

Time-Varying Inflation Risk and the Cross-Section of Stock Returns*

Martijn Boons[†] Fernando Duarte[‡] Frans de Roon[§]

Marta Szymanowska[¶]

November 30, 2016

*We thank 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, Laura Spierdijk, Adrien Verdelhan, Luis Viceira and Ivo Welch; conference participants at the 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. Nova School of Business and Economics, Lisbon, Portugal. E-mail: martijn.boons@novasbe.pt. Phone: +351-914246751.

[‡]Federal Reserve Bank of New York. E-mail: fernando.duarte@ny.frb.org

[§]Department of Finance, CentER - Tilburg University. E-mail: f.a.deroon@uvt.nl

[¶]Department of Finance, Rotterdam School of Management, Erasmus University. E-mail: mszymanowska@rsm.nl

Abstract

We show that inflation risk is priced in the cross-section of US stock returns with a price of inflation risk that is comparable in magnitude to that of the aggregate market. The inflation risk premium varies over time conditional on the nominal-real covariance, the time-varying relation between inflation and the real economy. Using a consumption-based equilibrium asset pricing model, we argue that inflation is priced because it predicts real consumption growth. The historical changes in the predictability of consumption with inflation, which are mediated by the nominal-real covariance, can account for the size, variability, predictability and sign-reversals –last observed in the 2000s– in the inflation risk premium.

JEL Classification Codes: G11, G12, G13

Keywords: Inflation, (Time-Varying) Inflation Risk Premium, Inflation Hedging, Cross-Sectional Asset-Pricing, Nominal-Real Covariance.

1 Introduction

In this paper, we show that inflation risk is priced in the cross-section of stock returns and that this price is strongly time-varying. We estimate an unconditional inflation risk premium of -4.2% for an inflation-mimicking portfolio, or 1.4% per unit of inflation beta, which corresponds to a Sharpe ratio of -0.30.¹ Thus, in magnitude, the price of inflation risk is comparable to the price of aggregate stock market risk. An unconditional negative price of risk means that, over long periods, high inflation corresponds to bad states of nature: investors are willing to accept lower returns when holding securities that are good hedges against inflation.

Recent studies show that inflation does not always signal a bad state of the economy (Bekaert and Wang (2010), Campbell et al. (2013), and David and Veronesi (2014)). Consistent with this observation, we find that the nominal-real covariance - which is strongly time-varying - is an economically and statistically significant predictor of the conditional price of inflation risk. The nominal-real covariance is measured as the (time-varying) slope coefficient of a rolling regression of real consumption growth on lagged inflation, and therefore indicates whether inflation shocks provide good or bad macroeconomic news. A one standard deviation change in the nominal-real covariance is associated with a change in expected return of 4.5% for the high-minus-low inflation beta portfolio, or 1.5% per unit beta, which is similar in magnitude to the unconditional price of inflation risk.

In addition to the price of risk, the inflation betas of stocks - the quantity of inflation risk - also vary distinctly with the nominal-real covariance. We find a wide cross-section of inflation exposures, with persistent and statistically significant inflation betas. For ten portfolios formed by sorting on inflation betas, a one standard deviation change in the nominal-real covariance leads to a change in inflation beta of 1.55 on average, which is about half of the cross-sectional spread in the betas of the ten decile portfolios. For individual stocks, or decile portfolios, a

¹This mimicking portfolio is the high-minus-low portfolio when sorting stocks into deciles based on their inflation exposure (beta).

change in beta of 1.55 translates into a change in expected return of about 2.2%. Although the effect of the nominal-real covariance on the price of risk and the quantity of risk cannot simply be added, as they are correlated, these numbers show that both channels are important as determinants of expected stock returns.

By studying the cross-section of stock returns, we not only uncover a new source of information about the inflation risk premium in the economy, but also provide insights about the distribution and pricing of inflation risk of individual firms. 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 Treasury Protected Inflation Securities (TIPS), estimates of the inflation risk in the bond market have become more reliable and widely available.² Another conventional way to estimate the inflation premium is to study the joint time-series behavior of inflation and aggregate market returns. A landmark example is Modigliani and Cohn (1979), who find a negative correlation between inflation and the S&P500 returns over the 1970's and propose an explanation based on inflation illusion. Other recent economic explanations of the inflation premium in the aggregate market are based on important contributions by Wachter (2006) using habit formation, Gabaix (2008) using rare disasters, and Bansal and Shaliastovich (2010) using long-run risk. If the fundamental mechanisms of the real effects of inflation originate at the level of individual firms, studying the cross-section of stocks can provide valuable additional information that is masked in the aggregate market and the yield curve. Our findings show that the cross-section of stock returns is a rich source of information about the inflation risk premium.

We develop an equilibrium model that rationalizes the observed inflation premium by arguing that inflation today predicts real consumption growth in the future. Unconditionally, as pointed out by Piazzesi and Schneider (2005), inflation predicts consumption growth negatively.

²Ang and Piazzesi (2003), Ang, Piazzesi and Wei (2006), Ang, Bekaert and Wei (2007, 2008), Singleton, Dai and Yang (2007), Singleton, Le and Qiang (2010), Singleton and Le (2010), Haubrick, Pennacchi, and Ritchken (2008), Gurkaynak, Sack, and Wright (2010), Chen, Liu, Cheng (2010), Campbell et al. (2013).

In our model, this generates the negative unconditional price of inflation risk observed in the data. Conditionally, the predictability of consumption growth with inflation is determined, in sign and magnitude, by the time-varying nominal-real covariance. Our model uses this time-varying relation between inflation and consumption to generate the dynamics of the price of inflation risk, as well as of the quantity of risk, consistent with our empirical estimates. The model takes the stochastic processes for consumption, inflation, and the nominal-real covariance as given and asset prices are then determined endogenously through the representative agent's Euler equation. Based on estimated processes for inflation, consumption, and their predictive regressions, and using reasonable values for preference parameters, we show that the model can quantitatively replicate the observed conditional and unconditional inflation risk premiums.

To generate inflation premiums consistent with the data, our model has four key ingredients, all of which are necessary. The first ingredient - as already mentioned - is that inflation predicts future real consumption growth. In full-sample regressions of consumption growth on inflation with constant coefficients, a one percent inflation this month is associated with a decrease of half a percent in consumption growth over the next 12 months. However, the R^2 of this regression is only slightly over 1%. Allowing for time-varying slope coefficients that may also change sign - the nominal-real covariance - the R^2 increases to over 10%.

The second ingredient is that inflation is persistent and follows an ARMA(1,1) process. Inflation persistence is widely documented in the literature, for example in Fuhrer and Moore (1995), Stock and Watson (2005), Campbell and Viceira (2001) and Ang, Bekaert and Wei (2007), and the ARMA(1,1) feature is similar to Fama and Gibbons (1984), Vassalou (2000), and Campbell and Viceira (2001). That inflation is persistent will be important in our model to quantitatively match the inflation premium: more persistent inflation induces a larger market price of inflation risk because it affects consumption growth for a longer period of time. The MA(1) part of the ARMA(1,1) structure is needed for inflation to predict future consumption over longer horizons in a time-varying way.

The third ingredient is that innovations to inflation and consumption have a time-varying

conditional covariance - the model's theoretical counterpart of the nominal-real covariance - which is needed for inflation to predict future consumption in a time-varying way.

The fourth ingredient is a representative agent with recursive Epstein-Zin-Weil (EZ) utility. With EZ preferences, shocks to expectations about future consumption growth are priced in addition to shocks to consumption growth itself. Since inflation predicts consumption growth, inflation shocks are priced in our model. This property of EZ utility is explored by many authors in the macro-finance literature.

Our empirical results are robust in a number of important dimensions. First, aside from sorting stocks on their exposure to inflation, we present consistent evidence on the time-varying price and quantity of inflation risk for a maximum-correlation inflation-mimicking portfolio (Breedon et al., 1989) as well as the inflation risk premium estimated using a Fama-MacBeth regression for individual stocks. Second, our evidence does not depend on a particular measure of the nominal-real covariance, as we show quantitatively and qualitatively similar results using either the covariance between inflation and future industrial production growth, which effectively treats inflation as a recession state variable of the type advocated in Cochrane (2005, Ch.9) and Kojien et al. (2013), or using the stock market beta of the long-term Treasury bond, as in Campbell et al. (2013). Third, although inflation betas contain a large industry component, our evidence on the inflation risk premium is robust when we use only within- or across-industry variation in exposures. This evidence highlights that even within an industry there is enough information to identify the inflation risk premium. Fourth, all of our results carry through when we control for the usual asset-pricing factors.³ Finally, we perform a truly out-of-sample exercise using real-time inflation, inspired by the approach of Ang et al. (2012).

Our contribution to the literature is in establishing that the cross-section of stock returns contains an inflation risk premium that is varying considerably over time. Chen, Roll and Ross

³The benchmark factors are MKT, SMB, HML, and MOM, combined in the CAPM (Sharpe (1964), Lintner (1965) and Mossin (1966)), the Fama-French three-factor model (Fama and French (1993)), and the Fama-French-Carhart model (Carhart (1997)).

(1986), Ang et al. (2012) and Duarte and Blomberger (2012) also estimate the inflation risk premium in the cross-section of portfolios or individual stocks, but do not analyze whether the risk premium is time-varying. We extend recent bond market evidence in arguing that the nominal-real covariance is an important driver of time series variation in the inflation risk premium and, in particular, of its recent reversal in sign. Campbell et al. (2013) show that nominal bond risks also vary in time with the nominal-real covariance and document a consistent change in sign of term premiums in U.S. government bond yields over the last decade. Campbell et al. (2014) analyze the monetary policy drivers of these changes in bond risk (premiums). Furthermore, Kang and Pflueger (2013) find that the nominal-real covariance affects corporate bond yields in six developed economies through a credit channel above and beyond government bond yields, whereby a firm's real liabilities and default rates change with inflation. Our findings show similar effects for the stock market.

The remainder of the paper is organized as follows. In Section I we describe how to measure inflation risk in the cross-section of stocks. In Section II we estimate the inflation risk premium, both unconditionally and conditional on the nominal-real covariance. In Section III we explain our results in the context of an equilibrium model in which the relation between inflation and consumption is time-varying. In Section IV we calibrate the model to the data. In Section V we present a number of robustness checks. We conclude in Section VI.

2 Inflation Risk

In this section we describe our data sources, measure inflation risk in the cross-section of stock returns, and analyze the time-varying relation between inflation and real consumption growth.

2.1 Data

Monthly inflation (Π_t) is the percentage change in the seasonally-adjusted Consumer Price Index for All Urban Consumers (CPI) available from the U.S. Bureau of Labor Statistics.⁴ We measure monthly nominal consumption (C_t) growth using the seasonally-adjusted aggregate nominal consumption expenditures on nondurables and services from the National Income and Product Accounts (NIPA) Table 2.8.5. Population numbers come from NIPA Table 2.6 and price deflator series from NIPA Table 2.8.4, which we use to construct the time series of per capita real consumption growth. Seasonally-adjusted industrial production growth is from the *FRED*[®] database of the St. Louis FED and the ten-year constant maturity treasury bond return is from CRSP. In our asset pricing tests, we use all ordinary common stocks traded on the NYSE, AMEX, and NASDAQ (excluding firms with negative book equity) from CRSP. The CRSP value-weighted market portfolio, the one-month t-bill return, benchmark asset pricing factors, and industry portfolios are from Kenneth French’s website. Table 1 presents descriptive statistics for the sample over which we run our asset pricing 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 most empirical studies of the cross-section.

2.2 Inflation Betas

At the end of each sample month t , we measure the exposure of firm i to inflation 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 (2001), we filter out these innovations, denoted $u_{\Pi,t}$, using an ARMA(1,1)-model.

We estimate inflation betas using a rolling weighted least-squares (WLS) regression over an expanding window that uses all observations from the first month the stock is included in

⁴In a robustness check (see Table 12), we perform a truly out-of-sample exercise using inflation in the real-time vintage CPI series, as in Ang et al. (2012).

the sample up to month t . The WLS weights are exponentially decaying in their distance to t . The expanding window ensures that we use as much information as possible, whereas an exponential decay in weights ensures that the estimated exposure gives most weight to the most recent information. We require that stocks have at least 24 out of the last 60 months of returns available and use past information only in the estimation. Thus, the estimator of a stock's inflation risk, $\beta_{\Pi,i,t}$, is given by

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

$$\text{with weights } K(\tau) = \frac{\exp(-|t - \tau| h)}{\sum_{\tau=1}^{t-1} \exp(-|t - \tau| h)} \text{ and } h = \frac{\log(2)}{60}, \quad (2)$$

so that the half-life of the weights $K(\tau)$ converges to 60 months for large t . Following Elton et al. (1978) and Cosemans et al. (2012), we transform the estimated $\widehat{\beta}_{\Pi,i,t}$ using a Vasicek (1973) adjustment

$$\widehat{\beta}_{\Pi,i,t}^v = \widehat{\beta}_{\Pi,i,t} + \frac{\text{var}_{TS}(\widehat{\beta}_{\Pi,i,t})}{\left[\text{var}_{TS}(\widehat{\beta}_{\Pi,i,t}) + \text{var}_{CS}(\widehat{\beta}_{\Pi,i,t})\right]} \left[\text{mean}_{CS}(\widehat{\beta}_{\Pi,i,t}) - \widehat{\beta}_{\Pi,i,t}\right], \quad (3)$$

where the subscripts TS and CS denote means and variances taken over the time-series and cross-sectional dimensions, respectively. The $\widehat{\beta}_{\Pi,i,t}^v$ is a weighted average of the stock-specific beta estimated in the time series and the average of all betas in month's t cross-section, where the former receives a larger weight when it is estimated more precisely.

We first create thirty value-weighted portfolios by performing a two-way sort of all stocks into portfolios at the intersection of ten inflation beta deciles and three size groups. The size groups are defined by the 20th and 50th percentiles of last month's NYSE market capitalizations (the Micro, Small, and Big groups of Fama and French (2008)). We then collapse the thirty portfolios into ten size-controlled inflation beta-sorted portfolios by averaging over the three size groups in each inflation beta decile. Our choice for size as a control variable is motivated

by Ang et al. (2012), who find that the best inflation hedgers are the smallest stocks.⁵ From this point forward, inflation betas refer to the WLS and Vasicek-adjusted betas and we drop the hat and superscript v .

We show in Section V that our conclusions are robust to estimating Equation (1) controlling for benchmark asset-pricing factors and using, instead of inflation innovations, either inflation itself (following Ang et al. (2012)) or the difference between inflation and the short-term Treasury-bill return (following Fama and Schwert (1977)). Estimating 60-month rolling window OLS betas also gives comparable, albeit noisier, results.⁶

2.3 The Time-Varying Relation Between Inflation and Consumption

Inflation is a risk variable when it indicates good or bad news about the state of the economy. To this end, in this section we analyze the relation between inflation and future consumption growth. In Panel A of Table 2, we present the results from simple regressions of future consumption growth from month $t + 1$ to $t + K$ on inflation over month t :

$$\Delta C_{t+1:t+K} = d_0^K + d_1^K \Pi_t + e_{t+1:t+K}. \quad (4)$$

We see that the unconditional relation between inflation and consumption is negative, consistent with previous evidence in, e.g., Piazzesi and Schneider (2006).⁷ Over our sample period from July 1967 to December 2014, however, the coefficient on inflation is not significantly different from zero at any frequency.

As already noted in Bekaert and Wang (2010) and Campbell et al. (2014), among others, this unconditional regression masks important variation over time. To quantify this time-variation

⁵We find that our conclusions are robust to controlling for book-to-market or momentum instead. These results are available upon request from the authors.

⁶These results are available upon request from the authors.

⁷In the data, we lag inflation by an additional month, to account for the reporting delay of CPI numbers.

and demonstrate its economic magnitude, we perform the following two-stage test:

$$\Delta C_{t+1:t+K} = d_0^{c,K} + d_1^{c,K}(\widehat{a_{t-1}^K} + \widehat{b_{t-1}^K}\Pi_t) + e_{t+1:t+K}, \text{ where} \quad (5)$$

$$\Delta C_{s+1:s+K} = a_{t-1}^K + b_{t-1}^K\Pi_s + e_{s+1:s+K}, \quad s = 1, \dots, t - K. \quad (6)$$

In the first stage, Equation (6) regresses consumption growth on lagged inflation over a backward-looking window using all data available up to month t (estimated using weighted least squares with weights identical to Equation (2)). Hence, the window s runs from 1 to $t - K$. In the second stage, Equation (5) uses the estimated coefficients and inflation observed at time t , i.e., $\widehat{a_{t-1}^K} + \widehat{b_{t-1}^K}\Pi_t$, to predict consumption growth from month $t + 1$ to $t + K$. This setup ensures that we use no forward-looking information when we predict consumption growth in the second stage. If this 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. First, $d_1^{c,K}$ is significantly larger than zero at all horizons and the R^2 increases from 3% at the one month horizon to a considerable 15% at the one year horizon, implying that inflation predicts consumption growth. In fact, for no horizon K can we reject the hypothesis that $d_0^{c,K} = 0$ nor that $d_1^{c,K} = 1$, which suggests that this structure succeeds in modelling the conditional expectation of consumption growth. We also calculate an out-of-sample R^2 ($R2 - OOS$) to ensure that the predictive performance does not come from the constant term a_{t-1}^K alone, which may measure lagged average consumption growth, but also from the nominal-real covariance, b_{t-1}^K . The $R2 - OOS$ is estimated as is usual in predictive regressions in the literature (see, e.g., Goyal and Welch (2008)):

$$R2 - OOS = 1 - \frac{Var(\Delta C_{t+1:t+K} - (a_{t-1}^K + b_{t-1}^K\Pi_t))}{Var(\Delta C_{t+1:t+K} - a_{t-1}^{*,K})}. \quad (7)$$

Here, $a_{t-1}^{*,K}$ is estimated in a backward-looking window regression of consumption growth on a constant following Equation 6, but leaving out lagged inflation. We find that the $R2 - OOS$ is

similarly increasing from 2% at the one month horizon to 10% at the one year horizon. These results imply that inflation is a potent predictor of consumption growth once accounting for a time-varying nominal-real covariance.

3 The Inflation Risk Premium

Our main objective is to analyze the inflation risk premium in the stock market. We use three standard approaches to measure the inflation risk premium and its variation over time. The first estimate of the inflation risk premium is the size-controlled High-minus-Low inflation beta portfolio (denoted HLIP).

The second estimate of the inflation risk premium is the maximum correlation inflation-mimicking portfolio (denoted MCIP) of Breeden et al. (1989). We estimate this portfolio by projecting inflation innovations on the space of ten size-controlled inflation beta-sorted portfolios:

$$u_{\Pi,t+1} = \textit{intercept} + \textit{weights}' \times R_{t+1} + e_{t+1}, \quad (8)$$

where $R_{t+1} = (R_{High,t+1}, \dots, R_{Low,t+1})'$, such that MCIP is the portfolio return $\textit{weights}' \times R_{t+1}$. We use the ten size-controlled inflation beta-sorted portfolios as they should contain a large share of the information in $u_{\Pi,t+1}$ that is relevant for the cross-section of stock returns.

Finally, we run a cross-sectional regression of monthly individual stock returns on lagged inflation betas, where we control for market capitalization to be consistent with our sort:

$$R_{i,t+1} = l_{0,t} + l_{\Pi,t} \beta_{\Pi,i,t} + l_{MV,t} MV_{n,t} + u_{t+1}. \quad (9)$$

The time-series of coefficient estimates $l_{\Pi,t}$ represent the third estimate of the inflation risk premium and the cross-sectional inflation premium (denoted CSIP). As shown in Fama (1976), CSIP captures the return of a zero-investment portfolio with pre-ranking inflation beta exactly equal to one.

3.1 The Unconditional Inflation Risk Premium

Table 3 describes the set of ten inflation beta-sorted portfolios, as well as the three estimates of the inflation risk premium: HLIP, MCIP and CSIP. MCIP and CSIP are scaled so that they have identical ex post inflation exposure to HLIP.⁸

Panel A first reports the ex-post inflation exposures of these portfolios. The exposures are estimated with a simple regression of portfolio returns on inflation innovations over the full sample. Analyzing whether these exposures are large, economically and statistically, is important to test whether inflation is a useless factor in the sense of Kan and Zhang (1999) and also as a reality check of the estimation procedure. The ex post exposures line up almost monotonically from High to Low and the dispersion is wide, giving a post-ranking beta of 3.0 for HLIP. This ex post exposure is significant and economically large, translating 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 for the same increase in $u_{\Pi,t}$. The remaining rows in Panel A report consistent evidence when we estimate portfolios' inflation exposures over an expanding window - by applying Equation (1) to the post-ranking returns of the portfolios - and calculate the average and standard deviation of these post-ranking rolling inflation betas. It is important to note that these rolling inflation betas are subject to substantial time-variation, which we will address in more detail below.

Panel B presents our estimates of the unconditional inflation risk premium. The results show that the average returns for the ten decile portfolios are almost monotonically decreasing in inflation beta from 9.49% for Low to 5.26% for High. This dispersion translates to a significant annualized average excess return for HLIP equal to -4.23% ($t = -2.08$). This estimate is robust in size and significance across the various methods with annualized average excess returns for MCIP and CSIP equal to -4.39% ($t = -3.06$) and -4.21% ($t = -2.22$), respectively. These

⁸For both MCIP and CSIP, the scaling factor is 3.00 divided by the post-ranking inflation beta of the respective portfolio.

estimates translate to a price of inflation risk as measured by these portfolio's Sharpe ratios equal to -0.30, -0.44 and -0.32, respectively. These numbers are comparable in magnitude to the Sharpe ratio of the market portfolio, which suggests that inflation is an important determinant of cross-sectional variation in stock returns.

A negative unconditional inflation risk premium is consistent with an unconditionally negative relation between inflation and consumption. Indeed, if a shock to inflation is bad news for investors, they should pay high prices for high beta stocks. As a result, high beta stocks should have lower expected returns. Previous estimates of the inflation risk premium in the literature are also negative using a small set of stock portfolios (Chen et al. (1986) and Ferson and Harvey (1991)) and using nominal and real bonds (e.g., Buraschi and Jiltsov (2005), Gurkaynak et al. (2010), Ang et al. (2008), and D'Amico et al. (2008)). Using the CRSP universe, Ang et al. (2012) estimate an unconditional inflation risk premium that is negative, but insignificant at minus one percent.

3.2 Time-Variation in the Inflation Risk Premium

Even though we estimate an unconditional inflation risk premium of about minus four percent per year, this estimate may mask significant variation over time. Piazzesi and Schneider (2006) show that investors dislike inflation not only because it lowers payoffs of nominal bonds, but also because it carries bad news for future consumption growth. Campbell et al. (2013) find that inflation risk premiums in bonds, as well as the correlations between stocks and bonds, vary over time, where the variation is driven by the nominal-real covariance of inflation with the real economy. Hence the risk premium for inflation risk should depend on whether inflation signals bad or good news for the economy. Consistent with the recent reversal in the covariance between inflation and the real economy (Bekaert and Wang (2010), Campbell et al. (2013), and David and Veronesi (2014)), the inflation risk premium in stocks, as documented in Table 3, may therefore also change with the nominal-real covariance. Therefore, we next estimate the

inflation risk premium conditional on the nominal-real covariance.

We report in Table 4 coefficient estimates from regressions of the inflation risk premium on the nominal-real covariance. Our main measure of the nominal-real covariance is the time-varying relation between inflation and future twelve-month consumption growth, as discussed above. We denote this as NRC_t^C , which is the slope coefficient b_t^K for $K = 12$ in the regression in Equation (6):

$$\Delta C_{s+1:s+12} = a_{t-1}^{12} + b_{t-1}^{12} \Pi_s + e_{s+1:s+12}, \quad s = 1, \dots, t - 12, \quad (10)$$

$$NRC_t^C = b_{t-1}^{12}.$$

Our setup ensures that we are not using any forward-looking information when predicting returns on inflation portfolios (compounded over horizons H of one, three, and twelve months) using:

$$R_{p,t+1:t+H} = L_0 + L_{NRC} NRC_t^C + \varepsilon_{t:t+H}. \quad (11)$$

For each horizon H , Panel A presents the estimated coefficients (annualized), Newey-West t -statistics (with H lags), and the adjusted R^2 's for the individual decile portfolios and the three inflation portfolios of interest ($p = HLIP, MCIP, CSIP$). NRC is standardized to mean equal to zero and standard deviation equal to one, so L_0 is identical to the average excess return of the respective portfolio for $H = 1$ and the coefficient estimate L_{NRC} measures the increase in annualized portfolio return for a one standard deviation increase in the nominal-real covariance.

First, L_{NRC} is decreasing monotonically from High to Low inflation beta. As a result, the effect of NRC_t^C on the inflation risk premium is positive and large, economically and statistically. For a one standard deviation decrease in the nominal-real covariance, the inflation risk premium as estimated using HLIP, MCIP, and CSIP decreases by about 3% to 4% for $H = 1$, which represents roughly a doubling of the portfolio's expected return and thus the price of inflation. The effect is slightly larger for longer horizons, growing to about 4% to 5% for

$H = 12$. The R^2 demonstrates an even stronger horizon effect, increasing from around 1% for $H = 1$ to over 7% for $H = 12$. We conclude that the inflation risk premium varies conditional on the nominal-real covariance.

In Panel B, we analyze whether our results extend 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}). NRC_t^{IP} is estimated by substituting industrial production for consumption on the left-hand side of Equation (10). NRC_t^{-BB} is estimated with a 60-month rolling window regression of the 10-year constant maturity treasury bond return on the CRSP value weighted stock market return, following Campbell et al. (2014). Figure 1 plots these three estimates of the nominal-real covariance and shows strong comovement at low frequencies. Consistent with Campbell et al. (2014), NRC_t^{-BB} has changed sign from negative to positive in the 2000s. Because the measures have different units, we standardize them in the predictive regressions. In short, our evidence is qualitatively and quantitatively robust. We find that each alternative measure predicts the inflation risk premium (as measured by HLIP, MCIP, and CSIP) with a positive coefficient that is comparable in magnitude, significance, and horizon pattern to what we have seen before.

Panel C shows that the predictive relation between the inflation risk premium and the nominal-real covariance is robust over time and exists in both sample halves. Although the effect is stronger in the second half of the sample, it follows from this finding that the time-variation in the inflation risk premium is not a recent phenomenon (see Figure 1).

We conclude that the inflation risk premium is strongly time-varying with the nominal-real covariance. The intuition is that high inflation beta stocks are attractive to hedge consumption risk when inflation predicts consumption with a (large) negative sign as it did historically. However, these same stocks are not attractive as a hedge anymore, and may even expose investors to additional consumption risk once the nominal-real covariance between inflation and consumption starts increasing, as it does towards the end of our sample. As a result,

expected returns should be increasing in the nominal-real covariance, which is what we find. In the following section, we formalize this intuition in a model.

3.3 Time-Variation in Inflation Betas

Given the impact of the nominal-real covariance on the inflation risk premium (the price of risk), a natural question is to ask whether the nominal-real covariance impacts inflation betas (the quantity of risk). To set the stage, Figure 2 depicts the histogram of inflation betas for four different time periods. We have selected December of 1971, 1983, 1994, and 2009 to portray the shape of the distribution of betas in different macroeconomic conditions and inflationary regimes. Inflation betas have significant dispersion in all four time periods, with values ranging from -20 to +10 typically, suggesting that there is a wide spectrum of ex ante inflation betas among individual stocks. Such wide cross-sectional dispersion in inflation betas is instrumental in identifying the inflation risk premium analyzed in the previous section. The mean of the distribution moves considerably through time.

Table 5 analyzes how inflation betas of our inflation portfolios vary over time and in the cross-section with the nominal-real covariance. To this end, we regress the rolling inflation beta of the ten decile portfolios (estimated by applying Equation (1) to the post-ranking returns of the portfolios) on the nominal-real covariance:

$$\hat{\beta}_{\Pi,p,t} = \beta_{p,0} + \beta_{p,NRC} NRC_t^C + \varepsilon_{p,t}. \quad (12)$$

As before, NRC_t^C is standardized to mean equal to zero and standard deviation equal to one, such that $\beta_{p,0}$ is equal to the average rolling inflation beta of the respective portfolio (as reported in Table 3). Analogous to Table 4, the table reports the estimated coefficients, with corresponding t -statistics based on Newey-West standard errors with 60 lags, and the R^2 from each regression.

In short, we find that inflation betas are increasing in the nominal-real covariance, as implied

by the model. For all decile portfolios, the effect is marginally significant and economically large, with an increase in inflation beta of 1.55 on average for a one standard deviation increase in NRC_t^C . Although the effect is not completely monotonic, we find that $\beta_{p,NRC}$ is larger for stocks with relatively low inflation betas. A change of 1.55 is about half of the cross-sectional spread in the ten decile betas reported in Table 3, illustrating the importance of the time-variation in betas. Based on the unconditional price of inflation risk of -4.2% in Table 3 for the HLIP portfolio, with an inflation beta of 3.0, the associated risk premium related to a change in beta of 1.55 is 2.2%.

4 Model

We present a model that takes as given the key empirical interconnections among inflation, consumption, and the nominal-real covariance uncovered in the last section and then reproduces - via pricing through the Euler equation of a representative agent - an equilibrium inflation risk premium that behaves like the one we empirically estimate. The 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, since our focus is on the asset pricing implications of inflation risk.

4.1 Preferences

The representative agent has preferences given by the recursive utility function of Epstein and Zin (1989) and Kreps and Porteus (1978),

$$U_t(W_t) = \left((1 - \delta) C_t^{1-1/\psi} + \delta E_t [U_{t+1}(W_{t+1})^{1-\gamma}]^{\frac{1-1/\psi}{1-\gamma}} \right)^{\frac{1}{1-1/\psi}}, \quad (13)$$

where W_t is real aggregate wealth and C_t is real aggregate consumption. The constant $\delta \in (0, 1)$ is the discount rate, $\gamma > 0$ is the coefficient of relative risk aversion and $\psi > 0$ is the elasticity

of intertemporal substitution (EIS). It is convenient to define the constant $\theta = \frac{1-\gamma}{1-1/\psi}$, which measures the magnitude of risk aversion relative to the EIS. The first order condition for the representative agent's problem implies that the gross return⁹ $R_{i,t+1}$ on any tradable asset i satisfies the Euler equation

$$1 = E_t [M_{t+1} R_{i,t+1}], \quad (14)$$

with a stochastic discount factor given by

$$\log M_{t+1} = m_{t+1} = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1) r_{c,t+1}, \quad (15)$$

where lowercase letters denote logarithms, so that $\Delta c_t = \ln C_t - \ln C_{t-1}$ and $r_{c,t} = \log(R_{c,t})$.

4.2 Dynamics of the economy

The processes for real consumption growth, Δc_t , inflation, π_t , the nominal-real covariance, φ_t , and real dividend growth for asset i , $\Delta d_{i,t}$, are exogenous and given by

$$\pi_{t+1} = \mu_\pi + \rho_\pi (\pi_t - \mu_\pi) + \phi_\pi u_{t+1} + \xi_\pi u_t, \quad (16)$$

$$\Delta c_{t+1} = \mu_c + \rho_c (\pi_t - \mu_\pi) + \sigma_c \eta_{t+1} + \varphi_t u_{t+1} + \xi_c \varphi_{t-1} u_t, \quad (17)$$

$$\varphi_{t+1} = \varphi_0 - v (\varphi_t - \varphi_0) + \sigma_w w_{t+1}, \quad (18)$$

$$\Delta d_{i,t+1} = \mu_i + \rho_i (\pi_t - \mu_\pi) + \sigma_i \eta_{t+1} + \phi_i \varphi_t u_{t+1} + \xi_i \varphi_{t-1} u_t. \quad (19)$$

The shocks u_t , η_t and w_t are i.i.d. standard normal. Equation (16) shows that, as in our empirical analysis, inflation follows an *ARMA*(1, 1) process. The constant μ_π is the unconditional mean of π_t , ρ_π is its *AR*(1) coefficient and ϕ_π , ξ_π control its *MA*(1) component. Estimating this univariate process to extract the path for u_t , as we did in the previous section and as we do in our calibration section below, gives positive values for ρ_π and ϕ_π , and a negative one for

⁹Note that in the empirical analysis $R_{i,t+1}$ refers to excess returns.

ξ_π .

The process for Δc_{t+1} in Equation (17) has several features. The term μ_c is the unconditional mean of Δc_{t+1} . The innovation η_{t+1} represents a shock to the real economy that is orthogonal to inflation and nominal-real covariance shocks, with homoskedastic impact on consumption growth controlled by σ_c . The rest of the terms in Equation (17) capture the effect that inflation and inflation shocks have on the real economy. These inflation non-neutralities have two components, an unpredictable and a predictable one. The unpredictable component is given by $\varphi_t u_{t+1}$, where the nominal-real covariance φ_t follows the mean-reverting process given by Equation (18). As is the case in many similar models, a stochastic φ_t implies that consumption growth has stochastic volatility. However, unlike standard stochastic volatility models, φ_t can change signs, so the conditional covariance between inflation and consumption growth,

$$cov_t(\Delta c_{t+1}, \pi_{t+1}) = \phi_\pi \varphi_t,$$

can also change signs (the conditional variance of Δc_{t+1} is guaranteed to always remain positive though). The predictable component of inflation non-neutralities in Equation (17) is $\rho_c(\pi_t - \mu_\pi) + \xi_c \varphi_{t-1} u_t$. All predictability of Δc_{t+1} comes from this component. Unconditional predictability - which empirically is found to be negative - can be seen from running, in the model, a regression of Δc_{t+k} on π_t analogous to the empirical one we run in Panel A of Table 2. The coefficient on π_t in this regression is

$$\frac{Cov(\Delta c_{t+k}, \pi_t)}{Var(\pi_t)} = q_k \rho_c + \frac{(1 - \rho_\pi^2) \xi_c \phi_\pi \varphi_0}{\phi_\pi^2 + \xi_\pi^2} \quad (20)$$

where q_k is a constant that depends on the parameters of the inflation process (ρ_π , ϕ_π , ξ_π) and is decreasing in the predictive horizon k .

The process for consumption also implies conditional predictability of consumption growth with inflation that is different from the unconditional predictability just examined. The con-

ditional predictability captures three key features of the relation between inflation and consumption discussed in the previous section. First, higher inflation can predict higher or lower future consumption growth depending on the time period examined. Second, the sign and magnitude of this conditional predictability are persistent over time. Third, at each period in time, predicting future consumption growth over horizons ranging from one to twelve months gives almost identical results. To see that these three features are indeed present in our setup, we run in the model a two-step predictive procedure analogous to the one in Panel B of Table 2. In the first step, we run a conditional regression of cumulative consumption growth over the next $k = 1, 2, \dots, 12$ months on inflation at time t . The model-implied coefficient on inflation in such a regression is

$$\theta_t^{(k)} = \frac{Cov_{t-1} \left(\sum_{j=1}^k \Delta c_{t+j}, \pi_t \right)}{Var_{t-1} (\pi_t)} = h_k + \frac{\xi_c}{\phi_\pi} \varphi_t, \quad (21)$$

where h_k is a constant that depends on the prediction horizon k . Equation (21) shows that $\theta_t^{(k)}$ can change signs and is persistent because φ_t has these exact properties. In addition, for reasonable parameters (such as the ones we use in the next section), h_k slowly decreases with k , producing coefficients $\theta_t^{(k)}$ that are roughly constant in k for the k 's we consider. This first-step regression makes clear why we call φ_t the nominal-real covariance: Under the null of the model, φ_t is a linear transformation of $\theta_t^{(k)}$, the same object we use to measure the nominal-real covariance in the data.

In the second step of the predictive procedure, we run an unconditional regression of cumulative consumption growth on its value predicted from the first step regression. In the model, the coefficient of this second-step regression is equal to one and the intercept is equal to zero. This result is consistent with the failure to reject the hypothesis that d_1^c is not equal to one and d_0^c is not equal to zero in Table 2.

The dynamics of real dividend growth for asset i is given by Equation (19). We model dividends as levered consumption, whereby dividends are subject to the same risks as consumption but with potentially different exposures. In the present context, asset i represents

one of the inflation-sorted portfolios from the last section. To derive the model-implied inflation risk premium from the cross-section of stock returns, it is enough to consider two assets, $i = H, L$. When we calibrate the model, these two assets will be mapped to the highest and lowest inflation beta portfolios that we constructed in the last section.

4.3 Wealth-consumption ratio

In equilibrium, the log wealth-consumption ratio is linear-quadratic in the state variables π_t , φ_t and u_t :

$$wc_t = A_0 + A_1\pi_t + A_2\varphi_{t-1}u_t + A_3u_t + A_4(\varphi_t - \tilde{\varphi}_0)^2, \quad (22)$$

where $\tilde{\varphi}_0$ is a constant close to φ_0 given in the Appendix.¹⁰ The loadings on the state variables are

$$A_1 = \left(1 - \frac{1}{\psi}\right) \frac{\rho_c}{1 - \kappa_1\rho_\pi}, \quad (23)$$

$$A_2 = \left(1 - \frac{1}{\psi}\right) \xi_c, \quad (24)$$

$$A_3 = \left(1 - \frac{1}{\psi}\right) \frac{\kappa_1\xi_\pi\rho_c}{1 - \kappa_1\rho_\pi}, \quad (25)$$

$$A_4 = \frac{1}{2\theta} \frac{\left(\left(\theta - \frac{\theta}{\psi}\right) + \theta A_2 \kappa_1\right)^2}{1 - \kappa_1 v^2 (\sigma_w + 1)}, \quad (26)$$

with $\kappa_1 \in (0, 1)$ a linearization constant.

The intuition behind A_1 , A_2 and A_3 is identical to that of the standard long-run risk model. When the EIS is greater than one, the intertemporal substitution effect dominates the wealth effect. In this case, which we henceforth assume, higher expected consumption growth $E_t[\Delta c_{t+1}]$ leads the representative agent to invest more, increasing the wealth-consumption

¹⁰It is not necessary to introduce this constant, but it substantially simplifies the framework. Indeed, having φ_t and φ_t^2 as two separate state variables instead of the single $(\varphi_t - \tilde{\varphi}_0)^2$ gives the same results by completing the square without the need for $\tilde{\varphi}_0$.

ratio. It follows that any state variable that increases (decreases) expected consumption growth has a positive (negative) loading in Equation (22). If $\rho_c < 0$, which is the relevant empirical case, higher π_t is bad news for expected consumption growth and hence $A_1 < 0$. The higher the persistence of inflation, ρ_π , and the higher the EIS, the stronger is the impact of π_t on wc_t (A_1 is more negative). Similarly, $\xi_c \varphi_{t-1} u_t$ is also part of expected consumption growth, which gives rise to the expression in A_2 . As u_t is *i.i.d.*, there is no multiplier akin to the $(1 - \kappa_1 \rho_\pi)^{-1}$ for A_1 . Even though A_2 is constant, because of the time-varying nominal-real covariance, a positive inflation shock u_t is sometimes good news and sometimes bad news for expected consumption and valuation ratios, which is captured by the cross-term $\varphi_{t-1} u_t$ in Equation (22). In our calibration, we find evidence that $\xi_c > 0$, which implies $A_2 > 0$. A positive inflation shock is then a good state of nature when the nominal-real covariance is positive and a bad state of nature otherwise.

The loading A_3 appears in Equation (22) because the shock u_t influences expected consumption growth indirectly through inflation: u_t affects π_{t+1} which, in turn, affects tomorrow's expected consumption growth $E_{t+1}[\Delta c_{t+2}]$. Because inflation is persistent, not only is tomorrow's expected consumption growth changed through this channel, but also the whole of its future path. The impact of u_t on inflation today is regulated by ξ_π while the impact of π_{t+1} on wc_{t+1} arising from all changes in future expected consumption is A_1 , as in our earlier discussion. The product of these two values, discounted by κ_1 to bring the effect on wc_{t+1} back to time t , produces A_3 .

The loading A_4 can be understood by recognizing that one of the roles of the nominal-real covariance is that of stochastic volatility of consumption. The particular expression for A_4 is close to the familiar one in the long-run risk literature, yet not identical. The differences arise because we model the volatility φ_t as an *AR*(1) instead of the variance φ_t^2 and because we have stochastic volatility in the moving average component in the consumption process (17) instead of in Bansal and Yaron's long-run risk process - which is absent in our model. On the other hand, the moving average component matters for A_4 for the same reason that stochastic

volatility of long-run risk matters for consumption: higher stochastic volatility for Δc_{t+1} today also means a more volatile expected consumption growth at $t + 1$ from the point of view of time t . Thus, despite the differences with the benchmark long-run risk model, the sign of the stochastic volatility loading A_4 is still determined by the sign of θ , so that if the EIS ψ and the CRRA coefficient γ are both greater than one, higher uncertainty is detrimental for asset prices and leads to a lower wealth-consumption ratio. The fact that the nominal-real covariance can change signs is unimportant for this channel.

4.4 Stochastic Discount Factor and Prices of Risk

Armed with the wealth-consumption ratio equation and its loadings, we can now examine the prices of risk in the economy. The innovation in the stochastic discount factor is given by

$$m_{t+1} - E_t[m_{t+1}] = -\lambda_{u,t}u_{t+1} - \lambda_{w,t}w_{t+1} - \lambda_{2w}w_{t+1}^2 - \lambda_\eta\eta_{t+1} \quad (27)$$

where the the prices of risk are

$$\lambda_{u,t} = \kappa_1(1 - \theta)(A_3 + \phi_\pi A_1) + [(1 - \theta)\kappa_1 A_2 + \gamma]\varphi_t \quad (28)$$

$$\lambda_{w,t} = -2\kappa_1\sigma_w(1 - \theta)A_4v(\varphi_t - \varphi_0) \quad (29)$$

$$\lambda_{2w} = \kappa_1\sigma_w^2(1 - \theta)A_4 \quad (30)$$

$$\lambda_\eta = \sigma_c\gamma \quad (31)$$

The price of inflation risk, $\lambda_{u,t}$, is a linear function of the nominal-real covariance φ_t , and so can change signs over time. In equilibrium, investors sometimes require compensation for bearing inflation risk because inflation shocks are a harbinger of poor future consumption growth. Other times, investors are willing to accept lower returns in order to hold inflation risk because it is a good hedge against bad macroeconomic outcomes. In addition to φ_t , the magnitude

and standard deviation of $\lambda_{u,t}$ are determined by preference parameters and the relationship between consumption growth, inflation and inflation shocks. The sign and magnitude of the unconditional mean of $\lambda_{u,t}$ depend on how strongly and in what direction inflation predicts expected consumption growth in an unconditional way. A negative mean for $\lambda_{u,t}$ reflects that, unconditionally, higher inflation is bad news for consumption growth and hence a bad state of the economy. The higher the persistence and volatility of inflation, or the stronger the negative predictability of consumption with inflation, the larger this effect. The same is true if risk aversion or the EIS increase, as the representative agent becomes less tolerant of expected consumption growth risk.

The time-varying part of $\lambda_{u,t}$ arises for two different reasons. First, because consumption growth has a direct exposure to contemporaneous inflation shocks, inflation shocks are priced. The last term in Equation (28), $\gamma\varphi_t$, captures the compensation for this type of risk. This term would be present even in the case of power utility, when the consumption-CAPM holds, as it represents short-term consumption volatility risk. In the case of power utility, however, none of the other terms of $\lambda_{u,t}$ would be present and the standard deviation of $\lambda_{u,t}$ would be much smaller than the one we estimate empirically unless risk aversion is set to be unreasonably high — another manifestation of the equity premium puzzle. The second reason for the presence of a time-varying term in $\lambda_{u,t}$ is that today's inflation shocks have a direct effect on tomorrow's expected consumption growth, a result of having the moving-average component of inflation be also present in the process for consumption growth. The size of the time-varying term, and hence of the volatility of $\lambda_{u,t}$, is determined by the volatility of φ_t , the risk aversion and EIS of the representative agent and the degree to which lagged inflation shocks move consumption growth.

The price of nominal-real covariance risk is given by $\lambda_{w,t}$ and λ_{2w} . The λ_{2w} component reflects compensation for pure stochastic volatility risk in consumption growth. This price of risk is analogous to the standard price of stochastic volatility risk in the baseline long-run risk model. It is the risk that the volatility of consumption will change in a persistent way, either

upwards or downwards. It is constant because the volatility-of-volatility σ_w is constant and because it does not capture the risks from sign changes in the nominal-real covariance, which are instead priced by $\lambda_{w,t}$. For example, the compensation for bearing the risk that, sometime in the future, a positive inflation shock will change from being a good shock for the economy to a bad one is included in $\lambda_{w,t}$. Compared to $\lambda_{u,t}$, which prices inflation risk for a given value of φ_t , $\lambda_{w,t}$ prices the risk of φ_t changing over time. Lastly, λ_η is the price of short-run consumption risk brought about by standard consumption CAPM logic.

4.5 Inflation-sorted Portfolios

Because dividends are exposed to the same risks as consumption, the equilibrium log price-dividend ratio for portfolio i has the same form as the log wealth-consumption ratio:

$$pd_{i,t} = D_{0i} + D_{1i}\pi_t + D_{2i}\varphi_{t-1}u_t + D_{3i}u_t + D_{4i}(\varphi_t - \tilde{\varphi}_{0i})^2, \quad (32)$$

where, as before, $\tilde{\varphi}_{0i}$ is a constant close to φ_0 . The loadings relevant to the analysis of the inflation risk premium in the cross-section of stocks are

$$D_{1i} = \left(\frac{\rho_i}{\rho_c} - \frac{1}{\psi} \right) \frac{\rho_c}{1 - \rho_\pi \kappa_{1,i}}, \quad (33)$$

$$D_{2i} = \left(\frac{\xi_i}{\xi_c} - \frac{1}{\psi} \right) \xi_c, \quad (34)$$

$$D_{3i} = \left(\frac{\rho_i}{\rho_c} - \frac{1}{\psi} \right) \frac{\kappa_{1,i} \xi_\pi \rho_c}{1 - \rho_\pi \kappa_{1,i}}, \quad (35)$$

and we give D_{4i} , the loading on the nominal-real covariance, in the Appendix. The intuition behind these loadings is similar to that of the wealth-consumption loadings. The only difference is that they depend on how inflation and inflation shocks affect dividend growth not in absolute terms but relative to how they affect consumption growth, since the representative agent evaluates stocks not by their outright exposures but by their ability to hedge consumption risk.

For example, in Equation (33), it is not the absolute exposure ρ_i of expected dividend growth to inflation that matters, but its magnitude relative to the exposure of consumption, ρ_c .

Furthermore, portfolio i is exposed to unexpected dividend growth risk, which does not influence the price-dividend ratio (as unexpected dividends have a one-to-one effect on the price) but does contribute to the overall risk of the portfolio. This risk — together with the expected dividend growth risk — is nevertheless manifested in returns space. The equilibrium risk premium for portfolio i implied by its Euler equation is

$$-Cov_t(m_{t+1}, r_{i,t+1}) = \lambda_{u,t}\beta_{ui} + \lambda_{w,t}\beta_{wi} + \lambda_{2w}\beta_{2wi} + \lambda_{\eta}\beta_{\eta i}, \quad (36)$$

where the β_i are portfolio-specific quantities of risk given by the correlation between a portfolio's return and the corresponding sources of risk. For our analysis, the most important is

$$\beta_{u,it} = \frac{\left(\rho_i - \frac{\rho_c}{\psi}\right) (\phi_{\pi} + \xi_{\pi}\kappa_{1,i}) \kappa_{1,i}}{1 - \rho_{\pi}\kappa_{1,i}} + \left(\phi_i + \left(\xi_i - \frac{\xi_c}{\psi}\right) \kappa_{1,i}\right) \varphi_t. \quad (37)$$

The term $\phi_i\varphi_t$ is the quantity of unexpected dividend growth risk that, as explained, adds to the riskiness of portfolio i but does not move its price-dividend ratio. It arises because dividend growth is directly exposed to inflation shocks through $\phi_i\varphi_t u_{t+1}$. The nominal-real covariance determines whether u_{t+1} shocks are good or bad states of nature, and hence also whether the quantity of risk $\phi_i\varphi_t$ is positive or negative. The other terms in $\beta_{u,it}$ are all related to the quantity of expected dividend growth risk. The first term in Equation (37) is the part that arises from inflation entering expected dividend growth through the term $\rho_i\pi_t$ in Equation (19). This risk operates through inflation shocks moving inflation and inflation subsequently altering expected dividend growth, so it is unrelated to the nominal-real covariance. On the other hand, the term $(\xi_i - \xi_c/\psi)\kappa_{1,i}\varphi_t$ gives the amount of u -risk generated by expected dividend growth's direct exposure to the moving-average component in inflation shocks through the term $\xi_i\varphi_{t-1}u_t$ in Equation (19). A positive surprise in u_t can lead to increases or decreases in the quantity of

portfolio i 's risk depending on the sign of φ_{t-1} .

5 Calibration

We calibrate the model in two steps. First, we estimate parameters for inflation and consumption in Equations (16) and (17) using only inflation and consumption data. The estimation by construction matches the persistence of inflation and the unconditional and conditional properties of the predictability of consumption with inflation. The exact procedure is in the Appendix.

In the second step, taking the parameters from the first step as given, we calibrate parameters for preferences, the nominal-real covariance in Equation (18) and the dividend growth processes for inflation portfolios in Equation (19). We calibrate dividends for two portfolios $i = H, L$ that map to the highest and lowest inflation beta portfolios constructed in the last section. The first target in this second step of calibration is the volatility of consumption growth, which depends on consumption parameters already determined in the first calibration step, but also on all parameters of the nominal-real covariance. The low observed volatility of consumption growth places tight restrictions on the volatility of the nominal-real covariance. The second set of targets for our calibration are the first and second moments of returns and inflation betas for the portfolios. Matching the returns of portfolios H and L ensures that the size of the inflation risk premium is consistent with the data. Matching the first moment of betas then gives the appropriate mean quantity of inflation risk. Together, betas and returns that match their empirical counterparts automatically give a model-implied price of risk that is consistent with the data. To capture the correlation between the H and L portfolios, we also target the standard deviation of returns and inflation beta for the HLIP portfolio. Finally, we focus on the coefficient $L_{NRC,HLIP}$ of a predictive regression of HLIP returns on lagged nominal real covariance analogous to the one in Equation (11) to test one of the key equilibrium results of our model, the connection between the exogenous nominal-real covariance and the

endogenous returns.

5.1 Results

Table 6 shows the calibrated parameters and Table 7 gives the resulting moments and their data equivalents. The parameters for inflation imply that it is persistent over business cycle frequencies but not in the longer run, with inflation shocks dissipating almost completely after three years. Expected consumption growth is negatively exposed to inflation, with a one percentage point increase in inflation resulting in a 12.6 basis-point decline in consumption growth next month (plus further declines in subsequent periods due to inflation’s persistence). The contemporaneous correlation between consumption and inflation is small and not a significant driver of our results, neither in the data, nor the model.

For preferences, we use $\psi = 2$ and $\gamma = 14.5$, which are within the range of values used in the literature.¹¹

Table 7 shows that the model can match means and standard deviations of returns and betas for the H and L portfolios. For the HLIP, the model generates a volatility of returns and betas lower than the H and L portfolios, as in the data. The spreads in mean returns and betas between the H and L portfolios are driven mainly by the differences in their ρ_i and ξ_i . For both portfolios, expected dividend growth is negatively exposed to inflation through $\rho_i < 0$ just as consumption growth, although both portfolios are much more exposed than consumption; the low beta portfolio’s exposure is more than one-for-one. On the other hand, the exposure of expected dividend growth to inflation shocks, given by ξ_i , is of opposite sign to ξ_c for both portfolios, which reduces the volatility of inflation betas and provides some hedging against inflation shocks. The constant term in the price of inflation risk $\lambda_{u,t}$ gives a Sharpe ratio of -0.23 , while its time-varying term implies an annualized volatility of 24%.

The last line of the table shows that the predictive regression coefficient $L_{NRC,HLIP}$ is in

¹¹For example, in a closely related model, Bansal and Shaliastovich (2013) use $\psi = 1.81$ and $\gamma = 20.9$.

line with the one found empirically in Table 4. To match this regression coefficient, it is key for this predictability regression to have a volatile and persistent enough nominal-real covariance. In contrast, to match the low observed volatility of consumption growth, ν and σ_w cannot be too large. Matching both moments simultaneously gives quantitative credence to a pricing mechanism operating through the consumption channel proposed in our model. The value of ν implies a half-life for φ_t of ten years, confirming the slow-moving nature of the nominal real covariance and the visual intuition of Figure 1. The calibrated values of ν and σ_w imply a volatility of φ_t comparable to that of real consumption growth, at 1.5% in annualized terms.

Our calibrated parameters differentiate the risks we are considering from those in the long-run risk literature despite the mathematical similarities in our models. The half-life of the predictable component of inflation is one quarter, while it is around ten years in most long-run risk calibrations. The half-life of the stochastic volatility of consumption growth in long-run risk models is usually more than fifty years, compared with ten years for the nominal-real covariance. On the other hand, the volatility of inflation is around a hundred times larger than that of long-run risk, and the volatility of the nominal real covariance is an order of magnitude larger than that of stochastic volatility. The lower persistences and higher volatilities of our setup generate a market price of risk similar to the one for the higher persistence, lower volatility, long-run risk model.

6 Industry Characterization and Robustness

In this section we analyze the relation between inflation sorted portfolios and industries, and we present robustness checks for our main results.

6.1 Characterizing Inflation-Sorted Portfolios

Having seen that inflation betas vary considerably over time and across stocks, we move to an important determinant of a stock's inflation beta: industry affiliation (see also Boudoukh et al. (1994) and Ang et al. (2012)). To this end, we use Kenneth French's classification into 48 industries. In each month of the sample and for each industry, we find the stocks that have inflation beta above and below the cross-sectional median inflation beta and calculate the fraction of the industry market capitalization allocated to each of these two groups. To conserve space, we present in Table 8 the ten industries with on average the best and worst inflation hedging capabilities. Our results are intuitive and in many ways consistent with Ang et al. (2012). Among the best inflation hedgers are industries such as oil, gold, utilities, and mining, for which industries over 64% of the market capitalization (on average) has above median inflation beta. Among the worst inflation hedgers are industries such as banks, insurance, clothes, and textiles, for which industries over 70% of the market capitalization (on average) has below median inflation beta. This industry composition is quite stable over time. Over half of the market capitalization of these same industries is allocated, respectively, to the above and below median inflation beta groups for over 70% of the months in our sample. We conclude that inflation betas contain an industry component. To address this evidence, we ask in the following whether our evidence on the inflation risk premium is driven by within- or across-industry variation, or both.

For the within-industry sort, we construct five market value-weighted stock portfolios within each industry 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. We then collapse the across-industry dimension by calculating the industry inflation beta as the value-weighted average inflation beta of all stocks in that industry and sorting the 48 industries into quintile portfolios. This leaves us with a 5-by-5 within- and across-industry sort. We exclude industry-months that contain fewer than ten stocks.

In Table 9, we replicate our main evidence on post-ranking betas and the inflation risk premium. We first collapse the five-by-five within- and across-industry sort into five within-industry portfolios and five across-industry portfolios. The within-industry portfolios are computed by averaging over five across-industry portfolios for each within-industry quintile. The across-industry portfolios are calculated as the equal-weighted average of the nine or ten industries that belong to the relevant quintile of the across-industry sort. The aggregate within-industry effect is presented in the sixth column and is the difference between the high and low within-industry portfolio return. The across-industry effect is in the twelfth column and is the difference between the high and low across-industry portfolio return.

In short, our main results extend in both dimensions. First, post-ranking beta is increasing in inflation beta both within- and across-industry. Second, the unconditional inflation risk premium is marginally negative both within- and across-industry at -1.71% ($t = -1.38$) and -2.89% ($t = -1.78$), respectively. Third, the nominal-real covariance predicts the returns with a positive and significant coefficient of 3.25% and 4.63% , respectively, at the twelve-month horizon. In both cases, these three results follow from monotonic effects from high to low inflation beta. In summary, these results suggest that although a sort on inflation beta in the CRSP universe contains a strong industry component, our evidence on the inflation risk premium is not solely driven by (across-) industry effects. Thus, variation in inflation beta within industries, perhaps due to differences in corporate hedging practices, market power (see also Weber (2016)), or the place of a firm in the supply chain, is priced in a manner consistent with our hypothesis, even when the industry at large is not strongly exposed to inflation.

6.2 Robustness

In this section, we present a range of robustness checks. From a practical point of view, it is relevant to see the persistence of the post-ranking inflation betas from our sort. To do so, Figure 3 plots the monthly post-ranking inflation beta one month, one year, two years,

five years, and ten years after the sorting date. To calculate these inflation betas, we fix the portfolio composition at the sorting date t and calculate value-weighted returns up to ten years after the sorting date. When a stock leaves the sample, we reallocate its market value across all remaining stocks. 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 monthly inflation exposure for an investor that rebalances infrequently with respect to inflation beta.

We find that our sort is very powerful, as post-ranking inflation beta is almost monotonically decreasing in pre-ranking beta up to ten years after the sort. As a result, the post-ranking beta of the high-minus-low spreading portfolio is also robust, falling from 3.00 one month-after sorting (as reported in Table 3) to a still large and significant 1.68 ten years after sorting. This finding suggests that our inflation betas are estimated with large precision and are persistent. This finding is consistent with the industry composition effect documented above.¹²

Next, we ask whether the inflation risk premium and its variation with the nominal-real covariance are robust to controlling for the standard asset pricing factors of the CAPM, Fama-French three factor model (FF3M) and the Fama-French-Carhart model (FFCM). To this end, we regress the inflation risk premium on NRC_t^C , controlling for contemporaneous exposure to these benchmark factors, following the approach of Baker and Wurgler (2006):

$$R_{p,t+1:t+H} = L_0 + L_{NRC}NRC_t^C + \beta_F'F_{t+1:t+H} + \varepsilon_{t:t+H}, \quad (38)$$

where $F_{t+1:t+H} = [R_{MKT,t+1:t+H}, R_{SMB,t+1:t+H}, R_{HML,t+1:t+H}, R_{MOM,t+1:t+H}]'$. The evidence reported in Table 10 shows qualitatively and quantitatively robust estimates for L_0 and L_{NRC} , which suggests that neither the negative unconditional inflation risk premium nor its positive

¹²This finding is seemingly inconsistent with the conclusion in Ang et al. (2012) that inflation betas are hard to estimate out-of-sample. However, that conclusion is based on a smaller sample of S&P500 stocks 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.

relation with nominal-real covariance is due to benchmark factor exposure. Consistent with this conclusion, Table 11 shows that our results also extend when we control for exposure to the benchmark factors when estimating stock's inflation exposures using Equation (1).

Finally, Table 12 asks whether our results are robust for measures of inflation risk other than ARMA(1,1) innovations. First, we find even stronger evidence that the inflation risk premium is unconditionally negative and time-varying with the nominal-real covariance when we estimate stock's exposures to raw inflation or the difference between inflation and the short-rate, which measure of inflation innovations is used in Fama and Schwert (1977). In fact, the final three columns show that our conclusions on the time-variation with NRC_t^C extend even when we perform a truly out-of-sample sort that uses stock's exposures to real-time vintage CPI inflation, as in Ang et al. (2012). For this exercise, we skip a month after portfolio formation, thus taking into account the reporting delay in real-time inflation data. In this exercise, the unconditional inflation risk premium is slightly smaller than what we have seen before, however.

7 Conclusion

This paper provides new information about inflation risk in the economy by analyzing the cross-section of stock returns. We find a substantial cross-sectional variation in inflation betas, a sizable inflation risk premium, and a strong time-series variation in both inflation exposures and risk premiums, that is driven by the variation in nominal-real covariance.

The unconditional price of inflation risk, a negative -4.2%, is consistent with the fact that inflation predicts consumption with a negative sign historically. As a result, investors are willing to hedge inflation by holding stocks with high inflation betas and hence accept their lower returns. Recent literature finds that the nominal-real covariance is strongly time-varying, and we are the first to show that betas and the inflation risk premium in the stock market are varying over time in a consistent and economically important fashion. This risk premium is determined, in sign and magnitude, by the nominal-real covariance such that a one standard deviation change

in the nominal-real covariance leads to a change in expected returns on inflation mimicking portfolios by 4.5%. Looking at the cross-sectional distribution of the inflation exposures, we find a spread in betas of around three among ten portfolios formed by sorting on inflation betas. Moreover, the betas of those ten portfolios also vary considerably over time, i.e., a one standard deviation change in the nominal-real covariance leads to a change in inflation beta of 1.55 on average.

We develop an equilibrium model which builds on the empirical observation that inflation today predicts real consumption growth in the future. Given that this predictive relation is determined by the time-varying nominal-real covariance, our model generates the dynamics of the prices of inflation risk as well as the quantity of risk consistent with our empirical results.

References

- Ang, A., G. Bekaert, and M. Wei, 2007, Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?, *Journal of Monetary Economics* 54, 1163-1212.
- Ang, A., G. Bekaert, and M. Wei, 2008, The Term Structure of Real Rates and Expected Inflation, *Journal of Finance* 63, 797-849.
- Ang, A., M. Briere, and O. Signori, 2012, Inflation and Individual Equities, *Financial Analysts Journal* 68, 36-55.
- Ang, A., and M. Piazzesi, 2006, A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables, *Journal of Monetary Economics* 50, 745-787.
- Ang, A., M. Piazzesi, and M. Wei, 2006, What does the yield curve tell us about GDP growth?, *Journal of Econometrics* 131, 359-403.
- Bansal, R., and I. Shaliastovich, 2010, A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets, Manuscript, Duke University.
- Bansal, R., and A. Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance* 59, 1481-1509.
- Baker, M., and J. Wurgler, 2006, Investor Sentiment and the Cross-Section of Stock Returns, *Journal of Finance* 61, 1645-1680.
- Bekaert G., and X. Wang, 2010, Inflation Risk and the Inflation Risk Premium, *Economic Policy* 25, 755-806.
- Boudoukh, J., M. Richardson, and R. F. Whitelaw, 1994, Industry Returns and the Fisher Effect, *Journal of Finance* 49, 1595-1615.

- Breeden D. T., M. R. Gibbons, and R. H. Litzenberger, 1989, Empirical Tests of the Consumption-Oriented CAPM, *Journal of Finance* 44, 231-262.
- Buraschi, A., and A. Jiltsov, 2005, Inflation Risk Premia and the Expectations Hypothesis, *Journal of Financial Economics* 75, 429-90.
- Campbell, J.Y., C. Pflueger, and L.M. Viceira, 2014, Monetary Policy Drivers of Bond and Equity Risks, Working Paper Harvard Business School.
- Campbell, J.Y., and R.J. Shiller, The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors, *Review of Financial Studies* 1, 195-228.
- Campbell, J.Y., A. Sunderam, and L.M. Viceira, 2016, Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds, *Critical Finance Review*, forthcoming.
- Campbell, J.Y., and L.M. Viceira, 2001, *Strategic Asset Allocation*, Oxford: Oxford University Press.
- Carhart, M., 1997, On Persistence of Mutual Fund Performance, *Journal of Finance* 52, 57-82.
- Cochrane, J.H., 2005, *Asset Pricing - Revised Edition*, Princeton: Princeton University press.
- Chen, R., B. Liu, and X. Cheng, 2010, Pricing the term structure of inflation risk premia : theory and evidence from TIPS, *Journal of Empirical Finance* 17, 702-721
- Chen, N., R.Roll, and S.A. Ross, 1986, Economic Forces and the Stock Market, *Journal of Business* 59,383-403.
- Cosemans, M., R. Frehen, P.C. Schotman, and R. Bauer, 2016, Estimating Security Betas Using Prior Information Based on Firm Fundamentals, *Review of Financial Studies*, forthcoming.
- Dai, Q., K. Singleton, and W. Yang, 2007, Regime Shifts in a Dynamic Term Structure Model of U.S. Treasury Bond Yields, *Review of Financial Studies* 5, 1669-1706.

- D'Amico, S., D. Kim, and M. Wei, 2008, Tips from TIPS: The Information content of Treasury inflation-protected security prices, Finance and Economics Discussion Series Working Paper 2008-30, Federal Reserve Board.
- David, A., and P. Veronesi, 2014, What Ties Return Volatilities to Price Valuations and Fundamentals?, *Journal of Political Economy* 121, 682-746.
- Duarte, F.M., and J.M. Blomberger, 2012, Cross-sectional Inflation Risk in Menu Cost Models with Heterogeneous Firms, Working Paper Federal Reserve Bank of New York.
- Elton, E.J., M.J. Gruber, and T.J. Urich, 1978, Are Betas Best?, *Journal of Finance* 33, 1375-1384.
- Epstein, L.G., and S.,E., Zin, 1989, Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, *Econometrica*, 57, 937-969.
- Fama, E.F., 1976. Foundations of finance. Basic Books, New York.
- Fama, E.F., and K.R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Econometrics* 33, 3-56.
- Fama, E.F., and K.R. French, 2008, Dissecting Anomalies, *Journal of Finance* 63, 1653-1678.
- Fama, E.F., and M.R. Gibbons, 1984, A Comparison of Inflation Forecasts, *Journal of Monetary Economics* 13, 1327-1348.
- Fama, E.F., and J.D. MacBeth, 1973, Risk, Return, and Equilibrium: Empirical Tests, *Journal of Political Economy* 81, 607-636.
- Fama, E.F., and G.W. Schwert, 1977, Asset Returns and Inflation, *Journal of Financial Economics* 5, 115-146.

- Ferson, W.E., and C.R. Harvey, 1991, The Variation of Economic Risk Premiums, *Journal of Political Economy* 99, 385-415.
- Fuhrer, Jeffrey, and Gerald Moore, 1995, Inflation Persistence, *Quarterly Journal of Economics* 110, 127-160.
- Gabaix X., 2008, Variable Rare Disasters: A Tractable Theory of Ten Puzzles in Macro-finance, *American Economic Review* 98, 64-67.
- Goyal, A., I. Welch, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455-1508.
- Gurkaynak, R., B. Sack, and J. Wright, 2010, The Tips Yield Curve and Inflation Compensation, *American Economic Journal: Macroeconomics* 2, 70-92.
- Haubrich, J., G. Pennacchi, and P. Ritchken, 2012. Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps, *Review of Financial Studies* 25, 1588-1629.
- Kan, R., and C. Zhang, 1999. Two-Pass Tests of Asset Pricing Models with Useless Factors, *Journal of Finance* 54, 203-235.
- Kang, J., and C. Pflueger, 2013, Inflation Risk in Corporate Bonds, *Journal of Finance*, Forthcoming.
- Koijen, R., H. Lustig, H., and S. Van Nieuwerburgh, 2013. The cross-section and time-series of stock and bond returns, Working Paper University of Chicago.
- , Kreps, D., and Porteus, E.L., 1978, Temporal Resolution of Uncertainty and Dynamic Choice Theory, *Econometrica*, 46, 185-200.
- Le, A., K. Singleton, and Q. Dai, 2010, Discrete-Time Dynamic Term Structure Models with Generalized Market Prices of Risk, *Review of Financial Studies* 23, 2184-2227.

- Le, A., and K. Singleton, 2010, An Equilibrium Term Structure Model with Recursive Preferences, *American Economic Review, Papers and Proceedings* May 2010.
- 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.
- Mossin, J., 1966, Equilibrium in a Capital Asset Market, *Econometrica* 34, 768–783.
- Modigliani, F., and R. Cohn, 1979, Inflation, Rational Valuation, and the Market, *Financial Analysts Journal* 35, 24-44.
- Newey, W., and K. West, 1987, A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703-708.
- Piazzesi, M., and M. Schneider, 2006, Equilibrium Yield Curves, in Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, NBER Macroeconomics Annual 2006, published in 2007, Cambridge MA: MIT Press p. 389-442.
- Sharpe, W.F., 1964, Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *Journal of Finance* 19, 425-442.
- Stock, James H., and Mark W. Watson, 2005, Implications of Dynamic Factor Models for VAR Analysis, NBER Working Papers 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.
- Wachter, J.A., 2006, A Consumption-Based Model of the Term Structure of Interest Rates, *Journal of Financial Economics* 79, 365-399.

8 Appendix

8.1 Model Solution

Conjecture that the log wealth-consumption ratio is

$$wc_t = A_0 + A_1\pi_t + A_2\varphi_{t-1}u_t + A_3u_t + A_4(\varphi_t - \tilde{\varphi}_0)^2, \quad (39)$$

where

$$\begin{aligned} \tilde{\varphi}_0 &= \frac{v^2\kappa_1 + v^2\kappa_1\sigma_w + 1}{v\kappa_1 + v\kappa_1\sigma_w + 1}\varphi_0 - \frac{\kappa_1\rho_c(\phi_\pi + \kappa_1\xi_\pi)}{(\phi_c + \kappa_1\xi_c)(\kappa_1\rho_\pi - 1)} \frac{(v^2\kappa_1 + v^2\kappa_1\sigma_w - 1)}{(v\kappa_1 + v\kappa_1\sigma_w + 1)} \\ &\approx \varphi_0 \text{ when } \nu, \kappa_1 \approx 1, \sigma_w \approx 0 \text{ and } \phi_\pi, \xi_\pi \text{ or } \rho_c \approx 0 \end{aligned}$$

A Campbell-Shiller approximation for the return on the aggregate wealth portfolio gives

$$r_{c,t+1} = \kappa_0 + \kappa_1wc_{t+1} - wc_t + \Delta c_{t+1} \quad (40)$$

where

$$\kappa_1 = \frac{e^{E[wc_t]}}{e^{E[wc_t]} + 1} \quad (41)$$

$$\kappa_0 = \log(e^{E[wc_t]} + 1) - \frac{e^{E[wc_t]}}{e^{E[wc_t]} + 1} E[wc_t] \quad (42)$$

are linearization constants and $E[wc_t]$ is the unconditional mean of the log wealth-consumption ratio. Using equations (17), (16), (18), (15), (39) and (40), the Euler equation for $r_{c,t+1}$

$$0 = E_t[m_{t+1} + r_{c,t+1}] + \frac{1}{2}Var_t[m_{t+1} + r_{c,t+1}] \quad (43)$$

gives

$$\begin{aligned}
0 = r_0^c & + \left(\theta \kappa_1 A_1 \rho_\pi - \theta A_1 + \left(\theta - \frac{\theta}{\psi} \right) \rho_c \right) \pi_t \\
& + (\theta \kappa_1 A_1 \xi_\pi - \theta A_3) u_t \\
& + \left(\left(\theta - \frac{\theta}{\psi} \right) \xi_c - \theta A_2 \right) \varphi_{t-1} u_t \\
& + \left(\frac{1}{2} \left(\phi_c \left(\theta - \frac{\theta}{\psi} \right) + \theta \kappa_1 A_2 \right)^2 - \theta A_4 + v^2 \theta \kappa_1 A_4 + v^2 \theta \kappa_1 \sigma_w A_4 \right) \varphi_t^2. \quad (44)
\end{aligned}$$

where

$$\begin{aligned}
r_0^c & = \theta (\kappa_0 + \ln \delta - A_0 - A_4 \tilde{\varphi}_0^2 + (A_0 + \mu_\pi A_1) \kappa_1) \\
& + \theta (v \varphi_0^2 (v+2) + \tilde{\varphi}_0 \varphi_0 (\tilde{\varphi}_0 \varphi_0 - 2(v+1)) + 2\sigma_w^2) \kappa_1 A_4 \\
& + \left(\theta - \frac{\theta}{\psi} \right) \mu_c + \frac{1}{2} \left(\theta - \frac{\theta}{\psi} \right)^2 \sigma_c^2 \\
& + \theta \kappa_1 \left(\frac{\theta \kappa_1}{2} (\phi_\pi A_1 + A_3)^2 + \sigma_w A_4 ((1+v) \varphi_0 - \tilde{\varphi}_0)^2 \right)
\end{aligned}$$

In order for equation (44) to be satisfied, the coefficients in front of the time-varying state variables must vanish, yielding Equations (23)-(26). The remaining constant A_0 is determined by setting $r_0^c = 0$, which we solve numerically.

For the inflation portfolios, we conjecture the log price-dividend ratio is

$$pd_{i,t} = D_{0i} + D_{1i} \pi_t + D_{2i} \varphi_{t-1} u_t + D_{3i} u_t + D_{4i} (\varphi_t - \tilde{\varphi}_{0i})^2,$$

where $\tilde{\varphi}_{0i}$ is a constant picked to ensure that the Euler equation has no terms linear in φ_t . We solve for the coefficients D_{ji} in the same way as for the coefficients A_j but using the Euler equation

$$0 = E_t [m_{t+1}] + E_t [r_{i,t+1}] + \frac{1}{2} Var_t [m_{t+1}] + \frac{1}{2} Var_t [r_{i,t+1}] + Cov_t [m_{t+1}, r_{i,t+1}] \quad (45)$$

for the return

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

where

$$\kappa_{i,1} = \frac{e^{E[pd_{i,t}]}}{e^{E[pd_{i,t}]} + 1} \quad (46)$$

$$\kappa_{i,0} = \log(e^{E[pd_{i,t}]} + 1) - \frac{e^{E[pd_{i,t}]}}{e^{E[pd_{i,t}]} + 1} E[pd_{i,t}] \quad (47)$$

The solution for D_{1i} , D_{2i} and D_{3i} are given by equations (33), (34) and (35) in the paper, while D_{4i} satisfies the quadratic equation

$$\begin{aligned} 0 = & \left(\frac{1}{2}\phi_c \left(\frac{\theta}{\psi} - \theta + 1 \right) - \frac{1}{2}\kappa_1 A_2 (\theta - 1) \right)^2 \\ & - \frac{1}{2} \left(\phi_c \left(\frac{\theta}{\psi} - \theta + 1 \right) - \kappa_1 A_2 (\theta - 1) \right) (\phi_i + \kappa_{i,1} D_{i,2}) \\ & + \left(\frac{1}{2}\phi_i + \frac{1}{2}\kappa_{i,1} D_{i,2} \right)^2 - A_4 (\theta - 1) \\ & - D_{i,4} + v^2 \kappa_{i,1} D_{i,4} + v^2 \sigma_w^2 \kappa_{i,1}^2 D_{i,4}^2 + v^2 \kappa_1 A_4 (\theta - 1) \\ & + 4v^2 \kappa_1^2 \sigma_w^2 A_4^2 (\theta - 1)^2 + 2v^2 \kappa_1 \sigma_w^2 A_4 (\theta - 1) \kappa_{i,1} D_{i,4} \end{aligned}$$

for which we take the positive root as our solution. Finally, we solve for D_{0i} numerically by setting the constant term in equation (45) to zero.

8.2 Estimation of Inflation and Consumption Growth Parameters

We choose parameters μ_π , ρ_π , ϕ_π , ξ_π by fitting an $ARMA(1,1)$ process, which corresponds exactly to the specification in Equation (16). Assuming $\{u_t\}$ is a white noise process, we use inflation data to estimate the constants m_0 , m_1 and m_2 in equation

$$\pi_{t+1} = m_0 + m_1\pi_t + u_{t+1} + m_2u_t \quad (48)$$

by maximum likelihood over the same time period as in our empirical section. We get

$$\hat{m}_0 = 0.0028, \quad (49)$$

$$\hat{m}_1 = 0.90, \quad (50)$$

$$\hat{m}_2 = 0.57. \quad (51)$$

Picking $u_0 = 0$ and plugging the estimates (49)-(51) into equation (48) gives an estimated path, $\{\hat{u}_t\}_{t=1}^T$, for inflation shocks. The resulting sequence $\{\hat{u}_t\}$ is identical to the one we use in Section 2 to construct inflation betas, as they are both generated by the same estimation procedure. The standard deviation of the estimated inflation shocks is

$$std(\hat{u}_t) = 0.0028 \quad (52)$$

and its mean is essentially zero. We map equations (49)-(52) to the model's inflation parameters as follows:

$$\mu_\pi = \hat{m}_0 = 0.0028,$$

$$\rho_\pi = \hat{m}_1 = 0.897,$$

$$\phi_\pi = std(\hat{u}_t) = 0.0028,$$

$$\xi_\pi = \hat{m}_2 \times std(\hat{u}_t) = -0.0016,$$

where we identify the model's u_t with $\hat{u}_t/std(\hat{u}_t)$ so that it has unit standard deviation.

To calibrate the consumption parameters, we use data on real consumption growth, the series for the nominal-real covariance NRC_t^C constructed in equation (10) and \hat{u}_t to run the following full-sample OLS regression:

$$\Delta c_{t+1} = g_0 + g_1 \pi_t + g_2 \times NRC_{t+1}^C \hat{u}_{t+1} + g_3 \times NRC_t^C \hat{u}_t + g_4 \hat{u}_{t+1} + g_5 \hat{u}_t + \varepsilon_{t+1}, \quad (53)$$

where g_i , $i = 0, 1, \dots, 5$, are the regression coefficients we estimate and ε_{t+1} is a mean-zero regression residual. To translate regression (53) into the model's consumption process, given by equation (17), we equate the data-constructed NRC_t^C with its model-implied analog $\theta_t^{(12)}$, so that

$$NRC_t^C = h_{12} + \frac{\xi_c}{\phi_\pi} \varphi_t. \quad (54)$$

With the aid of equation (54), comparing equation (17) to equation (53) gives a mapping from \hat{g}_i , $std(\hat{\varepsilon}_t)$ and $std(\hat{u}_t)$ to the model's consumption parameters:

$$\mu_c = \hat{g}_0 + \hat{g}_1 \mu_\pi = 0.0021,$$

$$\rho_c = \hat{g}_1 = -0.126,$$

$$\sigma_c = std(\hat{\varepsilon}_t) = 0.0032,$$

$$\xi_c = \frac{\phi_\pi}{std(\hat{u}_t) \hat{g}_2} = 0.0797.$$

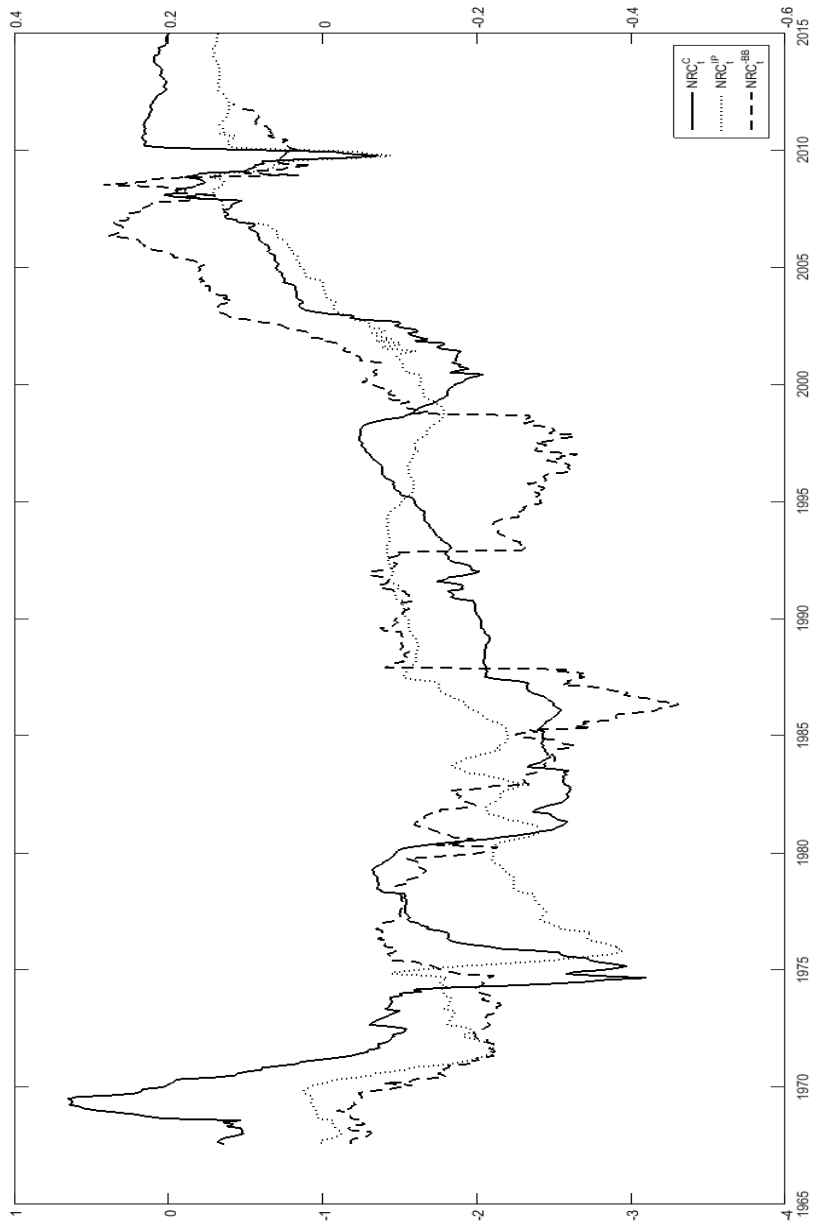


Figure 1: Alternative measures of the nominal-real covariance

This figure presents the time-series of three alternative measures of the nominal-real covariance. The main measure used in this study is the beta of future twelve-month consumption growth on lagged inflation, denoted NRC_t^C . The first alternative measure uses industrial production growth instead of consumption growth and is denoted NRC_t^{IP} . These measures are plotted on the left axis and the latter is divided by four to preserve scaling. The second alternative is the stock market beta of the 10-year constant maturity treasury bond (right axis). We multiply this measure by minus one to be comparable to the others and denote it NRC_t^{-BB} .

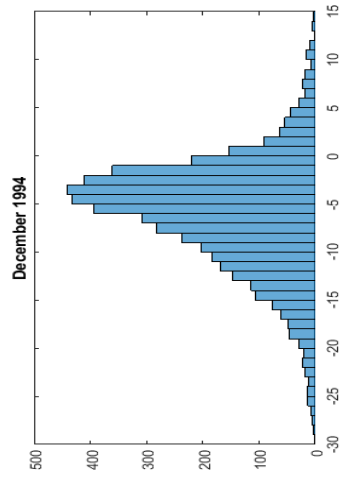
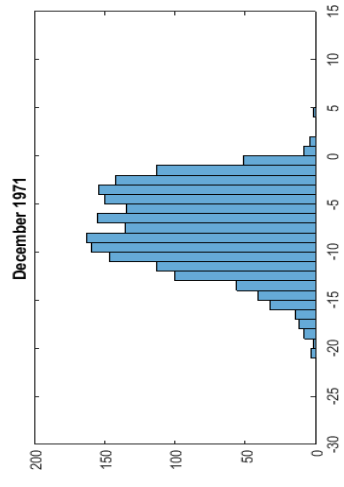
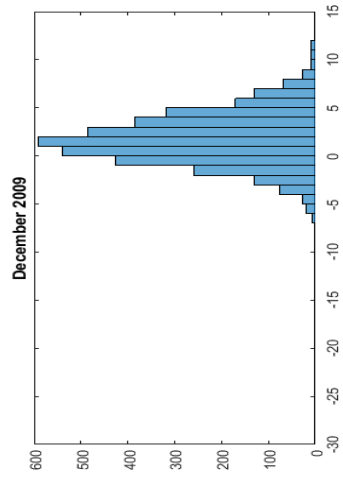
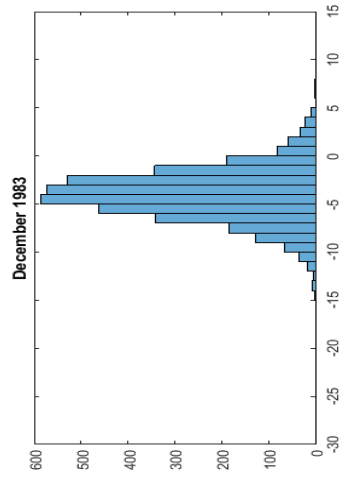


Figure 2: Histogram of inflation betas for different time periods

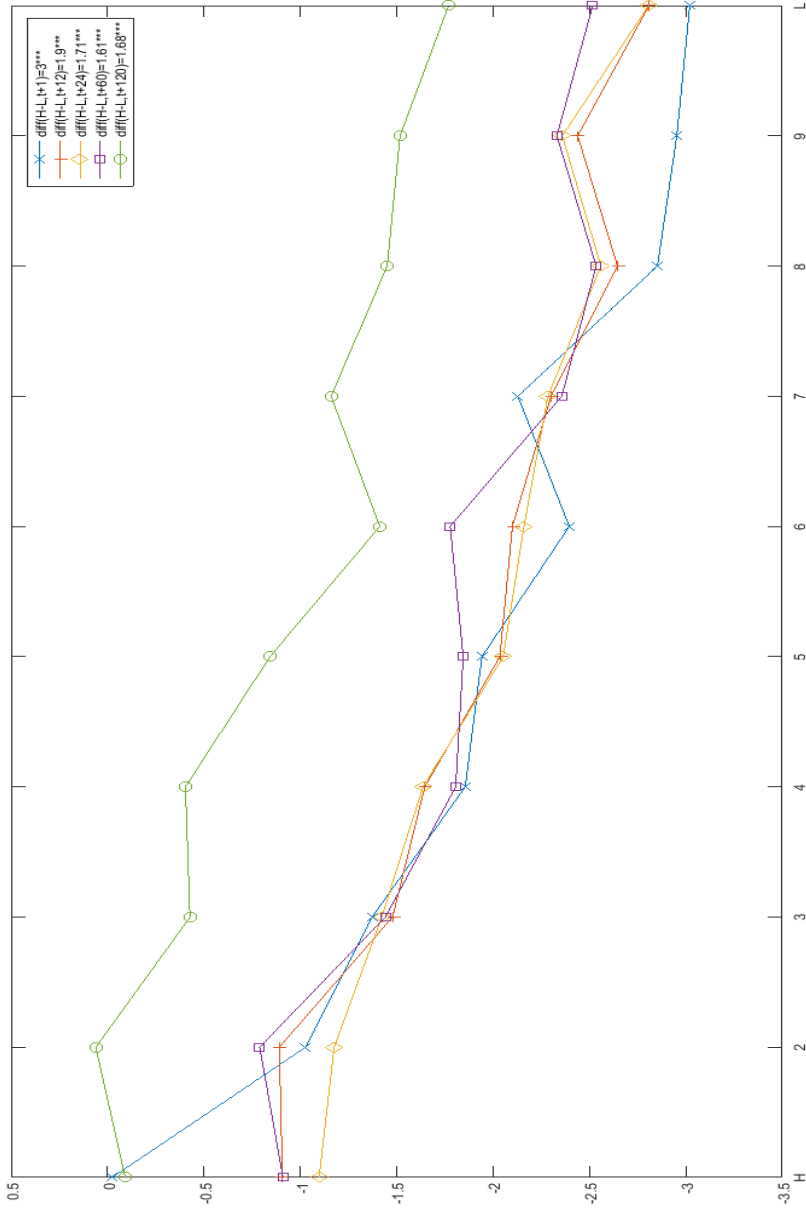


Figure 3: Post-ranking inflation beta: 1, 12, 24, 60, and 120 months after sorting

This figure plots the monthly post-ranking inflation beta one month, one year, two years, five years, and ten years after sorting. To calculate these inflation betas, we fix the portfolio composition at the sorting date t and calculate value-weighted returns up to ten years after, i.e., $t + 1, \dots, t + 12, \dots, t + 120$. When a stock leaves the sample, we reallocate its market value across all remaining stocks. 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. The legend includes the estimated inflation beta for the high minus low spreading portfolio. *, **, ***, **** indicate statistical significance at the 10, 5, and 1 %-level using Newey-West standard errors with lag length equal to one.

Table 1: Descriptive statistics

This table reports descriptive statistics in annualized percentages for CPI inflation (Π_t), $ARMA(1, 1)$ -innovations in inflation ($u_{\Pi,t}$), consumption growth (C_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, which adds up to 630 months.

	Π_t	$u_{\Pi,t}$	C_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)	0.62	0.15	-0.17	0.07	0.97
Correlations					
Π_t	1	0.79	-0.18	-0.13	0.50
$u_{\Pi,t}$		1	-0.17	-0.12	0.15
C_t			1	0.17	0.01
$R_{m,t}$				1	-0.09
$R_{f,t}$					1

Table 2: Unconditional and conditional predictive regressions of consumption growth on inflation

In this table we present results from predictive regressions of consumption growth over the full sample period from July 1967 to December 2014. Panel A presents results from an unconditional regression of consumption growth over horizon $K = 1, 3, 6, 12$, $C_{t+1:t+K}$, on lagged inflation Π_t : $C_{t+1:t+K} = d_0^{u,K} + d_1^{u,K} \Pi_t + e_{t+1:t+K}$. Panel B presents results from a two-stage conditional regression. In the first stage, we regress consumption growth on lagged inflation over a backward looking rolling window: $C_{s+1:s+K} = a_{t-1}^K + b_{t-1}^K \Pi_s + e_{s+1:s+K}$, $s = 1, \dots, t - K$. Combining the estimated coefficients with inflation in month t , we then predict future consumption growth in the full time series: $C_{t+1:t+K} = d_0^{c,K} + d_1^{c,K} (\widehat{a_{t-1}^K} + \widehat{b_{t-1}^K} \Pi_t) + e_{t+1:t+K}$. 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 in the backward looking rolling window regression. Denoting this constant \hat{a}_{t-1}^* , we have that $R2 - OOS = 1 - \frac{Var(C_{t+1:t+K} - (a_{t-1}^K + b_{t-1}^K \Pi_t))}{Var(C_{t+1:t+K} - \hat{a}_{t-1}^{*,K})}$. In both panels, standard errors are Newey-West with K lags and R^2 's are reported in percentage points.

Panel A: Unconditional					Panel B: Conditional				
Horizon K	1	3	6	12	1	3	6	12	
$d_0^{u,K}$	0.0018	0.0051	0.0099	0.0200	0.0002	0.0001	0.0001	0.0005	
(t)	(9.0962)	(8.7276)	(7.7807)	(7.5202)	(0.7596)	(0.1513)	(0.074)	(0.1022)	
$d_1^{u,K}$	-0.0651	-0.1516	-0.2001	-0.4960	0.7675	0.8432	0.8403	0.8114	
(t)	(-1.371)	(-1.0833)	(-0.7694)	(-1.0649)	(4.7405)	(5.5961)	(5.3036)	(4.5561)	
R^2	0.26	0.76	0.44	1.12	3.29	10.29	13.28	14.92	
					$R^2 - OOS$	1.80	6.70	10.41	

Table 3: Overview of inflation beta sorted portfolios

This table presents the set of ten inflation beta-sorted portfolios and our resulting estimates of the unconditional inflation risk premium, denoted HLIP, MCIP, and CSIP. The portfolios are constructed by sorting all CRSP stocks into 30 portfolios at the intersection of ten inflation beta-deciles and three size-groups. Then, the ten portfolios of interest are equal-weighted averages of the three portfolios in each inflation beta decile. HLIP is the High-minus-Low decile spreading portfolio from this sort and represents our first estimate of the inflation risk premium portfolio. MCIP is a maximum correlation mimicking-portfolio that is constructed from a multiple regression of inflation innovations on the ten inflation beta sorted portfolios. CSIP is the inflation risk premium estimated from a cross-sectional regression of individual stock returns on lagged inflation betas (controlling for market capitalization). Panel A presents ex post inflation betas, $\beta_{\Pi,post,t}$, that are estimated with a simple regression of portfolio returns on inflation innovations over the full sample. t -statistics are in parenthesis and use Newey-West standard errors with lag length one. Note that MCIP and CSIP are scaled to have the same post-ranking inflation beta as HLIP. Panel A also presents summary statistics for each portfolio's rolling inflation beta that is estimated as explained in Section I.B, $\beta_{\Pi,post,t}$. Panel B presents annualized performance statistics for the portfolio returns. The sample period is July 1967 to December 2014.

	H	2	3	4	5	6	7	8	9	L	HLIP	MCIP	CSIP
Panel A: Ex post inflation exposures													
$\beta_{\Pi,post}$	-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)
Average $\beta_{\Pi,post,t}$	-2.59	-2.83	-3.05	-3.58	-3.83	-4.14	-3.98	-4.71	-4.69	-4.87	2.28	2.67	2.71
St. Dev. $\beta_{\Pi,post,t}$	3.61	2.78	2.78	2.53	2.83	2.73	2.79	2.98	2.95	3.25	1.36	1.02	1.40
Panel B: The unconditional inflation risk premium													
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
(t)	(1.59)	(2.47)	(2.88)	(2.85)	(3.2)	(3.07)	(2.94)	(3.02)	(2.99)	(3)	(-2.08)	(-3.06)	(-2.22)
St. Dev.	22.72	20.36	18.79	18.38	18.34	18.61	18.98	19.38	20.20	21.78	14.04	9.88	13.06
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

Table 4: Predicting the inflation risk premium with the nominal-real covariance

This table presents simple predictive regressions for ten inflation beta-sorted portfolios as well as our three resulting estimates of the inflation risk premium (HLIP, MCIP, and CSIP) on the nominal-real covariance. We report for each regression the estimated coefficients, with corresponding t -statistics in parentheses based on Newey-West standard errors with H lags, and the adjusted R^2 (in percentage points). Panel A uses the full sample from July 1967 to December 2014 and our main measure of the nominal-real covariance that tracks the time-varying relation between inflation and consumption growth: NRC_t^C . The first five rows use overlapping three- and twelve-month returns ($H = 1$), whereas the second and third block of five rows use overlapping three- and twelve-month returns ($H = 3, 12$). Panel B presents results for $H = 12$ using two alternative measures of the nominal-real covariance: (i) the time-varying relation between inflation and industrial production growth, NRC_t^{IP} , and (ii) the negative of the time-varying stock market beta of the long-term bond, NRC_t^{-BB} . For this latter case, the sample ends in December 2011. Panel C presents results for $H = 12$ in the first- and second-half of the sample, split around February 1991.

$R_{t+1:t+H}$	High	2	3	4	5	6	7	8	9	Low	HLIP	MCIP	CSIP	
Predicting inflation portfolio returns with NRC_t^C from July 1967 to December 2014														
H=1	L_0	5.26	7.29	7.59	8.53	8.28	8.10	8.50	8.76	9.49	-4.23	-4.39	-4.21	
	(t)	(1.49)	(2.34)	(2.73)	(2.69)	(2.99)	(2.72)	(2.8)	(2.76)	(2.76)	(-1.95)	(-2.92)	(-2.08)	
	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
	(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)
	R^2	-0.11	-0.14	-0.17	-0.18	-0.14	-0.13	-0.07	-0.06	-0.07	0.67	0.46	0.59	
H=3	L_0	5.72	7.46	7.99	7.78	8.84	8.59	8.48	8.85	9.17	-4.27	-4.19	-3.93	
	(t)	(1.77)	(2.71)	(3.17)	(3.09)	(3.36)	(3.23)	(3.1)	(3.18)	(3.17)	(-2.03)	(-2.89)	(-2.2)	
	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
	(t)	(0.63)	(0.41)	(0.1)	(-0.15)	(-0.27)	(-0.49)	(-0.6)	(-0.89)	(-0.96)	(-0.89)	(2.45)	(2.37)	(2.32)
	R^2	0.01	-0.10	-0.17	-0.16	-0.14	0.01	0.23	0.28	0.20	2.48	2.18	2.16	
H=12	L_0	5.88	7.66	8.09	7.99	9.08	8.82	8.68	9.03	9.35	-4.37	-4.25	-3.94	
	(t)	(1.91)	(2.89)	(3.46)	(3.36)	(3.74)	(3.62)	(3.47)	(3.54)	(3.6)	(-1.68)	(-2.62)	(-2.31)	
	L_{NRC}	2.01	0.99	-0.20	-0.77	-1.14	-1.75	-2.31	-3.39	-3.63	-3.37	3.73	4.51	
	(t)	(0.59)	(0.34)	(-0.08)	(-0.27)	(-0.39)	(-0.6)	(-0.77)	(-1.15)	(-1.23)	(-1.07)	(2.22)	(2.58)	(2.78)
	R^2	0.41	0.01	-0.17	-0.04	0.11	0.94	2.13	2.32	1.55	6.65	7.64	9.68	

Table 4 continued

$R_{t+1:t+H}$	High	2	3	4	5	6	7	8	9	Low	HLIP	MCIP	CSIP
Panel B: Predicting inflation portfolio returns with alternative measures the nominal-real covariance													
Inflation and industrial production: NRC_t^{IP}													
H=12	L_0	5.88	7.66	8.09	7.99	8.82	8.68	9.03	9.35	10.24	-4.37	-4.25	-3.94
	(t)	(1.91)	(2.89)	(3.46)	(3.36)	(3.73)	(3.62)	(3.45)	(3.58)	(3.52)	(-1.67)	(-2.59)	(-2.32)
	L_{NRC}	1.94	1.18	-0.05	-0.23	-0.48	-1.02	-1.24	-2.30	-3.42	5.36	3.24	4.54
	(t)	(0.65)	(0.47)	(-0.02)	(-0.09)	(-0.19)	(-0.4)	(-0.47)	(-1.04)	(-1.17)	(2.39)	(2.21)	(2.74)
	R^2	0.36	0.09	-0.18	-0.17	-0.13	0.05	0.14	0.88	1.31	6.60	5.72	9.79
Stock market beta of long-term bond: NRC_t^{-BB}													
H=12	L_0	5.44	6.99	7.60	7.18	8.46	8.12	7.94	8.31	8.70	9.50	-3.85	-3.61
	(t)	(1.68)	(2.55)	(3.1)	(2.91)	(3.33)	(3.2)	(3.02)	(3.08)	(3.14)	(3.06)	(-2.19)	(-1.95)
	L_{NRC}	2.11	1.92	0.03	-0.55	-0.21	-1.00	-1.09	-1.77	-1.54	-1.49	2.96	3.72
	(t)	(0.72)	(0.79)	(0.01)	(-0.22)	(-0.09)	(-0.41)	(-0.43)	(-0.68)	(-0.55)	(-0.49)	(2.31)	(2.48)
	R^2	0.43	0.51	-0.19	-0.12	-0.18	0.02	0.05	0.42	0.24	2.73	4.48	6.23
Panel C: Subsamples using NRC_t^C													
First sample-half (July 1967 to February 1991)													
H=12	L_0	3.41	3.87	5.99	5.65	7.06	6.18	6.57	6.87	7.93	7.51	-4.09	-5.49
	(t)	(1.01)	(1.3)	(2.11)	(1.76)	(2.06)	(1.76)	(1.78)	(1.79)	(1.99)	(1.71)	(-1.39)	(-2.69)
	L_{NRC}	-4.08	-5.07	-5.38	-6.20	-6.01	-6.46	-7.20	-7.71	-7.43	-5.76	1.67	2.59
	(t)	(-1.06)	(-1.6)	(-1.84)	(-1.85)	(-1.63)	(-1.78)	(-1.88)	(-2.1)	(-1.96)	(-1.43)	(0.82)	(2.02)
	R^2	3.36	6.61	7.95	8.45	6.93	7.30	8.20	8.76	7.52	3.77	7.81	4.20
Second sample-half (March 1991 to December 2014)													
H=12	L_0	8.44	11.61	10.27	10.43	11.18	11.56	10.87	11.29	10.82	13.09	-4.66	-2.33
	(t)	(1.73)	(2.87)	(3)	(3.34)	(3.59)	(3.87)	(3.64)	(3.81)	(3.54)	(3.71)	(-1.19)	(-0.87)
	L_{NRC}	7.22	4.51	3.76	3.11	2.32	0.71	0.80	-1.29	-1.43	-4.19	11.41	6.26
	(t)	(1.47)	(1.12)	(1.14)	(1)	(0.77)	(0.24)	(0.27)	(-0.44)	(-0.47)	(-1.14)	(2.49)	(2.01)
	R^2	5.23	2.80	2.73	2.23	1.06	-0.22	-0.17	0.14	0.23	3.27	15.97	14.54

Table 5: Predicting inflation betas with the nominal-real covariance

This table analyzes how inflation betas of our inflation portfolios vary with the nominal-real covariance. To this end, we regress the rolling inflation beta of the ten decile portfolios (estimated by applying Equation (1) to the post-ranking returns of the portfolios) on the nominal-real covariance, NRC_t^C , which is standardized to have mean equal to zero and variance equal to one to accommodate interpretation: $\beta_{\Pi,p,t} = \beta_{p,0} + \beta_{p,NRC}NRC_t^C + \varepsilon_{p,t}$. We report for each regression the estimated coefficients, with corresponding t -statistics based on Newey-West standard errors with 60 lags, and the R^2 .

	High	2	3	4	5	6	7	8	9	Low
$\beta_{p,0}$	-2.59	-2.83	-3.05	-3.58	-3.83	-4.14	-3.98	-4.71	-4.69	-4.87
(t)	(-2.96)	(-4.27)	(-4.72)	(-6.51)	(-5.96)	(-6.98)	(-6.43)	(-7.37)	(-7.42)	(-7.12)
$\beta_{p,NRC}$	1.58	1.18	1.28	1.34	1.52	1.58	1.56	1.80	1.74	1.92
(t)	(1.68)	(1.72)	(2.03)	(2.32)	(2.31)	(2.67)	(2.44)	(2.98)	(3.06)	(3.18)
R^2	0.19	0.18	0.21	0.28	0.29	0.33	0.31	0.36	0.35	0.35

Table 6: 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	14.5
<i>Inflation</i>		
Mean of inflation rate	μ_π	0.0028
Inflation AR(1) coefficient	ρ_π	0.897
Inflation MA(1) coefficients	ϕ_π	0.0028
	ξ_π	-0.0016
<i>Consumption</i>		
Mean of real consumption growth	μ_c	0.0021
Exposure of expected consumption to inflation	ρ_c	-0.126
Volatility of consumption-specific shocks	σ_c	0.0032
Consumption exposure to inflation shocks	ξ_c	0.0797
<i>Nominal-real covariance</i>		
Mean of nominal-real covariance	φ_0	0.0017
Persistence of nominal-real covariance	ν	0.994
Volatility of nominal-real covariance	σ_w	4.86×10^{-4}
<i>High inflation beta portfolio</i>		
Mean of dividend growth	μ_H	7.6×10^{-4}
Exposure of expected dividends to inflation	ρ_H	-0.812
Exposure to consumption-specific shocks	σ_H	-0.074
Dividend exposure to inflation shocks	ϕ_H	4.048
	ξ_H	-2.26
<i>Low inflation beta portfolio</i>		
Mean of dividend growth	μ_L	0.0073
Exposure of expected dividends to inflation	ρ_L	-1.175
Exposure to consumption-specific shocks	σ_L	-0.051
Dividend exposure to inflation shocks	ϕ_L	4.85
	ξ_L	-3.34

Table 7: 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 high-minus-low portfolio on the nominal-real covariance.

	Data	Model
$\mathbb{E}[\pi_t]$	3.93	3.93
$\sigma(\pi_t)$	1.11	2.52
$\mathbb{E}[\Delta c_t]$	1.97	2.50
$\sigma(\Delta c_t)$	1.13	2.01
$corr(\pi_t, \Delta c_t)$	-17.8%	-3.6%
$\mathbb{E}[R_H]$	5.26	6.16
$\sigma(R_H)$	22.7	25.9
$\mathbb{E}[\beta_{u,H}]$	-2.59	-2.22
$\sigma(\beta_{u,H})$	3.61	3.17
$\mathbb{E}[R_L]$	9.49	9.00
$\sigma(R_L)$	21.78	18.29
$\mathbb{E}[\beta_{u,L}]$	-4.87	-4.25
$\sigma(\beta_{u,L})$	3.25	2.65
$\sigma(\beta_{u,HLIP})$	1.36	0.52
$\sigma(R_{HLIP})$	14.04	8.71
$L_{NRC,HLIP}$	4.47	4.33

Table 8: Industry composition of inflation beta sorted portfolios

This table presents the industry composition of our sort. For this exercise, we use the classification into 48 industries from Kenneth French’s web site. In each sample month, we find for each industry the stocks that have inflation beta below or above the median inflation beta in the full cross-section and translate this to the fraction of an industry’s total market capitalization that has below or above median inflation beta. In the first three columns we report results for the top 10 inflation hedgers, which are those industries with on average the largest fraction of market cap in the above median inflation beta portfolio. For these industries we report the average allocation to the above median portfolio (“% of market cap ($\beta_{\Pi,i,t} > \text{median}$)”) as well as the fraction of total sample months in which the allocation to the above median portfolio is larger than 50% (“% of months”). The next three columns present analogous evidence for the top 10 worst inflation hedgers, which are those industries with on average the largest fraction of market cap in the below median inflation beta portfolio. The sample period is July 1967 to December 2014.

Top 10 best inflation hedgers				Top 10 worst inflation hedgers		
		% of market cap ($\beta_{\Pi,i,t} > \text{median}$)	% of months		% of market cap ($\beta_{\Pi,i,t} \leq \text{median}$)	% of months
1	Oil	0.77	0.80	Other	0.82	0.94
2	Gold	0.73	0.76	Books	0.77	0.95
3	Utilities	0.66	0.68	Aerospace	0.76	0.89
4	Mines	0.64	0.70	Banks	0.74	0.84
5	Agriculture	0.62	0.71	Insurance	0.74	0.88
6	Paper	0.61	0.66	Meals	0.73	0.82
7	Computers	0.60	0.67	Lab Equipment	0.72	0.82
8	Steel	0.58	0.59	Clothes	0.71	0.79
9	Smoke	0.58	0.66	Textiles	0.71	0.82
10	Food	0.57	0.59	Personal Services	0.70	0.85

Table 9: The inflation risk premium within- and across-industry

This table asks whether our main results on the inflation risk premium are driven by within- or across-industry variation in inflation beta, or both. For the within-industry sort, we construct five market value-weighted stock portfolios within each industry by splitting at the quintiles of ranked inflation betas of the stocks within each of 48 industries. For the across-industry sort, we calculate the inflation beta of the 48 industries as the value-weighted average inflation beta of all stocks in that industry and sort the industries into quintile portfolios (that typically contain nine or ten industries each). This leaves us with a five-by-five within- and across-industry sort. We collapse this sort into five within-industry portfolios and five across-industry portfolios. The within-industry portfolios are computed by averaging over five across-industry portfolios for each within-industry quintile. The across-industry portfolios are calculated as the equal-weighted average of the nine or ten industries that belong to the relevant quintile of the across-industry sort. The aggregate within-industry effect is presented in the sixth column and is the difference between the high and low within-industry portfolio return. The across-industry effect is in the twelfth column and is the difference between the high and low across-industry portfolio return. For each portfolio, we present post-ranking inflation beta and average returns (as in Table 3) as well as the predictive regression of twelve-month compounded future returns on the lagged nominal-real covariance, NRC_t^C (as in Table 4).

	Within-industry effect					Across-industry effect						
	High β_{Π} stocks	2	3	4	Low β_{Π} stocks	High-low β_{Π} stocks	High β_{Π} industries	2	3	4	Low β_{Π} industries	High-Low
$\beta_{\Pi,post}(t)$	-1.23 (-0.69)	-2.05 (-1.22)	-2.25 (-1.44)	-2.82 (-1.81)	-2.85 (-1.83)	1.62 (3.42)	-1.70 (-1.25)	-2.11 (-1.57)	-2.72 (-1.68)	-3.24 (-2.28)	-2.78 (-1.67)	1.07 (1.76)
Avg. Ret. (t)	6.20 (2.19)	6.72 (2.54)	7.48 (2.86)	7.50 (2.76)	7.91 (2.8)	-1.71 (-1.38)	5.82 (2.44)	6.73 (2.7)	7.52 (2.93)	5.46 (2.04)	8.71 (3.08)	-2.89 (-1.78)
$L_0(t)$	6.56 (2.58)	7.16 (2.93)	7.78 (3.54)	7.89 (3.33)	8.28 (3.42)	-1.73 (-1.55)	5.93 (2.98)	6.99 (3.29)	7.94 (3.57)	5.78 (2.41)	9.29 (3.62)	-3.36 (-1.95)
$L_{NRC}(t)$	0.27 (0.09)	0.15 (0.05)	-1.09 (-0.41)	-2.10 (-0.75)	-2.97 (-1.07)	3.24 (3.29)	0.69 (0.3)	-0.67 (-0.28)	-1.23 (-0.48)	-1.51 (-0.54)	-3.94 (-1.3)	4.63 (2.7)
R^2	-0.17	-0.17	0.14	0.86	1.74	10.81	-0.02	-0.05	0.22	0.38	2.95	10.03

Table 10: Predicting the inflation risk premium controlling for benchmark factor exposure

This table asks whether our conclusions on the time-varying inflation risk premium extend when we control ex post for exposure of the inflation sorted portfolios to the benchmark asset-pricing factors. For this exercise, we regress returns of HLIP, MCIP, and CSIP on the nominal-real covariance (NRC_t^C) as well as on contemporaneous exposure to the CAPM (MKT), FF3M (MKT, SMB, HML) and FFCM (MKT, SMB, HML, and MOM). For the tests using overlapping three- and twelve-month returns, we compound returns on both the left-hand side inflation portfolios and the right hand side factors. We presents for each regression the estimated coefficients and the adjusted R^2 (in percentage points). To conserve space, we present t -statistics in parentheses (based on Newey-West standard errors with H lags) only for the intercept, L_0 , and coefficient on NRC_t^C , L_{NRC} .

$R_{t+1:t+H}$		CAPM			FF3M			FFCM		
		HLIP	MCIP	CSIP	HLIP	MCIP	CSIP	HLIP	MCIP	CSIP
H=1	L_0	-4.21	-3.18	-3.79	-3.47	-2.31	-2.63	-2.44	-2.24	-1.28
	(t)	(-1.84)	(-2.07)	(-1.81)	(-1.61)	(-1.55)	(-1.32)	(-1.07)	(-1.46)	(-0.6)
	L_{NRC}	4.47	2.68	3.93	4.18	2.34	3.47	3.84	2.31	3.03
	(t)	(2.11)	(1.92)	(1.98)	(1.99)	(1.72)	(1.75)	(1.85)	(1.72)	(1.53)
	β_{MKT}	0.00	-0.20	-0.07	-0.01	-0.21	-0.09	-0.03	-0.21	-0.11
	β_{SMB}				-0.07	-0.10	-0.10	-0.07	-0.10	-0.10
	β_{HML}				-0.12	-0.14	-0.19	-0.15	-0.14	-0.23
	β_{MOM}							-0.10	-0.01	-0.13
	R^2	0.50	10.78	1.12	0.98	12.87	3.02	1.81	12.72	4.85
H=3	L_0	-4.32	-3.20	-3.49	-3.69	-2.15	-2.51	-2.44	-1.81	-1.77
	(t)	(-1.92)	(-2.11)	(-1.82)	(-1.91)	(-1.61)	(-1.43)	(-1.06)	(-1.13)	(-0.93)
	L_{NRC}	5.16	3.25	4.28	4.89	2.79	3.86	4.48	2.68	3.61
	(t)	(2.43)	(2.42)	(2.31)	(2.37)	(2.13)	(2.12)	(2.21)	(2.09)	(1.97)
	β_{MKT}	0.01	-0.16	-0.07	0.09	-0.12	-0.01	0.07	-0.13	-0.02
	β_{SMB}				-0.34	-0.25	-0.32	-0.34	-0.25	-0.31
	β_{HML}				-0.08	-0.16	-0.14	-0.12	-0.17	-0.17
	β_{MOM}							-0.11	-0.03	-0.06
	R^2	2.32	8.18	2.70	6.95	15.24	8.25	7.80	15.25	8.54
H=12	L_0	-4.98	-3.56	-4.30	-3.85	-2.31	-3.09	0.48	-0.98	-1.43
	(t)	(-1.62)	(-1.86)	(-2.16)	(-1.74)	(-1.83)	(-2.04)	(0.21)	(-0.68)	(-0.91)
	L_{NRC}	5.46	3.65	4.56	4.96	3.08	4.01	3.45	2.62	3.43
	(t)	(2.21)	(2.52)	(2.76)	(2.07)	(2.25)	(2.48)	(1.64)	(2.09)	(2.25)
	β_{MKT}	0.10	-0.11	0.06	0.16	-0.09	0.08	0.09	-0.11	0.05
	β_{SMB}				-0.52	-0.31	-0.35	-0.54	-0.32	-0.36
	β_{HML}				-0.07	-0.12	-0.11	-0.20	-0.16	-0.16
	β_{MOM}							-0.35	-0.11	-0.14
	R^2	7.19	9.57	10.01	15.56	17.66	18.56	22.05	19.00	20.42

Table 11: Controlling for benchmark factors when estimating inflation beta

This table asks whether our results extend when we control ex ante, i.e., when estimating inflation beta, for stock's exposures to the benchmark asset-pricing factors of the CAPM (MKT), FF3M (MKT, SMB, HML) and FFCM (MKT, SMB, HML, and MOM). For each alternative sorting procedure, we calculate our three estimates of the inflation risk premium (HLIP, MCIP, and CSIP) and regress overlapping one-, three-, and twelve-month returns on the nominal-real covariance (NRC_t^C). We report for each regression the estimated coefficients, with corresponding t -statistics in parentheses based on Newey-West standard errors with H lags, and the adjusted R^2 (in percentage points).

$R_{t+1:t+H}$		CAPM			FF3M			FFCM		
		HLIP	MCIP	CSIP	HLIP	MCIP	CSIP	HLIP	MCIP	CSI
H=1	L_0	-4.23	-5.79	-4.41	-2.47	-3.69	-4.20	-2.23	-3.82	-3.92
	(t)	(-2.11)	(-3.97)	(-2.5)	(-1.47)	(-3.3)	(-2.58)	(-1.38)	(-3.23)	(-2.41)
	L_{NRC}	4.93	3.19	3.58	4.16	1.83	3.55	4.05	2.53	3.26
	(t)	(2.37)	(2)	(1.87)	(2.17)	(1.52)	(1.94)	(2.22)	(1.94)	(1.8)
	R^2	1.01	0.8	0.67	1.02	0.34	0.79	1.04	0.73	0.64
H=3	L_0	-4.15	-5.65	-4.12	-2.36	-3.65	-3.93	-2.15	-3.80	-3.65
	(t)	(-2.14)	(-4)	(-2.62)	(-1.49)	(-3.71)	(-2.68)	(-1.43)	(-3.5)	(-2.51)
	L_{NRC}	5.64	3.83	3.95	4.80	2.29	3.93	4.69	3.03	3.63
	(t)	(2.78)	(2.45)	(2.23)	(2.58)	(2.01)	(2.27)	(2.62)	(2.41)	(2.14)
	R^2	3.56	3.15	2.38	3.74	2.02	2.77	3.89	3.08	2.36
H=12	L_0	-4.47	-5.88	-4.17	-2.68	-3.94	-3.97	-2.50	-4.14	-3.71
	(t)	(-1.89)	(-3.72)	(-2.59)	(-1.54)	(-3.89)	(-2.62)	(-1.5)	(-3.51)	(-2.47)
	L_{NRC}	6.09	4.38	3.77	5.48	2.68	3.75	5.40	3.45	3.45
	(t)	(2.59)	(2.66)	(2.13)	(2.61)	(2.38)	(2.15)	(2.63)	(2.67)	(2.01)
	R^2	9.98	10.62	8	12.99	8.67	8.8	13.57	11.16	7.67

Table 12: Alternative measures of inflation risk

This table asks whether our results on the inflation risk premium extend when we estimate stock's exposures to alternative measures of inflation risk. First, we use raw inflation (Π_t), rather than its ARMA(1,1) innovations. Second, we use the difference between inflation and the short-rate as a measure of inflation innovations ($\Pi_t - R_{f,t}$). Third, we perform a truly out-of-sample exercise using real-time vintage CPI inflation (available from <http://research.stlouisfed.org/tips/alfred/>). For this exercise, we skip a month after portfolio formation, thus taking into account the reporting delay in inflation data. For each alternative measure of inflation risk, we calculate our three estimates of the inflation risk premium (HLIP, MCIP, and CSIP) and regress overlapping one-, three-, and twelve-month returns on the nominal-real covariance (NRC_t^C). We report for each regression the estimated coefficients, with corresponding t -statistics in parentheses based on Newey-West standard errors with H lags, and the adjusted R^2 (in percentage points). The sample period is from July 1967 to December 2014, except for the truly OOS exercise, which sample ends in December 2012.

$R_{t+1:t+H}$	Sorting on:	Inflation			Inflation-short rate			Truly OOS inflation		
		HLIP	MCIP	CSIP	HLIP	MCIP	CSIP	HLIP	MCIP	CSIP
H=1	L_0	-3.79	-4.37	-11.89	-5.13	-6.26	-7.16	-1.65	-2.81	-2.79
	(t)	(-1.66)	(-3.29)	(-1.78)	(-2.25)	(-3.12)	(-2.75)	(-0.72)	(-2.39)	(-1.27)
	L_{NRC}	6.03	1.84	14.81	6.88	4.51	7.15	6.88	1.14	7.55
	(t)	(2.59)	(1.35)	(2.24)	(3.01)	(2.2)	(2.77)	(2.9)	(0.91)	(3.5)
	R^2	1.19	0.22	0.82	1.49	0.78	1.27	1.61	0	2.27
H=3	L_0	-3.89	-4.37	-9.42	-5.30	-6.15	-6.79	-1.94	-2.68	-2.57
	(t)	(-1.86)	(-3.51)	(-1.6)	(-2.59)	(-3.3)	(-2.99)	(-0.91)	(-2.59)	(-1.29)
	L_{NRC}	6.69	2.18	15.17	7.69	5.30	7.61	7.35	1.44	7.62
	(t)	(2.99)	(1.65)	(2.47)	(3.53)	(2.74)	(3.29)	(3.1)	(1.33)	(3.59)
	R^2	3.86	1.13	2.5	5.34	3.23	4.28	4.68	0.64	5.86
H=12	L_0	-3.89	-4.58	-9.43	-5.48	-6.15	-6.65	-2.06	-2.70	-2.76
	(t)	(-1.74)	(-3.43)	(-1.89)	(-2.48)	(-3.24)	(-3.26)	(-0.96)	(-2.45)	(-1.6)
	L_{NRC}	7.35	2.93	14.73	8.55	6.42	7.85	7.84	1.61	7.92
	(t)	(3.25)	(2.23)	(2.96)	(3.98)	(3.57)	(4.15)	(3.35)	(1.57)	(5.11)
	R^2	14.3	6.94	10.52	18.83	14.65	18	16.12	3.01	22.95