



Federal Reserve Bank of Cleveland Working Paper Series

Post-COVID Inflation Dynamics: Higher for Longer

Randal J. Verbrugge and Saeed Zaman

Working Paper No. 23-06

January 2023

Suggested citation: Verbrugge, Randal J. and Saeed Zaman. 2023. "Post-COVID Inflation Dynamics: Higher for Longer." Working Paper No. 23-06. Federal Reserve Bank of Cleveland.
<https://doi.org/10.26509/frbc-wp-202306>.

Federal Reserve Bank of Cleveland Working Paper Series

ISSN: 2573-7953

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications.

See more working papers at: www.clevelandfed.org/research. Subscribe to email alerts to be notified when a new working paper is posted at: <https://www.clevelandfed.org/subscriptions>.

Post-COVID Inflation Dynamics: Higher for Longer

Randal Verbrugge

Federal Reserve Bank of Cleveland and NBER/CRIW, 1455 E. 6th St., Cleveland, OH 44114

Email: randal.verbrugge@clev.frb.org.

Saeed Zaman

Federal Reserve Bank of Cleveland, 1455 E. 6th St., Cleveland, OH 44114

Email: saeed.zaman@clev.frb.org.

January 13, 2023

Abstract

In the December 2022 Summary of Economic Projections (SEP), the median projection for four-quarter core PCE inflation in the fourth quarter of 2025 is 2.1 percent. This same SEP has unemployment rising by nine-tenths, to 4.6 percent, by the end of 2023. We assess the plausibility of this projection using a specific nonlinear model that embeds an empirically successful nonlinear Phillips curve specification into a structural model, identifying it via an underutilized data-dependent method. We model core PCE inflation using three components that align with those noted by Chair Powell in his December 14, 2022, press conference: housing, core goods, and core-services-less-housing. Our model projects that conditional on the SEP unemployment rate path and a rapid deceleration of core goods prices, core PCE inflation moderates to only 2.75 percent by the end of 2025: inflation will be higher for longer. A deep recession would be necessary to achieve the SEP's projected inflation path. A simple reduced-form welfare analysis, which abstracts from any danger of inflation expectations becoming unanchored, suggests that such a recession would not be optimal.

JEL codes: E31, E32, E52, C32

Keywords: nonlinear Phillips curve, frequency decomposition, supply price pressures, structural VAR, core PCE inflation components, welfare analysis

Acknowledgments: We thank Rick Ashley, David Berger, Mark Bills, Mark Bognanni, Andrew Figura, Ed Knotek, David Lebow, James Mitchell, Norm Morin, Michael Palumbo, Matthias Paustian, Katia Peneva, Juan Rubio-Ramirez, Jeremy Rudd, Arsenios Skaperdas, Brad Strum, and our other research colleagues at the Federal Reserve Bank of Cleveland, and seminar participants at the Board of Governors, for excellent comments that helped to greatly improve this paper. Any errors are our own. The views stated herein are those of the authors and not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

“...you can break inflation down into three sorts of buckets. The first is goods inflation, and we see now... goods inflation coming down ... Then you go to housing services ...that inflation will come down sometime next year. The third piece, which is something like 55 percent of the ... PCE core inflation index, is non-housing-related core services. And that's really a function of the labor market ... And we do see a very, very strong labor market, one where we haven't seen much softening, where job growth is very high, where wages are very high. Vacancies are quite elevated and ... there's an imbalance in the labor market between supply and demand. So that part of it, which is the biggest part, is likely to take a substantial period to get down. The other ... the goods inflation has turned pretty quickly now after not turning at all for a year and a half. Now it seems to be turning. But there's an expectation ... that the services inflation will not move down so quickly, so that we'll have to stay at it so that we may have to raise rates higher to get to where we want to go. And that's really why we are writing down those high rates and why we're expecting that they'll have to remain high for a time.”

FOMC Chair Jerome Powell, Press Conference, Dec. 14, 2022

1. Introduction

In his December 14, 2022, press conference, Jerome Powell, chair of the Federal Open Market Committee (FOMC), used a tripartite decomposition of core PCE inflation to explain why the FOMC expects that the federal funds rate will “have to remain high for a time.” This decomposition consists of core goods inflation, housing services inflation, and non-housing core services. In the press conference, Chair Powell noted that core goods inflation has “turned down pretty quickly” in recent months. He further noted that housing services inflation is expected to come down sometime in 2023. Finally, he noted that non-housing core services inflation is influenced by the “very, very strong labor market” and, for this reason, is “likely to take a substantial period to get down.”

In this paper, we use this tripartite decomposition of core PCE inflation to explore the path of inflation going forward.¹ We build upon recent work by Ashley and Verbrugge (2022a) and Verbrugge and Zaman (2023) – two studies providing compelling evidence in favor of a nonlinear Phillips curve – to construct a nonlinear structural vector autoregression (SVAR), a model suitable

¹ The idea of forecasting aggregate inflation by separately modeling and forecasting its underlying disaggregated components has a long tradition; see Tallman and Zaman (2017) and references therein.

for exploring counterfactual conditional inflation forecasts. We estimate our model over the 1985-2019 period and identify it using the data-determined method of Swanson and Granger (1997), which substantially reduces the role of subjective elements. When we feed in the December SEP's forecasted path for the unemployment rate (which has it increasing by 0.9 percentage point) into our model, we get a higher path for core PCE inflation than the SEP path that has core PCE inflation moderating to 2.1 percent by the end of 2025. Inflation is going to remain higher for longer: rather than core PCE inflation reaching 2.1 percent by the end of 2025, our model projects that it will be at 2.8 percent, with the 70 percent confidence interval spanning 2.4 to 3.2 percent. A key to this result is the fact that inflation is more persistent than commonly believed. We conclude that it would take a deep recession to reduce inflation faster. We investigate the claim of former Treasury Secretary Lawrence Summers (reported in Aldrick, 2022) and the supporting assessment of Ball, Leigh, and Mishra (2022) that it will require two years of 7.5 percent unemployment from its current low level of 3.6 to 3.7 percent to bring inflation down to its 2 percent target. We find that one year of 7.4 percent unemployment would accomplish this task.

But would such a recession be ideal? As a first pass at addressing this question, we perform a simple reduced-form welfare analysis using a quadratic loss function that equally penalizes quarterly deviations of inflation from 2 percent (the FOMC target level of inflation), and deviations of unemployment from 4 percent (the FOMC's estimate of the longer-run level of unemployment). In addition to producing inflation forecasts corresponding to the deep recession noted above and to the December SEP, we produce inflation forecasts corresponding to a moderate recession (defined by the path of unemployment taken in the 2001 recession) and to a soft landing for unemployment (which we define as the path of unemployment reported in the June SEP).² This welfare analysis generally prefers the December SEP, unless the weight on inflation is very low (in which case, it prefers the soft landing) or very high (in which case, it prefers the moderate recession). Importantly, this analysis abstracts from any danger of the unanchoring of inflation expectations that might be associated with inflation still being at 2.8 percent three years from now.

² Figura and Waller (2022) argue that a soft landing in the labor market is a plausible scenario.

2. Data, Methods and Model

2.1 Data

We use quarterly data spanning from 1985:Q1 through 2022:Q4, though we estimate the model using pre-COVID data.³ Most of the series are available at a monthly frequency, and we aggregate them up to a quarterly frequency. Following much precedent in the literature, we focus attention on the post-1984 period because inflation dynamics are thought to have changed markedly beginning in the mid-1980s onward, and this is the period associated with anchored inflation expectations.

Our model consists of six variables. The first is the PPI for core intermediate goods, denoted *PPI*. Verbrugge and Zaman (2023) find that PPI captures supply price pressures and is an important determinant of trimmed-mean PCE inflation. The next three variables are also inflation-specific, corresponding to the tripartite decomposition of Chair Powell. These include core goods and housing services. But rather than using non-housing core services, we instead construct, and use, *median* non-housing core services.⁴ We do this because non-housing core services are quite sensitive to outliers, particularly in non-market services. Verbrugge (2022) demonstrates that such sensitivity renders core inflation measures less reliable as indicators of trend inflation. Accordingly, we view median non-housing core services inflation as a more accurate estimate of the trend in non-housing core services, helping to more reliably capture both the persistence of non-housing core services and their sensitivity to labor market pressures.⁵ Figure 1 plots non-housing core services inflation alongside its “median” counterpart. As expected, the median series is smoother than the original series. Over the sample period displayed, the bias, defined as the gap between their respective inflation rates, is zero. However, over specific periods, there can be

³ At the time of this writing, we do not have complete 2022Q4 data. We use available monthly data to construct Q4 nowcasts for all variables. Our model is estimated using 1985-2019 data, but data from 2021Q3 onward are used for forecasting because the model specification includes five lags.

⁴Our choice of “median” variable is partly motivated by the successful track record of median CPI and median PCE variables constructed by the Federal Reserve Bank of Cleveland in tracking the trend in CPI and PCE inflation, respectively. The method used to construct the median series is similar to that of Carroll and Verbrugge (2019), who use all the available 190+ disaggregated price categories of the monthly PCE to construct the (weighted) median PCE series. To construct the (weighted) median non-housing core services, we use information about the price changes in the 82 disaggregated price categories of the PCE that are part of “PCE services excluding energy, food, and housing,” along with their respective nominal expenditure shares at a monthly frequency. Since we estimate the model with quarterly data, we aggregate up the monthly data to a quarterly frequency.

⁵ As with non-housing core services inflation, we find that median non-housing core services inflation has a statistically significant Phillips curve, but a weak relationship with nominal wage inflation (as measured either by average hourly earnings or by the employment cost index).

notable divergence, with more recent periods appearing as a prominent example. Accordingly, in computing forecasts of non-housing core services inflation, we apply bias adjustment to the forecasts of the *median* variable.⁶




Figure 1: Core services ex. housing inflation indicators

Following Ashley and Verbrugge (2022a) and Verbrugge and Zaman (2023), the final two variables are two “components” of the unemployment rate: a persistent (or low-frequency) gap component and a moderately persistent (or medium-frequency) component.⁷ The approach to filtering is described in Appendix A. Following Verbrugge and Zaman (2023), these components of the unemployment rate are derived from the *jobless* unemployment rate of Hall

⁶ The bias-adjustment procedure is informed by estimating two separate AR(1) processes on the historical wedge (i.e., the gap between the two series) and using the estimated processes to compute the estimates of the time-varying wedge over the forecast period. One of the AR(1) processes is estimated over the entire sample, based upon a 12-month moving average of the monthly series; the other one is estimated over the post-1985 sample (with an intercept change in 2010), based upon a 3-month moving average of the monthly series, resulting in two forecasts of the time-varying wedge that are averaged to construct a single series of the wedge. Forecasts of the median variable are then bias-adjusted using this forecast of the wedge, so as to obtain an unbiased forecast of core services less housing.

⁷ Specifically, the unemployment rate is split into “transient,” “moderately persistent,” and “persistent” components. But since the transient fluctuations were found to be unimportant predictors, to keep our model parsimonious, we abstract from these fluctuations. “Moderately persistent” refers to fluctuations that take 1-4 years to complete; “persistent” fluctuations last longer than that. To obtain valid inferences, frequency filtering must be done in a one-sided manner (see Ashley and Verbrugge 2022b). Hamilton (2018) recently introduced an alternative to HP filtering, but Ashley and Verbrugge (2022c) demonstrate that, for properly decomposing a time series into its lower-frequency and higher-frequency components, this procedure is inferior to the procedure used in Ashley and Verbrugge (2022a) and Ashley, Tsang, and Verbrugge (2020); see Appendix A for more details. We form a low-frequency *gap* by subtracting the Zaman (2022) U_t^* estimate from the low-frequency component. Our model forecasts even slower deceleration in inflation if we instead use the CBO natural rate estimate. The U^* estimates from Zaman’s model are available to download from <https://github.com/zamansaeed/macrostars>.

and Kudlyak (2022).⁸ The jobless unemployment rate is constructed by removing the temporary layoffs from overall unemployment. We relate inflation to the jobless unemployment rate rather than the overall unemployment rate, since during the pandemic collapse, temporary unemployment experienced a 20-standard-deviation shock. Such an extreme movement severely distorts coefficient estimates and frequency partitions. Even very modest nonlinearities in relationships are likely to dominate the comovement of variables for as long as temporary unemployment remains extremely elevated, and these data points will have extremely high leverage. Putting this differently, it seems likely that the ordinary relationship between overall unemployment and inflation would have broken down in the face of this extreme movement. Our approach is to sidestep these twin problems by a) focusing on the relationship of inflation to the jobless unemployment rate, since the jobless unemployment rate experienced fairly typical dynamics during the COVID recession, and b) by estimating the model over the 1985-2019 period.

These two components of the jobless unemployment rate are depicted in Figure 2. Our partitioning of the jobless unemployment rate into varying persistence components is motivated by the aforementioned previous findings of persistence-dependence in the Phillips curve relationship and by an emerging literature that is re-exploring the frequency domain to obtain clues about business cycle drivers and dynamics. In contrast to the previous work, which modeled the relationship between aggregate inflation (i.e., trimmed-mean PCE inflation), this paper separately models the nonlinear Phillips curve relationship for each of the inflation components using the two components of the unemployment rate. Accordingly, in our inflation equations, we admit sign asymmetry on the unemployment components. We find that each inflation variable is related only to the *negative* part of the persistent unemployment gap (i.e., when the persistent unemployment rate is below the natural rate of unemployment), and to the *positive* part of the moderately persistent unemployment component, which is generally consistent with the previous work focusing on aggregate inflation. Historically, these portions of the two components align closely with overheating and recession, respectively. As we explain below, this simple partition allows us to uncover very insightful nonlinear Phillips curve relationships in all of our inflation variables.

Because we are specifying a structural model, we accordingly specify and estimate an equation for each of these unemployment components separately.

⁸ The data necessary to construct the jobless unemployment rate are available from the Bureau of Labor Statistics.

2.2 Methods

Identification is achieved using the data-determined method of Swanson and Granger (1997), which substantially reduces the role of subjective elements. This method, which builds upon the correlation structure of the reduced-form residuals, is briefly discussed below and explained in Appendix B. We generate conditional forecasts by constructing nonlinear system forecasts that condition upon labor market variables and (as we explain below) upon structural shocks that allow us to impose near-term information about core goods inflation and housing services inflation. As has been long-established in the forecasting literature, overall forecast accuracy can be enhanced by conditioning upon near-term information (see, e.g., Faust and Wright, 2013; Tallman and Zaman, 2020).




Figure 2: Two most persistent components of the jobless unemployment rate

2.3 Specification

Our specification largely follows Verbrugge and Zaman (2023). We are ultimately interested in reliable forecasts, so model parsimony was a chief consideration. We used step-down testing, equation by equation, removing variable lags to obtain parsimonious equations that were favored by the Bayesian information criterion (BIC). In all inflation equations, we allowed for sign

asymmetry in the two unemployment rate components, but did not impose it. In each equation, we allow up to 5 quarterly lags in the dependent variable and up to 4 quarterly lags in each of the other variables. Allowing for the fifth lag is quite important for accurately assessing the persistence of each series, as demonstrated in Verbrugge and Zaman (2023).

In the PPI equation, the inclusion of all other inflation series was rejected. However, PPI has a significant Phillips curve relationship. We rejected sign asymmetry in both unemployment rate components. Subsequently, both components appeared to enter as first differences. We thus entered both as first differences, and this yielded an equation that fit the data almost equally well; furthermore, u^{lowgap} was no longer statistically significant. Dropping this term yielded a more parsimonious equation with almost no decline in fit, and so was favored by the BIC.

$$\pi_t^{PPI} = \alpha^{PPI} + \sum_{j=1}^4 \beta_j^{PPI} \pi_{t-j}^{PPI} + \delta \Delta u_{t-1}^{medfreq} + e_t^{PPI} \quad (1)$$

In Equation (1) and hereafter, π_t^{PPI} refers to 4-quarter (4Q) inflation in the PPI. Similarly, π_t^{CoreG} refers to 4Q inflation in core goods, $\pi_t^{MNHServ}$ refers to 4Q inflation in median non-housing core services, and π_t^{Hous} refers to 4Q inflation in housing services. Labor market variables are as follows: $\Delta u_t^{medfreq}$ refers to the 1-quarter change in the medium-frequency component, $u_t^{+medfreq}$ refers to the positive portion of the medium-frequency component, and $u_t^{-lowgap}$ refers to the negative portion of the low-frequency gap.

The core PCE component inflation rate equations are specified as

$$\begin{aligned} \pi_t^{CoreG} = & \alpha^{CoreG} + \phi_1^{CoreG} \pi_{t-1}^{CoreG} + \phi_2^{CoreG} \pi_{t-2}^{CoreG} + \phi_5^{CoreG} \pi_{t-5}^{CoreG} + \\ & + \beta_1^{CoreG} \pi_{t-1}^{PPI} + \beta_2^{CoreG} \pi_{t-2}^{PPI} + \lambda^{CoreG} u_{t-4}^{+medfreq} + \psi I^{1995} + e_t^{CoreG} \end{aligned} \quad (2)$$

$$\begin{aligned} \pi_t^{MNHServ} = & \alpha^{MNHServ} + \gamma_1^{MNHServ} \pi_t^{MNHServ} + \gamma_2^{MNHServ} \pi_t^{MNHServ} + \gamma_5^{MNHServ} \pi_t^{MNHServ} + \\ & + \lambda^{MNHServ} u_{t-1}^{+medfreq} + \mu^{MNHServ} u_{t-1}^{-lowgap} + e_t^{MNHServ} \end{aligned} \quad (3)$$

$$\pi_t^{Hous} = \alpha^{Hous} + \sum_{j=1}^5 \eta_j^{Hous} \pi_{t-j}^{Hous} + \lambda^{Hous} u_{t-1}^{+medfreq} + \mu^{Hous} u_{t-4}^{-lowgap} + e_t^{Hous} \quad (4)$$

where I^{1995} is a dummy variable that is 1 prior to 1995Q1. This variable allows us to capture an evident mean shift in core goods inflation in the mid-1990s; see Clark (2004).

Finally, our $u^{medfreq}$ equation was specified as

$$u_t^{medfreq} = \sum_{j=1}^2 \lambda_j^{med} u_{t-j}^{medfreq} + \sum_{j=1}^4 \mu_j^{med} u_{t-j}^{lowgap} + \beta^{med} \pi_{t-1}^{PPI} + e_t^{medfreq} \quad (5)$$

and our u^{lowgap} equation was specified as

$$u_t^{lowgap} = \alpha^{lowgap} + \sum_{j=1}^2 \mu_j^{low} u_{t-j}^{lowgap} + \sum_{j=1}^4 \lambda_j^{low} u_{t-j}^{medfreq} + \sum_{j=1}^4 \beta_j^{low} \pi_{t-j}^{PPI} + e_t^{lowgap} \quad (6)$$

Given the model's nonlinear nature, we construct forecasts and error bands via counterfactual simulations, following the procedure outlined in Kilian and Lütkepohl (2017), with shocks bootstrapped from estimated residuals. We compute the median response as well as the 15th and 85th percentiles from the simulations.

The forecast of core PCE inflation at time t for h quarters ahead is simply the composite forecast of the core goods inflation forecast, housing services inflation forecast, and the median services ex. housing inflation forecast (which is our proxy for the core services ex. housing forecast), combined using the share weights available as of time t . The weights reflect the relative shares of core goods inflation, housing services inflation, and core services ex. housing inflation in the overall core PCE inflation. Specifically, the weight for core goods inflation is computed as a nominal share of the personal consumption expenditures of core goods over the nominal PCE excluding energy and food, and similarly for the other two components.

3. Results

3.1 Identification

Structural identification requires us to model the correlations between the reduced-form residuals. Our procedure (taken from Verbrugge and Zaman (2023); see Appendix B) starts with identifying all statistically significant correlations between the reduced-form residuals.

Accordingly, we used simple OLS regressions (i.e., regressed the residuals from equation (1) on those from equation (2), etc.) and examined t-statistics. We found a significant correlation between PPI residuals and core goods residuals, between PPI residuals and median non-housing core services residuals, and between u^{lowgap} and $u^{medfreq}$ residuals; all other correlations were statistically insignificant. This left us with 8 possible models. On the basis of economic theory and *a priori* timing grounds, we assume that contemporaneously, PPI causes core goods, PPI

causes median non-housing core services, and $u^{medfreq}$ causes u^{lowgap} . Denoting our SVAR in matrix notation by

$$AM_t = B(L)M_t + V_t$$

where $M_t = [\pi_t^{PPI}, \pi_t^{CoreG}, \pi_t^{MNHserv}, \pi_t^{Hous}, u_t^{medfreq}, u_t^{lowgap}]^T$, and imposing that V_t is diagonal, our assumptions lead to the following loading matrix A (only nonzero entries are indicated):

$$AM_t = \begin{bmatrix} 1 & & & & & \\ -a_{21} & 1 & & & & \\ -a_{31} & & 1 & & & \\ & & & 1 & & \\ & & & & 1 & \\ & & & & -a_{65} & 1 \end{bmatrix} \begin{bmatrix} \pi_t^{PPI} \\ \pi_t^{CoreG} \\ \pi_t^{MNHserv} \\ \pi_t^{Hous} \\ u_t^{medfreq} \\ u_t^{lowgap} \end{bmatrix}$$

Maximum likelihood estimation of A , based on the variance-covariance matrix from the equation residuals and the zeroes of the loading matrix A , indicated that a_{21} , a_{31} and a_{65} were statistically significant.⁹

Given these results and the sparsity of the A matrix, to estimate the identified system, it suffices to modify the core goods and median non-housing core services equations by including a contemporaneous PPI term, modify the u^{lowgap} equation by adding a contemporaneous $u^{medfreq}$ term, and estimate the (now fully identified) nonlinear system equation by equation.¹⁰ Thus, the three respecified equations are

$$\begin{aligned} \pi_t^{CoreG} = & \alpha^{CoreG} + \phi_1^{CoreG} \pi_{t-1}^{CoreG} + \phi_2^{CoreG} \pi_{t-2}^{CoreG} + \phi_5^{CoreG} \pi_{t-5}^{CoreG} + \\ & + \beta_0^{CoreG} \pi_t^{PPI} + \beta_1^{CoreG} \pi_{t-1}^{PPI} + \beta_2^{CoreG} \pi_{t-2}^{PPI} + \lambda^{CoreG} u_{t-4}^{+medfreq} + \psi I^{1995} + v_t^{CoreG} \end{aligned} \quad (7)$$

$$\begin{aligned} \pi_t^{MNHserv} = & \alpha^{MNHserv} + \gamma_1^{MNHserv} \pi_t^{MNHserv} + \gamma_2^{MNHserv} \pi_t^{MNHserv} + \gamma_5^{MNHserv} \pi_t^{MNHserv} + \\ & + \beta_0^{MNHserv} \pi_t^{PPI} + \lambda^{MNHserv} u_{t-1}^{+medfreq} + \mu^{MNHserv} u_{t-1}^{-lowgap} + v_t^{MNHserv} \end{aligned} \quad (8)$$

$$u_t^{lowgap} = \alpha^{lowgap} + \sum_{j=1}^2 \mu_j^{low} u_{t-j}^{lowgap} + \sum_{j=1}^4 \lambda_j^{low} u_{t-j}^{medfreq} + \sum_{j=0}^4 \beta_j^{low} \pi_{t-j}^{PPI} + v_t^{lowgap} \quad (9)$$

⁹ There is some abuse of notation. Our full structural model has 11 equations, 5 of which are identities, as explained below. But what matters for identification is determining the contemporaneous causation structure among the variables, and none of these involve sign asymmetry

¹⁰ Results are qualitatively unchanged if we adopt the commonly used practice of adjusting the original reduced-form coefficients by multiplying by A^{-1} .

Further, in equations (1), (4), and (5), the reduced-form residuals e are relabeled as structural residuals v . Coefficient estimates are reported in Appendix C.

For simulating the system – necessary for estimation of forecasts and their error bands – we must augment these 4 equations with 5 additional equations: 4 equations that split each unemployment rate component projection into positive and negative parts, and a final one that defines the first difference of $u^{medfreq}$.

$$u_t^{+lowgap} \equiv \max(0, u_t^{lowgap}) \quad (10)$$

$$u_t^{-lowgap} \equiv \min(0, u_t^{lowgap}) \quad (11)$$

$$u_t^{+medfreq} \equiv \max(0, u_t^{medfreq}) \quad (12)$$

$$u_t^{-medfreq} \equiv \min(0, u_t^{medfreq}) \quad (13)$$

$$\Delta u_t^{medfreq} \equiv u_t^{medfreq} - u_{t-1}^{medfreq} \quad (14)$$

The full structural model consists of equations (1), (4), and (5) (with residuals v), and equations (7) through (14).

3.2 Forecasts

As has been long-established in the inflation forecasting literature, overall forecast accuracy can be enhanced by conditioning upon near-term information (see, e.g., Faust and Wright, 2013). The variables where such information is most useful for our purposes are core goods inflation (where monthly inflation has decelerated sharply) and housing inflation (where models relying on short-term information, discussed below, suggest that we will have at least one more quarter of inflation growth).

We incorporate the recent deceleration in core goods inflation by conditioning a path for 4Q core goods inflation over the next four quarters that leaves it slightly negative in 2023Q4.¹¹ If anything, doing so imposes a strong downward bias on our forecasts, since the model by itself (i.e., unconditionally) predicts a slower deceleration in core goods inflation.

¹¹ Following the nowcasting inflation work of Knotek and Zaman (2017), who found superior accuracy of core PCE nowcasts and short-term forecasts using simple models including AR processes, we construct the short-term forecast path for monthly core goods inflation using a simple AR(2) model estimated over our sample.

We incorporate short-term information in housing services by use of a short-term model, informed by Adams et al. (2022). This paper uses confidential CPI rent microdata to demonstrate that new-tenant rents lead official CPI rents (the ultimate source of the housing services inflation information in the core PCE) by about 4 quarters, and that the CoreLogic Single-Family Rent Index (SFRI) has historically tracked a CPI-microdata-based new-tenant rent index fairly well. We use a simple model for monthly housing services inflation¹² using lags of both housing services inflation and SFRI rent inflation to produce a forecast for housing services inflation for January, February, and March of 2023. This yields a 2023Q1 estimate of 7.7 percent (quarterly annualized or 7.9 percent 4Q-trailing basis), which we use as a starting condition for housing services inflation.

We first present the model projection for core PCE inflation through 2025, along with 70 percent confidence intervals, and the SEP projection in Figure 3. Our model projections are conditional on the December SEP path for unemployment. (For interpretive ease, we have interpolated between the SEP projected values for core PCE inflation, which are provided only for 2023Q4, 2024Q4, and 2025Q4.)




Figure 3: Projections of core PCE inflation

¹² We thank Mark Bognanni and Katia Peneva for advice in constructing this model.

The SEP projection initially lies above the model's mean projection, and outside the 70 percent confidence interval. This is because, in our model, the projected 2023 uptick in the unemployment rate in the SEP projection puts downward pressure on all of the inflation variables. Thereafter, however, the persistence of inflation reflected in our model estimates becomes evident, and progress toward the 2 percent target slows. Conversely, the SEP projection then continues its steady downward drift. This steady decline moves the SEP projection within the confidence interval, where it remains for most of 2024. However, thereafter, the SEP projection continues to move steadily lower, so that it moves outside of, and well below, the confidence interval. Hence, from late 2024 onward, the SEP projection is assessed as too optimistic relative to our model's assessment. Our model forecast is a touch below 2.7 percent by the end of 2025; it does not reach 2.1 percent inflation until several years later.

Figure 4 presents our model projections for our three components: core goods inflation, non-housing core services inflation, and housing inflation, conditional on their respective short-term conditions (as discussed above) and the SEP path for unemployment over the 2023-2025 period. Our model sees core goods inflation rebounding from -0.5 percent inflation in 2022Q3 to near 0 in early 2024, then slowing to a -0.40 percent pace in 2025. Non-housing core services inflation is projected to fall to 3.8 percent by the end of 2023, driven by downward pressure from the rising unemployment. Then its downward progress slows so that it decelerates to 3.4 percent by the end of 2025. Housing services inflation is projected to decline at a steady rate through late 2024, but then its downward progress stalls out, likely reflecting the sluggish dynamics of rent (see Adams et al., 2022 and Gallin and Verbrugge, 2019). During 2025, it settles in at a 4.5 percent pace, before continuing to decline briskly, reaching 3.8 percent by early 2027.




Figure 4: Projections of the components of core PCE inflation

We next provide a number of additional inflation projections, conditional on three alternative unemployment rate scenarios: a soft landing scenario, a moderate recession scenario, and a severe recession scenario. The soft landing scenario, which conditions on the projected unemployment path from the June SEP, has unemployment peaking at 4.1 percent by the end of 2024.¹³ The moderate recession scenario conditions on a path for unemployment from 2023Q1 onward that mimics the 2001 recession. For this path, unemployment tops out at 5.6 percent in 2025Q3. Finally, the severe recession scenario (inspired by the Summers/Ball/Leigh/Mishra assertions) conditions on a path for unemployment from 2023Q1 onward that mimics the 1973 recession. For this path, unemployment tops out at 7.8 percent in 2024Q2, although unemployment averages 7.4 percent over the year. Unemployment rates in all scenarios, with the exception of the severe recession, are assumed, after 2025Q4, to descend linearly to hit 4 percent by the end of 2029 (or in the case of the soft landing, by the end of 2025). All of these scenario paths are plotted below in Figure 5. In our specification, the exact path of unemployment taken after 2024 in its descent toward 4 percent is essentially immaterial for inflation, but these paths will impact the simple welfare analysis conducted below. The implied forecasts for core PCE inflation are shown in Figure 6.

¹³ The SEP projection reports the forecast of the overall unemployment rate. To back out the implied projection of the underlying jobless unemployment rate, we take the temporary-layoff rate reported by the BLS for the month of December 2022, and assume that it will persist into the future.




Figure 5: Projections of the unemployment rate




Figure 6: Alternative projections of core PCE inflation

Our model sees rapid deceleration of inflation over 2023, for all of these scenarios, driven by rapid deceleration in core goods prices and by initial movement back toward trend. Recessionary downward force, i.e., the deceleration pressure associated with the positive portion of the medium-frequency component of unemployment, amplifies this descent for all scenarios except the soft landing. This pressure eases in early 2024 for the December SEP path and the moderate recession path but continues for three more quarters in the severe recession scenario. Notice, however, that once the deceleration pressures ease, progress toward 2 percent slows markedly. Inflation is more persistent than is commonly believed.

Regarding that persistence, as noted in Verbrugge and Zaman (2023), allowing for the fifth lag in each of the three core PCE component inflation variables is quite important. In the appendix of their paper, Verbrugge and Zaman (2023) show that, effectively, the persistence of an autoregressive process with (say) a coefficient of 0.1 on the fifth lag is far greater than an autoregressive process with only one lag, even if the coefficient on that lag is equal to the sum of the autoregressive coefficients in the lag-5 process. These forecasts imply that it takes a very long time for inflation to return to trend. This fact is consistent with the inflation experience over the 2012-2019 period, when trend inflation moved a mere 0.5 percentage point (see Ashley and Verbrugge, 2022a). We further note that our core PCE projections are very similar to those of the model in Verbrugge and Zaman (2023), which is built upon trimmed-mean PCE inflation, and also to those from the headline PCE forecasts of the model in Verbrugge and Zaman (2022).

In Appendix D, we display a 10-year conditional recursive forecast from our model over the 2007-2016 period. Its accuracy leads us to believe that the forecasts from our model are generally reliable.

3.3 A Simple Welfare Analysis

Despite its higher inflation path, is a soft landing preferable? We conduct a simple welfare analysis, using a standard (though ad hoc) quadratic loss function. In some contexts, such loss functions are a second-order Taylor series approximation to the expected utility of the economy's representative household (Woodford, 2002), specified as

$$L\{u_t, \pi_t\}_{t=t_1}^{t_2} = \sum_{s=0}^{t_2-t_1} \left[w(u_{t_1+s} - u_t^*)^2 + (1-w)(\pi_{t_1+s} - \pi^*)^2 \right]$$

Guided by the December SEP and the FOMC’s inflation target, we set $u_t^* = 4.0$ and $\pi^* = 2.0$. We examine losses from $t_1 = 2023\text{Q1}$ to $t_2 = 2029\text{Q4}$. We compare the soft landing, moderate recession, severe recession, and December SEP scenarios. We report the losses in Table 1, for $w = \{0.1, 0.25, 0.5, 0.75, 0.9\}$.

Table 1: Welfare losses

Weight on unemployment	Soft landing	Moderate recession	Severe recession	December SEP
0.1	20.7	15.5	22.8	16.0
0.25	17.3	16.6	38.0	13.8
0.5	11.7	18.4	63.4	10.1
0.75	6.1	20.2	88.8	6.4
0.9	2.7	21.3	104.0	4.1

In Table 1, we have highlighted the minimum-loss scenario for each value of w . The estimates indicate that for moderate values of w , i.e., a weight of 0.25 or 0.50 on the unemployment rate, the projection conditioned on the SEP path results in a smaller welfare loss than does the soft landing. Only for higher values of w does the soft landing result in smaller welfare losses than does the SEP path, and only for very low values of w does the moderate recession result in smaller welfare losses than does the SEP path. In summary, this welfare analysis suggests that a December SEP is the preferred outcome. Importantly, however, this welfare analysis abstracts from any danger of the unanchoring of inflation expectations that might be associated with core PCE inflation still being nearly 2.7 percent three years from now.

4. Conclusion

This paper implements a nonlinear structural VAR model to jointly estimate the dynamics of inflation, as measured by three components of core PCE inflation, an indicator of supply-chain pressures, and two components of the jobless unemployment rate: a persistent component and moderately persistent component.

The model is estimated with post-1985 quarterly data and identification of structural shocks is achieved using the data-determined method of Swanson and Granger (1997), which substantially reduces the role of subjectivity.

Looking ahead, our model projects that inflation only very gradually falls back to 2 percent. Progress toward target is very much influenced by the path that unemployment will take over the next several years. Conditional on the December SEP median unemployment rate projections, inflation is projected to still be nearly 2.7 percent by the end of 2025, far above the SEP's median projection of 2.1 percent. A moderate recession (roughly equal to the recession of 2001) would put inflation at 2.4 percent by the end of 2025; conversely, a soft landing (which we define as the path of unemployment in the June SEP projection) would put inflation a touch above 2.8 percent by the end of 2025. What kind of recession would it take to hit the SEP projection for inflation, according to the model developed in this paper? We investigate the claim of former Treasury Secretary Lawrence Summers (reported in Aldrick, 2022) and the supporting assessment of Ball, Leigh, and Mishra (2022) that it will take two years of 7.5 percent unemployment from its current low level to bring inflation down to its 2 percent target. We find that one year of 7.4 percent unemployment would accomplish this task.

A simple welfare analysis based on a standard quadratic loss function overall favors the December SEP unemployment rate path. However, this welfare analysis abstracts from any danger of the unanchoring of inflation expectations that would be associated with core PCE inflation still being 2.8 percent three years from now.

Ashley and Verbrugge (2022a) summarize a large number of extant theoretical works whose predictions are consistent with their (and our) empirical results regarding the nonlinearity of the Phillips curve. We hope that the present paper provides further impetus for the development of structural models that are consistent with, and provide a theoretical explanation for, our findings.

References

- Adams, Brian, Lara Loewenstein, Hugh Montag, and Randal J. Verbrugge. 2022. “Disentangling Rent Index Differences: Data, Methods, and Scope.” Working paper 22–38. Federal Reserve Bank of Cleveland. <https://doi.org/10.26509/frbc-wp-202238>.
- Aldrick, Philip. 2022. “Larry Summers Says US Needs 5% Jobless Rate for Five Years to Ease Inflation.” *Bloomberg.Com*, June 20, 2022. <https://www.bloomberg.com/news/articles/2022-06-20/summers-says-us-needs-5-jobless-rate-for-five-years-to-ease-cpi>.
- Ashley, Richard, Kwok Ping Tsang, and Randal J. Verbrugge. 2020. “A New Look at Historical Monetary Policy (and the Great Inflation) through the Lens of a Persistence-Dependent Policy Rule.” *Oxford Economic Papers* 72 (3): 672–91. <https://doi.org/10.1093/oep/gpaa006>.
- Ashley, Richard, and Randal J. Verbrugge. 2008. “Frequency Dependence in Regression Model Coefficients: An Alternative Approach for Modeling Nonlinear Dynamic Relationships in Time Series.” *Econometric Reviews* 28 (1–3): 4–20. <https://doi.org/10.1080/07474930802387753>.
- Ashley, Richard, and Randal J. Verbrugge. 2022a. “The Intermittent Phillips Curve: Finding a Stable (but Persistence-Dependent) Phillips Curve Model Specification.” Manuscript. Virginia Tech. (Updated version of the Federal Reserve Bank of Cleveland Working paper 19-09R)
- Ashley, Richard, and Randal J. Verbrugge. 2022b. “Death to Regression Modeling on Bandpass-Filtered Data: New Results on the General Undesirability of Filtering Dynamic Regressions, with a Proposed Solution.” Manuscript, Virginia Tech.
- Ashley, Richard, and Randal J. Verbrugge. 2022c. “Time Series Filtering for Studying Frequency Dependence: What Works, What Doesn’t.” Manuscript.
- Ball, Laurence M., Daniel Leigh, and Prachi Mishra. 2022. “Understanding US Inflation during the COVID Era.” In *BPEA Conference Drafts, September 8-9, 2022*. <https://www.brookings.edu/bpea-articles/understanding-u-s-inflation-during-the-covid-era/>.
- Blalock Jr, Hubert M. 1961. *Causal Inferences in Nonexperimental Research*. Chapel Hill: University of North Carolina Press.
- Carroll, Daniel R., and Randal J. Verbrugge. 2019. “Behavior of a New Median PCE Measure: A Tale of Tails.” *Economic Commentary*, no. 2019–10 (July). <https://doi.org/10.26509/frbc-ec-201910>.
- Christiano, Lawrence J., and Terry J. Fitzgerald. 2003. “The Band Pass Filter.” *International Economic Review* 44 (2): 435–65. <https://doi.org/10.1111/1468-2354.t01-1-00076>.
- Clark, Todd E. 2004. “An Evaluation of the Decline in Goods Inflation.” *Economic Review* 89 (Q II): 19–51. <https://ideas.repec.org/a/fip/fedker/y2004iqiip19-51nv.89no.2.html>.

- Clark, Todd E., and Sharon Kozicki. 2005. "Estimating Equilibrium Real Interest Rates in Real Time." *The North American Journal of Economics and Finance* 16 (3): 395–413. <https://doi.org/10.1016/j.najef.2005.04.002>.
- Clark, Todd E., and Saeed Zaman. 2013. "Forecasting Implications of the Recent Decline in Inflation." *Economic Commentary*, no. 2013–15 (November). <https://doi.org/10.26509/frbc-ec-201315>.
- Dagum, Estela Bee. 1987. "Current Issues in Seasonal Adjustment." *Survey Methodology* 13 (1): 63–74. <https://www150.statcan.gc.ca/n1/en/catalogue/12-001-X198700114509>.
- Demiralp, Selva, and Kevin D. Hoover. 2003. "Searching for the Causal Structure of a Vector Autoregression." *Oxford Bulletin of Economics and Statistics* 65 (s1): 745–67. <https://doi.org/10.1046/j.0305-9049.2003.00087.x>.
- Demiralp, Selva, Kevin D. Hoover, and Stephen J. Perez. 2008. "A Bootstrap Method for Identifying and Evaluating a Structural Vector Autoregression." *Oxford Bulletin of Economics and Statistics* 70 (4): 509–33. <https://doi.org/10.1111/j.1468-0084.2007.00496.x>.
- Faust, Jon, and Jonathan H. Wright. 2013. "Chapter 1 - Forecasting Inflation." In *Handbook of Economic Forecasting*, 2, Part A:2–56. Elsevier. <https://doi.org/10.1016/B978-0-444-53683-9.00001-3>.
- Figura, Andrew, and Christopher J. Waller. 2022. "What Does the Beveridge Curve Tell Us about the Likelihood of a Soft Landing?" July. <https://doi.org/10.17016/2380-7172.3190>.
- Gallin, Joshua, and Randal J. Verbrugge. 2019. "A Theory of Sticky Rents: Search and Bargaining with Incomplete Information." *Journal of Economic Theory* 183 (September): 478–519. <https://doi.org/10.1016/j.jet.2019.06.003>.
- Glymour, Clark, and Peter Spirtes. 1988. "Latent Variables, Causal Models and Overidentifying Constraints." *Journal of Econometrics* 39 (1): 175–98. [https://doi.org/10.1016/0304-4076\(88\)90044-9](https://doi.org/10.1016/0304-4076(88)90044-9).
- Granger, Clive W. J. 1969. "Investigating Causal Relations by Econometric Models and Cross-Spectral Methods." *Econometrica* 37 (3): 424–38. <https://doi.org/10.2307/1912791>.
- Hall, Robert E., and Marianna Kudlyak. 2022. "The Unemployed with Jobs and without Jobs." *Labour Economics* 79 (December): 102244. <https://doi.org/10.1016/j.labeco.2022.102244>.
- Hamilton, James D. 2018. "Why You Should Never Use the Hodrick-Prescott Filter." *The Review of Economics and Statistics* 100 (5): 831–43. https://doi.org/10.1162/rest_a_00706.

- Iacobucci, Alessandra, and Alain Noullez. 2005. "A Frequency Selective Filter for Short-Length Time Series." *Computational Economics* 25 (1): 75–102. <https://doi.org/10.1007/s10614-005-6276-7>.
- Kaiser, Regina, and Agustín Maravall. 1999. "Estimation of the Business Cycle: A Modified Hodrick-Prescott Filter." *Spanish Economic Review* 1 (2): 175–206. <https://doi.org/10.1007/s101080050008>.
- Kilian, Lutz, and Helmut Lütkepohl. 2017. *Structural Vector Autoregressive Analysis*: 1st ed. Cambridge University Press. <https://doi.org/10.1017/9781108164818>.
- Knotek II, Edward S., and Saeed Zaman. 2017. "Nowcasting U.S. Headline and Core Inflation: MONEY, CREDIT AND BANKING." *Journal of Money, Credit and Banking* 49 (5): 931–68. <https://doi.org/10.1111/jmcb.12401>.
- Mise, Emi, Tae-Hwan Kim, and Paul Newbold. 2005. "On Suboptimality of the Hodrick–Prescott Filter at Time Series Endpoints." *Journal of Macroeconomics* 27 (1): 53–67. <https://doi.org/10.1016/j.jmacro.2003.09.003>.
- Moneta, Alessio. 2008. "Graphical Causal Models and VARs: An Empirical Assessment of the Real Business Cycles Hypothesis." *Empirical Economics* 35 (2): 275–300. <https://doi.org/10.1007/s00181-007-0159-9>.
- Stock, James H., and Mark W. Watson. 1999. "Forecasting Inflation." *Journal of Monetary Economics* 44 (2): 293–335. [https://doi.org/10.1016/S0304-3932\(99\)00027-6](https://doi.org/10.1016/S0304-3932(99)00027-6).
- Swanson, Norman R., and Clive W. J. Granger. 1997. "Impulse Response Functions Based on a Causal Approach to Residual Orthogonalization in Vector Autoregressions." *Journal of the American Statistical Association* 92 (437): 357–67. <https://doi.org/10.1080/01621459.1997.10473634>.
- Tallman, Ellis W., and Saeed Zaman. 2017. "Forecasting Inflation: Phillips Curve Effects on Services Price Measures." *International Journal of Forecasting* 33 (2): 442–57. <https://doi.org/10.1016/j.ijforecast.2016.10.004>.
- Tallman, Ellis W., and Saeed Zaman. 2020. "Combining Survey Long-Run Forecasts and Nowcasts with BVAR Forecasts Using Relative Entropy." *International Journal of Forecasting* 36 (2): 373–98. <https://doi.org/10.1016/j.ijforecast.2019.04.024>.
- Verbrugge, Randal. 2022. "Is It Time to Reassess the Focal Role of Core PCE Inflation in Assessing the Trend in PCE Inflation?" *Economia* 45 (89): 73–101. <https://doi.org/10.18800/economia.202201.004>.

Verbrugge, Randal J., and Saeed Zaman. 2022. “Improving Inflation Forecasts Using Robust Measures.” Working paper 22-23. Federal Reserve Bank of Cleveland. <https://doi.org/10.26509/frbc-wp-202223>.

Verbrugge, Randal J., and Saeed Zaman. 2023. “The Hard Road to a Soft Landing: Evidence from a (Modestly) Nonlinear Structural Model.” Working paper 23–03. Federal Reserve Bank of Cleveland. <https://doi.org/10.26509/frbc-wp-202303>.

Woodford, Michael. 2002. “Inflation Stabilization and Welfare.” *The B.E. Journal of Macroeconomics* 2 (1): 1–53. <https://ideas.repec.org/a/bpj/bejmac/vcontributions.2y2002i1n1.html>.

Zaman, Saeed. 2022. “A Unified Framework to Estimate Macroeconomic Stars.” Working paper 21-23R. Federal Reserve Bank of Cleveland. <https://doi.org/10.26509/frbc-wp-202123r>.

Appendix A: Partitioning the Jobless Unemployment Rate

A.1 Overview

To partition the jobless unemployment rate while applying the Ashley/Verbrugge (2022a,b) method (described below), we use the Iacobucci-Noullez (2005) filter, setting $k = 4$ (i.e., using 4 quarters of univariate forecasts in each rolling window). The Iacobucci-Noullez filter introduces no phase shift (unlike, e.g., the Christiano-Fitzgerald (2003) filter).¹⁴ Following Ashley and Verbrugge (2022a), we choose frequency cutoffs so that the jobless unemployment rate is partitioned into fluctuations lasting longer than 4 years (termed u_t^{lowgap} , for low-frequency gap), fluctuations lasting between 1 year and 4 years (termed $u_t^{medfreq}$, for medium-frequency), and transient fluctuations. Transient fluctuations were found to be unimportant drivers of inflation, and so were omitted. The two more persistent unemployment components are plotted in Figure 2. This figure also demonstrates how unusual the COVID collapse and recovery were. In particular, the low-frequency gap rose, notably much more sharply than usual. Meanwhile, the medium-frequency component also fell very sharply back to zero, dropping to historic lows.

The first step is to *partition* the real-time unemployment rate u_t into 3 persistence components – $u_t = \sum_{j=1}^3 u_{j,t}$ – which *by construction* add up to the original series. These 3 persistence components partition the variation in u_t into monotonically decreasing levels of persistence or, equivalently, increasing frequency levels. These components are obtained from the sample data using a moving window (augmented with a k -quarter forecast) to filter the u_t data at each time t in a one-sided (backward-looking) manner. This approach mitigates end-of-sample filter distortions, ensures that parameter estimates are consistent, and retains both the causality structure of the data-generating process and any orthogonality conditions that are present in the unfiltered data. The Ashley/Verbrugge persistence-dependent regression methodology then merely replaces $u_t - u_t^*$ with these 3 persistence components, estimating a separate coefficient for each. (We note in passing that we subtract u_t^* from the most persistent component. That way, the

¹⁴ Use of the CF filter in the AV method produced qualitatively similar results, though it calls for the predicted recession to begin a quarter or two later. For a comparison of the use of different filters for frequency-dependent regression (as well as the sensitivity of results to forecast parameters), see Ashley and Verbrugge (2020c). Using the CF filter with the Ashley/Verbrugge method mitigates its phase shift in any case.

3 components add up to an unemployment gap, the typical Phillips curve specification in the literature.) Simulation evidence in Ashley, Tsang, and Verbrugge (2020) and Ashley and Verbrugge (2022c) indicate that the method yields reliable coefficient estimates and inferences, for both linear and nonlinear data-generating processes. Below we summarize why partitioning, one-sided filtering, augmentation or ‘padding’ with forecasts, and restriction of the filtering solely to the u_t data are all *essential* for obtaining reliable inferences.

A.2 One-Sided Filtering Method of Ashley and Verbrugge




Figure A.1: One-sided filtering of unemployment rate data at date $s+\tau$ (using a two-sided filter, from time s to time $s+k+m$)

A2.1 Description of one-sided filtering

In brief, one undertakes the one-sided filtering by running a window through the data. Over each window, one saves the decomposition at the final *data point* in the window. Then one increments the window by one quarter. However, each window includes not just data but also a second component that is a forecast. In other words, each window includes data augmented with a forecast.

To explain this in more detail, consider Figure A.1. We wish to compute the decomposition of the unemployment rate at time $s+k$. As is well-known, obtaining the decomposition at $s+k$ by using a two-sided filter from time s to time $s+k$ would yield estimates with very poor properties. In particular, the resultant time series would (for most filters) incorporate a pronounced phase shift, in addition to being highly inaccurate; this inaccuracy is due to the well-known “edge effect” problem plaguing all filters.

Both the phase-shift and edge-effect problems are addressed by augmenting the data within a window with forecasts. In particular, as in Dagum (1987), Stock and Watson (1999), Kaiser and

Maravall (1999), Mise, Kim and Newbold (2005), and Clark and Kozicki (2005) – and as is done routinely in seasonal-adjustment procedures – one should augment the window sample data with forecasted data. In the situation depicted in Figure A.1 we have κ sample data points (from time period s to time period $s+\kappa$), and m months of projections, yielding a $(\kappa+m)$ -quarter window (from time period s to time period $s+\kappa+m$). We then use a two-sided filter to partition that window into persistence components, and then save the partition at date $s+\kappa$; notice that this is a one-sided partition, since no *data* after date $s+\kappa$ are used. To obtain the partition at date $s+\kappa+1$, we repeat this procedure, obtaining a forecast from data $s+\kappa+1$ to data $s+\kappa+1+m$, then use a two-sided filter over dates $s+1$ to $s+\kappa+m+1$ and saving the partition at date $s+\kappa+1$. This procedure also gracefully allows one to use real-time data.

A2.2 Testing for persistence-dependence

How does testing proceed? In the present case, we wish to test whether the Phillips curve is persistence-dependent. Thus, we partition the unemployment rate un into three components (say): un^1 , un^2 , and un^3 . Then we replace un in the Phillips curve specification with its 3 components. One may readily test for persistence-dependence using a standard Chow test. Since the components sum to the original series and are based upon one-sided filtering, the causality structure and the properties of the error term are preserved. For more details, see the appendix to Ashley, Tsang, and Verbrugge (2020).

A2.3 Sensitivity to forecasts and filter

Ashley and Verbrugge (2022c) demonstrate that the resultant persistence decomposition is not very sensitive to the number of forecast periods chosen, as long as at least a year of projections are used, nor to the frequency filter used (the Iacobucci-Noullez filter, the Christiano-Fitzgerald filter, or the Ashley-Verbrugge filter) nor to the details of how these forecasts are produced (as long as they are reasonably accurate).

What is crucial is to *partition* the explanatory variables into an interpretable small set of frequency/persistence components that add up to the original data, using moving windows passing through the data so that the filtering is done in a backward-looking or one-sided manner. The Ashley-Verbrugge filter has a key advantage: it can partition the time series into k components in a single pass and is thus more readily used for discovering the persistence-dependence in the original data. Other filters must be used in an iterative manner, and in our experience, results are disappointing if one attempts to partition the data into more than 3 components. Furthermore, Ashley and Verbrugge (2022c) demonstrate that the results using other filters are somewhat sensitive to the manner in which this iteration is done.

But with these details in mind, what is of practical macroeconomic importance is to allow for frequency/persistence dependence in the relationship, not – so long as one is mindful of the basic desiderata delineated above – the technical details of precisely how the explanatory variable is partitioned into its frequency/persistence components. Ashley and Verbrugge (2022a) report that

alternative techniques usually yield quite similar empirical results in practice; see Ashley and Verbrugge (2022c) for more details. RATS, Stata, and Matlab code to accomplish this type of one-sided decomposition (using simple univariate or multivariate forecasts) is available from the authors.

A2.4 Rationale for partitioning, one-sided filtering, and filtering only explanatory variables

Why are partitioning, one-sided filtering, and restriction of the filtering solely to the $u_t - u_t^*$ data all essential? Partitioning is necessary to ensure that these 3 components of the unemployment rate gap add up to the original data, making it easy to test the null hypothesis that the coefficients with which these 3 components enter a regression model for the inflation rate are all equal. One-sided filtering is necessary because two-sided filtering – such as ordinary HP filtering or ordinary spectral analysis – inherently mixes up future and past values of the unemployment rate gap in obtaining the persistence components, distorting the causal meaning of inference in the resulting inflation model and limiting its use for practical forecasting and/or policy analysis. These distortions from the use of two-sided filtering are particularly severe when the dependent variable is also filtered and when the key relationship likely (as here) involves feedback from the dependent variable (inflation) to the (filtered) components of $u_t - u_t^*$ being used as explanatory variables. Fundamentally, this is because filtering the dependent variable in a regression model implies that the model error term is similarly filtered. For more details, see Ashley and Verbrugge (2008, 2022b); for a “practical” comparison of methods, including the Hamilton (2018) filter, see Ashley and Verbrugge (2022c).

How about two-sided spectral estimates or filtering with wavelets? These are two-sided methods, so the same criticisms apply. Hence, two-sided cross-spectral estimates or filtering with wavelets are ruled out for analyses of the present sort. And regarding spectral methods, even absent feedback, transfer function gain and phase plots are substantially more challenging to interpret than our approach; even without the presence of feedback, Granger describes interpretation of such plots as “difficult or impossible” (Granger, 1969).

Appendix B: Identification

We adopt the Swanson and Granger approach to identification.¹⁵ This method is built upon the fact that most structural causal models, whether linear or nonlinear, imply overidentifying constraints. In particular, a given structural model implies partial correlation constraints on reduced-form regression residuals $(e_{X,t}, e_{Y,t}, e_{Z,t})$. These restrictions take the form $e_X \perp e_Y | e_Z$. Under fairly weak assumptions, such constraints may be tested using standard t -statistics and, if the test is rejected, one may thus reject that structural model.

But notice that *all* structural models that share such a constraint are also accordingly rejected. Hence, such tests may be used to restrict the class of models that are consistent with the data. By virtue of ruling out candidate models that are inconsistent with the data, tests of such overidentifying constraints thus prove useful in *specifying* a structural model. This procedure substantially reduces the subjective nature in the typical SVAR methodology.

To demonstrate how this works in practice, we provide a simple example. Consider the following structural model, an SVAR involving 3 variables, X , Y , and Z ; for simplicity, assume that each is standardized to have mean 0 and standard deviation 1. The model is a structural vector autoregression of order 2:

$$\begin{bmatrix} 1 & 0 & 0 \\ -a_{21} & 1 & 0 \\ -a_{31} & 0 & 1 \end{bmatrix} \begin{bmatrix} x_t \\ y_t \\ z_t \end{bmatrix} = \begin{bmatrix} b_{11} & 0 & b_{13} \\ 0 & b_{22} & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_{t-1} \\ y_{t-1} \\ z_{t-1} \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & c_{22} & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_{t-2} \\ y_{t-2} \\ z_{t-2} \end{bmatrix} + \begin{bmatrix} v_{X,t} \\ v_{Y,t} \\ v_{Z,t} \end{bmatrix} \quad (1)$$

In matrix notation, the SVAR is denoted

$$AM_t = B(L)M_t + V_t$$

where $M_t \equiv (X_t, Y_t, Z_t)'$, $B(L)$ is a matrix lag polynomial, and $V_t \equiv (v_{X,t}, v_{Y,t}, v_{Z,t})'$. The corresponding reduced-form model is given by

$$M_t = A^{-1}B(L)M_t + A^{-1}V_t \equiv \Phi(L)M_t + E_t$$

where $\Phi(L)$ is a matrix lag polynomial, and $E_t \equiv (e_{X,t}, e_{Y,t}, e_{Z,t})'$. Identification of the SVAR implies obtaining estimates of A and $B(L)$, with the variance-covariance matrix of V being diagonal.

The structural model errors ν are assumed to be distributed normally, with a diagonal covariance matrix (assumed, for simplicity, to be the identity matrix). This model may be graphically represented in Figure B.1, depicting time t variables as a function of other time t variables and lagged variables. In this figure, an arrow denotes a causal influence: a solid arrow represents a within-period influence, while a dashed arrow represents an intertemporal influence.

¹⁵ This method builds upon work in causal modeling (e.g., Glymour and Spirtes, 1988) and is extended in Demiralp and Hoover (2003) and Demiralp, Hoover, and Perez (2008); see also Moneta (2008). The method originated in Blalock (1961).

For simplicity, the influence of the exogenous structural shocks v_t^k on variables $k \in \{X, Y, Z\}$, is not depicted.




Figure B.1: A structural VAR, with causal influences depicted. Solid lines depict contemporaneous causation; dashed lines depict intertemporal causation. Thus, for example, at time t , variable Y is influenced by variable X contemporaneously, by its own value at time $t-1$, and by its own value at time $t-2$.

This model will be estimated in reduced-form, yielding the residuals (e_x, e_y, e_z) . Notice that if Equation (1) is the data-generating process, the reduced-form residuals will obey certain correlation and partial correlation restrictions. In particular, letting $\rho(e_j, e_k)$ denote the correlation between e_j and e_k , some of the correlation restrictions that these residuals must satisfy are $\rho(e_x, e_y) \neq 0$; $\rho(e_x, e_z) \neq 0$; $\rho(e_y, e_z) \neq 0$; and $\rho(e_y, e_z | e_x) = 0$.

Of course, in general, the model that generated the data is unknown. How can the data help us specify the model (or more specifically, the structure of the A matrix)?¹⁶ Suppose the model in Equation (1) is true, but the analyst does not know that. As will typically be the case with Normally distributed residuals, the data will not fully identify the model. But the power of the Swanson/Granger approach is that the data may nonetheless be used to sharply reduce the set of possible models. In a three-variable VAR with normal structural errors, ruling out structural models that are impossible to identify leaves 22 possible models (6 of which correspond to Cholesky identification schemes; see Figure A.1 in Appendix A). In the present case, as we will now demonstrate, the data will reject 19 of these. We describe the Swanson/Granger heuristic

¹⁶ The identification challenge in structural VARs of this form consists of restrictions on the A matrix. Of course, each unique set of restrictions on the A matrix corresponds to an entire class of models wherein intertemporal relationships are not restricted. However, intertemporal relationships may be estimated without ambiguity from the data, so identification consists of restrictions on the A matrix. For brevity, we refer to the class of models corresponding to a particular structure of the A matrix as “a” model.

procedure somewhat more formally below. Here we describe informally how one would reject the models that are inconsistent with data generated by Equation (1).

In the first step, by testing all pairwise correlations among the regression residuals (using simple t -tests), one would find that the three residuals are all pairwise correlated. This rules out the last 7 models in Figure A.1, namely, those in which at least one variable is neither caused by, nor causes, any other variable contemporaneously. As the next step, one would conduct all pairwise conditional residual tests, i.e., test $e_X \perp e_Y | e_Z$, $e_X \perp e_Z | e_Y$, and $e_Z \perp e_Y | e_X$ using OLS regressions. Given the data-generating process, the first two hypotheses will be rejected, but not the third. What does this additional information tell us? It rules out 12 of the remaining models, namely, those in which Y and Z have a relationship that is not intermediated in some fashion by X . Putting this differently, it tells us that there are only three possible models that are compatible with the data: those in which $Y \rightarrow X \rightarrow Z$, or $Z \rightarrow X \rightarrow Y$, or $Y \leftarrow X \rightarrow Z$. The data alone cannot be used to discriminate between these three models. However, prior economic information can now be used (in the usual manner) to select from among the three candidate models. For instance, economic theory can sometimes pin down a model based upon the *signs* of the partial correlations. Or one can use the usual timing restrictions – bearing in mind that at the micro level, agents may be responding to the micro data they currently observe, data that will later be aggregated up to data published by a statistical agency.

The heuristic search procedure involves three steps and relies upon the weak “faithfulness” assumption that if X causes Y (or vice versa) within the period, then their residuals will be correlated.¹⁷ First, compute all bivariate partial correlations and examine their statistical significance. If the correlation between e_X and e_Y is weak, and $e_X \perp e_Y$ cannot be rejected, then the data reject $X \rightarrow Y$ and $Y \rightarrow X$ within the period. In an SVAR, the corresponding entries in the impact matrix A would be set to 0. Second, for those variable pairs (X, Y) with significant correlation, construct trivariate partial correlations with all third variables Z , paying particular attention to those that are correlated with both. If $e_X \perp e_Y$ can be rejected, but if $e_X \perp e_Y | e_Z$ cannot be rejected, then we again conclude that the data reject $X \rightarrow Y$ and $Y \rightarrow X$; their correlation stems from a joint relationship with Z . Third, construct all models that are consistent with this evidence, and select the one that is in accord with economic theory priors. In our experience (and in the experience of Granger and Swanson), parsimonious models appear to agree with the data in most cases, and economic theory often plays a minor role in the selection

¹⁷ Swanson and Granger (1997) begin by forgoing unconditional correlation tests and start by examining all conditional correlations; this evidently mitigates reliance on the faithfulness assumption. This assumption will fail under “measure-zero” cases where X causes Y , but the two variables are uncorrelated because X causes Z and Z causes Y , and the two causal paths exactly cancel. In the literature, the “faithfulness-failure” examples occur when there is a decision maker who specifically exerts control over variable Z to accomplish this “cancellation.” If there is reason to believe that such a situation exists in a given context, we would recommend omitting conditional correlation tests (and using this information to help identify the model); otherwise, we recommend the usage of unconditional correlation tests for two reasons. First, our recommendations follow standard practice in the causal analysis literature (see, e.g., Moneta 2008). Second, in practice, what matters most for impulse response function estimates are the identifying assumptions made vis-à-vis variables whose residuals are strongly correlated.

of the final model.¹⁸ (In models with numerous variables, one may formally test higher-order partial correlation constraints implied by the model.)

While a joint testing procedure is unavailable, so that the usual size problems might arise, in practice this issue is often moot. This is because in many cases, the significance levels of tests can be adjusted significantly without any change in inferences. Furthermore, when a borderline case is “accommodated” – i.e., if the model is extended to specify either $X \rightarrow Y$ and $Y \rightarrow X$, when their partial correlation is modest – estimation typically yields impulse response functions that are insensitive to this choice.

¹⁸ If more than one model appears equally reasonable, one may investigate the sensitivity of, e.g., IRFs, to model choice.

Appendix C: Model Estimation Results

$$\pi_t^{PPI} = \alpha^{PPI} + \sum_{j=1}^4 \beta_j^{PPI} \pi_{t-j}^{PPI} + \delta \Delta u_{t-1}^{medfreq} + e_t^{PPI} \quad (1)$$

$$\begin{aligned} \pi_t^{CoreG} = & \alpha^{CoreG} + \phi_1^{CoreG} \pi_{t-1}^{CoreG} + \phi_2^{CoreG} \pi_{t-2}^{CoreG} + \phi_5^{CoreG} \pi_{t-5}^{CoreG} + \\ & + \beta_0^{CoreG} \pi_t^{PPI} + \beta_1^{CoreG} \pi_{t-1}^{PPI} + \beta_2^{CoreG} \pi_{t-2}^{PPI} + \lambda^{CoreG} u_{t-4}^{+medfreq} + \psi I^{1995} + v_t^{CoreG} \end{aligned} \quad (7)$$

Table A.1: Regression Results for Equations (1) and (7).

Regressor	Eq. (1); Dependent variable π_t^{PPI}		Eq. (7); Dependent variable π_t^{CoreG}	
	Coefficient estimate	Standard error	Coefficient estimate	Standard error
α	0.26	0.13	-0.12	0.05
π_{t-1}^{CoreG}			1.08	0.08
π_{t-2}^{CoreG}			-0.27	0.09
π_{t-5}^{CoreG}			0.05	0.04
π_t^{PPI}			0.08	0.03
π_{t-1}^{PPI}	1.49	0.08	-0.14	0.3
π_{t-2}^{PPI}	-0.78	0.15	0.10	0.02
π_{t-3}^{PPI}	0.04	0.15		
π_{t-4}^{PPI}	0.11	0.08		
$u_{t-4}^{+medfreq}$			-0.39	0.25
$\Delta u_{t-1}^{medfreq}$	-5.85	1.57		
I^{1995}			0.34	0.14
\bar{R}^2	0.87		0.94	

Note: We elected to retain $u_{t-1}^{+medfreq}$ and π_{t-5}^{CoreG} in the core goods equation since, absent I^{1995} , both clearly belong in the regression. Dropping π_{t-5}^{CoreG} would have little influence, given the size of the estimated coefficient, but dropping $u_{t-1}^{+medfreq}$ would make inflation a tad less responsive to recessionary pressure.

Table A.2: Regression Results for Equations (4) and (8).

$$\pi_t^{MNHserv} = \alpha^{MNHserv} + \gamma_1^{MNHserv} \pi_{t-1}^{MNHserv} + \gamma_2^{MNHserv} \pi_{t-2}^{MNHserv} + \gamma_5^{MNHserv} \pi_{t-5}^{MNHserv} + \beta_0^{MNHserv} \pi_t^{PPI} + \lambda^{MNHserv} u_{t-1}^{+medfreq} + \mu^{MNHserv} u_{t-1}^{-lowgap} + v_t^{MNHserv} \quad (8)$$

$$\pi_t^{Hous} = \alpha^{Hous} + \sum_{j=1}^5 \eta_j^{Hous} \pi_{t-j}^{Hous} + \lambda^{Hous} u_{t-1}^{+medfreq} + \mu^{Hous} u_{t-4}^{-lowgap} + e_t^{Hous} \quad (4)$$

Regressor	Eq. (8); Dependent variable $\pi_t^{MNHserv}$		Eq. (4); Dependent variable π_t^{Hous}	
	Coefficient estimate	Standard error	Coefficient estimate	Standard error
α	0.01	0.04	0.27	0.06
$\pi_{t-1}^{MNHserv}$	1.20	0.08		
$\pi_{t-2}^{MNHserv}$	-0.36	0.09		
$\pi_{t-5}^{MNHserv}$	0.14	0.04		
π_{t-1}^{Hous}			1.24	0.08
π_{t-2}^{Hous}			-0.26	0.13
π_{t-3}^{Hous}			0.06	0.13
π_{t-4}^{Hous}			-0.48	0.13
π_{t-5}^{Hous}			0.35	0.07
π_t^{PPI}	0.00	0.01		
π_t^{CoreG}			0.01	0.01
u_{t-1}^{lowgap}	-0.12	0.05		
u_{t-4}^{lowgap}			-0.15	0.06
$u_{t-1}^{medfreq}$	-0.19	0.12	-0.88	0.15
\bar{R}^2	0.98		0.97	

Table A.3: Regression Results for Equations (5) and (9).

$$u_t^{medfreq} = \sum_{j=1}^2 \lambda_j^{med} u_{t-j}^{medfreq} + \sum_{j=1}^4 \mu_j^{med} u_{t-j}^{lowgap} + \beta^{med} \pi_{t-1}^{PPI} + e_t^{medfreq} \quad (5)$$

$$u_t^{lowgap} = \alpha^{lowgap} + \sum_{j=1}^2 \mu_j^{low} u_{t-j}^{lowgap} + \sum_{j=1}^4 \lambda_j^{low} u_{t-j}^{medfreq} + \sum_{j=0}^4 \beta_j^{low} \pi_{t-j}^{PPI} + v_t^{lowgap} \quad (9)$$

Regressor	Eq. (5); Dependent variable $u_{t-1}^{medfreq}$		Eq. (9); Dependent variable u_t^{lowgap}	
	Coefficient estimate	Standard error	Coefficient estimate	Standard error
α	0.00	0.00	0.01	0.01
π_{t-1}^{PPI}	0.002	0.001	-0.01	0.01
π_{t-2}^{PPI}			0.03	0.01
π_{t-3}^{PPI}			-0.03	0.01
π_{t-4}^{PPI}			0.01	0.01
u_{t-1}^{lowgap}	0.44	0.03	0.58	0.14
u_{t-2}^{lowgap}	-0.44	0.04	0.37	0.14
u_{t-3}^{lowgap}	-0.18	0.05		
u_{t-4}^{lowgap}	0.18	0.04		
$u_t^{medfreq}$			2.08	0.25
$u_{t-1}^{medfreq}$	0.82	0.07	-1.40	0.28
$u_{t-2}^{medfreq}$	-0.18	0.06	1.32	0.29
$u_{t-3}^{medfreq}$			-0.77	0.27
$u_{t-4}^{medfreq}$			0.43	0.16
\bar{R}^2	0.97		0.99	

Appendix D: Historical Forecasts

To highlight the fit of our model to historical core PCE inflation data, in Figure D.1 below, we plot an inflation forecast from our model, conditional on the actual path of the unemployment rate over the 2007Q1-2016Q4 period, alongside the conditional forecast from a more conventional Phillips curve model. This latter model is specified as a *linear* bivariate model of core PCE inflation and the unemployment gap.¹⁹ As Ashley and Verbrugge (2022a) found in their study of trimmed mean PCE inflation, the conditional forecast from the present nonlinear model broadly captures the decline in inflation following the financial collapse and the very slow return of inflation to the inflation target. We think this serves as a demonstration of how the model responds appropriately to evolving slack and how it accurately captures inflation dynamics over an extended period. In contrast, the conventional model underestimates the strength of the Phillips curve relationship, thereby missing the large drop in inflation following the 2008 financial crisis (and the Great Recession), and its recovery after that. However, this simple model, just like our nonlinear model, correctly captures the high degree of persistence in core PCE inflation.

As discussed in Clark and Zaman (2013), these sorts of conditional forecasting exercises provide us some indication of how well a model is formulated. By conditioning on the historical path of the unemployment rate (thereby removing this source of uncertainty), we can assess how well the rest of the model (i.e., the inflation subcomponents) translates this information into inflation pressures and correctly captures the persistence of inflation. If the model is well-constructed and the unexpected shocks to the inflation components are not too large over the forecast horizon, then the conditional forecast of inflation should be close to the actual inflation path. As noted above, the conditional forecast of inflation coming from our nonlinear model does a respectable job tracking actual inflation over a 10-year period, and clearly outperforms its more standard counterpart, lending strong support to our model and providing reassurance about its ability to accurately provide conditional forecasts of the sort conducted in this study.

¹⁹ This model has 5 coefficients: a constant, coefficients on lags 1, 2, and 5, and the Phillips curve coefficient. For the importance of including the fifth lag, see the appendix to Verbrugge and Zaman (2023).




Figure D.1: Historical 10-year forecast from the model and from a conventional Phillips curve. Both are conditional recursive forecasts. The models see no inflation data after 2006Q4, but the forecasts are conditioned on the evolution of the unemployment rate.