One Asset Does Not Fit All: Inflation Hedging by Index and Horizon

Stefania D'Amico and Thomas B King*

July 16, 2024

Abstract

We examine the inflation-hedging properties of various financial assets and portfolios by estimating simple time-series models of the joint dynamics of each asset-inflation pair, for multiple inflation indices and at horizons from one month to 30 years. There is no one-size-fits-all approach to inflation hedging: the optimal hedge depends on the particular types of prices that an investor is exposed to and at which horizons. For example, food and energy prices are easy to hedge with commodities and certain stock portfolios, while non-housing service prices and wages are not highly correlated with any financial asset. Inflation-linked bonds and swaps are good hedges for the headline CPI at horizons matching their maturities, but they can perform poorly at other horizons and for other price indices. During the inflationary period of 2021–2023, many historical hedging relationships failed, as monetary policy tightening lagged inflation.

Keywords: Inflation, real assets, Treasury Inflation Protected Securities, hedging

*D'Amico: Federal Reserve Bank of Chicago (corresponding author, stefania.damico@chi.frb.org), King: Federal Reserve Bank of Chicago (thomas.king@chi.frb.org). The analysis in this paper builds on the results discussed in our companion chapter in the *Research Handbook of Financial Markets* (D'Amico and King, 2023). In particular, Appendix A contains the detailed tables referenced in that chapter. For helpful comments, discussions, and suggestions we thank Francois Gourio, Refet Gurkaynak, Ivan Shaliastovich, Shane Sherlund, Jonathan Wright, and participants at the 2024 Conference on Fixed Income Markets and Inflation. Corey Feldman and Andrew Souther provided superb research assistance. The views expressed are not official positions of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1 Introduction

Different economic agents have different beliefs about inflation. For instance, Figure 1 shows that, over the last ten years, consumer expectations have run consistently higher than those of professional forecasters and investment professionals, according to surveys. Figure 2 shows that consumers have also been considerably more uncertain about these forecasts. Even among investors with similar beliefs, inflation exposures may differ. Firms, for example, need to worry about the costs of wage bills and raw inputs; households need to consider large and distant nominal expenditures like the cost of college and retirement.

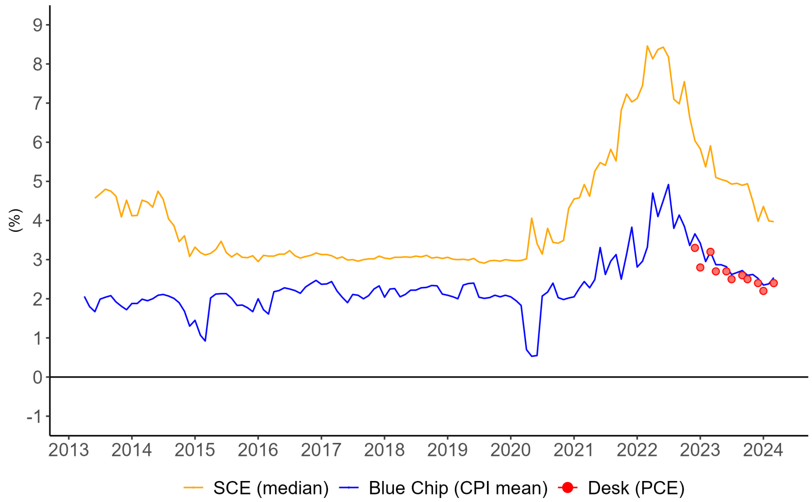

Figure 1: **One-year Inflation Expectations Across Different Types of Agents**. The yellow line shows the time series of the one-year inflation expectations from the Federal Reserve Bank of New York's Survey of Consumer Expectations, the blue line shows the one-year inflation expectations from the Blue Chip Economic Indicators, and the red dots shows the one-year inflation expectations from the Federal Reserve Bank of New York's Survey of Primary Dealers.

These observations suggest that the hedging needs of investors are likely diverse. Previous research suggests that the demand for inflation protection is increasing with respect to the expected inflation rate, and with respect to inflation uncertainty (e.g., Kwak and Lim, 2014). But the *kind* of inflation protection that is desired ought to depend on an investor's particular price exposures and investment horizon. The optimal portfolio for a firm that needs to hedge year-to-year wage inflation is much different from that of a retirement-planning household looking to insure its ability to afford consumer prices over a long horizon.

Ultimately, as stressed in D'Amico and King (2023), for each economic agent, the inflation risk originates from the gap between the inflation rates of assets and liabilities. If these inflation rates are expected to be similar, then the real value of future net worth should be protected. This implies that agents, to properly hedge inflation risk, have to pay attention to and form beliefs on different inflation measures and how these correlate with different financial instruments.

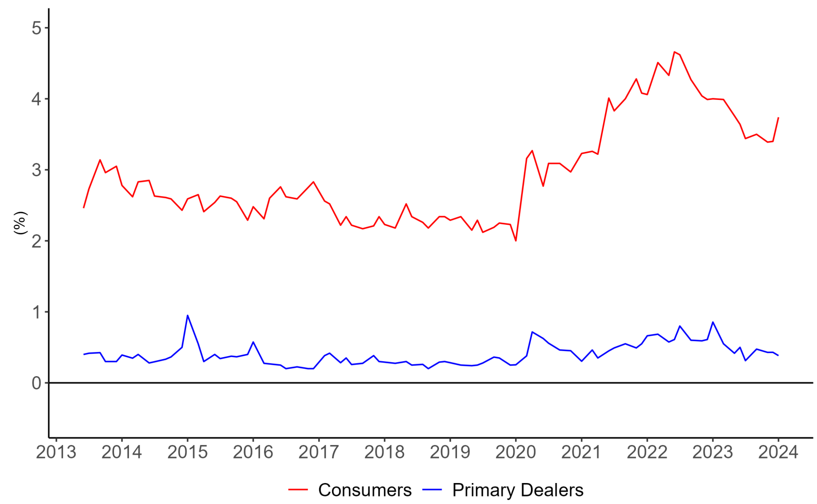

Figure 2: **One-Year Inflation Uncertainty Across Different Types of Agents**. The orange line shows the median interquartile range of the one-year inflation distributions from the Federal Reserve Bank of New York's Survey of Consumer Expectations; the blue line shows the median interquartile range from the Federal Reserve Bank of New York's Survey of Primary Dealers.

Motivated by the different needs of diverse types of investors, this paper explores how the returns on individual assets and portfolios are correlated with multiple inflation outcomes. We investigate both a range of different inflation *indices* and a spectrum of investment *horizons*. At any given horizon, for example, a particular asset may be highly correlated with headline PCE inflation, but not with wages. Similarly, it may be correlated with a given inflation measure at an annual frequency but only weakly correlated over a ten-year horizon, or vice versa. Our analysis looks exhaustively across *asset-index-horizon* triplets to uncover simple inflation-hedging strategies that can be customized to the specific requirements of individual investors.

While we consider a broad set of different assets, we pay particular attention to products whose payoffs are contractually linked to consumer-price indices, specifically, inflation swaps and Treasury Inflation-Protected Securities (TIPS). While the availability of such products in recent decades opened up new possibilities for obtaining inflation protection, we note several reasons that strategies based on these instruments may be less effective than hoped. First, the prices of these instruments are subject to deviations from fundamentals because of various market frictions, and trading in inflation-linked derivatives such as inflation swaps and TIPS asset swaps is generally unavailable to retail investors. Second, although the payoffs on TIPS and inflation swaps move nearly one-for-one with headline CPI inflation if held to maturity, they may differ significantly at shorter investment horizons or if they are rolled over after maturing. This is because real yields and inflation are dynamically correlated with each other as a result of monetary policy.

Finally, and most importantly, headline CPI is only weakly correlated with many other inflation measures, and these correlations also vary across investment horizons. For example, during the recent inflationary wave, there were large discrepancies between different measures of inflation—for example, headline versus core inflation, the CPI versus the PCE, and wages versus consumer prices. The prices of housing, non-housing services, and goods all experienced their own, separate dynamics, as monetary policymakers noted.[1] Thus, an investor wanting to hedge these kinds of inflation exposure may be better off using instruments other than CPI-linked products or, in some cases, may have no recourse at all. Indeed, we show that non-housing service prices and wages are not highly correlated with the returns on any financial asset.

If we were only interested in short-run comovements between various assets and inflation measures, we could simply examine their sample correlations in the data. At longer horizons, however, this approach becomes infeasible. Even using data going back to the 1970s (as far as most of our financial series extend), we have only five non-overlapping observations at the ten-year horizon, so raw sample correlations are only weakly identified. To overcome this problem, we estimate VARIMA models for the joint dynamics of each asset-price/price-index pair, with the order of each model chosen optimally to fit the data. Then, we use these models to project the correlations of asset returns and inflation at each horizon in a consistent manner across horizons, and in a way that exploits the full information contained in the data. We extend the analysis to consider "portfolios" of two assets, estimating a separate tri-variate VARIMA for each portfolio and inflation index.

We find that commodities are generally successful in hedging headline CPI, but

[1]See, for instance, Chair Powell's speech, "Inflation and the Labor Market," on November 30, 2022.

this mostly seems to reflect their significantly positive correlation with energy prices. Related assets, such as the stocks of oil-and-gas and metals-mining firms and some emerging-market (EM) currencies, share this property. Hedging core inflation is harder. At horizons of less than a year, there is little protection available, except for inflation swaps and TIPS. At longer horizons, short-term nominal bonds, real estate, and inflation swaps provide decent, but still imperfect, hedges. Certain stock-market strategies can also perform reasonably well, but here one has to be careful because outcomes vary substantially across different types of stocks. There are also some significant differences in how stock and bond strategies perform against core CPI vs. core PCE. For instance, post 1999, 2-year nominal bonds are a good hedge for core CPI at longer horizons, while certain stock-market sectors such as oil and gas are a good hedge for core PCE. We find that core producer prices (PPI) and wages are the most difficult types of inflation to hedge. While inflation swaps and real estate provide some protection for PPI core, short-term nominal bonds provide a decent amount of protection for wages inflation.

Most of consumer-price inflation stems from three sources: real estate costs, the passthrough of materials and energy prices to consumer goods, and the passthrough of labor costs to goods and services. Materials and energy prices can be hedged very effectively with commodities. Real estate can now be effectively hedged through Real Estate Investment Trusts (REITs) and other instruments that provide broad exposure to this sector. While the literature does not typically focus on hedging labor costs, our empirical exercise identifies short-term nominal bonds and real estate as reasonably good hedges in this dimension. And, since wages are a particularly large component of the cost of non-housing services, the same instruments do a decent job hedging service inflation. Overall, we find that many types of inflation are well hedged by rolling over short-term nominal Treasury bills, since the rates on bills usually adjust rapidly to the inflationary environment.

A separate question—perhaps more relevant for practicing investors—is how various assets and portfolios perform out of sample. Here, unfortunately, our results are not encouraging. At least during the out-of-sample period we look at, September 2021-August 2023, many of the hedging relations that had prevailed since at least 1972 broke down. Most of the individual assets and portfolios that we identified as moving closely with the various inflation indices in-sample fell short of both their own predicted values and the realized levels of inflation during this period. Most notably, short-term interest rates did not adhere to their usual pattern of closely tracking broad measures of inflation, as monetary policy tightening significantly lagged the increase in consumer prices. While some other assets, such as oil, did move higher as inflation rose, they did so by less than historical relationships would have suggested. These results highlight the dependence of assets' inflation-hedging ability on the type of shocks hitting the economy and the monetary-policy response, further reinforcing the difficulties of hedging any measure of inflation reliably.

Regarding the relation of this study to the literature, it should be noted that Parikh et al. (2019) already stress the idea that different investors might have different inflation-hedging needs. Kat and Oomen (2006), in analyzing the hedging properties of commodity futures, focus on three different types of inflation: CPI, PPI, and the employment cost index (ECI). And, in the case of CPI and PPI, they consider both headline and core. Further, Fang, Liu, and Roussanov (2021) also find that commonly used "real assets" such as stocks, commodity futures, and real estate are generally good hedges for energy inflation but not for core inflation. This is in line with Ajello et al. (2019), who show that there are different risks in core, food, and energy prices. Our discouraging out-of-sample findings are reminiscent of Ang et al. (2012), who show that trying to forecast inflation exposure at the individual stock level is difficult, as co-movements with inflation exhibit pronounced time variation, including a change in sign post 2008; they note that this makes it hard to construct portfolios of stocks that are good out-of-sample inflation hedges. Finally, the analysis in this paper extends the results discussed in our companion chapter in the *Research Handbook of Financial Markets* (D'Amico and King, 2023) by introducing additional assets and porfolios, further analyzing inflation-linked products, and distinguishing between the expected and unexpected components inflation. In addition, the out-of-sample tests in this paper are novel.

The remainder of the paper is organized as follows. Section 2 described the VARIMA models and the data. Section 3 discusses our in-sample results. Section 4 summarizes the out-of-sample results. Section 5 considers conditional correlations between unexpected inflation and buy-and-hold bond portfolios. Section 6 describes hedging strategies involving actual and hypothetical inflation-linked derivatives. Section 7 concludes. An extensive set of online appendix tables presents the full set of in- and out-of-sample performance results for all inflation-asset-horizon combinations and portfolios.

2 Empirical Approach

2.1 Model specification

2.1.1 Individual assets

Our measure of a given asset's inflation-hedging ability is its nominal price's simple (Pearson) correlation with each type of inflation (abstracting from transactions costs). Computing these correlations at short horizons would be straightforward. However, as noted above, at investment horizons of more than a few years the raw data are not sufficient to estimate correlations precisely. To overcome this problem, we estimate time-series models of each price index/asset price pair, and we use these models to project the correlations at different horizons. Specifically, for each inflation rate π^i and nominal asset return r^j we estimate a VARIMA(p,1,q) model of the form:

$$\begin{pmatrix} \pi_t^i \\ r_t^j \end{pmatrix} = \mathbf{a}^{i,j} + \sum_{k=1}^{p^{i,j}} \mathbf{A}_k^{i,j} \begin{pmatrix} \pi_{t-k}^i \\ r_{t-k}^j \end{pmatrix} + \sum_{k=0}^{q^{i,j}} \mathbf{B}_k^{i,j} \begin{pmatrix} e_{t-k}^i \\ e_{t-k}^j \end{pmatrix} \tag{1}$$

where $(e_t^i \quad e_t^j)' \sim N(\mathbf{0}, \mathbf{I})$ are iid error terms. To determine the lag orders, we search across all possible models in the range $p^{i,j} \in [0, 12]$ and $q^{i,j} \in [0, 3]$ for monthly data and $p^{i,j} \in [0, 4]$ and $q^{i,j} \in [0, 1]$ for the few quarterly series we consider. In each case, we select the best model using the AIC. Having chosen and estimated the best model for each asset-inflation pair, we simulate 1 million observations, and compute the correlations between log-levels at horizons of 1 month to 30 years.

Our use of VARIMA models, as opposed to simpler VAR specifications, reflect the stylized fact that there is a great deal of short-term noise in some measures of inflation, particularly those embedding food and energy prices. A low-order VAR is not likely to be able to fully distinguish between such movements and those of lower frequencies. Of course, a sufficiently high-order VAR can approximate any VARIMA process, but such models run the risk of over-parameterization. The VARIMA framework, though somewhat computationally intensive, provides a parsimonious way of allowing rich dynamics at both high and low frequencies. (Of course, the data are allowed to choose an MA order of $q = 0$ if they prefer, reducing the model to a VAR.)

One concern with this approach is that different models may in principle imply different conditional and unconditional distributions of inflation. In this case, our inflation projections would not be consistent across all assets. Another way of stating this concern is that some models might be significantly mis-specified, in the sense that they leave out information that is relevant for forecasting inflation. The resulting measurement error would appear as stochastic variation in inflation, which would affect our calculated correlations (presumably biasing them toward zero). To address this issue, we also ran versions of (1) in which we included ten-year Treasury yields as an additional regressor, as a control. These yields embed information about future inflation, and should help to refine the forecast and improve consistency across models. With some minor exceptions, however, this adjustment did not have material effects on our results.

2.1.2 Inflation-optimized portfolios

As discussed more fully below, several of the "assets" we consider in this exercise consist of indices or factors based on large groupings of underlying individual assets (e.g., commodity and equity indices, Fama-French factors). Thus, to an extent, the bivariate models in (1), already reflect potential portfolio strategies. However, such portfolios are not optimized for hedging inflation. To explore this problem, we also calculate returns on portfolios constructed from our initial set of assets (some of which themselves are portfolios) and compare them to inflation outcomes across different indices. This involves a two-stage process. In the first stage, we run three-variable VARIMA models of the form

$$\begin{pmatrix} \pi_t^i \\ r_t^j \\ r_t^l \end{pmatrix} = \mathbf{a}^{i,j,l} + \sum_{k=1}^{p^{i,j,l}} \mathbf{A}_k^{i,j,l} \begin{pmatrix} \pi_{t-k}^i \\ r_{t-k}^j \\ r_{t-k}^l \end{pmatrix} + \sum_{k=0}^{q^{i,j,l}} \mathbf{B}_k^{i,j,l} \begin{pmatrix} e_{t-l}^i \\ e_{t-l}^j \\ e_{t-l}^k \end{pmatrix} \tag{2}$$

where, as before, the lag structure is selected according to the AIC. In the second stage, we simulate the model for 1 million periods, and, at each horizon H, we regress the accumulated values of $\hat{\pi}_t^i$ on the accumulated values of $\hat{r}_t^j$ and $\hat{r}_t^l$:

$$\sum_{h=1}^{H} \hat{\pi}_{\tau+h}^i = w_i^j \sum_{h=1}^{H} \hat{r}_{\tau+h}^j + w_i^l \sum_{h=1}^{H} \hat{r}_{\tau+h}^l + \epsilon_\tau^{i,j,l,H} \tag{3}$$

where "hats" denote simulated data values and $\tau = \{1, ..., 1,000,000\}$ indexes the sequence of simulations. Equation (3) may be viewed as constructing a simple bivariate mimicking portfolio for each inflation index from each pair of assets at each horizon. The coefficients w_i^j and w_i^j estimate the optimal portfolio weights for a given pair of assets. We sort the portfolios of assets by their R^2 in these regressions, for each inflation measure at each horizon (equivalent to minimizing var $\left[\epsilon_\tau^{i,j,l,H}\right]$ over j and l), to evaluate which pairs of assets do the best job of hedging inflation in-sample.

2.1.3 Held-to-maturity bonds

The treatment of equities, commodities, and currencies in these exercises is relatively straightforward, but some difficulty arises when considering hedging using bonds in this context. A strategy of holding a bond for its entire life is very different from a strategy of continuously rolling over to maintain a bond portfolio of constant maturity. Consider a zero-coupon nominal bond of maturity m and an investment horizon $h \leq m$.[2] Denote the initial (time-t) yield on the nominal bond by $y_{t,m}^N$, and define the bond's return as the change in its log price, recalling that the price of a bond is just $\exp[-my_{t,m}^N]$. Then, the total return on the nominal bond over the investment horizon is:

$$r_{t,t+h}^{N(m)} = m\left(y_{t,m}^N - y_{t+h,m-h}^N\right) + hy_{t+h,m-h}^N. \tag{4}$$

When it comes to TIPS, the nominal return is:

$$r_{t,t+h}^{R(m)} = m\left(y_{t,m}^R - y_{t+h,m-h}^R\right) + hy_{t+h,m-h}^R + \pi_{t,t+h}^{CPI} \tag{5}$$

where $\pi_{t,t+h}^{CPI} = \log \frac{CPI_{t+h}}{CPI_t}$ is the log change in the headline CPI between periods t and $t+h$ and $y_{t,m}^R$ is the initial real yield on the TIPS. Thus, to calculate the return on a holding strategy for both types of bonds, one needs to know the m-period yields at the beginning of the investment and the $(m-h)$-period yields at the end. In the case of TIPS, one also needs to know the intervening rate of CPI inflation.[3]

With this in mind, for nominal bonds, we modify our strategy by extending the considered VARIMA models to include three variables rather than two: a price index,

[2]If $m > h$ some rollover will be necessary. We sidestep this more-complicated case.

[3]The calculation for TIPS abstracts from their embedded deflation floor and indexation lag. Adding these features will generally make the hedging performance worse, although their impact is typically small (See D'Amico et al., 2018).

an m-maturity bond yield, and an $(m-h)$-maturity bond yield:

$$\begin{pmatrix} \pi_t^i \\ y_{t,m}^N \\ y_{t,m-h}^N \end{pmatrix} = \mathbf{a}^{i,m,h} + \sum_{k=1}^{p^{i,m,h}} \mathbf{A}_k^{i,m,h} \begin{pmatrix} \pi_{t-k}^i \\ y_{t-k,m}^N \\ y_{t-k,m-h}^N \end{pmatrix} + \sum_{k=0}^{q^{i,m,h}} \mathbf{B}_k^{i,m,h} \begin{pmatrix} e_{t-k}^i \\ e_{t-k}^m \\ e_{t-k}^{m-h} \end{pmatrix} \quad (6)$$

Then, in our simulations, we calculate the returns on the bonds using equation (4) at each horizon and compute the correlations with the simulated inflation series over the same horizon. For TIPS, we include the headline CPI as a fourth variable in the model:

$$\begin{pmatrix} \pi_t^{CPI} \\ \pi_t^i \\ y_{t,m}^R \\ y_{t,m-h}^R \end{pmatrix} = \mathbf{a}^{i,m,h} + \sum_{k=1}^{p^{i,m,h}} \mathbf{A}_k^{i,m,h} \begin{pmatrix} \pi_{t-k}^{CPI} \\ \pi_{t-k}^i \\ y_{t-k,m}^R \\ y_{t-k,m-h}^R \end{pmatrix} + \sum_{k=0}^{q^{i,m,h}} \mathbf{B}_k^{i,m,h} \begin{pmatrix} e_{t-k}^{CPI} \\ e_{t-k}^i \\ e_{t-k}^m \\ e_{t-k}^{m-h} \end{pmatrix} \quad (7)$$

Then, in our simulations, we calculate the returns on the bonds using equation (5) at each horizon and compute the correlations with the simulated inflation series over the same horizon.

2.1.4 Inflation swaps and breakevens

Finally, we consider the hedging properties of inflation swaps and hypothetical "breakeven" positions constructed using long positions in TIPS and maturity-matched short positions in nominal bonds. In the former case, we run

$$\begin{pmatrix} \pi_t^{CPI} \\ \pi_t^i \\ s_{t,m} \end{pmatrix} = \mathbf{a}^{i,m,h} + \sum_{k=1}^{p^{i,m,h}} \mathbf{A}_k^{i,m,h} \begin{pmatrix} \pi_{t-k}^{CPI} \\ \pi_{t-k}^i \\ s_{t-k,m} \end{pmatrix} + \sum_{k=0}^{q^{i,m,h}} \mathbf{B}_k^{i,m,h} \begin{pmatrix} e_{t-k}^{CPI} \\ e_{t-k}^i \\ e_{t-k}^m \end{pmatrix} \quad (8)$$

where $s_{t,m}$ is the m-maturity inflation-swap rate. As in equation (5), the headline CPI is included here because inflation-swaps are explicitly indexed to that series. We run similar models with breakeven rates $(y_{t,m}^N - y_{t,m}^R)$ in place of $s_{t,\tau}$. These concepts, and the appropriate return calculations are discussed more fully in Section 5. Because of the complexity of these strategies, we limit our attention to positions that are held to maturity.

2.2 Data

We begin our baseline sample in 1999. The focus on this relatively short period allows us to bring in TIPS and many other inflation and asset-return series that cannot be considered in longer samples. It is also valuable to focus on relatively recent data, since there is evidence that correlations have shifted over time. We also consider a second sample beginning in 1972 using a subset of the series that are available since then. For our main set of results, we end both samples in August 2021. This gives us space to conduct our out-of-sample test over the post-COVID inflationary period September 2021-August 2023.

The inflation measures we consider are listed in Table 1, while the financial assets are listed in Table 2. Note that we include commodity prices in both categories, since, depending on the investor, a particular commodity may either represent a cost that needs to be hedged or a potential hedging instrument. (See Table B1 in the appendix for a detailed list of the data used and related sources.)

Table 1: Measures of Inflation
CPI: Headline; Core; Energy; Services; Durables; Nondurables
PCE: Headline and Core
PPI: Finished Goods and Finished Goods Core
Wage Inflation: Hourly Earnings
Broad Commodity Indices: BCOM and GSCI

Table 2: Asset price indices
Equities indices: Wilshire 5000; S&P 500
Equity portfolios: S&P industry sub-indices; Fama-French 5 factors
Bonds: GI EM-bond index; nominal Treasury securities and TIPS at various maturities
Commodity indices: BCOM index; GS index and sub-indices
Commodities: gold, silver, wheat, soybeans, hogs; WTI and Brent oil spot and futures
Real estate: Case-Shiller; Wilshire REIT ETF
Currencies: dollar vs. yen, euro, rand, ruble, real
Derivatives: Inflation swaps and hypothetical breakevens at various maturities

3 In-Sample Results

Here we summarize our most important in-sample results, starting with the post-1999 sample (Tables 1A through 3A in the Appendix) and then turning briefly to the post-1972 sample (Tables 4A and 5A in the Appendix).

3.1 Post-1999 Sample

3.1.1 Headline inflation

In general, we find that hedging food and energy inflation is relatively easy because these prices are closely linked to commodities. In particular, oil and natural gas spot and future prices displaying some of the highest correlations (between 70% and 95% beyond the one-year horizon). Because variation in food and energy prices makes up most of the variation in headline inflation indices, this also means that commodities are generally a good hedge for headline inflation. The broad commodity indices and oil futures have correlations as high as 70% at the 6-month horizon and beyond with headline CPI, PCE, and PPI. Contrary to conventional wisdom, we do not find any ability of gold to hedge headline inflation over the post-1999 sample.

Broad stock indices can also provide a good hedge for headline inflation, although much of the correlation is driven by energy-related stocks.[4] Currencies of emerging markets (EM) that are commodities exporters also provide some protection (i.e., the ruble displays correlations of about 50%). This is because those currencies typically gain when there is a rise in the prices of commodities exported by the country.[5] Finally, real estate, as proxied by the Case-Shiller index, does a surprisingly good job of hedging consumer prices at longer horizons (i.e., correlations range between 60% and 75%), while the Wilshire REIT index effectively hedges the PPI (i.e., correlations range between 50% and 60% beyond the one-year horizon).[6]

[4]The strong correspondence between energy-sector stocks and headline CPI is consistent with Ang et al. (2012) and Parikh et al. (2019). However, those studies also find that technology stocks are important. In contrast, we find only weak correlations between headline inflation and the semiconductor and telecommunications sectors and significant *negative* correlations with the software sector at longer horizons.

[5]We note that at very long horizons, the hedging ability of broad stock portfolios and currencies becomes substantially weaker when we include the ten-year Treasury yield as an additional control in the models.

[6]We caution that the results involving the Case-Shiller index in this sample are strongly influenced by the run-up to the 2008 housing crisis and its aftermath.

Generally speaking, returns on longer-term nominal bonds are negatively correlated with headline inflation, so that their inclusion in an investor's portfolio increases the exposure to inflation risk. This is true whether the bonds are held to maturity, continually rolled over to constant maturities, or held for an intermediate period.[7] On the other hand, a strategy of holding one-month Treasury bills (T-bills) provides relatively good protection against headline inflation at medium and longer horizons (correlations of roughly 50%). This is because these rates rise when monetary policy tightens in response to higher inflation.

Finally, short- and medium-term TIPS have performed reasonably well at protecting against headline CPI over short investment horizons (from 1 month to 1 year). At similar horizons, the 10-year TIPS has also provided some protection against headline PCE and PPI. Beyond those horizons, the correlations are rather weak. It may be counterintuitive that TIPS do not hedge inflation particularly well, even headline CPI inflation, to which they are indexed, and even at horizons that match their maturities. The reason is that we are considering *unconditional* correlations. In the data, it tends to be the case that periods of high inflation are preceded by low yields on TIPS. Thus, although they receive perfect CPI protection over their lifetimes, the overall return on a TIPS-based strategy is not particularly high in a high-inflation environment. In Section 5, we consider *conditional* correlations for comparison, and these are considerably higher, as one would expect.

3.1.2 Consumer and producer inflation components

While there are multiple attractive strategies for hedging non-core inflation, the prospects are somewhat dimmer when it comes to core. At *horizons of less than a year*, few of the assets we consider provide good protection. (Of course, one may wonder whether this is a very serious problem, given that core inflation displays very little variation at short horizons.) One exception is TIPS. The 10-year TIPS returns from the 3-month to the 1-year horizon have correlations of about 0.3-0.4 with core CPI, while the 2-year TIPS has a similar size correlation at the 1-month horizon. Correlations with core PCE are also positive but a bit smaller. The correlations with core PPI are mostly negative.

At *longer horizons*, there are substantial differences across the three core indices.

[7] In principle, a strategy of shorting nominal bonds provides protection against headline inflation, although such strategies can be costly to implement.

Core PCE behaves somewhat similarly to headline PCE. It is correlated with the broad commodity indices and oil futures, certain stock-market sectors including oil and gas, and real estate. However, in all of these cases the correlations are at most around 60%, which is significantly lower than the best-performing assets for headline inflation. For core CPI, on the other hand, the only assets that provide some hedging value are the 2- and 5-year nominal bonds rolling returns, some of the Fama-French factors, and the Case-Schiller price index. It is perhaps unsurprising that house prices perform well with core CPI as housing services constitute a large percentage of the core basket. Unlike with headline CPI, broad commodity indices and oil futures do not perform well with core CPI. Gold and most stock-market indices are significantly *negatively* correlated with core CPI.[8] Interestingly, although the overall stock market is negatively exposed to core inflation, the "robust-minus-weak" and "conservative-minus-aggressive" Fama-French factors seem to provide good protection at longer horizons, perhaps suggesting that profitable and conservative firms are more resilient to inflation.[9] To hedge core PPI is even harder. Only the 30-year nominal bond rolling returns, the "robust-minus-weak" Fama-French factor, and the Case-Shiller price index offer some protection (i.e., correlations vary between 20% and 40%).

The nondurable (ND) components of the CPI are dominated by food and energy, so their results are similar to the headline CPI correlations discussed above. Broad commodity indices and oil futures provide an almost perfect hedge against ND CPI at the 1-year horizon and beyond. Moreover, energy prices have high passthrough to the cost of durable goods, so CPI durables are also highly correlated with broad commodity indices and oil futures. Stock market sectors such as metals-mining, financials, and insurance also perform well with durables. In contrast, very few assets provide a decent hedge for CPI Services. Only the 2-year nominal bond rolling returns, T-Bills, and the Case-Shiller price index display positive correlation larger than 30%. (At very long horizons, the S&P Oil and Gas Exploration and Production sector returns and 2-year future on WTI also have correlations around 30%.) Thus, the weak correlations noted for core CPI inflation stems from the lack of a good hedge for the service sector inflation. Further, also TIPS can provide a decent protection against CPI services. In particular, the 2-year TIPS works well at the 3-month horizon, the 5-year TIPS at the 6-month horizon, and the 10-year TIPS for horizons longer than 6 months.

[8] The opposite signs on the correlation of the stock market with headline and core inflation are roughly consistent with the findings of Fang et al. (2022).

[9] However, this result is not robust to the longer sample discussed below.

Not surprisingly, strategies of holding nominal bonds for long periods, including to maturity, almost always generate additional exposure to inflation, rather than providing a hedge. Across bond maturities and holding periods, the correlations with various components of consumer and producer inflation range from slightly positive to -40%. The exception is CPI services, where we find that ten-year bonds held to maturity have a positive 50% correlation.

3.1.3 Wages and house prices

One reason that hedging prices in the service sector is difficult seems to be that labor costs constitute a large fraction of service prices. We find few good hedges for wage inflation. Indeed, most of the asset returns we consider display a negative correlation with average hourly earnings. And, since these are also generally small in magnitude, even potential short positions would not be successful in hedging wages. The main exception is rolling returns in shorter-term nominal Treasuries. T-Bills and 2-year bonds have correlations with average hourly earnings of more than 50% at horizons beyond one year, and ten-year bonds have a 34% correlation if held to maturity. Further, also TIPS provide some protection, with the 2-year TIPS hedging relatively well at the 3-month horizon, the 5-year TIPS at the 6-month horizon, and the 10-year TIPS at the 10-year horizon.[10]

Finally, although we have discussed real estate as a potential hedging instrument (and, indeed, we have shown that it performs well as a hedge in many cases), one might also want to hedge real-estate prices themselves. Our results using the Case-Shiller index show that there are a variety of ways of doing this successfully at horizons of one year and longer. Most components of the stock market, including the "small-minus-big" and "high-minus-low" Fama-French factors, as well as most commodities and some currencies are strongly correlated with house prices at these horizons. In contrast, rolling returns on longer-term bonds, display strong negative correlations with real-estate prices. (Of course, all of these results apply to national house prices at an aggregate level; since local housing markets are very idiosyncratic, no hedging strategy based on financial instruments is likely to offer much protection for an individual homeowner.)

[10] The Case-Shiller index also appears to provide a good hedge for wages at long horizons, but this result may be spurious since it does not hold for our other measure of real estate prices (the Wilshire REIT) and there is no obvious economic reason that it should be true.

3.2 Post-1972 sample

To examine the stability of the above results over time, we re-estimate the models using data since 1972, where possible. Before describing the results obtained in this sample period, it is worth noting that measures of headline and core inflation are considerably more highly correlated than in the 1999 sample. For example, at the five-year horizon, core and headline CPI have a correlation of 0.97 in the post-1972 sample, compared to just 0.55 in the post-1999 sample. This means that they are more likely to be hedged well by the same set of instruments. Most notably, commodities—in particular oil—do a better job of hedging core inflation over the longer sample period. This is consistent with oil's large role in driving business-cycle fluctuations throughout the 1970s and is quite different from what is emphasized in Fang et al. (2022).

Gold and silver perform quite well against most inflation measures in the post-1972 sample, suggesting that those commodities' inflation-hedging abilities must have been particularly good in the pre-1999 period. This behavior may have contributed to the common perception that precious metals are robust inflation hedges. However, as we showed above, that property seems to have disappeared over the last 20 years. In contrast, the yen and the rand—two currencies that we can track—do worse (against both headline or core) over the longer sample period. We continue to find relatively few possibilities for hedging wage inflation in the longer sample. Unlike in the post-1999 sample, average hourly earnings are positively correlated with commodity prices, but, except at very long horizons, those correlations are still quite modest.

Nominal bonds from 1972 display very low or negative correlations with most measures of inflation. However, over the longest holding period (10 years), the 10-year nominal bond is a reasonably good hedge against CPI Services (40%) and Non Durable (58%), PCE headline (43%) and Core (31%), as well as PPI Final Goods (45%).

4 Out-of-Sample Results

Tables 6A and 7A in the Appendix report the results of our out-of-sample tests during the inflationary period following the COVID-19 crisis. For the purposes of this exercise, we rely on the sample that begins in 1972, since this sample includes inflation readings that are similar to those experienced in 2021-23, although this limits the range of assets we can consider.

For single assets, the out-of-sample test is conducted by using the VARIMA simulated data to run a regression of cumulative 2-year changes in monthly returns on 2-year changes in inflation, for each asset/inflation-index pair:

$$\sum_{h=1}^{24} \hat{r}^{j}_{\tau+h} = \alpha^{i}_{j} + \beta^{i}_{j} \sum_{h=1}^{24} \hat{\pi}^{i}_{\tau+h} + u^{i,j}_{\tau} \tag{9}$$

The coefficients of this regression β^i_j are then used to obtain the returns predicted by the inflation that occurred, in each index, in the out-of-sample period of September 2021–August 2023. The predicted returns are then compared to the realized returns over the inflationary period, which allows us to understand whether each asset fell short of its own predicted value, and to the realized level of inflation, to understand whether the asset returns at least protected from inflation despite performing worse or better than expected.

4.1 Post-COVID results for individual assets

Table 3 reports some highlights of the results of this exercise. Specifically, we focus first on the best performing single-asset "portfolios" over an investment horizon of two years. By comparing the first and second columns, it can be noted that, based on the in-sample results, the best inflation-hedging security for almost all measures of inflation is expected to be the risk-free one, that is, the one-month T-Bill on a rolling basis. And indeed, the expected returns from this single-asset strategy are quite close to the inflation realized from September 2021 to August 2023, as shown in the 5th and 7th columns. However, the actual returns reported in the 4th column are very different from the expected returns, in most cases falling short by 4 to 6 percentage points, or more than one standard error (shown in the last column).

The poor out-of-sample performance of the one-month T-Bill indicates that the post-COVID inflationary period was characterized by a profound break with the past. This finding is reminiscent the broad observation of changing correlation between nominal yields and inflation over various business cycles analyzed in Cieslak and Pflueger (2023).[11] In the particular case of the risk-free rate, the correlations observed in the past break down in the post-COVID period because of the substantial departure of the monetary-policy response from prior cycles. As previously discussed, short-term

[11]This result is also in line with Ang et al. (2012), who study the out-of-sample performance of equities in a different sample period, but it is in contrast with Bampinas and Panagiotidis (2016).

rates, such as the one-month T-Bill rate, tend to rise when inflation increases because monetary policy tightens, providing a good inflation hedge. But, when the monetary policy tightening lags inflation, as occurred in 2021, rolling T-Bills suddenly becomes a bad strategy.

A few inflation measures are best hedged (in an univariate sense) by assets other than the T-Bill. In particular, oil future prices typically provide a good hedge for CPI Energy, wheat futures prices a good hedge for the Agriculture sub-index price, and the Wilshire 5000 index for house prices. However, during the period in question, these assets too underperformed. The only asset in the table that performed better than expected was the Case-Shiller index, which is the best-performing univariate hedge for CPI durables prices.

However, some assets, such as oil prices, did protect against multiple measures of inflation even if they performed worse than expected. This can be seen by comparing, in Tables 4 through 8, the realized returns to the inflation rate reported in parenthesis next to the inflation index, which is in the top left corner of each table and sub-panel. Further, other assets, such as the Fama-French portfolios did perform better than expected and their returns more than offset the increase in inflation, but based on their past performance would not have been chosen by investors. Finally, only the Case-Shiller index performed as expected and did not fall short of inflation across many measures of inflation.

4.2 Post-COVID results for two-asset portfolios

Table 9 reports highlights of the results for the post-COVID inflationary period related to the best performing two-asset portfolios, over an investment horizon of two years. The weights for each of the two assets are reported in the 3rd and 5th columns.

Again, it is possible to observe that the risk-free security is one of the two assets in almost all the portfolios, the only exceptions being commodities and housing. Oil futures, the broad commodity index, and the Case-Shiller index are the other assets that show up frequently in the top portfolios, although they tend to receive relatively little weight outside of the hedging regressions for commodity prices themselves. One counterintuitive result is that the best portfolio for hedging house prices (as measured by Case-Shiller) turns out to be composed of the Fama-French conservative-minus-aggressive factor and the Wilshire 5000 index. This seems to reflect in part the out-performance of the CMA strategy in the early 2000s, around the time of the housing boom. Again, however, the performance of strategies to hedge house-price inflation are limited compared to some other cases—the in-sample R^2 of the portfolio is only 0.51.

If the portfolios had performed as predicted during the out-of-sample period, they would have matched the respective changes in inflation fairly well—to within one or two percentage points, in most cases. However, as in the single-asset case, most portfolio returns fell significantly short of their predicted values and therefore failed to hedge inflation. This is not surprising given the high weight placed on the rolling T-Bill as a hedging instrument. Overall, forecast errors are slightly smaller in the two-asset models than in the single-asset models, both in absolute terms and relative to the regression standard errors. Nonetheless, they are still relatively large. Again, this highlights the dependence of assets' inflation-hedging ability on the particular structural regime underlying the data, and on the monetary-policy reaction function in particular.

5 Hedging Unexpected Inflation

To this point, we have measured the inflation-hedging properties of assets using the unconditional correlations implied by our time-series models. These are the appropriate measures for investors who seek equal protection against both the expected and unexpected components of inflation. Such an objective implies that investors ignore potential predictability in inflation and asset returns; thus, it is equivalent to assuming that investors do not update their portfolio decisions based on the state of the economy. Of course, in some cases, it may be that investors only wish to hedge movements in inflation that are *unexpected*, given the information available at the time the assets are purchased. In this case, the relevant metrics are *conditional* correlations—that is, the correlations between the deviations of asset returns and inflation over the period t to $t + h$ from the predicted paths of those variables based on time-t information.

To see how the differences between expected and unexpected outcomes—or, equivalently, between conditional and unconditional correlations—can matter, consider a five-year inflation-indexed Treasury bond held to maturity. By construction, this asset perfectly hedges unexpected movements in headline CPI over the holding period (abstracting from the indexation lag and deflation floor). Thus, its conditional correlation with headline CPI at that horizon is 1. But, in the data, the *unconditional* correlation in this case is only 0.50 (see Appendix Table 3A). The difference between the conditional and unconditional correlations reflects the expected component of the CPI. In the data, times of low real interest rates have preceded times of high inflation. Consequently, a strategy of simply rolling over five-year TIPS would not have generated a return that perfectly hedged total CPI inflation.

More generally, for any inflation index and asset return, the law of total covariance implies

$$\begin{aligned} \text{cov}[\pi_{t+h}, r_{t+h}] &= \text{E}\left[\text{cov}[\pi_{t+h}, r_{t+h}|\mathcal{I}_t]\right] + \text{cov}\left[\text{E}[\pi_{t+h}|\mathcal{I}_t], \text{E}[r_{t+h}|\mathcal{I}_t]\right] && (10) \\ &= \text{cov}[\pi_{t+h}, r_{t+h}|\mathcal{I}_t] + \text{cov}[\text{E}[\pi_{t+h}|\mathcal{I}], \text{E}[r_{t+h}|\mathcal{I}]] && (11) \end{aligned}$$

where $\mathcal{I}_t$ is the information set at time-t, and the second equality follows because our VARIMA models are all homoskedastic. The second term in (11) represents the comovement between the expected components of inflation and asset returns. All else equal, assets that have predictable returns that are positively correlated with the predictable component of inflation will thus have higher unconditional correlations.

For equities, commodities, and currencies, the predictable component of returns ($\text{E}[r_{t+h}|\mathcal{I}_t]$) is nearly constant at most horizons. Consequently, the second term in (11) is small, and the unconditional covariance is a good approximation to the conditional one. For bonds, the situation is quite different. Bond returns are highly predictable, especially at horizons close to their remaining maturities. The five-year TIPS example just discussed is a case in point. Similarly, the returns on nominal bonds held to maturity are known perfectly in advance and therefore their conditional correlations with any random variable are undefined.

Using our estimated VARIMA models, we can compute conditional correlations between bond returns and inflation indices at various horizons up to the bond's maturity, providing a metric for how these assets hedge the unexpected component of inflation. For expositional simplicity, we report the results only for headline and core PCE inflation at the 6-month, 2-year, and 10-year horizons. These are shown in Table 10. For nominal bonds, the conditional correlations are only slightly more negative than the unconditional ones that are shown in the Appendix. For TIPS, they are quite a bit stronger. For example, when held to maturity, a 2-year TIPS is a near-perfect hedge for headline PCE. This compares to an unconditional correlation of just 36%.

Still, even in a conditional sense, TIPS are quite imperfect hedges for horizons and price indices that do not exactly match their contractual terms. For example, even when held to maturity, the 10-year TIPS has a conditional correlation of just 53% with core PCE, and at shorter investment horizons this correlation is even lower. Again, this points to the difficulty of using specific inflation-indexed products to hedge against all types of inflation.

6 Inflation-linked derivatives: with and without market frictions

So far we have analyzed strategies that, in theory, are available to any investor, including retail investors, for instance, through Exchange Traded Funds (ETFs). Flows into TIPS ETFs reached record-high levels in early 2022 (about $60bn) and, over the month of April 2024, because of sticky inflation readings, broad-based commodity ETFs saw an inflow of nearly $970mn.[12] However, conditional on having different demands for inflation protection, market participants have also different levels of financial sophistication and different access to financial instruments (e.g., over-the-counter derivatives). Since large and more sophisticated investors, such as dealers and hedge funds, are usually the participants in the market for inflation risk, and are crucial to its existence and functioning, we consider two additional inflation-hedging instruments usually used by these investors.[13]

In particular, we consider a strategy that replicates TIPS breakeven inflation (BEI) rates, the spread between Treasury nominal and TIPS yields of comparable maturity, and inflation-linked swap (ILS) rates. ILS are over-the-counter derivatives in which one party pays a fixed rate, the ILS rate, and the other party pays a floating rate tied to an inflation index. This index references the non-seasonally-adjusted CPI for all urban consumers and matches the TIPS' inflation index, including the same indexation lag of about 2.5 months. Hence, BEI and ILS rates should provide the same inflation protection in the absence of market frictions. However, as we show shortly, the prevalence of certain market frictions implies that those strategies have different costs and therefore their hedging performance can diverge at times.

Regarding the strategy that replicates BEI rates, here we describe a simplified

[12]Based on Bloomberg data and article of April 30, 2024, by Alex Longley.

[13]See, for instance, Bahaj et al (2023) and Barria and Pinter (2023) for more detail about the participants in the market for inflation risk and their respective demands for inflation protection.

version and refer to Campbell at al. (2009) and Christensen and Gillan (2012) for more detail. The strategy consists of combining two fixed-income asset swaps.[14] A long asset swap position for the zero-coupon TIPS where the dealer agrees to paying SOFR plus the TIPS ASW spread in order to receive the fixed accrued TIPS coupon and the accrued inflation compensation, and a short asset swap position in the nominal Treasury bond where the dealer agrees to paying the fixed accrued nominal coupon and receives SOFR plus the nominal ASW spread.

The cost of this strategy is determined by the difference between the TIPS and nominal asset swap (ASW) spread.[15] Since this is also the most common hedging strategy used by the few dealers that in the U.S. ILS market sell inflation protection, its cost gets transferred to the buyers of inflation protection (i.e., the receivers of the floating rate in the ILS), pushing the ILS rate above the BEI rate. In other words, the difference between the ILS and BEI rate, known as the inflation basis, is almost always positive because it is determined by the difference between the TIPS and nominal ASW spread. As shown in Figure 3, over the last 10 years, the inflation basis at the 10-year maturity has been averaging around 29 basis points. And, when using zero-coupon BEI rates, the size of the inflation basis is similar across different maturities.

In turn, the reason why the TIPS ASW spread is almost always larger than the nominal ASW spread is because TIPS are usually considered less liquid than Treasury nominal securities, and sellers of the TIPS ASW account for the liquidity risk involved in holding TIPS on their balance sheets. Hence, as noted by Campbell et al. (2009), in the absence of default risk, the difference between the TIPS and nominal ASW spread can be regarded as an ideal measure of the relative illiquidity of TIPS.[16] The relative illiquidity of TIPS implies that BEI rates might deviate from their frictionless counterparts because of a premium for illiquidity.

[14]In a fixed-income asset swap (ASW), one party exchanges the fixed-rate cash flows from the underlying security for a floating-rate cash flow, where the floating rate is typically quoted as 6-month LIBOR until March 2022 and SOFR after that, plus a spread.

[15]We should also consider the cost of acquiring and holding the nominal securities, but since usually dealers have them in their inventories, we ignore it for simplicity.

[16]Dittmar, Hsu, Roussellet and Simasek (2019) more recently show that the inflation basis might be affected also by the default risk of the U.S. Treasury.

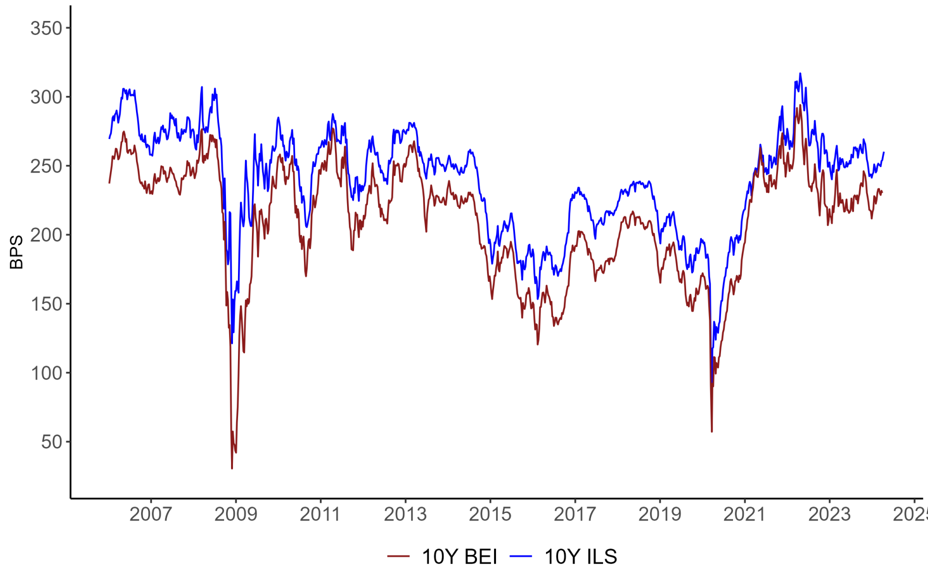

Figure 3: **BEI versus ILS at the 10-year maturity**. The plot shows the time series of the zero-coupon 10-year BEI rate and ILS rate. Source: GSW (2007) dataset and Bloomberg.

Further, as discussed in detail in D'Amico et al. (2018), there are other frictions that can affect BEI rates, such as demand and supply imbalances between nominals and TIPS, the embedded put option priced in TIPS because of the deflation floor, the inflation lag premium, etc. These market frictions and related premiums imply that BEI rates can deviate from investor inflation expectations not only because of the inflation risk premium (e.g., Breach et al., 2020), but also because of a market-frictions premium, potentially affecting the performance of BEI as inflation hedge. To test this hypothesis, in the next section, we build a frictionless BEI and verify whether this hypothetical hedging instrument would provide better inflation protection. This helps illustrate how the limitations of the U.S. market for inflation risk affect investors' ability to properly hedge against inflation.

To compare these strategies, the table below summarizes their respective cash flows (CF). Acquiring inflation protection either through ILS, or by replicating BEI, or even the frictionless BEI provides the same benefit: the CPI inflation accrued over the life of the strategy ($I_{t+\tau}$).[17] In terms of their costs, ILS rates embed the inflation basis, which increases the cost of this instrument relative to BEI and is mostly due to the TIPS-Nominal ASW spread ($ASW^{R,N}$). But, replicating BEI also requires paying the

[17]For simplicity we are ignoring the lag in the inflation index, which is common to all three strategies.

TIPS-Nominal ASW spread, hence these two strategies should deliver identical cash flows. However, BEI rates also embed a market frictions premium (MFP_t), which is usually positive due in part to their relative illiquidity. If this premium turns negative and becomes a discount, which as we discuss below seems to happen during inflationary periods as TIPS are in high demand, then replicating BEI becomes more expensive. In other words, the cash flows generated by replicating the frictionless BEI would be higher than those generated by replicating the actual BEI.

Returns of Actual and Hypothetical Inflation-Linked Derivatives
ILS $CF_{t+\tau}^{ils} = I_{t+\tau} - (1 + r_t^{ils})^{\tau}$ $r_t^{ils} \approx r_t^{bei} + ASW_t^{R,N}$
Replicating BEI $CF_{t+\tau}^{bei} = I_{t+\tau} - (1 + r_t^{bei})^{\tau}$ adjusting for the cost of the strategy: $CF_{t+\tau}^{bei} = I_{t+\tau} - (1 + r_t^{bei} + ASW_t^{R,N})^{\tau}$
Hypothetical BEI* $CF_{t+\tau}^{bei^*} = I_{t+\tau} - (1 + r_t^{bei^*})^{\tau}$ $r_t^{bei^*} = r_t^{bei} + MFP_t$ if $MFP_t < 0 \rightarrow CF_{t+\tau}^{bei^*} > CF_{t+\tau}^{bei}$

6.1 A hypothetical inflation hedge: Frictionless BEI

As shown in multiple studies,[18] TIPS contain a market friction premium, which pushes the TIPS yield above the frictionless real yield and, therefore, the actual BEI below the frictionless BEI. This would imply that replicating BEI is usually cheaper than what ought to be. Figure 4 provides an example of the sign and size of this market friction premium using the estimates from D'Amico et al. (2018) at the 2- and 10-year maturities.

[18]See, for instance, D'Amico et al. (2018); Abrahams et al. (2016); Andreasen et al. (2021).

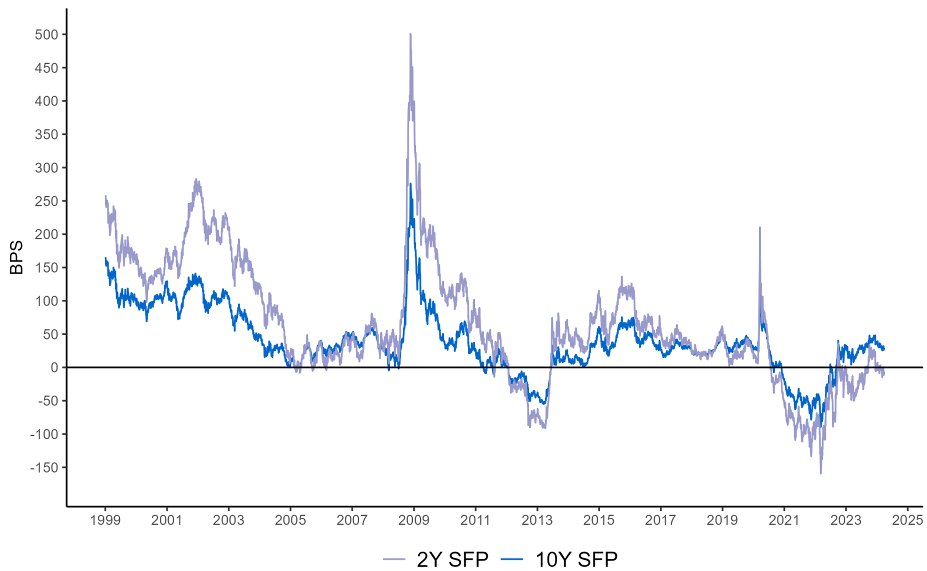

Figure 4: **Estimated Market Friction Premium at 2- and 10-year maturity**. The plot shows the time series of market friction premiums estimated using the model of D'Amico, Kim, and wei (2018). Source: GSW (2007) dataset and Bloomberg.

According to these estimates, the market friction premium turned negative during the inflationary wave and returned into positive territory toward the end of 2023. This is not unique to the estimates of D'Amico et al. (2018), as this premium turned negative also in the case of European inflation-linked bonds, according to the estimates of Christensen and Mouabbi (2024), and in the case of UK ILS rates, as estimated by Baja et al. (2023). A negative market friction premium makes TIPS more expensive, widens the BEI rate relative to its frictionless counterpart, and therefore constitutes an additional cost to the inflation hedging strategy consisting of replicating BEI.

We build a frictionless BEI rate (BEI*) by adding the estimated market friction premium to the actual BEI rate. Figure 5 depicts the comparison at the 2-year maturity, between the ILS rate, BEI rate, and BEI* rate. The ILS rates are often closer to BEI* than the raw BEI rates, especially during disinflation periods. However, during the inflationary wave, ILS and BEI have been very close to each other, providing a similar approximation of BEI*. And, since 2023, the actual BEI has shown small deviations from its frictionless counterpart.

Overall, this suggests that during the inflationary wave, when investors increase their demand for inflation protection, TIPS have become more expensive, likely due to the limited supply of inflation protection in the US market. Next, we analyze how the limited supply of inflation protection has affected the performance of the inflation-linked derivatives as inflation hedge.

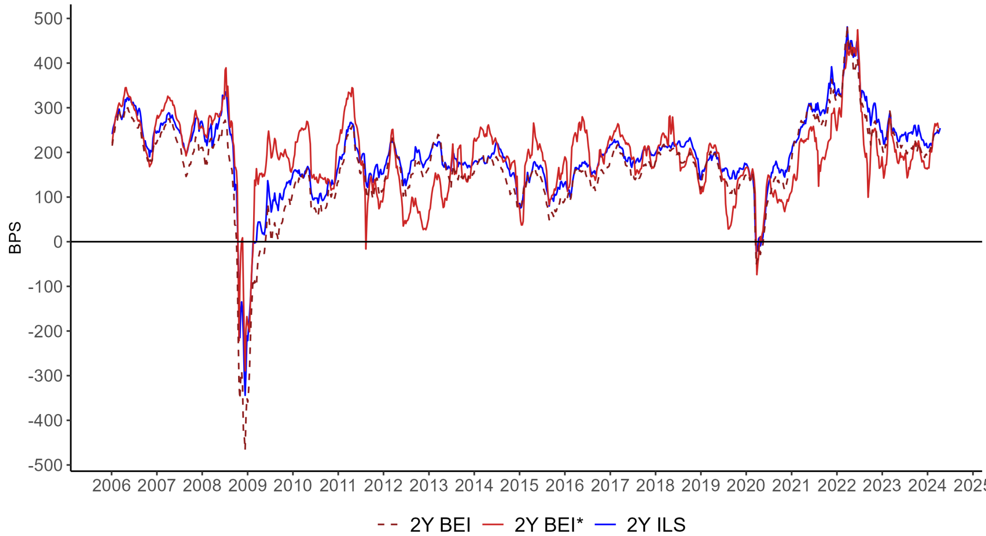

Figure 5: **Comparison of 2-year ILS rate vs BEI rate vs frictionless BEI**. The plot shows the time series of frictionless BEI obtained adding to the raw BEI the market friction premium estimated using the model of D'Amico, Kim, and wei (2018), together with the raw 2-year BEI and 2-year ILS rate. Source: GSW (2007) dataset and Bloomberg.

6.2 In-sample performance of inflation-linked derivatives

In this section, we consider only strategies in which the holding period coincides with the maturity of the instrument.

Table 8A in the Appendix reports the in-sample correlations for inflation swap rates estimated over the sample period starting in 2004, when those rates are available in Bloomberg, and ending in August 2021. We find that inflation swaps display very high correlations with energy and commodity prices (between 66% and 83%) as well as CPI non-durable good prices. They are also quite good at hedging Headline PCE and PPI, with correlations between 65% and 79%, but the correlations with CPI and PCE core are at most as high as 55-58% at the 5- and 10-year maturities. Similarly to other inflation hedges, ILS correlations with CPI Services and wages are rather weak.

Table 9A in the Appendix reports the in-sample correlations for BEI rates estimated over the sample period starting in 1999 and ending in August 2021. We exclude the first two years of the TIPS existence (1997 and 1998) to be able to include zero-coupon 2-year BEI rates. It can be noted that most of the correlations are lower than those estimated for the ILS rates, except for CPI Durable goods prices, suggesting that there are market frictions that make BEI rates somewhat worse hedges than ILS rates.

In fact, when we analyze the hedging ability of the frictionless BEI rate, whose estimated correlations are reported in Table 10A in the Appendix, it is possible to see that many of the estimated correlations are higher than those obtained for the actual BEI. Further, as shown in Table 11, some of the correlations are also higher than those reported for the ILS. In particular, the frictionless BEI seems to provide a better hedge for all core price indices, especially at the 10-year maturity. At this maturity, it also displays a correlation as high as 50% with CPI Services prices and 22% with wages. The frictionless BEI also seems to provide a better hedge for housing prices. But, it definitely performs worse than ILS in the case of energy and commodity prices.

These findings indicate that the market friction premium and its dynamic weaken the ability of actual BEI rates to provide inflation protection, especially when this protection is the most needed. The lack of natural sellers of inflation protection in the ILS market, and the low Treasury issuance of TIPS during the inflationary wave, have likely prevented investors from having the "optimal" inflation protection, that is, the protection that could have been obtained in the absence of certain market frictions.

We do not analyze the out-of-sample performance of the actual and frictionless inflation-linked derivatives described in this section because their performance in the case of Headline CPI should be very good by construction, and, relative to all the other indices considered in this study, it will just depend on the correlation of the CPI inflation rate with all the other inflation rates.

7 Conclusion

There is no one-asset-fits-all approach to inflation hedging; the optimal hedge depends on the particular types of prices that an investor is exposed to and at which horizons. Over the period 1999 to 2021, we find that commodities and commodity-related stocks and currencies were generally successful in hedging headline consumer inflation, but this mostly seems to reflect their significantly positive relation with energy prices. Hedging core inflation is harder. At horizons of less than a year, there was little protection available, except for inflation swaps and TIPS. At longer horizons, short-term nominal bonds and real estate provided decent hedges, and there is some evidence supporting certain stock-market strategies. Core producer prices and wages were the most difficult types of inflation to hedge, although real estate and short-term nominal bonds provided some protection. And, since wages are a particularly large component of the cost of non-housing services, the same instruments did a decent job hedging service inflation. Finally, house prices were effectively hedged through REITs and other instruments that provide broad exposure to this sector.

Those statements apply entirely to the historical, *in-sample* experience. A separate question—perhaps more relevant for practicing investors—is how various assets and portfolios perform out of sample. Here, unfortunately, our results are less encouraging. At least during the out-of-sample period we look at, September 2021-August 2023, many of the hedging relations that had prevailed since at least 1972 broke down. In particular, short-term interest rates did not adhere to their usual pattern of closely following broad measures of inflation. These results highlight the dependence of assets' inflation-hedging ability on the structural regime and the monetary-policy rule and further reinforce the difficulties of hedging any measure of inflation reliably.

Further, our findings indicate that TIPS market frictions premium and its dynamic weaken the ability of actual BEI rates to provide inflation protection, especially when this protection is the most needed. The lack of natural sellers of inflation protection in the inflation swap market, and the low Treasury issuance of TIPS during the inflationary wave, have likely prevented investors from having the "optimal" inflation protection, that is, the protection that could have been obtained in the absence of market frictions.

Bibliography

Abrahams, M., T. Adrian, R. K Crump, E. Moench, and R. Yu, 2016. Decomposing real and nominal yield curves. Journal of Monetary Economics, 84, 182–200.

Ajello, A.; L. Benzoni; and O. Chyruk. 2019. Core and 'Crust': Consumer Prices and the Term Structure of Interest Rates. Review of Financial Studies 33 (8), 3719-3765.

Andreasen, M.M., Christensen J. H. E., and Riddell S, 2021. The TIPS Liquidity Premium. Review of Finance 25(6), 2021, 1639–1675.

Ang, A., Brière, M., and Signori, O., 2012. Inflation and individual equities. Financial Analysts Journal, 68(4), 36–55.

Bahaj, S., Czech R., Ding S., and Reis R., 2023. The Market for Inflation Risk. Staff Working Paper No. 1028. Bank of England, June.

Barria, R., and G. Pinter, 2023. Mispricing in Inflation Markets. Staff Working Paper No. 1034. Bank of England, August.

Breach, T., S., D'Amico, and A., Orphanides, 2020. The term structure and inflation uncertainty. Journal of Financial Economics 138, 388-414.

Campbell, J.Y., R. J. Shiller, and L. M. Viceira, 2009. Understanding Inflation-Indexed Bond Markets. In Brookings Papers on Economic Activity, ,79–120.

Christensen, J. H. E., and J. M. Gillan, 2012. Could the U.S. Treasury Benefit from Issuing More TIPS? Federal Reserve Bank of San Francisco Working Paper 2011–16.

Christensen, J. H.E., and S, Mouabbi, 2024. The Natural Rate of Interest in the Euro Area: Evidence from Inflation-Indexed Bonds. Federal Reserve Bank of San Francisco Working Paper 2024-08.

Cieslak, A., and C. Pflueger, 2023. Inflation and Asset Returns. NBER Working Paper 30982, March.

D'Amico, S., D. Kim and M. Wei, 2018. Tips from TIPS: the information content of Treasury inflation-protected security prices. Journal of Financial and Quantitative Analysis 53, 395-436.

D'Amico, S., and T. B. King, 2023. Inflation hedging products, in the *Handbook of Financial Markets*, edited by Refet Gürkaynak and Jonathan Wright, Edward Elgar Press, May 18, 2023, 470-489.

Dittmar, R. F., A. Hsu, G. Roussellet, and P. Simasek, 2023. "Default Risk and the Pricing of U.S. Sovereign Bonds." Georgia Tech Scheller College of Business Research Paper No. 18-20.

Fang, X., Y., Liu, and N. Roussanov. 2021. Getting to the core: Inflation risks within and across asset classes. Working Paper, May.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007. The U.S. Treasury Yield Curve: 1961 to the Present. Journal of Monetary Economics, 54(8), 2291-2304.

Kat, H. M., and R. C. A. Oomen, 2006. What every investor should know about commodities (Part II). Alternative Investment Research Center Working Paper, No.0033.

Kwak, M, and B. H. Lim, 2014. Optimal portfolio selection with life insurance under inflation risk. Journal of Banking and Finance, 46, 59–71.

Parikh, H., R. K. Malladi, and F. J. Fabozzi, 2019. Preparing for higher inflation: Portfolio solutions using U.S. equities. Review of Financial Economics, 542-554.

Powell, Jerome H., 2022. Inflation and the Labor Market. Speech at the Hutchins Center on Fiscal and Monetary Policy, Brookings Institution, Washington, D.C. November 30, 2022.