SLACK AND CYCLICALLY SENSITIVE INFLATION

James H. Stock
Mark W. Watson

ABSTRACT

We investigate the flattening Phillips relation by making two departures from standard specifications. First, we measure slack using real activity variables that are bandpass filtered or year-over-year changes in activity (these are similar), instead of gaps. Second, we study the components of inflation instead of the standard aggregates. We find that some inflation components have strong and stable correlations with the cyclical component of real activity; these components tend to be relatively well-measured and domestically determined. Other components, typically prices that are poorly measured or internationally determined, have weak and/or unstable correlations with cyclical activity. We construct a new inflation index, Cyclically Sensitive Inflation, that weights the components by their joint cyclical covariation with real activity. The index has strong and stable correlations with cyclical activity and provides a real-time measure of cyclical movements in inflation.

James H. Stock
Department of Economics
Harvard University
Littauer Center M26
Cambridge, MA 02138
and NBER
James_Stock@harvard.edu

Mark W. Watson
Department of Economics
Princeton University
Princeton, NJ 08544-1013
and NBER
mwatson@princeton.edu

1   Introduction

Figure 1 summarizes the puzzle of the flattening Phillips curve in the United States. From 1960 to 1983, a one percentage point increase in the annual average unemployment gap, as measured by the unemployment rate minus the Congressional Budget Office's (CBO) estimate of the natural rate of unemployment, was associated with a -0.48 (standard error of 0.10) percentage point change in the year-over-year change in the rate of core PCE inflation. In 1984-1999, this slope flattened to -0.26 (SE = 0.08). From 2000 through the first quarter of 2019, by this measure, the Phillips curve was essentially flat, with a slope of -0.03 (SE = 0.03). As discussed in Section 3, this flattening is also found using other gap measures of slack, in Phillips curve forecasting regressions, and in New Keynesian Phillips curves.

Figure 1.
The flattening U.S. Phillips curve: Year-over-year change in the rate of inflation vs. the four-quarter average of the CBO unemployment gap. Numerical values are slopes. Left: Headline PCE (total); right: core PCE.
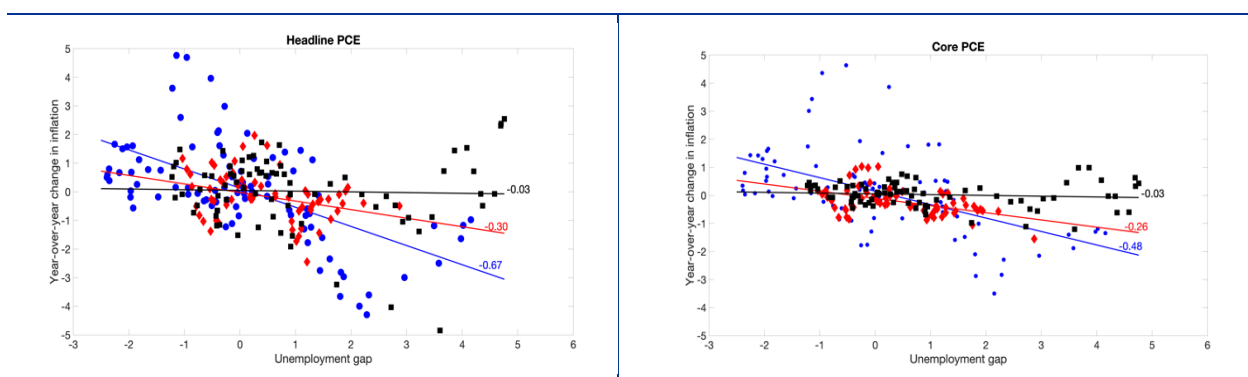
Notes: 1960-83 (blue circles); 1984-99 (red diamonds); 2000-2019q1 (black squares). Data are quarterly. The year-over-year change in inflation plotted on the vertical axis is the four-quarter change of the (backwards-looking) four-quarter moving average of the inflation rate. The horizontal axis plots the (backwards-looking) four-quarter moving average of the CBO unemployment gap.

This apparent disconnect between the rate of inflation and labor market slack raises new questions for monetary policy. Is this flattening of the Phillips curve a new and permanent feature of modern economies with credible monetary authorities? Is it the consequence of structural changes, such as the increasing importance of international markets in setting prices? Or is it in some sense a measurement artifact, so that tight economic conditions are building inflationary pressures that simply have not yet been observed?

The inflation puzzle is typically expressed, as it is in Figure 1, as a relation between inflation aggregates, typically headline or core, and an activity gap, such as the unemployment gap. In this paper, we make two departures from this standard approach.

First, we go beyond the aggregates and consider the Phillips curve properties of the components of PCE inflation. There are multiple reasons why the sensitivity of inflation to real activity might differ from one component to the next. For example, the extent to which the price of a given good or service responds to domestic cyclical pressures depends in part on the extent to which that price is set based on international or domestic market conditions. At one extreme, the prices of commodities such as oil are set in world markets, so the link between economic activity in any one country and the change in the oil price will be attenuated. In contrast, many services, such as recreational services or food served at restaurants, are largely nontradable and have prices that are set in local markets, so should be more subject to local and national cyclical pressures. More generally, price-setting dynamics and thus cyclical

variation would be expected to vary across sectors based on market structure, wage-setting practices, and so forth. In addition, the quality of price measurement varies considerably across component: for some components, measurement problems are sufficiently severe that measurement error could overwhelm cyclical movements.

Second, we take a nonstandard approach to filtering that leads us to consider alternative activity measures. Instead of focusing on gaps, which can be thought as the output of passing a real activity variable through a high-pass filters with a very low frequency cutoff, we consider real activity variables that have been filtered using filters that concentrate their gain at business cycle frequencies. Concretely, we consider two filters, a business cycle bandpass filter with pass-band of 6 to 32 quarters and the year-over-year filter (the four-quarter difference of four-quarter average), which is the filter applied to the inflation rate in Figure 1. As we show in Section 2, these two filters have similar low-frequency properties, although the year-over-year filter passes more high-frequency noise than the bandpass filter. We apply these filters to both inflation and real activity. Using these filters allows us to sidestep the problem of estimating low-frequency trends, such as aggregate trend inflation, sectoral trend inflation, potential output, or the natural rate of unemployment.

It turns out that the 17 components of the PCE price index do, in fact, exhibit a wide range of cyclicality. When we compare sectoral inflation to a bandpass-filtered index of real activity, some components, such as health care, transportation services, and financial services & insurance have very low or even negative correlations with our activity index. Other components, however, such as housing excluding gas & electric utilities and food services & accommodations, have positive and much larger cyclical correlations.

Motivated by these results, we construct a new price index designed to maximize the cyclical variation in the price index. This index, which we call Cyclically Sensitive Inflation (CSI), estimates the weights on the component prices to maximize the correlation of the CSI with our bandpass measure of aggregate cyclical variation. It turns out that this index places low weights on tradeable goods, such as energy, motor vehicles & parts, and durable household equipment. The index also places low weight on the least well-measured sectors, such as clothing & footwear and final consumption of nonprofit institutions serving households (NPISH). The sectors that receive the greatest weight – housing excluding gas & electric utilities, followed by food & beverages for off-premises consumption, and recreational services – tend to be both locally determined (nontradeable) and relatively well-measured.

Our empirical work yields four main conclusions.

First, in contrast to the declining Phillips correlation as normally measured, the correlation between the CSI inflation rate and our bandpass activity measure is high and has not declined over time. As we explain in Section 5, both of our two changes (using inflation components and using band-pass or year-over-year filtered activity) are quantitatively important to finding this stability.

Second, once the weights are estimated, the CSI inflation index can be computed in real time. It therefore provides a new indicator of the response of inflation to cyclical tightness, which can be used for real-time monitoring. This is not to say, however, that the CSI index would be an appropriate inflation target because it does not measure the overall (share-weighted) cost of living. While our approach could perhaps be useful to refining technical aspects of monetary policy, more research is needed before recommending that step.

Third, our analysis of the different cyclical behavior of the inflation components sheds light on the behavior of two other inflation measures, the median CPI produced by the Federal Reserve Bank of Cleveland and the trimmed mean PCE produced by the Federal Reserve Bank of Dallas. These two series are more highly correlated with the CSI than they are with core PCE, which (as we explain in Section 5) is not surprising because they end up placing much of their weight on components, such as housing, which are cyclically sensitive and which receive considerable weight in the CSI. Thus, the trimming involved with those two inflation series does more than simply provide statistical robustness: it ends up placing

considerable weight on relatively well-measured, cyclically sensitive components, which explains why some have found that Phillips relations hold up better for the trimmed series than for core PCE (Ball and Mazumder (2019)).

Fourth, the resilient Phillips relation we find depends both on using well-measured cyclically sensitive components and on replacing gaps with bandpass-filtered inflation or year-over-year changes in real activity. In Section 2, we use Cogley and Sbordone's (2008) generalization of the New Keynesian Phillips curve with an inflation trend to motivate applying the same filter (with the property that it eliminates a random walk stochastic trend) to both inflation and the marginal cost proxy. But more work is needed to reconcile the resiliency of band-pass slack with a theory of price setting.

Related literature. This paper contributes to a large literature that proposes multiple explanations for the apparent flattening of the Phillips curve. One set of explanations focuses on the role and formation of inflation expectations. A commonly proposed explanation is the success of monetary policy in anchoring expectations, however it is difficult to reconcile that theory with the US evidence without also having a reduction in the Phillips curve slope coefficient (e.g. Fuhrer (2012)). Coibion and Gorodnichenko (2015) suggest that firms' inflation expectations moved *countercyclically* during the recession and recovery because they are overly influenced by oil prices, which increased from 2009 to 2011 and (extending their argument) fell from 2014 through 2017. A second set of explanations focuses on special features of the financial crisis. For example, Gilchrist et. al. (2017) suggest that special features of the financial crisis affecting the pricing behavior of liquidity-constrained firms, countyeracting the expected downward pressure on inflation during the recession and early recovery. A third set of explanations focuses on structural changes that could lead to a reduction of the Phillips curve coefficient. For example, to the extent that prices of tradeable goods are set in international markets, the increasing prevalence of tradeable goods restrains aggregate inflation even when domestic labor markets are tight (e.g., Peach, Rich, and Lindner (2013), Tallman and Zaman (2017), and Forbes (2018)). In addition, technological developments have made it easier to substitute capital (robots, Web sites) for labor, further restraining wages and thus prices. A fourth set of explanations, dating to Kareken and Solow (1963) and Goldfeld and Blinder (1972) and recently summarized by McLeay and Tenreyro (2019), is that it is a consequence of the Fed credibly and successfully targets the aggregate inflation rate, which results in a flat reduced-form Phillips relation even though there is a steep and stable structural relation. In support of this view, they point to evidence of steeper and more stable Phillips relations in wages (e.g., Galí and Gambetti (2019)) and in regional inflation data (Hooper, Mishkin and Sufi (2019), and Babb and Detmeister (2017)), neither of which are explicitly targeted by the Fed. A fifth set of explanations relates to measurement problems: perhaps the apparent flattening of the Phillips Curve is, at least in part, an artefact of mismeasurement of economic slack, of the rate of price inflation, or both.

The papers most closely related to this one also focus on sectoral inflation. Peach, Rich, and Lindner (2013) propose different price-setting mechanisms for goods and services inflation (the former being more trade-sensitive) and use goods and services separately to forecast inflation. Tallman and Zaman (2017) use inflation components to forecast aggregate inflation. At least two groups have developed experimental cyclically sensitive indexes, the Federal Reserve Bank of San Francisco (Mahedy and Shapiro, 2017) and Goldman Sachs economic research (Struyven, 2017). Dées and Güntner (2017) find improvements to Euro Area inflation forecasts by disaggregating to four sectors (industry, services, construction, and agriculture). The ECB also has investigated the cyclical properties of HICP components as described in a box in the ECB Monthly Bulletin (ECB (2014)).

This paper is also related to work on core inflation, which uses inflation components to construct a less noisy measure of trend inflation. Research on core and on the use of inflation components to measure trend inflation includes the early papers of Gordon (1975) and Eckstein (1981), and more recently Cristadoro, Forni, Reichlin, Veronese (2005), Boivin, Giannoni, and Mihov (2009), and Amstad, Potter, and Rich (2017); see Stock and Watson (2016b) for additional references and discussion

of this literature. Papers on the apparent flattening of the Phillips curve in the 2000s, and especially since the financial crisis recession includes (among others) Stock and Watson (2010), Ball and Mazumder (2011, 2014), Stock (2011), Gordon (2013), Watson (2014), Kiley (2015), Blanchard (2016), and Bell and Blanchflower (2018); see McLeay and Tenreyro (2019) and Hooper, Mishkin, and Sufi (2019) for recent reviews. This literature focuses on the United States. Mazumder (2018) finds a stable Phillips curve for the Euro area using short-term professional survey expectations data, and he attributes the weakening of EA inflation to a decline in expected inflation.

The rest of the paper is organized as follows. Section 2 discusses the data transformations associated with 'gaps,' business cycles and trends. Section 3 documents the inflation puzzle and the flattening of the Phillips curve. The main finding is that the inflation puzzle remains across a range of slack measures, both real-time and ex-post. Section 4 turns to the cyclical components of inflation and digresses to discuss their construction and measurement challenges (material that proves useful for interpreting the CSI weights). Section 5 presents the CSI index, both methods and results. Section 6 concludes.

2   Gaps, Cycles, and Filters

Is the puzzling absence of a Phillips relation in the recent US data simply an artefact of mismeasuring slack? In this section, we examine Phillips correlations, Phillips slopes, and inflation forecasting relations using multiple measures of slack. We find that the results for these additional slack measures mirror those for the unemployment gap: for all these slack measures, the Phillips correlation has fallen over time, the Phillips slope has flattened, and inflation forecasts using the candidate slack measure are unstable.

2.1   Slack and gaps

Slack is an economic construct that is not measured directly. Slack is commonly estimated using an activity gap computed as the difference between an activity variable measured in real time and an unobserved level of that variable that represents full utilization of productive resources. These full-utilization levels are unobserved but can be estimated.

We refer to gap measures in which the full-utilization value is estimated using retrospective (full-sample) data as *ex-post* gap measures, in contrast to gap measures that are available in real time (real time gaps). As new data become available, the *ex-post* estimates of the full-utilization value, and thus of the gap, are revised. These revisions tend to be largest towards the end of the sample, where the newly available data have the greatest influence. As a result, *ex-post* gaps can be useful for understanding historical patterns, but are noisy and potentially misleading indicators of real-time economic conditions (Orphanides and Norden [2002]).

In this section, we consider seven ex-post gaps. The first two are from the Congressional Budget Office (CBO): the unemployment gap, which is the difference between the unemployment rate and the CBO long-term NAIRU, and the output gap, which is the log difference between GDP and CBO's estimate of potential GDP (both in logarithms).

The remaining five gap measures are constructed using time series estimates of the full-utilization value. The premise of the time series approach is that, over a period of a decade or longer, a given activity measure fluctuates around a long-term value that tracks the full-utilization value. Thus the long-term mean, or more precisely the estimated mean constructed using a low-frequency filter, of the activity measure can serve as a proxy for the full-utilization value, and deviations from this long-term mean provide estimates of the gap. Concretely, we estimate the low-frequency mean using a two-sided

biweight filter with a bandwidth of 60 quarters, and the gap is the deviation of the activity measure from this low-frequency mean.[1]

The five activity gaps estimated using the time series approach are the unemployment rate, the short-term unemployment rate (those unemployed 26 weeks or less as a fraction of the labor force), the employment-population ratio (household survey), the employment-population ratio for ages 25-54, and the capacity utilization rate.[2] To facilitate comparisons, we transform each gap to have the same mean and standard deviation as, and to be positively correlated with, the CBO unemployment gap.

Figure 2
Ex-post gaps and their first principal component), 1984-2019q1
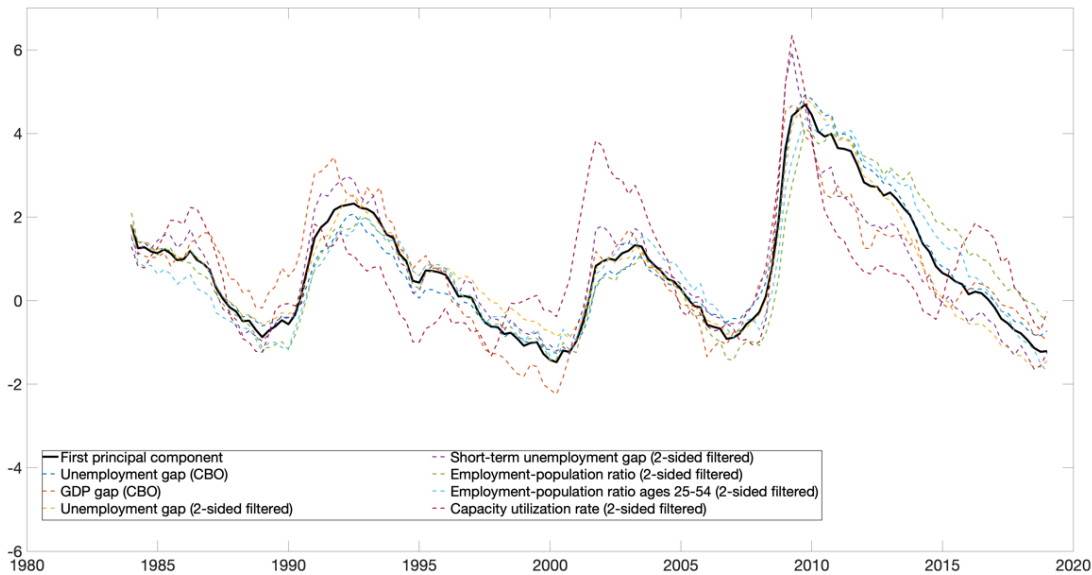
Notes: Variables are transformed to have the same mean, standard deviation, and sign as the CBO unemployment gap.

The seven standardized gaps are plotted for the period 1984-2019q1 in Figure 2. Most of the seven measures are highly correlated, with 12 of the 21 correlations exceeding 0.85 and the smallest correlation being 0.48. Figure 2 also plots the first principal component of these seven standardized gaps. As can be seen in the figure, there is considerable dispersion of the individual gap measures around their first principal component, and at any given time one of the gap measures might be giving a

---

[1] For the unemployment rate, we can compare the CBO estimate of the gap to our time series estimate. Over 1984-2019q1, the two unemployment gap measures have a correlation is 0.95. The two measures differ the most at the end of the sample (where the low-frequency filter must be mainly one-sided, and the CBO NAIRU estimate lacks data on future inflation and the unemployment rate); over 1990-2005, the correlation between the two unemployment gaps rises to 0.98.

[2] Stock (2011), Gordon (2013), Ball and Mazumder (2014), Krueger, Cramer, and Cho (2014), and Watson (2014) generally find that the short-term unemployment rate is a more stable activity variable in empirical Phillips curves than the long-term unemployment rate, using aggregate time series data for the US, however Kiley (2015) finds no advantage to using the short-term unemployment rate over the standard unemployment rate using state data. The capacity utilization rate received attention as a possible slack measure in Phillips curve research in the 1990s (e.g. Garner (1994) and Franz and Gordon (1993). The employment-population ratio is a less commonly used slack measure, but can be thought of as a broad unemployment rate because it incorporates those not in the labor force, including those who might have dropped out of the labor force because of absence of work but would want to work if a job were on offer.

different signal than the other or than the composite index. This dispersion in part reflects the difficulty of estimating full-utilization values, and thus gaps, at the end of the sample. Despite this dispersion, as of this writing, for the most recent data (first quarter of 2019), all the indexes are at historically low levels.

2.2   Trends, filtering, and the NKPC

We motivate our examination of year-over-year (yoy) changes and bandpass-filtered inflation and output by using a simplified version of Cogley and Sbordone's (2008) generalized New Keynesian Phillips Curve. The standard NKPC is derived by log-linearizing around a zero rate of inflation. Cogley and Sbordone (2008) extend the NKPC to slowly-varying trends. Their general NKPC includes additional, forward-looking terms that do not enter the standard NKPC, and has parameters that are nonlinear functions of the trends. They provide a special case, however, in which they assume that non-resetting firms' prices are fully indexed to a mixture of current trend inflation and one-period lagged inflation. In this case, their general NKPC with trends simplifies to[3],

$$\hat{\pi}_t = \gamma_f E_t \hat{\pi}_{t+1} + \gamma_b \hat{\pi}_{t-1} + \kappa \hat{x}_t + \gamma_b \left(\beta E_t \Delta \bar{\pi}_{t+1} - \Delta \bar{\pi}_t \right) \tag{1}$$

where the "^"s denote deviations from long-run trends, so that $\hat{\pi}_t = \pi_t - \bar{\pi}_t$ and $\hat{x}_t = x_t - \bar{x}_t$, where $\bar{\pi}_t$ is trend inflation and $\bar{x}_t$ is the trend value of x_t. As usual, theory suggests that x_t is marginal cost, which is unobserved so in practice a gapped activity variable is used, such as the output gap or (with the sign of κ reversed) the unemployment gap.

The generalized NKPC in (1) is a relation between gaps: the inflation gap and the unemployment gap. As discussed in the previous section, the gaps can be estimated by first estimating the full-utilization path of x and the local trend rate of inflation. However, because these trends are well-modeled as being integrated of order 1 (I(1))[4], the estimation error in the trend typically will impart an I(1) component to the error term, raising estimation and inference issues. Another approach is to model the latent trends as part of the empirical exercise (not using a two-step approach), for a recent example see Crump et. al. (2019). But doing so requires an explicit subsidiary model for the trends.

In this paper, we take a different approach to handling the trends in π and x, which is to filter the data by applying the same linear filter to both sides of (1). We choose the filter to satisfy two conditions. First, the filter should eliminate an I(1) trend in $\bar{\pi}_t$ and $\bar{x}_t$. This requirement leads us to consider filters of the form $b(L) = a(L)(1-L) = a(L)\Delta$, where $0 < |a(1)| < \infty$. Second, because we are particularly interested in the Phillips relation over the course of the business cycle, we want the filter to maximize its gain over business cycle frequencies, which we take to be periods of 6 to 32 quarters.

Applying the filter $b(L)$ to both sides of (1) yields,

$$b(L)\pi_t = \gamma_f b(L) E_t \pi_{t+1} + \gamma_b b(L)\pi_{t-1} + \kappa b(L)x_t + \zeta_t, \tag{2}$$

where ζ_t is integrated of order zero and is a function solely of $\Delta \bar{\pi}_t$ and $\Delta \bar{x}_t$, their expected future values, and associated leads and lags.

[3] Also see Ascari and Sbordone (2014) and Mavroeidis, Plagborg-Møller, and Stock (2014).
[4] See for example Stock and Watson (2007) for inflation and Gordon (1997) for the natural rate of unemployment; for more recent references to I(1) natural rates see Crump et. al. (2019).

Equation (2) expresses both the dependent variable and the marginal cost variable in terms of the same filter, b(L), which eliminates the low-frequency trends in the two series and which maximizes its gain at business cycle frequencies.

With this as motivation, the rest of this paper focuses on the relation between filtered inflation and filtered marginal cost proxies. Equation (2) provides additional guidance on the dynamics (as does the more complicated general expression in Cogley and Sbordone (2008)), however because we also wish to incorporate sectoral cyclicality, for the purposes of this paper we do not attempt to model those dynamics and instead focus on contemporaneous relationships between b(L)π_t and b(L)x_t.

Choice of filter. We consider two filters b(L). The first is a business cycle bandpass filter with pass band of 6-32 quarters.[5] The bandpass filter eliminates local trends (trends with periodicities exceeding 32 quarters) and smooths out high-frequency noise. A disadvantage of the bandpass filter is that it is two-sided so, like the gaps considered in the previous subsection, is least reliable at the end of the sample.

Figure 3
Gain of ideal bandpass, truncated bandpass, year-over-year, and Hodrick-Prescott filters
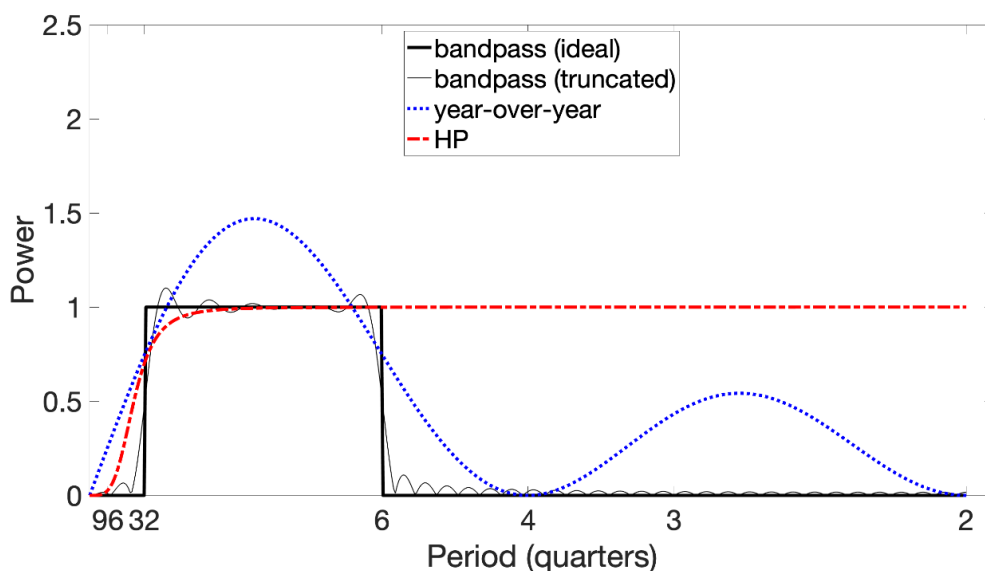
Notes: Gains are for the 6-32 ideal bandpass, the (feasible) bandpass filter truncated at 40 leads/lags, the $(1-L^4)(1+L+L^2+L^3)/4$ (year-over-year changes) filter, and the Hodrick Prescott filter with smoothing parameter 1600.

The second filter we consider is the four-quarter difference of the four-quarter average, that is, $(1-L^4)(1+L+L^2+L^3)/4$ (which can be written as a(L)Δ, where a(L) = $(1+L+L^2+L^3)^2/4$). Following convention, we refer to this filter as the "year-over-year" (yoy) filter because, when evaluated in the fourth quarter, it provides the annual average of the quarterly values of the series over the current calendar year, minus the annual average for the previous calendar year. Two disadvantages of this filter are that it passes more high-frequency variation (noise) and has a phase shift, relative to the bandpass filter. A significant advantage of this filter is that it is one-sided so does not suffer from end-point problems nor does it induce revisions, other than those induced by revisions to the underlying data themselves.

---

[5] We use the optimal 6-32 bandpass filter truncated at 40 leads and lags (see Baxter and King (1999)) augmented with forecasts and backcasts of the series computed from univariate autoregressions. For related methods see Christiano and Fitzgerald (2003).

Figure 3 plots the gains of the truncated bandbass filter, the yoy filter, and the ideal bandpass filter with pass band from 6 to 32 quarters. The truncated bandpass and yoy filters clearly have similar properties for low-frequency components of signals (periodicities exceeding four quarters), however the yoy filter passes more high-frequency noise. As an additional comparison, the figure also includes the gain of the Hodrick-Prescott filter (smoothing parameter 1600). The HP filter eliminates somewhat less low-frequency trend than the bandpass filter, but somewhat more than the yoy filter. Unlike the other filters, the HP filter does not attenuate higher frequency noise.

Figure 4 illustrates the differences of gaps and business cycle filters using the CBO unemployment gap and the filtered unemployment rate using the bandpass, yoy, and HP filters. The most notable feature of this figure is that the three cyclically filtered series are more like each other than they are the unemployment gap, a feature confirmed by computing correlations among the series. Some of the time, the cyclically filtered series and the unemployment gap track each other, such as during the late 1970s, but most of the time they give different readings on slack. This divergence between the unemployment gap on the one hand and the three cyclically filtered variables on the other is most striking in the late 1960s, the mid 1980s, and since 2012.

Figure 4
Unemployment gap and three filtered versions of the unemployment rate: business cycle-bandpass (BP), year-over-year changes (YOY), and Hodrick-Prescott (HP)
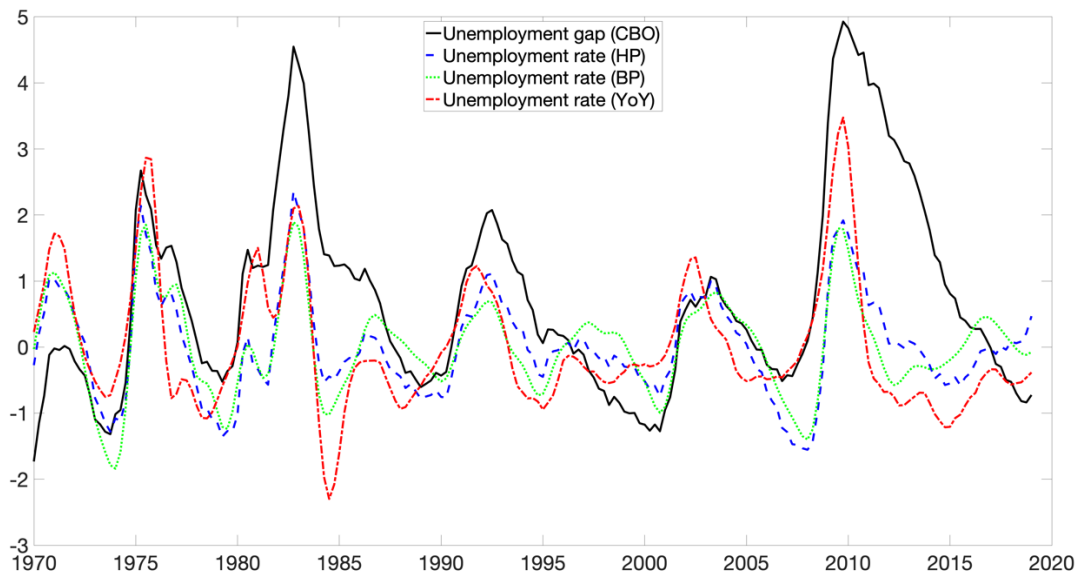
Notes: The filtered unemployment rate series are standardized. The Hodrick Prescott filter has smoothing parameter 1600.

We begin with bandpass filtered values of activity variables because the bandpass filter eliminates high-frequency noise more successfully than the yoy filter. We repeat some of our analysis using yoy-filtered activity variables as a sensitivity check.

2.3   A composite index of real cyclical activity

We consider seven activity variables: Gross Domestic Output (GDO, the geometric average of GDP and Gross Domestic Income, see Nalewaik, 2010), the capacity utilization rate, establishment

employment, the overall employment-to-population ratio (household survey), the employment-population ratio for ages 25-54, the unemployment rate, and the short-term unemployment rate. The bandpass filtered cyclical measures computed from these seven variables are plotted in Figure 5. To facilitate subsequent visual comparisons with inflation, the cyclical activity variables are standardized to have mean zero and unit standard deviation, and the unemployment rate activity variables are multiplied by -1 to co-vary positively with output. (Note that this "output gap" sign convention is the opposite of the "unemployment gap" sign convention in the Section 2.1.)

Figure 5
Bandpass cyclical activity measures for the US
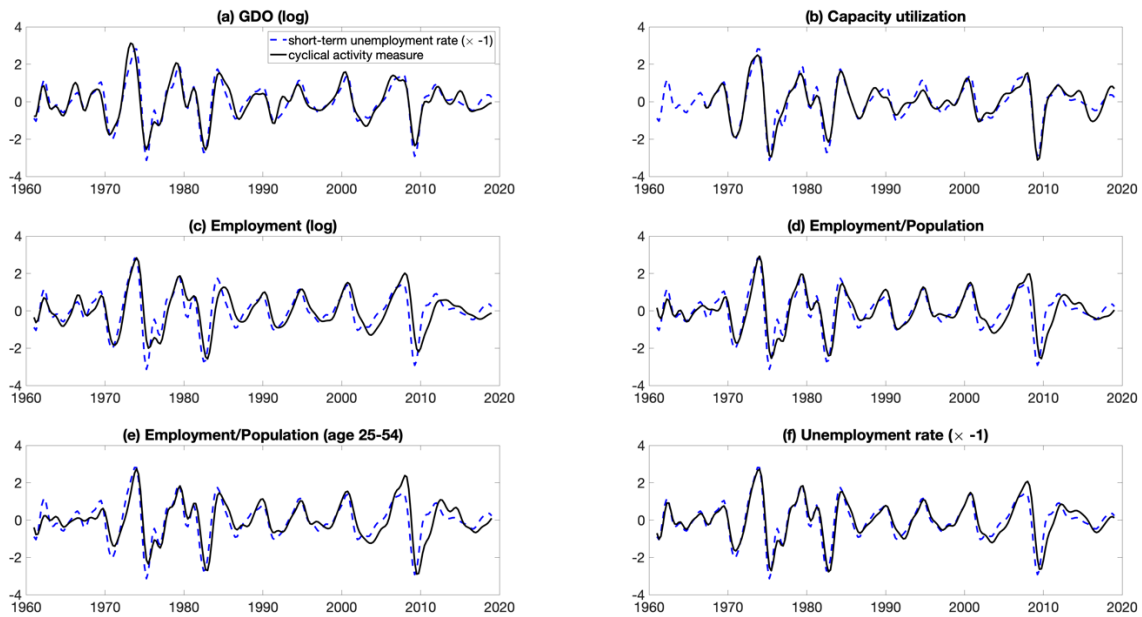
Notes: The cyclical activity measures are bandpass filtered of the various activity variables, using a pass band of 6-32 quarters. The bandpass filtered series are standardized to have mean zero and unit variance. The unemployment rates are multiplied by -1 so that they co-vary positively with the output gap.

The seven cyclical activity measures are evidently very similar, however they exhibit different timing, as can be seen by comparing each measure to the cyclical component of the short-term unemployment rate (shown for reference in each panel). The cyclical components of the short-term unemployment rate, GDO, and capacity utilization are approximately contemporaneous, however establishment employment, the employment-population ratios, and the unemployment rate each lag the short-term unemployment rate by 2 quarters.

We use these seven series to construct a composite index of cyclical activity, computed as the first principal of the second lag of the short-term unemployment rate, GDO, and capacity utilization, and the unlagged value of the other four cyclical measures, all bandpass filtered and standardized. This composite activity index (CAI) is plotted in Figure 6, along with the seven constituent cyclical activity measures (in three cases, lagged two quarters). The composite index explains 92% of the variation (trace R^2) of its constituent cyclical activity measures.

Figure 6

Seven bandpass measures of cyclical activity and their first principal component (the Cyclical Activity Index)
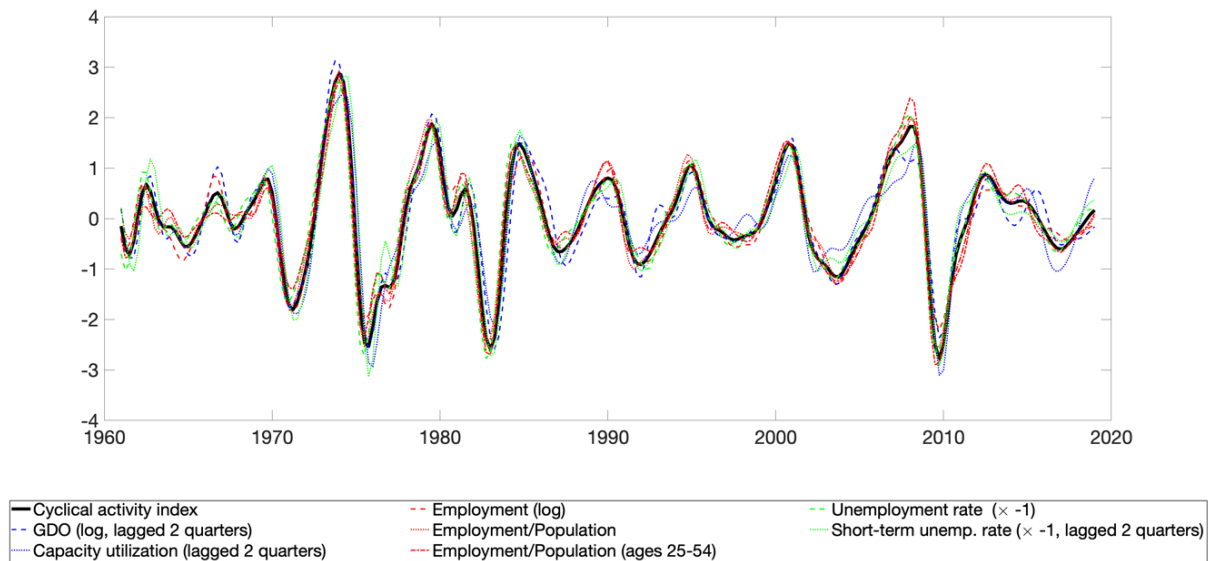
Notes: The seven cyclical activity measures are the bandpass-filtered activity variables listed in the legend. The cyclical activity index is the first principal component of the cyclical activity measures. The capacity utilization rate, GDO, and the short-term unemployment rate are lagged two quarters, and the unemployment rate and short-term unemployment rate are multiplied by -1 so they co-vary positively with the output gap.

Figure 7

Activity index computed using bandpass- and yoy-filtered activity variables
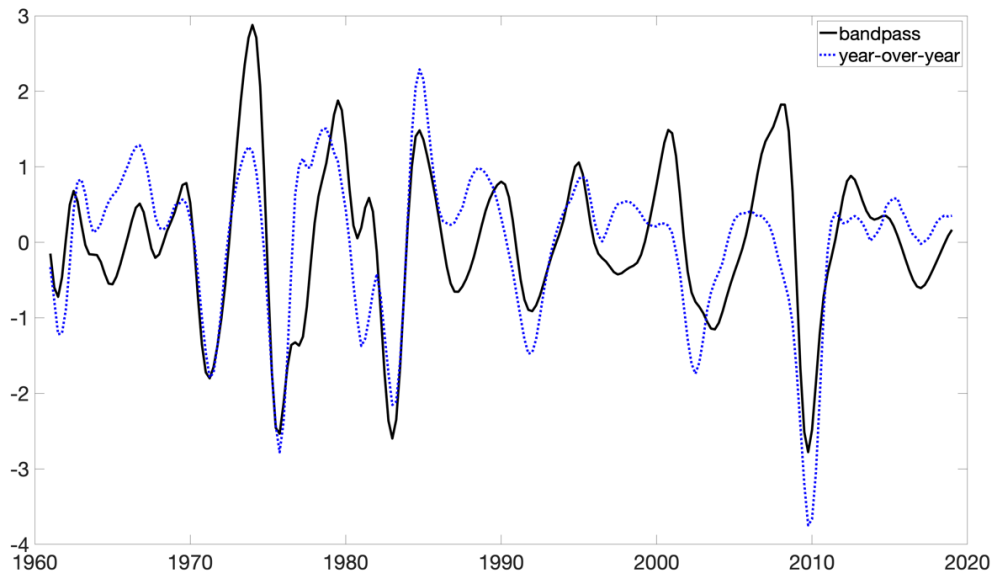
Notes: The bandpass-filtered version (solid line) is the CAI.

An alternative is to compute the activity index using yoy-filtered activity variables instead of bandpass-filtered activity variables. Figure 7 shows the CAI and its yoy counterpart, the first principal

component of the yoy-filtered activity variables. The two indexes have generally similar behavior over business cycles and have a correlation of 0.68.

3 The Changing Phillips Correlation

One possibility is that the Phillips correlation is stable if one uses the right gap measure. This turns out, however, not to be the case: the decline is robustly found across gap measures. Moreover, the decreasing usefulness of gaps for explaining movements in inflation extends to simple Phillips curve forecasts.

3.1 Evolution of the Phillips correlation

Table 1provides the correlation between the four-quarter change in four-quarter core PCE inflation, i.e., the yoy-filtered inflation rate ($\Delta_4 \pi_t^4 = \pi_t^4 - \pi_{t-4}^4$, where $\pi_t^4 = (\pi_t + \pi_{t-1} + \pi_{t-2} + \pi_{t-3})/4$) and the contemporaneous standardized four-quarter moving average of seven ex-post gap measures ($x_t^4 = (x_t + x_{t-1} + x_{t-2} + x_{t-3})/4$). The first three columns show these correlations for three periods, 1960-1983, 1984-1999, and 2000-2019q1; the final three columns show the slopes of the Phillips relation, $\Delta_4 \pi_t^4 = \beta_0 + \beta_1 x_t^4 + u_t^4$, estimated by OLS. In addition, results are shown for the (not gapped) unemployment rate and the short-term unemployment rate. For these two measures over these sub-samples, the variation in the estimated full-utilization values is fairly small relative to the variation in the activity measure, so that most of the variation in the activity measure is variation in the gap.

Consistent with the scatterplot in Figure 1, by each of these slack measures the US accelerationist Phillips correlation has been getting weaker and its slope has been getting flatter. This conclusion is robust to using shorter or longer temporal aggregation and to deviating π_t^4 from a t–4 – dated univariate forecast.

3.2 Inflation forecasts using slack over the financial crisis recession and recovery

Although our primary focus is on the contemporaneous Phillips relation, we note that there is a comparable deterioration in the contribution of activity gaps to Phillips curve forecasts. To illustrate, we conduct a small comparison of Phillips curve vs. univariate forecasts of four-quarter ahead inflation. We use real-time (one-sided) gaps, where the full-utilization values are computed as a one-sided exponentially-weighted moving average, with a weight with half-life of 15 years.[6] These real-time gaps were computed for the unemployment rate, the short-term unemployment rate, the capacity utilization rate, and the two employment-population ratios. We also use two non-gapped variables, the unemployment rate and the short-term unemployment rate. We consider the Phillips curve forecasting model, $\Delta_4 \pi_t^4 = \beta_0 + \beta_1 x_{t-4} + \beta_2 \pi_{t-4}^4 + e_t^4$, where x_t is the candidate real-time gap.

Table 2 summarizes results for two illustrative forecasting exercises. The first column reports the sup-Wald test of the hypothesis that the coefficients in the forecasting regression are stable over the 1984q1-2019q1 period. The second column summarizes the results of a pseudo out-of-sample forecasting exercise, in which the forecasting model was estimated using pre-recession data (from 1984q1-2007q1) and used to forecast inflation during the recession and recovery (from 2008q1-2019q1; 2008q1 is the first fully out-of-sample date for the four-quarter ahead forecast). The table reports the

---

[6] The exponential moving average filter yields real time gaps with correlations with the two-sided biweight smoothing gaps between 0.88 and 0.96 for the two unemployment rates and the capacity utilization rate; these correlations are lower (.72 and .79) for the employment-population ratio gaps, which have large nonstationary components. Similar results obtain using one-sided 15-year equal-weighted moving averages to construct the gaps, although those gaps generally have a lower correlation with the two-sided biweight gaps.

root mean square forecasting error (RMSFE) in the out-of-sample period from the model including slack, relative to the RMSFE of the model with the slack measure excluded, so a relative RMSFE less than one indicates that the slack measure improved inflation forecasts over the period 2008q1-2019q1.

Table 1
Phillips correlations and slopes for core PCE inflation and various gaps

Slack Measure	Correlation			Slope (SE)		
	1961-1983	1984-1999	2000-2019q1	1961-1983	1984-1999	2000-2019q1
Ex-post gaps						
Unemployment gap (CBO)	-0.53	-0.47	-0.11	-0.47 (0.10)	-0.25 (0.08)	-0.03 (0.03)
GDP gap (CBO)	-0.49	-0.34	-0.23	-0.37 (0.06)	-0.14 (0.07)	-0.07 (0.05)
Unemployment gap (two-sided filtered)	-0.57	-0.52	-0.08	-0.54 (0.12)	-0.26 (0.08)	-0.02 (0.03)
Short-term unemployment gap (two-sided filtered)	-0.53	-0.52	-0.23	-0.41 (0.09)	-0.24 (0.08)	-0.07 (0.04)
Employment-population ratio (two-sided filtered)	-0.57	-0.47	-0.02	-0.61 (0.14)	-0.20 (0.07)	-0.00 (0.03)
Employment-population ratio ages 25-54 (two-sided filtered)	-0.45	-0.47	-0.05	-0.53 (0.14)	-0.22 (0.08)	-0.01 (0.03)
Capacity utilization rate (two-sided filtered)	-0.64	-0.45	-0.17	-0.53 (0.10)	-0.23 (0.08)	-0.05 (0.03)
Real-time activity						
Unemployment rate	-0.49	-0.41	-0.10	-0.42 (0.09)	-0.19 (0.07)	-0.02 (0.03)
Short-term unemployment rate	-0.44	-0.36	-0.23	-0.35 (0.08)	-0.15 (0.06)	-0.09 (0.05)

Notes: Regressions: $\Delta_4 \pi_t^4 = \beta_0 + \beta_1 x_t^4 + u_t^4$, where x_t is the slack variable and the superscript "4" denotes four-quarter moving average. Because of data availability, the first period for the capacity utilization rate is 1967q4-1983. All slack measures have been standardized to have the same mean and standard deviation as the CBO unemployment gap, and multiplied by -1 when needed to be positively correlated with the unemployment gap; thus the slope coefficients have the same units so their magnitudes are comparable. Standard errors (in parentheses in the final three columns) are Newey-West with 8 lags.

Table 2
Forecasting annual changes in core PCE inflation using gap variables.

Predictor slack variable	Sup-Wald test statistic (*p-value*)	Pseudo out-of-sample RMSFE ratio 2008q1-2019q1
Unemployment rate	12.28 (<0.01)	1.32
Short-term unemployment rate	9.04 (<0.01)	0.98
Unemployment rate (real-time gap)	12.33 (<0.01)	1.27
Short-term unemployment rate (real-time gap)	7.82 (<0.01)	0.97
Employment-population ratio (real-time gap)	25.52 (<0.01)	1.32
Employment-population ratio ages 25-54 (real-time gap)	20.68 (<0.01)	1.00
Capacity utilization rate (real-time gap)	21.17 (<0.01)	1.04

Notes: Four-quarter ahead forecasting regression: $\Delta_4 \pi_t^4 = \beta_0 + \beta_1 x_{t-4} + \beta_2 \pi_{t-4}^4 + e_t^4$, where x_t is the real-time gap variable. The first column reports the Sup-Wald statistic (15% trimming) testing the null hypothesis that all three coefficients in the forecasting regression are stable, when estimated over the period 1984q1-2019q1. The second column is the ratio of the pseudo out-of-sample root mean squared forecast errors of the direct forecasting regression in the table header, to the RMSFE for the restricted version without the slack variable, where all regressions are estimated over 1983q1-2007q1 and the RMSFEs are computed over 2008q1-2019q1.

The results in Table 2 are striking. Only two of the gap variables, the short-term unemployment rate and its real-time gap version, improve out-of-sample performance, compared with using no gap variable at all, and the performance improvement is economically negligible. For four of the eight slack variables, the forecasting performance is *worse* using the variable than not. For all the gap measures, the hypothesis of coefficient stability is rejected at the 1% significance level. This finding of instability, illustrated here for simple forecasting models, is in line with the literature on inflation forecasting, which stresses the prevalence of time-variation in forecasting relations using activity variables (e.g. Groen, Paap, and Ravazzolo (2013)).

The conjecture that motivated the investigation of alternative gap measures in this section was that perhaps the apparent flattening of the Phillips curve was an artefact of focusing on the unemployment gap, so that the apparent flattening would be resolved if we found the "right" gap measure. The evidence, however, does not support this conjecture.

4   Cyclically Properties of the PCE Inflation Components

We now turn to the cyclical properties of inflation components. Recently there has been increasing attention to the possibility of mismeasuring prices and, as a result, inflation and productivity growth. Our interest here is in whether measurement problems could be obscuring the cyclical movements in inflation. After listing the components, we therefore provide a brief discussion of some price measurement challenges at the components level. We then examine the cyclical properties of the inflation components.

4.1   Components of PCE inflation

Personal consumption expenditures are expenditures on final purchases of goods and services consumed by persons. PCE inflation measures the rate of price inflation of those goods, weighted by

their share in final consumption. The US Bureau of Economic Analysis (BEA) uses 16 third-tier components of consumption (four components of durable goods, four of nondurable goods, seven of household services expenditures, and final consumption expenditures by nonprofit institutions serving households (NPISH) that pay for services then provide them to households without charge. We further decompose housing services into two components, housing excluding energy and housing energy services, for a total of 17 components.

These 17 components are listed in the first column of Table 3.[7] The second column gives the component expenditure shares in total PCE (average over 2000s). The components with the largest shares (16% each) are housing ex utilities and health care; the percentage share weights of all other components are in the single digits. The quarterly rates of inflation for the 17 components are plotted in Figure 8.

Table 3

Third-tier components of PCE inflation and their shares

Component	Share (2000s)	Subtotals
A. Well-measured		
Housing ex utilities	0.16	
Recreation services	0.04	
Food and beverages for off-premises consumption	0.08	**0.34**
Food services and accommodations	0.06	
Housing - energy utilities component	0.02	**0.05**
Gasoline and other energy goods	0.03	
B. Some information content		
Other services	0.09	
Other nondurable goods	0.08	
Transportation services	0.03	
Motor vehicles and parts	0.04	**0.29**
Other durable goods	0.02	
Furnishings and durable household equipment	0.03	
Health care	0.16	**0.16**
C. Poorly measured		
Recreational goods and vehicles	0.03	
Clothing and footwear	0.03	
Financial services and insurance	0.08	**0.17**
NPISH	0.03	

Sources: US BEA and FRED for the data, and author's judgement for the A, B, and C categories.

---

[7] Data on nominal PCE and price indexes for the United States are from the US NIPA Tables 2.3.4U and 2.3.5U.

Figure 8

The 17 PCE inflation components and core PCE inflation (dashed) (four-quarter inflation)
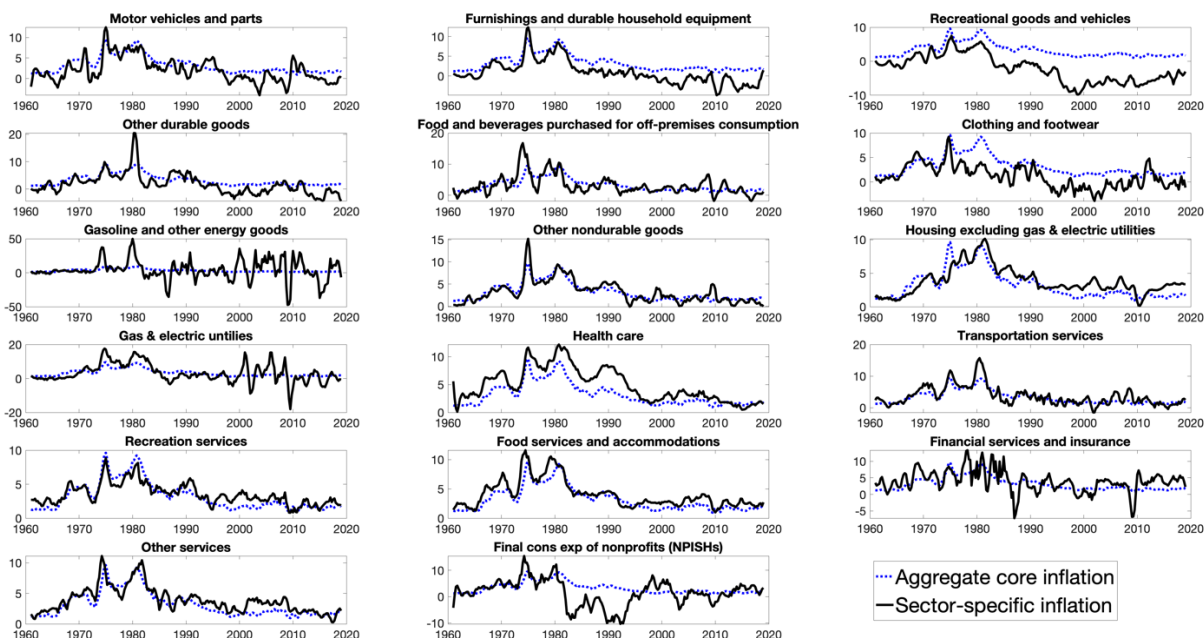
4.2   A digression on measurement

The PCE price concept is the price paid for final consumption of a good or service. This price could be paid by the final consumer directly or on behalf of the consumer by a company or institution (e.g. an insurance company or a nonprofit serving individuals). Price measurement confronts a number of well-known challenges, of which we focus on two: the estimation of prices when market prices are not available, and the challenge of rolling in prices on new or improved goods or services. Additional challenges include substitution bias, incomplete historical revisions for some sectors when methods change[8], updating sampling procedures (e.g. incorporating new outlets), and (perhaps) introducing prices for non-priced goods provided for free to consumers by businesses (e.g. Google searches). We keep the discussion here brief and refer the reader to Moulton (2018) and US BEA (2017) for details and references.

When available, posted market prices are used. Posted market prices are typically available for goods, but not for many services. For example, in the US, health care prices typically are negotiated prices not posted market prices (negotiated between health care provider organizations and insurance companies), in which case BEA and BLS attempt to estimate prices for specific packages of health services. In other cases, such as some legal services sold as final consumption (wills, real estate closings,

---

[8] For example, the 2013 PCE revision introduced a number of changes to the imputation of prices for financial services, including the use of a less volatile interest rates to measure foregone interest in accounts at commercial banks that provided unpriced conveniences. The BEA revised the series using the new methodology back to 1985, but before 1985 the series is unrevised. The large break in volatility evident in this component of inflation in 1985 in Figure 8 is due to this partial revision (Hood (2013)).

personal legal defense fees, etc.), prices are in part estimated based on a cost approach using billable hourly rates and estimated numbers of hours for a service. An extreme example of this is the price index for unpriced services provided to the public by nonprofits, such as religious institutions, where the price for religious services (say) is estimated based on the cost of providing those services. Another example of imputation of prices where none exist (either negotiated or market) is many financial services. For example, the price of convenience services provided by a bank for checking accounts is imputed using the interest income forgone by holding a balance in a checking account instead of a non-checkable asset with a higher rate of interest; implementing this concept requires estimating the interest rate on the foregone (counterfactual) investment.

Another challenge for price measurement concerns new goods and quality improvements. The problem with quality improvements arises when a good reaches the end of its life cycle and is replaced by a similar, but improved, good. The new goods problem is an extreme version that arises when a new type of good becomes available, such as the introduction of smart phones. BEA has a number of strategies for addressing the new/improved goods problem. In some cases, the value of the quality improvements can be estimated using hedonic methods. In other cases, the quality improvements are estimated based on changes in production costs, however this method conflates efficiencies in production with quality improvements. In yet other cases, new goods are chained in without an attempt to quality-adjust. The challenges posed by new/improved goods problem is often raised in the context of IT goods, but it includes low-tech as well as high-tech goods. For example, clothing typically has a short life cycle stemming from changing fashions, and prices for a given good (say, a specific shirt) decline over time as it gets marked down; at some point, the good disappears as new goods (new shirts) are introduced.[9]

Based on these and related considerations, and on discussions with experts on price measurement in the US government and elsewhere, we categorized the 17 PCE components into three working categories, A, B, and C, and grouped the components in Table 3 accordingly.

Category A consists of components that have relatively well measured prices. Prices in these categories tend to be market prices, and the new goods problem (while present) is relatively less pronounced than in other categories. For example, rents (the basis for the housing inflation index) are measured using a rotating survey of a panel of housing rental units with low turnover, and are adjusted for improvements in the units.[10]

Category B contains components which in our judgement have some information content, but for which either the new goods or non-market price problems are potentially substantial. For example, health care prices are measured using (typically negotiated) prices actually paid for specific representative health care goods, but are not adjusted for quality based on outcomes so arguably understate quality improvements.

---

[9] A third challenge, which has been the subject of considerable attention recently, is the free goods problem. This issue is frequently raised in the context of IT services provided for free, such as services provided by free apps or Google searches. The free goods problem also is not new: television provides free goods too. Whether to address the free goods problem raises basic questions about whether NIPA accounting measures welfare (if so, they should be included) or market-based economic activity (if so, they should not). Here we stick to the standard concept of market-based activity so do not venture into the realm of free goods.

[10] For owner-occupied housing, the housing services component treats the price the owner pays as the rents the owner would pay to herself, where those rents are imputed based on rents for comparable homes in the local market. This imputation introduces imputation error, especially for more expensive homes for which the rental market is thin. Nevertheless, the imputation is based on actual rental prices so the imputation simply places greater weight on some rental units than others.

Category C components are ones that in our judgement have very significant measurement issues, including new/improved goods problems (IT equipment, which falls under recreational goods and vehicles, and clothing) and/or rely mainly on imputed nonmarket prices (like the price index for services provided for free by nonprofit institutions serving households [NPISH]).

4.3   Cyclical properties of inflation components

We begin our examination of the variation in cyclical properties of sectoral inflation by comparing movements in the yoy inflation to the composite index of cyclical activity (the CAI). These series are plotted in Figure 9 for the 17 components. The correlations between the inflation components and the cyclical index are given in Table 4 for bandpass filtered inflation (first column) and for yoy inflation (second column). Recall that the CAI sign convention is the "output gap" sign convention, so positive comovement (procylical inflation) corresponds to a downward-sloping Phillips relation.

Figure 9

Seventeen inflation components and the CAI
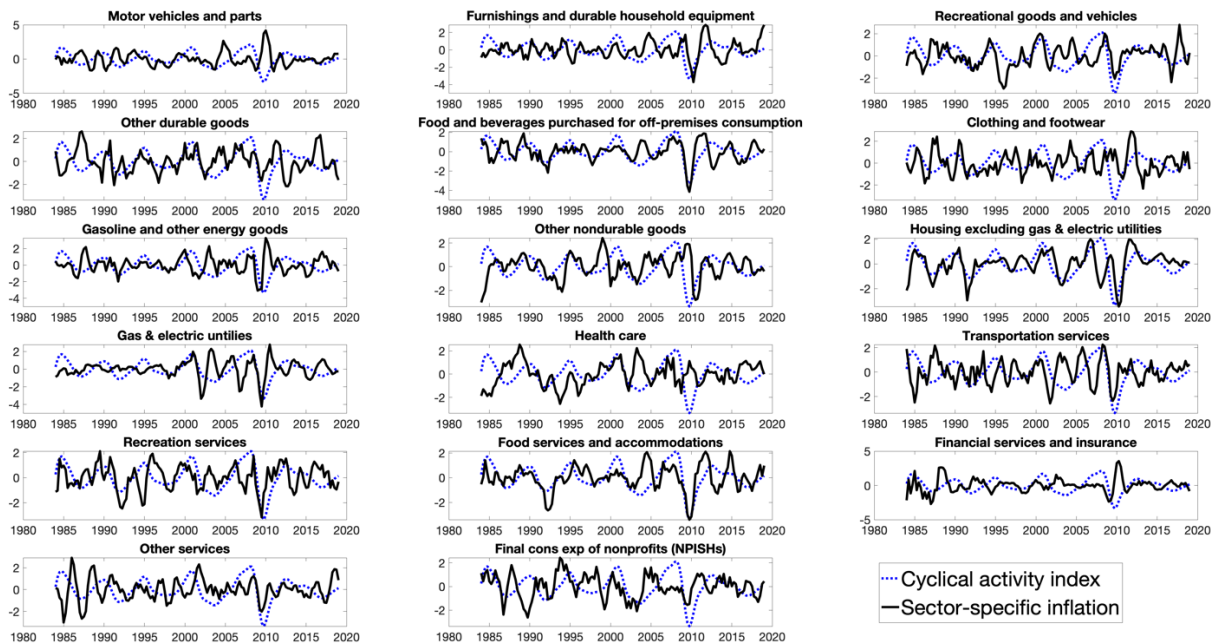
Notes: YOY inflation ($\Delta_4 \pi_{it}^4$), standardized to have mean zero and unit standard deviation; the CAI is in red.

The variation across components in the cyclicality of inflation is striking. For some components, cyclical inflation (bandpass filtered) is very highly correlated with the cyclical activity index; these sectors include food services and accommodations (correlation = 0.67) and housing excluding energy (also 0.62), Other components, however, either exhibit little cyclical variability or vary countercyclically. These noncyclical components include other nondurable goods, transportation services, health care, gasoline and other energy goods, clothing and footwear, and financial services and insurance. Motor vehicles and parts is countercyclical, a feature that is largely driven by the price jump in used cars in October 2009 following the "cash for clunkers" program. For most components, correlations for yoy inflation are lower than for bandpass inflation, however both filters show the same pattern across components.

These correlations and plots are consistent both with cyclical sensitivity varying across sectors and with the quality of measurement varying across sectors. The sectors with the highest cyclical correlations tend to be dominated by services that have prices determined in local (non-tradable) markets and which are relatively well-measured: housing services, recreational services, and food services and accommodations. Food and beverages off-premises is relatively well-measured and although raw commodity prices are set internationally, there is a substantial local (non-tradeable) component of food prices.

Table 4
Correlations between inflation components and the composite slack measure, 1984-2019q1

Component	Correlation between cyclical slack measure and:	
	Bandpass inflation	4Q change in 4Q inflation
Motor vehicles and parts	-0.27	-0.33
Furnishings and durable household equipment	0.31	0.08
Recreational goods and vehicles	0.35	0.25
Other durable goods	0.17	0.08
Food and beverages purchased for off-premises consumption	0.58	0.42
Clothing and footwear	-0.02	-0.09
Gasoline and other energy goods	-0.06	-0.10
Other nondurable goods	-0.11	-0.03
Housing excluding gas & electric utilities	0.62	0.40
Gas & electric utilities	0.24	0.13
Health care	-0.04	-0.11
Transportation services	0.03	0.01
Recreation services	0.39	0.27
Food services and accommodations	0.67	0.45
Financial services and insurance	-0.12	-0.17
Other services	0.13	0.16
Final cons. exp. of nonprofits (NPISHs)	0.31	0.15

The sectors with the smallest cyclical correlations tend to be internationally traded goods (e.g. gasoline); sectors with prices that are heavily influenced by internationally traded goods (e.g. transportation services, which relies on refined petroleum products); sectors with managed, negotiated, or regulated prices (health care and some transportation services); and/or sectors with prices that are poorly measured (financial services and insurance and clothing & footwear). The components of other services prices are in many cases estimated using costs (e.g. attorneys' hourly costs), and the low correlation of that sector might be a consequence of the cost-based imputation

missing cyclical variation in markups. One surprising finding is the procyclicality of NPISH inflation, which presumably stems from procyclicality of the costs used to impute NPISH prices (recall that by definition prices do not exist for NPISH consumption because it is provided to end consumers without charge).

5 Cyclically Sensitive Inflation

5.1 Benchmark CSI index

The CSI index is a weighted average of the 17 component rates of inflation in Table 4. The weights are chosen to maximize the correlation between the composite index of cyclical activity and the yoy change in the index, subject to the constraint that the weights are positive and add to one. These weights are estimated by nonlinear least squares estimation of the regression,

$$CAI_t = \beta_0 + \beta_1 \sum_{i=1}^{17} w_i \Delta_4 \pi_{it}^4 + u_t, \text{ subject to } 0 \leq w_i \leq 1 \text{ and } \sum_{i=1}^{17} w_i = 1, \tag{3}$$

where CAI_t is the cyclical index. The quarterly CSI rate of inflation is $\pi_t^{CSI} = \sum_{i=1}^{17} \hat{w}_i \pi_{it}$.[11]

The CSI weights on sectoral inflation, estimated over the 1984-2019q1 sample, are reported in the first column of Table 5. The estimates place nonzero weight on only a few sectors: half of the weight is placed on housing ex energy, with the remaining weight spread over food and beverages consumed off-premises, other services, recreation services, and recreational goods and vehicles. Only two goods categories, food & beverages off-premises and recreational goods and vehicles, receive weight. Notably, 89% of the weight in the CSI index is on the relatively well-measured Category A series, even though those components comprise only 39% of consumption.

Figure 10 plots the yoy CSI inflation index and the CAI over the period 1960-2019q1. The vertical line in the figure marks the start of the 1984-2019q1 sample over which the weights were estimated. For the 1984-2019q1 sample, the CSI index in Figure 10 is the in-sample predicted value from estimation of regression (3). In the 1960-1983 period, the CSI was computed by applying the 1984-2018 weights in Table 5 to the historical values of the PCE components.

Because the CSI weights were estimated over the 1984-2019q1 sample, the 1960-1983 sample provides an opportunity to assess the cyclical stability of CSI inflation. Inspection of Figure 10 suggests that the cyclical properties of CSI inflation are stable in the pre-estimation sample. The correlation between the two series in Figure 10 is 0.66 in the pre-estimation sample (1960-1983), *greater* than its value of 0.62 in the estimation sample (1984-2019q1). A regression test of the stability of this relationship in and out of sample does not reject stability at the 10% significance level. Similar stability results are found for the other bandpass filtered activity variables. There are a number of reasons why these correlations might be smaller in the 1960-1983 out-of-sample period than in the estimation period, including the supply-side sources of the inflation shocks of the 1970s, differences in monetary policy regimes, and changes in the relative quality of measurement of the components. In this light, this

---

[11] Our baseline CSI estimate in (3) rather incongruously uses yoy inflation and the CAI, which is derived from bandpass activity data. This mixture of filters is done with the practical purpose in mind of constructing an inflation index that can be computed in real time (the yoy filter is one-sided but the bandpass filter is not). Other alternatives are to use yoy inflation and yoy activity or bandpass inflation and bandpass activity, then use those weights to construct CSI as the weighted average of current-quarter component inflation. We treat these other choices as sensitivity checks below.

stability of the cyclical behavior of the CSI index in the pre-estimation period suggests that its cyclical behavior could be stable in the post-estimation period as well.

Figure 11 plots CSI inflation (in levels) along with headline PCE and PCE-xFE inflation. Three features are noteworthy.

First, CSI has more pronounced cyclical movements than the other measures, especially towards the end of the last three expansions: CSI rises as the cyclical peak approaches and subsequently falls during the recession and the early recovery. This pattern is evident in every recession since 1960, except for the brief first recession of the twin recessions of the 1980s.

Second, the most persistent divergences of CSI inflation from headline and core appear since 1990. During the 1990s core and headline declined while CSI inflation remained constant, then CSI inflation rose substantially towards the end of the 1990s expansion. CSI inflation also shows stronger cyclical behavior than core before and during the financial crisis recession.

Third, the CSI index seems to be less sensitive to energy prices than headline or even core inflation. For example, CSI inflation rose less than headline and core during the oil price jump of 1973, nor did it fall by as much as headline or core during the oil price collapse of 1986. Neither CSI nor core PCE inflation fell during the oil price decline of 2014-15.

Table 5

Correlations between inflation components and the cyclical activity index, and CSI weights, 1984-2019q1

	CSI Weights				Share-weighted Indexes				MUCSVO
	1	2	3	4	5	6	7	8	9
(a) Weights									
Motor vehicles and parts	0.000	0.00	0.00	0.01	0.05	0.06	0.13	0.00	0.04
Furnishings and durable household equipment	0.000	0.00	0.00	0.00	0.03	0.03	0.08	0.00	0.03
Recreational goods and vehicles	0.087	0ᵃ	0.11	0.12	0.03	0.04	0.08	0.00	0.04
Other durable goods	0.000	0.00	0.00	0.00	0.02	0.02	0.05	0.00	0.02
Food and beverages purchased for off-premises consumption	0.137	0.17	0.03	0.10	0.09	0.00	0.24	0.00	0.10
Clothing and footwear	0.000	0ᵃ	0.00	0.00	0.04	0.05	0.11	0.00	0.02
Gasoline and other energy goods	0.000	0.00	0.00	0.00	0.03	0.00	0.08	0.00	0.00
Other nondurable goods	0.000	0.00	0.00	0.01	0.08	0.09	0.22	0.00	0.08
Housing excluding gas & electric utilities	0.496	0.56	0.51	0.20	0.16	0.19	0.00	0.25	0.17
Gas & electric utilities	0.016	0.02	0.03	0.01	0.02	0.00	0.00	0.04	0.01
Health care	0.000	0.00	0.00	0.00	0.15	0.17	0.00	0.23	0.15
Transportation services	0.000	0.00	0.00	0.00	0.03	0.04	0.00	0.05	0.03
Recreation services	0.100	0.08	0.02	0.17	0.04	0.04	0.00	0.06	0.05
Food services and accommodations	0.033	0.08	0.28	0.19	0.06	0.07	0.00	0.10	0.10
Financial services and insurance	0.000	0ᵃ	0.01	0.00	0.07	0.08	0.00	0.11	0.05
Other services	0.103	0.08	0.00	0.16	0.08	0.10	0.00	0.13	0.07
Final cons. exp. of nonprofits (NPISHs)	0.028	0ᵃ	0.00	0.03	0.02	0.03	0.00	0.04	0.03
(b) Correlation with the composite activity index									
1961-1983	0.66	0.66	0.64	0.77	0.72	0.59	0.74	0.62	0.72
1984-1999	0.56	0.53	0.59	0.45	0.15	0.04	0.14	0.06	0.19
2000-2019q1	0.66	0.64	0.61	0.64	0.02	0.10	-0.09	0.24	0.38

Notes: CSI weights are estimated by nonlinear least squares estimation of the regression in Equation (3). The CSI weights are as follows: 1. Benchmark (see text), 2. Excludes four poorly measured sectors (see Table 3), 3. Uses bandpass inflation in place of yoy inflation, 4. Imposes maximum weight of 0.20. The share weights columns report weights and correlations for share-weighted indexes, where the indexes constructed for: 5. All sectors, 6. Excluding food and energy, 7. Goods, 8. Services (for columns 5-8, the indexes are constructed using 1984-2019q1 average share weights). The final column (F) reports weights for the 17-component index computed using the Stock-Watson (2016) MUCSVO model. ᵃRestricted to equal zero.
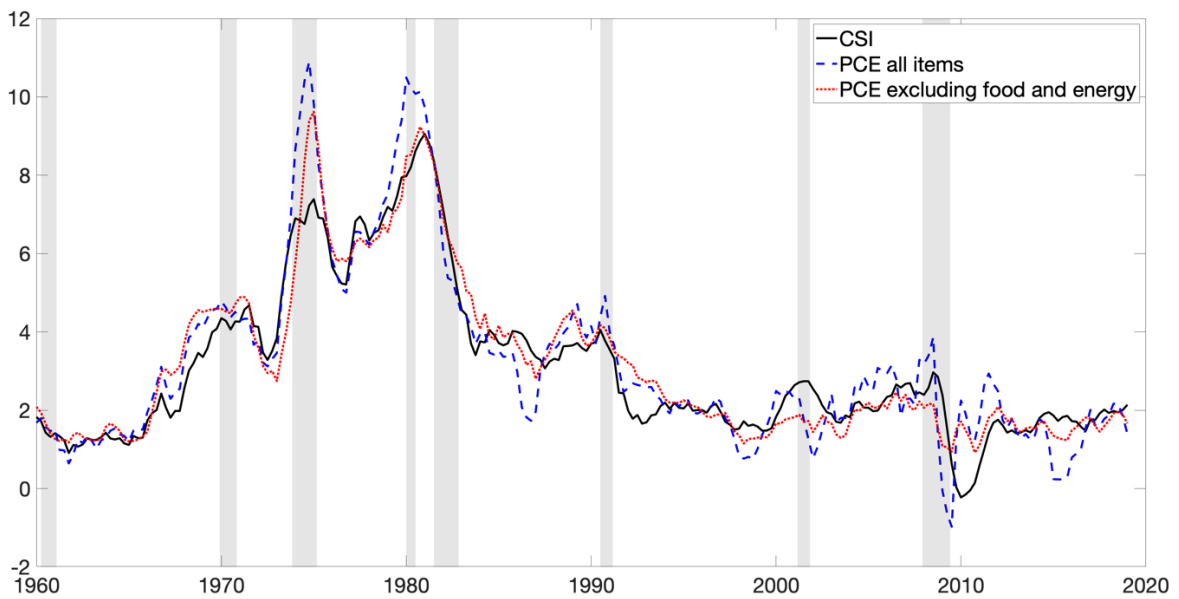
Figure 10

YOY CSI inflation ($\Delta_4 \pi_t^{CSI,4}$) and the CAI
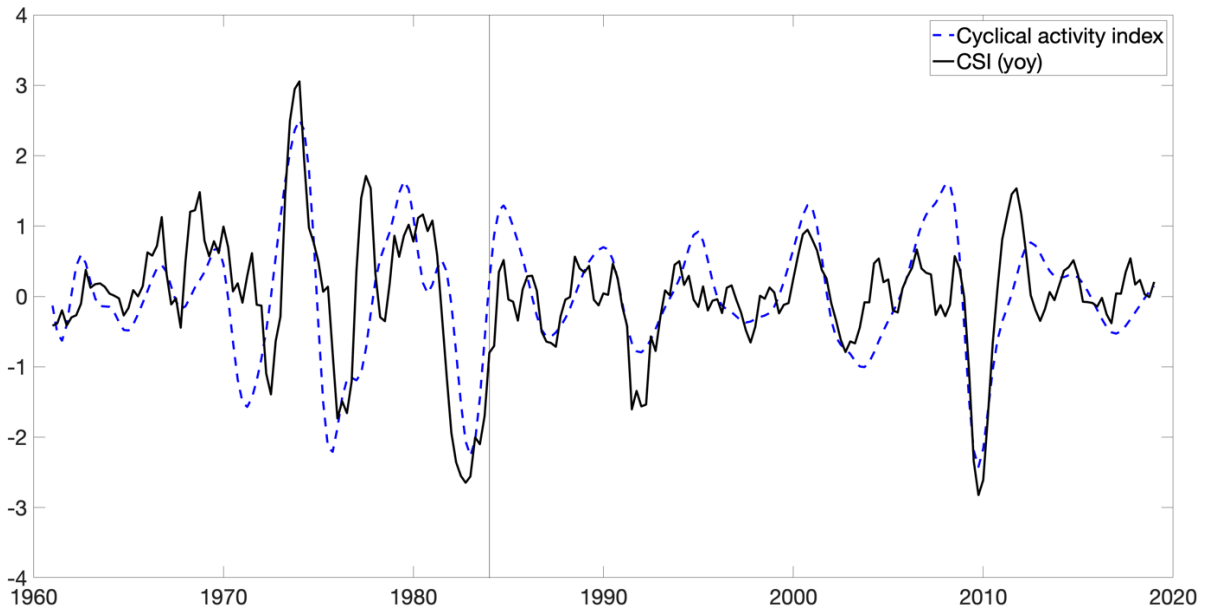
Notes: CSI inflation is computed using weights estimated over 1984-2019q1 (after the vertical line). For this figure, the CAI is rescaled to have the same standard deviation as $\Delta_4 \pi_t^{CSI,4}$.

Figure 11

US four-quarter inflation rates for the US: PCE, PCExFE, and CSI


Note: Shading denote NBER recessions.

One of the motivations for this investigation was the flattening of the Phillips curve and the declining Phillips correlation. Figure 12 plots the relationship between the CAI and yoy CSI inflation (left) and bandpass CSI (right). Unlike the Phillips scatterplots in Figure 6, the Phillips scatterplot using CSI and the composite cyclical slack measure is stable.

Figure 12

Evolution of the US CSI inflation Phillips correlation: filtered CSI inflation vs. the CAI (inverted). 1961-83 (blue circles), 1984-99 (red diamonds), 2000-2019q1(black squares).
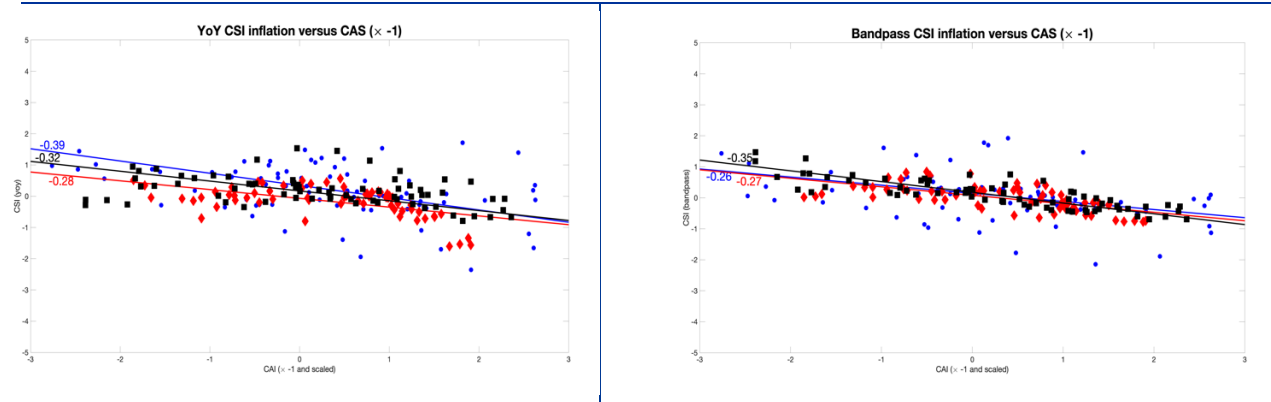
Notes: The CAI is normalized to have the same standard deviation as the CBO unemployment gap.

5.2   Sensitivity analysis

We summarize five sets of sensitivity checks.

First, the benchmark CSI weights were computed using the full 1984-2019q1 sample, and it is of interest to whether and how the weights and the resulting CSI inflation have been stable over time. We therefore recomputed the CSI measure by estimating Equation (1) using rolling regressions with a 60-quarter window. The resulting rolling CSI inflation is compared with the full-sample CSI index in Figure 13, which plots both series as yoy inflation. Although there is substantial time variation in the rolling weights themselves, the components that receive weights do not differ substantially over time (most weight is put on housing, food & accommodation services, food & beverages off-premises, and recreation services), and the predicted changes in CSI inflation differ little between the full- and rolling-sample estimates. This finding that the weights are unstable, but the CSI inflation estimate is not, seems to be a consequence of the relatively high correlation among those components that receive weight.

Second, the benchmark CSI weights used all 17 sectors, but as noted above, three of the four sectors which we had judgmentally considered to be the least well measured in fact received essentially zero weight. We therefore re-estimated the CSI weights, dropping the four "Category C" series in Table 3. The results are reported in the second column of Table 5. Doing so increases the weight on housing, but otherwise makes little difference to the weights, and makes only a negligible change to the correlations with the composite activity index. These results reinforce our conclusion that one of the things that the CSI weighting is doing is zeroing out the most poorly measured series because the measurement issues mask any cyclical properties they might have.
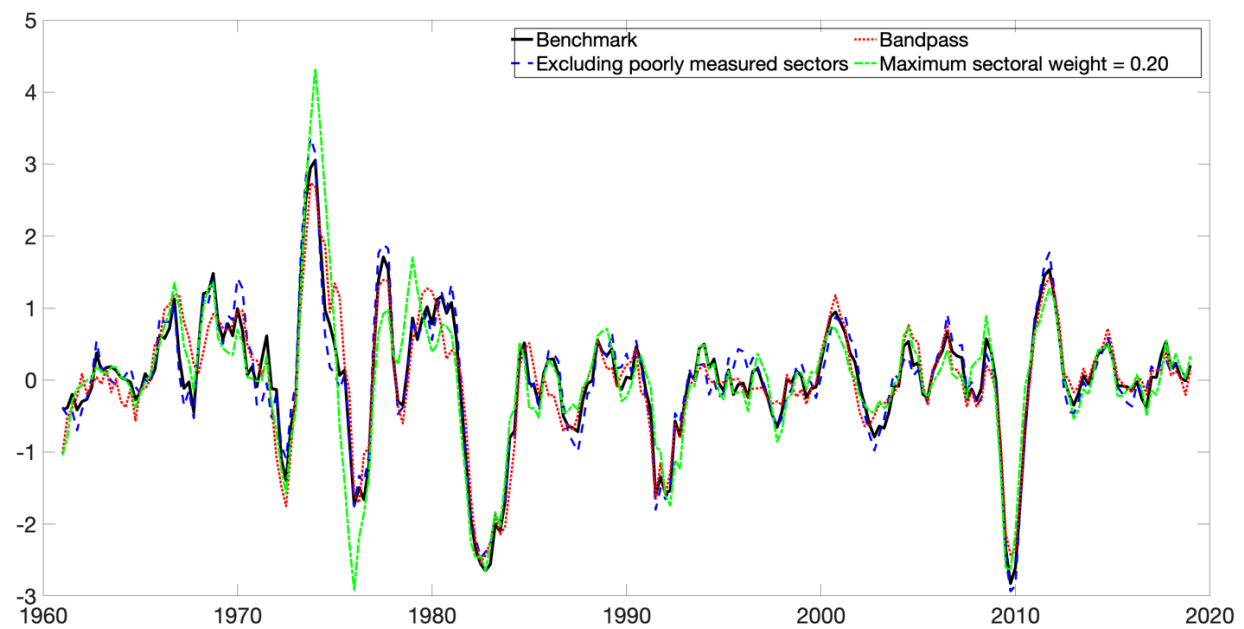
Figure 13

Sensitivity check: 4-quarter changes in 4-quarter CSI inflation using benchmark (solid) and rolling (dashed) weights
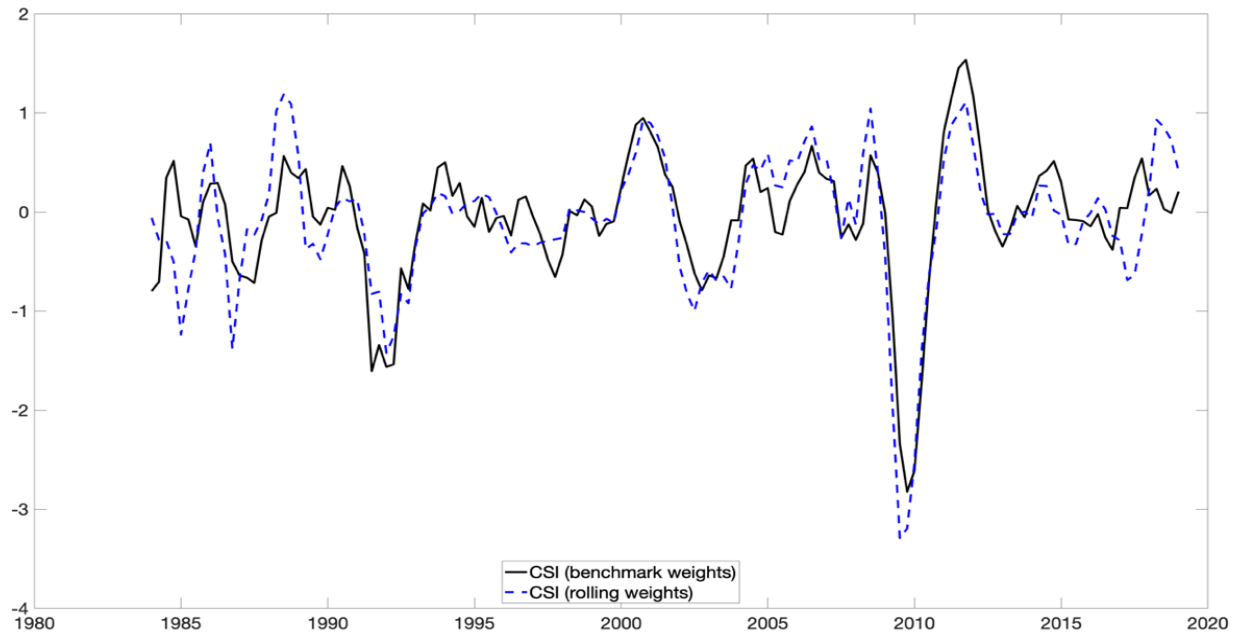
Notes: The rolling weights were estimated using rolling regressions with a 60-quarter window, with the first sample period 1962q1-1984q1.

Figure 14

Sensitivity check: Alternative CSI weights

Notes: The alternative CSI weights are in columns 2-4 of

Table 5.

Third, the benchmark CSI weights are estimated using yoy sectoral inflation and the CAI, which is computed from bandpass activity variables. An alternative approach is to estimate the weights using bandpass sectoral inflation instead, then using those weights to compute CSI from the component quarterly inflation series. The results are given in the third column of Table 5. Evidently, using bandpass inflation instead of yoy inflation to estimate the weights makes little difference (we return to this point in Section 5.4).

Fourth, one feature of the CSI index is that just over half the weight is placed on housing ex energy, a component that has a share weight of 16%, which is large compared to the other components but small compared to the CSI weight. Colum 4 of Table 5 therefore reports results in which all 17 components are used, but the CSI weights are restricted to be at most 0.20. With this restriction, housing, food services & accommodation, recreation services, and other services all receive large and approximately equal weight. Notably, these final three sectors fit into the category of locally-determined prices, and two of the three are in the well-measured Category A in Table 3. Otherwise, the weights on the other components do not change much from the benchmark CSI. The resulting index is more highly correlated with the CAI in the pre-estimation sample, and has only a slightly lower correlation with the CAI in the 2000-2019q1 sample, than the benchmark CSI. Although we retain the benchmark CSI, one could make a case for preferring this restricted CSI because it more evenly distributes weight among multiple sectors.

Fifth, the single cyclical activity index imposes either no or second lags (only) of the component cyclical activity variables. As an alternative, we estimated the CSI weights to maximize the correlation between the 13 component well-measured (Category A and B) inflation series (yoy and, alternatively, bandpassed) and the 6 real activity variables including 0-3 lags each for a total of 24 activity indicators. The weights were restricted to be between 0 and 1 and each set of weights (on inflation, and on activity) were restricted to sum to one, so this method corresponds to maximizing the restricted canonical correlation. The resulting activity index is numerically very close to the composite cyclical activity index used in our benchmark estimation, as is the resulting CSI (results not shown).

5.3   Comparison with other inflation indexes

The CSI has both similarities and differences with trimmed mean PCE inflation, median CPI inflation, services-only PCE inflation, and the Stock-Watson (2016) sectoral estimate of trend inflation. Year-over-year changes in the CSI, the trimmed mean PCE, median CPI, and core PCE inflation indexes are plotted in Figure 15, and correlations among these series are given in Table 6. As is evident from both the figure and the table, the CSI, median CPI, and trimmed mean PCE are more highly correlated with each other than they are with core PCE. The reason for these similarities is that the three series place most of their weight on similar components.
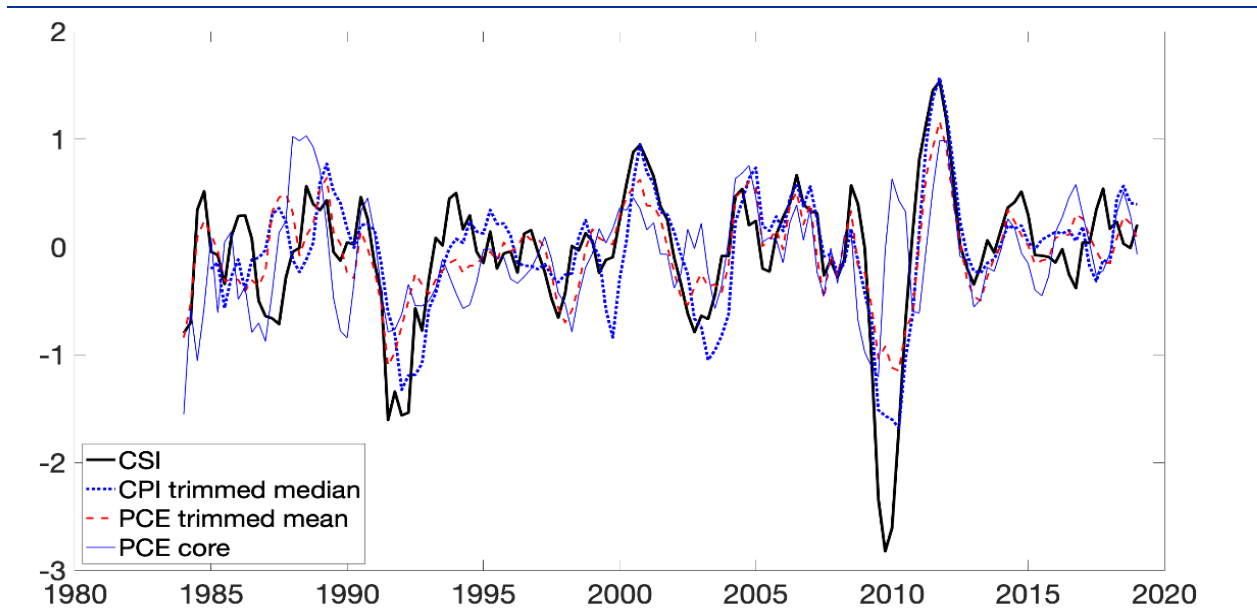
Table 6

Correlations among various inflation indexes, year-over-year changes, 1984q1-2019q2

Index	CSI	Median CPI	Trimmed mean PCE	Core PCE
CSI	1	0.80	0.77	0.34
Median CPI		1	0.83	0.39
Trimmed mean PCE			1	0.60
Core PCE				1

Figure 15

Year-over-year changes in inflation indexes: Core PCE, PCE trimmed mean, median CPI, and CSI

Median CPI. The Federal Reserve Bank of Cleveland reports a weighted median CPI rate of inflation, which is computed as the median of the share-weighted monthly changes in the CPI (using CPI share weights).[12] Because housing (owner-equivalent rent of primary residence) receives a large share weight of CPI expenditures and because it has a large sample size and temporal smoothing built into its estimates, the median CPI is a housing price inflation measure in nearly 57% of months, and in another 17% of months the median CPI is either food away from home, the CPI category corresponding to food services, or recreation, which is mainly recreational services.[13] Thus, in nearly three-quarters of months,

---

[12] Median CPI is computed by first computing a weighted empirical cumulative distribution function for the monthly percentage changes in each of 45 CPI components, using CPI share weights. The median CPI is the monthly percentage change at the 50th percentile of this empirical cdf. See https://www.clevelandfed.org/en/our-research/indicators-and-data/median-cpi/background-and-resources.aspx.

[13] In the CPI, computers appear as education and communication commodities, not under recreational goods as in the PCE. The weights reported here are for the of months from January 1988 through July 2009, the most recent

the median CPI inflation equals one of the well-measured, locally determined prices that receive high weight in the CSI.

Trimmed mean **PCE.** The Federal Reserve Bank of Dallas reports a weighted trimmed-mean PCE rate of inflation, which is computed as the share-weighted average of inflation after trimming out the share-weighted highest and lowest one-sixths of the monthly percentage changes in the PCE components.[14] . On average from 1984-2018, housing has received approximately 30% of the weight in the trimmed mean PCE, and food services & accommodations and food & beverages for off-premises consumption have received approximately 15% of the weight. However, the trimmed mean PCE places weight on some sectors that receive very little weight in the CSI, for example over 1984-2018 health care receives approximately 20% of the trimmed mean PCE weight. The trimmed mean PCE places more weight on cyclically sensitive components than core PCE, but less than the CSI.

Services-only PCE. Because the CSI places most of its weight on services, one question is whether the CSI is essentially replicating the services component of PCE inflation. If so, it would be natural just to use the services component, as Tallman and Zaman (2017) do, rather than CSI.

The first set of comparisons, given in columns 5-8 of Table 5, therefore examines the weights and Phillips correlations for alternative share-weighted measures computed using the PCE inflation components: all PCE components (headline PCE), PCE excluding food & energy (core PCE), PCE for goods, and PCE for services. Each of these share-weighted composites has correlations with the CAI that are much less stable than the CSI, and which are also much lower in both the 1984-1999 and post-2000 samples. The share-weighted index with the strongest correlation with the CAI over the post-2000 period is the services index, but inspection of the weights indicates that, among other things, this index puts substantial weight on poorly-measured series and on series, like health care, that have very limited cyclicality.

Stock-Watson (2016) sectoral trend inflation measure. The final column of Table 5 compares the CSI to a measure, taken from Stock and Watson (2016), that uses the 17 components to create an alternative index of trend inflation. As discussed in Stock and Watson (2016), that index is similar to share-weighted PCE excluding energy. Interestingly, it is more highly correlated with the CAI than the share-weighted indexes, but less correlated and less stably correlated than the CSI.

5.4   Crosswalk from core PCE on gaps to the CSI on band-passed activity index

We have departed from most Phillips curve empirical research in two ways: by examining the components, and by using filters that focus on cyclical correlations. A third departure is using an index of cyclical activity instead of a single cyclical indicator. Here, we pull together our results and assess which of these changes are consequential for finding a stable Phillips correlation, and which are not.

The results are summarized in Table 7. Each row of Panel A represents a different combination of inflation and slack, starting with the standard specification of yoy core PCE and the CBO unemployment gap and ending with bandpass-filtered CSI and the band-pass activity index, the CAI. Inspection of the correlations and slopes reveals that, as long as core PCE is used, the Phillips curve puzzle remains for all the slack measures. If, however, CSI inflation is used, the Phillips correlations and

---

period for which these weighting data have been compiled. See https://www.clevelandfed.org/~/media/content/our%20research/indicators%20and%20data/median%20cpi/revisions%20to%20mcpi%20and%20trimmed%20mean%202007.pdf?la=en.

[14] Specifically, a weighted empirical cumulative distribution function is computed for the monthly percentage changes in each of 178 PCE components, using PCE share weights. The lower 24% and upper 31% of changes in this empirical cdf are removed, and the share-weighted average of the remaining 45% is computed to produce the trimmed mean. See https://www.dallasfed.org/research/pce and, for trimming details, https://www.dallasfed.org/-/media/Documents/research/pce/tech.pdf?la=en.

slopes are somewhat more stable using the CBO unemployment gap than for core PCE, but the full stability finding depends both on using CSI and bandpass activity. Whether one uses the bandpass unemployment rate or the CAI is not consequential.

Panel B of the table provides some additional results for CSI: first, between yoy CSI and a yoy version of the CAI (where the bandpass activity variables are instead yoy filtered), then between bandpass CSI and the CAI. These results are generally similar to the results in the final row of Panel A, for yoy CSI and the CAI, confirming that whether one uses yoy filtering or bandpass filtering is not essential to the finding of stability.

Panel C presents results for the median CPI. Recently, Ball and Mazumder (2019) found that Phillips curve regressions between median CPI (yoy) and HP-filtered unemployment are more stable than the standard yoy core PCE – unemployment gap Phillips curve.[15] As discussed in the previous subsection, median CPI is quite similar to the CSI because the median CPI typically is determined by the same cyclically sensitive series that receive much of the weight in the CSI. It is therefore not surprising that the slopes and correlations between median CSI and the CAI, or the bandpass unemployment gap, are stable, like the CSI and CAI correlation. The stability assessment for the median CPI is hampered because the median CPI does not extend into the 1961-1983 sample period.

Table 7
Crosswalk from standard Phillips curve to CSI-CAI Phillips curve

Inflation	Slack	Correlation			Slope (SE)		
		1961-1983	1984-1999	2000-2019q1	1961-1983	1984-1999	2000-2019q1
A. Crosswalk							
PCE-core, yoy	Unemployment rate, CBO gap	-0.527	-0.473	-0.109	-0.468	-0.250	-0.027
PCE-core, yoy	Unemployment rate, bandpass	-0.462	0.015	-0.191	-0.352	0.007	-0.053
PCE-core, yoy	CAI	-0.590	-0.027	-0.120	-0.446	-0.014	-0.034
CSI, yoy	Unemployment rate, CBO gap	-0.654	-0.308	-0.295	-0.459	-0.16	-0.125
CSI, yoy	Unemployment rate, bandpass	-0.626	-0.474	-0.637	-0.377	-0.236	-0.302
CSI, yoy	CAI	-0.656	-0.556	-0.658	-0.392	-0.281	-0.315
B. Additional CSI specifications							
CSI, yoy	Cyclical index of yoy activity	-0.531	-0.572	-0.785	-0.354	-0.235	-0.376
CSI, bandpass	CAI	-0.571	-0.644	-0.899	-0.262	-0.274	-0.347
C. Median CPI							
median CPI, yoy	Unemployment rate, CBO gap	--	-0.461	-0.297	--	-0.229	-0.106
median CPI, yoy	Unemployment rate, bandpass	--	-0.511	-0.608	--	-0.253	-0.243
median CPI, yoy	CAI	--	-0.456	-0.622	--	-0.216	-0.251

Notes: The CAI is the bandpass cyclical activity index. All slack variables are normalized to have the same standard deviation as the CBO unemployment gap over 1961-2019q1. The cyclical index of yoy activity is the first principal component of the activity variables used to construct the CAI, except filtered using yoy changes instead of bandpass filtered. The median CPI is only available starting 1983. yoy denotes four-quarter change of four-quarter moving average.

6   Conclusions

Different components of inflation have very different cyclical properties. Goods that are traded in international markets tend to have little cyclical variability. Health care prices also have only a small

---

[15] Ball and Mazumder (2019) also consider median PCE and find similar results of stability as they do using median CPI. They constructed median PCE inflation using the raw data from the Dallas Fed trimmed mean PCE. For their regressions, they deviate inflation from the Survey of Professional Forecasters 10-year expected inflation.

cyclical component, perhaps because they are poorly measured or because they are, in many cases, negotiated prices paid on behalf of consumers. In contrast, prices that are determined largely in local markets, such as housing and prices at restaurants and hotels, have large cyclical components. Such prices get the most weight in the CSI index. In addition, some components of inflation are better measured than others, and our results suggest that cyclical movements in headline and core inflation are, in part, masked by noise imparted by the poorly measured components.

By using both inflation components and filters that eliminate trends and focus on cyclical variation, a different picture of the stability of the Phillips curve emerges. Whereas the standard accelerationist relationship between changes in inflation and gaps has flattened, the relationship between the weighted cyclical components and cyclical activity is substantially more stable.

We see the main use of the CSI index as an early indicator that tight – or loose – economic conditions are having an effect on the rate of inflation. Given a set of historically estimated weights, the CSI index can be computed in real time, and in principle can be computed monthly. Given the challenges of estimating slack and cyclical activity in real time, the CSI index provides a new window on movements in the rate of inflation. Because the CSI index tends to focus its weights on sectors with locally determined prices, it provides a way to separate out prices that are domestically determined from prices that are heavily influenced by international conditions.

References

Amstad, M., Potter, S., and Rich, R. (2017). "The New York Fed Staff Underlying Inflation Gauge", Federal Reserve Bank of New York Economic Policy Review 23(2), December.

Ball, L. and Mazumder, S. (2011). "Inflation Dynamics and the Great Recession", Brookings Papers on Economic Activity, Spring 2011, 337-405.

Ball, L. and Mazumder, S. (2014). "A Phillips Curve with Anchored Expectations and Short-Term Unemployment", NBER Working Paper 20715.

Ball, L. and Mazumder, S. (2019). "The Nonpuzzling Behavior of Median CPI Inflation", NBER Working Paper 25512.

Bell, D.N.F. and Blanchflower, D.G. (2018). "The Lack of Wage Growth and the Falling NAIRU", NBER Working Paper 24502.

Blanchard, O. (2016). "The United States Economy: Where To from Here?", *American Economic Review: Papers & Proceedings*, 106(5), 31-34.

Boivin, J., Giannoni, M.P., and Mihov, I. (2009). "Sticky Prices and Monetary Policy: Evidence from Disaggregated Data", *American Economic Review*, 99(1): 350–84.

Coibion, O. and Gorodnichenko, Y. (2015). "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation", *American Economic Journal: Macroeconomics*, 7(1), 197-232.

Cristadoro, R., Forni, M., Reichlin, L. and Veronese, G. (2005), "A Core Inflation Indicator for the Euro Area", *Journal of Money, Credit, and Banking*, Vol. 37, No. 3, pp. 539-560.

Crump, R.K., S. Eusepi, M. Giannoni, and A. Sahin (2019). "A unified approach to measuring u*." *Brookings Papers on Economic Activity*, Spring 2019 (forthcoming).

Dées, S. and Güntner, J. (2017). "Forecasting Inflation across Euro Area Countries and Sectors: A Panel VAR Approach", *Journal of Forecasting*, 36, 431-453.

Eckstein, O. (1981). Core Inflation. New York: Prentice Hall, 1981.

European Central Bank (2014). "The Responsiveness of HICP Items to Changes in Economic Slack", *ECB Monthly Bulletin*, September 2014, 65-67.

Forbes, K. (2018). "Time for a New Astrolabe?" *ECB Forum on Central Banking*, Sintra, Portugal.

Franz, W., and Gordon, R.J. (1993), "German and American Wage and Price Dynamics", *European Economic Review*, May, 719– 62.

Fuhrer, J.C. (2012). "The Role of Expectations in Inflation Dynamics", *International Journal of Central Banking* 8 (Supplement 1), 137-166.

Gilchrist, S., Schoenle, R., Sim, J. and Zakrajšek, E. (2017). "Inflation Dynamics during the Financial Crisis", *American Economic Review* 107(3), 785-823.

Galí, J. (2011). "The Return of the Wage Phillips Curve", *Journal of the European Economic Association* 9(3), 436-461.

Garner, A.C. (1994), "Capacity Utilization and US Inflation", *Federal Reserve Bank of Kansas City Economic Review*, Fourth Quarter, 1–21.

Goldfeld, S.M. and A.S. Blinder (1972). "Some Implications of Endogenous Stabilization Policy," *Brookings Papers on Economic Activity* 3:1972, 585-644.

Gordon, R. (2013). "The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery", NBER Working Paper 19390.

Groen, J.J.J., Paap, R. and Ravazzolo F. (2013). "Real-Time Inflation Forecasting in a Changing World," *Journal of Business & Economic Statistics* 31(1), 29-44.

Hood, K.K. (2013). "Measuring the Services of Commercial Banks in the National Income and Products Accounts: Changes in Concepts and Methods in the 2013 Comprehensive Revision", *Survey of Current Business*, February 2013, 8-19.

Hooper Peter, Frederic S. Mishkin and Amir Sufi (2019). "Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating? National Bureau of Economic Research working paper 25792.

Kareken, J. and R.M. Solow (1963), "Lags in Monetary Policy," in E.C. Brown et. al., *Stabilization Policies*, Prepared for the Commission on Money and Credit, 14-96.

Kiley, M.T. (2015). "An Evaluation of the Inflationary Pressure Associated with Short- and Long-term Unemployment", *Economics Letters* 137, 5-9.

Krueger, A.B., Cramer, J. and Cho, D. (2014). "Are the Long-Term Unemployed on the Margins of the Labor Market", *Brookings Papers on Economic Activity*, Spring 2014, 229-299.

Mahedy, T. and Shapiro, A. (2017). "What's Down with Inflation?" *Federal Reserve Bank of San Francisco Economic Letter*, 2017-35, November 27, 2017.

Mazumder, S. (2018). "Inflation in Europe after the Great Recession," *Economic Modelling* 71, 202-213.

McLeay, Michael, and Silvana Tenreyro (2019). "Optimal inflation and the identification of the Phillips curve," forthcoming, *NBER Macroeconomics Annual* 2019.

Moulton, B.R. (2018). "The Measurement of Output, Prices, and Productivity: What's Changed since the Boskin Commisson?", manuscript, Hutchins Center, Brookings Institution.

Nalewaik, J.J. (2010), "The Income- and Expenditure-Side Estimates of US Output Growth", *Brookings Papers on Economic Activity*, Spring 2010, 71-127.

Peach, R., Rich, R. and Lindner, H.M. (2013). "The Parts are More than the Whole: Separating Goods and Services to Predict Core Inflation", *Current Issues in Economics and Finance*, 19(7), Federal Reserve Bank of New York.

Orphanides, A. and van Norden, S. (2002). "The Unreliability of Output-Gap Estimates in Real Time", *The Review of Economics and Statistics*, 84(4), 569-583.

Stock, J.H. (2011). "Discussion of 'Inflation Dynamics and the Great Recession' by Ball and Mazumder", *Brookings Papers on Economic Activity*, Spring 2011, 387-402.

Stock, J.H. and Watson, M.W. (2010). "Modeling Inflation after the Crisis", in Macroeconomic Policy: Post-Crisis and Risks Ahead, *Proceedings of the Federal Reserve Bank of Kansas City 2010 Jackson Hole Symposium*, 173-220.

Stock, J.H. and Watson, M.W. (2016a). "Sectoral Inflation Measures and Cyclically Sensitive Inflation", Presentation deck, Federal Reserve Bank of Cleveland conference *2016 Inflation: Drivers and Dynamics,* September 30, 2016.

Stock, J.H. and Watson, M.W. (2016b). "Core Inflation and Trend Inflation", *The Review of Economics and Statistics* 98(4) 770-784.

Struyven, D. (2017). "US Daily: Which Prices Still Respond to Slack?", Goldman Sachs Economic Research.

Tallman, E. and Zaman, S. (2017), "Forecasting Inflation: Phillips Curve Effects on Services Price Measures", *International Journal of Forecasting* 33(2), 442-457.

US Bureau of Economic Analysis (2017). *NIPA Handbook: Concepts and Methods of the US National Income and Product Accounts*, Chapter 5.

Watson, M.W. (2014). "Inflation Persistence, the NAIRU, and the Great Recession", *American Economic Review: Papers & Proceedings* 104(5), 31-36.