflation risk premium estimated using cross-sectional regressions, IP_t^{CS} , represents the return of a portfolio strategy with an ex ante inflation beta fixed at one (Fama, 1976). Hence any variation in the returns of IP_t^{CS} must follow from variation in the price of inflation risk.

²²In Table OA.5 of the Online Appendix we present similar results when we proxy for the market portfolio using the average of the ten decile portfolios.

to zero to 8.59%. The coefficient estimate on the interaction term is economically large and significant at 2.85% ($t = 3.55$). Similar to Panel B of Table 6 we also present the predicted market risk premium in four distinct cases. When both the nominal-real covariance and the market's inflation beta are low (that is, large and negative), stocks are most risky, and the predicted market risk premium is 13.61% ($t = 4.08$), which is about double its unconditional mean. Furthermore, the difference in the predicted market risk premium when the market's inflation beta goes from low to high (i.e., from one standard deviation below the market's mean inflation beta to one standard deviation above) equals 6.65% when the nominal-real covariance is high but -17.64% when the nominal-real covariance is low.

In Panel B, we present nonparametric estimates of the market risk premium, which are similar in absolute and relative magnitude to those reported in Panel A in all four cases. These estimates are calculated by averaging realized excess market returns in sample months $t + 1$ when the market's inflation beta and the nominal-real covariance in month t are either below the 40th or above the 60th percentile. In conclusion, the results in Table 7 confirm that the interaction between the market's inflation beta and the nominal-real covariance contains information about future market returns.

In Table OA.6 of the Online Appendix, we find that the interaction term between inflation beta and the nominal-real covariance, $\beta_{\Pi,p,t} \times NRC_t^C$, remains large and significant in a pooled regression that controls for market beta as well as the interaction between market beta and the nominal-real covariance. The same result obtains for the portfolio- and time-specific components of inflation betas. We conclude that variation in market betas does not fully explain the variation in risk premia we uncover.

5.2. Industry affiliation and inflation risk

Having seen that inflation betas vary persistently and considerably over time and in the cross-section of stocks, we ask to what extent our results are driven by across-industry variation in inflation risk and returns. Among others, Boudoukh, Richardson, and Whitelaw (1994) and Ang, Brière, and Signori (2012) show that inflation betas vary across industries.

To start, we construct long-short industry factors using the top and bottom $K = 1, \dots, 5$ industries as ranked by their unconditional full sample exposure to innovations in inflation. Consistent with common intuition, Panel A of Table 8 shows that coal, oil, and gold are among the best unconditional inflation hedges while meals, textiles, and clothes are among the worst.²³ In every sample month, we then sort individual stocks into decile portfolios based on their rolling exposure to each industry factor, leaving out all stocks that belong to those top and bottom K industries. In Panel B of Table 8, we first show that the High-minus-Low returns from these sorts are significantly exposed ex post to ARMA(1,1)-innovations in inflation for all K . In the remainder of Panel B, we see that the nominal-real covariance significantly predicts the returns of these

²³Table OA.7 of the Online Appendix shows that similar industries are also among the best (worst) conditional inflation hedges and that the industry decomposition is quite stable over time.

High-minus-Low portfolios in the time-series, again for all K .

These results confirm that inflation betas vary across industries and that such variation contributes to time-variation in returns. However, these results do not fully separate the across-industry from the within-industry contribution to the variation in inflation premia that we uncover in Section 4. To isolate these within- and across-industry components, we calculate the inflation beta of each of the 48 industries from Kenneth French's website as the value-weighted average inflation beta of all stocks in that industry (estimated conditionally as in Eqs. (4) and (5)). We exclude industry-months that contain fewer than ten stocks. To obtain the across-industry component, we sort the industries into quintile portfolios, such that each across-industry portfolio contains nine or ten value-weighted industries. Using this sort, we calculate equal-weighted across-industry quintile returns. For the within-industry component, we first construct five value-weighted portfolios within each industry by splitting at the quintiles of ranked inflation betas. This gives us a total of 48-by-5 value-weighted portfolios. Within-industry quintile portfolio returns are then calculated as the equal-weighted average over the 48 industries. Table 9 presents the predictive regression of 12-month compounded returns from each of these sorts on the nominal-real covariance. The last column presents the aggregate within-industry (across-industry) effect calculated as the difference between the High and Low within-industry (across-industry) inflation beta portfolios.

The unconditional inflation risk premium, as estimated by the intercept in the predictive regression, is negative for both the within- and across-industry sorts. This unconditional premium is larger in magnitude across-industry (at -3.36% with $t = -1.95$) than within-industry (at -1.73% with $t = -1.55$). The coefficient estimate for the nominal-real covariance is positive, economically large, and significant for both sorts at 4.63% ($t = 2.70$) across-industry and 3.24% ($t = 3.29$) within-industry. Both coefficient estimates follow from monotonic effects from High to Low inflation beta and translate to an R^2 that is slightly over 10%. Monotonicity in the across-industry sort is important, because it suggests that our results are not solely driven by industries that are highly exposed to a particular component of aggregate inflation, such as oil. The within-industry result suggests that variation in inflation risk is priced in a time-varying way independent of the industry to which a stock belongs.

We conclude that although inflation risk certainly varies across industries, this variation is not crucial for generating the time-variation in inflation risk premia that we find.

5.3. Inflation risk in cash flows versus discount rates

In this section, we ask whether the cross-sectional variation in inflation risk in our portfolios is driven by cash flows or discount rates. We use the vector autoregression (VAR) approach of Campbell (1991) and Campbell and Vuolteenaho (2004) to decompose the unexpected component of inflation portfolios' monthly returns into discount rate and cash flow news, $R_{p,t+1} - E_t[R_{p,t+1}] = -N_{DR,t+1} + N_{CF,t+1}$. With this decomposition in hand, inflation beta, measured as the coefficient in a regression of $R_{p,t+1} - E_t[R_{p,t+1}]$ on

ARMA(1,1)-innovations in inflation $u_{\Pi,t+1}$, can be decomposed into a discount rate inflation beta plus a cash flow inflation beta:

$$\beta_{\Pi} = \frac{Cov(R_{p,t+1} - E_t[R_{p,t+1}], u_{\Pi,t+1})}{Var(u_{\Pi,t+1})} \quad (10)$$

$$= \frac{Cov(-N_{DR,t+1}, u_{\Pi,t+1})}{Var(u_{\Pi,t+1})} + \frac{Cov(N_{CF,t+1}, u_{\Pi,t+1})}{Var(u_{\Pi,t+1})} \quad (11)$$

$$= \beta_{\Pi}^{-DR} + \beta_{\Pi}^{CF}. \quad (12)$$

We consider three different specifications that include different state variables in the VAR. For the first specification, we use the state variables of Campbell and Vuolteenaho (2004), which are the term spread, the market price-earnings ratio, and the value spread. For the second, we use the market dividend yield, default spread, and term spread (following, e.g., Maio and Santa-Clara, 2012; Boons, 2016). For the third, we add the nominal-real covariance to the first specification. To ensure that the VAR models are correctly specified, we additionally include in all three specifications the dividend yield of the respective inflation beta-sorted portfolio as a state variable.²⁴

Table 10 reports the inflation beta, β_{Π} , and its decomposition into β_{Π}^{-DR} and β_{Π}^{CF} , for the High and Low inflation beta portfolios. The table also shows the share of the variance of $R_{p,t+1} - E_t[R_{p,t+1}]$ explained by $-N_{DR,t+1}$, $N_{CF,t+1}$, and twice their covariance.²⁵ Consistent with the evidence discussed in Section 4.1.4, we find a large difference in inflation exposure in unexpected returns between the High and Low portfolios.

In all three specifications, we find a High-minus-Low beta, $\beta_{\Pi,High} - \beta_{\Pi,Low}$, close to 1.8. The discount rate component of beta, β_{Π}^{-DR} , is negative for the High and Low portfolios. In contrast, the cash flow component β_{Π}^{CF} is positive for the High portfolio but negative for the Low one. It follows that in the inflation beta of the High-minus-Low portfolio, the discount rate component of the High and Low components offset each other, while the cash flow components reinforce each other. In the last row of the table, we present the fraction of the High-minus-Low inflation beta that can be attributed to exposure to cash flows, $(\beta_{\Pi,High}^{CF} - \beta_{\Pi,Low}^{CF}) / (\beta_{\Pi,High} - \beta_{\Pi,Low})$. In the second specification, this fraction equals 80%. In the other two specifications, this fraction is above 100%, which implies that the discount rate component actually reduces the spread in inflation betas in the cross-section of stocks. Thus, the evidence suggests that the inflation beta of the High-minus-Low portfolio is largely driven by exposure to cash flows. Importantly, this conclusion applies even though the relative contribution to the variance of returns (column labeled “% Var”) attributed to discount rate and cash flow news varies considerably across the three specifications.

²⁴See Campbell and Ammer, 1993; Engsted, Pedersen, and Tanggaard, 2012.

²⁵To reduce noise, the High and Low portfolios in this section are top and bottom terciles. These terciles simply average returns over the top and bottom three deciles of our original inflation beta-sorted portfolios.

6. Model

To motivate our empirical findings, we model the processes for consumption, inflation, and the nominal-real covariance as exogenous, generating the time-varying predictability of consumption with inflation uncovered in the data in Section 3. The model then produces — via pricing through the Euler equation of a representative agent with Epstein-Zin utility — equilibrium quantities of inflation risk for the cross-section of stocks and the aggregate market, together with a market price of inflation risk. These quantities and prices behave qualitatively and quantitatively like the ones we find empirically in Sections 4 and 5.²⁶

6.1. Exogenous processes and stochastic discount factor

The exogenous processes for real consumption growth (Δc_{t+1}), inflation (π_{t+1}), the nominal-real covariance (φ_{t+1}), and real dividend growth for asset i ($\Delta d_{i,t+1}$) are given by

$$\pi_{t+1} = \mu_\pi + \rho_\pi (\pi_t - \mu_\pi) + \xi_\pi u_t + \phi_\pi u_{t+1}, \quad (13)$$

$$\Delta c_{t+1} = \mu_c + \rho_c (\pi_t - \mu_\pi) + \xi_c \varphi_{t-1} u_t + \sigma_c \eta_{t+1}, \quad (14)$$

$$\varphi_{t+1} = \varphi_0 + v (\varphi_t - \varphi_0) + \sigma_w w_{t+1}, \quad (15)$$

$$\Delta d_{i,t+1} = \mu_i + \rho_i (\pi_t - \mu_\pi) + \xi_i \varphi_{t-1} u_t + \sigma_i \eta_{t+1}, \quad (16)$$

$$u_t, \eta_t, w_t \sim iid \mathcal{N}(0, 1). \quad (17)$$

Consistent with our empirical analysis, inflation in Eq. (13) is an ARMA(1,1)-process with innovation $u_{\Pi,t+1} \equiv \phi_\pi u_{t+1}$. Inflation and inflation shocks affect the real economy via expected consumption growth in Eq. (14). Inflation predicts consumption growth unconditionally and in a time-invariant way with sign and magnitude determined by ρ_c . In contrast, the inflation shock u_t predicts consumption growth in a time-varying manner, where the sign and magnitude of the predictability are given by $\xi_c \varphi_{t-1}$. The nominal-real covariance, φ_t , follows the mean-reverting process given in Eq. (15) and can change sign over time. Dividends in Eq. (16) are subject to the same risks as consumption, but with potentially different exposures. The shocks u_{t+1} , η_{t+1} and w_{t+1} are standard normal.

The representative agent has Epstein-Zin utility, which yields the (log) stochastic discount factor (SDF)

$$m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t+1} - (1 - \theta) r_{c,t+1}, \quad (18)$$

where $r_{c,t+1}$ is the (log) return on the wealth portfolio, that is, the claim on aggregate consumption; $\delta \in (0, 1)$ is the discount rate; $\psi > 0$ is the elasticity of intertemporal substitution (EIS); and $\theta = (1 - \gamma) / (1 - 1/\psi)$, with $\gamma > 0$ the coefficient of relative risk aversion.

²⁶Detailed derivations for all the results presented in this section can be found in Section 2 of the Online Appendix.

6.2. The nominal-real covariance

The nominal-real covariance, φ_t , plays a dual role in our model. It determines (i) the time-variation in the predictability of consumption with inflation and (ii) the stochastic volatility of expected consumption growth.

To see that the nominal-real covariance determines the time-variation in the predictability of consumption, we run in the model the same two-stage procedure that we run in the data in Eqs. (2) and (3). The model-implied coefficient on inflation in the first-stage regression (2) is

$$b_t^K = \frac{Cov_t \left(\sum_{j=1}^K \Delta c_{t+1+j}, \pi_{t+1} \right)}{Var_t(\pi_{t+1})} = \rho_c h_K + \frac{\xi_c}{\phi_\pi} \varphi_t, \quad (19)$$

where h_K is a constant that depends on the prediction horizon K and parameters of the inflation process.²⁷ If $\xi_c \neq 0$, Eq. (19) shows that inflation predicts cumulative consumption growth in a time-varying manner determined by φ_t . Eq. (19) clarifies why we call φ_t the nominal-real covariance. Under the null of the model, φ_t is a linear transformation of $\widehat{b}_t^{12} = NRC_t^C$, the same variable we use in the data to measure the nominal-real covariance. The second-stage regression in Eq. (3) also holds exactly in the model with $d_0^{c,K} = 0$ and $d_1^{c,K} = 1$ (for any K and any calibration of the model).

To see that the nominal-real covariance also plays the role of stochastic volatility of expected consumption growth, it is instructive to compare our model to the long-run risk model of Bansal and Yaron (2004). Replacing $\pi_t - \mu_\pi$ with long-run risk and setting $\xi_\pi = \xi_c = \xi_i = 0$, we recover the basic version of the long-run risk model that does not have stochastic volatility. Thus, the new elements in our model, in addition to considering inflation instead of long-run risk, are the terms $\xi_\pi u_t$, $\xi_c \varphi_{t-1} u_t$, and $\xi_i \varphi_{t-1} u_t$ in Eqs. (13), (14), and (16). In the version of the long-run risk model with stochastic volatility, there is a single mean-reverting process that drives the stochastic volatilities of both expected and unexpected consumption growth (where we recall that long-run risk is exactly equal to expected consumption growth). In our model, we have stochastic volatility for neither inflation nor unexpected consumption growth, but we do have stochastic volatility for expected consumption growth through the term $\xi_c \varphi_{t-1} u_t$.²⁸

²⁷Using changes in inflation ($\pi_{t+1} - \pi_t$) or inflation innovations ($u_{\Pi,t+1}$) instead of inflation in Eq. (19) gives b_t^K that are also affine in φ_t . When standardized to have mean zero and variance one, these different measures of the nominal-real covariance are identical. Empirically, using inflation shocks (for instance, from an ARMA(1,1)-process) instead of inflation (Π_t) in Eqs. (2) and (3) gives similar estimates of the nominal-real covariance b_t^{12} with correlations above 90%.

²⁸An additional advantage of modeling φ_t (instead of φ_t^2 as in the long-run risk literature) as an AR(1) process is that volatilities are guaranteed to always be positive.

6.3. The price of inflation risk

To understand the asset pricing implications of the model, we study innovations in the SDF by using a decomposition similar to that in Campbell, Giglio, Polk, and Turley (2018):

$$m_{t+1} - E_t m_{t+1} = -\gamma (E_{t+1} - E_t) \Delta c_{t+1} - (1 - \theta) \kappa_1 (E_{t+1} - E_t) w c_{t+1} \quad (20)$$

$$= -\gamma N_{C,t+1} - (1 - \theta) \kappa_1 \left(\left(1 - \frac{1}{\psi}\right) N_{PATH,t+1} + \frac{\theta \kappa_1^2}{2} N_{RISK,t+1} \right), \quad (21)$$

where we define

$$\begin{aligned} N_{C,t+1} &\equiv (E_{t+1} - E_t) \Delta c_{t+1}, \\ N_{PATH,t+1} &\equiv (E_{t+1} - E_t) \sum_{j=0}^{\infty} \kappa_1^j E_{t+1+j} [\Delta c_{t+2+j}], \\ N_{RISK,t+1} &\equiv (E_{t+1} - E_t) \sum_{j=0}^{\infty} \kappa_1^j Var_{t+1+j}(w c_{t+2+j}), \end{aligned}$$

and κ_1 is an approximation constant generally close to one. Eq. (20) decomposes news to the SDF into news to contemporaneous consumption growth ($N_{C,t+1} = (E_{t+1} - E_t) \Delta c_{t+1}$) and news to the (log) wealth consumption ratio ($(E_{t+1} - E_t) w c_{t+1}$). Eq. (21) further decomposes news to the wealth consumption ratio into $N_{PATH,t+1}$, which is news about the future expected path of consumption, and $N_{RISK,t+1}$, which is “news about risk” (revisions in the future conditional variance of the wealth-consumption ratio).

Using Eq. (14), news to contemporaneous consumption growth can be written as $N_{C,t+1} = \sigma_c \eta_{t+1}$, which shows that $N_{C,t+1}$ depends on the consumption shock η_{t+1} , but not on any of the other shocks. $N_{C,t+1}$ enters the SDF by standard consumption-CAPM logic. $N_{PATH,t+1}$ and $N_{RISK,t+1}$ enter the SDF because with non-time separable utility, today’s marginal utility depends not only on contemporaneous consumption growth but also on its entire future path. As can be seen directly from their definitions, $N_{PATH,t+1}$ and $N_{RISK,t+1}$ depend only on future consumption growth and not on contemporaneous consumption growth, so they do not depend on the consumption shock η_{t+1} . It follows that news to the wealth-consumption ratio arises not from changes in contemporaneous consumption but exclusively from changes in wealth, that is, from changes in the pricing of the aggregate consumption claim.

When the EIS is greater than one, which we assume, the intertemporal substitution effect dominates the wealth effect. Positive news to the expected path of consumption growth ($N_{PATH,t+1} > 0$) leads the representative agent to invest more in the aggregate wealth portfolio, increasing its price and hence also the wealth-consumption ratio. $N_{PATH,t+1}$ depends only on the inflation shock, $u_{\Pi,t+1} (= \phi_\pi u_{t+1})$, because by Eq. (14) it is the only shock that provides news about expected consumption growth (shocks to the nominal-real covariance that affect expected consumption growth at t arrive at $t - 1$, so they provide no news). Using Eq. (13), we find that

$$N_{PATH,t+1} = \left(\frac{\rho_c \phi_\pi}{1 - \rho_\pi \kappa_1} + \frac{\rho_c \kappa_1 \xi_\pi}{1 - \rho_\pi \kappa_1} + \xi_c \varphi_t \right) \frac{u_{\Pi,t+1}}{\phi_\pi}. \quad (22)$$

Eq. (22) will be key to understanding the inflation risk premium. The first two terms inside the parentheses in Eq. (22) arise through the term $\rho_c(\pi_t - \mu_\pi)$ in Eq. (14). The first term reflects that news to inflation is news to expected consumption growth. The second term reflects a more indirect channel: in Eq. (13), u_t predicts future inflation π_{t+1} , which in turn predicts expected consumption growth one period later. Both of these effects reflect the unconditional (non-time-varying) effect of inflation on consumption and are stronger when inflation has a larger or longer-lasting impact on future consumption, which is when inflation is more persistent (higher ρ_π), more volatile (higher ϕ_π or ξ_π), or more predictive of consumption (higher ρ_c).

The third term in Eq. (22) reflects the conditional (time-varying) relation between consumption and inflation. It appears because the term $\xi_c \varphi_{t-1} u_t$ in Eq. (14) is part of expected consumption growth. Inflation shocks can be good or bad news for the expected path of consumption. When $\xi_c/\phi_\pi > 0$, a positive inflation shock ($u_{\Pi,t+1} > 0$) is good news for the path of expected consumption if and only if $\varphi_t > 0$, just as in the data and our model calibration below. The volatility of $N_{PATH,t+1}$ increases when ξ_c is higher and when φ_t is more volatile (higher σ_w). The persistence of $N_{PATH,t+1}$ increases when φ_t is more persistent (higher v).

When the representative agent has a preference for early resolution of uncertainty ($\theta < 0$), which we assume, news about risk increasing ($N_{RISK,t+1} > 0$) represents a bad state of nature. Even though the expected path of consumption growth remains unchanged, the consumption claim has become riskier. In equilibrium, today's price of the wealth portfolio must go down for the representative agent to be willing to hold the riskier consumption claim, which results in a lower wealth-consumption ratio and a higher marginal utility. $N_{RISK,t+1}$ depends only on the nominal-real covariance shock w_{t+1} , because $N_{RISK,t+1}$ is a function of conditional second moments only and the nominal-real covariance is the only source of heteroskedasticity.

Because each shock η_{t+1} , $u_{\Pi,t+1}$, and w_{t+1} drives a separate news component of the SDF, we can give an intuitive interpretation for the price of risk for each of them. Innovations to the SDF as a function of shocks and their respective prices of risk are

$$m_{t+1} - E_t[m_{t+1}] = -\lambda_\eta \eta_{t+1} - \lambda_{u,t} u_{\Pi,t+1} - \lambda_{w,t} w_{t+1} - \lambda_{2w} \left(\frac{w_{t+1}^2 - 1}{\sqrt{2}} \right)$$

where the shocks η_{t+1} , $u_{\Pi,t+1}$, w_{t+1} , and $(w_{t+1}^2 - 1)/\sqrt{2}$ are all mean zero and uncorrelated (but not independent) from each other.

Comparing to Eq. (21), and using Eq. (22), we can identify the price of inflation risk $\lambda_{u,t}$ in

$$\lambda_{u,t} u_{\Pi,t+1} = \kappa_1 \left(\gamma - \frac{1}{\psi} \right) N_{PATH,t+1} = \kappa_1 \left(\gamma - \frac{1}{\psi} \right) \left(\frac{\rho_c(\phi_\pi + \kappa_1 \xi_\pi)}{1 - \rho_\pi \kappa_1} + \xi_c \varphi_t \right) \frac{u_{\Pi,t+1}}{\phi_\pi}. \quad (23)$$

The price of inflation risk $\lambda_{u,t}$ is, by definition, the compensation that the representative agent requires to hold one unit of inflation risk ($u_{\Pi,t+1}$ -risk). Our model provides the economic interpretation that inflation risk is priced because inflation shocks provide news to the expected future path of consumption growth, and changes in this path change the marginal utility of the representative agent. In addition, the model provides

an equilibrium explanation for why the price of inflation risk moves over time together with the nominal-real covariance. When $\varphi_t > 0$, the representative agent requires positive compensation for bearing inflation risk, because positive inflation shocks contain good news about future consumption growth. Since higher future consumption growth is a good state of nature, inflation has a positive covariance with marginal utility. On the other hand, if $\varphi_t < 0$, the representative agent is willing to accept lower expected returns to hold inflation risk, because inflation shocks come with bad news about future consumption growth, making inflation a good consumption and marginal utility hedge. Since $\lambda_{u,t}$ is proportional to $N_{PATH,t+1}$, its dependence on parameters is the same as for $N_{PATH,t+1}$, already discussed above. In addition to $N_{PATH,t+1}$, $\lambda_{u,t}$ also depends on preference parameters. In particular, because inflation shocks generate risk over the entire future path of expected consumption growth, a stronger preference for early resolution of uncertainty (higher γ or ψ) magnifies the effect of news about the expected path of consumption on the price of inflation risk.

The same logic that we used for inflation can be used to find the prices of risk for consumption and nominal-real covariance shocks and to interpret them, respectively, as coming from news to contemporaneous consumption, $N_{C,t+1}$, and from news about risk, $N_{RISK,t+1}$. Nominal-real covariance risk (w_t -risk) is priced because of both of the roles that the nominal-real covariance plays in this model (it controls the time-variation in the predictability of consumption with inflation and it also drives the stochastic volatility of expected consumption growth). However, the nominal-real covariance interacts with inflation risk ($u_{\Pi,t}$ -risk) only due to its role in the predictability of consumption with inflation, and not because of its role in inducing stochastic volatility.²⁹

6.4. The quantity of inflation risk and inflation betas

A Campbell-Shiller approximation implies that (log) returns for asset i , $r_{i,t+1}$, can be written as

$$r_{i,t+1} = \kappa_{i,0} + \kappa_{i,1}pd_{i,t+1} - pd_{i,t} + \Delta d_{i,t+1}, \quad (24)$$

where $pd_{i,t}$ is the (log) price-dividend ratio of i , and $\kappa_{i,0}$, $\kappa_{i,1}$ are approximation constants with $\kappa_{i,1}$ generally close to one.³⁰ Decomposing news to returns in a manner similar to the SDF, we have

$$\begin{aligned} (E_{t+1} - E_t)r_{i,t+1} &= \kappa_{i,1}(E_{t+1} - E_t)pd_{i,t+1} + (E_{t+1} - E_t)\Delta d_{i,t+1} \\ &= \kappa_{i,1} \left(N_{PATH,t+1}^i + N_{m,t+1}^i + \frac{1}{2}N_{RISK,t+1}^i \right) + N_{D,t+1}^i, \end{aligned} \quad (25)$$

²⁹Section 2 of the Online Appendix has expressions and a deeper discussion of consumption and nominal-real covariance risks, which we omit here since our main focus is on inflation risk.

³⁰Eq. (24) and all formulas below apply also to the consumption portfolio (with $i = c$) if we replace $pd_{i,t}$, $\Delta d_{i,t}$, $\kappa_{i,0}$, and $\kappa_{i,1}$ with wc_t , Δc_t , κ_0 , and κ_1 , respectively.

where

$$\begin{aligned}
 N_{PATH,t+1}^i &\equiv (E_{t+1} - E_t) \sum_{q=0}^{\infty} \kappa_{i,1}^q E_{t+1+q} [\Delta d_{i,t+2+q}], \\
 N_{m,t+1}^i &\equiv (E_{t+1} - E_t) \sum_{q=0}^{\infty} \kappa_{i,1}^q E_{t+1+q} [m_{t+2+q}], \\
 N_{RISK,t+1}^i &\equiv (E_{t+1} - E_t) \sum_{q=0}^{\infty} \kappa_{i,1}^q Var_{t+1+q} [m_{t+2+q} + \kappa_{i,1} pd_{i,t+1} + \Delta d_{i,t+1}], \\
 N_{D,t+1}^i &\equiv (E_{t+1} - E_t) \Delta d_{i,t+1}.
 \end{aligned}$$

Using the same line of reasoning used to interpret Eq. (21), we find that $N_{PATH,t+1}^i$, which is news to the expected future path of dividends of asset i , depends only on the inflation shock $u_{\Pi,t+1}$; $N_{RISK,t+1}^i$, which is news to the conditional future variance of $m_{t+1} + \kappa_{i,1} pd_{i,t+1} + \Delta d_{i,t+1}$, depends only on the nominal-real covariance shock w_{t+1} ; and $N_{D,t+1}^i$, which is news to contemporaneous dividend growth of asset i , depends only on the consumption shock η_{t+1} . In addition to these three news components, news to the returns of asset i also depends on $N_{m,t+1}^i$, which is news to the expected path of the SDF (news to expected discount rates). This news depends on the inflation shock, $u_{\Pi,t+1}$, and the nominal-real covariance shock, w_{t+1} . In news for returns, $N_{m,t+1}^i$ enters because positive news to the future expected path of the SDF ($N_{m,t+1}^i > 0$) is positive news for today's returns: future marginal utility is expected to be higher, so the same stream of dividends is more valuable, pushing up the price-dividend ratio and realized returns.³¹

The quantity of inflation risk of asset i is given by its inflation beta. As in our empirical analysis, the inflation beta $\beta_{\Pi,i,t}$ is defined as the coefficient in a regression of returns, $r_{i,t+1}$, on inflation shocks, $u_{\Pi,t+1}$, conditional on time t information. Using Eq. (25), we have

$$\beta_{\Pi,i,t} \equiv \frac{Cov_t(u_{\Pi,t+1}, r_{i,t+1})}{Var_t(u_{t+1})} = \kappa_{i,1} Cov_t(u_{\Pi,t+1}, N_{m,t+1}^i) + \kappa_{i,1} Cov_t(u_{\Pi,t+1}, N_{PATH,t+1}^i). \quad (26)$$

Intuitively, when news to the path of either dividends, $N_{PATH,t+1}^i$, or discount rates, $N_{m,t+1}^i$, covary more with inflation shocks, the price of asset i is more exposed to inflation and the quantity of inflation risk increases. The reason that news to only expected dividends and expected discount rates enter $\beta_{\Pi,i,t}$ is that inflation shocks only affect the expected paths of dividends and discount rates, but not their higher (conditional) moments. $N_{RISK,t+1}^i$ and $N_{D,t+1}^i$ are uncorrelated to inflation shocks $u_{\Pi,t+1}$, so they do not contribute to $\beta_{\Pi,i,t}$.

³¹Given that the risk-free rate is defined by $r_{f,t} = -E_t[m_{t+1}] - \frac{1}{2}Var_t(m_{t+1})$, this news component is related to, but not the same as, news to the path of risk-free rate. $N_{m,t+1}^i$ is different for different assets because the expected path of the SDF is discounted by the asset-specific constant $\kappa_{i,1}$.

Computing the covariances in Eq. (26), we find that³²

$$\beta_{\Pi,i,t} = -\kappa_{i,1} \underbrace{\frac{\phi_\pi}{\psi} \left(\frac{\rho_c (\kappa_{i,1} \xi_\pi + \phi_\pi)}{1 - \kappa_{i,1} \rho_\pi} + \xi_c \varphi_t \right)}_{Cov_t(u_{\Pi,t+1}, N_{m,t+1}^i)} + \kappa_{i,1} \phi_\pi \underbrace{\left(\frac{\rho_i (\kappa_{i,1} \xi_\pi + \phi_\pi)}{1 - \kappa_{i,1} \rho_\pi} + \xi_i \varphi_t \right)}_{Cov_t(u_{\Pi,t+1}, N_{P_{ATH,t+1}}^i)}. \quad (27)$$

The intuition behind the dependence of $\beta_{\Pi,i,t}$ on preference parameters and parameters of the exogenous processes in Eqs. (13) to (16) is similar to the intuition for the terms in Eq. (23) discussed earlier. More importantly, all time-variation in $\beta_{\Pi,i,t}$ comes from φ_t .

6.5. Risk premia of inflation beta-sorted portfolios

The risk premium on asset i is

$$-Cov_t(m_{t+1}, r_{i,t+1}) = \lambda_{u,t} \beta_{\Pi,i,t} + \lambda_{w,t} \beta_{w,i,t} + \lambda_{2w} \beta_{2w,i} + \lambda_\eta \beta_{\eta,i}, \quad (28)$$

where the inflation risk premium component is

$$\lambda_{u,t} \beta_{\Pi,i,t} = \underbrace{\kappa_1 \left(\gamma - \frac{1}{\psi} \right) \left(\frac{\phi_\pi + \kappa_1 \xi_\pi}{1 - \kappa_1 \rho_\pi} \rho_c + \xi_c \varphi_t \right)}_{\lambda_{u,t}} \underbrace{\kappa_{i,1} \phi_\pi \left(\frac{\phi_\pi + \kappa_{i,1} \xi_\pi}{1 - \kappa_{i,1} \rho_\pi} \left(\rho_i - \frac{\rho_c}{\psi} \right) + \left(\xi_i - \frac{\xi_c}{\psi} \right) \varphi_t \right)}_{\beta_{\Pi,i,t}}. \quad (29)$$

Because the price of inflation risk, $\lambda_{u,t}$, and the quantity of inflation risk, $\beta_{\Pi,i,t}$, are each linear functions of the nominal-real covariance, the inflation risk premium is quadratic in φ_t . The inflation risk premium is positive when $\lambda_{u,t}$ and $\beta_{\Pi,i,t}$ have the same sign. If $\lambda_{u,t}$ and $\beta_{\Pi,i,t}$ are both negative, the returns of asset i comove negatively with inflation shocks that contain bad news for future growth. If $\lambda_{u,t}$ and $\beta_{\Pi,i,t}$ are both positive, the returns of asset i comove positively with inflation shocks that contain good news. In both cases, the payoff of asset i is low in bad times and high in good times, exacerbating consumption risk. In order for the representative agent to be willing to hold asset i in equilibrium, the inflation risk premium must be positive. If $\lambda_{u,t}$ and $\beta_{\Pi,i,t}$ have different signs, the opposite is true. Asset i has high payoff in bad states and low payoff in good states, and the representative agent is willing to accept a negative premium to hold this asset because it provides a consumption hedge.

To derive the model-implied inflation risk premium from the cross-section of stock returns, we model two assets that represent the High (H) and Low (L) inflation beta portfolios from our empirical analysis. The High-minus-Low inflation beta, $\beta_{u,HL,t}$, is simply the difference $\beta_{u,Ht} - \beta_{u,Lt}$, because betas are linear in returns (see Eq. (26)). If $\kappa_{H,1} \approx \kappa_{L,1}$, as is the case in any reasonable calibration of the model, a good

³²Using changes in inflation ($\pi_{t+1} - \pi_t$) or inflation levels (π_{t+1}) instead of inflation innovations in Eq. (26) also give $\beta_{\Pi,i,t}$ that are affine in φ_t and differ from each other only in scaling. Thus, as in the data (see Section 8 and Table OA.13 in the Online Appendix), these different ways to measure inflation risk capture the same time-variation driven by the nominal-real covariance.

approximation for $\beta_{u,HL,t}$ is

$$\begin{aligned}\beta_{u,HL,t} &= Cov_t(u_{\Pi,t+1}, N_{PATH,t+1}^H) - Cov_t(u_{\Pi,t+1}, N_{PATH,t+1}^L) \\ &= \frac{\phi_\pi + \xi_\pi}{1 - \rho_\pi} (\rho_H - \rho_L) + (\xi_H - \xi_L) \varphi_t\end{aligned}\quad (30)$$

with associated inflation risk premium $\lambda_{u,t}\beta_{u,HL,t} = \lambda_{u,t}(\beta_{u,Ht} - \beta_{u,Lt})$. The High-minus-Low inflation beta is determined by how differently inflation shocks covary with news to the expected path of dividends for asset H versus asset L . The assumption that $\kappa_{H,1} \approx \kappa_{L,1}$ makes the discount-rate components $\kappa_{H,1}Cov_t(u_{\Pi,t+1}, N_{m,t+1}^H)$ and $\kappa_{L,1}Cov_t(u_{\Pi,t+1}, N_{m,t+1}^L)$ of $\beta_{u,H,t}$ and $\beta_{u,L,t}$ identical, so they cancel out when taking the difference.

7. Calibration

Table 11 shows the calibrated parameters of the model and Table 12 shows the empirical and model-implied moments that we target.³³ The model has 26 parameters (three preference parameters; 11 parameters for the joint process of inflation, consumption, and the nominal-real covariance; and four parameters for each of the three stocks we model) and we consider 37 moments from Tables 1 to 7.

For preferences, we use $\delta = 0.99$, $\psi = 2$, and $\gamma = 15$, which are within the range of values commonly used in similar studies.³⁴ We calibrate the exogenous processes for inflation, consumption, and the nominal-real covariance in Eqs. (13) to (15) to match the summary statistics of Table 1 and the predictability regressions of Table 2 for $K = 12$. The model-implied moments for π_t and Δc_t match the data well. The one exception is the standard deviation of Δc_t , which at 2.49% is somewhat higher than the 1.13% in the data. This value is within the range of values used in other studies, however.³⁵ The relatively high volatility of Δc_t is needed to match the relatively high volatility and persistence of the nominal-real covariance under the parsimonious $AR(1)$ specification for φ_t . We prioritize matching moments of the nominal-real covariance, because that is the mechanism we propose to explain the time-variation in the inflation risk premium. In any case, by standard consumption-CAPM logic, whether the standard deviation of consumption is 1.13% or 2.49% has little impact on asset prices, even with $\gamma = 15$. The $AR(1)$ coefficient for Δc_t is essentially zero in the model and small (point estimate -16.84%) and not statistically significant in the data at the 1%-level. The moments concerning the inflation shock $u_{\Pi,t}$ also match the summary statistics of Table 1 well. In the model, the mean and $AR(1)$ coefficient of the inflation shock $u_{\Pi,t}$ are equal to zero by assumption (for any choice of parameters), as is the correlation between $u_{\Pi,t}$ and Δc_t . In the data, the values of these three moments are close to zero and not statistically significant at the 5%-level.³⁶

³³We solve for all moments in closed form. Formulas are in Section 2 of the Online Appendix.

³⁴For example, in a closely related model, Bansal and Shaliastovich (2012) use $\psi = 1.81$ and $\gamma = 20.9$.

³⁵For example, Bansal and Yaron (2004) use a standard deviation of Δc_t of 2.93%.

³⁶The correlation between consumption and inflation is negative in monthly data but usually positive for samples that use

From Panel A of Table 2, we match the coefficients $d_0^{u,12}$ and $d_1^{u,12}$ together with the low R^2 , which guarantees that our model has the same low degree of unconditional predictability as the data. We also match the coefficients $d_0^{c,12}$ and $d_1^{c,12}$ and the higher R^2 of Panel B, which shows that the model has the same strong conditional predictability as the data. As mentioned before, we design our exogenous processes so that $d_0^{c,K} = 0$ and $d_1^{c,K} = 1$ for any calibration of the model and any K , consistent with the failure to reject the hypothesis that $d_0^{c,K} = 0$ and $d_1^{c,K} = 1$ in the data. In addition to the moments in Table 2, we match the volatility and $AR(1)$ coefficient of the nominal-real covariance, as measured by the first-stage coefficient $b_t^{12} = NRC_t^C$ from Eq. (3).

From Table 3, we match the first and second moments of returns and inflation betas for the two inflation portfolios with the highest and lowest inflation betas. Matching the returns of the High and Low portfolios ensures that the size of the inflation risk premium is consistent with the data. Matching the first moment of betas gives the appropriate mean quantity of inflation risk, while matching the second moment of betas gives the same degree of time-variation in the quantity of risk as in the data. Matching the mean returns and inflation betas of the High and Low portfolios automatically implies that we match the mean return and inflation beta of the High-minus-Low inflation beta-sorted portfolio, IP_t^{HL} , which for both returns and betas is simply their difference. Together, matching the inflation beta and returns for IP_t^{HL} implies that the price of inflation risk in the model is consistent with the data. The model also matches two patterns from Fig. 3: the correlation between $\beta_{\Pi,H,t}$ and $\beta_{\Pi,L,t}$ is high (0.93 in the data and 1.0 in the model), and the volatility of the inflation beta of the IP_t^{HL} portfolio is around one-third of the volatility of the inflation betas of either the High or Low portfolio. From Table 5, we match the coefficient L_{NRC} for the IP_t^{HL} portfolio at the one-year horizon ($K = 12$), which shows that the model is capable of reproducing the empirical predictability in the price of inflation risk with the nominal-real covariance.

When we calibrate the aggregate market portfolio to have the same quantity of inflation risk as in the data, the model-implied aggregate stock market risk premium contains an inflation component that has the same magnitude and time-variation properties as the data. To calibrate the quantity of inflation risk of the market, we match the mean and standard deviation of the market's inflation beta, $\beta_{\Pi,m,t}$, as well as its correlation with $\beta_{\Pi,H,t}$ and $\beta_{\Pi,L,t}$. Once we have a calibrated $\beta_{\Pi,m,t}$, we use the NRC_t^C implied by the model to run regression (9). We find that the key coefficient on the interaction term $\beta_{\Pi,m,t} \times NRC_t^C$ is 2.69 in the model, which is close to the empirical value of 2.85 in Table 7.

Last, the calibration we propose generates a timing premium for the early resolution of uncertainty that is considerably lower than in the long-run risk model, thus mitigating the criticism in Epstein, Farhi, and Strzalecki (2014). The timing premium is defined as the maximum share of consumption that the

quarterly or annual data. Our model allows for parameter combinations that give either sign for this correlation, although what matters most for asset pricing is the predictive relation between inflation and consumption rather than the contemporaneous correlation.

representative agent is willing to forgo every period to resolve all uncertainty at $t = 1$. The first column in Table OA.8 of the Online Appendix shows that for our calibration the timing premium is 14.8%, less than half the 31% timing premium in the baseline calibration of the long-run risk model.³⁷ The last two columns of Table OA.8 show the low timing premium in our model arises because inflation is not as persistent as long-run risk, and not because of the persistence or volatility characteristics of the nominal-real covariance process φ_t .

8. Robustness checks

We discuss a range of robustness checks for which we report results in Section 3 of the Online Appendix.

Benchmark asset pricing factors and predictors

We test whether the inflation risk premium and its variation with the nominal-real covariance are robust to controlling for the factors in the CAPM, Fama-French three- and five-factor models (FF3M, FF5M), and the Carhart four-factor model (FFCM). First, we include these factors as control variables when estimating stocks' ex ante inflation exposures using Eq. (4). Table OA.10 presents the predictive regressions for these alternative estimates of the inflation risk premium. We see that both the unconditional inflation risk premium (as measured by the intercept of the predictive regression) as well as the coefficient on the nominal-real covariance are robust, both in magnitude and significance.

Next, we control for these benchmark factors ex post in the predictive regression in Eq. (7). Table OA.11 shows that the unconditionally negative inflation risk premium is captured in the FFCM and FF5M. In contrast, we find that the coefficient on the nominal-real covariance remains economically large and at least marginally significant in all models. We thus conclude that time-variation in the inflation risk premium is not due to exposure to benchmark factors, neither ex ante nor ex post.

Table OA.12 runs the predictive regressions of Eq. (7) controlling for standard Intertemporal CAPM predictors: dividend yield (DY), default spread (DS), and term spread (TS) as well as the consumption-wealth ratio (CAY). We find that the inflation risk premium varies significantly with the nominal-real covariance even after controlling for these benchmark predictors and conclude that the nominal-real covariance represents a new source of conditioning information.

Alternative measures of inflation and sorting procedures

Table OA.13 analyzes alternative measures of inflation risk. We consider raw inflation; an AR(1)-innovation in inflation; the monthly change in annual inflation; and, real-time vintage CPI inflation. The latter test represents a truly out-of-sample exercise, as we skip a month after portfolio formation to take

³⁷For comparison, an i.i.d. model has a premium of 10%, the long-run risk model without stochastic volatility has a premium of 29%, and the long-run risk model with stochastic volatility and preference parameters equal to the ones we use ($\gamma = 15$ and $\psi = 2$) has a premium of 44% (Table OA.9 in the Online Appendix).

into account the reporting delay in inflation data. Focusing on predictability of the High-minus-Low inflation beta-sorted portfolio return, we find that our benchmark ARMA(1,1) specification represents a lower bound. Two alternative measures provide particularly strong evidence for time-variation in the inflation risk premium with the nominal-real covariance. Using raw inflation or real-time inflation, the R^2 more than doubles to about 15% with a coefficient estimate that is over 7% and a t -statistic over three. Moreover, the unconditional inflation risk premium is negative in all cases.

Table OA.14 shows that our results are also robust to changes in the method of estimating inflation betas. We again focus on High-minus-Low portfolios that are constructed by sorting stocks on inflation betas estimated using (i) weighted least squares plus shrinkage (WLS) or ordinary least squares (OLS); (ii) a sort where each portfolio comes from a double sort on inflation beta and size (Size) or a single sort on beta (No Size), and (iii) ARMA(1,1) innovations in inflation (ARMA) or raw inflation (Inflation). In all cases, the coefficient on the nominal-real covariance is positive, economically large, and at least marginally significant. In three out of eight cases, the t -statistic on the nominal-real covariance is well above three. Finally, Table OA.15 replicates our main analysis using quarterly data. We find that the High-minus-Low inflation risk premium (estimated by sorting stocks on exposure in quarterly returns to ARMA(1,1)-innovations in quarterly inflation) is also predictable at this alternative frequency by the nominal-real covariance (estimated as the rolling coefficient in a regression of quarterly consumption growth on lagged quarterly inflation).

Time-varying inflation betas

In the model of Section 6, inflation betas are linear in the nominal-real covariance. To check if our results are robust to this specification of inflation betas, we adjust Eq. (4) as follows:

$$\left(\widehat{\alpha}_{p,t}, \widehat{\beta}_{\Pi,p,t}\right) = \arg \min_{\alpha_{p,t}, \beta_{\Pi,p,t}^0, \beta_{\Pi,p,t}^1} \sum_{\tau=1}^t K(\tau) \left(R_{i,\tau} - \alpha_{p,t} - (\beta_{\Pi,p,t}^0 + \beta_{\Pi,p,t}^1 NRC_{t-1}^C) u_{\Pi,\tau}\right)^2. \quad (31)$$

We find that the linear conditional inflation betas for each decile portfolio p , $\beta_{\Pi,p,t} = \beta_{\Pi,p,t}^0 + \beta_{\Pi,p,t}^1 NRC_{t-1}^C$, are highly correlated to the betas we use in Section 5.1 (at a correlation of 0.79 on average). Moreover, Table OA.16 shows that the results from our pooled regressions are, if anything, stronger than what we report already in Table 6. Both unconditional variation in inflation betas across portfolios as well as time-variation in the average inflation beta interact with the nominal-real covariance to generate large and significant time-variation in the returns of inflation beta-sorted portfolios.

Monetary policy (risk)

We perform two tests to assess the impact of monetary policy on our results. First, Campbell, Pflueger, and Viceira (2015) link time-variation in the nominal-real covariance to three periods: the period of rising inflation up to the first quarter of 1977, the inflation-fighting period under former Federal Reserve Board chairmen Volcker and Greenspan up to the last quarter of 2000, and the recent period of increased central bank transparency, gradualism, and renewed attention to output stabilization. Table OA.17 shows that our

conclusions on the interaction between inflation beta and the nominal-real covariance are robust to including in the pooled regression a dummy for each of these subperiods. This finding highlights that our empirical evidence on the inflation risk premium is richer than a sample split into long-lasting monetary policy regimes.

Next, we analyze whether monetary policy risk can drive our results. Table OA.18 presents the exposures of our inflation-beta sorted portfolios to Romer and Romer (2004) monetary policy shocks and changes in the federal funds rate (also used in, e.g., Neuhierl and Weber, 2017; Velikov, 2017). Whereas exposure to changes in the federal funds rate is typically negative and significant, there is no clear pattern from High to Low inflation beta, neither in the full sample nor within any of the three monetary policy regimes. As a result, exposures for the High-minus-Low portfolio (IP_t^{HL}) are small and insignificant. Similar results obtain for the Romer and Romer (2004) shocks. We conclude that monetary policy risk is unlikely to explain the variation in the inflation risk premium that we uncover. This conclusion is consistent with Velikov (2017), who argues that exposure to monetary policy shocks is transient and best observed in high frequency stock returns. In contrast, our exposures to inflation risk are persistent.

Bond returns

Although our focus is on the inflation risk premium in the stock market, we can examine whether similar results apply in the bond market. We use returns on constant maturity Treasury bonds (with maturities from one to 30 years) in excess of the one-month T-bill return. In Panel A of Table OA.19, we see that average excess returns are increasing in maturity, whereas inflation beta is decreasing in maturity. The ten-minus-one-year maturity difference in average excess return is 1.44%. Dividing this number by the ten-minus-one-year difference in inflation beta gives an estimate of the unconditional price of inflation risk: $1.44/(-1.16) = -1.24\%$. This price is close to what we estimate in the stock market in Table 3: $-4.23/3.00 = -1.41\%$.

Having said that, the evidence in Panel B shows that the nominal-real covariance is not a strong predictor of excess bond returns in the time-series. We report coefficient estimates from a regression of annual excess bond returns on the lagged nominal-real covariance. Depending on maturity, a one standard deviation increase in the nominal-real covariance predicts an increase in excess bond return between -0.81% and -0.41%. Although non-negligible economically, the estimated coefficient is only statistically significant at the 1%-level for the one-year maturity bond and at the 10%-level for the two-year maturity bond.

9. Conclusion

We show that inflation risk is priced in the stock market in a time-varying way. We study a key driver of this time-variation: the nominal-real covariance, which is the time-varying relation between inflation and real future consumption growth.

Our main empirical finding is that the inflation risk premium increases by a large and significant 5.3% for a standard deviation increase in the nominal-real covariance. Given an unconditional inflation risk premium

of -4.2%, this result implies that the inflation risk premium can change sign. The economic intuition is that when inflation predicts consumption growth with a negative sign as it did, for example, in the 1970s, stocks with a high inflation beta are attractive as a hedge against bad states of nature and thus command relatively low expected returns. When inflation predicts consumption growth with a positive sign, as it did since the early 2000s, high inflation beta stocks are risky (as they covary negatively with marginal utility) and the inflation risk premium switches sign.

By decomposing the time-variation in the inflation risk premium into variation in the price and quantity of risk, we obtain a rich set of results that are new to the literature. The time-variation in the inflation risk premium between high and low inflation beta stocks is driven by time-variation in the price of inflation risk. Quantities do matter for the average inflation risk premium across stocks, however. The reason is that time-variation in the quantity of inflation risk is largely common across stocks. Consequently, the market risk premium contains an inflation component that is time-varying due to both the price and quantity of risk.

We develop an equilibrium model that builds on the empirical observation that inflation today predicts real consumption growth in the future in a time-varying way. In the model, this time-variation is controlled by the nominal-real covariance, which allows us to closely match the observed dynamics in the price and quantities of inflation risk in both the cross-section of stocks and the aggregate market. Our theoretical contribution is to extend the long-run risk framework by using an observable state variable that is related to future consumption growth in a time-varying way.

References

- Ang, A., Bekaert, G., Wei, M., 2007. Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics* 54, 1163–1212.
- Ang, A., Bekaert, G., Wei, M., 2008. The term structure of real rates and expected inflation. *Journal of Finance* 63, 797–849.
- Ang, A., Brière, M., Signori, O., 2012. Inflation and individual equities. *Financial Analysts Journal* 68, 36–55.
- Ang, A., Piazzesi, M., 2003. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics* 50, 745–787.
- Ang, A., Piazzesi, M., Wei, M., 2006. What does the yield curve tell us about GDP growth? *Journal of Econometrics* 131, 359–403.
- Bansal, R., Shaliastovich, I., 2012. A long-run risks explanation of predictability puzzles in bond and currency markets. *Review of Financial Studies* 26, 1–33.
- Bansal, R., Yaron, A., 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59, 1481–1509.

- Bekaert, G., Wang, X., 2010. Inflation risk and the inflation risk premium. *Economic Policy* 25, 755–806.
- Bilal, M. A., 2017. Zeroing in: Asset pricing at the zero lower bound. Unpublished working paper. New York University.
- Boons, M., 2016. State variables, macroeconomic activity, and the cross section of individual stocks. *Journal of Financial Economics* 119, 489–511.
- Boudoukh, J., Richardson, M., Whitelaw, R. F., 1994. Industry returns and the Fisher effect. *Journal of Finance* 49, 1595–1615.
- Breeden, D. T., Gibbons, M. R., Litzenberger, R. H., 1989. Empirical tests of the consumption-oriented CAPM. *Journal of Finance* 44, 231–262.
- Bretschler, L., Hsu, A., Tamoni, A., 2017. Level and volatility shocks to fiscal policy: Term structure implications. Unpublished working paper. London Business School.
- Campbell, J. Y., 1991. A variance decomposition for stock return. *The Economic Journal* 101, 157–179.
- Campbell, J. Y., Ammer, J., 1993. What moves the stock and bond markets? A variance decomposition for long-term asset returns. *The Journal of Finance* 48, 3–37.
- Campbell, J. Y., Giglio, S., Polk, C., Turley, R., 2018. An intertemporal CAPM with stochastic volatility. *Journal of Financial Economics* 128, 207–233.
- Campbell, J. Y., Pflueger, C., Viceira, L. M., 2015. Monetary policy drivers of bond and equity risks. Unpublished working paper. Harvard University.
- Campbell, J. Y., Sunderam, A., Viceira, L. M., 2017. Inflation bets or deflation hedges? The changing risks of nominal bonds. *Critical Finance Review* 6, 263–301.
- Campbell, J. Y., Viceira, L. M., 2002. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press, New York.
- Campbell, J. Y., Vuolteenaho, T., 2004. Bad beta, good beta. *American Economic Review* 94, 1249–1275.
- Carhart, M. M., 1997. On persistence in mutual fund performance. *Journal of Finance* 52, 57–82.
- Chen, H., Dou, W. W., Kogan, L., 2015. Measuring the “dark matter” in asset pricing models. Unpublished working paper. Massachusetts Institute of Technology.
- Chen, N.-F., Roll, R., Ross, S. A., 1986. Economic forces and the stock market. *Journal of Business* 59, 383–403.

- Chen, R.-R., Liu, B., Cheng, X., 2010. Pricing the term structure of inflation risk premia: Theory and evidence from TIPS. *Journal of Empirical Finance* 17, 702–721.
- Chordia, T., Goyal, A., Shanken, J. A., 2015. Cross-sectional asset pricing with individual stocks: betas versus characteristics. Unpublished working paper. Emory University, University of Lusanne.
- Cochrane, J. H., 2005. *Asset Pricing*, vol. 1. Princeton University Press, Princeton, NJ.
- Cochrane, J. H., 2017. Macro-finance. *Review of Finance* 21, 945–985.
- Cosemans, M., Frehen, R., Schotman, P. C., Bauer, R., 2016. Estimating security betas using prior information based on firm fundamentals. *Review of Financial Studies* 29, 1072–1112.
- Dai, Q., Singleton, K. J., Yang, W., 2007. Regime shifts in a dynamic term structure model of US Treasury bond yields. *Review of Financial Studies* 20, 1669–1706.
- David, A., Veronesi, P., 2013. What ties return volatilities to price valuations and fundamentals? *Journal of Political Economy* 121, 682–746.
- Driscoll, J. C., Kraay, A. C., 1998. Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics* 80, 549–560.
- Duffee, G., 2018. Expected inflation, real rates, and stock-bond comovement. Unpublished working paper. Johns Hopkins University.
- Elton, E. J., Gruber, M. J., Urich, T. J., 1978. Are betas best? *Journal of Finance* 33, 1375–1384.
- Engsted, T., Pedersen, T. Q., Tanggaard, C., 2012. Pitfalls in VAR based return decompositions: A clarification. *Journal of Banking & Finance* 36, 1255–1265.
- Epstein, L. G., Farhi, E., Strzalecki, T., 2014. How much would you pay to resolve long-run risk? *American Economic Review* 104, 2680–2697.
- Fama, E. F., 1976. *Foundations of Finance: Portfolio Decisions and Securities Prices*. Basic Books (AZ).
- Fama, E. F., French, K. R., 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3–56.
- Fama, E. F., French, K. R., 2008. Dissecting anomalies. *Journal of Finance* 63, 1653–1678.
- Fama, E. F., French, K. R., 2015. A five-factor asset pricing model. *Journal of Financial Economics* 116, 1–22.

- Fama, E. F., Gibbons, M. R., 1984. A comparison of inflation forecasts. *Journal of Monetary Economics* 13, 327–348.
- Fama, E. F., MacBeth, J. D., 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81, 607–636.
- Ferson, W. E., Harvey, C. R., 1991. The variation of economic risk premiums. *Journal of Political Economy* 99, 385–415.
- Fuhrer, J., Moore, G., 1995. Inflation persistence. *Quarterly Journal of Economics* 110, 127–159.
- Gabaix, X., 2008. Variable rare disasters: A tractable theory of ten puzzles in macro-finance. *American Economic Review* 98, 64–67.
- Gorodnichenko, Y., Weber, M., 2016. Are sticky prices costly? Evidence from the stock market. *American Economic Review* 106, 165–199.
- Gourio, F., Ngo, P., 2016. Risk premia at the ZLB: a macroeconomic interpretation. Unpublished working paper. Federal Reserve Bank of Chicago.
- Gürkaynak, R. S., Sack, B., Wright, J. H., 2010. The TIPS yield curve and inflation compensation. *American Economic Journal: Macroeconomics* 2, 70–92.
- Harvey, C. R., Liu, Y., Zhu, H., 2016. and the cross-section of expected returns. *Review of Financial Studies* 29, 5–68.
- Haubrich, J., Pennacchi, G., Ritchken, P., 2012. Inflation expectations, real rates, and risk premia: Evidence from inflation swaps. *Review of Financial Studies* 25, 1588–1629.
- Kan, R., Zhang, C., 1999. Two-pass tests of asset pricing models with useless factors. *Journal of Finance* 54, 203–235.
- Kang, J., Pflueger, C. E., 2015. Inflation risk in corporate bonds. *Journal of Finance* 70, 115–162.
- Koijen, R. S., Lustig, H., Van Nieuwerburgh, S., 2017. The cross-section and time series of stock and bond returns. *Journal of Monetary Economics* 88, 50–69.
- Le, A., Singleton, K. J., 2010. An equilibrium term structure model with recursive preferences. *American Economic Review* 100, 557–561.
- Le, A., Singleton, K. J., Dai, Q., 2010. Discrete-time affine q term structure models with generalized market prices of risk. *Review of Financial Studies* 23, 2184–2227.

- Lettau, M., Ludvigson, S., 2001a. Consumption, aggregate wealth, and expected stock returns. *Journal of Finance* 56, 815–849.
- Lettau, M., Ludvigson, S., 2001b. Resurrecting the (C) CAPM: A cross-sectional test when risk premia are time-varying. *Journal of Political Economy* 109, 1238–1287.
- Levi, Y., Welch, I., 2017. Best practice for cost-of-capital estimates. *Journal of Financial and Quantitative Analysis* 52, 427–463.
- Lintner, J., 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* 47, 13–37.
- Maio, P., Santa-Clara, P., 2012. Multifactor models and their consistency with the ICAPM. *Journal of Financial Economics* 106, 586–613.
- Modigliani, F., Cohn, R. A., 1979. Inflation, rational valuation and the market. *Financial Analysts Journal* 35, 24–44.
- Mossin, J., 1966. Equilibrium in a capital asset market. *Econometrica: Journal of the Econometric Society* 34, 768–783.
- Neuhierl, A., Weber, M., 2017. Monetary policy slope and the stock market. Unpublished working paper. University of Chicago Booth School of Business.
- Patton, A. J., Timmermann, A., 2010. Monotonicity in asset returns: New tests with applications to the term structure, the CAPM, and portfolio sorts. *Journal of Financial Economics* 98, 605–625.
- Piazzesi, M., Schneider, M., 2006. Equilibrium yield curves. *NBER Macroeconomics Annual* 21, 389–472.
- Romer, C. D., Romer, D. H., 2004. A new measure of monetary shocks: Derivation and implications. *American Economic Review* 94, 1055–1084.
- Sharpe, W. F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19, 425–442.
- Song, D., 2017. Bond market exposures to macroeconomic and monetary policy risks. *Review of Financial Studies* 30, 2761–2817.
- Stock, J. H., Watson, M. W., 2005. Implications of dynamic factor models for VAR analysis. NBER Working Paper No. 11467.
- Vasicek, O. A., 1973. A note on using cross-sectional information in Bayesian estimation of security betas. *Journal of Finance* 28, 1233–1239.

- Vassalou, M., 2000. Exchange rate and foreign inflation risk premiums in global equity returns. *Journal of International Money and Finance* 19, 433–470.
- Velikov, M., 2017. FOMC announcements and predictable returns. Unpublished working paper. Federal Reserve Bank of Richmond.
- Wachter, J. A., 2006. A consumption-based model of the term structure of interest rates. *Journal of Financial Economics* 79, 365–399.
- Weber, M., 2015. Nominal rigidities and asset pricing. Unpublished working paper. University of Chicago Booth School of Business.
- Welch, I., Goyal, A., 2007. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 21, 1455–1508.

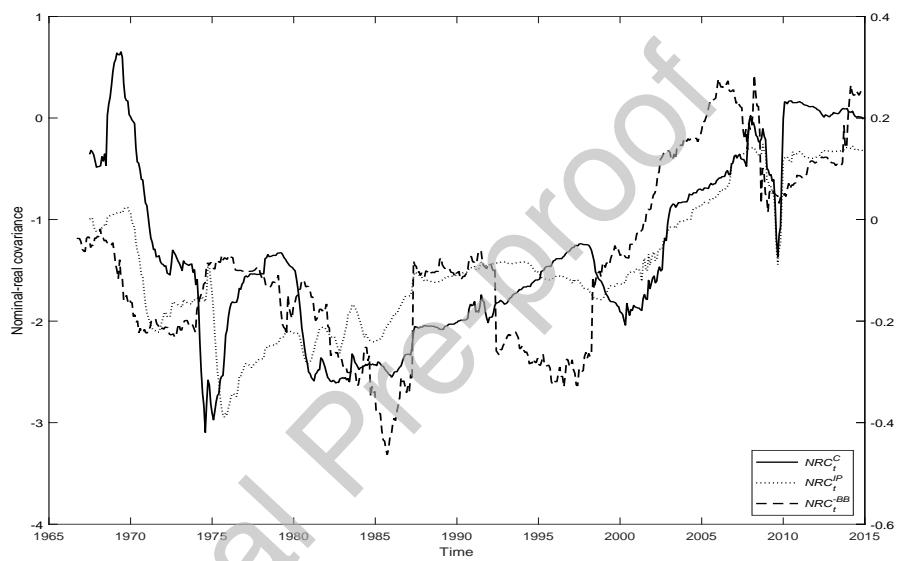


Fig. 1. Alternative measures of the nominal-real covariance. We present three measures of the nominal-real covariance. The first is the coefficient in a rolling regression of 12-month consumption growth on lagged inflation (NRC_t^C). The second measure substitutes industrial production for consumption (NRC_t^{IP}). These measures are plotted on the left axis and NRC_t^{IP} is divided by four to preserve scaling. The third measure is the negative of the stock market beta of the 10-year constant maturity Treasury bond (NRC_t^{BB} , right axis).




Fig. 2. Post-ranking inflation betas. We plot the post-ranking inflation beta (estimated over the full sample from July 1967 to December 2014) for the High-minus-Low portfolio one month, one year, two years, five years, and ten years after sorting. The legend includes the estimated inflation beta and *, **, *** indicate statistical significance at the 10%, 5%, and 1%-level using Newey-West(1) standard errors.

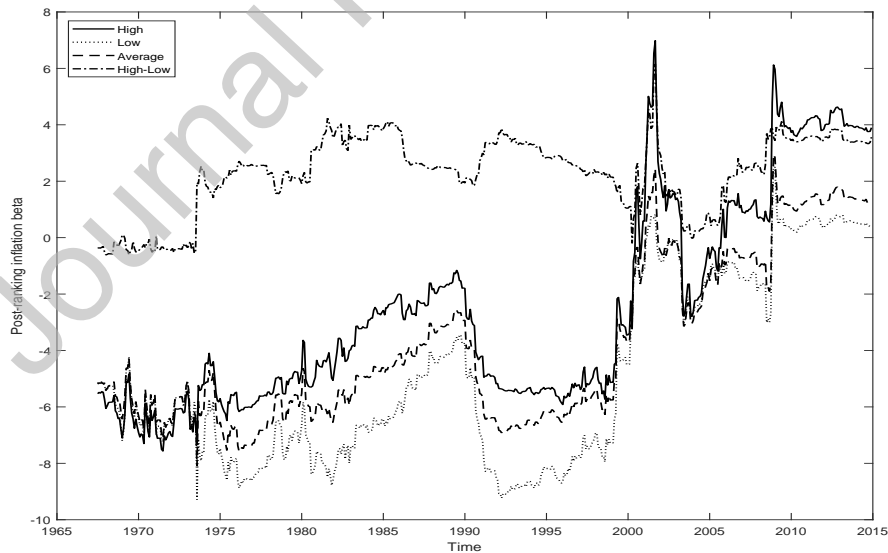


Fig. 3. Post-ranking inflation betas over time. We present the time-series of inflation betas for the High, Low, and High-minus-Low inflation beta decile portfolio as well as the average inflation beta over the ten portfolios.

Table 1

Descriptive statistics.

This table presents descriptive statistics in annualized percentages for CPI inflation (Π_t), ARMA(1,1)-innovations in inflation ($u_{\Pi,t}$), consumption growth (ΔC_t), the aggregate stock market excess return ($R_{m,t}$), and the one-month T-bill return ($R_{f,t}$). AR(1) is the first-order autocorrelation coefficient. The sample period is from July 1962 to December 2014.

	Π_t	$u_{\Pi,t}$	ΔC_t	$R_{m,t}$	$R_{f,t}$
Mean	3.93	0.11	1.97	6.45	4.85
St. dev.	1.11	0.85	1.13	15.44	0.89
AR(1)	62.47	15.35	-16.84	7.47	96.90
Correlations					
Π_t	1	79.03	-17.75	-13.35	50.35
$u_{\Pi,t}$		1	-17.41	-11.72	15.35
ΔC_t			1	17.45	1.26
$R_{m,t}$				1	-8.68
$R_{f,t}$					1

Table 2

Unconditional and conditional predictive regressions of consumption growth on inflation.

In this table we present predictive regressions of consumption growth, $\Delta C_{t+1:t+K}$ over horizons $K = 1, 3, 6, 12$, on lagged inflation, Π_t , from July 1967 to December 2014. Panel A presents an unconditional regression: $\Delta C_{t+1:t+K} = d_0^{u,K} + d_1^{u,K} \Pi_t + e_{t+1:t+K}$. Panel B presents a two-stage conditional regression. In the first stage, we regress consumption growth on lagged inflation over a backward-looking rolling window: $\Delta C_{s+1:s+K} = a_t^K + b_t^K \Pi_s + e_{s+1:s+K}$, for $s = 1, \dots, t - K$. Combining the estimated coefficients with inflation in month t , we then predict future consumption growth in the full time series: $\Delta C_{t+1:t+K} = d_0^{c,K} + d_1^{c,K} (\widehat{a}_t^K + \widehat{b}_t^K \Pi_t) + e_{t+1:t+K}$. For this conditional setup, we present also an out-of-sample R^2 , which compares the performance of the conditional model that includes inflation with a conditional model that includes only a constant. To be precise, $R^2 - OOS = 1 - \frac{Var(\Delta C_{t+1:t+K} - (\widehat{a}_t^K + \widehat{b}_t^K \Pi_t))}{Var(\Delta C_{t+1:t+K} - \widehat{a}_t^{*,K})}$, where $\widehat{a}_t^{*,K}$ is estimated through a backward-looking regression of consumption growth on a constant at each point in time (thus setting $b_t^K = 0$ in the first stage). In both panels, we report t -statistics using Newey-West standard errors with K lags. In Panel B, we also report block-bootstrapped t -statistics using standard errors calculated as the standard deviation of coefficient estimates in 1,000 bootstrap replications. We report these t -statistics for two null hypotheses: $H_0 : d_1^{c,K} = 0$ and $H_0 : d_1^{c,K} = 1$. R^2 's are reported in percentage points.

<i>Panel A: Unconditional</i>					<i>Panel B: Conditional</i>				
Horizon K	1	3	6	12		1	3	6	12
$d_0^{u,K}$	0.00	0.01	0.01	0.02	$d_0^{c,K}$	0.00	0.00	0.00	0.00
t	(9.10)	(8.73)	(7.78)	(7.52)	t	(0.76)	(0.15)	(0.07)	(0.10)
					t_{boot}	(0.42)	(0.09)	(0.04)	(0.05)
$d_1^{u,K}$	-0.07	-0.15	-0.20	-0.50	$d_1^{c,K}$	0.77	0.84	0.84	0.81
$t(d_1^{u,K} = 0)$	(-1.37)	(-1.08)	(-0.77)	(-1.06)	$t(d_1^{c,K} = 0)$	(4.74)	(5.60)	(5.30)	(4.56)
					$t_{boot}(d_1^{c,K} = 0)$	(2.57)	(2.74)	(2.24)	(1.77)
					$t(d_1^{c,K} = 1)$	(-1.44)	(-1.04)	(-1.01)	(-1.06)
					$t_{boot}(d_1^{c,K} = 1)$	(-0.78)	(-0.51)	(-0.43)	(-0.41)
R^2	0.26	0.76	0.44	1.12	R^2	3.29	10.29	13.28	14.92
					$R^2 - OOS$	1.80	4.60	6.70	10.41

Table 3

Characterizing inflation beta-sorted portfolios and the inflation risk premium.

This table presents the set of ten inflation beta-sorted portfolios and our resulting estimates of the inflation risk premium, denoted IP_t^{HL} , IP_t^{MC} , and IP_t^{CS} . IP_t^{HL} is the High-minus-Low inflation beta decile portfolio from the sort. IP_t^{MC} is a maximum correlation mimicking-portfolio that is constructed from a multiple regression of inflation innovations on the inflation beta decile portfolios. IP_t^{CS} is the inflation risk premium estimated from a cross-sectional regression of individual stock returns on lagged inflation betas (controlling for size, book-to-market, and momentum). Panel A presents ex post inflation betas, $\beta_{\Pi, post}$, that are estimated with a simple regression of portfolio returns on inflation innovations over the full sample. t -statistics are in parentheses and use Newey-West standard errors with lag length one. IP_t^{MC} and IP_t^{CS} are scaled to have the same post-ranking inflation beta as IP_t^{HL} . Next, we present annualized performance statistics (average returns, Sharpe ratio, and CAPM alpha) over the full sample from July 1967 to December 2014 (Panel B) and split around December 2002 (Panel C). For average returns, Sharpe ratios, and alphas, we present in the last column also the p -value from the monotonicity test of Patton and Timmermann (2010). Inference from this test is based on a block-bootstrap with a block size of one year. The null hypothesis is a non-increasing (non-decreasing) relation between returns and inflation beta whenever the High portfolio outperforms (underperforms) the Low portfolio.

	H	2	3	4	5	6	7	8	9	L	IP_t^{HL}	IP_t^{MC}	IP_t^{CS}	p -val	
<i>Panel A: Ex post inflation exposures</i>															
$\beta_{\Pi, post}$	-0.02	-1.03	-1.37	-1.86	-1.94	-2.40	-2.13	-2.85	-2.95	-3.02	3.00	3.00	3.00		
t	(-0.01)	(-0.68)	(-0.91)	(-1.26)	(-1.34)	(-1.59)	(-1.36)	(-1.82)	(-1.85)	(-1.68)	(4.38)	(5.82)	(4.69)		
<i>Panel B: The inflation risk premium (full sample)</i>															
Avg. ret.	5.26	7.29	7.85	7.59	8.53	8.28	8.10	8.50	8.76	9.49	-4.23	-4.39	-4.21	(0.046)	
t	(1.59)	(2.47)	(2.88)	(2.85)	(3.20)	(3.07)	(2.94)	(3.02)	(2.99)	(3.00)	(-2.08)	(-3.06)	(-2.22)		
Sharpe	0.23	0.36	0.42	0.41	0.46	0.44	0.43	0.44	0.43	0.44	-0.30	-0.44	-0.32	(0.086)	
α_{CAPM}	-2.03	0.52	1.39	1.23	2.14	1.81	1.55	1.85	1.87	2.18	-4.21	-3.18	-3.79	(0.097)	
t	(-1.15)	(0.38)	(1.31)	(1.21)	(2.07)	(1.68)	(1.34)	(1.48)	(1.35)	(1.38)	(-1.83)	(-2.06)	(-1.80)		

Table 3 continued

	H	2	3	4	5	6	7	8	9	L	IP_t^{HL}	IP_t^{MC}	IP_t^{CS}	p -val Mono.
Pre-2002														
Avg. ret.	1.84	4.42	5.74	5.87	6.97	6.96	6.95	7.84	8.38	9.33	-7.49	-5.85	-6.81	(0.005)
t	(0.50)	(1.31)	(1.85)	(1.91)	(2.27)	(2.19)	(2.11)	(2.32)	(2.35)	(2.45)	(-3.24)	(-3.51)	(-3.12)	
Sharpe	0.08	0.22	0.31	0.32	0.38	0.37	0.35	0.39	0.40	0.41	-0.54	-0.59	-0.52	(0.022)
α_{CAPM}	-3.58	-0.73	0.80	0.98	2.01	1.87	1.72	2.51	2.79	3.46	-7.04	-4.71	-6.02	(0.002)
t	(-1.81)	(-0.44)	(0.61)	(0.78)	(1.67)	(1.45)	(1.23)	(1.69)	(1.73)	(1.90)	(-2.94)	(-2.96)	(-2.73)	
Post-2002														
Avg. ret.	15.37	15.76	14.11	12.67	13.14	12.20	11.53	10.46	9.89	9.94	5.42	-0.06	3.49	(0.036)
t	(2.12)	(2.61)	(2.51)	(2.35)	(2.47)	(2.38)	(2.35)	(2.13)	(2.02)	(1.81)	(1.30)	(-0.02)	(0.93)	
Sharpe	0.61	0.75	0.72	0.68	0.71	0.69	0.68	0.61	0.58	0.52	0.38	-0.01	0.27	(0.826)
α_{CAPM}	0.70	3.04	2.09	1.03	1.76	1.09	0.91	-0.15	-0.59	-1.67	2.37	0.66	0.93	(0.227)
t	(0.20)	(1.34)	(1.10)	(0.63)	(1.00)	(0.71)	(0.60)	(-0.09)	(-0.34)	(-0.79)	(0.55)	(0.21)	(0.25)	
Post-minus-Pre-2002 difference														
Avg. ret.	13.53	11.34	8.37	6.80	6.17	5.25	4.58	2.62	1.50	0.62	12.92	5.79	10.30	(0.000)
t	(1.67)	(1.64)	(1.30)	(1.10)	(1.01)	(0.87)	(0.77)	(0.44)	(0.25)	(0.09)	(2.71)	(1.79)	(2.37)	
Sharpe	0.53	0.53	0.41	0.35	0.33	0.32	0.32	0.22	0.19	0.11	0.92	0.58	0.79	(0.002)
α_{CAPM}	4.28	3.76	1.29	0.05	-0.25	-0.78	-0.81	-2.66	-3.38	-5.13	9.41	5.37	6.96	(0.007)
t	(1.08)	(1.34)	(0.56)	(0.02)	(-0.12)	(-0.39)	(-0.39)	(-1.21)	(-1.43)	(-1.84)	(1.91)	(1.50)	(1.60)	

Panel C: The inflation risk premium in subsamples

Table 4

Exposures to breakeven inflation.

This table presents the exposure of inflation beta-sorted portfolios to first-differences in breakeven inflation (denoted ΔBE_t). Breakeven inflation is the difference between the yield on a ten-year constant maturity nominal Treasury bond and its TIPS equivalent. The exposures are estimated with the OLS time-series regression $R_{p,t} = a_p + \beta_{BE,p} \Delta BE_t + e_{p,t}$. In Panel A, we use the returns of portfolios sorted on exposure to ARMA(1,1)-innovations in inflation ($u_{\Pi,t}$), as in Table 3. In Panel B, we use the returns of decile portfolios sorted on exposure to ΔBE_t . The exposure of each stock to ΔBE is estimated with an OLS regression over a 60-month rolling window. We report the coefficients $\beta_{BE,p}$ (with t -statistics in parentheses) and the R^2 for each of the regressions. The sample for both panels is January 2003 (when TIPS data become available) to December 2014.

	High	2	3	4	5	6	7	8	9	Low	IP_t^{HL}
<i>Panel A: Portfolios sorted on exposure to ARMA(1,1)-innovations in inflation ($u_{\Pi,t}$)</i>											
$\beta_{BE,p}$	17.84	11.43	9.36	9.31	9.68	8.52	8.87	8.09	6.98	8.05	9.79
	(2.25)	(1.73)	(1.46)	(1.49)	(1.66)	(1.45)	(1.49)	(1.42)	(1.08)	(1.13)	(4.12)
R^2	13.09	7.46	5.60	6.07	6.85	5.61	6.76	5.46	3.95	4.20	11.87
<i>Panel B: Portfolios sorted on exposure to changes in breakeven inflation (ΔBE_t)</i>											
$\beta_{BE,p}$	21.09	16.88	15.38	12.82	12.86	10.53	9.27	7.29	8.19	8.61	12.47
	(2.22)	(1.93)	(1.98)	(1.54)	(1.63)	(1.60)	(1.38)	(1.09)	(1.31)	(1.03)	(4.82)
R^2	16.40	13.66	13.35	10.43	11.21	9.88	7.58	5.28	7.09	5.08	20.01

Table 5

Predicting the inflation risk premium with the nominal-real covariance.

This table presents time-series predictive regressions of ten inflation beta-sorted portfolios as well as our three estimates of the inflation risk premium (IP_t^{HL} , IP_t^{MC} , and IP_t^{CS}) on the lagged nominal-real covariance. The nominal-real covariance is proxied by the time-varying relation between current inflation and future 12-month consumption growth: NRC . The left-hand side returns are compounded over horizons of one, three, and 12 months. We report for each regression the estimated coefficients, with corresponding t -statistics based on either Newey-West standard errors with K lags or standard errors that are derived from 1,000 block-bootstrapped coefficient estimates, and the adjusted R^2 (in percentage points). We also present the p -value from a Patton and Timmermann (2010) monotonicity test. The null hypothesis is a non-increasing (non-decreasing) relation between inflation beta and the main coefficient of interest L_{NRC} (the intercept L_0). The sample period is July 1967 to December 2014.

	H	2	3	4	5	6	7	8	9	L	IP_t^{HL}	IP_t^{MC}	IP_t^{CS}	p -val
$K = 1$														
L_0	5.26	7.29	7.85	7.59	8.53	8.28	8.10	8.50	8.76	9.49	-4.23	-4.39	-4.21	(0.045)
t	(1.49)	(2.34)	(2.73)	(2.69)	(2.99)	(2.85)	(2.72)	(2.80)	(2.76)	(2.76)	(-1.95)	(-2.92)	(-2.08)	
t_{boot}	(1.79)	(2.87)	(3.28)	(3.24)	(3.65)	(3.50)	(3.31)	(3.36)	(3.23)	(3.18)	(-1.80)	(-2.10)	(-2.10)	
L_{NRC}	2.06	1.36	0.53	-0.16	-0.54	-1.14	-1.46	-2.19	-2.37	-2.41	4.47	2.73	3.94	(0.001)
t	(0.57)	(0.44)	(0.18)	(-0.05)	(-0.18)	(-0.37)	(-0.46)	(-0.68)	(-0.72)	(-0.68)	(2.11)	(1.85)	(1.98)	
t_{boot}	(0.67)	(0.51)	(0.21)	(-0.05)	(-0.22)	(-0.45)	(-0.57)	(-0.84)	(-0.85)	(-0.78)	(1.90)	(1.37)	(1.99)	
R^2	-0.11	-0.14	-0.17	-0.18	-0.17	-0.14	-0.13	-0.07	-0.06	-0.07	0.67	0.46	0.59	
$K = 3$														
L_0	5.72	7.46	7.99	7.78	8.84	8.59	8.48	8.85	9.17	9.99	-4.27	-4.19	-3.93	(0.040)
t	(1.77)	(2.71)	(3.17)	(3.09)	(3.36)	(3.23)	(3.10)	(3.18)	(3.17)	(3.18)	(-2.03)	(-2.89)	(-2.20)	
t_{boot}	(1.93)	(2.90)	(3.31)	(3.29)	(3.74)	(3.59)	(3.43)	(3.46)	(3.34)	(3.30)	(-1.75)	(-1.95)	(-2.01)	
L_{NRC}	2.17	1.20	0.28	-0.43	-0.79	-1.46	-1.84	-2.73	-3.03	-2.99	5.16	3.33	4.32	(0.003)
t	(0.63)	(0.41)	(0.10)	(-0.15)	(-0.27)	(-0.49)	(-0.60)	(-0.89)	(-0.96)	(-0.89)	(2.45)	(2.37)	(2.32)	
t_{boot}	(0.71)	(0.45)	(0.12)	(-0.16)	(-0.32)	(-0.57)	(-0.72)	(-1.04)	(-1.08)	(-0.96)	(2.14)	(1.65)	(2.19)	
R^2	0.01	-0.10	-0.17	-0.16	-0.14	-0.05	0.01	0.23	0.28	0.20	2.48	2.18	2.16	
$K = 12$														
L_0	5.88	7.66	8.09	7.99	9.08	8.82	8.68	9.03	9.35	10.24	-4.37	-4.25	-3.94	(0.059)
t	(1.91)	(2.89)	(3.46)	(3.36)	(3.74)	(3.62)	(3.47)	(3.54)	(3.60)	(3.51)	(-1.68)	(-2.62)	(-2.31)	
t_{boot}	(1.84)	(2.74)	(3.08)	(3.10)	(3.53)	(3.39)	(3.23)	(3.25)	(3.13)	(3.14)	(-1.69)	(-1.85)	(-1.97)	
L_{NRC}	2.01	0.99	-0.20	-0.77	-1.14	-1.75	-2.31	-3.39	-3.63	-3.37	5.38	3.73	4.51	(0.040)
t	(0.59)	(0.34)	(-0.08)	(-0.27)	(-0.39)	(-0.60)	(-0.77)	(-1.15)	(-1.23)	(-1.07)	(2.22)	(2.58)	(2.78)	
t_{boot}	(0.64)	(0.35)	(-0.07)	(-0.28)	(-0.43)	(-0.66)	(-0.85)	(-1.21)	(-1.21)	(-1.02)	(2.13)	(1.77)	(2.32)	
R^2	0.41	0.01	-0.17	-0.04	0.11	0.50	0.94	2.13	2.32	1.55	6.65	7.64	9.68	

Table 6

Time-varying inflation risk premia in pooled regressions.

Panel A of this table presents the estimated coefficients [with asymptotic Driscoll and Kraay (1998) standard errors in parentheses] from pooled predictive regressions of annual returns of the inflation beta-sorted decile portfolios on their lagged inflation beta ($\beta_{\Pi,p,t}$), the nominal-real covariance (NRC_t^C), and an interaction term: $R_{p,t+1:t+12} = L_0 + L_{\beta_{\Pi}}\beta_{\Pi,p,t} + L_{NRC}NRC_t^C + L_{\beta_{\Pi} \times NRC}(\beta_{\Pi,p,t} \times NRC_t^C) + \varepsilon_{p,t+1:t+12}$. In Model [1], we set $L_{\beta_{\Pi} \times NRC} = 0$, whereas Model [2] estimates all three coefficients freely. For the latter model, Panel B presents the predicted risk premia (with asymptotic and bootstrapped standard errors in parentheses) in four distinct cases. That is, when $\beta_{\Pi,p,t}$ and NRC_t^C are at plus or minus one standard deviation from their respective means in the pool. Models [3] to [5] in Panel A analyze which components of inflation betas interact with the nominal-real covariance. Model [3] replaces the interaction term, $\beta_{\Pi,p,t} \times NRC_t^C$, with the interaction between the nominal-real covariance and the portfolio-specific component of inflation betas, $\widehat{\beta}_{\Pi,p} = T^{-1} \sum_{t=1}^T \beta_{\Pi,p,t}$. Model [4] replaces the interaction term, $\beta_{\Pi,p,t} \times NRC_t^C$, with the interaction between the nominal-real covariance and the time-specific component of inflation betas, $\widehat{\beta}_{\Pi,t} = 10^{-1} \sum_{p=1}^{10} \beta_{\Pi,p,t}$. Model [5] includes both component-wise interaction terms. To accommodate interpretation, NRC_t^C is standardized in the time-series to have mean zero and standard deviation equal to one, whereas $\beta_{\Pi,p,t}$ is demeaned in the pool. The sample period is July 1967 to December 2014.

Panel A: Coefficient estimates						Panel B: Predicted risk premia from [2]		
Model	[1]	[2]	[3]	[4]	[5]		High NRC_t^C	Low NRC_t^C
L_0	8.48	5.10	8.48	4.96	5.00	High $\beta_{\Pi,p,t}$	7.65	-0.50
t	(3.48)	(2.06)	(3.48)	(1.96)	(1.97)	t	(2.02)	(-0.11)
$L_{\beta_{\Pi}}$	0.51	-0.51	0.56	-0.62	-0.57	t_{boot}	(1.65)	(-0.09)
t	(0.50)	(-0.65)	(0.56)	(-0.77)	(-0.70)	Low $\beta_{\Pi,p,t}$	-3.33	16.56
L_{NRC}	-2.09	-2.93	-2.17	-2.86	-2.92	t	(-0.54)	(4.09)
t	(-0.64)	(-1.09)	(-0.66)	(-1.06)	(-1.09)	t_{boot}	(-0.52)	(3.18)
$L_{\beta_{\Pi} \times NRC}$		2.36				Diff	10.98	-17.05
t		(3.63)				t	(1.90)	(-2.69)
$L_{\beta_{\Pi,p} \times NRC}$			2.45		2.17	t_{boot}	(1.40)	(-2.27)
t			(2.80)		(2.50)			
$L_{\beta_{\Pi,t} \times NRC}$				2.45	2.43			
t				(3.45)	(3.42)			
R^2	0.71	8.80	1.40	8.26	8.80			

Table 7

The nominal-real covariance and the market risk premium (MRP).

Panel A of this table presents coefficient estimates from time-series predictive regressions of CRSP value-weighted excess stock market returns on the market's inflation beta ($\beta_{\Pi,m,t}$), the nominal-real covariance (NRC_t^C), and their interaction. The right block of results in Panel A presents the predicted MRP from Model [2] in four cases, that is, when $\beta_{\Pi,m,t}$ and NRC_t^C equal plus or minus one standard deviation from their respective means. We present this regression for the 12-month horizon and standard errors are Newey-West(12). To accommodate interpretation, NRC_t^C is standardized and $\beta_{\Pi,m,t}$ is demeaned. Panel B presents nonparametric estimates of the MRP in four similar cases. These estimates are calculated by averaging realized excess market returns in sample months $t+1$ when $\beta_{\Pi,m,t}$ and NRC_t^C in month t are either below the 40th or above the 60th percentile. The sample period is July 1967 to December 2014.

<i>Panel A: Time-series predictive regressions</i>					
Model	Coefficient estimates		Predicted MRP from [2]		
	[1]	[2]		High NRC_t^C	Low NRC_t^C
$L_{0,m}$	6.39	3.02	High $\beta_{\Pi,m,t}$	4.57	-4.03
t	(2.95)	(1.26)	t	(1.46)	(-0.68)
$L_{\beta_{\Pi,m}}$	0.28	-1.29	Low $\beta_{\Pi,m,t}$	-2.08	13.61
t	(0.23)	(-1.16)	t	(-0.41)	(4.08)
$L_{NRC,m}$	-1.14	-1.77	Diff	6.65	-17.64
t	(-0.39)	(-0.72)	t	(1.50)	(-2.54)
$L_{\beta_{\Pi} \times NRC,m}$		2.85			
t		(3.55)			
R^2	-0.08	8.59			

<i>Panel B: Nonparametric estimates of the MRP</i>		
	$NRC_t^C \geq P_{60}$	$NRC_t^C \leq P_{40}$
$\beta_{\Pi,m,t} \geq P_{60}$	8.10	-2.06
t	(1.77)	(-0.25)
$\beta_{\Pi,m,t} \leq P_{40}$	4.01	16.14
t	(0.50)	(3.03)
Diff	4.09	-18.20
t	(0.44)	(-1.84)

Table 8

Sorting stocks on exposure to inflation-hedging industry factors.

For this table, we rank 48 industries on their full sample, unconditional exposure to ARMA(1,1)-innovations in inflation ($u_{\Pi,t}$). From this ranking, we construct long-short industry factors that equal-weight the top (bottom) K industries, where K ranges from one to five. We then sort individual stocks in decile portfolios on their (60-month rolling window) exposure to these industry factors. We leave out all stocks that belong to the top and bottom K industries, which are presented in Panel A. We then construct a High-minus-Low decile portfolio using the market value-weighted portfolios of the stocks with the highest and lowest exposures to the K th industry factor. For each of these K sorts, we present in Panel B: the ex post, or post-ranking, exposure to $u_{\Pi,t}$; and, the estimates from a predictive regression at the annual horizon of the High-minus-Low portfolio returns on the lagged nominal-real covariance (see Eq. (7) of the paper). t -statistics in parentheses are calculated using Newey-West standard errors. The sample period is July 1967 to December 2014.

	K	1	2	3	4	5
<i>Panel A: The 5 best and worst inflation hedging industries</i>						
Top (Best hedges)	1	Coal	Coal	Coal	Coal	Coal
	2		Oil	Oil	Oil	Oil
	3			Gold	Gold	Gold
	4				Mines	Mines
	5					Agriculture
Bottom (Worst hedges)	1	Meals	Meals	Meals	Meals	Meals
	2		Textiles	Textiles	Textiles	Textiles
	3			Clothes	Clothes	Clothes
	4				Transportation	Transportation
	5					Personal Svcs.
<i>Panel B: Inflation exposure and predictability of industry-sorted portfolios</i>						
	$\beta_{\Pi,post}$	5.18 (4.99)	4.70 (4.57)	4.41 (4.33)	4.27 (4.49)	4.21 (4.32)
	L_0	-4.10 (-1.36)	-4.97 (-2.17)	-5.27 (-2.30)	-4.78 (-2.12)	-5.05 (-2.20)
	L_{NRC}	9.03 (3.16)	5.86 (2.77)	6.01 (3.01)	6.26 (3.19)	5.64 (2.96)
	R^2	12.13	8.72	9.15	9.94	8.04

Table 9

The inflation risk premium within- and across-industry.

This table presents time-series predictive regressions for two alternative inflation risk premia (derived from a within- and an across-industry sort) on the lagged nominal-real covariance. For the across-industry sort, we calculate the inflation beta of each of 48 industries as the value-weighted average inflation beta of all stocks in that industry. We sort the industries into quintile portfolios, such that each across-industry quintile portfolio contains either nine or ten industries, which we equal-weight to calculate returns. For the within-industry sort, we first construct five value-weighted stock portfolios within each of the 48 industries by splitting at the quintiles of ranked inflation betas of the stocks within that industry. This gives us a total of 48-by-5 value-weighted portfolios. Within-industry quintile portfolio returns are then calculated as the equal-weighted average over the 48 industries. The aggregate across- and within-industry effects are presented in the last column as the difference between the High and Low across- and within-industry portfolio returns, respectively. For each portfolio, we present coefficient estimates from the predictive regression of 12-month compounded future returns on the lagged nominal-real covariance (as in Table 5). The sample period is July 1967 to December 2014.

<i>Panel A: Across-industry</i>						
	High β_{Π} industries	2.00	3.00	4.00	Low β_{Π} industries	High-Low
L_0	5.93	6.99	7.94	5.78	9.29	-3.36
t	(2.98)	(3.29)	(3.57)	(2.41)	(3.62)	(-1.95)
L_{NRC}	0.69	-0.67	-1.23	-1.51	-3.94	4.63
t	(0.30)	(-0.28)	(-0.48)	(-0.54)	(-1.30)	(2.70)
R^2	-0.02	-0.05	0.22	0.38	2.95	10.03
<i>Panel B: Within-industry</i>						
	High β_{Π} stocks	2.00	3.00	4.00	Low β_{Π} stocks	High-Low
L_0	6.56	7.16	7.78	7.89	8.28	-1.73
t	(2.58)	(2.93)	(3.54)	(3.33)	(3.42)	(-1.55)
L_{NRC}	0.27	0.15	-1.09	-2.10	-2.97	3.24
t	(0.09)	(0.05)	(-0.41)	(-0.75)	(-1.07)	(3.29)
R^2	-0.17	-0.17	0.14	0.86	1.74	10.81

Table 10

Inflation risk in cash flows versus discount rates.

We use the procedure in Campbell and Vuolteenaho (2004) to decompose the unexpected return of the High and Low inflation beta-sorted portfolios into discount rate and cash flow news: $R_{t+1} - E_t(R_{t+1}) = -N_{DR,t+1} + N_{CF,t+1}$ (over the sample from July 1967 to December 2014). We regress each of these components separately on ARMA(1,1)-innovations in inflation ($u_{\Pi,t}$) to decompose inflation betas: $\beta_{\Pi} = \beta_{\Pi}^{-DR} + \beta_{\Pi}^{CF}$. We consider three sets of state variables. In the first set, we include the term spread (TS), price-earnings ratio (PE), and value spread (VS). The second set includes the dividend yield (DY), default spread (DS), and term spread. In the third set, we add the nominal-real covariance (NRC^C) to the first specification. To ensure that the VAR-models are correctly specified, we include in all three specifications the dividend yield of the respective inflation beta-sorted portfolio ($DY_{H/L}$) as a state variable. We report below the estimated exposures. The t -statistics use either Newey-West standard errors with one lag (t_{NW}) or standard errors simulated under the null of no stock return predictability as in Campbell (1991) (t_{sim}). In this simulation, we allow (innovations in) the state variables and inflation to be correlated, just like in the data. For each portfolio, we also report the resulting return variance decomposition (column % Var) that shows the share of return variance explained by (the negative of) discount rate news, cash flow news, and two times their covariance. In the last row, we report the fraction of the High-minus-Low difference in inflation beta coming from exposure in cash flows.

		Model 1: $TS, PE, VS, DY_{H/L}$				Model 2: $DY_{MKT}, DS, TS, DY_{H/L}$				Model 3: $NRC^C, TS, PE, VS, DY_{H/L}$			
		β_{Π}	t_{NW}	t_{sim}	% Var	β_{Π}	t_{NW}	t_{sim}	% Var	β_{Π}	t_{NW}	t_{sim}	% Var
High	$R_{t+1} - E_t[R_{t+1}]$	-0.89	(-0.54)	(-0.83)	1.00	-0.94	(-0.57)	(-0.90)	1.00	-0.87	(-0.52)	(-0.81)	1.00
	$-N_{DR,t+1}$	-1.78	(-1.57)	(-1.94)	0.43	-1.66	(-1.32)	(-1.06)	0.57	-2.82	(-2.14)	(-2.84)	0.87
	$N_{CF,t+1}$	0.89	(1.24)	(0.66)	0.34	0.71	(1.01)	(0.39)	0.32	1.95	(1.67)	(1.40)	1.08
	$2 \times Cov(-N_{DR,t+1}, N_{CF,t+1})$				0.23				0.11				-0.95
Low	$R_{t+1} - E_t[R_{t+1}]$	-2.71	(-1.65)	(-2.58)	1.00	-2.73	(-1.66)	(-2.60)	1.00	-2.70	(-1.65)	(-2.57)	1.00
	$-N_{DR,t+1}$	-1.20	(-1.66)	(-1.29)	0.37	-2.02	(-3.73)	(-0.17)	0.37	-1.05	(-1.40)	(-0.88)	0.35
	$N_{CF,t+1}$	-1.51	(-1.45)	(-1.15)	0.42	-0.71	(-0.48)	(-0.06)	0.62	-1.65	(-1.70)	(-1.13)	0.39
	$2 \times Cov(-N_{DR,t+1}, N_{CF,t+1})$				0.22				0.02				0.26
$\frac{\beta_{\Pi}^{CF}}{\beta_{\Pi}^{High} - \beta_{\Pi}^{Low}}$		1.32				0.80				1.96			

Table 11

Calibrated parameters of the model.

This table reports the configuration of the parameters used in the calibration of the model. The model is calibrated on a monthly decision interval.

<i>Preferences</i>		
Discount factor	δ	0.99
Elasticity of intertemporal substitution	ψ	2
Risk aversion coefficient	γ	15
<i>Inflation</i>		
Mean of inflation rate	μ_π	0.0033
Inflation AR(1) coefficient	ρ_π	0.7992
Inflation MA(1) coefficients	ϕ_π	-0.0025
	ξ_π	7.30×10^{-4}
<i>Consumption</i>		
Mean of real consumption growth	μ_c	0.0016
Exposure of expected consumption to inflation	ρ_c	-0.0521
Volatility of consumption-specific shocks	σ_c	0.0020
Consumption exposure to inflation shocks	ξ_c	-0.1983
<i>Nominal-real covariance</i>		
Mean of nominal-real covariance	φ_0	-0.0062
Persistence of nominal-real covariance	ν	0.9963
Volatility of nominal-real covariance	σ_w	0.0029
<i>High inflation beta portfolio</i>		
Mean of dividend growth	μ_H	3.44×10^{-4}
Exposure of expected dividends to inflation	ρ_H	-0.2345
Exposure to consumption-specific shocks	σ_H	0.0318
Exposure to inflation shocks	ξ_H	0.1942
<i>Low inflation beta portfolio</i>		
Mean of dividend growth	μ_L	0.0060
Exposure of expected dividends to inflation	ρ_L	0.6989
Exposure to consumption-specific shocks	σ_L	-0.0517
Exposure to inflation shocks	ξ_L	0.1082
<i>Aggregate market</i>		
Mean of dividend growth	μ_{MKT}	1.72×10^{-4}
Exposure of expected dividends to inflation	ρ_{MKT}	0.7099
Exposure to consumption-specific shocks	σ_{MKT}	0.0337
Exposure to inflation shocks	ξ_{MKT}	-0.3076

Table 12

Calibrated moments.

This table shows that the model can match means and standard deviations of returns and betas, as well as the slope coefficient from regressing returns on the High-minus-Low portfolio (as well as the aggregate market portfolio) on the nominal-real covariance (NRC_t^C).

	Data			Model		
<i>Descriptive statistics</i>						
	π_t	$u_{\pi,t}$	Δc_t	π_t	$u_{\pi,t}$	Δc_t
Mean	3.93	0.11	1.97	3.93	0.00	1.97
St. dev.	1.11	0.85	1.13	1.11	0.85	2.49
AR(1)	62.47	15.35	-16.84	62.47	0.00	0.19
<i>Correlations</i>						
π_t	1	79.03	-17.75	1	76.71	-8.10
$u_{\pi,t}$		1	-17.41		1	0.00
Δc_t			1			1
<i>Predictive regressions of consumption growth on inflation</i>						
	Uncond.	Cond.		Uncond.	Cond.	
$d_0^{I^2}$	0.02	0.00		0.02	0.00	
$d_1^{I^2}$	-0.50	0.81		-0.50	1.00	
R^2	1.12	14.92		0.41	15.40	
AR(1) NRC_t^C		99.63			99.63	
St. dev. NRC_t^C		0.87			0.79	
<i>Characterizing inflation beta-sorted portfolios</i>						
	High	Low		High	Low	
Mean $\beta_{u,it}$	-0.02	-3.02		-0.02	-3.04	
St. dev. $\beta_{u,it}$	3.61	3.25		4.03	2.86	
Avg. ret.	5.26	9.49		5.28	9.42	
St. dev.	22.72	21.78		22.74	18.97	
<i>Characterizing the inflation risk premium (IP_t^{HL})</i>						
St. dev. $\beta_{u,HLt}$		1.36			1.00	
$L_{NRC,IP_t^{HL}}$		5.38			5.41	
<i>The nominal-real covariance and the market risk premium</i>						
Mean $\beta_{u,mt}$		-1.96			-1.95	
St. dev. $\beta_{u,mt}$		2.19			2.85	
$corr(\beta_{u,H}, \beta_{u,m})$		95.75			100.00	
$corr(\beta_{u,L}, \beta_{u,m})$		97.12			100.00	
$L_{\beta_{H \times NRC,m}}$		2.85			2.69	