

RELATIVE GOODS' PRICES, PURE INFLATION, AND THE PHILLIPS CORRELATION

By Ricardo Reis and Mark W. Watson*

This paper uses a dynamic factor model for the quarterly changes in consumption goods' prices in the U.S. since 1959 to separate them into three independent components: idiosyncratic relative-price changes, a low-dimensional index of aggregate relative-price changes, and an index of equiproportional changes in all inflation rates, that we label "pure" inflation. We use the estimates to answer two questions. First, what share of the variability of inflation is associated with each component, and how are they related to conventional measures of monetary policy and relative-price shocks? Second, what drives the Phillips correlation between inflation and measures of real activity? (JEL E31, C43, C32)

*Reis: Department of Economics, Columbia University, 1002 International Affairs Building, 420 West 118th Street, New York, NY 10027 (e-mail: rreis@columbia.edu); Watson: Department of Economics and Woodrow Wilson School, Princeton University, Princeton, NJ 08544-1013 (e-mail: mwatson@princeton.edu). We are grateful to several colleagues for their valuable input to this research project. For their support and hospitality, Reis thanks the Hoover Institution at Stanford University under a Campbell national fellowship, Columbia University's Program for Economic Research, the University of Chicago, and Princeton University where much of the work for this paper was completed. The National Science Foundation provided support through grant SES-0617811.

One of the goals of macroeconomics is to explain the aggregate sources of changes in goods' prices. If there was a single consumption good in the world, as is often assumed in models, describing the price changes of consumption would be a trivial matter. But, in reality, there are many goods and prices, and there is an important distinction between price changes that are equiproportional across all goods (absolute-price changes) and changes in the cost of some goods relative to others (relative-price changes). The goal of this paper is to empirically separate these sources of price changes and to investigate their relative size, their determinants, and their role in the macroeconomic Phillips relation between inflation and output.

Our data are the quarterly price changes in each of 187 sectors in the U.S. personal consumption expenditures (PCE) category of the national income and product accounts from 1959:1 to 2006:2. Denoting the rate of price change for the i 'th good between dates $t-1$ and t by π_{it} , and letting $\boldsymbol{\pi}_t$ be the $N \times 1$ vector that collects these goods' price changes, we model their co-movement using a linear factor model:

$$(1) \quad \boldsymbol{\pi}_t = \boldsymbol{\Lambda} \mathbf{F}_t + \mathbf{u}_t$$

The k factors in the $k \times 1$ vector \mathbf{F}_t capture common sources of variation in prices. These might be due to aggregate shocks affecting all sectors, like changes in aggregate productivity, government spending, or monetary policy, or they might be due to shocks that affect many but not all sectors, like changes in energy prices, weather events in agriculture, or exchange rate fluctuations and the price of tradables. The $N \times k$ matrix $\boldsymbol{\Lambda}$ contains coefficients ("factor loadings") that determine how each individual good's price responds to these shocks. The $N \times 1$ vector \mathbf{u}_t is a remainder that captures good-specific relative-price variability associated with idiosyncratic sectoral events or measurement error.

We see the empirical model in (1) as a useful way to capture the main features of the covariance matrix of changes in good's prices. To the extent that the factors in \mathbf{F}_t explain a significant share of the variation in the data, then changes in goods' prices provide information on the aggregate shocks that macroeconomists care about. We separate this aggregate component of price changes into an absolute-price component and possibly several relative-price components. Denoting these by the scalar a_t and the \mathbf{R}_t vector of size $k-1$ respectively, this decomposition can be written as:

$$(2) \quad \Delta \mathbf{F}_t = \mathbf{I}a_t + \mathbf{\Gamma}\mathbf{R}_t$$

Absolute price changes affect all prices equiproportionately, so \mathbf{I} is an $N \times 1$ vector of ones, while relative price changes affect prices in different proportions according to the $N \times (k-1)$ matrix $\mathbf{\Gamma}$. The first question this paper asks is whether the common sources of variation, $\Delta \mathbf{F}_t$, can be decomposed in this way.

One issue is that \mathbf{I} may not be in the column space of $\mathbf{\Lambda}$; that is, there may no absolute-price changes in the data. Given estimates of the factor model, we can investigate this empirically using statistical tests and measures of fit. Another issue is that the decomposition in (2) is not unique; that is, a_t and \mathbf{R}_t are not separately identified. The key source of the identification problem is easy to see: for any arbitrary $(k-1) \times 1$ vector $\boldsymbol{\alpha}$, we have that $\mathbf{I}a_t + \mathbf{\Gamma}\mathbf{R}_t = \mathbf{I}(a_t + \boldsymbol{\alpha}'\mathbf{R}_t) + (\mathbf{\Gamma} - \mathbf{I}\boldsymbol{\alpha}')\mathbf{R}_t$, so that (a_t, \mathbf{R}_t) cannot be distinguished from $(a_t + \boldsymbol{\alpha}'\mathbf{R}_t, \mathbf{R}_t)$. The intuition is that the absolute change in prices cannot be distinguished from a change in “average relative prices” $\boldsymbol{\alpha}'\mathbf{R}_t$, but there are many ways to define what this average means.¹

¹ One natural way is to assume that relative price changes must add up to zero across all goods. Reis and Watson (2007) use this restriction to define a numeraire price index that measures absolute price changes. A further identification issue in the model is that $\mathbf{\Gamma}\mathbf{R}_t = \mathbf{\Gamma}\mathbf{A}\mathbf{A}^{-1}\mathbf{R}_t$ for arbitrary non-singular matrix \mathbf{A} . For our purposes we will not need to separately identify the elements of \mathbf{R}_t so this final issue is not important.

We overcome this identification problem by focusing instead on two independent components: “pure” inflation v_t , and a low-dimensional relative price index $\boldsymbol{\rho}_t$ defined as:

$$(3) \quad v_t = a_t - E[a_t | \{\mathbf{R}_\tau\}_{\tau=1}^T]$$

$$(4) \quad \boldsymbol{\rho}_t = E[\mathbf{F}_t | \{\mathbf{R}_\tau\}_{\tau=1}^T]$$

Pure inflation is identified, and it has a simple interpretation: it is the common component in price changes that has an equiproportional effect on all prices and is uncorrelated with changes in relative prices at all dates. We label it “pure” because, by construction, its changes are uncorrelated with relative-price changes at any point in time, and because it corresponds to the famous thought experiment that economists have used since David Hume (1752): “imagine that all prices increase in the same proportion, but no relative price changes.”² The relative-price index captures all of the aggregate movements in goods’ price changes that are associated with some change in relative prices at some date. In an economic model, these components map into different fundamental shocks. For instance, an exogenous but anticipated increase in the money supply that leads all price-setters to raise their prices in the same proportion leads to pure inflation, while an unanticipated increase in money to which some firms respond, but others do not, leads to a change in the relative-price index.³ In this dichotomy, inflation due to changes in the money supply that are reactions to relative price changes also lead to changes in the relative-price index, since pure inflation is uncorrelated with any relative-price change.

The first contribution of this paper consists in estimating the empirical model in (1)-(4) providing a decomposition of inflation into three independent components:

² The definition of pure inflation also appears in current textbooks (Olivier J. Blanchard, 2003: 33).

³ In some sticky-price models, like a Calvo model in which the frequency of price adjustments differs across sectors, an anticipated increase in the money supply would not lead to pure inflation, as the price-setters are, by assumption, unable to all raise their prices at once.

$$(5) \quad \pi_t = \mathbf{1}v_t + \Theta\rho_t + \mathbf{u}_t$$

Our estimates show that these three components have differed markedly over the last 40 years, and allow us to address two issues. First, we are able to quantify the share of inflation's variability associated with each of the components. We find that for a typical good, its idiosyncratic relative-price component accounts for roughly 70 percent of its variability, so that macroeconomic shocks account for almost as much as 1/3 of the movement in sectoral prices. Within aggregate sources of variation, pure inflation accounts for about 15-20 percent of the variability in PCE inflation. Researchers must be cautious when comparing the predictions for inflation from models with a single consumption good to the data, because most of the variation in standard aggregate inflation indices is associated with relative-price movements, which these models ignore. Second, we relate our estimates to other variables. At business-cycle frequencies, pure inflation is barely correlated with money growth, while it has a correlation of around 0.5 with nominal interest rates. The relative-price index is weakly related to food and energy prices, but it is strongly related to the relative price of non-durable and services. However, even considering as many as four conventional measures of relative-price changes, the two relative-price factors in our baseline specification appear to be a more comprehensive measure of relative price movements.

The second contribution of this paper is to re-examine the correlation between inflation and real activity. Alban W. H. Phillips (1958) famously first estimated it, and a vast subsequent literature confirmed that it is reasonably large and stable (James H. Stock and Watson, 1999). This correlation has posed a challenge for macroeconomists because it signals that the classical dichotomy between real and nominal variables may not hold. The typical explanation for the Phillips correlation in economic models involves movements in

relative prices. For instance, models with sticky wages but flexible goods prices (or vice-versa), explain it by movements in the relative price of labor. Models of the transaction benefits of money or of limited participation in asset markets explain the Phillips correlation by changes in the relative price of consumption today vis-à-vis tomorrow, or asset returns. Models with international trade and restrictions on the currency denomination of prices explain it using the relative price of domestic vis-à-vis foreign goods, or exchange rates. We show that, after controlling for all of these relative prices, the Phillips correlation is still quantitatively and statistically significant. Then, using our estimates, we control instead for the relative price of different goods. This would be suggested by models with many consumption goods, as is the case in modern sticky-price or sticky-information models. We find that, controlling for relative goods prices, the Phillips correlation becomes quantitatively negligible. This suggests a more important role for rigidities in goods markets and a less important role for rigidities in labor or asset markets.

The paper is organized as follows. Section I outlines the methods that we use to estimate the factor model and to compute the inflation components and their correlation with other variables. Section II presents a stylized structural model of inflation dynamics that generates the decomposition in equation (5) and relates its three components to fundamental economic shocks. Sections III and IV present estimates of the factor model, the factors, and their relation to observables. Section V investigates the Phillips correlation, and section VI discusses the robustness of the conclusions in the previous two sections to different specifications. Section VII concludes, summarizing our findings and discussing their implications.

Relation to the literature

There has been much research using statistical models to define and measure inflation

(see the survey by E. Anthony Selvanathan and D. S. Prasada Rao, 1994) but, as far as we are aware, there have been relatively few attempts at separating absolute from relative-price changes. An important exception is Michael F. Bryan and Stephen G. Cecchetti (1993), who use a dynamic factor model in a panel of 36 price series to measure what we defined above as a_t . They achieve identification and estimate their model imposing strong and strict assumptions on the co-movement of relative prices, in particular that relative prices are independent across goods. Moreover, while they use their estimates to forecast future inflation, we use them to separate inflation into components and to assess the Phillips correlation.⁴

In methods, our use of large-scale dynamic factor models draws on the literature on their estimation by maximum likelihood (e.g., Danny Quah and Thomas J. Sargent, 1993, and Catherine Doz, Domenico Giannone and Lucrezia Reichlin, 2008) and principal components (e.g., Jushan Bai and Serena Ng, 2002, Mario Forni, Marc Hallin, Marco Lippi and Reichlin, 2000, and Stock and Watson, 2002). We provide a new set of questions to apply these methods.

Using these methods on price data, Riccardo Cristadoro, Forni, Reichlin and Giovanni Veronesi (2005) estimate a common factor on a panel with price and quantity series and ask a different question: whether it forecasts inflation well. Marlene Amstad and Simon M. Potter (2007) address yet another issue, using dynamic factor models to build measures of the common component in price changes that can be updated daily. Marco Del Negro (2006) estimates a factor model using sectoral PCE data allowing for a single common component and relative price factors associated with durable, non-durable, and services goods sectors. Finally, Filippo Altissimo, Benoit Mojon, and Paolo Zaffaroni (2009) estimate a common factor model using disaggregated Euro-area CPI indices and use the model to investigate the

⁴Bryan, Cecchetti and Roisin O’Sullivan (2002) use a version of the Bryan-Cecchetti (1993) model to study the importance of asset prices for an inflation index.

persistence in aggregate Euro-area inflation. The common factor in these papers is not a measure of pure inflation, since it affects different prices differently.

Closer to our paper, in the use of dynamic factor models to extract a measure of inflation that is then used to assess macroeconomic relations suggested by theory, is Jean Boivin, Marc Giannoni and Ilian Mihov (2009). They extract a macroeconomic shock using many series that include prices and real quantities, estimate the impulse response of individual prices to this shock, and then compare their shape to the predictions of different models of nominal rigidities. In contrast, we use only price data (and no quantity data) to separate different components of inflation, so that we can later ask how they relate to quantities. Moreover, we apply our estimates to assess unconditional correlations of real variables with inflation, whereas they focus on the link conditional on identified monetary shocks. Finally, we separate relative prices and pure inflation, while their inflation measure is a mix of the two, so we ask a different set of questions.

I. Measuring the Components of Inflation and Calculating Macro-Correlations

The model in (1)-(4) is meant to capture the key properties of the inflation series as they pertain to the estimation of their separate components, with an eye on the applications that we discussed in the introduction. We use a factor model for the covariance between sectoral inflation rates because past research focusing on the output of different sectors, and macroeconomic variables more generally, found that this model is able to flexibly and parsimoniously account for the main features of the economic data (Stock and Watson, 1989, 2005, Forni et al, 2000).

A. Estimating the Dynamic Factor Model

The strategy for estimating the model can be split in two steps. First, we choose the number of factors (k). Second, we estimate the factors (a_t, \mathbf{R}_t) and the factor loadings ($\mathbf{\Gamma}$), and examine the restriction that the factor loading on a_t is equal to unity. We discuss each of these in turn.

Choosing the number of factors, that is the size k of the vector \mathbf{F}_t , involves a trade-off. On the one hand, a higher k implies that a larger share of the variance in the data is captured by the aggregate components. On the other hand, the extra factors are increasingly harder to reliably estimate and are less quantitatively significant. Bai and Ng (2002) have developed estimators for k that are consistent (as $\min(N, T) \rightarrow \infty$) in models such as this. We compute the Bai-Ng estimators, which are based on the number of dominant eigenvalues of the covariance (or correlation) matrix of the data. We complement them by also looking at a few informative descriptive statistics on the additional explanatory power of the marginal factor. In particular, we estimate an unrestricted version of (1) that does not impose the restriction in (2) that the first factor has a unit loading. We start with one factor and successively increase the number of factors, calculating at each step the incremental share of the variance of each good's inflation explained by the extra factor. If the increase in explained variance is large enough across many goods, we infer it is important to include at least these many factors. These pieces of information lead us to choose a benchmark value for k . In section VI, we investigate the robustness of the results to different choices of k .

To estimate the factor model, we follow two approaches. The first approach estimates (1)-(2) by restricted principal components. It consists of solving the least-squares problem:

$$(6) \quad \min_{\mathbf{\Gamma}, (a, \mathbf{R})} \sum_{i=1}^N w_i \sum_{t=1}^T (\pi_{it} - a_t - \gamma_i' \mathbf{R}_t)^2$$

where γ_i denotes the i 'th row of Γ (from (2)) and w_i are weights. We set w_i equal to the inverse of the sample variance of π_{it} so that the solution to (6) yields the restricted principal components associated with the sample correlation matrix of the inflation series (C. Radhakrishna Rao (1973), section 8g.2). When N and T are large and the error terms u_{it} are weakly cross-sectionally and serially correlated, the principal components/least squares estimators of the factors have two important statistical properties that are important for our analysis (Stock and Watson, 2002, Bai, 2003, Bai and Ng, 2006). First, the estimators are consistent. Second, the sampling error in the estimated factors is sufficiently small that it can be ignored when the estimates, say \hat{a}_t and $\hat{\mathbf{R}}_t$, are used in regressions in place of the true values of a_t and \mathbf{R}_t .

The second approach makes parametric assumptions on the stochastic properties of the three latent components (a_t , \mathbf{R}_t , and u_{it}), estimates the parameters of the model by maximum likelihood, and then computes estimates of the factors using signal extraction formulae. In particular, we assume that (a_t, \mathbf{R}_t) follow a vector autoregression, while u_{it} follow independent autoregressive processes, all with Gaussian errors.⁵ The resulting unobserved-components model is:

$$(7) \quad \pi_{it} = a_t + \gamma_i' \mathbf{R}_t + u_{it}$$

$$(8) \quad \Phi(L) \begin{pmatrix} a_t \\ \mathbf{R}_t \end{pmatrix} = \boldsymbol{\varepsilon}_t$$

$$(9) \quad \beta_i(L) u_{it} = c_i + e_{it}$$

⁵ One concern with assuming Gaussianity is that disaggregated inflation rates are skewed and fat-tailed. In general, skewness is not a major concern for Gaussian MLEs in models like this (Watson, 1989), but excess kurtosis is more problematic. To mitigate the problem, we follow Bryan, Cecchetti and Rodney L. Wiggins II (1997) and pre-treat the data to eliminate large outliers (section III has specifics).

with $\{e_{it}\}, \{e_{jt}\}_{j \neq i}, \{\varepsilon_t\}$ being mutually and serially uncorrelated sequences that are normally distributed with mean zero and variances $\text{var}(e_{it}) = \sigma_i^2$, $\text{var}(\boldsymbol{\varepsilon}_t) = \mathbf{Q}$. To identify the factors, we use the normalizations that the columns of $\boldsymbol{\Gamma}$ are mutually orthogonal and add up to zero, although the estimates of v_t and ρ_t do not depend on this normalization.

Numerically maximizing the likelihood function is computationally complex because of the size of the model. For example, our benchmark model includes a_t and two additional relative price factors, a VAR(4) for (8), univariate AR(1) models for the $\{u_{it}\}$, and $N = 187$ price series. There are 971 parameters to be estimated.⁶ Despite its complexity, the linear latent variable structure of the model makes it amenable to estimation using an EM algorithm with the “E-step” computed by Kalman smoothing and the “M-step” by linear regression. The web appendix to this paper describes this in more detail.

While this exact dynamic factor model (7)-(9) is surely misspecified – for instance, it ignores small amounts of cross-sectional correlation among the u_{it} terms, conditional heteroskedasticity in the disturbances, and so forth – it does capture the key cross sectional and serial correlation patterns in the data. Doz, Giannone and Reichlin (2008) study the properties of factors estimated from an exact factor structure as in (7)-(9) with parameters estimated by Gaussian MLE, but under the assumption that the data are generated from an approximate factor model (so that (7)-(9) are misspecified). Their analysis shows that when N and T are large, the factor estimates from (7)-(9) are consistent despite potential misspecification in the model.

We carry out our analysis using the principal components estimates of the factors and the estimates from (7)-(9). To save space, unless noted otherwise, the results reported in sections III-V are based on the estimates from the parametric factor model (7)-(9); results with the principal components estimates of the factors are shown in section VI, which

⁶The number of unknown parameters is $186 + 185 (\gamma_i) + 187 (\beta_i) + 187 (c_i) + 187 (\text{var}(e_i)) + 36 (\boldsymbol{\Phi}) + 3 (\text{var}(\boldsymbol{\varepsilon})) = 971$, where these values reflect the normalizations used for identification.

focuses on the robustness of the empirical conclusions.

The model in (1)-(2) imposes the restriction that the loading on the absolute-price factor must be one for all goods. To investigate how restrictive this is, we calculate the increase in fit that comes from dropping the restriction, measured as the fraction of (sample) variance of π_i explained by the factors. Moreover, we estimate the value of ς_i in the N regressions:

$$(10) \quad \pi_{it} = \varsigma_i a_t + \boldsymbol{\xi}_i' \mathbf{R}_t + u_{it},$$

using \hat{a}_t and $\hat{\mathbf{R}}_t$ in place of a_t and \mathbf{R}_t , as explained above. When $\varsigma_i = 1$, this corresponds to our restricted model, so we can use the estimates of ς_i to judge how adequate is this restriction.

B. Computing the Aggregate Components of Inflation

To separate the components of inflation and obtain time-series for pure inflation and the relative-price index (v_t, ρ_t) , we need to calculate the expectation of absolute-price changes conditional on relative-price changes in (3)-(4). This requires a model of the joint dynamics of a_t and \mathbf{R}_t . We model this as a VAR, as in (8), which is estimated by Gaussian MLE in the parametric factor model, or by OLS using the principal-component estimators for the factors as in the two-step approach taken in factor-augmented VARs (Bernanke, Boivin, Elias, 2005). Finally, given estimates of $\Phi(L)$, we compute the implied projection in (3) and (4) to obtain pure inflation and the relative-price index. Details are provided in the web appendix.

C. Computing Macro-Correlations at Different Frequencies

As described in the introduction, we are interested in the relationship between pure

inflation and the relative-price index, v_t and ρ_t , and other macro variables such as the PCE deflator, food and energy prices, the unemployment rate, or the nominal interest rate. Let x_t denote one of these macro variables of interest and consider the projection of x_t onto leads and lags of v_t (or ρ_t)

$$(11) \quad x_t = \delta(L)v_t + e_t.$$

The fraction of variability of x_t associated with $\{v_t\}$ (or $\{\rho_t\}$) can be computed as the R^2 from this regression. Adding additional control variables, say z_t , to the regression makes it possible to compute the partial R^2 of x with respect to leads and lags of v_t (or ρ_t) after controlling for z_t .

We will compute frequency-domain versions of these variance decompositions and partial R^2 's (squared coherences or partial squared coherences). One of their virtues is that they allow us to focus on specific frequency bands, like business cycle frequencies. Another virtue is that they are robust to the filter used to define the variables (e.g., levels or first differences). In particular we report the squared coherence (the R^2 at a given frequency) between x and v (or ρ) averaged over various frequency bands. When it is relevant, we also report partial squared coherences controlling for (leads and lags) of a vector of variables \mathbf{z} (the partial R^2 controlling for \mathbf{z} at a given frequency), again averaged over various frequency bands.

We are also interested in the relationship between \mathbf{R}_t and standard measures of relative prices such the relative price of consumer durables, food, energy, and so forth. Let \mathbf{q}_t denote a vector of these variables. We summarize the correlation between \mathbf{R} and \mathbf{q} using canonical correlations, again implemented using frequency domain methods. In particular we report the squared canonical coherences between \mathbf{R} and \mathbf{q} averaged over various

frequency bands.

These various spectral R^2 measures are computed using VAR spectral estimators, where the VAR is estimated, in the first instance, using x_t , \hat{a}_t , $\hat{\mathbf{R}}_t$, and (if appropriate) \mathbf{z}_t , and in the second instance using $\hat{\mathbf{R}}_t$ and \mathbf{q}_t . The standard errors for the spectral measures are computed using the delta-method and a heteroskedastic-robust estimator for the covariance matrix of the VAR parameters. Details are provided in the web appendix.

II. A Theoretical Framework

The statistical decomposition in (5) expresses π_{it} in terms of three components: v_t , which we have labeled pure inflation, \mathbf{p}_t , which is a function of aggregate relative-price shocks, and u_{it} , which captures sector-specific relative-price changes or measurement error. Structural macro models give rise to an analogous representation for inflation, where the components depend on the various shocks in the macro model. The specifics of the structural model determine the relative variances of the components and their correlation with non-price variables such as real output, money, interest rates, and so forth.

This section presents a simple economic model of inflation that relates key structural shocks to the statistical constructs v , \mathbf{p} and u . The goal is to help guide the reader's intuition about the forces underlying the statistical factors that we estimate, so the model is kept as simple as possible. We believe that its main conclusions are robust to the specific modeling choices (like the source of nominal rigidities, or the particular functional forms), and that this model could serve as the starting point for a more structural empirical analysis of inflation's components in future work.

A. The Model

The general-equilibrium setup follows Blanchard and Nobuhiro Kiyotaki (1987) and Michael Woodford (2003) allowing for many sectors. A representative consumer maximizes:

$$(12) \quad E_0 \left[\sum_{t=0}^{\infty} \delta^t \left(\ln C_t - \frac{L_t^{1+\psi}}{1+\psi} \right) \right]$$

where $\delta < 1$ is the discount factor and ψ is the elasticity of labor supply. Total labor supplied is L_t and C_t is a consumption aggregator with elasticity γ across N sectors, indexed by i , and across a continuum of varieties within each sector, indexed by j :

$$(13) \quad C_t = \left(N^{-1/\gamma} \sum_{i=1}^N C_{it}^{(\gamma-1)/\gamma} \right)^{\gamma/(\gamma-1)}, \text{ with } C_{it} = \left(\int_0^1 C_{it}(j)^{(\gamma-1)/\gamma} dj \right)^{\gamma/(\gamma-1)}.$$

At every date, the consumer purchases each good at price $P_{it}(j)$ for a total spending of $S_t = \sum_i \int P_{it}(j) C_{it}(j) dj$, earns a wage W_t for labor services, and pays taxes at rate T_t . She has three other sources income that are lump-sum: profits D_t from firms, transfers G_t from the fiscal authorities, and money injections H_t from the monetary authority. Finally, the consumer holds money M_t to save and to purchase consumption goods. The budget and cash-in-advance constraints are:

$$(14) \quad S_t + M_t = (1 - T_t)W_t L_t + M_{t-1} + D_t + G_t + H_t,$$

$$(15) \quad S_t \leq M_{t-1} + H_t.$$

Firms are monopolistically competitive, each hiring labor $L_{it}(j)$ to produce output $Y_{it}(j)$ with productivity $X_{it}(j)$ subject to decreasing returns to scale at rate $\eta < 1$:

$$(16) \quad Y_{it}(j) = X_{it}(j)L_{it}(j)^\eta$$

Finally, the two government authorities simply return their funds to consumers, so that $H_t = M_t - M_{t-1}$ and $G_t = T_t W_t L_t$. Market clearing in the goods and labor market require $Y_{it}(j) = C_{it}(j)$ and $L_t = \sum_i \int L_{it}(j) dj$ respectively.

In this simple economy, there is uncertainty about taxes T_t , the money supply M_t , and productivity $X_{it}(j)$, each of which depends on shocks. Letting small letters denote the natural logarithm of the corresponding capital letters:

$$(17) \quad x_{it}(j) = \bar{x} + \xi_{it}(j) + \chi_{it} + \theta_i \zeta_t$$

$$(18) \quad m_t = \bar{m}t + \omega_t + \mu_t + \varpi \zeta_t$$

$$(19) \quad \ln(1 - T_t) = \bar{t} - \tau_t$$

The six independent shocks (in Greek letters with time subscripts) are: firm-specific productivity (ξ), sectoral productivity (χ), aggregate productivity with a sector-specific impact (ζ), anticipated monetary policy (ω), unanticipated monetary policy (μ) and anticipated tax changes (τ). In (18) note that the monetary policy rule responds systematically to productivity shocks. For simplicity, we assume that each shock follows an independent random walk, and that the variety-specific and sector-specific shocks approximately average to zero in each time period, $\int_0^1 \xi_{it}(j) dj \approx 0$ for all i and $N^{-1} \sum_{i=1}^N \chi_{it} \approx 0$.

We further assume that m_t grows over time at a rate \bar{m} that is large enough so that

$\delta E_t(M_t/M_{t+1}) < 1$, ensuring that the cash-in-advance constraint always binds.

We model price rigidity through imperfect information as in Robert E. Lucas Jr. (1973) and N. Gregory Mankiw and Reis (2002). In particular, we assume that at the beginning of every period, all firms learn about the past values of the six shocks, as well as the current values of the anticipated fiscal and monetary shocks. However, only a randomly drawn fraction ϕ_t of firms in each sector observe the contemporaneous realization of the other four shocks before making their pricing decisions. The remaining $1-\phi_t$ fraction of firms learn these shocks only in the following period, and we denote their expectations with this incomplete information by $\hat{E}(\cdot)$. This assumption of imperfect information has a long tradition in macroeconomics (Woodford, 2003, chapter 3, labels it the “neoclassical” case) and is in line with the recent work on sticky-information Phillips curves.

The web appendix solves for the equilibrium in this economy, showing that it can be reduced to the following equations for p_{it} , p_t , and y_t (ignoring constants):

$$(20) \quad p_{it} = \phi_i \left[p_t + \alpha(y_t - x_t) + \kappa\eta\tau_t - \kappa(x_{it} - x_t) \right] + (1 - \phi_i) \hat{E} \left[p_t + \alpha(y_t - x_t) + \kappa\eta\tau_t - \kappa(x_{it} - x_t) \right],$$

$$(21) \quad p_t = N^{-1} \sum_{i=1}^N p_{it},$$

$$(22) \quad m_t = p_t + y_t,$$

where α and κ are two positive parameters that depend on the preference and production parameters, $x_{it} = \int x_{it}(j) dj$ is sectoral productivity and $x_t = N^{-1} \sum_{i=1}^N x_{it}$ is average productivity. The first equation is the fundamental pricing equation in new Keynesian models, relating sectoral prices to marginal costs, which in this model depend on aggregate output, sectoral productivity, and taxes. The second equation is a log-linear approximation to

the static cost-of-living price index, which we denote by p_t . The third equation is the quantity theory relation that follows from the cash-in-advance constraint. This basic reduced-form structure is shared by many modern models of inflation dynamics.

A few steps of algebra show that sectoral price changes in this economy follow the same linear dynamic factor model in (5), $\boldsymbol{\pi}_t = \mathbf{I}v_t + \boldsymbol{\Theta}\boldsymbol{\rho}_t + \mathbf{u}_t$, that we will estimate in the data. The three components are:

$$(23) \quad v_t = \Delta\omega_t + \left(\frac{\kappa\eta}{\alpha}\right)\Delta\tau_t,$$

$$(24) \quad \begin{aligned} \boldsymbol{\Theta}_i\boldsymbol{\rho}_t = & \Delta\mu_{t-1} + \left(\frac{\alpha\phi_i}{1-(1-\alpha)\bar{\phi}}\right)\Delta^2\mu_t \\ & + (\varpi - \bar{\theta})\Delta\zeta_{t-1} - \phi_i \left(\frac{\alpha(\bar{\theta} - \varpi) + \bar{\phi}(1-\alpha)(1-\bar{\theta})(\alpha - \kappa)}{1-(1-\alpha)\bar{\phi}}\right)\Delta^2\zeta_t \\ & - \kappa(\theta_i - \bar{\theta})[\phi_i\Delta\zeta_t + (1-\phi_i)\Delta\zeta_{t-1}], \end{aligned}$$

$$(25) \quad u_{it} = -\kappa(\phi_i\Delta\chi_{it} + (1-\phi_i)\Delta\chi_{it-1}).$$

where $\bar{\theta}$ and $\bar{\phi}$ are the sectoral averages of θ_i and ϕ_i respectively, $\Delta = (1-L)$ is the first-difference operator, and $\boldsymbol{\Theta}_i$ on the left-hand side of (24) denotes the factor loading for π_{it} . Aggregate output in turn is:

$$(26) \quad y_t = \left[\frac{1-\bar{\phi}}{1-(1-\alpha)\bar{\phi}}\right]\Delta\mu_t - \left(\frac{\eta\kappa}{\alpha}\right)\tau_t + \bar{\theta}\zeta_{t-1} + \bar{\phi}\left[\frac{\alpha + \kappa(\bar{\theta} - 1)}{1-(1-\alpha)\bar{\phi}}\right]\Delta\zeta_t$$

B. Relation Between the Theory and the Estimates

We ask, in the model, the same two questions that we will ask in the data: what part of price changes is due to pure inflation, relative-price aggregate shocks, and

idiosyncratic shocks? And, which shocks account for the Phillips correlation between real activity and inflation? Answering these questions shock by shock (table 1 summarizes the results):

Anticipated money (ω): When money grows by 1 percent and all price-setters know about it, then all raise their prices by exactly 1 percent once and for all. No relative prices change and there is pure inflation. The quantity theory in turn implies that output is unchanged so there is no Phillips correlation. This is a result of the absence of money illusion in this model, as rational consumers and firms realize that nothing but units have changed, so there is no reason to change any real actions.

Unanticipated money (μ): In this case, only a fraction ϕ_i of the firms in each sector change their prices in response to a shock while the others remain inattentive. As a result, sectoral inflation is different across sectors, depending on the share of attentive firms in each sector, so there is an aggregate relative-price change. A monetary expansion raises output, because of the information stickiness of prices, and thus there is a Phillips correlation.

Aggregate productivity (ζ): This shock has a similar effect as an unanticipated money shock, both through its direct effect and through the systematic response of monetary policy. Either because it affects the productivity of different sectors differently, or because of different information stickiness across sectors, the shock induces a change in relative prices. Output increases when firms become more productive so there is a Phillips correlation.

Sectoral productivity (χ): These shocks map directly into the idiosyncratic shocks to relative prices in our measurement model. Because we assumed that an approximate law of large numbers holds, they do not affect output, so they do not generate a Phillips correlation.

Firm-specific productivity (ξ): These shocks are an example of something that the statistical model will miss. While these shocks induce relative-price changes within each sector, they wash out at the sectoral level so they do not affect relative sectoral prices.

Likewise, while they affect the allocation of production across firms, within each sector, they aggregate to zero on aggregate, so they neither move aggregate output nor prices. Whether these shocks are important or not cannot be answered without more disaggregated data.

Anticipated tax changes (τ): When taxes increase, the representative agent raises her wage demand. As this affects all firms equally, all raise their price in the same amount. Therefore, there is pure inflation. However, output falls as the return to work has fallen. The Phillips correlation results not from a change in the relative price of goods, but rather from a change in the relative price of labor versus consumption.

To conclude, our empirical estimates are informative about the role of different shocks in this economy. Our first empirical contribution, the estimates of the variability of overall inflation due to pure inflation, aggregate sources of relative prices, and idiosyncratic shocks, pins down the relative variance of anticipated versus unanticipated shocks and of aggregate versus sectoral shocks.⁷ Our second empirical contribution, whether the Phillips correlation is still present after we control for goods' relative prices, tells us whether this famous relation is due to monetary and productivity shocks via goods' relative prices as models of monopolistic competition emphasize, or whether it is due instead to shocks to other relative prices like real wages.

III. The Estimated Factor Model

A. The Data

The price data are monthly chained price indices for personal consumption expenditures by major type of product and expenditure from 1959:1 to 2006:6. Inflation is

⁷Our definition of pure inflation only allows us to gauge the relative variance of the unsystematic parts of money supply. To identify the policy rule, in this case the coefficient ϖ , requires more structure just as in the VAR literature on identifying monetary policy shocks. Also, as in that literature, a small role for pure inflation should not be confused for a small role for monetary policy.

measured in percentage points at an annual rate using the final month of the quarter prices: $\pi_{it} = 400 \times \ln(P_{it}/P_{it-1})$, where P_{it} are prices for March, June, September, and December.⁸ Prices are for goods at the highest available level of disaggregation that have data for the majority of dates, which gives 214 series. We then excluded series with unavailable observations (9 series), more than 20 quarters of 0 price changes (4 series), and series j if there was another series i such that $Cor(\pi_{it}, \pi_{jt}) > 0.99$ and $Cor(\Delta\pi_{it}, \Delta\pi_{jt}) > .99$ (14 series). This left $N = 187$ price series. Large outliers were evident in some of the inflation series, and these observations were replaced with centered 7-quarter local medians. A detailed description of the data and transformations are given in the web appendix.

As the economic model from section II makes clear, the level of aggregation across goods and time affects the interpretation of the estimated model. For example, as stressed in section II, the sectoral data provide no information about the relative prices of goods within a sector. The hope, therefore, is that the sectoral information is rich enough to capture important aggregate shocks. Furthermore, as with all models of information flows and discrete actions, the definition of the appropriate time period is important. The use of quarterly data means, for example, that equiproportional changes in all sectoral price indexes within the quarter are included in a_t , even if these changes occur at different times throughout the quarter. Said differently, the relative price factors, \mathbf{R}_t , capture only those relative prices changes that persist for at least one quarter. Because most macroeconomic analyses focusing on aggregate shocks use quarterly data, we are not departing from tradition.

One feature of these data is the constant introduction of new goods within each sector (Christian Broda and David Weinstein, 2007). Insofar as our statistical factor model of sectoral price changes remains a good description of their co-movement during the sample period, this should not affect our results. Another common concern with price data is the

⁸We considered using monthly, rather than quarterly, price changes, but found that the extra idiosyncratic error in monthly price changes outweighed the benefit of more observations.

need to re-weight prices to track expenditure shares and measure their effects on welfare. Our model in (1)-(5) does not require any expenditure shares, since the objective of measuring pure inflation is not to measure the cost of living, but rather to separate absolute from relative price changes.

B. *The Number of Factors and the Estimated Parameters*

Panel (a) of figure 1 shows the largest twenty eigenvalues of the sample correlation matrix of the inflation data. It is clear that there is one large eigenvalue, but it is much less clear how many additional factors are necessary. The Bai-Ng estimates confirm this uncertainty: their ICP_1 , ICP_2 and ICP_3 estimates are 2 factors, 1 factor, and 11 factors respectively. Panel (b) of figure 1 summarizes instead the fraction of variance explained by unrestricted factor models with 1 through 4 factors for each of the 187 inflation series.⁹ To make the figure easier to read, the series have been ordered by the fraction of variance explained by the 1-factor model. The uncertainty in the appropriate number of factors is evident here as well: the second factor improves the fit for several series, but it is unclear whether a third, fourth or fifth factor is necessary. In our benchmark model we will use 3 factors (a_t and two relative price factors in \mathbf{R}_t). We summarize the key results for other choices in section VI.¹⁰

We use the parametric factor estimates from (5)-(7) in our benchmark calculations; results using the principal components estimators are similar and are summarized in Section VI. The VAR for the factors in the benchmark specification uses 4 lags, guided by a few

⁹ These measures were computed as $R_i^2 = 1 - [\text{var}(u_i) / s_{\pi_i}^2]$, where $\text{var}(u_i)$ is the estimated variance of u_i implied by the estimated model and $s_{\pi_i}^2$ is the sample variance of π_i .

¹⁰ There is also uncertainty about the number of dynamic factors, which corresponds to the rank of the covariance matrix of $\boldsymbol{\varepsilon}$ in (8). The estimator developed in Bai and Ng (2007) indicates that the number of dynamic factors is the same as the number of static factors, while the estimators discussed in Dante Amengual and Watson (2007) and Hallin and Roman Liska (2007) suggest one dynamic factor. In our parametric model we will not constrain the rank of the covariance matrix of $\boldsymbol{\varepsilon}$.

diagnostic tests (not reported). It is well-known that inflation series are quite persistent and it is difficult to reject the null hypothesis that they have a unit root in the autoregressive representation (Pivetta and Reis, 2007). When we estimate the VAR in (6), we find that there are several large roots in $\Phi(L)$ and one that is very close to unity. In our benchmark model, we impose two unit roots in $\Phi(L)$; that is, a_t and one of the relative price factors are treated as I(1) processes. Results in which these unit roots are not imposed turn out to be very similar, and again, we summarize results for these models in section VI. Finally, we use only one lag in the univariate autoregressions of u_{it} , as suggested by diagnostic tests. The estimated AR(1) coefficients for u_{it} are typically small, suggesting I(0) variation in the idiosyncratic relative inflation rates.

Values for the estimated parameters for the benchmark model are given in the appendix.

C. The Unit Coefficient on a_t

Panel (a) of figure 2 summarizes the fit of unrestricted factor models that do not impose the unit restriction on the loading of the absolute-price factor. It shows that the increase in fit, measured by R^2 is less than 3 percent for 80 percent of the series. The median increase is less than 1 percent. The unrestricted model appears to fit appreciably better only for a small number of price series: for 10 series the increase in R^2 exceeds 10 percent. Panel (b) of figure 2 shows the ordered values of the estimates of ς_i from (10), that is the least-squares coefficient from regressing π_{it} on \hat{a}_t controlling for $\hat{\mathbf{R}}_t$. Most of the estimates are close to 1. Panel (c) shows the ordered values of the (4-lag Newey-West) t -statistic testing that $\varsigma_i = 1$. There are far more rejections of the restriction than would be expected by sampling error, with over 30 percent of the t -statistics above the standard 5 percent critical values and over 20 percent above the 1 percent critical values. These results suggest that, as

a formal matter, the unit factor loading restriction in (2) appears to be rejected by the data. That said, the results in panels (a) and (b) suggest that little is lost by imposing this restriction.

IV. Decomposing Sectoral Inflation

Figure 3 shows the historical decomposition of headline PCE inflation (top panel) and a representative sector, “major household appliances” (bottom panel). By construction, the pure-inflation (v) component is identical in the two plots (note the difference in the scales), while the idiosyncratic (u) component differ across goods, and the aggregate relative prices components (ρ) differs in its impact (Θ). Because $v_t = a_t - E[a_t | \{\mathbf{R}_\tau\}_{\tau=1}^T]$, we have plotted the data from 1965-1999 to eliminate uncertainty associated with pre-sample and post-sample values of \mathbf{R}_t .

Pure inflation is somewhat smoother than the other series and less volatile. The standard deviation of Δv_t is 0.3 percent, while the standard deviation of PCE inflation changes is 1.7 percent. Sectoral inflation is more volatile: the standard deviation for changes in “Major Household Appliances” is 4.1 percent, and the median across all 187 sectors is 5.9 percent. Evidently, aggregate relative price changes (ρ) explain much of the low-frequency variability in headline PCE. For example, much of the increases in inflation in the early 1970s and the declines in inflation in the 1990s were associated with changes in the relative price factor. That said, pure inflation (v) did account for over 2 percent of the increase in inflation from 1970-1980 and over 2 percent of the subsequent decline from 1980-1983.

A. The Relative Size of the Components

Table 2 shows the fraction of the variability of overall inflation associated with each

of its components, either averaged over all frequencies or just over business-cycle frequencies. The first row of the table uses the PCE deflator as the measure of overall inflation and shows that, at business-cycle frequencies, 15 percent of the movements in the series are accounted for by pure inflation, 76 percent is accounted for by the relative-price index, and the remainder is accounted for by the idiosyncratic sectoral shocks. The second and third row look at two other commonly used measures of overall inflation, the GDP deflator and the Consumer Price Index, and show similar results. The 2-dimensional relative-price index captures most of the variance in aggregate measures of inflation, while pure inflation plays a smaller but not negligible role. Including all frequencies, the role of pure inflation rises (with the exception of the CPI) while the relative-price index is significantly less prevalent.

These results have implications for macroeconomic models. For example, in terms of the model of section II, this 5-to-1 ratio in the relative variances of the relative-price index and pure inflation would say that a weighted average of the variance of anticipated shocks is significantly less volatile than an average of the unanticipated shocks. More generally, it is customary to compare the predictions of models with a single good for inflation with, for example, the data on the PCE deflator. The results in table 2 show that it is dangerous to do so since as much as 85 percent of the movements in the PCE deflator are driven by changes in the relative prices of different goods. For some questions, it might be better to compare the predictions of these models with our estimated series for pure inflation.

Two common approaches to strip relative-price movements from inflation are to exclude the prices of food and energy or to look at the median inflation across the different sectors. The next two rows in table 2 shows that these rough attempts at controlling for relative-price changes go in the right direction but remain quite far from excluding all relative-price changes. Core inflation is less closely tied to the relative-price index and more

related to pure inflation, but the squared coherences are still only slightly different, 69 percent and 21 percent respectively. For median CPI inflation, the idiosyncratic component is higher, but the two aggregate components are also only slightly lower.

The last section of table 2 summarizes the distribution of variance decompositions for the 187 sectoral inflation rates. Looking at the 25th and 75th quartile, the relative-price index accounts for between 15 percent and 42 percent of the overall variability of sectoral inflation rates, and pure inflation between 2 percent and 8 percent. As expected, the idiosyncratic relative-price shocks account for a much bigger share of sectoral price movements than they do for aggregate inflation measures. More remarkable, at the median, almost 1/3 of relative-price movements at the sectoral level are accounted for by the aggregate measures of pure inflation and the 2-dimensional relative-price index. Using sectoral price data, these findings confirm a result found for different macroeconomic datasets, countries, and time periods: a few aggregate factors (in our case three) can account for a large share of the variability in the economy (Stock and Watson, 1989, 2005, Forni et al, 2000).

B. Components of Inflation and other Observables

Table 3 compares the 2-dimensional index of relative prices with several conventional measures of relative-price changes. In the first row is the change in the price of durables relative to the headline PCE. The squared canonical coherence of this measure of relative prices with the relative price factors is high, around 0.5, but this single indicator falls short of capturing all of the variability in relative prices. The next two rows look at the relative prices of non-durables and services. The link between these and the two relative-price factors is higher, but they are still quite far from being a comprehensive indicator for relative-price changes. The next two rows show the relative price of food and energy,

popular measures of relative-price shocks in the macro literature. These are still statistically significant, but they perform significantly worse. In spite of the attention devoted to the price of energy, Table 3 suggests that it can account for only roughly one third of the relative-price shocks hitting the U.S. economy at business cycle frequencies.

Figure 4 illustrates these results by showing the projection of the change in the relative prices of services and energy onto 2 leads and lags of $\hat{\mathbf{R}}_t$, the estimated vector of relative price factors. For services, the regression's adjusted R^2 is 0.56, but for energy it falls to 0.22. Both series can deviate quite significantly from the relative-price index, but energy prices provide a particularly poor fit to the aggregate movements in relative prices.

Table 3 indicates that combining food and energy captures a larger share of the movements in relative prices, but still only comparable to the share accounted for by services. Finally, the resulting 4-dimensional index of relative prices (durables, nondurables, food and energy) can only account for at most 87 percent of the variability of relative prices captured by the two relative-price factors. These results suggest that, given its parsimony and comprehensiveness, the two relative-price factors estimated from the statistical model provide a useful summary of relative-price shocks in the U.S. economy.

The bottom panel of Table 3 investigates the correlation of pure inflation with measures of monetary policy and the term spread. Milton Friedman and Anna J. Schwartz (1963) famously observed that in the long run, money growth and inflation are tightly linked. Equally famously, Irving Fisher (1930) and many that followed showed that there is an almost as strong link between nominal interest rates and inflation in the long run. At business-cycle frequencies though, these correlations are much weaker. The correlation between money growth and inflation is unstable and typically low (Stock and Watson, 1999), while the correlation between inflation and nominal interest rates is typically higher, but well below its level at lower-frequencies (Frederic S. Mishkin, 1992). Panel b of Table 3 shows

the average squared coherence of pure inflation and measures of money growth (M0, M1, and M2) and different short-term nominal interest rates (the federal funds rate and the 3-month Treasury bill rate). The correlation between money growth and pure inflation is very close to zero for all measures. The correlation between nominal interest rates and pure inflation is significantly higher and statistically significant at conventional significance levels, especially at business-cycle frequencies. These correlations are much like correlations found by other researchers using overall measures of inflation. The final row of Table 3 shows the correlation of pure inflation with the term spread (the difference between to yield on 10-year Treasury bonds and 3-month Treasury bills), where the results look much like the results for short-term rates.

In terms of the model of section II, these estimates again provide useful information on the relative size of different shocks. Identifying some sectors in the model with services, non-durables, food or energy, the results in panel (a) of table 3 provide information on the relative size of the sectoral-specific productivity shocks. In turn, the results in panel (b) of the table indicate the relative weight of anticipated monetary shocks vis-à-vis unanticipated monetary shocks and fiscal shocks.

V. The Phillips Correlation

One of the most famous correlations in macroeconomics, due to Phillips (1958), relates changes in prices with measures of real activity. The first panel of table 4 displays the Phillips correlation using our measures of squared coherence. At business-cycle frequencies, measuring inflation with the PCE deflator and real activity with GDP, the average squared coherence (R^2) is 0.28, corresponding to a “correlation” of roughly 0.5. The Phillips correlations for industrial production, consumption, employment or the unemployment rate

are all similarly large.

The second and third panels in table 4 show that the usual controls for relative prices reduce the strength of this correlation. Controlling for intertemporal relative prices (using short-term interest rates and stock returns), for the relative price of labor and consumption (using real wages), or for the relative price of domestic and foreign goods (using the real exchange rate) cuts the Phillips correlations in approximately half. Still, these correlations remain quantitatively large and statistically significant.

The fourth and fifth panels in table 4 include instead two of the conventional measures of relative prices that we discussed in the previous section. Controlling for food and energy relative prices, the Phillips curve relation falls significantly, but the squared coherences remain sizeable and at least 0.10 for two of the five real series. Including all four relative-price indicators drives down the Phillips relation to between 0.03 and 0.08 (although with four relative price series included in the VAR used to estimate the coherences, one might conjecture that some of this decline is associated with over-fitting).

The last panel of table 4 introduces as controls instead the two relative price factors from the estimated model. Strikingly, controlling for ρ_t , the Phillips correlation disappears over business cycle frequencies. The largest squared coherence point estimate between PCE inflation and measures of real activity, controlling for our relative-price index, is 0.03 and the point estimates are statistically insignificant at the 10 percent level for all measures of real activity. Apparently, the empirical regularity that Phillips first brought attention to is essentially entirely explained by the two relative-price factors.

Table 5 provides a different perspective by decomposing the Phillips relation into the inflation components that we have separated. The first panel shows that removing the idiosyncratic sources of inflation variation makes the Phillips relation much stronger than it was with headline PCE inflation. At business-cycle frequencies, the squared coherence

between the aggregate components of inflation and measures of real activity is as high as 70 percent, and it is highly statistically significant at conventional significance levels. The second panel controls for the relative-price index, so it shows the squared coherence between pure inflation and measures of real activity. Again, controlling for relative prices essentially eliminates the Phillips correlation, with the squared coherences falling by a factor of roughly one-seventh. According to the model in section II, the little that remains of the relation between pure inflation and real activity could be due to omitted relative prices like wages. Panels (c) and (d) control for real wages, asset prices and exchange rates, which cuts the squared coherences a little further.

The results in these tables suggest that a large part of the Phillips correlation, that has puzzled macroeconomists for half a century, is explained by changes in good's relative prices. Changes in the unit of account, as captured by pure inflation, do not seem to affect real variables, consistent with anticipated money shocks accounting for most of pure inflation. However, note that a few of the estimates in table 5 are statistically significant, even if small, even after controlling for other relative prices. This suggests that some money illusion may be present, although it seems to explain very little of the variability of real activity.

VI. The Robustness of the Results

Table 6 investigates the robustness of the key empirical conclusions to four aspects of the model specification: (i) the number of estimated factors, (ii) the method for estimating the factors (signal extraction using the parametric factor model (5)-(7) versus principal components on (4)), (iii) the imposition of unit roots in the factor VAR for the parametric model, and (iv) the number of lags and imposition of unit roots in the VAR spectral estimator

used to compute the various coherence estimates. The table focuses on seven key results described below.

The first row of the table shows results for the benchmark model, where the first column provides details of the factor estimates and where “(1,1,0)” denotes a parametric $k=3$ factor model where the first and second factor are $I(1)$ processes and the third is $I(0)$. The next column, labeled “VAR”, summarizes the specification of the VAR used to compute the spectral estimates, which for the benchmark model involves 4 lags of (a_t, \mathbf{R}_t) with a_t and the first element of \mathbf{R}_t entered as first differences (D,4).

Results shown in the column labeled (1) are for the fraction of the business cycle variability of headline PCE inflation explained by the relative price factors (ρ) and pure inflation (v); for the benchmark model these are taken from the first row of Table 2. Results shown in the column labeled (2) are the average squared canonical coherences between the relative inflation factors \mathbf{R}_t and relative inflation rates for durables, nondurables, food, and energy (benchmark model from Table 3, panel a, final row). Columns (3) and (4) show the average squared coherence between pure inflation and the growth rate of M2 and the 3-month Treasury bill rate (benchmark model from Table 3, panel b, rows 3 and 5). Column (5) shows the average coherence between real GDP and headline PCE inflation after controlling for the estimated relative inflation factors (benchmark model, Table 4, panel f, first row). The final two columns show the fraction of business cycle variability explained by the factors (a_t, \mathbf{R}_t) (Column 6) and the fraction explained by pure inflation, v_t (benchmark models results from Table 5, row 1 of panels a and b).

Looking across the entries in the table, the key quantitative conclusions from Tables 2-5 appear to be robust to the changes in specification studied in Table 6. From column (1), relative inflation factors explain a significant fraction of the business cycle variability of aggregate inflation (as measured by the headline PCE deflator), while pure inflation (v)

explains a smaller, but non-negligible fraction. Observable measures of relative price inflation are reasonably highly correlated with one of the relative inflation factors (the first canonical coherence is roughly 0.90), but less highly correlated with the other factor (the other canonical coherences are generally less than 0.50). For all of the specifications the estimates of pure inflation are very weakly correlated with M2, but more highly correlated with nominal interest rates. Finally, in all of the specifications, controlling for the relative inflation factors essentially eliminates the correlation between PCE inflation and real GDP (column 5), and, while the estimated factors are highly correlated with real GDP (column 6), the pure inflation factor is very weakly correlated with real GDP (column 7).

VII. What Have we Done and Why Does it Matter?

In this paper, we decomposed the quarterly change in sectoral goods' prices into three components: pure inflation, an aggregate relative-price index, and idiosyncratic relative prices. We used different estimation techniques and specifications to estimate these components, proposed a simple method to compute their correlations with other macroeconomic variables, and presented a stylized structural model that showed how these components relate to different economic shocks.

Our first finding was that pure inflation, the relative-price index, and conventional measures of inflation, like the PCE deflator or its core version, can all differ markedly. Pure inflation is smoother and less volatile than the others, and much of the low-frequency swings in standard inflation measures are associated with changes in relative prices. More concretely, a large part of the increase in inflation in the early 1970s and the decrease in inflation in the 1990s was associated with changes in relative prices, while some of the increase in the late 1970s and the decrease in inflation in the early 1980s was associated with

changes in pure inflation.

Second, we found that aggregate shocks account for roughly 90 percent of the variability of aggregate inflation, and a still sizable 1/3 of the variability of sectoral inflation rates. Within aggregate shocks, the relative-price index dominates, but pure inflation is also quantitatively significant, accounting for 15-20 percent of the variability in inflation measured by conventional price indices, like the PCE deflator, the GDP deflator, or the CPI. This finding has at least two implications for the work of economic theorists building models to explain inflation. First, it shows that comparing the predictions of one-good models with common measures of inflation is flawed. Changes in the relative prices of goods are large enough that they can easily lead to mistakenly accepting or rejecting models that ignore this feature of the data. Second, our estimates provide statistics that can be used to calibrate the relative variances of anticipated versus unanticipated shocks, and aggregate versus sectoral shocks.

Our third finding was that conventional measures of relative-price inflation, such as the relative inflation of non-durables, food and energy, or combinations of several of them, all fall short of capturing most of the relative-price inflation in the data. Our 2-dimensional relative-price index provides a parsimonious yet comprehensive measure of relative-price inflation that we hope will be useful in other studies that either need to statistically control for relative-price changes, or that seek to provide economic models of the main sources of relative-price movements. Pure inflation is only partly related to monetary policy variables. The link to the growth rate in monetary aggregates is weak, but the correlation with nominal interest rates at business cycle frequencies is stronger (approximately 0.5).

Our most striking finding was perhaps that, once we controlled for the two relative price factors, the Phillips correlation became quantitatively insignificant. Therefore, the correlation between real quantity variables and nominal inflation variables observed in the

data can be accounted for by changes in goods' relative prices. This implies that models that break the classical dichotomy via nominal rigidities in good's price adjustment are likely more promising than models that rely on money illusion on the part of agents. Moreover, changes in the relative prices of labor and intertemporal prices were less successful in explaining the Phillips correlation, suggesting a less important role for rigidities in the labor and asset markets.

To conclude, the distinction between absolute and relative prices is a central one in economic theory. Models of inflation have strong predictions on the relative sizes of pure and relative-price inflation and on what accounts for the Philips correlation. However, separating absolute and relative-price movements is naturally difficult, since the two concepts themselves are more a fruit of thought experiments than something easily observed. As a result, there have been few systematic attempts to measure and separate them in the data. The goal of this paper was to make some progress on this decomposition and on understanding its effects. Our estimates are certainly not perfect. We hope, however, that they are sufficiently accurate that future research can look deeper into the time-series and the moments that we provide, and that by stating the challenges and putting forward a benchmark, we can motivate future research to come up with better estimators. Likewise, we are sure that our findings will not settle the debates around the Phillips correlation. Our more modest hope is that they offer a new perspective on how to bring data to bear on this long-standing question.

REFERENCES

- Altissimo, Filippo, Benoit Mojon, and Paolo Zaffaroni.** 2009. "Can Aggregation Explain the Persistence of Inflation?" *Journal of Monetary Economics*, 56(2): 231-241.
- Amengual, Dante and Mark W. Watson.** 2007. "Consistent Estimation of the Number of Dynamic Factors in a Large N and T Panel." *Journal of Business and Economic Statistics*, 25(1): 91-96.
- Amstad, Marlene and Simon M. Potter.** 2007. "Real Time Underlying Inflation Gauges for Monetary Policy Makers." Unpublished.
- Bai, Jushan.** 2003. "Inferential Theory for Factor Models of Large Dimensions." *Econometrica*, 71(1): 135-171.
- Bai, Jushan and Serena Ng.** 2002. "Determining the Number of Factors in Approximate Factor Models." *Econometrica*: 70(1): 191-221.
- Bai, Jushan and Serena Ng.** 2006. "Confidence Intervals for Diffusion Index Forecasts and Inference for Factor Augmented Regressions." *Econometrica*, 74(4): 1133-1150.
- Bai, Jushan and Serena Ng.** 2007. "Determining the Number of Primitive Shocks in Factor Models." *Journal of Business and Economic Statistics*, 25(1): 52-60.
- Bernanke, Ben , Jean Boivin and Piotr S. Elias.** 2005. "Measuring the Effects of Monetary Policy: A Factor-augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics*, 120(1): 387-422.
- Blanchard, Olivier J.** 2003. *Macroeconomics*. Upper Saddle River NJ: Prentice-Hall,
- Blanchard, Olivier J. and Nobuhiro Kiyotaki.** 1987. "Monopolistic Competition and the Effects of Aggregate Demand." *American Economic Review*, 77(4): 647-666.
- Boivin, Jean, Marc Giannoni and Ilian Mihov.** 2009. "Sticky Prices and Monetary Policy: Evidence from Disaggregated Data." *American Economic Review*, 99(1): 350-384.
- Broda, Christian and David Weinstein.** 2007. "Product Creation and Destruction:

- Evidence and Price Implications.” National Bureau of Economic Research Working Paper 13041.
- Bryan, Michael F. and Stephen G. Cecchetti.** 1993. “The Consumer Price Index as a Measure of Inflation.” Federal Reserve Bank of Cleveland Economic Review, 15-24.
- Bryan, Michael F., Stephen G. Cecchetti and Rodney L. Wiggins II.** 1997. “Efficient Inflation Estimation.” National Bureau of Economic Research Working Paper 6183.
- Cristadoro, Riccardo, Mario Forni, Lucrezia Reichlin and Giovanni Veronese.** 2005. “A Core Inflation Index for the Euro Area.” *Journal of Money, Credit and Banking*, 37(3): 539-560.
- Del Negro, Marco.** 2006. “Signal and Noise in Core PCE.” Unpublished.
- Doz, Catherine, Domenico Giannone and Lucrezia Reichlin.** 2008. “A Quasi Maximum Likelihood Approach for Large Dynamic Factor Models.” Unpublished.
- Fisher, Irving.** 1930. *The Theory of Interest*. New York: MacMillan.
- Forni, Mario, Marc Hallin, Marco Lippi, and Lucrezia Reichlin.** 2000. “The Generalized Dynamic Factor Model: Identification and Estimation.” *Review of Economics and Statistics*, 82(4): 540-554.
- Friedman, Milton and Anna J. Schwartz.** 1963. *A Monetary History of the United States*. Princeton: Princeton University Press.
- Hallin, Marc and Roman Liška.** 2007. “Determining the Number of Factors in the Generalized Dynamic Factor Model.” *Journal of the American Statistical Association*, 102: 603-617.
- Hume, David.** 1752. *Political Discourses*. Reprinted in “Essays, Moral, Political, and Literary” edited by Eugene F. Miller, Liberty Fund.
- Lucas, Robert E., Jr.** 1973. “Some International Evidence on Output-Inflation Trade-Offs.” *American Economic Review*, 63(3): 326-334.

- Mankiw, N. Gregory and Ricardo Reis.** 2002. “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve.” *Quarterly Journal of Economics*, 117(4): 1295-1328.
- Mishkin, Frederic S.** 1992. “Is the Fisher Effect for Real? A Reexamination of the Relationship between Inflation and Interest Rates.” *Journal of Monetary Economics*, 30(2): 195-215.
- Phillips, Alban W. H.** 1958. “The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957.” *Economica*, 25(2): 283-299.
- Pivetta, Frederic and Ricardo Reis.** 2007. “The Persistence of Inflation in the United States.” *Journal of Economic Dynamics and Control*, 31(4): 1326-1358.
- Quah, Danny and Thomas J. Sargent.** 1993. “A Dynamic Index Model for Large Cross-Sections.” In *Business Cycles, Indicators, and Forecasting*, ed. James H. Stock and Mark W. Watson, 161–200. Chicago: University of Chicago Press and NBER.
- Rao, C. Radhakrishna.** 1973. *Linear Statistical Inference and Its Applications*, 2nd Edition. New York: John Wiley and Sons.
- Reis, Ricardo and Mark W. Watson.** 2007. “Measuring Changes in the Value of the Numeraire.” Unpublished.
- Selvanathan, E. Anthony and D. S. Prasada Rao.** 1994. *Index Numbers: A Stochastic Approach*. Michigan: University of Michigan Press.
- Stock, James H. and Mark W. Watson.** 1989. “New Indexes of Leading and Coincident Indicators.” *NBER Macroeconomics Annual*, 4: 351-394.
- Stock, James H. and Mark W. Watson.** 1999. “Business Cycle Fluctuations in U.S. Macroeconomic Time Series.” In *Handbook of Macroeconomics*, ed. John B. Taylor and Michael Woodford. Elsevier.

Stock, James H. and Mark W. Watson. 2002. "Forecasting Using Principal Components from a Large Number of Predictors." *Journal of the American Statistical Association*, 97: 1167-1179.

Stock, James H. and Mark W. Watson. 2005. "Implications of Dynamic Factor Models for VAR Analysis." National Bureau of Economic Research Working Paper 11467.

Watson, Mark W. 1989. "Recursive Solution Methods for Dynamic Linear Rational Expectations Models." *Journal of Econometrics*, 41(1): 65-91.

Woodford, Michael. 2003. *Interest and Prices*. Princeton: Princeton University Press.

Table 1. Fundamental shocks, inflation components and the Phillips correlation

Fundamental shocks	Inflation component	Phillips correlation?
Anticipated money (ω)	v	No
Unanticipated money (μ)	ρ	Yes
Aggregate productivity (ζ)	ρ	Yes
Sectoral productivity (χ)	u	No
Firm-level productivity (ξ)	\emptyset	No
Anticipated tax changes (τ)	v	Yes

Table 2. Fraction of Variability of Inflation associated with Aggregate Components

Average squared coherence over frequencies (standard errors in parentheses)

Inflation measure	All frequencies		$\pi/32 \leq \omega \leq \pi/6$	
	ρ_t	v_t	ρ_t	v_t
Aggregate Inflation Rates				
Headline PCE	0.51 (0.05)	0.16 (0.04)	0.76 (0.10)	0.15 (0.07)
Headline GDP	0.35 (0.06)	0.21 (0.04)	0.71 (0.11)	0.15 (0.07)
Headline CPI	0.47 (0.04)	0.12 (0.03)	0.76 (0.09)	0.15 (0.06)
Core PCE	0.32 (0.05)	0.24 (0.05)	0.69 (0.11)	0.21 (0.09)
Median CPI	0.39 (0.08)	0.14 (0.04)	0.64 (0.13)	0.18 (0.08)
187 Sectoral Inflation Rates				
25 th Percentile	0.13	0.03	0.15	0.02
Median	0.19	0.05	0.25	0.05
75 th Percentile	0.25	0.07	0.42	0.08

Notes: PCE is the Personal Consumption Expenditures deflator, GDP is the Gross Domestic Product deflator, and CPI is the Consumer Price Index. Median CPI inflation is from the Federal Reserve Bank of Cleveland and these data are available for $t \geq 1967:2$. For the last three rows, we computed the fraction of variability explained by pure inflation for each of the 187 goods' series, and report the 25, 50, and 75 percent values.

Table 3. The Components of Inflation and Other Observables
Average squared canonical coherence over frequencies (standard errors in parenthesis)

Observable	Frequencies			
	All		$\pi/32 \leq \omega \leq \pi/6$	
<i>a. Relative-Price Index ρ_t</i>				
Durables	0.42 (0.06)		0.58 (0.09)	
Nondurables	0.47 (0.05)		0.72 (0.09)	
Services	0.48 (0.05)		0.75 (0.08)	
Food	0.20 (0.05)		0.55 (0.14)	
Energy	0.30 (0.05)		0.37 (0.11)	
Food, Energy	0.53 (0.04)	0.06 (0.03)	0.78 (0.08)	0.10 (0.08)
Durables, Nondurables, Food, Energy	0.62 (0.04)	0.25 (0.04)	0.87 (0.05)	0.42 (0.10)
<i>b. Pure Inflation v_t</i>				
M0	0.04 (0.02)		0.01 (0.02)	
M1	0.06 (0.03)		0.01 (0.02)	
M2	0.03 (0.02)		0.01 (0.02)	
Federal Funds Rate	0.11 (0.04)		0.27 (0.10)	
3-Month T-bill Rate	0.12 (0.03)		0.27 (0.12)	
Term Spread (10Y-3Month)	0.08 (0.04)		0.27 (0.11)	

Table 4. Fraction of Variability of Real Variables associated with PCE Inflation
Average squared coherence over frequencies (standard errors in parenthesis)

Real Variable	Frequencies	
	All	$\pi/32 \leq \omega \leq \pi/6$
<i>a. No Controls</i>		
GDP	0.11 (0.05)	0.28 (0.12)
Industrial Production	0.13 (0.06)	0.27 (0.14)
Consumption	0.15 (0.06)	0.28 (0.13)
Employment	0.19 (0.06)	0.32 (0.12)
Unemployment Rate	0.22 (0.07)	0.34 (0.15)
<i>b. Controls: Interest Rates, Stock Returns, Wages</i>		
GDP	0.09 (0.05)	0.14 (0.07)
Industrial Production	0.13 (0.05)	0.12 (0.05)
Consumption	0.07 (0.04)	0.12 (0.06)
Employment	0.15 (0.04)	0.24 (0.09)
Unemployment Rate	0.14 (0.04)	0.18 (0.07)
<i>c. Controls: Interest Rates, Stock Returns, Wages, Exchange Rates ($t \geq 1973$)</i>		
GDP	0.14 (0.05)	0.17 (0.08)
Industrial Production	0.15 (0.05)	0.14 (0.06)
Consumption	0.10 (0.05)	0.18 (0.08)
Employment	0.12 (0.04)	0.24 (0.10)
Unemployment Rate	0.13 (0.04)	0.20 (0.08)
<i>d. Controls: Relative Inflation Rates of Food and Energy</i>		
GDP	0.03 (0.02)	0.05 (0.04)
Industrial Production	0.07 (0.03)	0.08 (0.05)
Consumption	0.07 (0.03)	0.04 (0.04)
Employment	0.12 (0.04)	0.10 (0.06)
Unemployment Rate	0.10 (0.04)	0.12 (0.06)
<i>e. Controls: Relative Inflation Rates of Durable, Non-durables, Food and Energy</i>		
GDP	0.02 (0.02)	0.03 (0.03)
Industrial Production	0.04 (0.02)	0.04 (0.04)
Consumption	0.05 (0.02)	0.03 (0.03)
Employment	0.09 (0.03)	0.06 (0.04)
Unemployment Rate	0.07 (0.03)	0.08 (0.05)
<i>f. Controls: Relative-Price Index ρ_t</i>		
GDP	0.02 (0.02)	0.01 (0.02)
Industrial Production	0.03 (0.02)	0.01 (0.02)
Consumption	0.06 (0.03)	0.03 (0.02)
Employment	0.08 (0.03)	0.03 (0.03)
Unemployment Rate	0.08 (0.03)	0.03 (0.03)

Notes: The results in panel (c) use only data only from 1973 onwards because of data availability for the weighted U.S. real exchange rate series.

Table 5. Fraction of Variability of Real Variables associated with Inflation Components
Average squared coherence over frequencies (standard errors in parenthesis)

Real Variable	Frequencies	
	All	$\pi/32 \leq \omega \leq \pi/6$
<i>a. Aggregate Inflation Components, v_t and ρ_t</i>		
GDP	0.26 (0.05)	0.60 (0.10)
Industrial Production	0.28 (0.06)	0.59 (0.12)
Consumption	0.28 (0.06)	0.62 (0.11)
Employment	0.35 (0.05)	0.65 (0.10)
Unemployment Rate	0.42 (0.06)	0.70 (0.11)
<i>b. Pure Inflation v_t</i>		
GDP	0.05 (0.02)	0.09 (0.05)
Industrial Production	0.06 (0.02)	0.09 (0.06)
Consumption	0.08 (0.03)	0.08 (0.04)
Employment	0.07 (0.02)	0.12 (0.06)
Unemployment Rate	0.12 (0.03)	0.14 (0.07)
<i>c. Pure inflation v_t, control for Interest Rates, Stock Returns, Wages</i>		
GDP	0.04 (0.02)	0.05 (0.03)
Industrial Production	0.05 (0.02)	0.04 (0.02)
Consumption	0.05 (0.02)	0.06 (0.03)
Employment	0.06 (0.02)	0.10 (0.04)
Unemployment Rate	0.12 (0.03)	0.07 (0.03)
<i>d. Pure inflation v_t, control for Interest Rates, Stock Returns, Wages, Exchange Rates</i> ($t \geq 1973$)		
GDP	0.03 (0.02)	0.07 (0.04)
Industrial Production	0.04 (0.02)	0.04 (0.03)
Consumption	0.04 (0.02)	0.06 (0.04)
Employment	0.06 (0.02)	0.16 (0.07)
Unemployment Rate	0.10 (0.02)	0.09 (0.05)

Notes: The results in panel (d) use only data only from 1973 onwards because of data availability for the weighted U.S. real exchange rate series.


Table 6: The Robustness of the Conclusions
Average squared coherences over business cycle frequencies (standard errors in parenthesis)

Factor	VAR	(1)	(2)	(3)	(4)	(5)	(6)	(7)			
Estimates											
<i>Benchmark Model</i>											
(1,1,0)	D,4	0.76 (0.10)	0.15 (0.07)	0.87 (0.05)	0.42 (0.10)	0.01 (0.02)	0.27 (0.12)	0.01 (0.02)	0.60 (0.10)	0.09 (0.05)	
<i>Alternative Parametric Factor Estimates</i>											
(1,1,0)	D,6	0.85 (0.06)	0.11 (0.04)	0.90 (0.04)	0.37 (0.12)	0.06 (0.07)	0.20 (0.09)	0.07 (0.05)	0.67 (0.08)	0.11 (0.06)	
(0,0,0)	D,4	0.74 (0.10)	0.16 (0.08)	0.87 (0.04)	0.42 (0.10)	0.01 (0.02)	0.30 (0.11)	0.02 (0.02)	0.61 (0.10)	0.09 (0.06)	
(0,0,0)	D,6	0.84 (0.06)	0.12 (0.05)	0.90 (0.04)	0.36 (0.12)	0.05 (0.07)	0.20 (0.09)	0.06 (0.04)	0.68 (0.08)	0.11 (0.06)	
(1,1,0,0)	D,4	0.80 (0.08)	0.12 (0.06)	0.89 (0.04)	0.56 (0.08)	0.14 (0.07)	0.01 (0.01)	0.21 (0.10)	0.01 (0.01)	0.62 (0.10)	0.08 (0.05)
(1,1,0,0)	D,6	0.87 (0.05)	0.09 (0.03)	0.92 (0.03)	0.50 (0.12)	0.13 (0.09)	0.05 (0.06)	0.15 (0.09)	0.06 (0.04)	0.70 (0.08)	0.12 (0.06)
(1,1,0)	L,4	0.80 (0.08)	0.11 (0.05)	0.87 (0.05)	0.53 (0.11)		0.02 (0.03)	0.20 (0.11)	0.01 (0.01)	0.62 (0.10)	0.12 (0.07)
(1,1,0)	L,6	0.86 (0.05)	0.09 (0.04)	0.90 (0.05)	0.50 (0.12)		0.03 (0.05)	0.20 (0.11)	0.05 (0.04)	0.71 (0.09)	0.17 (0.08)
(0,0,0)	L,4	0.79 (0.08)	0.11 (0.05)	0.88 (0.05)	0.52 (0.12)		0.02 (0.03)	0.19 (0.11)	0.00 (0.01)	0.62 (0.10)	0.12 (0.07)
(0,0,0)	L,6	0.85 (0.05)	0.09 (0.04)	0.90 (0.04)	0.49 (0.12)		0.03 (0.05)	0.19 (0.11)	0.05 (0.04)	0.72 (0.08)	0.16 (0.08)
(1,1,0,0)	L,4	0.82 (0.07)	0.10 (0.05)	0.90 (0.04)	0.61 (0.09)	0.15 (0.07)	0.02 (0.03)	0.24 (0.10)	0.01 (0.01)	0.64 (0.09)	0.11 (0.06)
(1,1,0,0)	L,6	0.87 (0.05)	0.08 (0.04)	0.92 (0.03)	0.56 (0.11)	0.16 (0.09)	0.03 (0.04)	0.22 (0.10)	0.05 (0.04)	0.72 (0.08)	0.16 (0.07)
<i>Using Principal Component Factor Estimates</i>											
PC-3	D,4	0.70 (0.11)	0.19 (0.08)	0.82 (0.07)	0.37 (0.11)		0.00 (0.01)	0.30 (0.12)	0.01 (0.01)	0.53 (0.11)	0.05 (0.04)
PC-3	D,6	0.80 (0.07)	0.15 (0.06)	0.86 (0.05)	0.40 (0.11)		0.06 (0.06)	0.28 (0.10)	0.04 (0.03)	0.64 (0.10)	0.06 (0.05)
PC-4	D,4	0.71 (0.10)	0.19 (0.08)	0.85 (0.05)	0.50 (0.09)	0.10 (0.06)	0.01 (0.02)	0.36 (0.10)	0.01 (0.02)	0.55 (0.11)	0.05 (0.04)
PC-4	D,6	0.80 (0.07)	0.14 (0.05)	0.87 (0.04)	0.46 (0.11)	0.09 (0.07)	0.08 (0.06)	0.33 (0.09)	0.04 (0.03)	0.68 (0.08)	0.08 (0.04)
PC-2	D,4	0.69 (0.11)	0.17 (0.07)	0.57 (0.13)			0.00 (0.01)	0.30 (0.13)	0.01 (0.01)	0.40 (0.12)	0.03 (0.03)
PC-2	D,6	0.81 (0.07)	0.13 (0.05)	0.74 (0.09)			0.06 (0.06)	0.29 (0.10)	0.05 (0.05)	0.49 (0.12)	0.07 (0.06)
PC-3	L,4	0.81 (0.07)	0.07 (0.04)	0.86 (0.06)	0.50 (0.11)		0.01 (0.01)	0.15 (0.09)	0.00 (0.00)	0.56 (0.11)	0.04 (0.04)
PC-3	L,6	0.87 (0.05)	0.06 (0.03)	0.89 (0.05)	0.44 (0.11)		0.05 (0.07)	0.16 (0.09)	0.04 (0.04)	0.65 (0.10)	0.07 (0.04)
PC-4	L,4	0.82 (0.07)	0.08 (0.04)	0.89 (0.04)	0.58 (0.10)	0.11 (0.06)	0.00 (0.01)	0.26 (0.10)	0.00 (0.00)	0.61 (0.10)	0.05 (0.04)
PC-4	L,6	0.87 (0.04)	0.06 (0.03)	0.90 (0.04)	0.47 (0.11)	0.09 (0.06)	0.07 (0.07)	0.24 (0.10)	0.02 (0.02)	0.71 (0.09)	0.09 (0.05)
PC-2	L,4	0.78 (0.08)	0.07 (0.04)	0.78 (0.08)			0.02 (0.04)	0.18 (0.10)	0.01 (0.01)	0.40 (0.13)	0.02 (0.03)
PC-2	L,6	0.85 (0.06)	0.07 (0.03)	0.81 (0.08)			0.04 (0.06)	0.21 (0.11)	0.04 (0.05)	0.50 (0.12)	0.08 (0.06)


Notes: The first column describes the factor estimates, where the parametric estimates are based on signal extraction applied to (5)-(7) with parameters estimated by Gaussian MLE, and the numbers in parenthesis indicate the number of factors and whether the relevant factor is modeled as an I(1) or I(0) process. For example, “(1,1,0)” is a three factor model modeled as I(1), I(1), and I(0) processes. PC- k denotes a k -factor model estimated by principal components. The column labeled VAR shows the specification of the VAR used to compute the VAR-spectral estimates, where “D” and “L” denote first differences and levels specifications and the numbers 4 and 6 denote the number of lags in the VAR. Results shown in the column (1) are the average squared coherences between Headline PC and \mathbf{p}_t and v_t (benchmark model results from Table 2 first row); (2) canonical coherences between the relative prices of (Durables, Nondurables, Food, Energy) and \mathbf{R}_t (benchmark model results from Table 3 row 7); (3) coherence between M2 and v (Table 3 row 10); (4) coherence between Federal Funds Rate and v_t (Table 3 row 11); (5) coherence between PCE Inflation and GDP controlling for \mathbf{R}_t (Table 4, panel f, row 1); (6) coherence between $v_t + \mathbf{\Omega p}_t$ and GDP (Table 5, panel a, row 1), and (7) between v_t and GDP (Table 5, panel b, row 1).

Figure 1. Choosing the number of factors

Panel A. Eigenvalues of the correlation matrix




Panel B. Contribution of more factors to the R^2 of each good




Notes: Panel a) shows the eigenvalues of the $N \times N$ sample correlation matrix of inflation rates. Panel b) shows the fraction of sample variance of inflation explained by k factors, where k varies from $k=1$ to $k=4$.

Figure 2. Comparison with unrestricted factor model


Panel A. Increase in R^2 from moving to unrestricted model



Panel B. Estimates of ζ_i , the coefficient on the absolute-price component




Panel C. Individual t-statistics for hypothesis $\zeta_i=1$



Notes: The horizontal axis in each panel goes from $i = 1$ to $i = 187$. In each panel, the goods are organized in increasing order.

Figure 3. Estimates of inflation and its components

Panel A. Headline PCE inflation



Panel B. Major household appliances inflation





Figure 4. The Relative Price Index and Other Observables

Panel A. Relative services inflation (thick line) and projection onto ρ (thin line).



Panel B. Relative energy inflation (thick line) and projection onto ρ (thin line).

