

Duarte, Fernando M.

Working Paper

Inflation risk and the cross section of stock returns

Staff Report, No. 621

Provided in Cooperation with:

Federal Reserve Bank of New York

Suggested Citation: Duarte, Fernando M. (2013) : Inflation risk and the cross section of stock returns, Staff Report, No. 621, Federal Reserve Bank of New York, New York, NY

This Version is available at:

<https://hdl.handle.net/10419/93630>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Federal Reserve Bank of New York
Staff Reports

Inflation Risk and the Cross Section of Stock Returns

Fernando M. Duarte

Staff Report No. 621
May 2013



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the author and are not necessarily reflective of views at the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the author.

Inflation Risk and the Cross Section of Stock Returns

Fernando M. Duarte

Federal Reserve Bank of New York Staff Reports, no. 621

May 2013

JEL classification: E00, E31, E44, G00

Abstract

I establish that inflation risk is priced in the cross section of stock returns: Stocks that have low returns during inflationary times command a risk premium. I estimate a market price of inflation risk that is comparable in magnitude to the price of risk for the aggregate market. Inflation is therefore a key determinant of risk in the cross section of stocks. The inflation premium cannot be explained by either the Fama-French factors or industry effects. Instead, I argue the premium arises because high inflation lowers expectations of future real consumption growth. To formalize and test this hypothesis, I develop a consumption-based general equilibrium model. The model generates a price of inflation risk consistent with my empirical estimates, while simultaneously matching the joint dynamics of consumption and inflation, the aggregate equity premium, and the level and slope of the yield curve. My model suggests that the costs of inflation are significant: A representative agent would be willing to give up 1.5 percent of lifetime consumption to eliminate all inflation risk.

Key words: inflation, stock returns

Duarte: Federal Reserve Bank of New York (e-mail: fernando.duarte@ny.frb.org). The author is grateful to Ricardo Caballero and Leonid Kogan for invaluable comments and discussions. He also thanks Hengjie Ai, Christine Breiner, David Cesarini, Hui Chen, Maya Eden, Xavier Gabaix, Jonathan Goldberg, Jennifer La'O, Guido Lorenzoni, Gustavo Manso, Marti Mestieri, Matt Notowidigdo, Sahar Parsa, Michael Powell, Jenny Simon, Alp Simsek, Ivo Welch, and seminar participants at the Massachusetts Institute of Technology, the MIT Sloan School of Management, Duke University's Fuqua School of Business, the Tepper School of Business at Carnegie Mellon, the Smith School of Business at the University of Maryland, and the Federal Reserve Bank of New York, and particularly Adrien Verdelhan for helpful comments and discussion. 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.

1 Introduction

In this paper, I document how and why inflation risk is priced in the cross-section of stock returns. Using a Fama-MacBeth (1973) procedure, I find that stocks whose returns are negatively correlated to inflation shocks command a risk premium. I estimate the market price of inflation risk to be -0.33 when measured as the Sharpe ratio of an inflation-mimicking portfolio. The price of inflation risk is therefore comparable in magnitude to the price of risk for the aggregate market. The negative price of risk means that inflationary periods correspond to bad states of nature: investors are willing to accept lower unconditional returns when holding securities that are good hedges against inflation. I argue the premium arises because high inflation today predicts low real consumption growth over many subsequent periods. I develop a model that uses the relationship between inflation and consumption to generate a price of inflation risk consistent with my empirical estimates.

By studying the cross-section of stock returns, I not only uncover a new source of information about the inflation 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 models of the term structure and the emergence of Treasury Protected Inflation Securities (TIPS), estimates of the inflation risk in the bond market have become more reliable and widely available¹. 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 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)

¹Ang 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).

using rare disasters and Bansal and Shaliastovich (2010), who use long-run risk. If the fundamental mechanisms of the real effects of inflation originate at the level of individual households or firms, studying the cross-section of stocks can provide valuable additional information that is masked in the aggregate market and the yield curve.

The variation cross-sectional returns associated with inflation is not well described by any of the risk factors most commonly used to price assets. For example, the Fama-French factors have a pricing error of 2.8% per year when confronted with portfolios sorted on exposure to inflation. Industry effects also fail to explain a significant fraction of the spread in returns of inflation-sorted portfolios. Consequently, firm characteristics that differ across sectors of the economy —like menu costs, leverage, tax liabilities or labor relations— although important, should be supplemented by further factors to fully understand stocks’ cross-sectional heterogeneity in inflation risk.

I propose an explanation of the cross-sectional inflation premium by arguing that high inflation is a bad state of nature because it predicts low future real consumption growth. I formalize and test this hypothesis by developing a consumption-based equilibrium model. The model takes the stochastic processes for consumption and inflation as given and asset prices are then determined endogenously through the representative agent’s Euler equation. After estimating parameters using generalized method of moments (GMM), I show that the model can quantitatively replicate the observed inflation premium while simultaneously matching key empirical moments of consumption, inflation, bond yields and the aggregate stock market.

To generate an inflation premium consistent with the data, my model has three key ingredients, all of which are necessary. The first ingredient —as already mentioned— is that high inflation predicts low future real consumption growth. I estimate that an increase of one percentage point in inflation this month is associated with a decrease of 2.3 percentage points in real consumption growth over the next two years. Additionally, I show that several lags of inflation are useful in predicting consumption, even after controlling for current inflation and current consumption growth. Piazzesi and

Schneider (2005) also find that inflation is a leading indicator for consumption and use this relationship to rationalize the inflation premium in the yield curve.

The second ingredient is that inflation is persistent. 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). Ang et al. show that the first-order autocorrelations of inflation at the monthly and quarterly frequencies are 0.92 and 0.77 respectively. Furthermore, inflation persistence decays slowly over the business cycle, with a first-order autocorrelation of 0.35 at 10 quarters. That inflation is quite persistent will be important in my model to quantitatively match the large inflation premium: more persistent inflation induces a larger market price of inflation risk because it affects consumption growth negatively for a longer period of time.

The third 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 my model. This property of EZ utility is explored by many authors in the macro-finance literature².

To give brief intuition of the model, consider what happens when a positive inflation shock hits a two period economy. Inflation unexpectedly jumps up and remains above its initial value in the second period. Consumption is unchanged in the first period and predictably decreases in the second period. The price of the wealth portfolio—which is simply a claim to future consumption—will change due to income and substitution effects. The income effect makes the price of period-2 consumption go down since the representative agent’s wealth has decreased and therefore demands less consumption. The substitution effect makes the price of period-2 consumption increase

²Bansal and Yaron (2004) use it in the context of long-run risk. Piazzesi and Schneider (2005) use it for bond pricing. Binsbergen et al. (2008), Caldara et al. (2009), Darracq et al. (2010), and the references there, use it in the context of dynamic stochastic general equilibrium models (DSGEs). Andreasen (2009), Rudebusch et al. (2009), Guvenen (2009), Amisano et al. (2009) use it in New Keynesian frameworks. Levin et al. (2008) use it for optimal Ramsey allocation.

because the representative agent would like to smooth her consumption path by shifting consumption away from period 1 and into period 2. When the elasticity of intertemporal substitution (EIS) is greater than one, which is the relevant case in my model, the representative agent is not willing to pay a high price to smooth consumption and the income effect dominates the substitution effect. The wealth portfolio now has lower returns. An inflation-mimicking portfolio, since it co-varies negatively with the return on the wealth portfolio, reduces the volatility of expected consumption growth. If the representative agent is averse to risk in expected consumption growth³, inflation shocks have a negative market price of risk.

Conceptually, my model builds on Parker and Julliard's (2005) idea that ultimate consumption risk depends on an asset's correlation not only with present consumption but also with consumption growth over many subsequent periods. The particular mechanism I consider is most closely related to Bansal and Yaron's (2004) long-run risk model. In their model, asset prices are driven by a small, persistent long-run predictable component of consumption. Long-run risk is priced in their model for the same reasons that inflation is priced in mine. However, there are three important differences. First, predictability of consumption using inflation, and inflation persistence itself, operate at business cycle frequencies. Bansal and Yaron's long-run risk operates at substantially lower frequencies, of 10 years or more. Second, inflation shocks have a higher variance than long-run risk shocks. The combination of higher volatility and lower persistence of inflation shocks makes their inflation premium comparable in magnitude to the premium earned by the low volatility and high persistence long-run risk shocks. Third, inflation is directly observable while long-run risk must be inferred from asset prices using the model's assumptions. The observability of inflation provides key additional moments to test my model. In particular, a successful model must match, as I do, the correlation between asset returns and inflation, while having a realistic process for

³Aversion to risk in expected consumption growth is equivalent to having a preference for early resolution of uncertainty. For EZ preferences, this happens when the product of the coefficient of relative risk aversion and the EIS is greater than one.

inflation with several lags and heteroskedasticity.

Related literature Almost all studies of the inflation premium look at the time series of aggregate stock returns and the yield curve instead of the cross-section of stock returns. Bansal and Shaliastovich (2010) attach to the canonical long-run risk model a process for inflation and use it to price nominal bonds. The inflation premium in their model arises not because inflation feeds back into consumption, but because inflation is exposed to the same real shocks that drive consumption and long-run risk. Gabaix (2008) explains the inflation premium in a model with rare disasters. Inflation is priced because when a disaster occurs, inflation tends to increase. Wachter (2006) explains the inflation premium using i.i.d. consumption and habit-formation. All three models, unlike the one I propose, were originally designed to explain classic pricing puzzles such as the equity premium puzzle and the failure of the expectations hypothesis.

A notable exception to using time-series estimates is Chen, Roll and Ross (1986). They postulate a model with many macroeconomic and aggregate factors, including inflation innovations. The emphasis is not in estimating risk premia precisely but to find plausible state variables for asset prices. They find that inflation is priced only for the 1968-1977 subsample and, in contrast to my results, that stocks are weak hedges against inflation. Their study differs from mine in several respects. They use 20 portfolios sorted on size as their test assets, while I use individual stocks and portfolios sorted on inflation risk, which is the relevant variable. Their sample ends in 1984 and contains many fewer securities than mine. Finally, they use yearly instead of monthly data and a different Fama-MacBeth procedure.

Piazzesi and Schneider (2005) analyze how the fact that inflation predicts future consumption growth affects the pricing of nominal bonds. Their paper can be viewed as the counterpart of my paper in the bond market. They argue, consistent with my findings, that inflation is bad news for future consumption, producing an upward sloping yield curve. While I use rational expectations throughout, they study the impact of changing investors' beliefs. They find that learning is important in an environment in

which investors cannot easily distinguish permanent and transitory movements in inflation. While I do not allow for learning or endogenously changing beliefs, I do analyze exogenous structural changes in the relationship between inflation and consumption over my sample. Another important difference is that I allow for heteroskedasticity in the inflation process.

2 Measuring Inflation Risk

In this section, I estimate that the price of risk of inflation shocks is -0.33 using a two-step Fama-MacBeth procedure. The differences in stock returns arising from inflation risk are captured neither by the Fama-French factors nor by other standard pricing models. Industry effects play a limited role in explaining the large heterogeneity in inflation exposures present in the cross-section of stocks.

2.1 Data

I use monthly data for the period 1959-2009. For inflation, I use the consumer price index (CPI) from the Bureau of Labor Statistics. For consumption growth, I use real personal consumption expenditures (PCE) in non-durables and services from the Bureau of Economic Analysis. Individual stock returns are from the Center for Research in Security Prices (CRSP) and I use the CRSP value weighted index for aggregate market returns. I use the entire universe of CRSP, which for my sample has 27,688 companies represented and 3,262,429 total month-company observations⁴. Yields for bonds are obtained from the Fama-Bliss bond files, and the risk-free rate is from the Fama risk-free rate files, both available in CRSP. Fama-French and momentum factors, industry portfolios, short and long-term reversal factors are all obtained from Professor French's website. The Cochrane-Piazzesi factor is from Professor Piazzesi's website. Oil prices are from the International Monetary Fund. I chose 1959 as the start of my sample to coincide with availability of PCE data.

⁴In the following section, I will eliminate 1% of the sample due to outliers.

2.2 Inflation betas

In the first step of the Fama-MacBeth procedure, I measure firms' exposure to inflation by estimating their "beta", just as one would do in a CAPM setting. Instead of market returns, I use inflation innovations as the risk factor. I only use past information when estimating risk exposures to eliminate look-ahead bias when I later form inflation portfolios —an investor living in any period of my sample could have replicated my results in real time. For each stock $n = 1, \dots, N$ and each time period $t = 1, \dots, T$, I find an estimate of the inflation beta $\hat{\beta}_{n,t}$ by running a weighted least-squares regression of excess returns⁵ $R_{n,t}^e$ on inflation innovations $\Delta\pi_t = \pi_t - \pi_{t-1}$, using all observations in the interval $[1, t-1]$. Since the dependent variables are excess returns, I am considering the exposure of *real* returns to inflation⁶. I use weights that decay exponentially with the distance between observations and have a half-life of five years. This estimator efficiently captures time variation in betas by using all available past information. I choose to have decaying weights because recent observations are more likely to contain information about inflation exposure going forward. The weighted-least squares estimator resembles the original 5-year rolling window estimator used in Fama and French (1992), with the advantage of using more information to produce smoother estimates⁷. The estimator is given by

$$\left(\hat{\alpha}_{n,t}, \hat{\beta}_{n,t}\right) = \arg \min_{\alpha, \beta} \sum_{\tau=1}^{t-1} K(t-\tau) \left(R_{n,\tau} - R_{\tau}^f - \alpha - \beta \Delta\pi_{\tau}\right)^2 \quad (1)$$

with weights

$$K(t-\tau) = \frac{\exp(-|t-\tau-1|h)}{\sum_{\tau=1}^{t-1} \exp(-|t-\tau-1|h)}. \quad (2)$$

⁵Throughout the paper, I will use a superscript "e" to denote excess returns, so for example, if the risk free rate is R_t^f , excess returns for stock n at time t is $R_{n,t}^e = R_{n,t} - R_t^f$.

⁶ $R_{n,s}^{\text{nominal}} - R_{f,s}^{\text{nominal}} = (R_{n,s}^{\text{nominal}} - \pi_s) - (R_{f,s}^{\text{nominal}} - \pi_s) = R_{n,s}^{\text{real}} - R_{f,s}^{\text{real}}$

⁷I later report that estimates for the market price of inflation risk are similar when using the weighted-least squares and the 5-year rolling window. The main difference is in the standard error. The weighted least squares estimator also performs better when predicting ex-post exposures, which I attribute to the reduction in noise coming from using more observations.

I use $h = \log(2)/60$ to get a half-life of 5-years. The least squares estimator in (1) can also be thought of as a kernel estimator with exponential kernel given by (2) and bandwidth⁸ h . The 5-year rolling window estimator also satisfies (1) but uses a flat kernel that becomes zero after 5 years. The interpretation of the estimates is straightforward: A value of $\hat{\beta}_{n,t} = -2$, for example, means that a change in inflation of one percentage point is associated with a decrease in excess returns of two percentage points over the same time period. Ex-ante (backward-looking) betas are useful as a measure of risk insofar as they also capture risk exposure going forward. Elton, Gruber, and Thomas (1978) show that making a simple Vasicek adjustment to the ex-ante betas can make ex-ante exposures better predictors of ex-post exposures. The Vasicek adjustment is a Bayesian updating procedure in which the prior distribution is given by the beta $\hat{\beta}_{n,t}$ estimated from the time-series and the posterior distribution is obtained by incorporating information about the cross-sectional distribution of $\hat{\beta}_{n,t}$ for fixed t . The formula is:

$$\hat{\beta}_{n,t}^{adj} = w_{n,t}\hat{\beta}_{n,t} + (1 - w_{n,t})\mathbb{E}_{XS}[\hat{\beta}_{n,t}] \quad (3)$$

$$w_{n,t} = 1 - \frac{var_{TS}(\hat{\beta}_{n,t})}{var_{TS}(\hat{\beta}_{n,t}) + var_{XS}(\hat{\beta}_{n,t})}, \quad (4)$$

where the subscripts TS and XS denote means and standard deviations taken over the time series (over the variable t) and the cross section (over the variable n) respectively. Vasicek betas are a weighted sum of each stock's beta and the mean beta in the cross-section. The adjustment places higher weight on individual betas that are more precisely estimated and higher weight on the cross-sectional mean when the cross-section has less dispersion. From this point forward, inflation betas refer to estimated Vasicek-adjusted betas and I will drop the superscript *adj* and the hat.

Figure 1 depicts the histogram of inflation betas for four different time periods. I have selected January of 1979, 1983, 1994 and 2009 to portray the shape of the distribution of betas in different macroeconomic conditions and inflationary regimes. Betas

⁸See Ang and Christensen (2010).

have significant dispersion in all four time periods, with values⁹ ranging from -25 to +15. The mean of the distribution moves considerably through time. Strikingly, during the downturn of 2009 the mean inflation beta is positive. Campbell, Sunderam and Viceira (2010) document, consistent with Figure 1, that the “nominal-real” covariance of inflation with the real economy is positive on average but has been negative since the downturn of 2001.

Figure 2 shows the time series of inflation beta for the aggregate market, a five-year zero-coupon nominal Treasury bond and four well-known firms representative of different sectors of the economy. The market’s inflation beta is a good proxy for the mean of the distribution of betas shown in Figure 1. Figure 2 displays mostly negative betas for the aggregate market in the 1980’s and 1990’s, with positive values at the beginning and end of the time series. Compared to the market, the excess returns on the 5-year bond have a small and almost constant beta. This means that the spread of the 5-year real risk-less rate over the 1-month real risk-less rate has little exposure to inflation. Figure 2 also shows that individual companies’ betas tend to move together with the market, especially at lower frequencies, yet still exhibit considerable cross-sectional heterogeneity. The correlation between firms’ betas is also time varying. For example, Coca-Cola and General Electric move in lockstep in the 1970’s but move in opposite directions in the 1980’s. I focus exactly on exploiting this type of cross-sectional variation to identify the inflation premium. In this respect, my research differs substantially from Campbell, Sunderam and Viceira (2010) and from most other studies of the inflation premium, as explained in the introduction.

2.3 Inflation-sorted portfolios

Figure 2 makes clear that individual estimates of inflation betas have substantial high frequency variation. If some of that variation is due to noise (e.g. measurement error), statistical inference of the inflation premium or the distribution of betas can be chal-

⁹I cut from the sample stocks with betas in the top and bottom 0.5% of the distribution because their betas are extreme. Results are robust to windorizing with a threshold for betas of ± 25 .

lenging. Following Fama and MacBeth (1973), and Black, Jensen and Scholes (1972), I address this problem by forming portfolios, in the hope that idiosyncratic variation will average out within each portfolio. To form portfolios, I perform a double-sort based on inflation betas and size. Size is measured by market equity (price multiplied by shares outstanding) in June of the previous year, just as in the construction of the Fama-French factors. At time t , each firm belongs to one of 10 deciles of the inflation beta distribution and one of 10 deciles of the size distribution. I create 100 value-weighted portfolios by grouping stocks that belong to the same beta and size deciles. This procedure implies rebalancing portfolios every month. In practice, however, around 80% of the firms remain in the same portfolio after one year and about 30% after 5 years, which is not surprising given that estimates in betas for two consecutive months only differ by one observation. As when creating size and book-to-market factors in Fama-French, I reduce the 100 inflation-and-size portfolios to 10 inflation-only portfolios by collapsing the size dimension. Specifically, I create 10 new value-weighted portfolios from the original 100 value-weighted portfolios by grouping together portfolios that belong to the same decile of inflation betas. The resulting portfolios have differential ex-ante exposure to inflation innovations but little variation in size. Ideal test assets have identical exposure to every risk factor except for inflation. In that case, any difference in mean returns can be interpreted as compensation for inflation risk. When reducing the original 100 portfolios to the new 10 portfolios, I eliminate most of the differential exposure to size. Conveniently, size smoothing also makes exposures to the market and other risk factors much more homogeneous across portfolios. The resulting portfolios are therefore much closer to ideal test assets and allow me to better isolate the effects of inflation¹⁰.

One main feature of the 10 inflation portfolios is that they exhibit a spread in returns:

¹⁰There are other advantages to averaging along size. The model I develop does not have any size or book-to-market effects, so using assets that have size and book-to-market exposure will only complicate estimation and make results difficult to interpret.

Having 10 instead of 100 portfolios also makes it feasible to estimate non-linear standard errors with GMM when I test factor models.

the highest beta portfolio has a mean return of 5.16% per year, compared to 6.91% for the lowest beta portfolio. To put the spread of 1.75% in perspective, the analogous spread induced by size and book-to-market differences are 2.6% and 4.8% respectively. Columns 1 and 3 of Table 1 show the mean ex-ante inflation betas and the mean excess returns for all portfolios. Column 2 shows the mean post-ranking betas. To construct the post-ranking beta of a portfolio at time t , I freeze the time- t portfolio weights and regress the excess returns of the fixed-weights portfolio on inflation innovations, using the five years of data starting at $t + 1$. The post-ranking betas can be thought of as an out of sample test for the estimates of the ex-ante betas. Table 1 shows that ex-ante betas align neatly with post-ranking betas, showing that portfolios do capture ex-post exposure to inflation. Post-ranking betas are squeezed together compared to ex-ante betas, which is a well-known effect in this set-up¹¹. The range of post-ranking betas is also more reasonable than for the noisy ex-ante betas; it is difficult to imagine a stock whose returns have a systematic 10-fold reaction to inflation.

A first pass “long-short” estimator of the inflation premium can be found by looking at the last row of Table 1, which computes the spread in betas and returns between the highest and lowest beta portfolios. Dividing the 1.75% spread in returns by the difference in their ex-post betas, I find a slope of $\lambda^{Long-Short} = -0.74$. I divide this crude non-linear slope estimator by the standard deviation of returns of the long-short portfolio to find a market price of inflation risk of -0.23. Higher inflation beta is associated with lower mean returns, which implies a negative market price of risk. The price of risk obtained in this way is not statistically significant. However, combined with how well returns align with ex-ante and post-ranking betas, the obtained value for $\lambda^{Long-Short}$ provides further motivation to more deeply analyze how inflation is priced in the cross-section of stock returns. The estimator $\lambda^{Long-Short}$ is inefficient because it discards changes in the cross-section of stocks that occur every period – it simply averages across time first and then ignores all but the two corner portfolios. In the next

¹¹See Elton and Gruber (1995).

section, I will formally statistically test whether the spread in returns can be attributed to the differences in betas by using the entire cross-sectional variation of the 10 portfolios over time. The test will confirm that inflation is priced in the cross-section of stock returns in a statistically significant manner, with a market price of risk of -0.33.

Table 2 analyzes other characteristics of the portfolios. The first four columns show that portfolios are not systematically different in terms of their exposure to market, size, book-to-market or momentum. The last three columns show that portfolios are only slightly different in terms of their industry composition. Industries are defined by the first two digits of the Standard Industry Classification code (SICCD). Column 5 is a Herfindahl industry concentration index obtained by summing the squares of the shares of firms in each industry within a portfolio. A value of 1 for this index means that all companies in the portfolio belong to the same industry, and the closer the index is to 0, the more diversified the portfolio is across industries. Column 6 provides a measure of distance in the distribution of industries between a given portfolio and the remaining 9 portfolios. The index is normalized so that a value of 1 means that the portfolio in question has the same exact distribution of industries as the remaining 9 portfolios, and the measure decreases toward 0 when there is no intersection between industries in the portfolios. Column 7 uses the same measure as Column 6 to compare each portfolio's industry distribution with the distribution of the same portfolio five years later. Low beta, high return portfolios have slightly higher industry concentration and persistence but the message of Columns 5-7 is that portfolios are by and large well-diversified, similar to each other and not very persistent in their industry composition. Hence, there are no large industry differences within portfolios, across portfolios or along different time periods. The observed pattern implies that the bulk of the heterogeneity of inflation risk in the cross-section of stock returns cannot be ascribed to industry effects. Table 2 has important implications for any theory that attempts to explain the variation of stock returns induced by their differential exposures to inflation. For example, menu costs, taxes, leverage, or labor relations between a firm and its employees, although im-

portant, cannot provide a complete explanation of why different firms react differently to inflation, as these characteristics vary strongly by industry while inflation-sorted portfolios don't.

To further confirm that the spread in returns of inflation portfolios are not driven by market exposure, size, industry effects or other standard factors that are commonly used to explain returns, I run time-series regressions of inflation portfolios' returns on different risk factors. I consider the Fama-French factors, momentum, short and long-term reversal factors, liquidity, oil, industry portfolios and the Cochrane-Piazzesi factor. Table 3 shows results for different combinations of the factors. The mean absolute pricing errors –the average of the absolute value of the intercepts or “alphas”– are on the order of 1.88% to 4.28% per year, which is of the same order of magnitude as the difference in returns between the lowest and highest inflation beta portfolios. In addition to being economically sizable, I show in Table 3 that the pricing errors are also statistically different from zero by performing a Gibbons-Ross-Shanken (GRS) test (1989). The GRS test is an F-test adjusted for finite sample bias for the hypothesis that the pricing errors for the 10 portfolios are jointly zero. Oil and industry portfolios perform better than other factors under the GRS measure but still have large pricing errors that are statistically different from zero at the 1% level.

2.4 Market price of inflation risk

In this section, I use the inflation portfolios of the last section to perform the second step of the Fama-MacBeth procedure. The goal of this section is to produce an estimate for the market price of inflation risk implied by the cross-section of stock returns.

To do so, I start by running one cross-sectional regression for each time period t . The dependent variables are the time t returns for the 10 inflation portfolios and the independent variables are the estimated time t post-ranking inflation betas obtained in

the first stage of Fama-MacBeth ¹²:

$$R_{p,t}^e = a_t + \lambda_t \beta_{p,t} + \varepsilon_t \quad (5)$$

$$p = 1, \dots, 10.$$

The estimated coefficient $\hat{\lambda}_t$ measures the average extra returns earned by assuming one extra unit of inflation beta at time t . Table 4 reports the average annualized price of inflation risk in my sample. I show both the average $\bar{\lambda}$ of λ_t , which gives the risk premium per unit of inflation beta, and $\bar{\lambda}/std(\Delta\pi_t)$, which gives the risk premium associated with a one standard deviation shock in inflation innovations. Since λ_t is persistent over time, I use Newey-West standard errors with 12 lags to construct standard errors and verify that both estimates for the price of risk are statistically different from zero.

The negative value of $\bar{\lambda}$ implies that inflation shocks correspond, on average, to bad states of nature. Holding assets that have low excess returns when inflation is increasing must offer higher mean returns as compensation for bearing inflation risk. Another way to understand the inflation premium is to imagine that each period we move from the first to the last decile in the distribution of ex-post betas. In this case, the associated expected increase in returns is 8.79%. This exercise is not the same as moving from portfolio 1 to 10 every period. When moving from decile to decile in the distribution of individual betas, we do it unconditionally, while portfolios are conditioned on size because of the initial double-sort. Stocks with lower inflation beta are also smaller on average. When moving unconditionally across the distribution, both effects are captured in the extra returns. The 8.79% can then be thought of as a total derivative, while the 1.75% spread in returns is closer to a partial derivative.

To compare the market price of inflation risk to the aggregate market's price of risk, it is more useful to look at the normalized $\bar{\lambda}/std(\Delta\pi_t)$. Table 4 shows that under this measure, the price of risk of inflation is comparable to the market's, for which 0.3 is a

¹²Thus, the independent variables in this regression are themselves regression coefficients.

good approximation. This means that an inflation-mimicking portfolio has about the same Sharpe ratio as the market. Inflation is therefore an important component of risk in the cross-section of stock returns.

An alternative to forming portfolios when measuring prices of risk, advocated by Ang, Liu and Schwarz (2010), is to use individual betas in the second stage regression (5). The rationale is that while betas may be more precisely estimated when forming portfolios, efficiency in the estimate for $\bar{\lambda}$ is increased when no information about the cross-section is destroyed. Column 2 of Table 4 reports the estimates for the inflation premium using individual stocks. As a robustness check, Column 3 reports the estimate for $\bar{\lambda}$ obtained when creating portfolios using a 5-year rolling window instead of an exponential kernel. All three measures are similar, especially the ones in Columns 2 and 3. The estimator using individual stocks does have a smaller variance, confirming the message of Ang, Liu and Schwarz (2010).

3 Model

In this section, I present a consumption-based model that can explain and quantitatively match the inflation premium I estimated in the last section. I first consider a simple version to illustrate how assets are priced. This version is similar in its mechanics to Bansal and Yaron's (2004) long-run risk model. I then present a version of the model that has richer processes for inflation, consumption and dividends, that is more suitable for quantitatively estimating and testing of the model.

3.1 Set-up

Environment The model is an exchange economy with a single representative agent. Time is discrete and indexed by $t \in \{0, 1, \dots\}$. For each period t , there is one consumption good denoted by C_t which represents the economy's real aggregate consumption. I will use lower-case letters to denote the logarithm of the corresponding variable, so for example $c_t = \ln C_t$.

The joint process for consumption growth $\Delta c_{t+1} = \ln C_{t+1} - \ln C_t$ and inflation π_t is exogenous and given by

$$\pi_{t+1} = \mu_\pi + \rho_\pi (\pi_t - \mu_\pi) + \sigma_\pi \varepsilon_{t+1} \quad (6)$$

$$\Delta c_{t+1} = \mu_c + \rho_c (\pi_t - \mu_\pi) + \sigma_c \eta_{t+1} \quad (7)$$

The stochastic disturbances ε_s, η_r are i.i.d. standard normal for all $s, r \in \{0, 1, \dots\}$. Eq. (6) shows that inflation follows an $AR(1)$ process with constant auto-regressive coefficient $\rho_\pi \in (-1, 1)$ and unconditional moments controlled by the constants μ_π and σ_π^2 :

$$\mathbb{E}\pi_{t+1} = \mu_\pi \quad (8)$$

$$Var(\pi_{t+1}) = \frac{\sigma_\pi^2}{1 - \rho_\pi^2} \quad (9)$$

The $AR(1)$ specification for inflation captures, in a stylized way, the persistent nature of inflation. Eq. (7) models real consumption growth as being a predictable function of past inflation. When an inflation shock ε_t hits the economy, inflation reacts contemporaneously and the effect extends into subsequent periods. However, consumption growth only starts reacting to the inflation shock the next period. Thus, inflation leads consumption growth and inflation shocks translate not into changes in current consumption, but into shocks to consumption expectations. The sign and magnitude of the predictability is given by the constant ρ_c . When $|\rho_c|$ is large and when ρ_π is close to 1, inflation shocks have a large, persistent effect on future consumption growth. When σ_π^2 is large, inflation is very volatile and small inflation shocks also have a large effect on future consumption and consumption expectations.

Table 5 justifies my choice for process (6). Column 1 of Table 5 shows that eq. (6) is not unreasonable as a first approximation, although it does mask many features of inflation dynamics. Column 2 shows that inflation has more inertia than hinted by its first autoregressive coefficient, with several lags significant and comparable in magnitude to the first lag. Ang, Bekaert and Wei (2007) show that at the ten quarter

horizon, the first order autocorrelation of inflation is still 0.35. Including multiple lags in the monthly process, or looking at longer horizons, makes clear that inflation shocks are active throughout the entire business cycle, spanning a window of 2 to 4 years. A potentially restrictive assumption in (6) is that the parameters for inflation are not time-varying. Taylor (1998) argues for a break in inflation regimes before and after the Volcker era¹³. Without trying to produce sophisticated econometric analysis of breaks and switching regimes, Panel A of Table 6 shows basic evidence that inflation was more volatile and persistent before 1980. Under the interpretation in Taylor (1998), the reason is that the Federal Open Market Committee accommodated inflation before Volcker but started leaning against it in the early 1980's. For both the simple model and my main specification, I will keep the inflation parameters constant. However, I will later exploit the time-variation in inflation parameters to test my model. This will be an important validation of my model, because it focuses on the key mechanism generating the inflation premium. I will show that the model successfully replicates the change in inflation premium observed in the data when inflation persistence and volatility change.

Table 7 addresses the empirical evidence of consumption predictability using inflation. Column 1 shows that an increase in one percentage point in inflation this month is associated with an expected decrease of 1.5 percentage points in real consumption growth over the next year. Column 2 shows that up to three lags of inflation contribute in predicting consumption¹⁴. Column 3 shows that if enough lags of inflation are included, past consumption need not be included in (7). The last three columns of the table show the same regressions for a two-year horizon.

I take the process for consumption as exogenous and therefore do not attempt to explain why inflation predicts consumption. There is already a large and sophisticated body of literature in Macroeconomics that puts forward several mechanisms that gen-

¹³On the other side of the argument, Orphanides (2004) uses information available to the FMOOC in real time to argue there was no change in regimes. I consider both cases.

¹⁴The two subsequent lags are not significant and get smaller in magnitude.

erate real effects of inflation and real consumption predictability. For example, Clarida, Gali and Gertler (1999) provide a reduced-form model that captures a large class of dynamic equilibrium models of nominal rigidities. After computing expectations that ultimately come from agents' optimization, their process for inflation and consumption can be mapped to eqs. (6) and (7) if I allow for contemporaneous correlation between inflation and consumption (which I do below). Other explanations for the real effects of inflation include menu costs, rational inattention and informational frictions¹⁵. I take no position as to which explanations are correct or quantitatively important, but instead empirically estimate a joint process for consumption and inflation that is flexible enough to accommodate any of these models and capture their main dynamic characteristics.

If I substitute inflation for long-run risk in eqs. (6) and (7), I obtain the same basic specification as Bansal and Yaron (2004). However, my model differs conceptually from theirs in significant ways. First, inflation is observable, so its process and its relation to consumption can be estimated directly. In the long-run risk model, the variable predicting expected consumption growth is inferred from asset prices and assumptions about preference parameters. In the next sections, I will use the observable properties of consumption and inflation, such as the presence of multiple lags, to depart from Bansal and Yaron's (2004) specification. This departure will provide me with additional moment restrictions to test my model in a way that would not be possible in the long-run risk model. Second, even though inflation and long-run risk have a similar functional form, their stochastic properties are drastically different. Long-run risk has an extremely long half-life, operating at frequencies of 10 to 30 years instead of the 1 to 2 years for inflation. The counterpart for a long half-life of long-run risk is that its volatility is about 200 times smaller than inflation's. The lower persistence and higher volatility of inflation generates a market price of risk similar to the one for the higher persistence, lower volatility long-run risk.

¹⁵See the introduction for references

Representative Agent The representative agent has recursive Epstein-Zin-Weil (EZ) preferences. If she enters period t with wealth W_t , then her utility is defined by

$$U_t(W_t) = \left((1 - \delta) C_t^{1-1/\psi} + \delta E_t [U_{t+1}(W_{t+1})^{1-\gamma}]^{\frac{1-1/\psi}{1-\gamma}} \right)^{\frac{1}{1-1/\psi}} \quad (10)$$

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 relative magnitude of risk aversion against the EIS. The EZ utility function is a generalization of the familiar constant relative risk-aversion (CRRA) utility function, which is obtained when $\theta = 1$. With CRRA utility, the coefficient of relative risk aversion is the inverse of the elasticity of intertemporal substitution. EZ disentangles these two conceptually different parameters—there is no reason to assume that the desire to smooth over time is the same as the desire to smooth over different states of nature. In fact, the EIS is important in a dynamic deterministic economy while the lack of uncertainty makes risk-aversion irrelevant. Similarly, γ is relevant in a static economy with uncertainty, while the EIS is not. Because inflation operates through the predictable component of consumption, the EIS will play a crucial role in determining the asset pricing implications of inflation.

Another important trait of the EZ utility is that it is not time-separable (it cannot be written as a sum of period utilities). To understand why non-separability is important, consider the following example from Duffie and Epstein (1992). An agent picks between two consumption plans for a long number of periods before any consumption is realized. Consumption plan A is obtained by tossing a fair coin every period and giving the agent high or low consumption in each period depending on the outcome of the toss in that period. Consumption plan B is obtained by a *single* coin toss before all consumption takes place and gives the agent high consumption in every period if heads and low consumption in every period if tails. In plan B, uncertainty about consumption is resolved early, while for plan A uncertainty is resolved gradually. When $\theta < 1$, which is the case I consider in my model, the representative agent prefers early resolution

of uncertainty —i.e. likes to plan ahead— and prefers plan B. This case occurs when $\gamma\psi > 1$, requiring high risk aversion or high EIS.

Early resolution of uncertainty can also be understood in terms of aversion to risk in consumption growth. Expected consumption growth is mean-reverting in plan A and constant for plan B. When the agent prefers early resolution of uncertainty, she also has a preference for less risk in expected consumption growth and plan B is more desirable than A. In my model, positive inflation shocks will command a premium because inflation induces this type of risk.

The representative agent's budget constraint is given by

$$W_{t+1} = R_{c,t+1} (W_t - C_t), \quad (11)$$

where $R_{c,t+1}$ is the return on the wealth portfolio. The wealth portfolio is the asset that pays consumption C_t each period as its dividends. Therefore, the agent consumes C_t out of wealth W_t and invests the remainder in the economy's aggregate endowment (consumption) technology.

Assets There are $1 + N$ assets in the economy indexed by n . The first asset is the wealth portfolio described in the last paragraph. The other N assets are defined to be levered claims to consumption as in Abel (1999). They pay exogenous dividends given by

$$\Delta d_{n,t+1} = \mu_{n,d} + l_n \rho_c (\pi_t - \mu_\pi) + \varphi_n \omega_{n,t+1} \text{ for } n = 1, \dots, N, \quad (12)$$

with $\omega_{n,t+1}$ i.i.d. standard normal across time, across assets and with respect to all other shocks in the economy. The process (12) for dividend growth has the same form as the process (7) for consumption growth. These processes can have different mean growth rates given by $\mu_{n,d}$, different volatilities given by φ_n and, more importantly, different exposures to inflation given by $l_n \beta_c$, where l_n is an asset-specific leverage parameter. The difference between the sum of dividends and aggregate consumption is assumed to come from other sources of income such as human wealth, which I do not explicitly model. When I estimate the parameters of the model with GMM, the N assets will be

mapped to the 10 inflation portfolios that I constructed in the empirical section of this paper.

3.2 Asset pricing

Representative agent's problem The representative agent's problem in period t is to pick a consumption path $\{C_s\}_{s=t}^{\infty}$ to maximize utility (10) subject to the budget constraint (11) and the exogenous processes for consumption (7) and inflation (6).

Stochastic discount factor and inflation-CCAPM The first order condition for the representative agent's problem implies that the return $R_{n,t+1}$ of any tradable asset n satisfies the Euler equation

$$1 = E_t [SDF_{t+1} R_{n,t+1}] \quad (13)$$

with a stochastic discount factor given by

$$\log SDF_{t+1} = sdf_{t+1} = \theta \log \delta - \frac{\theta}{\psi} \Delta c_{t+1} + (\theta - 1) r_{c,t+1}. \quad (14)$$

Rearranging equations (13) – (14) and using the log-normal structure of the set-up, I find that expected excess returns $E_t^e [r_{n,t+1}]$ follow a two factor model

$$E_t^e [r_{n,t+1}] \equiv -Cov_t (sdf_{t+1}, r_{n,t+1}) \quad (15)$$

$$= \frac{\theta}{\psi} Cov_t (\Delta c_{t+1}, r_{n,t+1}) + (1 - \theta) Cov_t (r_{n,t+1}, r_{c,t+1}) \quad (16)$$

Equation (16) states that the risk premium of asset n depends on the covariance of its returns $r_{n,t+1}$ with two factors. The first one is consumption growth Δc_{t+1} , just as in the consumption-CAPM. The second one is the return on the wealth portfolio $r_{c,t+1}$. The wealth portfolio arises with non-separable utility because, as explained above, the shape of the entire path of consumption matters when computing utility, rather than just the sum of expected utilities across periods. Because the return on the wealth portfolio $r_{c,t+1}$ is the price of the stream of all future consumption, it incorporates information about future consumption that is not included in Δc_{t+1} . For

an agent with non-separable utility, this additional information should be useful when computing marginal utilities and asset prices. For the EZ specification, it turns out that the return on the wealth portfolio $r_{c,t+1}$ is a sufficient statistic for the entire future path of expected consumption growth and hence the only other pricing factor beyond contemporaneous consumption. When $\theta < 1$, the representative agent is averse to risk in expected consumption growth. Assets that covary positively with the return on the wealth portfolio induce more expected consumption growth risk and have a positive risk premium.

Linearizing $r_{c,t+1}$ around the mean wealth-consumption ratio (Campbell 1991), the pricing equation (16) can be re-written as an inflation-consumption-CAPM:

$$E_t [r_{n,t+1}^e] = \gamma Cov_t (\Delta c_{t+1}, r_{n,t+1}) + \frac{(\gamma - 1/\psi)(\rho_c - 1/\psi)}{(\kappa_1 - \rho_\pi)(1 - 1/\psi)} Cov_t (\pi_{t+1}, r_{n,t+1}) \quad (17)$$

where κ_1 is a linearization constant that depends on the mean wealth-consumption ratio¹⁶. In a broader model, other state variables that determine returns on the wealth portfolio should be included. The prediction that inflation is priced because it predicts consumption should be robust to the inclusion of any other factors as long as inflation does not cease to have predictive power. The inflation-CCAPM (16) summarizes all the asset pricing content of the model. The market price of risk of consumption shocks is positive and given by $\gamma\sigma_c$ as in the CCAPM.

The magnitude and sign of the market price of risk for inflation shocks depend on the model's parameters. The larger the variance of inflation shocks σ_π , the larger the premium. In addition, if

- (i). inflation predicts consumption growth negatively ($\rho_c < 0$) and inflation is persistent ($\rho_\pi > 0$),
- (ii). the substitution effect dominates the income effect ($\psi > 1$), and
- (iii). the representative agent dislikes uncertainty in expected consumption growth

¹⁶In practice, the constant is close to 1 for most parameter values.

$$(\gamma - 1/\psi > 0),$$

then assets that have low returns when inflation shocks are positive will command a risk premium. The larger any of the three effects, the larger the premium.

I now give intuition for these components. The product of components (i) and (ii), given by $\frac{1}{1-1/\psi} \frac{\rho_c - 1/\psi}{\kappa_1 - \rho_\pi}$, captures how inflation shocks affect returns to the wealth portfolio. It “translates” eq. (16) to eq. (17). When $\rho_c < 0$, positive inflation shocks are bad news for future consumption growth. If $\psi > 1$, the substitution effect is larger than the income effect, and the adverse shock to expected consumption growth leads to smaller returns of the wealth portfolio. To see this, I use the budget constraint to write the log consumption-wealth ratio as

$$c_t - w_t + a = (1 - \psi) \mathbb{E}_t \left[\sum_j b^j r_{c,t+j} \right] = \mathbb{E}_t \left[\sum_j b^j (r_{c,t+j} - \Delta c_{t+j}) \right]$$

for some constants a and b . The first equality shows that when $\psi > 1$, today’s consumption decreases relative to wealth when expected returns rise – the substitution effect dominates. The second equality shows that a fall in expected returns is associated with a fall in expected consumption growth.

Component (iii) determines how shocks to the wealth portfolio are compensated in equilibrium. If the representative agent prefers earlier resolution of uncertainty and is therefore averse to risk in expected consumption growth, holding assets that covary with the returns on the wealth portfolio must have a positive risk premium, since they increase the volatility of expected consumption growth.

Combining ingredients (i), (ii) and (iii) implies that when an inflation shock hits the economy, expected consumption growth decreases, the returns on the wealth portfolio decrease and assets that pay off poorly in those bad states of nature command a risk premium.

Inflation Betas and Unconditional Returns To compare the model to the data, it is useful to understand the Fama-MacBeth procedure in the model. Inflation betas are endogenous in the model and given by the coefficient of the univariate regression of

excess returns on inflation innovations

$$\beta_{n,t} = \frac{Cov_t(\Delta\pi_{t+1}, r_{n,t+1}^e)}{Var_t(\Delta\pi_{t+1})} = \frac{l_n \rho_c - 1/\psi}{\kappa_{n,1} - \rho_\pi}. \quad (18)$$

The betas are not time-varying, but will be when I introduce stochastic volatility. Given that $\rho_c < 0$ and $\kappa_{n,1} - \rho_\pi > 0$, the sign of $\beta_{n,t}$ depends on the relative magnitudes of firms' leverage and the EIS. Positive leverage makes betas negative because they inherit the consumption risk induced by inflation. When the EIS is small, the agent is reluctant to change her consumption path after an inflationary shock. In that case, the adjustment to the new economic conditions requires large changes in prices.

Combining the inflation-CCAPM eq. (17) and the inflation betas, I obtain the cross-sectional regression in the second stage of the Fama-MacBeth procedure

$$r_{n,t}^e = \gamma \sigma_c^2 + \frac{(\gamma - 1/\psi)(\rho_c - 1/\psi)\sigma_\pi^2}{(\kappa_1 - \rho_\pi)(1 - 1/\psi)} \beta_{n,t} + \xi_t, \quad (19)$$

where ξ_t is a random disturbance. Comparing this equation to its empirical counterpart (5), we can identify the coefficient

$$\lambda_t = \frac{(\gamma - 1/\psi)(\rho_c - 1/\psi)\sigma_\pi^2}{(\kappa_1 - \rho_\pi)(1 - 1/\psi)} \quad (20)$$

Fama-MacBeth is therefore the right procedure to estimate inflation risk in my model. This fact not only helps justify the empirical methodology I employ, but also makes straightforward the comparison of the model and the data.

3.3 A general version

The model presented in the last section, although simple, can quantitatively generate an inflation premium as large as the one estimated in Table 4. I do a back of the envelope calculation of the premium with the following reasonable parameters: $\gamma = 3$, $\psi = 1.5$, $\rho_\pi = 0.6$, $\rho_c = -0.1$ and $Var(\pi_t) = (1 - \rho_\pi^2)\sigma_\pi^2 = 1.5\%$ per year. The market price of risk in this case is $\bar{\lambda} = -0.316$ which is very much in line with my empirical estimates.

However, as discussed in the previous section, the dynamics for consumption and

inflation are more elaborate than $AR(1)$. A more compelling model should generate the empirically observed price of inflation risk using processes for inflation and consumption that more closely adjust to the data. I therefore consider a richer version of my model and show that it can indeed price inflation in the cross-section of stock returns in accordance to my empirical estimates. The generalized processes for inflation, consumption and dividends are given by

$$\pi_{t+1} = \mu_\pi + \sum_{s=0}^2 \rho_{\pi,s} (\pi_{t-s} - \mu_\pi) + \sigma_{\pi,t+1} \varepsilon_{t+1} + \varphi_{\pi c} \sigma_{c,t+1} \eta_{t+1} \quad (21)$$

$$\sigma_{\pi,t+1}^2 = \sigma_\pi^2 + \sum_{s=0}^1 \nu_{\pi,s} (\sigma_{\pi,t-s}^2 - \sigma_\pi^2) + \sigma_{\pi w} u_{t+1} \quad (22)$$

$$\Delta c_{t+1} = \mu_c + \sum_{s=0}^2 \rho_{c,s} (\pi_{t-s} - \mu_\pi) + \sigma_{c,t+1} \eta_{t+1} \quad (23)$$

$$\sigma_{c,t+1}^2 = \sigma_c^2 + \nu_c (\sigma_{c,t-s}^2 - \sigma_c^2) + \sigma_{cw} w_{t+1} \quad (24)$$

$$\Delta d_{n,t+1} = \mu_{n,d} + l_n \sum_{s=0}^2 \rho_{s,c} (\pi_{t-s} - \mu_\pi) + \varphi_{nc} \sigma_{c,t+1} \omega_{n,t+1} \quad (25)$$

where the shocks $\varepsilon_t, \eta_t, u_t, w_t$ and $\omega_{n,t}$, are i.i.d. normal across time and across processes.

Consumption and inflation now have three lags of inflation and stochastic volatility. Dividends are still levered consumption with parameter l_n . The inclusion of lags is justified by Tables 5 and 7, and the discussion following eqs. (6) and (7). Stochastic volatility for inflation follows an $AR(2)$ process, creating GARCH-like effects, including heteroskedasticity and persistence. GARCH effects in inflation are documented in Bollerslev, Russell and Watson (2009) and Bollerslev (1986). Stochastic volatility for consumption and dividends are also $AR(2)$, which allows the model to match the widely documented time-varying volatility of stock returns.

The inclusion of stochastic volatility plays a dual role. As explained by Campbell and Beeler (2009) in the context of long-run risk, stochastic volatility in dividends

and consumption help generate a realistic aggregate equity premium. In my model, it will also help match the average level of returns in inflation portfolios. Stochastic volatility in consumption and dividends plays no role in generating a spread in returns across stocks with different exposures to inflation. Stochastic volatility of inflation, on the other hand, does increase the mean inflation premium. This is most easily seen in the context of the simple model. Using Jensen's inequality and replacing σ_π by a time-varying process $\sigma_{t,\pi}$ with mean $\bar{\sigma}_\pi$, we have

$$\bar{\lambda} = E[\lambda_t] = \frac{(\gamma - 1/\psi)(\rho_c - 1/\psi)}{(\kappa_1 - \rho_\pi)(1 - 1/\psi)} E[\sigma_{t,\pi}^2] \geq \frac{(\gamma - 1/\psi)(\rho_c - 1/\psi)}{(\kappa_1 - \rho_\pi)(1 - 1/\psi)} \bar{\sigma}_\pi^2. \quad (26)$$

As a last modification from the simple model, I allow correlation between contemporaneous inflation and consumption, parametrized by $\varphi_{\pi c}$. In many models, supply (e.g. productivity) shocks will tend to increase consumption and lower inflation. A demand shock, on the other hand, will tend to increase both. In those models, the correlation between consumption and inflation captures the relative strength of these two competing effects. Apart from theoretical considerations, there is a negative contemporaneous correlation between inflation and consumption at the monthly frequency that could not be otherwise captured. After estimating the parameters of the model, I find that this correlation does not contribute significantly to the inflation premium.

4 GMM Estimation

In this section, I estimate the parameters of the general model by using generalized method of moments. I find that the model can reproduce the observed market price of risk for inflation, the aggregate market's risk premium and its volatility, the level and volatility of the risk free rate and the level and slope of the yield curve, while simultaneously matching the processes for consumption and inflation. I use a standard 2-step feasible GMM. In the first step, I use the identity matrix as the weighting matrix. In the second step, I use as weighting matrix the inverse of the variance-covariance matrix estimated using the parameters found in the first step.

4.1 Moments

I classify the 59 moment conditions that I use into four groups:

1. **Consumption and inflation** (11 moments). I estimate all the OLS moments of eqs. (21) and (23), together with the variance-covariance matrix of same-period inflation and consumption growth. These are the natural moments to estimate for the linear exogenous processes for inflation and consumption.
2. **Inflation portfolios** (30 moments). I include the mean and variance of returns of the 10 assets of the model, together with their inflation betas. The empirical moments corresponding to these assets naturally come from the inflation portfolios constructed in the empirical section.
3. **Aggregate market** (6 moments). I match the mean of aggregate dividend growth and the mean, variance and inflation beta of the market's return. In addition, I include the mean and variance of the price-dividend ratio as a moment condition to highlight the model's ability to match a property of the aggregate market that proves difficult to match in other models. The volatility of the price-consumption ratio also helps identify consumption and inflation's stochastic volatility.
4. **Bonds** (12 moments). I incorporate as moment conditions the means and variances of nominal bond yields for 1, 2, 3, 4 and 5 year maturities together with the mean and variance of the risk-free rate.

4.2 Parameters

I divide the 53 parameters to be estimated into five categories:

1. **Consumption and inflation** (11 parameters). The vector of parameters is $\Theta_{c\pi} = (\mu_c, \mu_\pi, \varphi_{\pi c}, \rho_{\pi, s}, \rho_{c, s}, \sigma_c, \sigma_\pi)$ with $s = 0, 1, 2$. If I were estimating just the process for consumption and inflation (without stochastic volatility), the 11

parameters would be exactly identified from their corresponding 11 moment conditions and could be estimated by OLS.

2. **Inflation portfolios** ($3 \times 10 = 30$ parameters). The vector of parameters is $\Theta_p = (\mu_{n,d}, l_n, \varphi_{nc})$ for $n = 1, \dots, 10$. It includes the mean $\mu_{n,d}$, leverage parameter l_n and volatility σ_n of dividend growth for each stock. If these were the only parameters to be estimated, they would be exactly identified (through a non-linear transformation) by the 30 moments for inflation portfolios discussed above.
3. **Aggregate market** (3 parameters). The vector of parameters is $\Theta_m = (\mu_m, l_m, \sigma_m)$, which are the mean, leverage parameter and volatility of the market's dividend process.
4. **Stochastic volatility** (6 parameters). The vector of parameters is $\Theta_{vol} = (\sigma_{\pi w}, \sigma_{cw}, \nu_{\pi,s}, \nu_{c,s})$ with $s = 0, 1$. These are the variances and auto-regressive coefficients for volatility.
5. **Preferences** (3 parameters). The three preference parameters $\Theta_u = (\delta, \gamma, \psi)$ are the discount rate, the coefficient of risk aversion and the EIS. The small number of preference parameters and the functional form of the stochastic discount factor consistent with equilibrium is one of the main reasons why the GMM system is overidentified.

4.3 Estimation results

The preference parameters obtained from the calibration are $\gamma = 8.46$ and $\psi = 1.44$. The EIS is the key preference parameter to match the slope of returns with respect to inflation betas. The EIS that I estimate is very close to the one used in the LRR literature, which is generally calibrated to be $\psi = 1.5$. A relatively high level of risk aversion γ contributes to the large inflation premium, but is more important in determining the overall level of returns than their sensitivity to inflation. Because dividends and consumption are exposed to the same underlying volatility shocks, GMM

faces a tradeoff between matching the high returns of stocks and the low volatility of consumption. Increasing gamma reduces this tradeoff. If my model included long-run risk, disaster risk, or some other additional source of consumption risk, it would be able to match the equity premium with a lower γ . Table 8 reports the other parameters.

The main goal of the parameter estimation of the model is to test whether the model can generate an inflation premium consistent with the data using realistic processes for inflation and consumption. Column 4 of Table 4 shows the model's results when performing the same Fama-MacBeth procedure that I used in the data. Comparing to the empirical estimates in columns 1, 2 and 3, we see that the model can reproduce all of the inflation premium. Table 10 goes deeper into the model's predictions for inflation portfolios. The table shows that the model can closely match the individual average betas and returns of the inflation portfolios. Table 9 shows that the source of heterogeneity in the model's betas and returns comes mostly from having a different exposure to consumption and inflation and not from their difference in volatility loadings.

Table 11 shows that the model successfully matches basic moments for inflation and consumption, while Table 12 shows the degree of consumption predictability and inflation persistence in the model. The standard deviation for inflation and consumption are 1.34% and 2.25% in the model and 1.14% and 2.14% in the data. The slightly higher volatility of inflation is important to match the observed price of risk for inflation. The other moments for consumption and inflation are accurately aligned to their empirical counterparts.

Table 11 also shows that the model replicates the Sharpe ratio and the mean and standard deviation of the price-dividend ratio for the aggregate market. The level of the nominal risk-free rate is also closely matched, although its volatility in the model is less than half of what we observe in the data. I find the same pattern for the yield curve: the model produces an upward sloping yield curve, but the volatility of yields is too small and decays with horizon faster than in the data.

Table 14 shows interesting moments that were not targets of my GMM calibration.

The table verifies that consumption is as persistent in the data as in the model. Kojien, Lustig, Van Nieuwerburgh and Verdelhan (2010) emphasize the moments of the wealth-consumption ratio as a means to differentiate between asset pricing models. Table 14 corroborates that my model performs well in this dimension.

Tables 11-14 compare my results to Bansal and Shaliastovich (2010). Their model is a standard long-run risk model with an exogenous process for inflation attached to it. Unlike my model, inflation does not feed back into consumption or any other real variables and acts just as a conversion factor between nominal and real prices. In their model, an inflation premium arises because inflation itself is exposed to consumption and long-run risk shocks.

Before I compare their results to mine, two caveats are in order. First, I estimate my model with GMM, while they resort to calibration to choose parameters. Picking parameters using GMM may give their model a better fit and provide a more even comparison between the two. Second, their calibration is for a slightly different period and done at the quarterly frequency after time-aggregating monthly series from their model.

I also emphasize that Bansal and Shaliastovich's (2010) model was not designed to match the inflation premium in the cross-section of stock returns. It is therefore not surprising to find that the inflation premium in their model is about half the size of the one I find in the data, as can be seen in Table 4. Their model is designed to explain predictability puzzles in bond and currency markets while matching the level and volatility of nominal yields and the market's return. As can be seen from Tables 11-14, they succeed at matching the means and variances of aggregate market returns, the risk-free rate, bond yields of all maturities, consumption and inflation. My model, on the other hand, is designed to explain neither the equity premium puzzle nor the predictability puzzles of bond and currency markets. In this respect, the principal objective of our models is different and they can be regarded as complementary.

Another concern that I address is the possibility that the inflation process has

changed throughout my sample, perhaps due to a change in monetary policy. I re-estimate the model in two separate subsamples, one before and one after 1980. Panel B of Table 2 shows that in both the model and the data, the higher persistency of inflation before 1980 is associated with a higher inflation premium. This is an important validation for the model: the change in key parameters determining the inflation premium in the data and the model imply the same reaction of the inflation premium. This exercise is perhaps the closest we can get to a “natural experiment” in models of this kind.

5 Conclusion

A stock’s inflation risk can be written as the product of the market price of inflation risk and the stock’s quantity of risk. In this paper, I estimate both by using a two-step Fama-MacBeth procedure. Inflation betas measure the quantity of risk. The coefficient in a cross-sectional regression of excess returns on betas measures the market price of risk. I find that stocks whose returns covary negatively with inflation shocks have unconditionally higher returns. This implies that the average market price of risk of inflation shocks is negative: periods with positive inflation shocks tend to be bad states of nature and investors are willing to pay insurance in the form of lower mean returns when holding an inflation-mimicking portfolio. I estimate that holding such a portfolio gives the agent a Sharpe ratio of -0.33.

I argue that the negative price of inflation risk arises because high inflation today predicts low growth in future real consumption. I develop a model that is able to match the observed inflation market price of risk when estimated by GMM to have the same level of consumption predictability and inflation persistence as in the data.

A limitation of the model is that it takes the distribution of betas —the distribution of the quantity of risk— as given. Full understanding of inflation risk in the cross-section of stocks requires also explaining why the quantity of risk varies from firm to firm. There are four classic explanations in the literature: (i) Summers (1981) and

Feldstein (1980) argue that taxes are responsible for firms' inflation risk; (ii) Fama (1981) points out that positive supply shocks increase future expected dividends but lowers the current price level, inducing a spurious correlation between stock returns and inflation; (iii) Cohen, Polk and Voulteenaho (2005), based on Modigliani and Cohn (1979), propose inflation illusion; and (iv) Mundell and Tobin put forward expected inflation and shoe leather costs. Other macroeconomic models, although not specifically designed to address the stock market's heterogeneity in inflation risk, can also provide important insights. For example, the work of Nakamura and Steinsson (2008) implies that firms' inflation risk is heterogeneous due to differences in menu costs and the variance of idiosyncratic productivity shocks. In this paper, I begin the exploration of why firms have different inflation betas and find that a sizable amount of heterogeneity in firm's inflation riskiness does not depend on what industry the firm belongs to, its size, book-to-market or exposure to fluctuations in oil prices. Theories that rely solely on the aforementioned effects will most likely need additional ingredients to provide a comprehensive explanation of inflation risk.

6 References

- Abel, Andrew, 1999 , Risk Premia and Term Premia in General Equilibrium, *Journal of Monetary Economics* 43, 3-33.
- Amisano, Gianni, and Oreste Tristani, 2009, A DSGE model of the term structure with regime shifts, Working Paper.
- Andreasen, Martin Møller, 2009, Explaining Macroeconomic and Term Structure Dynamics Jointly in a Non-Linear DSGE Model, CREATES Working Paper.
- Ang, Andrew, Geert Bekaert, and Min Wei, 2007, Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?, *Journal of Monetary Economics* 54, 1163-1212.
- Ang, Andrew, Sen Dong, and Monika Piazzesi, 2007, No-Arbitrage Taylor Rules, NBER Working Papers 13448.
- Ang, Andrew, and Dennis Kristensen, 2009, Testing Conditional Factor Models, CREATES Research Papers, School of Economics and Management, University of Aarhus.
- Ang, Andrew, Jun Liu, and Krista Schwarz, 2010, Using Individual Stocks or Portfolios in Tests of Factor Models, Working Paper, Columbia University.
- Ang, Andrew, Monika Piazzesi, and Min Wei, 2006, What does the yield curve tell us about GDP growth?, *Journal of Econometrics* 131, 359-403.
- Bansal, Ravi, and Ivan Shaliastovich, 2010, A Long-Run Risks Explanation of Predictability Puzzles in Bond and Currency Markets, Manuscript, Duke University.
- Bansal, Ravi, and Amir Yaron, 2004, Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles, *Journal of Finance* 59, 1481-1509.
- Beeler , Jason, and John Y. Campbell, 2009, The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment, NBER Working Papers 14788.
- Bekaert, Geert, and Xiaozheng Wang, 2010, Inflation Risk and the Inflation Risk Premium, *Economic Policy* 25, 755- 806.

van Binsbergen, Jules H., Jesús Fernández-Villaverde, Ralph S. J. Koijen and Juan F. Rubio-Ramírez, 2008, Working with Epstein-Zin Preferences: Computation and Likelihood Estimation of DSGE Models with Recursive Preferences, Working Paper.

Black, Fischer, Michael C. Jensen, and Myron Scholes, 1972, The Capital Asset Pricing Model: Some Empirical Tests, in *Studies in the Theory of Capital Markets*, Michael C. Jensen, ed, New York: Praeger, 79-121.

Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307-327.

Bollerslev, Tim, Jeffrey R. Russell, and Mark W. Watson, 2009, Volatility and time series econometrics: essays in honor of Robert Engle, *Advanced Texts in Econometrics*. Oxford University Press.

Buraschi, Andrea, and Alexei Jiltsov, 2005, Inflation Risk Premia and the Expectations Hypothesis, *Journal of Financial Economics* 75, 429-490.

Caldara, Dario, Jesús Fernández-Villaverde, Juan F. Rubio-Ramírez and Wen Yao, 2009, Computing DSGE Models with Recursive Preferences, NBER Working Paper Series 15026.

Campbell, John Y., 1991, A Variance Decomposition for Stock Returns, *Economic Journal, Royal Economic Society* 101, 157-179.

Campbell, John Y. and Luis M. Viceira, 2001, Who Should Buy Long-Term Bonds?, *American Economic Review* 91, 99-127.

Campbell, John Y., Adi Sunderam, and Luis M. Viceira, 2010, Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds, NBER Working Papers 14701.

Chen, Nai-Fu, Richard Roll, and Stephen A. Ross, 1986, Economic Forces and the Stock Market, *The Journal of Business* 59, 383-403.

Chernov, Mikhail, and Philippe Mueller, 2007, The Term Structure of Inflation Forecasts, Working Paper.

Clarida, Richard, Jordi Gali, and Mark Gertler, 1999, The Science of Monetary Policy: A New Keynesian Perspective, *Journal of Economic Literature* 37, 1661-1707.

Cochrane, John H. and Monika Piazzesi, 2005, Bond Risk Premia, *American Economic Review* 95, 138-160.

Cohen, Randolph, Christopher Polk, and Tuomo Vuolteenaho, 2005, Money Illusion in the Stock Market: The Modigliani-Cohn Hypothesis, *Quarterly Journal of Economics* 120, 639-668.

D'Amico, Stefania, Don H. Kim, and Min Wei, 2008, Tips from tips: The Information content of Treasury inflation-protected security prices, Finance and Economics Discussion Series Working Paper, Federal Reserve Board.

Darracq Pariès, Matthieu, and Alexis Loublier, 2010, Epstein-Zin Preferences and their Use in Macro-finance Models: Implications for Optimal Monetary Policy, Working Paper Series 1209, European Central Bank.

Duffie, Darrell, and Larry G. Epstein, 1992, Asset Pricing with Stochastic Differential Utility, *Review of Financial Studies* 5, 411-436.

Elton, Edwin J., and Martin J. Gruber, 1995, *Modern Portfolio Theory and Investment Analysis* 5th ed., New York, Wiley & Sons, Inc.

Elton, Edwin J., Martin J. Gruber, and Thomas Ulrich, 1978, Are Betas Best?, *Journal of Finance* 13, 1375-1384.

Fama, Eugene F., 1981, Stock Returns, Real Activity, Inflation, and Money, *American Economic Review* 74, 545-565.

Fama, Eugene F., and Kenneth R. French, 1992, The Cross-Section of Expected Stock Returns, *Journal of Finance* 47, 427-467.

Fama, Eugene F. and Kenneth R. French, 1996, Multifactor Explanations of Asset Pricing Anomalies, *Journal of Finance* 51, 55-84.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, Return, and Equilibrium: Empirical Tests, *Journal of Political Economy* 81, 607-636.

Feldstein, Martin, 1980, Inflation and the Stock Market, *American Economic Review* 70, 839-847.

Fuhrer, Jeffrey, and Gerald Moore, 1995, Inflation Persistence, *Quarterly Journal of Economics* 110, 127-160.

Gabaix, Xavier, 2008, Variable Rare Disasters: A Tractable Theory of Ten Puzzles in Macro-finance, *American Economic Review* 98, 64-67.

Gibbons, Michael R., Stephen A. Ross and Jay Shanken, 1989, A Test of the Efficiency of a Given Portfolio, *Econometrica* 57, 1121-1152.

Grishchenko, Olesya, and Jing-zhi Huang, 2008, Inflation Risk Premium: Evidence from the TIPS Market, Working Paper.

Guvenen, F., 2009, A Parsimonious Macroeconomic Model for Asset Pricing, *Econometrica* 77, 1711-1750.

Koijen, Ralph S.J, Hanno Lustig, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2010, Long Run Risk, the Wealth-Consumption Ratio, and the Temporal Pricing of Risk, *American Economic Review* 100, 552-556.

Levin, Andrew T., J. David López-Salido, Edward Nelson, and Tack Yun, 2008, Macroeconometric equivalence, microeconomic dissonance, and the design of monetary policy, *Journal of Monetary Economics* 55, 48-62.

Lustig, Hanno, Stijn Van Nieuwerburgh, and Adrien Verdelhan, 2008. The Wealth-Consumption Ratio. NBER Working Papers 13896.

Modigliani, Franco, and Richard Cohn, 1979, Inflation, Rational Valuation, and the Market, *Financial Analysts Journal* 35, 24-44.

Nakamura, Emi, and Jón Steinsson, 2008, Monetary Non-Neutrality in a Multi-Sector Menu Cost Model, NBER Working Papers 14001.

Orphanides, Athanasios, 2004, Monetary Policy Rules, Macroeconomic Stability and Inflation: A View from the Trenches, *Journal of Money, Credit and Banking* 36, 151-175.

Parker, Jonathan A., and Christian Julliard, 2005, Consumption Risk and the Cross Section of Expected Returns, *Journal of Political Economy* 113, 185-222.

Piazzesi, Monika, and Martin Schneider, 2007, Equilibrium Yield Curves, *NBER Macroeconomics Annual 2006* 21, 389-472.

Rudebusch, Glenn, and Eric Swanson, 2009, The bond premium in a DSGE model with long-run real and nominal risks, Working Paper, Federal Reserve Bank of San Francisco.

Sentana, Enrique, 2004, Factor Representing Portfolios in Large Asset Markets, *Journal of Econometrics* 119, 257-289.

Stock, James H., and Mark W. Watson, 2005, Implications of Dynamic Factor Models for VAR Analysis, NBER Working Papers 11467.

Summers, Lawrence H., 1981, Inflation and the Valuation of Corporate Equities, NBER Working Papers 0824.

Taylor, John B., 1998, An Historical Analysis of Monetary Policy Rules, NBER Working Paper.

Wachter, Jessica A, 2006, A Consumption-Based Model of the Term Structure of Interest Rates, *Journal of Financial Economics* 79, 365-399.

7 Tables and Figures

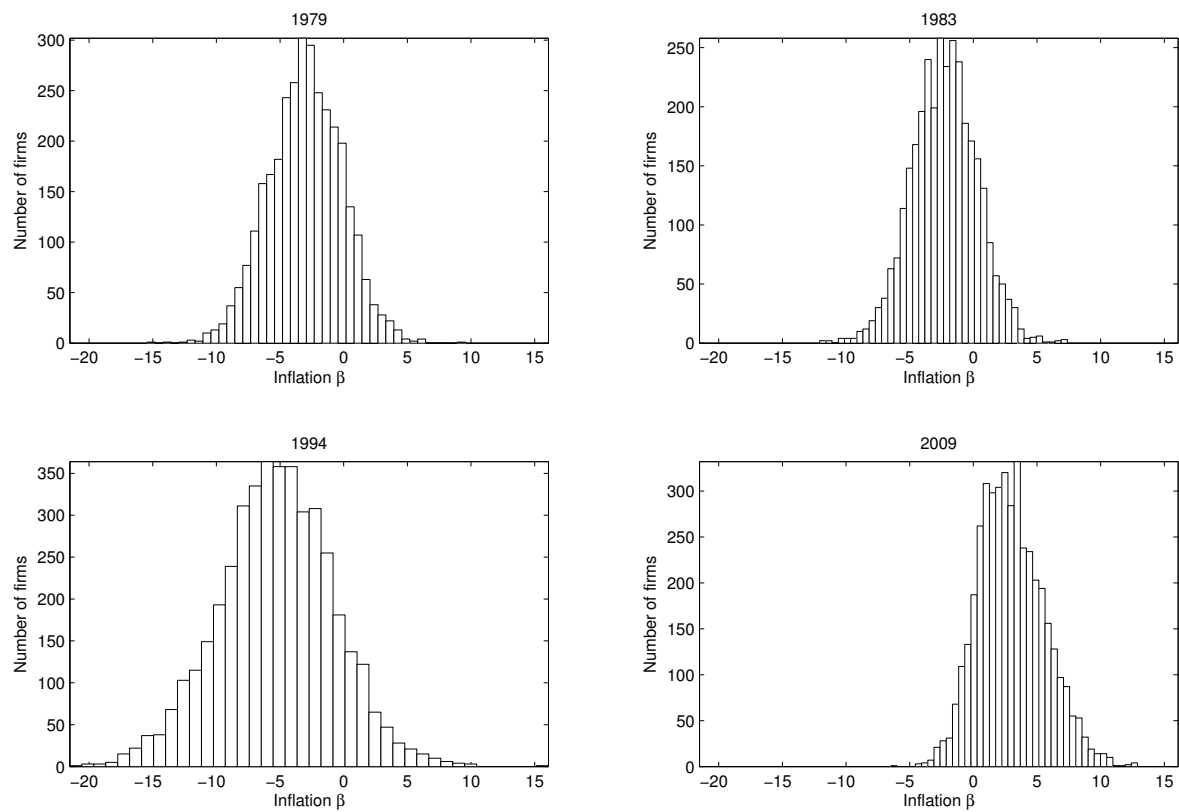


Figure 1: HISTOGRAM OF INFLATION BETAS FOR DIFFERENT TIME PERIODS.




Figure 2: TIME SERIES OF INFLATION BETAS FOR THE AGGREGATE STOCK MARKET, THE FIVE YEAR NOMINAL BOND, AND FOUR REPRESENTATIVE COMPANIES.

TABLE 1: INFLATION-SORTED PORTFOLIOS HAVE RETURNS
WELL ALIGNED WITH INFLATION BETAS

	Ex-ante β	Post-ranking β	$\mathbb{E}[R_t^e]$	$\sigma(R_t)$
<i>Portfolios</i>				
$p = 1$	-5.61	-2.32	6.91	15.8
2	-3.52	-1.71	6.87	13.3
3	-2.56	-1.51	6.44	16.9
4	-1.80	-1.10	6.29	14.7
5	-1.17	-0.874	6.36	14.0
6	-0.555	-0.762	6.17	12.8
7	0.104	-0.597	5.58	14.8
8	0.853	-0.365	5.53	15.0
9	1.86	-0.007	5.56	15.6
$p = 10$	4.10	0.064	5.16	13.4
<i>Spread</i> (1 minus 10)	-9.71	-2.38	1.75	2.40

Notes: To construct portfolios, I first find stock n 's inflation beta at time t , $\beta_{n,t}$, by regressing its excess returns on inflation innovations, only using observations that occurred before t . I give smaller weight to more distant observations by using an exponential kernel with a half-life of five years. I construct 10 inflation portfolios by initially double-sorting stocks on 10 groups according to size (market equity) and 10 groups according to $\beta_{n,t}$, and then averaging across size. The ex-ante betas are the averages across time of $\beta_{p,t}$ for each portfolio p . I find post-ranking betas by freezing portfolio weights at time t and regressing the excess returns of this fixed-weights portfolio on inflation innovations, using the five years of data starting at $t + 1$. The second column shows the average across time of portfolios' post-ranking beta. The third and fourth columns report mean and standard deviation of excess returns in percentage points per year. I use all stocks in CRSP. Observations are monthly from February 1959 to December 2009. Even though returns align well with inflation betas, the spread in returns between portfolios is not statistically significant. Table 4 shows that using the more efficient Fama-MacBeth procedure leads to an inflation premium that is similar in magnitude but also statistically significant.

TABLE 2: RISK EXPOSURES AND INDUSTRY PROPERTIES OF INFLATION PORTFOLIOS

	Risk-factor exposures				Industry properties		
	Market	Size	Book-to-Market	Momentum	Concentration	Correlation	Persistence
<i>Portfolios</i>							
$p = 1$	0.992	0.990	0.568	-0.086	0.142	0.316	0.377
2	0.997	0.979	0.525	-0.101	0.143	0.315	0.377
3	1.00	0.985	0.599	-0.088	0.130	0.315	0.378
4	0.999	0.983	0.581	-0.093	0.118	0.312	0.356
5	0.987	0.964	0.557	-0.085	0.131	0.313	0.371
6	1.00	0.981	0.609	-0.055	0.130	0.316	0.375
7	1.01	0.985	0.568	-0.105	0.119	0.311	0.361
8	1.03	0.995	0.562	-0.116	0.114	0.313	0.367
9	1.03	0.992	0.580	-0.098	0.130	0.307	0.340
$p = 10$	1.01	0.985	0.557	-0.098	0.127	0.319	0.354
<i>Spread</i> (1 minus 10)	-0.018	0.005	0.011	0.012	0.015	-0.003	0.023

Notes: Risk-factor exposures are the coefficients from a regression of excess returns of inflation portfolios on the Fama-French factors. Industry concentration measures how diversified a portfolio is – a value of 0 means very diversified and a value of 1 means that all firms belong to the same industry. Industry correlation is a measure of distance between portfolio p and the remaining 9 portfolios. A value of 0 means that portfolio p does not share any industries with the other portfolios and a value of 1 means that the industry distribution of p is identical to the distribution of all other portfolios combined. Industry persistence is analogous to industry correlation but compares portfolio p to itself 5 years later.

TABLE 3: STANDARD FACTOR MODELS HAVE LARGE ERRORS WHEN
PRICING INFLATION PORTFOLIOS

	$R_{p,t}^e = a_p + b_p X_t + e_{p,t}$						
	Model number						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Factors X_t</i>							
Market	Yes	Yes	Yes	Yes	Yes	×	Yes
HML / SMB / Mom	×	Yes	Yes	×	×	×	Yes
ST rev + LT rev	×	×	Yes	×	×	×	Yes
Oil	×	×	×	Yes	×	×	Yes
CP factor	×	×	×	×	Yes	×	Yes
Industry	×	×	×	×	×	Yes	Yes
<i>Test $H_0 : \text{all } a_p = 0$</i>							
Mean $ a_p $	2.08	2.80	2.68	1.88	3.11	2.31	4.28
p-value	0.005	0.00	0.00	0.027	0.00	0.043	0.00
R^2	58.3%	60.9%	61.2%	58.6%	57.2%	62.3%	63.8%
N	540	540	540	540	528	540	528

Notes: This table reports the results of regressing inflation portfolio's excess returns $R_{p,t}^e$ on asset pricing factors X_t . Monthly observations, ending in December of 2009 and starting depending on availability of factors X_t . The Fama-French factors, short and long-term reversal factors, and industry portfolios are obtained from Professor French's website. Oil returns are from the IMF. The Cochrane-Piazzesi factor is from Professor Piazzesi's website. The mean $|a_p|$ (mean absolute pricing error) is in percentage points per year. The p-values are for the the null hypothesis H_0 that all pricing errors are zero using a Newey-West variance-covariance matrix with 60 lags and the GRS statistic (which is an F-test adjusted for finite sample bias).

TABLE 4: INFLATION PRICE OF RISK IN THE CROSS-SECTION OF STOCKS:
FAMA-MACBETH ESTIMATES

	Results from $R_{p,t}^e = a_t + \lambda_t \beta_{p,t} + \varepsilon_t$				
	10 portfolios	All stocks	Flat kernel	Model	B.S. (2010)
$\bar{\lambda} = \frac{1}{T} \sum \hat{\lambda}_t$	-0.368** (0.024)	-0.340** (0.019)	-0.343** (0.031)	-0.377** (0.033)	-0.184* (0.023)
$\bar{\lambda}/\sigma_\pi$	-0.323	-0.298	-0.300	-0.285	-0.101

Notes: (*,**) Significant at the 5%, 1% level.

The estimates $\hat{\lambda}_t$ are the annualized coefficients of a cross-sectional regression of excess returns at time t on the estimated inflation betas for the same period. The estimate $\bar{\lambda}$ is the average over time of the cross-sectional estimates $\hat{\lambda}_t$. The second row normalizes $\bar{\lambda}$ by the standard deviation of inflation innovations. Column 1 corresponds to my main specification with 10 inflation-sorted portfolios. Column 2 uses individual stocks as advocated in Ang, Liu and Schwarz (2010). Column 3 is the same as column 1 but uses a simple 5-year rolling window to estimate inflation betas. Columns 4 and 5 show the price of inflation risk implied by my model and Bansal and Shaliastovich (2010). The first three columns show in parenthesis Newey-West standard errors with 12 lags and Shanken's adjustment. Column 4 reports GMM asymptotic standard errors. Observations are monthly from February 1959 to December 2009.

TABLE 5: INFLATION IS PERSISTENT

	Regression of π_t on its lags	
π_{t-1}	0.629** (0.031)	0.519** (0.040)
π_{t-2}	×	0.091** (0.046)
π_{t-3}	×	0.107** (0.040)
R^2	39.6%	41.6%
N	611	609

Notes: (**) Significant at the 1% level.

Inflation is seasonally adjusted CPI. Monthly observations, February 1959 to December 2009. OLS standard errors are in parenthesis.

TABLE 6: INFLATION REGIMES
BEFORE AND AFTER 1980

Panel A: Regression of π_t on its lag			
	Full Sample	Pre-1980	Post-1980
π_{t-1}	0.629** (0.031)	0.679** (0.047)	0.547** (0.044)
R^2	39.6%	45.2%	30.0%
N	611	251	360
$\sigma(\pi_t)$	1.14	1.14	1.07
Panel B: Inflation premium $\bar{\lambda}$			
	Full Sample	Pre-1980	Post-1980
Data	-0.368	-0.371	-0.317
Model	-0.377	-0.401	-0.324

Notes: (**) Significant at the 1% level.

Inflation is seasonally adjusted CPI. Monthly observations, February 1959 to December 2009. OLS standard errors in parenthesis.

TABLE 7: INFLATION PREDICTS CONSUMPTION GROWTH

	Regression of consumption growth $\Delta c_{t \rightarrow t+k}$ on lags of inflation and consumption growth					
	$k = 12$			$k = 24$		
<i>Inflation lags</i>						
π_{t-1}	-1.15* (0.507)	-0.440 (0.389)	-0.452 (0.430)	-2.28* (0.934)	-1.41** (0.439)	-1.55** (0.453)
π_{t-2}	×	-0.616* (0.250)	-0.610* (0.249)	×	-0.719 (0.384)	-0.656 (0.353)
π_{t-3}	×	-0.695* (0.298)	-0.701* (0.290)	×	-0.800 (0.506)	-0.842 (0.498)
<i>Consumption lags</i>						
Δc_{t-1}	×	×	-0.044 (0.186)	×	×	-0.398 (0.206)
R^2	6.12%	9.71%	9.73%	8.92%	10.9%	11.4%
N	599	599	599	587	587	587

Notes: (*,**) Significant at the 5%, 1% level.

Monthly observations, February 1959 to December 2009. Inflation is seasonally adjusted CPI. Consumption is non-durables and services components of PCE. Newey-West standard errors with $2k$ lags are in parenthesis.

TABLE 8A: GMM ESTIMATES OF PARAMETERS FOR
PREFERENCES AND CONSUMPTION

Preference parameters

Discount factor	δ	0.989 (0.04)
Elasticity of intertemporal substitution	ψ	1.44 (0.18)
Risk aversion coefficient	γ	8.46 (0.93)

Consumption growth parameters

Mean of consumption growth	μ_c	0.0026 (0.0012)
Consumption loadings on inflation	$\rho_{c,0}$	-0.140 (0.23)
	$\rho_{c,1}$	-0.083 (0.21)
	$\rho_{c,2}$	-0.052 (0.15)
Consumption volatility level	σ_c	0.0065 (0.003)
Consumption volatility persistence	$\nu_{c,0}$	0.79 (0.21)
	$\nu_{c,1}$	0.31 (0.23)
Consumption volatility of volatility	σ_{cw}	1.1×10^{-5} (2.4×10^{-5})

Notes: GMM asymptotic standard errors in parenthesis.

TABLE 8B: GMM ESTIMATES OF PARAMETERS FOR
INFLATION

Inflation parameters

Mean of inflation rate	μ_π	0.0038 (0.00064)
Inflation auto-regressive coefficients	$\rho_{\pi,0}$	0.621 (0.05)
	$\rho_{\pi,1}$	1.21 (0.08)
	$\rho_{\pi,2}$	1.02 (0.14)
Inflation volatility level	σ_π	0.0039 (0.003)
Inflation volatility persistence	$\nu_{\pi,0}$	0.89 (0.21)
	$\nu_{\pi,1}$	0.31 (0.23)
Inflation volatility of volatility	$\sigma_{\pi w}$	3.8×10^{-5} (1.6×10^{-5})
Volatility loading on consumption shocks	$\varphi_{\pi c}$	-0.18 (0.84)

Notes: GMM asymptotic standard errors in parenthesis.

TABLE 9: GMM ESTIMATES OF PARAMETERS FOR INFLATION
PORTFOLIOS AND THE AGGREGATE MARKET

	Mean of dividend growth	Dividend leverage on consumption	Volatility loading of dividend growth
	$\mu_{p,d}$	l_p	φ_{pc}
<i>Portfolios</i>			
$p = 1$	0.0028	-1.12	0.896
2	0.0026	-1.06	0.874
3	0.0022	-1.01	0.830
4	0.0031	-0.997	0.858
5	0.0022	-0.909	0.916
6	0.0026	-0.838	0.860
7	0.0021	-0.944	0.929
8	0.0023	-0.891	0.881
9	0.0026	-0.900	0.862
$p = 10$	0.0028	-0.773	0.812

TABLE 10: ESTIMATES FOR INFLATION PORTFOLIOS

	Data		Model	
	β	$\mathbb{E}[R_t^e]$	β	$\mathbb{E}[R_t^e]$
<i>Portfolios</i>				
$p = 1$	-2.32	6.91	-1.89	6.95
2	-1.71	6.87	-1.74	6.54
3	-1.51	6.44	-1.61	6.33
4	-1.10	6.29	-0.980	6.17
5	-0.874	6.36	-0.892	6.12
6	-0.762	6.17	-0.725	5.97
7	-0.597	5.58	-0.632	5.71
8	-0.365	5.53	-0.438	5.62
9	-0.007	5.56	-0.105	5.45
$p = 10$	0.064	5.16	-0.101	5.28
<i>Spread</i> (1 minus 10)	-2.38	1.67	-1.78	1.83

Notes: The data section reproduces the ex-post betas and returns of Table 1. I compute Columns 3 and 4 using model parameters estimated via GMM.

TABLE 11: MOMENTS FOR INFLATION, CONSUMPTION, THE
AGGREGATE MARKET AND THE RISK FREE RATE

	Data	Model	B.S. (2010)
$\mathbb{E}[\pi_t]$	4.47	4.52	3.30
$\sigma(\pi_t)$	1.14	1.34	1.82
$\mathbb{E}[\Delta c_t]$	3.14	3.14	1.92
$\sigma(\Delta c_t)$	2.14	2.25	1.35
$corr(\pi_t, \Delta c_t)$	-0.26	-0.26	-0.34
$\mathbb{E}[R_t^{m,e}]$	6.65	7.25	5.01
$\sigma(R_t^{m,e})$	15.5	16.8	15.2
$\mathbb{E}[P_t/D_t]$	26.97	25.42	21.71
$\sigma(P_t/D_t)$	7.32	8.32	12.17
$\mathbb{E}[R_t^f]$	1.18	1.26	1.19
$\sigma(R_t^f)$	0.97	0.45	—

Notes: For the data column, I report annualized estimates from monthly observations for February 1959 to December 2009. I compute moments for the model using parameters estimated via GMM for the same period. Column 3 reports the results in Bansal and Shaliastovich (2010).

TABLE 12: ESTIMATES OF CONSUMPTION GROWTH
PREDICTABILITY AND INFLATION PERSISTENCE

Panel A: Regression of consumption growth Δc_t on lags of inflation			
	Data	Model	B.S. (2010)
π_{t-1}	0.02	-0.14	-0.30
π_{t-2}	-0.11	-0.08	-0.08
π_{t-3}	-0.07	-0.05	-0.01
Panel B: Regression of inflation π_t on its lags			
	Data	Model	B.S. (2010)
π_{t-1}	0.52	0.62	0.65
π_{t-2}	0.09	0.12	0.63
π_{t-3}	0.11	0.10	0.52

Notes: I compute Columns 2 using the GMM estimates of my model. Column 3 reports results from Bansal and Shaliastovich's (2010) model.

TABLE 13: ESTIMATES OF THE NOMINAL YIELD CURVE

	Data		Model		B.S. (2010)	
	$\mathbb{E}[y_t^{(n)}]$	$\sigma(y_t^{(n)})$	$\mathbb{E}[y_t^{(n)}]$	$\sigma(y_t^{(n)})$	$\mathbb{E}[y_t^{(n)}]$	$\sigma(y_t^{(n)})$
<i>Bond Maturity</i>						
1 year	6.40	0.86	6.25	0.52	5.60	2.92
2 years	6.63	0.87	6.26	0.38	5.85	2.81
3 years	6.81	0.87	6.27	0.15	6.28	2.71
4 years	6.95	0.87	6.29	0.02	6.82	2.61
5 years	7.03	0.88	6.35	0.00	7.43	2.53

Notes: Bond yield data are from the Fama-Bliss bond files. I compute Column 2 using the GMM estimates of my model. Column 3 reports the results in Bansal and Shaliastovich (2010).

TABLE 14: MOMENTS NOT TARGETED IN GMM
ESTIMATION

	Data	Model	B.S. (2010)
$corr(\Delta c_t, \Delta c_{t-1})$	-0.26	-0.26	0.35
$\mathbb{E}[W_t/C_t]$	88.59	26.42	48.97
$\sigma(W_t/C_t)$	14.11	17.23	12.59

Notes: Columns 1 and 3 for the wealth-consumption ratio W_t/C_t are from Koijen, Lustig, Van Nieuwerburgh and Verdelhan (2010). I compute Column 2 using the GMM estimates of my model.