Federal Reserve Bank of Cleveland Working Paper Series

The Hard Road to a Soft Landing: Evidence from a (Modestly) Nonlinear Structural Model

Randal J. Verbrugge and Saeed Zaman

Working Paper No. 23-03

January 2023

Suggested citation: Verbrugge, Randal J. and Saeed Zaman. 2023. "The Hard Road to a Soft Landing: Evidence from a (Modestly) Nonlinear Structural Model." Working Paper No. 23-03. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202303.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

The Hard Road to a Soft Landing: Evidence from a (Modestly) Nonlinear Structural Model

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.

December 30, 2022

Abstract

What drove inflation so high in 2022? Can it drop rapidly without a recession? The Phillips curve is central to the answers; its proper (nonlinear) specification reveals that the relationship is strong and frequency dependent, and inflation is very persistent. We embed this empirically successful Phillips curve – incorporating a supply-shocks variable – into a structural model. Identification is achieved using an underutilized data-dependent method. Despite imposing anchored inflation expectations and a rapid relaxation of supply-chain problems, we find that absent a recession, inflation will be more than 3 percent by the end of 2025. A simple welfare analysis supports a mild recession as preferred to an extended period of elevated inflation, under a typical loss function.

JEL codes: E31, E32, E52, C32
Keywords: nonlinear Phillips curve, frequency decomposition, supply price pressures, structural VAR, nonlinear impulse response functions, welfare analysis

Acknowledgments: We thank Rick Ashley, David Berger, Mark Bils, Oli Coibion, Andrew Figura, Ana Galvao, Ed Knotek, Loretta Mester, Katia Peneva, Juan Rubio-Ramirez, Ayşegūl Șahin, Arsenios Skaperdas, Jim Stock, Brad Strum, Michael Weber, 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.

1. Introduction

This paper offers compelling answers to two prominent and important questions facing macroeconomists and policymakers today. The first question is: What caused the sharp run-up in inflation in 2021-22? Reis (2022) argues, "The rise in inflation in 2021-22 is such a dramatic event that it will likely spur a large literature and a heated debate in trying to explain it over the next many years." The second question is an important and classic one in macroeconomics: How long will it take for inflation to get back to target, and what are the prospects for a soft landing?[1] There is great concern today that, barring a fairly rapid return of inflation to target, inflation expectations might become unanchored.[2] US monetary policymakers cited worsening inflation expectations ahead of the June Federal Open Market Committee (FOMC) meeting as part of the reason for the 75 basis point increase in the fed funds rate. And, as we note below, there is sharp disagreement as to whether inflation will return to target rapidly without a recession.

These issues are intimately related. For instance, in late 2021 it was widely believed that transitory pandemic-related supply-chain pressures were the "smoking gun" that was chiefly responsible for driving up core inflation. Once these pressures eased, it was thought, inflation in the price of goods would thereby ease as well, and overall inflation would return rapidly to target. But in early 2022, it became evident that not only were supply-chain pressures more persistent than expected, but also inflation had become much more broad-based, extending beyond tradable goods and gasoline into services. For instance, the shelter service inflation rate – a strongly cyclical and persistent series that is a reliable signal of the central tendency of inflation – rose in 2022 to levels not seen in decades, amid a backdrop of GDP growth that was tepid at best.

Will a recession be necessary, or even more than mildly helpful, in achieving a return to low and stable inflation? One viewpoint is that long-term inflation expectations exert a strong force on inflation, so that a recession is not necessary; and a complementary view is that the Phillips curve relationship is weak, so that a recession is not very helpful. At least in early June, FOMC Chair Jerome Powell was optimistic about the prospects of a return to low inflation absent a recession. In congressional testimony on June 22, he stated, "We're not trying to provoke, and don't think we need to provoke, a recession."[3]

But some others are not nearly so upbeat. For instance, former Treasury Secretary Lawrence Summers believes that, just like in the early 1980s, a very long, or very sharp, recession will be

[1] In Appendix 4, we also demonstrate that the same model can forecast/explain the inflation dynamics of the financial crisis and recovery (2007-2019) – and thus, for example, it explains the missing disinflation, and has no missing inflation puzzle. We further find that the Phillips curve is quite strong and that it did not weaken after 2006.
[2] Binder, Janson, and Verbrugge (2022) find that the degree of anchoring of long-run inflation expectations in the Survey of Professional Forecasters is overstated. See also Reis (2021, 2022). The present paper, however, proceeds under the assumption that inflation expectations are (and remain) anchored, in a specific sense to be explained below.
[3] Figura and Waller (2022) also argue that a soft landing in the labor market is possible.

necessary to cool off inflation.[4] Similarly, Ball, Leigh, and Mishra (2022) argue that absent the realization of "quite optimistic assumptions," two years of 7.5 percent unemployment will be necessary.[5]

We provide evidence that neither viewpoint is correct. Inflation is very persistent, and anchored inflation expectations exert only a modest force; so, contra those who believe a soft landing is achievable, we demonstrate that a recession will be necessary to bring PCE inflation down to anywhere near 2 percent by the end of 2025. However, the recession need not be "very long or very sharp."

This evidence is informed from a simple structural model. The model incorporates two elements that we view as particularly crucial for providing reliable answers to our questions (under the assumption that inflation expectations remain anchored).[6] First, to address "how we got here," to capture current or recent upward or downward forces on inflation, and to determine whether inflation is likely to reverse quickly, our model incorporates a variable that reflects supply-chain price pressures.[7] Second, our framework carefully accounts for the relationship between the state of the labor market and inflation, which, following Ashley and Verbrugge (2022a), we find to be nonlinear. Getting this relationship approximately right is absolutely crucial for understanding both how we got here and the prospects going forward, and, more generally, for understanding how monetary policy is related to inflation.

Why? Because conventional Phillips curve specifications, which are linear, yield estimates that indicate a rather weak and unstable relationship between unemployment and inflation, and underestimate the persistence of inflation. These findings, if true, have fundamental policy implications. If the unemployment rate is only weakly connected to the inflation rate, then overheating has little influence on inflation, and a recession likewise does little to bring down high inflation. As noted above, many economists assert that well-anchored inflation expectations will accomplish a fairly rapid deceleration of inflation. We agree that *without* well-anchored expectations, it will be challenging for inflation to return to target; see Mankiw and Reis (2018) and Reis (2021). But as both recent history (2012-2018) and our findings below indicate, well-anchored expectations are not sufficient to imply that inflation returns with necessary speed to the inflation target. Hence, to get inflation to descend back to target in a timely fashion, a recession is *necessary* – and if the Phillips curve is weak, then this suggests (à la Larry Summers) that this recession must either be very sharp or very prolonged.

[4] See Bloomberg.com, June 20, 2022: "Larry Summers Says US Needs 5% Jobless Rate for Five Years to Ease Inflation", and Domash and Summers (2022).
[5] Blanchard, Domash, and Summers (2022) also state that the unemployment rate must increase beyond 4.9 percent in order to reduce inflation, but they do not attempt to quantify exactly how far. For a contrary view, see Figura and Waller (2022).
[6] Weakened anchoring would almost certainly further slow inflation's projected sluggish deceleration. Binder, Janson, and Verbrugge (2022) provide evidence that among SPF respondents, long-run inflation expectations have become less strongly anchored in the recent period.
[7] Reis (2022), and many others, call attention to the importance of this factor. Harding, Lindé, and Trabandt (2022) do not include intermediate inputs in their New Keynesian model, though this model does allow for cost-push shocks to play a role.

However, these linear model specifications are at odds with the data, as demonstrated by Stock and Watson (2010) and Ashley and Verbrugge (2022a). These authors find compelling evidence, verified in the present study, that the Phillips curve is not a single (linear) relationship. Rather, it consists of two distinct relationships: one associated with a recessionary force, and the other with an overheating force.[8] In between these business-cycle phases – i.e., during the recovery period – the Phillips curve effectively disappears.[9] Furthermore, absent the influences of supply-chain problems and unemployment fluctuations, inflation is otherwise estimated to be extremely persistent. These features of the data are of central importance to current monetary policy decision-making. A specification that wrongly insists on a single (linear) relationship between inflation and unemployment is bound to underestimate the persistence of inflation, the inflationary impact of the labor market overheating, and the disinflationary impact of a recession, and overestimate the speed at which inflation will return to target. It will thus provide a highly misleading picture of the current situation and of prospects going forward.

Our model, a nonlinear analogue of a structural vector autoregression, consists of four equations in four variables: a supply price pressure variable; inflation (as measured by trimmed mean PCE inflation, an indicator that dominates core PCE along many dimensions[10]); and two "components" of the jobless unemployment rate: a persistent (or low-frequency) component and a moderately persistent (or medium-frequency) component.[11] Our partitioning of the jobless unemployment rate into varying persistence components is motivated by 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. As we explain below, this simple partitioning allows us to uncover a very insightful nonlinear relationship.

---

[8] Stock and Watson (2010) locate the recessionary force but not the overheating force. Ashley and Verbrugge (2022a) demonstrate that the Phillips curve relationship is frequency-dependent – i.e., that inflation responds differently to highly persistent movements in the unemployment rate than it does to less persistent fluctuations. Frequency-dependence may sound exotic, but it has deep roots in empirical macroeconomics. Already in the 1960s, Friedman (1968) and Phelps (1968) established that the most persistent fluctuations in the unemployment rate (i.e., natural rate fluctuations) are unrelated to inflation. Below we review a nascent body of research that is re-exploring the frequency domain for clues about business cycle drivers and dynamics. For instance, Beaudry et al. (2020) highlight medium-frequency cycles and argue that it is crucial to examine the properties of the data at different frequencies in order to discriminate across classes of models.

[9] Much extant theory (reviewed in Ashley and Verbrugge, 2022a) is consistent with these findings. For instance, models with capacity constraints (e.g., Kuhn and George, 2019) naturally deliver overheating price pressures, while models such as Gilchrist et al. (2017) and Hong (2019) produce sharp price drops at the onset of recessions.

[10] See Verbrugge (2022), who provides evidence and cites a wide range of research.

[11] We considered the remaining component, transient fluctuations, but since these were found to be unimportant predictors, to keep our model parsimonious, we abstract from these fluctuations. Similarly, we eschew the use of a fed funds rate variable (either real or nominal) for two reasons. First, small models are typically better for forecasting (and, accordingly, offer more reliable inferences); hence parsimony is preferred. Second, in the aftermath of the Great Recession, the nominal fed funds rate was at zero, but the FOMC made use of unconventional monetary policy. Since we omit policy variables, our results are conditional on monetary policy acting as it has in the past.

Each equation in our model is linear in its arguments, except for the trimmed mean PCE equation. In that equation, we admit sign asymmetry and establish (following Ashley and Verbrugge, 2022a) that inflation 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. As we explain below, historically, these portions of the two components align closely with overheating and recession, respectively. We model trimmed mean PCE inflation in terms of its deviations from the Survey of Professional Forecasters' median 10-year PCE inflation expectations. This serves several purposes: it effectively imposes anchored long-run inflation expectations; it serves to focus attention on shorter-term inflation dynamics (which are influenced by supply shocks and unemployment fluctuations and the like) and abstract from slow-moving inflation trends (which are ultimately determined by the longer-run goals of monetary policy); and it leads to better-performing inflation models; see Verbrugge and Zaman (2022a). Identification is achieved using the data-determined method of Swanson and Granger (1997), which substantially reduces the role of subjective elements. Nonlinear impulse response functions (IRFs) are generated following Kilian and Lütkepohl (2017).

Our five main findings are as follows. First, supply-chain pressures are found to exert a notable influence on inflation. Partly for this reason, our model is able to better explain the recent run-up in inflation; in particular, out-of-sample forecasts from our model over the past four quarters are far more accurate than those from alternative benchmark models. Because supply-chain pressures are important determinants of inflation, the projected easing of these price pressures plays a role in reducing inflation, although it is not, by itself, able to bring inflation back to target.

Second, our baseline forecast calls for a recession in the near term.

Third, while most of the IRFs are essentially linear, our estimated IRFs of inflation to identified structural shocks in the unemployment components are highly nonlinear and provide useful insights into the dynamics of inflation. For instance, consider a shock to the medium-frequency component of unemployment. For this variable, a positive shock that occurs when that component is positive yields an IRF that, mainly reflecting its direct influence on inflation, indicates a notable deceleration of inflation that peaks at 6 quarters. Conversely, a negative shock to this component, when the component is negative, yields an IRF that increases inflation and that peaks a year later (at 10 quarters). This upward movement in inflation is driven almost entirely by the movement this shock induces in the *low*-frequency component of the unemployment rate.

Fourth, our estimates imply that inflation is quite persistent and quite sensitive to (both components of) the unemployment rate. In particular, our model projects that inflation only very gradually falls back to 2 percent. Moreover, this baseline inflation projection is strongly influenced by the projected imminent recession. This recession will both *remove* an upward force on inflation – one that stems from overheating – and *impose* a downward force on inflation – one that stems from recessionary pressures.[12] In contrast, a counterfactual "soft landing" simulation, described immediately below, indicates that without this recession, inflation would remain much more elevated for longer. This latter path is accompanied, of course, by an increased risk of de-anchoring inflation expectations.

Lastly, we perform a simple welfare analysis using a standard quadratic loss function that abstracts from any risk of expectations de-anchoring. We first compare the baseline forecast, which calls for a recession, to a "soft landing" alternative, and to a "realistic hard landing" alternative. The soft landing scenario takes the June Summary of Economic Projections' (SEP) projected path of the unemployment rate; we condition on this path (and on our baseline projection for the supply-chain price pressure variable) and form projections of inflation. Inflation is very persistent. Under this scenario, inflation is projected to decline very slowly, so that it does not fall below 3 percent until mid-2026, and is still nearly 2.6 percent by the end of 2028. We also explore a "realistic hard landing" scenario because, from a historical perspective, the baseline recession projected by our model is quite pessimistic and unlikely. Accordingly, we choose a more likely recession scenario, one that is informed by the recession of 2001. This scenario has a projected path of inflation that generally lies between the baseline path and the soft landing path. Under this scenario, inflation hits 3 percent by early 2025, although given the persistence of inflation, its progress toward the policy target slows dramatically after that. The realistic hard landing path is preferred by our welfare criterion to both the baseline path and the soft landing path.

2. Data

We use quarterly data spanning from 1985:Q1 through 2022:Q3. 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

To effectively capture the supply-chain pressures, we consider two possible variables. The first is a domestic transportation price index, which is constructed using seven monthly variables: 5 producer price index (PPI) and two non-PPI variables. The PPI variables are obtained from the Federal Reserve Economic Data (FRED) database of the Federal Reserve Bank of Saint Louis (fred.stlouis.org). The two non-PPI variables are the Cass Freight Index: Expenditures and the

[12] While our model features a strong inflation response to overheating, it is possible that – owing to an unusual shift in the Beveridge curve – the persistent unemployment gap is being understated at the moment. For instance, Blanchard, Domash, and Summers (2022) argue that a shift in the Beveridge curve has caused the natural rate to rise by 1.3 percentage points from its pre-COVID level. If we were to use this higher estimate of the natural rate rather than the rate from Zaman (2022), the estimated upward force on inflation from our model would be notably larger.

Cass Freight Index: Shipments, both available from Cass Information Systems, Inc. starting in 1990:M1 onward. We construct an implicit price index by dividing the expenditures by shipments. The PPI variables are all PPI by industry: Scheduled Freight Air Transportation Services (1987:M12 onward), Line-Haul Railroads (1984:M12 onward), Deep Sea Freight Transportation Services (1988:M6 onward), General Freight Trucking: Long-Distance Truckload, Truckload (1992:M6 onward), and General Freight Trucking: Long-Distance Less Than Truckload (1992:M6) onward. The second supply-chain pressures variable we considered is the PPI for core intermediate goods (PPI-IG). As discussed below, we found the latter variable to be modestly superior.

Our inflation variable is the trimmed mean PCE (trPCE) inflation rate, from the Federal Reserve Bank of Dallas (based upon data underlying the PCE price index from the Bureau of Economic Analysis). This variable, found to be a very useful signal of trend inflation in many studies,[13] has a more stable relationship to unemployment than does core PCE.[14] We impose anchored long-run inflation expectations by modeling trPCE as the gap between trPCE inflation and the Survey of Professional Forecasters' (SPF)10-year inflation expectations.[15]

Both the persistent and moderately persistent components of the unemployment rate are derived from the *jobless* unemployment rate of Hall and Kudlyak (2022).[16] 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. Indeed, a small literature has developed regarding how best to address the extreme outliers in the 2020 data. 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.

[13] See, e.g., Mertens (2016), Verbrugge (2022), and Verbrugge and Zaman (2022a). Limited-influence inflation estimators are often found to be good signals of trend inflation; see, e.g., Meyer, Venkatu, and Zaman (2013).
[14] This is partly because core PCE can be dominated by anomalous price movements in nonmarket goods, as it was between late 2009 and 2010; Verbrugge (2022). See also Ball and Mazumder (2020) and Ashley and Verbrugge (2022a) for evidence regarding the stability of the Phillips curve when specified in terms of trimmed mean PCE or median PCE.
[15] SPF data are available from the Federal Reserve Bank of Philadelphia.
[16] The data necessary to construct the jobless unemployment rate are available from the Bureau of Labor Statistics. Ball, Leigh, and Mishra (2022) provide an alternative approach to answering the big questions in our paper. Inspired by their work, we explored the usefulness of the ratio of job vacancies (V) to the overall unemployment (U) rate as an additional variable. (Historical vacancy data were obtained from Barnichon, 2010) We found, however, that the V/U ratio did not provide additional explanatory power for inflation over the 1985-2019 period; so we did not pursue this avenue further.

After extracting the aforementioned components from the jobless unemployment data, we transform the persistent component into a gap, where the gap is defined as the deviation of the persistent component of the jobless rate from the estimate of the natural rate of unemployment (i.e., U*(t)) obtained from Zaman (2022). Our results are robust had we instead used the estimate of U*(t) from the Congressional Budget Office (CBO).[17] Similar to Ashley and Verbrugge (2022a), this transformation permits us to explore sign asymmetry in the inflation response to the permanent component.

3. Econometric Method and Identification

3.1 Construction and selection of the supply-chain price pressures variable

We construct our transportation index using a robust principal components method on the seven price components over the sample period 1994:M6-2022:M2. We first compute 12-month growth rates for each price index, then normalize each series by subtracting its median and dividing by its median absolute deviation (multiplied by 1.48). Next, we compute a robust covariance matrix based upon series correlations estimated using the robust rom-rho method (see Chakhchoukh et al. 2010, and its refinement in Garciga and Verbrugge 2021).[18]

The loadings (parameter values) on the seven series are close to equal in size, and the resulting transport price index is similar whether we use unequal loadings or equal loadings (their correlation is 0.96), so we choose equal weights for simplicity. To extend the index before 1994:M6, in each month, we apply equal loadings on all series available at the time. For instance, before 1989:M6, only the line haul railroad price index was available, so it receives 100 percent of the weight over that period.

We also examine the properties of the producer price index for core intermediate goods (PPI-IG) as a price-pressures variable. Presumably, supply-chain price pressures will be quickly reflected in the prices of intermediate goods. As it turns out, this variable is indeed highly correlated with the domestic transport price index (rho = 0.77). The two indexes have similar dynamics (see Figure 1, which plots the PPI-IG and the transportation index, rescaled so that both have the same variance), and both indexes have predictive content for inflation. However, we found that the PPI-IG has more predictive content. The PPI-IG also features no implicit break in the sample caused by missing price indexes before 1994:M6. Hence, we settle on the PPI-IG as our measure of supply-chain price pressures.

---

[17] The Zaman (2022) estimate pertains to aggregate unemployment, not jobless unemployment. We estimate the persistent part of the time-varying gap between these series, abstracting from the massive spike in temporary layoffs in the pandemic period by replacing these observations with a projection starting in 2020:Q1. We subtract this gap estimate from U*(t) to obtain a "jobless U*(t)" estimate. This estimate is available from the authors upon request.
[18] The covariance matrix we obtained was positive semi-definite, so we did not need to use the renormalization step.

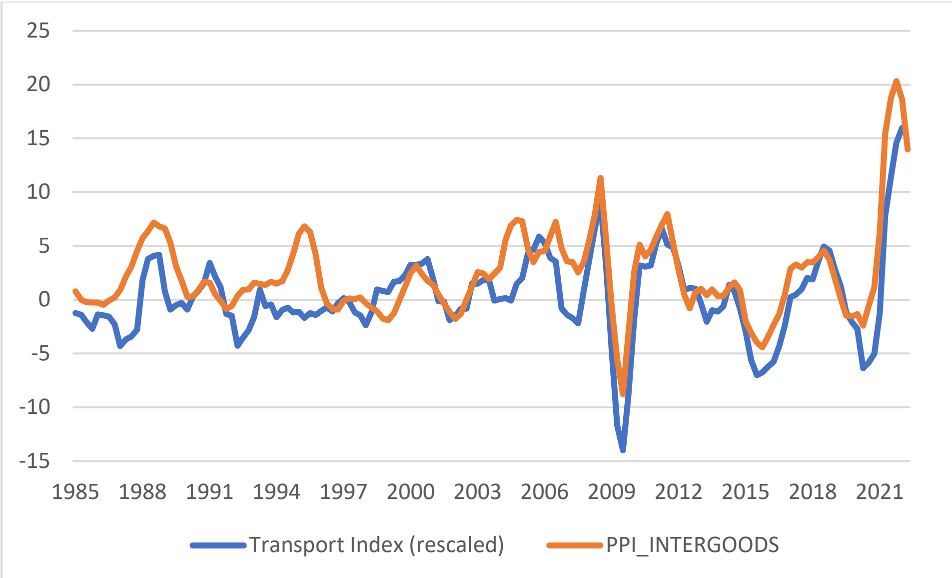

Figure 1: PPI-IG and the (rescaled) transportation index.

3.2 Frequency-dependent regression

Frequency-dependence (or persistence-dependence) in a relationship between two time-series variables *Y* and *X* implies that *Y* responds differently to persistent fluctuations in *X* than it does to transitory fluctuations in *X*. Theory suggests that many relationships in macroeconomics are frequency-dependent. For instance, the permanent income hypothesis implies that consumption responds much more strongly to persistent movements in income than it does to transient movements. "Real-business-cycle" modeling was built upon the presumption that business cycle relationships are distinct from low-frequency relationships, and this idea has recently regained traction.[19] The motivation for seasonal adjustment is the notion that the relationship between time series is different at seasonal frequencies than at other frequencies. Transient measurement error also induces frequency-dependence. For instance, suppose that the relationship between *Y* and *X* is linear, but the econometrician only observes $Z = (X + e)$, where *e* is a transient measurement error. Then the relationship between *Y* and *Z* will be frequency-dependent, approaching the *Y*-*X* relationship for low frequencies but approaching 0 for higher frequencies (see Hannan (1963), who first highlighted this issue, and Cochrane (2018) and Ashley and Verbrugge (2015) for recent applications).

[19] A nascent body of research – building upon earlier work by Comin and Gertler (2006) – is exploring the frequency domain for clues about business cycle drivers and dynamics. Angeletos et al. (2020) have recently argued that assessing the root causes of business cycles requires delving into the drivers of unemployment, output, consumption, and the like at different frequencies. Beaudry et al. (2020) state: "Therefore, in order to evaluate business cycle properties, one needs to find a way to extract properties of the data that are unlikely to be contaminated by the lower-frequency forces that are not of direct interest." Below, we note how standard approaches will typically lead one astray.

Thus, persistence-dependent regression (or frequency-dependent regression) has a long history.[20] But until Ashley and Verbrugge (2009), it was subject to a rather severe drawback: as argued in that study, and proved in Ashley and Verbrugge (2022b),[21] previous approaches led to inconsistent parameter estimates. In particular, the application of a two-sided filter, when the series are in a feedback relationship, yields coefficient estimates that are biased and inconsistent.[22] This distortion occurs irrespective of whether the analysis is conducted in the time domain or in the frequency domain.[23]

The Ashley/Verbrugge "persistence-dependent regression" methodology used herein was developed in Ashley and Verbrugge (2009) and Ashley et al. (2020) and is briefly reviewed in Appendix 2. This paper is the first application of this method to a structural model.

To fix ideas, begin with a standard reduced-form Phillips curve specification, where the dependent variable is trPCE inflation, and the explanatory variables are lags of inflation and an unemployment rate gap, and we wish to decompose the coefficient on the unemployment rate gap by persistence level (frequency).

The first step is to *partition* the real-time unemployment rate u_t into *m* persistence components – $u_t = \sum_{j=1}^{m} u_{j,t}$ – which *by construction* add up to the original series. These *m* 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 *m* persistence components, estimating a separate coefficient for each. (We note in passing that we subtract u_t^* from the most persistent component.) Simulation evidence in Ashley, et al. (2020) and Ashley and Verbrugge (2022c) indicates that the method yields reliable coefficient estimates and inferences, for both linear and nonlinear data-generating processes. (See Appendix 2 for more details, which summarizes why

[20] See Ashley and Verbrugge (2022a) for a review of numerous recent studies locating evidence for persistence-dependence in macroeconomic relationships.
[21] This paper builds upon a result in Sargent (1987). See also Doppelt (2021) for a related result; this author demonstrates the rather severe drawbacks of seasonally adjusted data for structural modeling.
[22] Even absent feedback, transfer function gain and phase plots are substantially more challenging to interpret than the Ashley/Verbrugge approach; and in the presence of feedback, Granger describes the interpretation of such plots as "difficult or impossible" (Granger, 1969). For interesting examples of the rather dramatic failures that two-sided filtering can induce, see Ashley and Verbrugge (2022c). In that paper, we also compare the Ashley and Verbrugge (2009) method to that of Hamilton (2018) and demonstrate that the Ashley/Verbrugge method is generally superior.
[23] Frequency-dependence in a regression is distinct from variation of gain, phase, or coherence by frequency. Those can vary by frequency even with a linear data-generating process. For more details, see Ashley and Verbrugge (2022a,c).

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).

In this paper, to partition the jobless unemployment rate while applying the Ashley/Verbrugge method, we use the Iacobucci-Noullez (2005) filter, setting $k = 4$ (i.e., using four quarters of univariate forecasts in each rolling window). The Iacobucci-Noullez filter introduces no phase shift (unlike, e.g., the Christiano-Fitzgerald (2003) filter).[24] Following Ashley and Verbrugge (2022a), we set m =3 and choose frequency cutoffs so that the jobless unemployment rate is partitioned into fluctuations lasting longer than 4 years (termed u_t^{lowgap}, for a 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 and, especially, much more sharply than usual. Meanwhile, the medium-frequency component also fell very sharply back to zero, and is currently at historic lows.

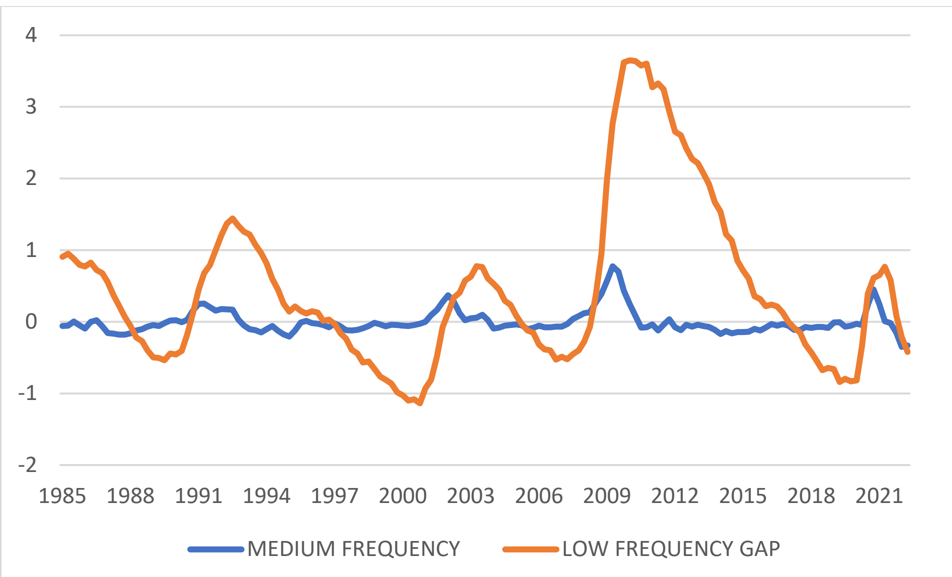

Figure 2: Two most persistent components of the jobless unemployment rate.

Because we are specifying a structural model, we accordingly specify and estimate a law-of-motion for each of these unemployment components separately. Similarly, we project each separately, and provide impulse response functions for each, separately. To obtain an estimated projection for the (overall) jobless unemployment rate, we simply add the two projections.

[24] 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.

3.3 Identification

We adopt the Swanson and Granger (1997) approach to identification.[25] 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 $e_X \perp e_Y \mid e_Z$. 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, a 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:

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

In matrix notation, the SVAR is denoted

$$A\mathrm{M}_t = B(L)M_t + V_t$$

where $\mathrm{M}_t \equiv (X_t, Y_t, Z_t)'$, $B(L)$ 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

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

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

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

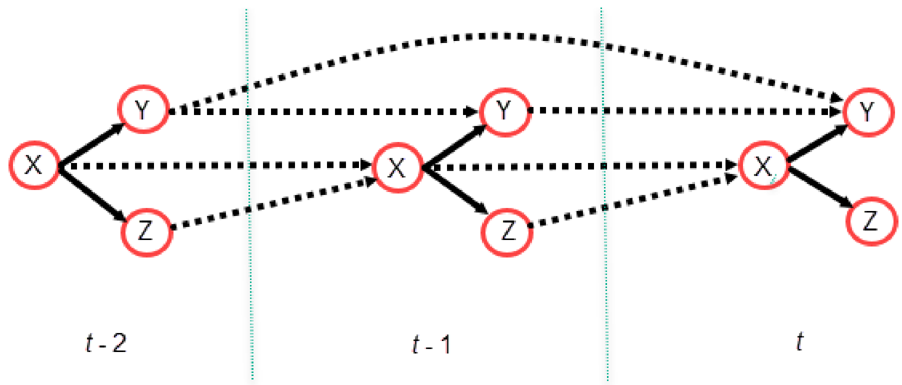

Figure 3: A structural VAR, with causal influences depicted. Solid lines depict contemporaneous causation; dashed lines depict intertemporal causation. Thus, for example, at time *t*, variable *Y* is influenced by variable *X* contemporaneously, by its own value at time *t*-1, and by its own value at 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 $\rho(e_Y, e_Z \mid e_X) = 0$.

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

---

[26] 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.

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 $e_X \perp e_Y \mid e_Z$, $e_X \perp e_Z \mid e_Y$, and $e_Z \perp e_Y \mid e_X$ 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.[27] 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 $e_X \perp e_Y \mid e_Z$ 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 of the final model.[28] (In models with numerous variables, one may formally test higher-order partial correlation constraints implied by the model.)

---

[27] 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 "cancelation." 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. There are 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.
[28] If more than one model appears equally reasonable, one may investigate the sensitivity of, e.g., IRFs to model choice.

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.

3.4 Forecasts, nonlinear impulse response functions, and error bands

Given the model's nonlinear nature, we construct impulse response functions, forecasts, and error bands via counterfactual simulations, following the procedure outlined in Kilian and Lütkepohl (2017), with shocks bootstrapped from estimated residuals. We plot the mean response, as well as the 10^{th} and 90^{th} percentiles from the simulations, as explained below.

In a nonlinear model, the IRFs will generally depend upon the size of the shock, its sign, and the initial conditions. Hence, theory should dictate which of these numerous IRFs to investigate/estimate, for any given variable *X*. Several types of nonlinear IRFs exist in the literature. Given that the nonlinearity in the present model relates to the different impacts of fluctuations in unemployment by frequency and sign, of most interest is a particular subset of IRFs:

- IRFs to positive and negative shocks to the medium-frequency component, conditioned on starting in a positive or negative initial condition, respectively.
- IRFs to positive and negative shocks to the low-frequency (highly persistent) component, conditioned on starting in a positive or negative initial condition, respectively.

Our estimation suggests that IRFs to other variables are more-or-less linear.

For computing the nonlinear IRFs to variable *X*, we make use of the notion that an IRF is the difference between a forecast with a particular shock and the forecast without one.

As noted above, we first decide upon the set of initial conditions that will be considered (e.g., for the medium-frequency component of the unemployment rate, one case is $u^{medFreq}(t) > 0$). Given that decision, our next step is to identify all dates satisfying the initial conditions. For each of *N* replications, we randomly sample from these dates as starting conditions. Then we produce *k*-period forecasts, conditional on bootstrap draws from the estimated regression residuals (for each period, drawing the vector of residuals over the four equations). This produces a "baseline" forecast for this initial condition and this set of bootstrap draws. Then we produce a second forecast, where we replace the initial (period *s*) bootstrap draw for the variable *X* by a given shock. In particular, we replace the bootstrap draw at date *s* by *scale**(δ_X), where *scale* is the sign and scaling factor, and δ_X corresponds to a specific shock for *X*, whose size is chosen to be sensible for *X*. The difference between the two forecasts, for each draw, is the impulse response function for that draw. This IRF is stored; then we begin the next replication. When all replications have been completed, we compute the mean and the 10th and 90th percentiles of the impulse response functions, to estimate error bands and their central tendency.[29]

Baseline forecasts and their error bands are computed similarly, although the "initial conditions" are fixed at the end of the sample, and there is no additional step to where an initial residual is replaced.

3.5 Specification

Given our aims in the present paper, and the fact that we use quarterly rather than monthly data, we use a modification of the baseline inflation equation of Ashley/Verbrugge (2022a). 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 BIC. In the PPI-IG and trPCE 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.

In the *PPI-IG* equation, the inclusion of *trPCE* inflation was rejected. However, *PPI-IG* 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 BIC.

$$\pi_t^{PPI} = \alpha^{PPI} + \sum\nolimits_{j=1}^{4} \beta_j^{PPI} \pi_{t-j}^{PPI} + \delta \Delta u_{t-1}^{medfreq} + e_t^{PPI} \tag{2}$$

In Equation (2) and hereafter, π_t^{PPI} refers to 12-month inflation in PPI-IG. Similarly, $\left(\pi_t^{trPCE} - \pi_t^{PTR}\right)$ refers to 12-month trPCE inflation, minus PTR, the 10-year expected PCE inflation rate in the SPF. Coefficient estimates are reported in Appendix 3, Tables A.2 and A.3.

In the trPCE equation, all of the other three variables were found to have a significant relationship. In these quarterly data, we found the same sort of sign asymmetry in the unemployment components as did Ashley and Verbrugge (2022a) (see Appendix 3); a single lag of each component sufficed. Our trPCE equation was specified as

$$\begin{aligned}\left(\pi_t^{trPCE} - \pi_t^{PTR}\right) = \alpha^{trPCE} + \beta_1^{trPCE}\left(\pi_{t-1}^{trPCE} - \pi_{t-1}^{PTR}\right) + \beta_5^{trPCE}\left(\pi_{t-5}^{trPCE} - \pi_{t-5}^{PTR}\right) \\ + \sum\nolimits_{j=1}^{2} \lambda_j^{trPCE} \pi_{t-j}^{PPI} + \gamma u_{t-1}^{-lowgap} + \phi u_{t-1}^{+medfreq} + e_t^{trPCE}\end{aligned} \tag{3}$$

where $u_{t-1}^{-lowgap}$ refers to the negative portion of the low-frequency unemployment gap, and $u_{t-1}^{+medfreq}$ refers to the positive portion of the medium-frequency component.

[29] In this study, mean IRFs were fairly insensitive to the size of the shock. However, IRFs to larger shocks were estimated more precisely. Kilian and Lütkepohl (2017) do not discuss this issue.

In both of our unemployment component equations, asymmetry in unemployment terms was rejected. Thus, the only equation featuring nonlinearity *in its dependent variables* is Equation (3), although all equations are nonlinear in the (overall) unemployment rate. Our $u^{medfreq}$ equation was specified as

$$u_t^{medfreq} = \sum_{j=1}^{4} \gamma_j^{med} u_{t-j}^{lowgap} + \sum_{j=1}^{2} \phi_j^{med} u_{t-j}^{medfreq} + \beta^{med} \left(\pi_{t-1}^{trPCE} - \pi_{t-1}^{PTR} \right) + \lambda^{med} \pi_{t-1}^{PPI} + e_t^{medfreq} \tag{4}$$

Finally, our u^{lowgap} equation was specified as

$$u_t^{lowgap} = \alpha^{lowgap} + \sum_{j=1}^{2} \gamma_j^{low} u_{t-j}^{lowgap} + \sum_{j=1}^{4} \phi_j^{low} u_{t-j}^{medfreq} + \beta^{low} \left(\pi_{t-1}^{trPCE} - \pi_{t-1}^{PTR} \right) + \lambda^{low} \pi_{t-1}^{PPI} + e_t^{lowgap} \tag{5}$$

4. Results

4.1 Identification

We estimated all pairwise correlations of the reduced form residuals. We found a significant correlation between *PPI-IG* and *trPCE* residuals, and between u^{lowgap} and $u^{medfreq}$ residuals; all other correlations were insignificant. This left us with 4 possible models. On the basis of economic theory and *a priori* timing grounds, we assume that PPI-IG→trPCE and $u^{medfreq}$ → u^{lowgap} contemporaneously, leading to the following loading matrix *A*:

$$A\mathrm{M}_t = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -a_{21} & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -a_{43} & 1 \end{bmatrix} \begin{bmatrix} PPI\text{-}IG_t \\ trPCE_t \\ u_t^{medfreq} \\ u_t^{lowgap} \end{bmatrix}$$

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 both a_{21} and a_{43} were statistically significant.

Given these results and the sparsity of the *A* matrix, to estimate the identified system, it suffices to modify the *trPCE* equation by including the contemporaneous *PPI-IG* term, modify the u^{lowgap} equation by adding contemporaneous $u^{medfreq}$ term, and estimate the (now fully identified) nonlinear system equation by equation.[30] Thus, the two respecified equations are

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

$$\left(\pi_t^{trPCE} - \pi_t^{PTR}\right) = \alpha^{trPCE} + \beta_1^{trPCE}\left(\pi_{t-1}^{trPCE} - \pi_{t-1}^{PTR}\right) + \beta_5^{trPCE}\left(\pi_{t-5}^{trPCE} - \pi_{t-5}^{PTR}\right) + \sum_{j=0}^{2} \lambda_j^{trPCE} \pi_{t-j}^{PPI} + \gamma u_{t-1}^{-lowgap} + \phi u_{t-1}^{+medfreq} + v_t^{trPCE} \tag{6}$$

$$u_t^{lowgap} = \alpha^{lowgap} + \sum_{j=1}^{2} \gamma_j^{low} u_{t-j}^{lowgap} + \sum_{j=0}^{4} \phi_j^{low} u_{t-j}^{medfreq} + \beta^{low}\left(\pi_{t-1}^{trPCE} - \pi_{t-2}^{PTR}\right) + \lambda^{low} \pi_{t-1}^{PPI} + v_t^{lowgap} \tag{7}$$

Further, in equations (2) and (4), the reduced-form residuals *e* are relabeled as structural residuals *v*.

For simulating the system – necessary for estimation of forecasts, IRFs, 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 $u^{medfreq}$.

$$u_t^{+lowgap} \equiv \max\left(0, u_t^{lowgap}\right) \tag{8}$$

$$u_t^{-lowgap} \equiv \min\left(0, u_t^{lowgap}\right) \tag{9}$$

$$u_t^{+medfreq} \equiv \max\left(0, u_t^{medfreq}\right) \tag{10}$$

$$u_t^{-medfreq} \equiv \min\left(0, u_t^{medfreq}\right) \tag{11}$$

$$\Delta u_t^{medfreq} \equiv u_t^{medfreq} - u_{t-1}^{medfreq} \tag{12}$$

The full structural model consists of equations (2) and (4) (with residuals *v*), and equations (5) through (12).

4.2 The recent run-up in inflation: A simple forecast comparison

A salient question is: How well does our model explain the recent run-up in inflation? This surge was a notable upside surprise to most forecasting models[31] and raises questions such as: If a given model did not predict well the recent surge in inflation, can we then be confident in its projections of a rapid deceleration in inflation going forward? We compare the absolute values of the one-quarter-ahead forecast errors from the present model to three standard benchmark forecasting alternatives: a simple random walk forecast (RW); a random walk forecast based upon the four-quarter moving average (RW-4Q avg); and a Stock and Watson (2007) unobserved components stochastic volatility (UCSV) model, estimated on trimmed mean PCE data.[32]

We display absolute forecast errors in Figure 4. While the present model had some notably large forecast misses in the last three quarters, they were far smaller than those of the benchmark

[31] See Lusompa and Sattiraju (2022), who find that sophisticated cutting-edge time-varying-parameter models performed worse over the run-up than did simpler forecasting models.
[32] The UCSV model is one of the two best-performing models in the forecasting comparison of Lusompa and Sattiraju (2022).

models. Thus, there is reason to place some confidence in our model's projections going forward.

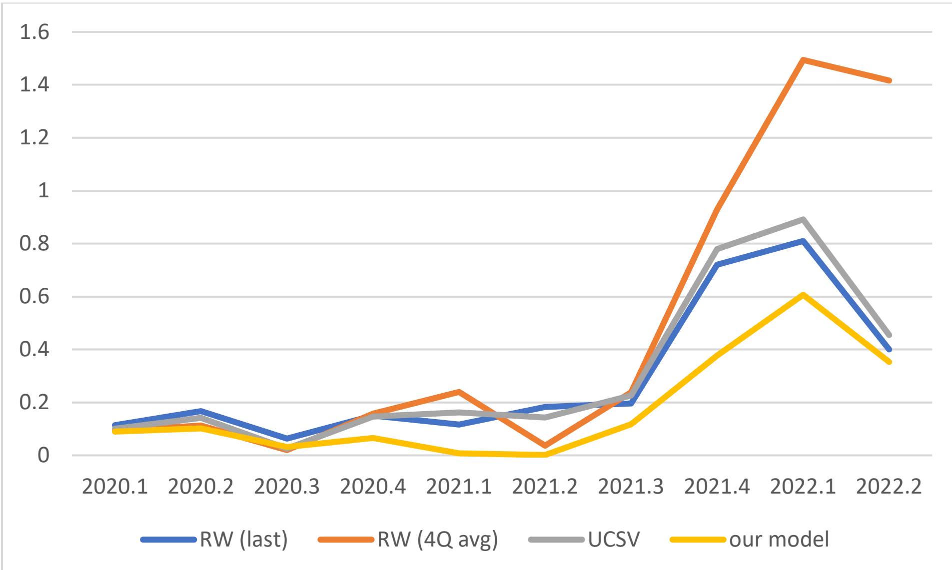

Figure 4: Absolute values of one-quarter-ahead forecast errors from four models

Table 1: RMSEs over this period

Model	RW-last	RW-4Q avg	UCSV	Our model
RMSE	0.39	0.73	0.42	0.26

We next display impulse response functions (IRFs), to give the reader a sense of how the model understands the historical data.

4.3 Impulse response functions

Most of the IRFs in the model are nonlinear, but only modestly so. The chief exceptions are the IRFs of trimmed mean PCE inflation to shocks to unemployment components, just as one would expect.

We begin by plotting the IRF of inflation to a large shock to PPI-IG inflation. In 2021:Q4, PPI-IG inflation reached 20.3 percent. We estimate the IRF of trPCE inflation to a shock to PPI-IG inflation of size +14 percentage points. We choose this size because when PPI-IG is hit with a shock of this magnitude, it rises to +20.4 percent in the following period. (As will be evident below, by period 7, PPI-IG inflation is back to near zero.) Our objective in this exercise is to assess how big a role the notable run-up in supply-chain price pressures in 2021 may have played in the surge in trimmed mean inflation in 2022.

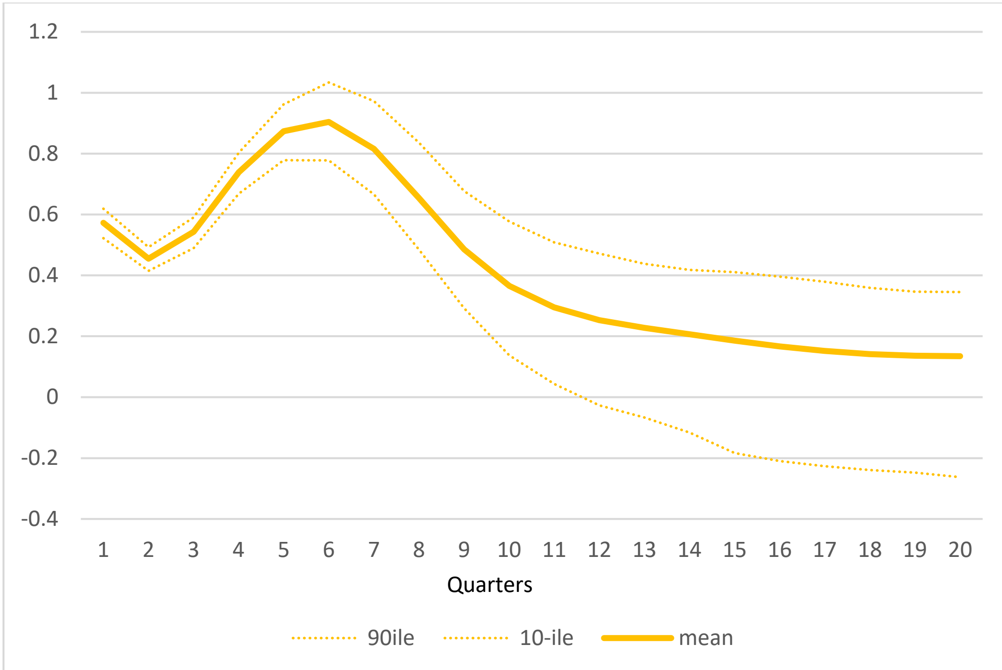

Figure 5: Response of trPCE inflation to a shock to PPI-IG inflation that eventually raises PPI-IG inflation by 20 percent.

In response to a shock that eventually raises PPI-IG inflation by 20 percent, trimmed mean PCE inflation rises and peaks at 6 quarters by about 0.9 percent, before falling gradually, over the next 6 quarters or so, back to its trend. The IRF suggests that supply-chain price pressures played a notable role in trPCE inflation in 2022; similarly, as price pressures ease, they will have a notable influence on reducing trPCE inflation.[33]

In Figure 6 we plot the IRFs of all four variables to a shock of size +2.0 percentage points to PPI-IG inflation; in panel (d), note that this eventually causes PPI-IG inflation to rise, in the next quarter, to about 3 percent. For inflation, because IRFs are slightly nonlinear, we also depict the mean IRFs of other shocks (+4.0 percentage points, -2.0 percentage points, and -4.0 percentage points), scaled to match a 2.0 percentage point shock, and reflected about the axis for negative shocks. We found that in contrast to mean responses, IRF error bands were only slightly nonlinear, so we only depict the error bands for a shock of size +2.0 percentage points. In response to PPI-IG shocks, trimmed mean inflation moves in the same direction. In response to a 2.0 percentage point shock to PPI-IG inflation, trPCE inflation rises immediately upon impact, dips slightly, then rises more to a peak response of about 0.13 percentage points at 6 quarters, and finally begins falling gradually – slightly lagging the PPI-IG path. For the mean response of inflation, nonlinearity is apparent from quarter 13 onward, depending upon size and whether the shock is positive or negative. This reflects the differential impact of positive versus negative movements in the low-frequency component of the unemployment rate. In response to a *positive*

[33] We performed a counterfactual conditional forecast, in which we set PPI shocks to 0. This forecast indicated that PPI shocks accounted for nearly half of the run-up of trimmed mean inflation by 2022Q3. In this context, it is worth noting that a higher U* estimate in the recent period would attribute more of the run-up in inflation to overheating.

PPI-IG shock, the low-frequency gap gradually rises, peaking 13 or 14 quarters after the shock, then gradually falls. Being positive, this movement has no influence on the inflation path. (The medium-frequency component rises as well, but evidently not enough to have much influence on the inflation path.) But in response to a *negative* PPI-IG shock, the low-frequency gap *falls*, which elevates inflation, i.e., moves it farther from zero. Upon reflection about the x-axis, it is thus more negative.

IRFs of other variables to this PPI-IG shock are essentially linear. In particular, responses of other variables to negative shocks were essentially mirror images of the responses to the positive shock, and this is true for error bands as well. We conclude that, aside from inflation, the other IRFs are essentially linear.

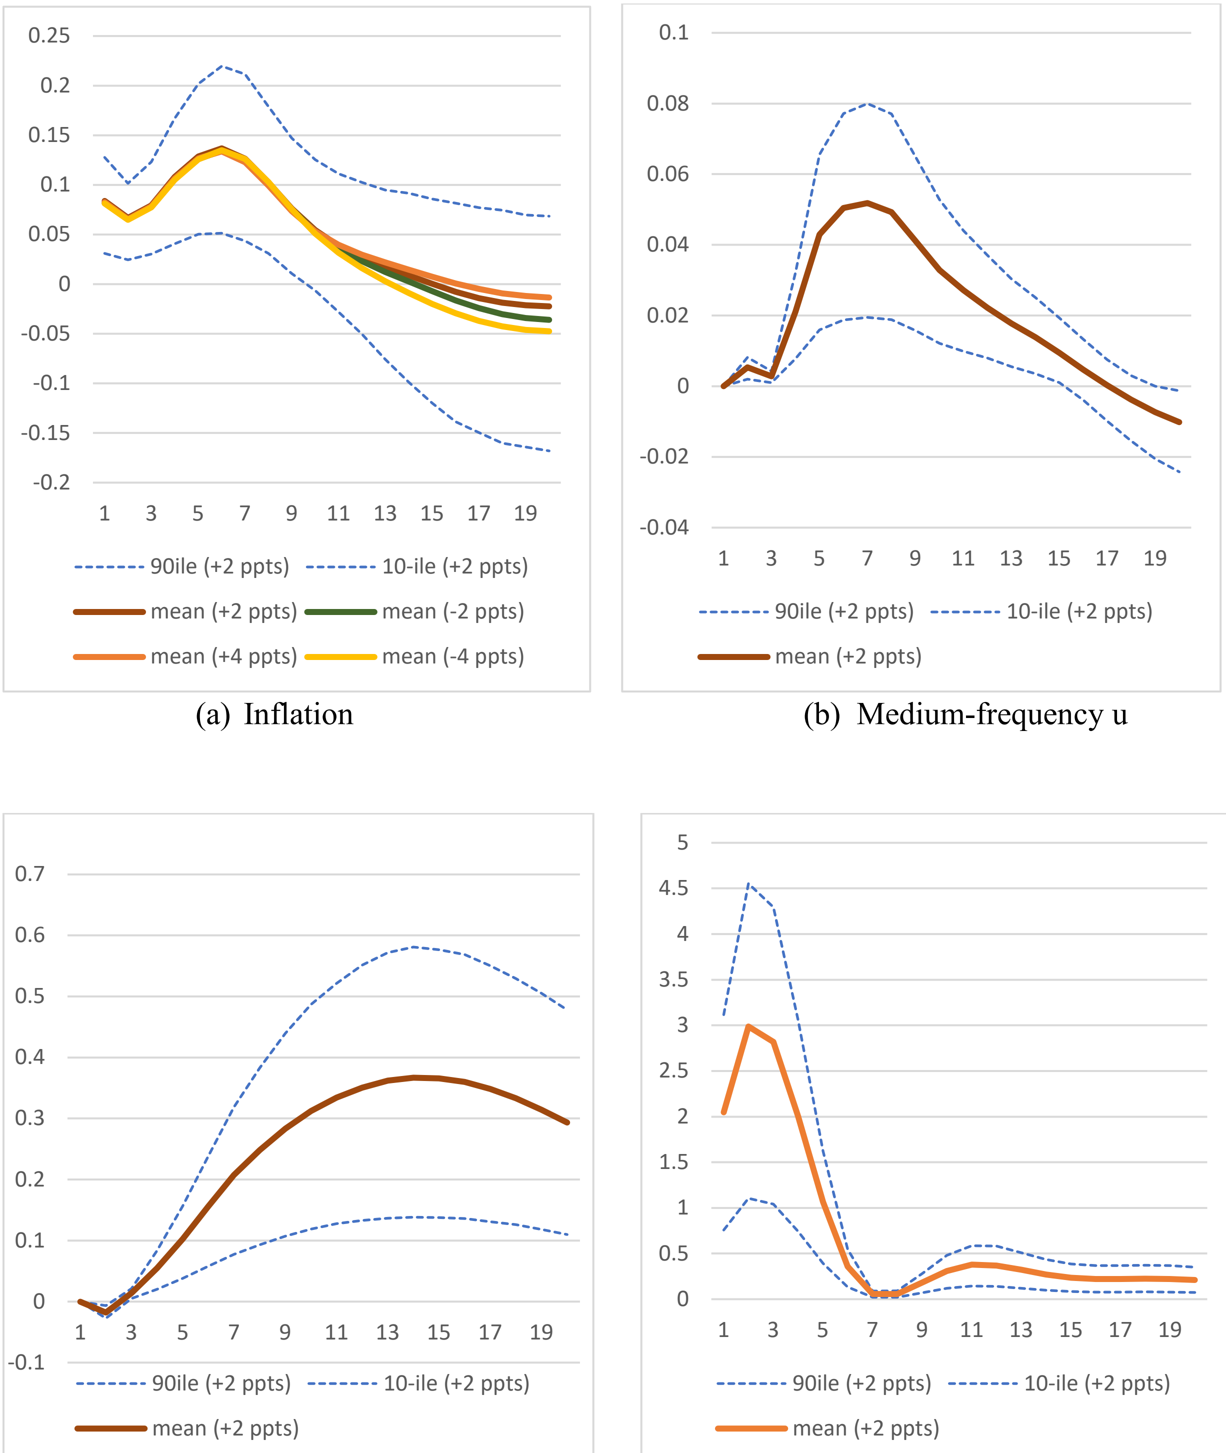

(a) Inflation

(b) Medium-frequency u

(c) Low-frequency u gap

(d) PPI-IG (+2 ppts)

Figure 6: IRFs of each variable to a +2 ppt. shock to PPI-IG inflation. The IRFs to negative shocks, in panel (a), are reflected about the x axis. IRFs (and error bands) are effectively linear, except for the mean response of inflation (panel (a)), which is slightly nonlinear.

In Figure 7, we plot the IRFs of all four variables to a shock of size +0.05 percentage points to the medium-frequency component of the unemployment rate, drawing from initial conditions such that this component is already positive and rising. We also depict IRFs to a shock of size -0.05 percentage points (drawing from initial conditions such that this component is negative); and in all panels except panel (a), we reflect these over the x axis, to highlight the modest nonlinearity.

In response to its own shock, the medium-frequency component rises further in the next quarter, remains elevated for an additional quarter, and then starts to fall. By quarter 9, it is back to zero, and it continues to fall so that it reaches about -0.03 by quarter 17. The shock to the medium-frequency component causes the low-frequency component to gradually rise; it tops out at about 0.42 percent in the 8^{th} quarter following the shock, and then gradually falls so that it is almost back to zero by quarter 20. (The response to a negative 0.05 percentage point shock causes the low-frequency component to fall by about 0.44 percent in the 8^{th} quarter.)

The PPI-IG inflation IRF is close to a mirror image of that of the medium-frequency component. There is a Phillips curve influence, as PPI-IG inflation initially decelerates in response to the positive medium-frequency shock, bottoming out at -0.75 percent by quarter 4, then begins to rise gradually so that it peaks at +0.44 percent by quarter 9.

In contrast to the other IRFs, the trPCE inflation IRF is clearly nonlinear. In response to a positive shock of +0.05 percentage points to the medium-frequency component, inflation falls, bottoming out at -0.17 percent at 6 quarters, after which it very gradually moves back to 0. This trajectory mainly reflects the direct influence on inflation of the medium-frequency component. Conversely, in response to a negative shock of the same magnitude, inflation begins to slowly rise. It reaches a maximum of +0.17 percent, 10 quarters after the shock. In this case, the inflation trajectory reflects the influence of the *low*-frequency component. A negative shock to the medium-frequency component has no direct influence on inflation but causes the low-frequency gap to become negative, which exerts upward pressure on inflation.

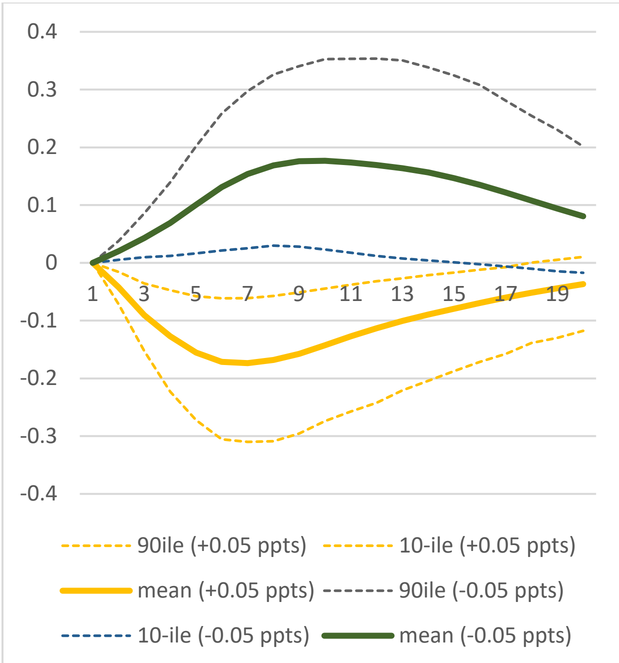

(a) Inflation

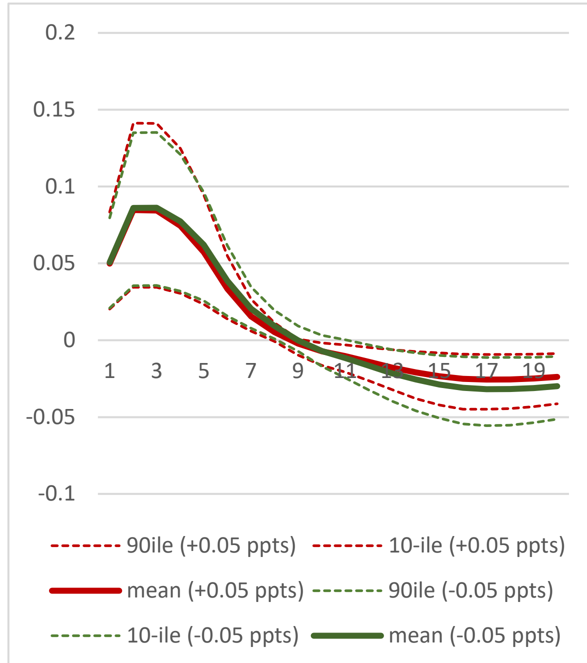

(b) Medium-frequency u

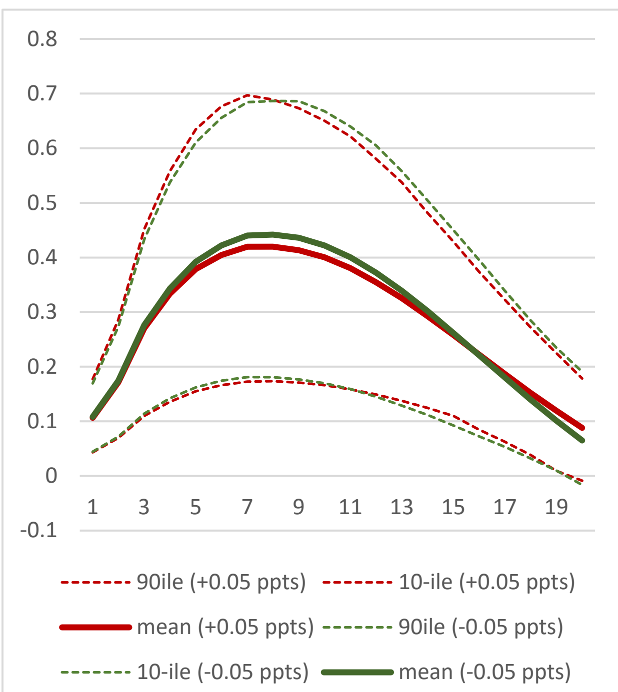

(c) Low-frequency u gap

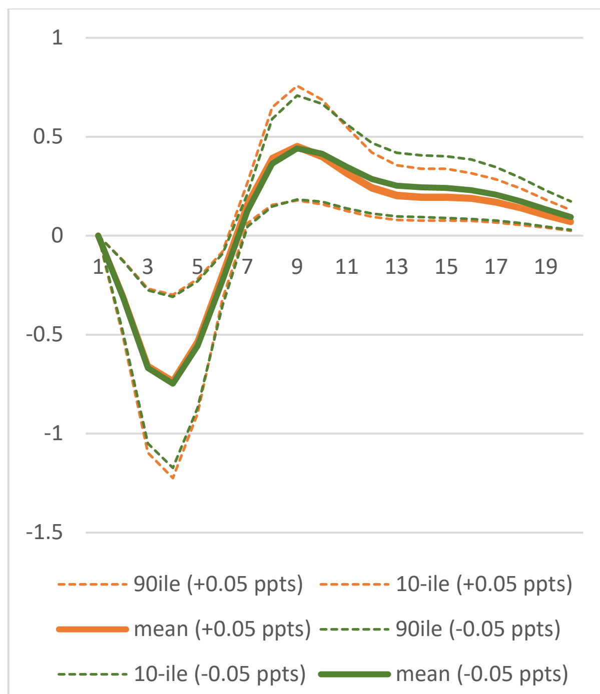

(d) PPI-IG

Figure 7: IRFs of each variable to a +/-0.05 shock to medium-frequency unemployment rate. Initial conditions for + shocks: medium-frequency component positive and rising; for – shocks, medium-frequency component negative. The IRFs to negative shocks, except in panel (a), are reflected about the x axis. All IRFs and error bands display nonlinearity, to a greater or lesser extent.

In Figure 8 we plot the IRFs of all four variables to a shock of size -0.20 percentage points to the low-frequency component of the unemployment rate, drawing from initial conditions such that this component is already negative. We also depict IRFs to a shock of size +0.20 percentage points (drawing from initial conditions when this component is positive); and in all panels except panel (a), we reflect these over the x axis, to highlight the modest nonlinearity.

In response to its own shock, the low-frequency component falls further over the next quarter, with a peak response of -0.52 percent in quarter 5; then the component starts to move, very gradually, back to 0. The IRF to a +0.2 percentage point shock is not quite a mirror image, in that the peak response is +0.54 percent, coming in quarter 7, and the component moves even more slowly back to 0 in this case. The reason for this "discrepancy" is that, as discussed below, in response to a negative shock, inflation rises firmly; and this movement in inflation tends to pull low-frequency unemployment up, thereby causing it to return more rapidly to 0. In response to a positive shock, inflation falls by a lesser extent, and so low-frequency unemployment is subject to a weaker (downward) influence from inflation.

The -0.2 percentage point shock to the low-frequency component causes the medium-frequency component to fall as well, with a peak response of -0.12 percent in quarter 3, after which this component rapidly returns to 0 by quarter 7, and overshoots modestly, peaking at 0.05 percent in quarter 16. (The IRF to a +0.2 percentage point shock is not quite a mirror image: the overshoot is a tad more modest and peaks in quarter 18.) The IRFs of PPI-IG inflation are nearly mirror images; in response to a -0.2 percentage point shock to the low-frequency component, PPI-IG inflation rises to peak at +1 percentage point in quarter 4, then falls to 0 by quarter 7 and overshoots to -0.6 percent in quarter 9, before gradually falling to 0.

The trPCE inflation IRF is clearly nonlinear and features strong Phillips curve dynamics. In response to a shock of -0.2 percentage points, inflation rises to peak at 0.27 percent in quarter 8, before falling gradually. This trajectory mainly reflects the direct influence on inflation of the low-frequency component, although PPI-IG inflation plays a supporting role. Conversely, in response to a positive shock of the same magnitude, inflation begins to slowly fall, to a lesser extent. It reaches a maximum of -0.12 percent, 8 quarters after the shock. In this case, the inflation trajectory reflects the influence of the *medium*-frequency component, with PPI-IG inflation playing a supporting role.

(a) Inflation (b) Medium-frequency u

(c) Low-frequency u gap (d) PPI-IG

Figure 8: IRFs of each variable to a +/-0.20 ppt shock to low-frequency unemployment rate. Initial conditions for + shocks: low-frequency component positive; for – shocks, low-frequency component negative. The IRFs to negative shocks, except in panel (a), are reflected about the x axis. All IRFs and error bands display nonlinearity, to a greater or lesser extent.

In Figure 9, we plot the IRFs of all four variables to a shock of size +0.25 percentage points to trimmed mean PCE inflation. We also depict IRFs to a shock of size -0.25. In all panels, we reflect these over the x axis, to highlight the very modest nonlinearity.

In response to its own shock, trimmed mean PCE inflation declines briskly for the first 5 quarters (so that it is at 0.10 percent). After that, it declines much more slowly, so that it remains significantly positive until 14 quarters later. The shock to inflation induces modest rises in the medium-frequency and low-frequency components of unemployment. There is no direct impact of trimmed mean PCE inflation on PPI-IG inflation, but the induced unemployment movements drive PPI-IG inflation down notably for 8 quarters, after which it moves modestly positive for the remainder of the IRF horizon.

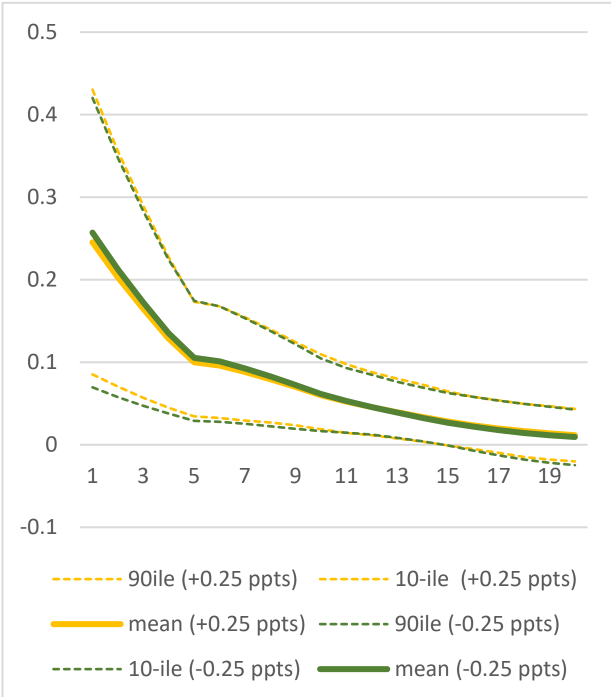

(a) Inflation

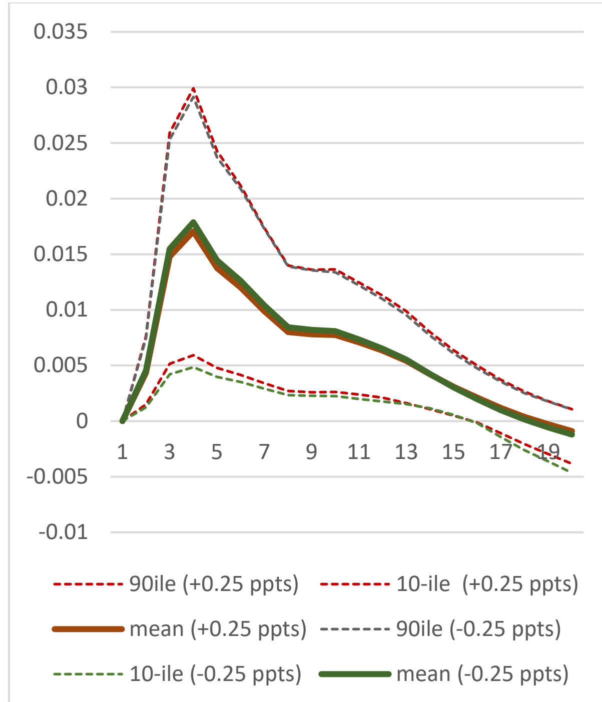

(b) Medium-frequency u

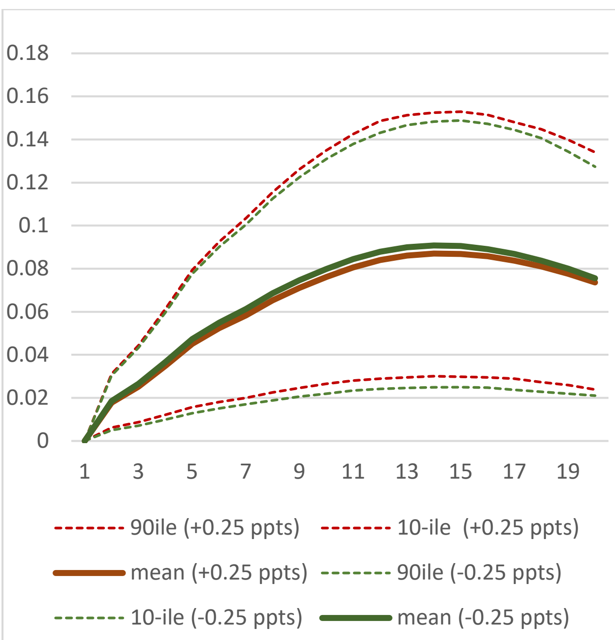

(c) Low-frequency u gap

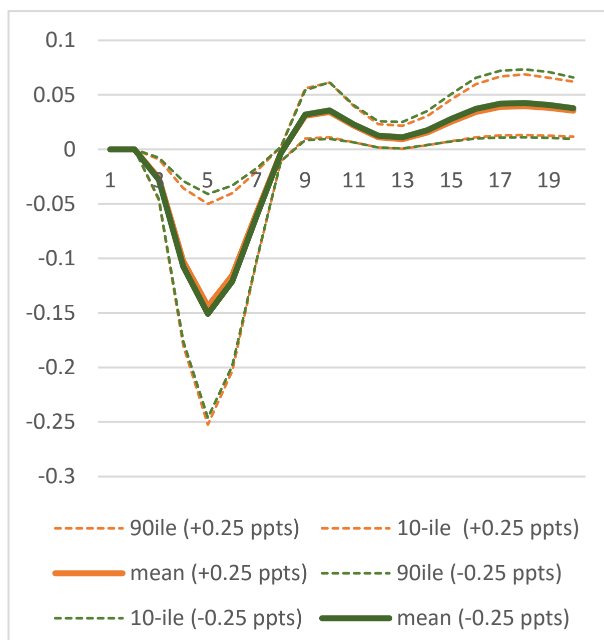

(d) PPI-IG

Figure 9: IRFs of each variable to a +/-0.25 ppt shock to trPCE inflation. The IRFs to negative shocks are reflected about the x axis. All IRFs and error bands display very modest nonlinearity.

4.4 Baseline forecast and counterfactual forecasts

As inflation is modeled in terms of deviations from PTR, in order to produce an inflation forecast, we must also forecast PTR. Our approach is to take the SPF 2022Q3 value for the five-year/five-year forward headline PCE inflation rate, namely 2.04 percent, as the 2027Q3 value for PTR, and assume that PTR linearly descends from its 2022Q3 value of 2.45 towards that value over the next five years.

Our baseline forecast calls for a recession (with unemployment gradually rising to peak at 6.5 percent at the end of 2026), and – despite the recession – for inflation to remain above the 2 percent PCE inflation target through the end of 2025.[34] The baseline forecast also calls for a fairly rapid return of PPI-IG inflation to trend; we explore sensitivity to this projection below.

Inflation is projected to be very persistent. This projection is at odds with that of the FOMC. The median projection from the FOMC's June SEP, issued about a month before the end date of our data, put core inflation at the end of 2024 at 2.3 percent and unemployment at 4.1 percent. Our model regards this outcome as exceedingly unlikely. Out of 1000 forecast simulation paths, none achieved this combination. This is not surprising: at the end of 2024, the lower confidence interval for unemployment is 4.35 percent, while the lower confidence interval for inflation is 2.6 percent. As is evident in Figure 10, unemployment is projected to climb very gradually, peaking at 6.5 percent at the end of 2026, before gradually falling.

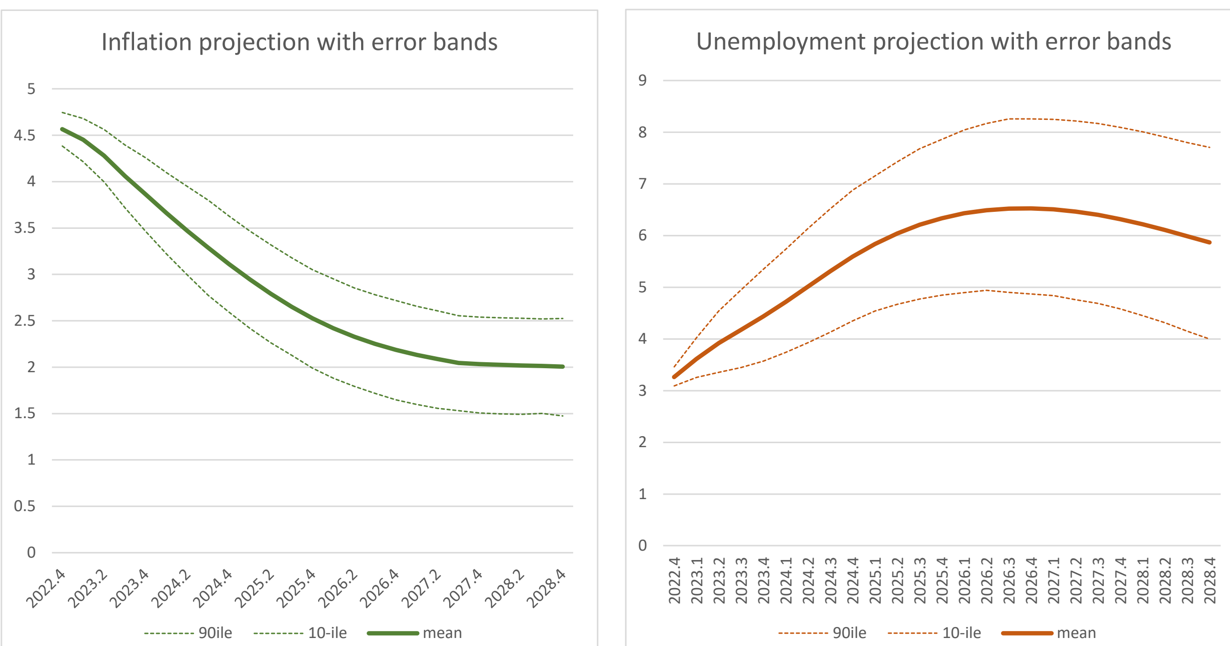

Figure 10: Baseline inflation and unemployment projections, with error bands.

[34] At this date, the lower confidence interval just includes 2%.

As our model is structurally identified, we can meaningfully construct counterfactual forecasts. We explore three such counterfactuals.

The first counterfactual relates to the path of PPI-IG inflation going forward. The baseline model calls for PPI-IG inflation to have fallen to 2.6 percent by 2023Q1, suggesting that supply-chain problems are completely gone shortly thereafter. However, many analysts see a risk of continued supply-chain problems. To assess this risk, we condition on a path for PPI-IG inflation that calls for a more gradual progress in the resolution of supply-chain issues. In particular, we replace the baseline PPI-IG forecast with one generated from a univariate model that uses lags 1, 2, and 5 of PPI-IG. Forecasts from this model feature an elevated path for 2023 and 2024, compared to the baseline forecast; see Figure 11, panel (a). The slower return to steady-state PPI-IG inflation implies that trimmed mean PCE inflation takes longer to attain its steady-state; see panel (b).

(a) (b)

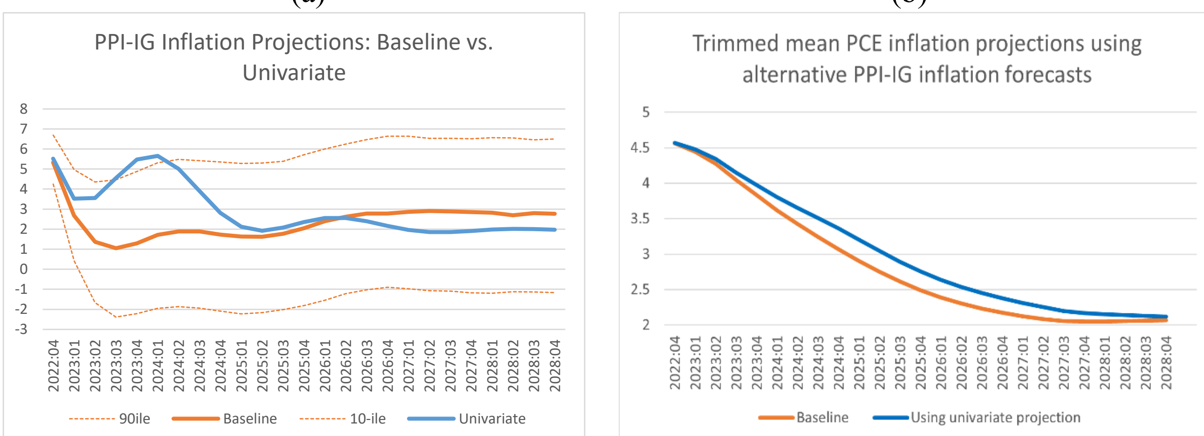

Figure 11: Counterfactual 1: Univariate projection for PPI-IG inflation.

The remaining two counterfactuals focus on alternative unemployment dynamics, and condition on the baseline path for PPI-IG. The first of these, "Soft Landing," imposes the June SEP path for unemployment – which implies no recession – and explores the implications of such an unemployment path for inflation. If the Phillips curve were quite flat, the absence of a recession would have little influence on inflation; but as we have demonstrated above, the Phillips curve is anything but flat.

The final counterfactual, "Realistic Hard Landing," explores the implications of a more realistic recession. Our model's baseline projections for both medium-frequency and low-frequency unemployment are both somewhat unrealistic. In both cases, the rise in unemployment is projected to last far too long relative to the 1985-2019 experience. As will be seen below, this generates a welfare loss that is, accordingly, unrealistically high. This counterfactual asks: What would happen if we imposed more typical and realistic recession dynamics? In particular, we select and impose the recession of the early 2000s (whose trough occurred in November 2001). We chose this recession because it is a moderate recession (in terms of overall size) over our estimation period. We condition our forecast from 2022:Q4 onward on a path for medium-frequency and low-frequency unemployment that matches these components' movements in the early 2000s recession. In the counterfactual, the medium-frequency component peaks a little earlier (4 quarters after it becomes nonnegative, rather than 6) and peaks somewhat higher (at 0.37 percent rather than 0.20 percent) but remains nonzero over a shorter period (11 quarters, rather than 22). As will be evident below, in the counterfactual, the medium-frequency component exerts less downward force on inflation, on net. The low-frequency gap in the counterfactual peaks lower (at 0.78 versus 2.5) and much earlier (at 6 quarters after becoming nonnegative, rather than 14).

The Soft Landing scenario involves a much higher path of inflation. For instance, at the end of 2025, inflation is 3.2 percent, 0.70 percentage points above the baseline. Under this scenario, inflation does not fall below 3 percent until mid-2026, and is still above 2.5 percent at the end of 2028. The Realistic Hard Landing scenario sees inflation fall a tad faster than the baseline through 2024, but slower thereafter, reflecting how persistent inflation is when it is not being influenced by unemployment. Both the baseline and the Realistic Hard Landing scenario project trimmed mean PCE inflation to be in the 3.1-3.2 percent range by the end of 2024. The gap between these inflation paths peaks in early 2027, at 0.5 percentage points. At the end of 2028, the Realistic Hard Landing inflation projection is at 2.45 percent, about 0.4 percentage points above the baseline.

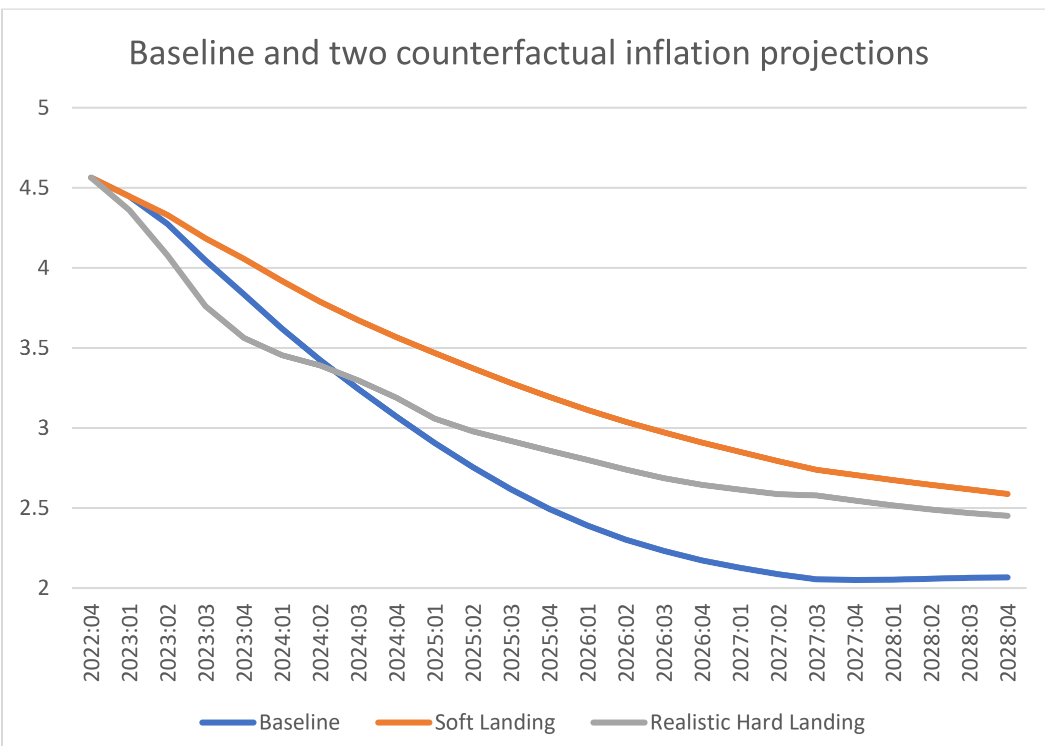

Figure 12: Trimmed mean PCE inflation projections from the Baseline and two counterfactuals.

4.5 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

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

We set $\beta = 0.995$, and guided by the June SEP and the FOMC inflation target, we set $u_t^* = 4.0$ and $\pi^* = 2.0$. We examine losses from $t_1 =$ 2022Q4 to $t_2 =$2028Q4. We compare the baseline forecast, the soft landing, and the realistic hard landing. We report the losses in Table 2, for $w = \{0.25, 0.5, 0.75\}$.

Table 2: Welfare analysis

w	Baseline	Soft Landing	Realistic Hard Landing
0.25	42.9	39.5	27.6
0.50	51.6	26.9	19.3
0.75	60.4	14.1	11.1

Even for a relatively low weight on the unemployment rate, the baseline projection results in a larger loss than does the soft landing; both feature higher losses than the realistic hard landing. As the weight on the unemployment rate increases relative to that of inflation, the baseline forecast losses grow, while those of the soft landing and the realistic hard landing shrink. Still, even for a weight of 0.75, the realistic hard landing is preferred to the soft landing. In short, this welfare analysis suggests that a recession of moderate intensity is preferable to a soft landing.

5. Conclusion

This paper implements an analogue of a nonlinear structural VAR model to jointly estimate the dynamics of inflation, as measured by the trimmed mean 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.

As does the model of Ashley and Verbrugge (2022a), the present model does a respectable job of explaining inflation dynamics over the 2007-2019 period, and in explaining the recent run-up in inflation; supply shocks were found to play a notable role in boosting trend inflation. Furthermore, during the inflation run-up, the model's forecasts were found to be more reliable than some standard benchmarks. These facts give us some confidence that the model's projections are reliable.

In this model, the estimated impulse response functions of inflation to identified structural shocks in unemployment components are highly nonlinear. For example, a positive shock to the medium-frequency component of unemployment when this component is positive (i.e., above the natural rate of unemployment) yields an IRF that indicates a notable deceleration of inflation that peaks at 6 quarters. Conversely, a negative shock to this component, when the component is negative (i.e., when it is below the natural rate), yields an IRF that increases inflation and that peaks a year later (at 10 quarters). This upward drift in inflation is driven almost entirely by the movement this shock induces in the low-frequency component of the unemployment rate.

Looking ahead, our model projects that inflation only very gradually falls back to 2 percent. Moreover, this inflation projection (baseline) is strongly influenced by the projected recession. This recession will both remove an upward force on inflation – one that stems from overheating – and impose a downward force on inflation – one that stems from recessionary pressures. In contrast, a counterfactual "soft landing" simulation indicates that without this recession, inflation would remain much more elevated for longer. Lastly, a simple welfare analysis based on a standard quadratic loss function favors a realistic hard landing over both the soft landing and the baseline scenarios. Any concern about the unanchoring of inflation expectations will only reinforce this conclusion.

Ashley and Verbrugge (2022a) summarize a large number of extant theoretical works whose predictions are consistent with their (and our) empirical results. 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

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2020. "Business-Cycle Anatomy." *American Economic Review* 110 (10): 3030–70. https://doi.org/10.1257/aer.20181174.

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.

———. 2015. "Persistence Dependence in Empirical Relations: The Velocity of Money." Working paper 15–30. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-201530.

———. 2022a. "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.

———. 2022b. "Time Series Filtering for Studying Frequency Dependence: What Works, What Doesn't." Manuscript, Federal Reserve Bank of Cleveland.

———. 2023. "The Intermittent Phillips Curve: Finding a Stable (but Persistence-Dependent) Phillips Curve Model Specification." Manuscript. Virginia Tech.

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/.

Ball, Laurence M., and Sandeep Mazumder. 2020. "The Nonpuzzling Behavior of Median Inflation." In *Changing Inflation Dynamics, Evolving Monetary Policy*, edited by Gonzalo Castex, Jordi Gali, and Diego Saravia, 27:49–70. Series on Central Banking Analysis and Economic Policies. Santiago, Chile: Banco Central de Chile. https://hdl.handle.net/20.500.12580/4880.

Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters* 109 (3): 175–78. https://doi.org/10.1016/j.econlet.2010.08.029.

Beaudry, Paul, Dana Galizia, and Franck Portier. 2020. "Putting the Cycle Back into Business Cycle Analysis." *American Economic Review* 110 (1): 1–47. https://doi.org/10.1257/aer.20190789.

Binder, Carola Conces, Wesley Janson, and Randal J. Verbrugge. forthcoming. "Out of Bounds: Do SPF Respondents Have Anchored Inflation Expectations?" *Journal of Money, Credit and Banking*. https://doi.org/10.1111/jmcb.12968.

Blalock Jr, Hubert M. 1961. *Causal Inferences in Nonexperimental Research*. Chapel Hill: University of North Carolina Press.

Blanchard, Olivier J., Alex Domash, and Lawrence H. Summers. 2022. "Bad News for the Fed from the Beveridge Space." Policy Brief PB22-7. Peterson Institute for International Economics. https://ideas.repec.org//p/iie/pbrief/pb22-7.html.

Chakhchoukh, Yacine. 2010. "A New Robust Estimation Method for ARMA Models." *IEEE Transactions on Signal Processing* 58 (7): 3512–22. https://doi.org/10.1109/TSP.2010.2046413.

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., 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.

Cochrane, John H. 2018. "A Brief Parable of Over-Differencing." Manuscript, University of Chicago, Booth School of Business. https://www.johnhcochrane.com/s/overdifferencing-889r.pdf.

Comin, Diego, and Mark Gertler. 2006. "Medium-Term Business Cycles." *American Economic Review* 96 (3): 523–51. https://doi.org/10.1257/aer.96.3.523.

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.

Domash, Alex, and Lawrence H. Summers. 2022. "A Labor Market View on the Risks of a U.S. Hard Landing." Working paper 29910. National Bureau of Economic Research. https://doi.org/10.3386/w29910.

Doppelt, Ross. 2021. "Should Macroeconomists Use Seasonally Adjusted Time Series? Structural Identification and Bayesian Estimation in Seasonal Vector Autoregressions."

Manuscript, Michigan State University. https://www2.gwu.edu/~forcpgm/Doppelt_Seasonality_2021.pdf.

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.

Friedman, Milton. 1968. "The Role of Monetary Policy." *The American Economic Review* 58 (1): 1–17. https://www.jstor.org/stable/1831652.

Garciga, Christian, and Randal J. Verbrugge. 2021. "Robust Covariance Matrix Estimation and Identification of Unusual Data Points: New Tools." *Research in Economics* 75 (2): 176–202. https://doi.org/10.1016/j.rie.2021.03.001.

Gilchrist, Simon, Raphael S. Schoenle, Jae Sim, and Egon Zakrajšek. 2017. "Inflation Dynamics during the Financial Crisis." *American Economic Review* 107 (3): 785–823. https://doi.org/10.1257/aer.20150248.

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.

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.

Hannan, Edward J. 1963. "Regression for Time Series with Errors of Measurement." *Biometrika* 50 (3/4): 293–302. https://doi.org/10.2307/2333900.

Harding, Martin, Jesper Lindé, and Mathias Trabandt. 2022. "Understanding Post-COVID Inflation Dynamics." Staff Working Paper 22-50. Bank of Canada. https://ideas.repec.org//p/bca/bocawp/22-50.html.

Hong, Sungki. 2019. "Customer Capital, Markup Cyclicality, and Amplification." In *2019 Meeting Papers*. Society for Economic Dynamics. https://ideas.repec.org/p/red/sed019/959.html.

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.

Kuhn, Florian, and Chacko George. 2019. "Business Cycle Implications of Capacity Constraints under Demand Shocks." *Review of Economic Dynamics* 32 (April): 94–121. https://doi.org/10.1016/j.red.2019.01.001.

Lusompa, Amaze, and Sai A. Sattiraju. 2022. "Cutting-Edge Methods Did Not Improve Inflation Forecasting during the COVID-19 Pandemic." *Economic Review (Federal Reserve Bank of Kansas City)* 107 (3): 21–36. https://doi.org/10.18651/ER/v107n3LusompaSattiraju.

Mankiw, N. Gregory, and Ricardo Reis. 2018. "Friedman's Presidential Address in the Evolution of Macroeconomic Thought." *Journal of Economic Perspectives* 32 (1): 81–96. https://doi.org/10.1257/jep.32.1.81.

Mertens, Elmar. 2016. "Measuring the Level and Uncertainty of Trend Inflation." *Review of Economics and Statistics* 98 (5): 950–67. https://doi.org/10.1162/REST_a_00549.

Meyer, Brent H., Guhan Venkatu, and Saeed Zaman. 2013. "Forecasting Inflation? Target the Middle." *Economic Commentary*, no. 2013-05 (April). https://doi.org/10.26509/frbc-ec-201305.

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.

Phelps, Edmund S. 1968. "Money-Wage Dynamics and Labor-Market Equilibrium." *Journal of Political Economy* 76 (4, Part 2): 678–711. https://doi.org/10.1086/259438.

Reis, Ricardo. 2022a. "Losing the Inflation Anchor." *Brookings Papers on Economic Activity* 2021 (Fall): 307–61. https://www.brookings.edu/bpea-articles/losing-the-inflation-anchor/.

———. 2022b. "The Burst of High Inflation in 2021–22: How and Why Did We Get Here?" Working paper 1060. Bank for International Settlements. https://www.bis.org/publ/work1060.htm.

Sargent, Thomas J. 2009. *Dynamic Macroeconomic Theory*. Harvard University Press.

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.

———. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (s1): 3–33. https://doi.org/10.1111/j.1538-4616.2007.00014.x.

———. 2010. "Modeling Inflation After the Crisis." Working paper 16488. National Bureau of Economic Research. https://doi.org/10.3386/w16488.

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.

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. 2022a. "Post-COVID Inflation Dynamics: Higher for Longer." Manuscript, Federal Reserve Bank of Cleveland.

———. 2022b. "Improving Inflation Forecasts Using Robust Measures." Working paper 22-23. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202223.

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 1

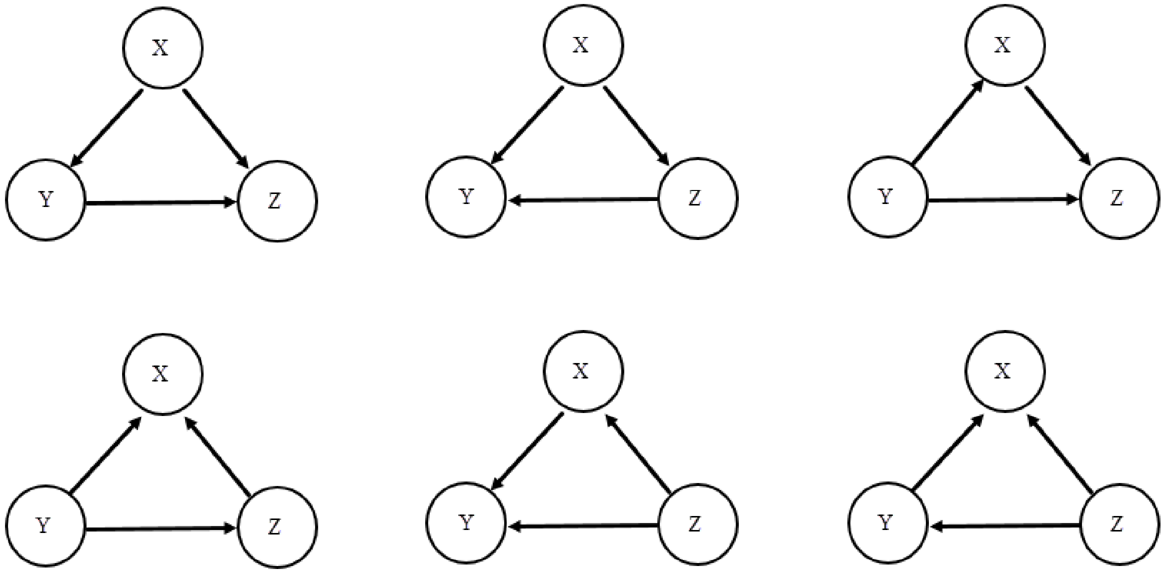

Cholesky identification models

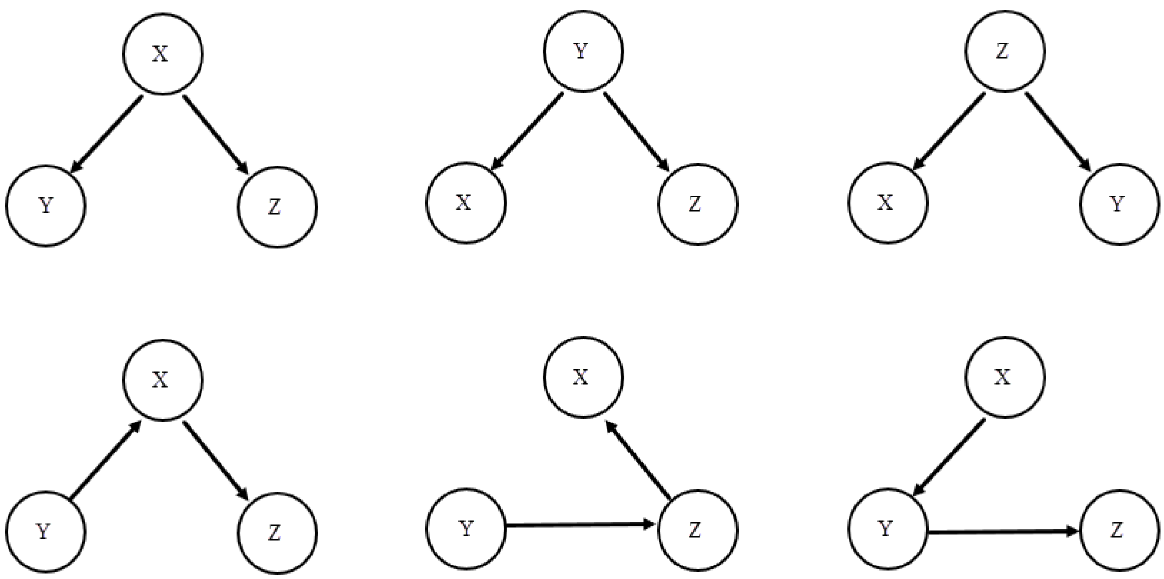

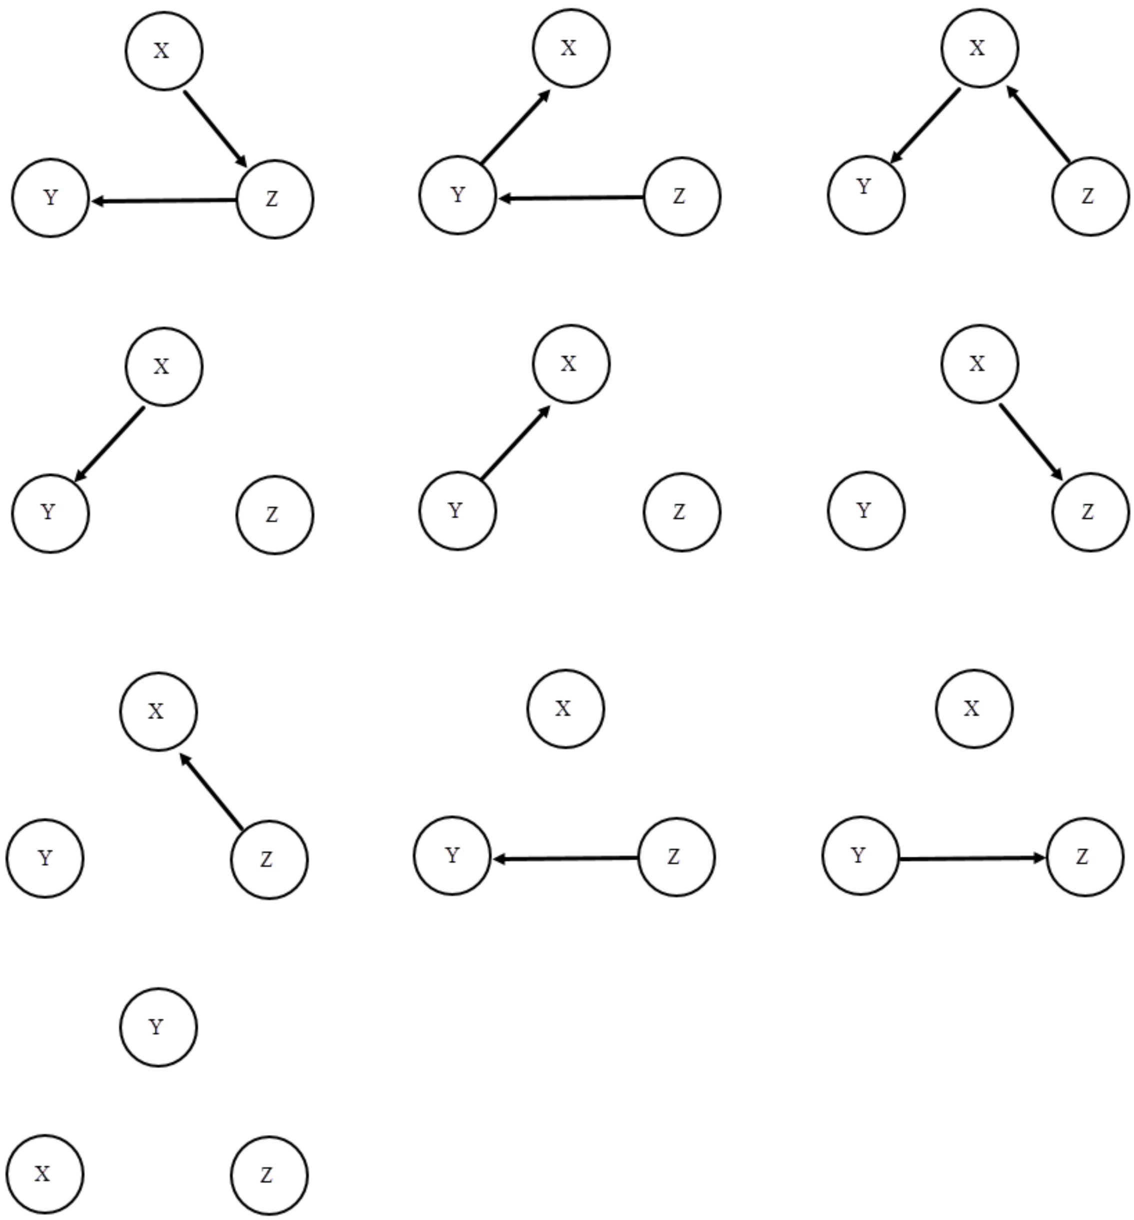

Figure A.1: All identifiable structural models involving three variables: *X*, *Y*, and *Z*.

Appendix 2: One-Sided Filtering Method of Ashley and Verbrugge

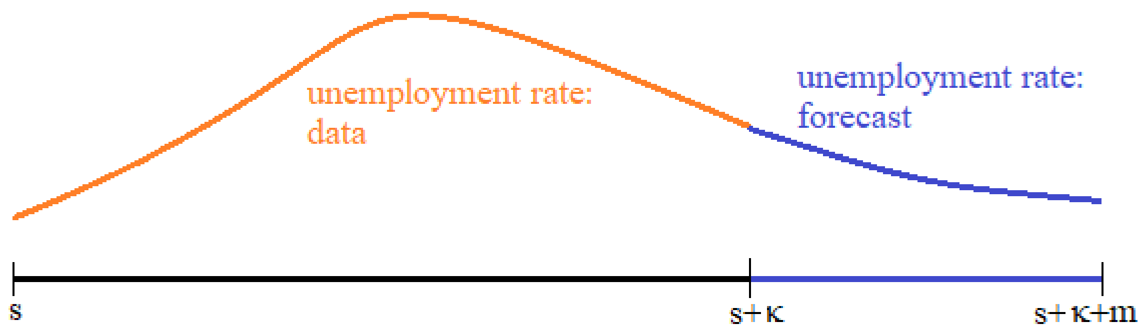

Figure A.2: One-sided filtering of unemployment rate data
(using a two-sided filter, from time s to time s+k+m)

A2.1 Description of one-sided filtering

In brief, one undertakes 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.2. We wish to compute the decomposition of the unemployment rate at time s+κ. As is well-known, obtaining the decomposition at s+κ by using a two-sided filter from time s to time s+κ 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 the edge effect problems are addressed by augmenting the data within a window with forecasts. In particular, as in Dagum (1978), 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 9, we have κ sample data points (from time period s to time period s+κ), and m months of projections, yielding a (κ+m)-quarter window (from time period s to time period s+κ+m). We then use a two-sided filter to partition that window into persistence components, and then save the partition at date s+κ; notice that this is a one-sided partition, since no *data* after date s+κ are used. To obtain the partition at date s+κ+1, we repeat this procedure, obtaining a forecast from data s+κ+1 to data s+κ+1+m, then using a two-sided filter over dates s+1 to s+κ+m+1 and saving the partition at date s+κ+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 three 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 interpretably 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 three 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 macroeconometric 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 so that these three components of the unemployment rate gap add up to the original data, and therefore, it will be easy to test the null hypothesis that the coefficients with which these three 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 (2009, 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 the interpretation of such plots as "difficult or impossible" (Granger, 1969).

Appendix 3: Estimation Results

In Table A.1, we report estimation results from the immediate precursor of Equation (3) in the text. Note that the coefficient estimates on the 4th lag of trPCE are not statistically significant at conventional significance levels. Also note that, in keeping with the findings of Ashley and Verbrugge (2022a), the Phillips curve has an asymmetric and frequency-dependent relationship to inflation: the coefficient estimates pertaining to $u_{t-1}^{+lowgap}$ and $u_{t-1}^{-medfreq}$ are not statistically significant, and the coefficient estimates on the remaining two unemployment terms are both fairly precisely estimated, and differ by a factor of 4.

(Estimation sample: 1985Q1 to 2019Q4)

Table A.1: Regression Results for a Precursor to Equation (3).

	Dependent variable $\left(\pi_t^{trPCE} - \pi_t^{PTR}\right)$	
Regressor	Coefficient estimate	Standard error
α	-0.04	0.031
$\left(\pi_{t-1}^{trPCE} - \pi_{t-1}^{PTR}\right)$	0.83	0.053
$\left(\pi_{t-4}^{trPCE} - \pi_{t-4}^{PTR}\right)$	-0.14	0.089
$\left(\pi_{t-5}^{trPCE} - \pi_{t-5}^{PTR}\right)$	0.22	0.074
π_{t-1}^{PPI}	-0.00	0.008
π_{t-2}^{PPI}	0.02	0.008
$u_{t-1}^{+lowgap}$	-0.00	0.015
$u_{t-1}^{-lowgap}$	-0.12	0.049
$u_{t-1}^{+medfreq}$	-0.51	0.135
$u_{t-1}^{-medfreq}$	-0.00	0.274
$\bar{R}^2$	0.92	

Table A.2: Regression Results for Equations (2) and (6).

Regressor	Eq. (2); Dependent variable π_t^{PPI} Coefficient estimate	Standard error	Eq. (6); Dependent variable $\left(\pi_{t-1}^{trPCE} - \pi_{t-1}^{PTR}\right)$ Coefficient estimate	Standard error
α	0.26	0.122	-0.06	0.025
$\left(\pi_{t-1}^{trPCE} - \pi_{t-1}^{PTR}\right)$			0.82	0.043
$\left(\pi_{t-5}^{trPCE} - \pi_{t-5}^{PTR}\right)$			0.11	0.037
π_t^{PPI}			0.04	0.010
π_{t-1}^{PPI}	1.46	0.083	-0.06	0.016
π_{t-2}^{PPI}	-0.74	0.151	0.04	0.010
π_{t-3}^{PPI}	0.04	0.149		
π_{t-4}^{PPI}	0.10	0.082		
$u_{t-1}^{-lowgap}$			-0.12	0.046
$u_{t-1}^{+medfreq}$			-0.51	0.116
$\Delta u_{t-1}^{medfreq}$	-6.11	1.528		
$\overline{R}^2$	0.87		0.93	

Table A.3: Regression Results for Equations (4) and (7).

Regressor	Eq. (4); Dependent variable $u_{t-1}^{medfreq}$ Coefficient estimate	Standard error	Eq. (7); Dependent variable $\left(\pi_{t-1}^{trPCE}-\pi_{t-1}^{PTR}\right)$ Coefficient estimate	Standard error
α	0.00	0.004	0.00	0.012
$\left(\pi_{t-1}^{trPCE}-\pi_{t-1}^{PTR}\right)$	0.014	0.005	0.05	0.065
$\left(\pi_{t-2}^{trPCE}-\pi_{t-2}^{PTR}\right)$			-0.08	0.063
π_{t-1}^{PPI}	0.002	0.001	-0.02	0.008
π_{t-2}^{PPI}			0.04	0.013
π_{t-3}^{PPI}			-0.03	0.013
π_{t-4}^{PPI}			0.01	0.007
u_{t-1}^{lowgap}	0.45	0.025	0.54	0.143
u_{t-2}^{lowgap}	-0.41	0.043	0.41	0.140
u_{t-3}^{lowgap}	-0.18	0.049		
u_{t-4}^{lowgap}	0.14	0.039		
$u_{t}^{medfreq}$			2.14	0.262
$u_{t-1}^{medfreq}$	0.76	0.072	-1.35	0.279
$u_{t-2}^{medfreq}$	-0.22	0.059	1.30	0.278
$u_{t-3}^{medfreq}$			-0.74	0.264
$u_{t-4}^{medfreq}$			0.41	0.156
$\overline{R}^2$	0.97		0.99	

Appendix 4: Explaining the Inflation Dynamics after the Financial Crisis

We follow Ashley and Verbrugge (2022a) and conduct a conditional recursive forecast exercise. In particular, we estimate equation (3) using data up through 2006Q4, and then fix the coefficients. We then use this equation to recursively forecast inflation from 2007Q1 through 2019Q4. No inflation data are used after 2006Q4, but the forecasts are conditional on the realizations of unemployment and the PPI for intermediate goods. For comparison, we present a parallel forecast from a more traditional specification, one in which the two unemployment terms in Equation (3) are replaced by a gap term, defined as $u_t^{traditional} := \left(u_t^{lowgap} + u_t^{medfreq} - u_t^*\right)$.

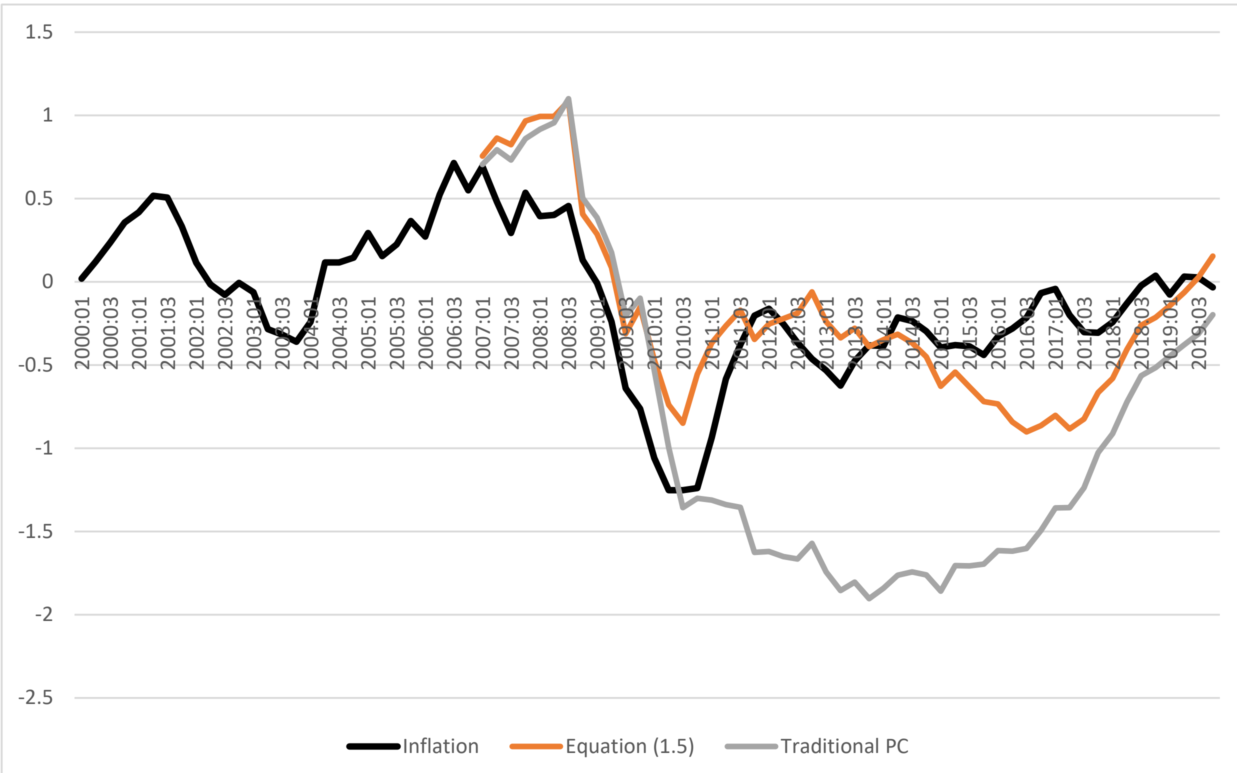

While not perfect, our quarterly equation (3) model does a far better job tracing out the dynamics of inflation over the financial crisis and the recovery than does the traditional Phillips curve (and for this type of exercise, our model would also easily outperform a Stock-Watson (2007) UCSV forecast or a random walk forecast). The RMSE of our model versus inflation is 0.31, compared to that of the traditional Phillips curve, at 0.76. Notice that with our specification, missing disinflation is all but absent (indeed, our model predicts a tad too *little* disinflation in 2010). We performed a Chow test on Equation (3) to see if the Phillips curve coefficients changed post-2006. The p-value of this test is 0.37, indicating that the Phillips curve did not weaken after 2006. Conversely, the Phillips curve coefficient in the traditional specification weakens notably (as many others have found), with a Chow test p-value of 0.013.

Appendix 5: Persistence of AR Processes
As a rule of thumb, one commonly uses the sum of the estimated lag coefficients on a univariate autoregressive process to summarize its persistence. However, this is only a rule of thumb. In fact, different autoregressive processes with the same sum of estimated lag coefficients can feature notably different persistence, as (for example) measured by half-life, depending upon initial conditions.

We illustrate this using an example. We compare the following two autoregressive processes.

$$(y_t - 2) = 0.92(y_{t-1} - 2) + e_t \qquad 'AR(1)'$$
$$(y_t - 2) = 0.81(y_{t-1} - 2) + 0.11(y_{t-5} - 2) + u_t \qquad 'AR(1,5)'$$

We begin each process at y=4, for the first 5 periods, then simulate each process for the next 42 periods (suppressing shocks). Notice that the AR(1) process reaches 2.5 in period 22, while the AR(1,5) process does not reach 2.5 until period 31. Obviously such comparisons depend upon initial conditions. If we begin each process at 4 for only the first period, then the AR(1,5) process moves toward 0 more quickly for the first 14 periods, then more slowly after that. For example, the AR(1) process descends below 0.5 in period 26, while the AR(1,5) process does not do so until period 34.

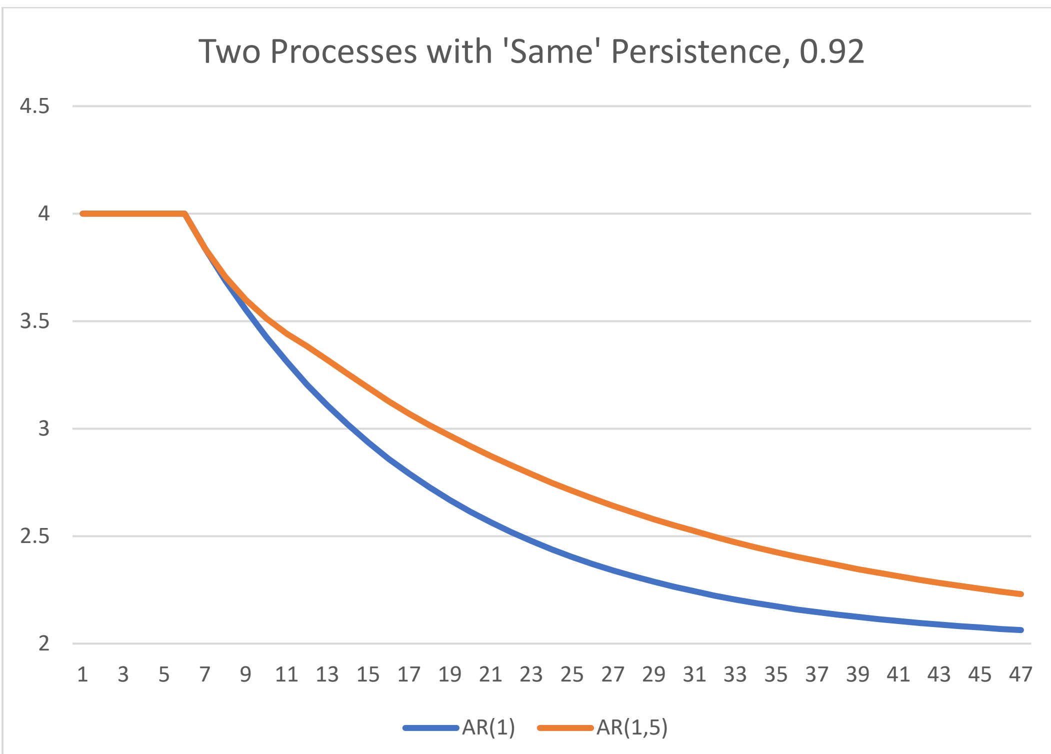

Appendix 6: Alternative Inflation Variables (and Specifications)
We estimate specifications similar to that of Equation (6) for median PCE and core PCE.

Estimation sample: 1985Q1 to 2019Q4.

Coefficients in bold are statistically significant at the 5 percent level.

Sum of AR coeffs	Dependent Variable	Constant	Lag 1	Lag 5	π_t^{PPI}	π_{t-1}^{PPI}	π_{t-2}^{PPI}	$u_{t-1}^{-lowgap}$	$u_{t-1}^{+medfreq}$
0.92	TrPCE-PTR	-0.06	**0.81**	**0.11**	0.04	-0.06	0.04	**-0.12**	**-0.51**
0.93	TrPCE-PTR	-0.02	**0.97**	-0.04					
0.85	TrPCE-PTR(2,5)	-0.10	**0.74***	**0.11**	0.04	-0.05	0.05	**-0.22**	**-0.75**
0.96	TrPCE	0.06	**0.93**	0.03	0.05	-0.06	0.04	**-0.07**	**-0.49**
0.96	TrPCE	0.08	**1.04**	-0.07					
0.96	CorePCE	0.02	**0.92**	0.04	0.10	-0.13	0.05	-0.07	-0.07
0.88	CorePCE - PTR	-0.11	**0.86**	0.02	0.10	-0.13	0.05	**-0.10**	0.00
0.93	MedPCE - PTR	-0.02	**0.87**	**0.06**	0.03	-0.05	0.03	**-0.13**	**-0.48**

* Refers to coefficient on lag 2 rather than lag 1.

Items worthy of note:

- All three variables (trimmed mean PCE, median PCE, and core PCE) are roughly equally persistent, as measured by the rule-of-thumb sum of AR coefficients. However, core PCE does not have a statistically significant coefficient at lag 5, indicating that it is a less persistent series. Median PCE is, perhaps, a tad less persistent than trimmed mean PCE, in that its estimated coefficient on lag 5 is smaller, but its coefficient estimates are otherwise very similar. Ball and Mazumder (2020) find that median PCE is "non-puzzling" in its relationship to unemployment fluctuations.
- When lag 2 rather than lag 1 is used in a specification otherwise identical to that used in the paper, the coefficients on the unemployment rate variables rise in magnitude. See Ashley and Verbrugge (2022a), who suggest that sensitivity to both unemployment terms is notably larger.
- Core PCE "has no Phillips curve" if it is not specified in terms of deviations from PTR; if it is specified in deviations from PTR, only the overheating term is statistically significant. We view this as a notable defect of core PCE, related to defects noted in Verbrugge (2022). However, in a companion paper (Verbrugge and Zaman, 2022b), we note that if one partitions core PCE into three subcomponents – housing, core services less housing, and core goods – each subcomponent displays sensitivity to either one or both unemployment rate components.[35]

[35] Verbrugge and Zaman (2022b) use median core services less housing to better capture inflation trend dynamics in core services less housing. Core services less housing is highly sensitive to outliers, which is particularly problematic for nonmarket services. See Verbrugge (2022) for a discussion of this issue.