Breaks in the Phillips Curve: Evidence from Panel Data

Simon C. Smith[a], Allan Timmermann[b], Jonathan H. Wright[c]

[a] *Federal Reserve Board*
[b] *University of California, San Diego*
[c] *Johns Hopkins University*

Draft: October 25, 2023

Abstract

We revisit the Phillips curve, applying new Bayesian panel methods with structural breaks to US and EU disaggregate data. Our approach lets us estimate both the number and timing of breaks and to determine the existence of clusters of industries, cities, or countries whose Phillips curves display similar patterns. We find evidence of a flattening for US sectoral data and among EU countries, particularly poorer ones. Evidence of flattening is weaker for MSA-level data and the wage Phillips curve. We find evidence of a kink in the Phillips curve which remains relatively steep when the economy is running hot. We find that this can explain about half of the pandemic-era runup in inflation. Our methods are useful for monitoring changes in the Phillips curve going forward.
Keywords: Phillips curve, Inflation, Unemployment, Panel data, Structural breaks, Bayesian analysis, Covid pandemic, real time inflation monitoring

JEL classifications: C11, C22, E51, E52

Email addresses: `simon.c.smith@frb.gov` (Simon C. Smith), `atimmermann@ucsd.edu` (Allan Timmermann), `wrightj@jhu.edu` (Jonathan H. Wright)

Comments from Francesco Furlanetto (discussant), Refet Gürkaynak, and Callum Jones (discussant) as well as seminar and conference participants at the 2021 IAAE webinar, Deutsche Bundesbank, 2023 Bank of Canada/University of Toronto conference "The Return of High Inflation: Challenges for Monetary Policy", Federal Reserve Board, 2023 North American Summer Meeting of the Econometric Society, 2023 CEBRA Annual Meeting, and the 2023 Federal Reserve System Meeting of Econometrics are gratefully acknowledged. Any remaining errors are our own. The views expressed in this paper are those of the authors and do not necessarily reflect the views and policies of the Board of Governors or the Federal Reserve System.

1. Introduction

"There was a time where there was a tight connection between unemployment and inflation. That time is long gone." (Jerome Powell, 2021.)[1]

The Phillips curve is a key element of the New-Keynesian macroeconomic model and is critical in how central banks think of the macroeconomy. Recently there has been much debate about a potential flattening of the Phillips curve, which could, in turn, hinder the central banks' ability to control inflation and, thus, have major policy implications. The goal of this paper is to examine the strength of the empirical evidence on changes to the slope of the Phillips curve across multiple panel data sets covering industries, cities, states and countries, and prices as well as wages. Coupled with a novel, flexible panel estimation methodology, we provide the most complete picture available on instabilities in the Phillips curve.

Examining disaggregate data, whether categorized by industry, by region, or by country, offers notable advantages. First, the presence of cross-sectional heterogeneity can shed light on the causes of changing Phillips curves. Second, since different regions and sectors experience different business cycles, there is extra information in disaggregate data that enables us to identify slope coefficients and regime changes more precisely than using aggregate data alone. Third, several recent papers (e.g. Hooper *et al.* (2020), Fitzgerald *et al.* (forthcoming) and McLeay and Tenreyro (2020)) have pointed out that if the central bank is successfully targeting inflation, then this creates an endogeneity bias in the slope of the Phillips curve, biasing the coefficient towards zero. The use of disaggregate data in conjunction with the inclusion of time fixed effects avoids this problem, because the central bank does not specifically target inflation in any one particular region or sector.[2]

Building on these insights, in this paper we apply a novel Bayesian panel estimation approach to study breaks in Phillips curve inflation dynamics. Our analysis offers three central new insights. First, all the existing literature on regional or sectoral Phillips curves requires a common slope coefficient (complete pooling). If series display heterogeneous dynamics, however, pooled estimates will be biased (Canova forthcoming). Alternatively, estimating

[1]This quote is from Federal Reserve Chair Jerome Powell's press Conference, March 17, 2021; https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20210317.pdf.

[2]The problem would not be solved with disaggregate data without time fixed effects, because in that case some of the identification would come from the time series dimension where there is endogeneity.

the Phillips curve for each industry or region separately gives up a lot of information. Our approach instead allows us to pursue a middle ground, partial pooling, while allowing for cross-sectional variation in the slope coefficients. For example, our methodology allows us to consider groupings by industry or geographic region, with different slope coefficients applying to each group. We can impose the groupings *a priori*, or the grouping structure can be estimated as part of the modeling process. If the data support a homogeneous Phillips curve that is identical across all units, only a single group will be identified. Conversely, strong heterogeneity in Phillips curves across industries or regions will lead to a model in which each group comprises a single unit. Our methodology endogenously determines whether any of these special cases or an intermediate scenario with multiple units in each cluster, is supported by the data, adapting to the degree of heterogeneity found. We find that partial pooling with a structure comprising a small number of groups is supported empirically in most cases as parameters are not completely homogeneous.

Secondly, while many researchers have considered time-variation in Phillips curves, they do so by assuming that the parameters of the Phillips curve follow a random walk (Ball and Mazumder 2011; Matheson and Stavrev 2013; Blanchard 2016; Inoue *et al.* 2022), or are subject to breaks at pre-specified dates or dates determined by the single-breakpoint test of Andrews (1993) or based on regressions with rolling windows (Roberts 2006; Coibion *et al.* 2013; Coibion and Gorodnichenko 2015; Leduc *et al.* 2017; Ball and Mazumder 2019; Galí and Gambetti 2019; Gilchrist and Zakrajšek 2019; Del Negro *et al.* 2020; Fitzgerald *et al.* forthcoming; Hooper *et al.* 2020; Cerrato and Gitti 2022; Hazell *et al.* 2022).[3] However, break tests conducted on individual inflation series have low power, making it difficult to detect breaks in the Phillips curve. As argued by Bai *et al.* (1998) and Smith and Timmermann (2021), imposing common timing of breaks for panel data can increase not only the power of tests for breaks but also the precision of break date estimates. Our study is the first to formally consider panel data estimation of Phillips curves with multiple breaks.[4] We estimate the number of breaks to the Phillips curve and the time of their occurrence, which would be infeasible to do with any precision with aggregate data or individual inflation series. We discuss interpretations of changes in the slope of the Phillips curve. Ball *et al.* (1988) argue that a higher level and volatility of inflation leads to a steeper Phillips curve. A complex set of other factors could be at play, including changes in unionization and

[3] Barnichon and Mesters (2021) use a subsample split to track time-variation in the Phillips *multiplier*. Table 1 contains a list of some of the main studies on variation in the Phillips curve.

[4] Allowing for multiple breaks is also crucial when considering long time series.

wage indexation, exposure to international trade, and even labor market integration. We also consider an extension that allows breaks to affect individual inflation series at different times, in a "non-common breaks" specification, uncovering lead-lag dynamics that can help identify the underlying economic causes of instability.

Thirdly, we demonstrate a close tie between instability in the slope of the Phillips curve and nonlinearity reflecting the steepness of this curve at low levels of unemployment. Consistent with the recent inflationary experience, we find that the Phillips curve steepens notably in very tight labor markets and show that this effect tends to be stronger over the most recent period.

In our empirical work, we consider US price Phillips curves using disaggregation at the industry and MSA level, and wage Phillips curves at the state level. We also examine Phillips curves at the country level within the European Union.

In US industry data covering the sample 1959Q1-2022Q3, we find two regime changes in the Phillips curve; a steepening around 1972 and a flattening in 2001. Moreover, the recent flattening of the Phillips curve is more pronounced for goods prices than for services prices. The steepening around 1972 comes after a period when inflation had been trending up for some years and when indexation of wage contracts, either implicit or explicit, became more common. This would in turn steepen the Phillips curve. Meanwhile, the subsequent flattening corresponds to a time of greater import penetration, especially from China, with China joining the World Trade Organization in 2001.[5] While economic intuition might suggest that changes to the slope of the Phillips curve would occur gradually, China's accession to the WTO may have caused more of a sudden break with quite sharp effects documented in studies such as Bena and Simintzi (2022).[6] The rise in the level and volatility of inflation in the 1970s, and their subsequent declines in the 1990s can also be thought of as drivers of the steepening and subsequent flattening of the Phillips curve as higher and more volatile inflation leads firms to adjust prices more frequently (Ball *et al.* 1988). Declining unionization is also a possible partial explanation for the flattening of the Phillips curve.[7] These

[5] Auer *et al.* (2017), Gilchrist and Zakrajšek (2019), Firat (2020), and Stock and Watson (2020) all show how greater trade openness can flatten the Phillips curve.

[6] While our breakpoint approach assumes that the Phillips curve undergoes an instant shift, it does incorporate uncertainty surrounding the break date into the parameter estimate. Greater uncertainty surrounding the break date will lead to a smoother evolution of the Phillips curve slope which might be observationally approximate to a slow-evolving parameter approach such as a time-varying parameter model. Of course, the two approaches are philosophically distinct.

[7] While we do find a flattening break in the wage Phillips curve (estimated with state data over the sample 1980Q1 - 2019Q4), it is of smaller magnitude than for the price Phillips curve.

regime changes that we detect are broadly consistent with much of the existing literature (e.g. Hooper *et al.* (2020)), although we estimate the break dates much more precisely.

US regional (MSA) data are not available as far back in time, spanning the shorter sample 1980-2022. This means that we cannot examine the presence of Phillips curve breaks in the 1970s for this data. Still, even with this shorter sample coverage, we manage to identify a regime change around 2000. Further, we find that MSAs with above (below) median rates of import penetration from China have experienced a considerably stronger (weaker) flattening of their price Phillips curve. These findings are consistent with more goods competition from China explaining a part of the flattening of the price Phillips curve.

Broadly similar patterns are found in the EU for a sample that begins in 1986 and ends in 2021. For this data we find evidence of a single break which we estimate occurs in 2004 at which point the slope of the Phillips curve flattens significantly. Using our clustering methodology, we find that the Phillips curve used to be particularly steep in poorer (mostly East European) countries prior to the 2004 break, but has flattened by more in those countries, consistent with clear evidence of Phillips curve convergence across countries that, early in our sample, used to display a very different inflation-unemployment trade-off.

We also study nonlinearity of the Phillips curve, which, as noted by Hooper *et al.* (2020), is much easier to do with disaggregate data since the national labor market has not really been tight since the late 1960s, whereas many individual MSAs have had tight labor markets in this time period. Since the current policy debate is focused on such tight values of the labor market, regional data seem likely to be helpful here. We consider a kink in the Phillips curve at threshold unemployment rates of 5 or 4.2 percent.[8] Using these thresholds, we find that the Phillips curve is steeper in a tight labor market. Hooper *et al.* (2020), Babb and Detmeister (2017) and Leduc *et al.* (2019) also find that the Phillips curve is steeper in a tight labor market but do not consider subsample instability.

Next, we explore some implications of our Phillips curve estimates for the recent inflation experience. Our estimates for both the US and the EU imply essentially no missing disinflation during the Great Recession and no missing reinflation during the subsequent recovery years. In addition, we find that a steeper (nonlinear) Phillips curve in hot labor markets combined with a higher natural rate of unemployment (Crump *et al.* 2022) can explain almost half of the surge in U.S. inflation between 2020 and 2022.[9] Our methodology provides

[8]For comparison, Stock and Watson (2009) define a tight labor market as an unemployment gap below minus 1.5 percent while Babb and Detmeister (2017) use the same thresholds as we do.

[9]The kinked (nonlinear) Phillips curve effects are an important part of this explanation.

a tool for early detection shifts in the Phillips curve going forward, which is especially useful if we are entering a period of greater structural change as conjectured by Lagarde (2023).

Our analysis adopts a Bayesian approach so a natural concern is to what extent our empirical results hinge on our choice of priors. We choose uninformative priors centered on zero for the critical slope parameters of the Phillips curve, thus, if anything stacking the odds against finding breaks.[10] The Bayesian approach does, however, offer key advantages over more traditional panel estimation methods. First, by explicitly accounting for parameter and model (break) uncertainty, the Bayesian approach allows us to characterize important features such as the degree of precision with which we estimate the number of breaks to the Phillips curve and their location. Second, it allows us to combine a lot of features such as group structure and heterogeneity in the timing of breaks that would be very hard to do in a frequentist setting. A third advantage is that our approach can be used to examine possible breaks at different frequencies by varying the expected time between breaks. Our main analysis assumes a regime duration of twenty years and so focuses on infrequent, "secular" shifts, filtering out noise in the short-term relationship between inflation and labor market tightness. However, we also implement our analysis with a prior expected regime duration of only five years. We show that this facilitates better real-time stability monitoring and allows us to identify a break and significant steepening of the Phillips curve in 2020Q1 for the sectoral CPI and PCE data.

The remainder of the paper proceeds as follows. Section 2 introduces the panel data sets used in our analysis while Section 3 explains our Bayesian panel approach, including estimation, model selection and choice of priors. Section 4 presents our main empirical results on breaks in the industry and regional Phillips curves, and Section 5 discusses implications of our results for the recent inflation experience. Section 6 conducts a set of robustness exercises, while Section 7 concludes. Additional empirical results are presented in an Appendix.

2. Data

This section introduces our data along with the data sources used in our empirical analysis. We begin by introducing our price and wage series before describing our aggregate unemployment gap measures, unemployment rate, and NAIRU measures.

[10] Moreover, Jones *et al.* (2021) note that priors have relatively little influence on posteriors when estimating regional/sectoral Phillips curves (with pooled parameters) as we do.

2.1. Price and Wage Data

2.1.1. MSA level

We source monthly total CPIs for 22 MSAs from the BLS. We construct annual levels as the average of all monthly observations in the corresponding year and compute annual inflation rates as $log(CPI_{it}/CPI_{it-1}) \times 100$ in which CPI_{it} denotes the level for the ith MSA in year t.[11] Our sample for these data begins in 1980 and ends in 2022, but for many MSAs the data only start in 1990.

2.1.2. Industry level

We use quarterly Personal Consumption Expenditures price indexes (PCE) for 16 industry components, similar to those analyzed by Stock and Watson (2020), sourced from the Bureau of Economic Analysis (BEA).[12] Our sample is 1959:Q1 - 2022:Q3. We construct annualized quarterly inflation rates as $log(PCE_{i,t}/PCE_{i,t-1}) \times 400$.

From the BLS, we source monthly CPI inflation for 31 "level 3" industries, as currently formulated, beginning in January 1954 and ending in September 2022, though not all series go all the way back. We construct our annualized quarterly inflation rate observations from end of quarter monthly observations as $log(CPI_{i,t}/CPI_{i,t-4}) \times 400$.

2.1.3. EU data

We source headline (as well as total goods and total services) annual inflation rates for our 28 countries (the 27 current members and the UK) from the ECB statistical warehouse. Our sample begins in 1986 and ends in 2021.

2.1.4. Wage data

Following Hooper *et al.* (2020), we compute average hourly earnings (AHE) for each of the 50 states and the District of Columbia using the latest (2019) CEPR uniform extract

[11] Data for all but a few MSAs are collected only in either odd or even months. See https://www.bls.gov/opub/hom/cpi/pdf/cpi.pdf for details of the complete methodology and https://www.bls.gov/cpi/additional-resources/geographic-sample.htm. for the geographic definitions.

[12] The two categories – Housing and Household utilities – have since been replaced by one: Housing and utilities.

from the Current Population Survey (CPS)[13]. Aggregating from monthly data, we construct quarterly data from 1980:Q1 through 2019:Q4, from which we construct quarterly annualized wage inflation.

2.2. Unemployment rates, NAIRU, and Inflation Expectations

We use the end-of-quarter monthly aggregate unemployment gap, measured as the difference between the unemployment rate from the U.S. Bureau of Labor Statistics (BLS) and the NAIRU estimate (from the Congressional Budget Office). These data begin in January 1949 and end in September 2022.

We source the annual country-level unemployment rate and NAIRU estimates for the 28 EU member countries (the current 27 plus the UK which was a member until recently), and hence the unemployment gaps, for the sample period 1965-2021 from the DG ECFIN/AMECO—the European Commission's macroeconomic database.[14]

For the regional analysis, we obtain annual unemployment rate data from 1980 to 2022 for 22 MSAs from the BLS. We also use the end of quarter monthly unemployment rate for all 51 states (including the District of Columbia), also obtained from the BLS. These data begin in January 1980 and end in December 2019.

We source four-quarter-ahead Consumer Price Index (CPI) inflation expectations from Blue Chip Economic Indicators. These data go back to 1985. Between 1980 and 1985, we use Producer Price Index (PPI) inflation expectations from the same source. Before 1980, we use data from Livingston which is only updated every six months and so we simply repeat observations in the two corresponding quarters, effectively assuming that inflation expectations remain the same in each 6-month period. Because U.S. inflation expectations are only measured for the aggregate price index as opposed to at the regional or sectoral level, we can only use these inflation expectations data in specifications without time fixed effects.

2.3. Group structure

We will be interested in group heterogeneity, with either the group allocation imposed according to pre-determined selection criteria, or determined by the Bayesian algorithm as part

[13] The data are available from https://ceprdata.org/cps-uniform-data-extracts/.

[14] We thank Michele Lenza for helping us access these country-level NAIRU estimates.

of the estimation process.

The 16 PCE sectors are split into goods – Motor vehicles and parts, Furnishings and durable household equipment, Recreational goods and vehicles, Other durable goods, Food and beverages purchased for off-premises consumption, Clothing and footwear, Gasoline and other energy goods, and Other nondurable goods – and services – Housing and utilities, Health care, Transportation services, Recreation services, Food services and accommodations, Financial services and insurance, Other services, and NPISH.

We also split the 28 EU countries into rich and poor countries with rich countries defined as countries with real GDP per capita deflated by PPP in 2019 above the EU average and poor countries defined as the rest. The rich countries include Luxembourg, Ireland, Denmark, Netherlands, Austria, Germany, Sweden, Belgium, Finland, France, and UK.[15]

3. Methodology

Our analysis examines three different Bayesian panel specifications. The first is our baseline pooled panel model with multiple breakpoints. This model applies the methodology developed by Smith and Timmermann (2021) to exploit information in the cross-section and obtain increased power to detect structural breaks. Breaks are assumed to be common, i.e., they hit every series in the cross-section at the same time. To summarize, this model assumes homogeneity both in the timing of any breaks and in their impact on individual variables.

To gain further insight into the break dynamics, our second model relaxes the common break-timing assumption, allowing series to be hit at different times. We accomplish this using the methodology developed by Smith (2018) which is designed to detect lead-lag relations in the impact of breaks across different variables in the cross-section. This approach can, thus, shed light on the diffusion of breaks and the speed at which different sectors, regions, or countries are affected by breaks to their Phillips curves.

Our third model endogenously estimates both the number of groups and the assignment of each series to a group using the methodology developed by Smith (2023). Relative to the baseline model that pools parameters across the entire cross-section, this model pools parameters across all series within a group, but allows the parameters to differ across groups. This provides an effective way to allow for heterogeneity in the impact of breaks on individual

[15]The poor countries therefore include Malta, Italy, Czech Republic, Spain, Cyprus, Slovenia, Slovakia, Romania, Portugal, Poland, Bulgaria, Estonia, Lithuania, Latvia, Hungary, Greece, and Croatia.

variables.[16] The baseline homogeneous (pooled) panel model arises as a special case of this specification when the data only identifies a single group. At the other extreme, a model where each series in the cross-section gets assigned to its own individual group would allow for complete heterogeneity.

3.1. Common breakpoint model

The first–baseline–model we take to the data allows for an unknown number of K breaks occurring at unknown times $\tau = (\tau_1, \ldots, \tau_K)$ which are assumed to be common to all $i = 1, \ldots, N$ series in the cross-section.[17] Our first specification is for the Phillips curve at the MSA level. The data are annual, and the model for the kth regime takes the form (for $k = 1, \ldots, K+1$):

$$\pi_{it} = \alpha_i + \gamma_t + \rho_k \pi_{it-1} + \lambda_k URATE_{it} + \epsilon_{it}, \qquad t = \tau_{k-1}+1, \ldots, \tau_k \tag{1}$$

in which π_{it} denotes the inflation rate for the ith series at time t, α_i and γ_t denote two-way fixed effects, π_{it-1} is the lagged inflation rate for variable i, $URATE_{it}$ denotes the unemployment rate for the ith series at time t, and ϵ_{it} is the residual for the ith series at time t which is assumed to be normally distributed $\epsilon_{it} \sim N(0, \sigma_{ik}^2)$, so we allow volatility to vary across individual variables.[18]

The parameters ρ_k, λ_k, and σ_{ik}^2 are all allowed to shift across regimes separated by a break, but the former two are assumed to be identical across all series within a given regime, effectively following step functions that shift at τ_k. Hall (2023) argues that the Phillips curve is steeper in times of high volatility because volatile price determinants reduce price stickiness as a larger fraction of sellers elect to reset their prices. Allowing volatility to vary across breaks could thus be important in identifying shifts in the steepness of the Phillips curve slope.

[16]We generally condition on the regimes identified by the baseline model when implementing this third model in our empirical analysis.

[17]For simplicity our notation uses N as the cross-sectional dimension, but our approach can readily handle unbalanced panels with a time-varying dimension, N_t.

[18]This specification corresponds to Equation (19) of Hazell *et al.* (2022) and so we are estimating "ψ" in their terminology rather than "κ" from their Equation (17). Like Cerrato and Gitti (2022), we choose this specification because estimating the latter requires dropping the final 5-10 years of data which is problematic when estimating breaks in the Phillips curve, particularly when the latter part of the sample includes the inflationary surge unleashed by the COVID-19 pandemic. Note that we omit the relative price of nontradeables and replace it with lagged inflation.

Our baseline model assumes that the residuals ϵ_{it} are cross-sectionally and serially uncorrelated. This assumption means that we are not required to estimate the $N(N-1)/2$ covariance terms in each break segment but may not be empirically valid in some empirical applications. Section 6 discusses how to test the validity of this assumption. More broadly, we can allow for cross-sectional correlation in ϵ_{it} through a common factor structure that allows for heterogeneity in factor loadings across units but assumes that the idiosyncratic shocks that remain, after accounting for the common factors, are orthogonal across i.

The specification in Equation (1) uses the unemployment *rate* rather than the unemployment *gap* as the slack measure. At the MSA level, there are no estimates of the natural rate of unemployment and while we could HP detrend the city-level unemployment data, such estimates would be sensitive to the bandwidth parameter. We instead rely on the two-way fixed effects to absorb variation in the natural rate across time and cities. Common time variation in inflation expectations in Equation (1) is also absorbed by the time fixed effects.

The same model is applied to the EU-level data, except that the unemployment gap replaces the unemployment rate since we have NAIRU estimates for EU countries unlike for the US MSAs:

$$\pi_{it} = \alpha_i + \gamma_t + \rho_k \pi_{it-1} + \lambda_k UGAP_{it} + \epsilon_{it}, \qquad t = \tau_{k-1}+1, \ldots, \tau_k. \tag{2}$$

For the US industry-level data, using either PCE or CPI, we do not observe industry-level unemployment rates, let alone a NAIRU estimate.[19] For this case, we substitute the aggregate unemployment gap, $UGAP_t$, for the disaggregate unemployment gap in Equation (2). This means we must drop the time fixed effects which are not separately identifiable from the aggregate unemployment gap. Finally, we include four-quarter-ahead CPI inflation expectations, BC_t, which are identified in the absence of time fixed effects, yielding the model:

$$\pi_{it} = \alpha_i + \rho_k \pi_{it-1} + \lambda_k UGAP_t + \psi_k BC_t + \epsilon_{it}, \qquad t = \tau_{k-1}+1, \ldots, \tau_k. \tag{3}$$

Note that in this specification, the data are at a quarterly frequency.

[19] BLS does have some data on industry-level unemployment, in the sense of breaking out unemployment by the sector of the unemployed worker's last job, but this data only goes back to 2000, and the concept of unemployment by sector is inherently hard to measure.

3.1.1. Implied aggregate Phillips curve slopes

Hazell *et al.* (2022) show that the regional Phillips curve slope can be divided by the expenditure share on nontradeables to obtain the national Phillips curve slope. We use the 31 CPI industry weights to compute the expenditure share on nontradeables. We follow Hazell *et al.* (2022) by assigning the following series to nontradeables: Full Service Meals and Snacks, Limited Service Meals and Snacks, Food at employee sites and schools, Food from vending machines and mobile vendors, Other food away from home, Electricity, Utility (piped) gas service, Water and sewer and trash collection services, Household operations, Medical care services, Transportations services, Recreation services, Education and communication services, Other personal services, and Shelter. The expenditure share on nontradeables is therefore 69.1 percent.

This scaling approach is applied to our regional Phillips curve estimates using MSA-level data to infer the U.S. aggregate Phillips curve slope coefficient when we consider the aggregate implications of our estimates in Section 5.[20] We apply the same scaling approach to infer the EU aggregate Phillips curve from our EU country-level estimates.[21]

3.2. Noncommon breakpoint model

For parsimony, we only formally exposit the noncommon breakpoint model that generalizes the common breakpoint model detailed in Equation (1).[22] The only difference is that the break timing, which was previously common (τ_k), is now allowed to differ across series (τ_{ik}). Formally, for the MSA-level annual data the model is (for regimes $k = 1, \ldots, K+1$)

$$\pi_{it} = \alpha_i + \gamma_t + \rho_k \pi_{it-1} + \lambda_k URATE_{it} + \epsilon_{it}, \qquad t = \tau_{ik-1} + 1, \ldots, \tau_{ik}. \tag{4}$$

[20] We do not infer the U.S. aggregate Phillips curve slope from our U.S. sectoral estimates for the aggregate implications because the absence of time fixed effects means that the argument for identifying the Phillips curve from disaggregate data does not go through.

[21] Hazell *et al.* (2022)'s theory requires that all regions/countries belong to a single monetary union. Since several EU countries are not in the eurozone, one might be concerned that monetary policy might in fact respond to their country-specific shocks. However, some of these countries effectively peg their currency to the euro (e.g. Denmark) and others are heavily influenced by it. The assumption that monetary policy does not respond to their country-specific shocks therefore seems reasonable. Also, our results that follow are essentially unchanged when we omit the UK, which is the country in our sample that most obviously conducts independent monetary policy.

[22] The models that use either EU data – displayed in Equation (2) – or U.S. industry-level data – displayed in Equation (3) – generalize in the obvious way.

While this specification does not impose that the timing of the breaks is identical across all variables, we control the degree of heterogeneity in the timing of breaks across units by effectively only considering "local" variation in the break timing, i.e., breaks whose occurrence is close to the break date for the majority of variables. This prevents our approach from identifying idiosyncratic breaks in the individual series and enables us to use cross-sectional information to more accurately identify clusters of breaks whose impact can spread across units at different speeds.

Intuitively, the approach works by identifying break *windows* rather than single break points. Variables can be hit at any time within a given break window. For example, a common break approach might identify a break around the Global Financial Crisis in September 2008 when Lehman Brothers failed. The break window approach, however, might identify a local break window of, say, 6-12 months during which firms were hit by the break at different times as the financial crisis cascaded through the economy. We control the degree of heterogeneity in the timing of breaks across series through the prior, detailed in Section 3.4.

3.3. Grouped heterogeneity model

So far we assumed homogeneity in the regression coefficients and, consequently, in the effect of breaks on individual variables. However, in many cases both the slope coefficients and the impact of breaks may differ across sectors, regions, or countries. For such cases, it is important to allow for heterogeneous parameters. We accomplish this by assuming the existence of G_k groups or clusters of variables and allowing parameters to vary across groups while they are the same within groups. Each unit in the cross-section belongs to a single group ($i \in g_k$) and, by estimating this algorithm conditional on the regimes identified by the baseline breakpoint model, both the group membership and the number of groups is allowed to vary across regimes. This approach offers a flexible specification. For example, we can allow for full heterogeneity in a given regime by setting $G_k = N$, whereas homogeneity within the regime corresponds to $G_k = 1$. Values of G_k between these extremes indicate some degree of clustering within that regime. Moreover, variation across regimes in the number of clusters can provide important information about issues such as convergence (or lack thereof) in the Phillips curves across units.

Using the model for the EU-level data as our lead example, we estimate the following model in each of the $k = 1, \ldots, K + 1$ regimes identified by the baseline model[23]

[23]The models that use industry-level data are not formally exposited for simplicity, but follow the same

$$\pi_{it} = \alpha_i + \gamma_t + \rho_{g_k}\pi_{it-1} + \lambda_{g_k} UGAP_{it} + \epsilon_{it}, \qquad t = \tau_{k-1}+1, \ldots, \tau_k, \tag{5}$$

where $\epsilon_{it} \sim N(0, \sigma^2_{ik})$. The parameters ρ_{g_k} and λ_{g_k} are pooled across all series within the g_kth group, but differ across the G_k different groups. The number of groups and the series assigned to each group can be either specified *a priori* or alternatively determined as part of the estimation. In the latter case, our priors lean against identifying groups that contain only a single series, thus reducing the likelihood of simply identifying outliers in the data.[24]

3.4. Prior distributions

Our Bayesian panel break approach requires us to specify priors on the regime durations and regression parameters. Using the baseline model as our lead example, we next explain how these priors are set. We further specify our priors on the break lags in our second, noncommon breaks, model and our priors on the clustering (grouping) model.

3.4.1. Prior on regime durations

Following Koop and Potter (2007), the regime durations, $l_k = \tau_k - \tau_{k-1}$, follow a Poisson prior distribution

$$p(l_k \mid \zeta_k) = Po(\zeta_k), \qquad k = 1, \ldots, K+1, \tag{6}$$

where the intensity parameter ζ_k follows a conjugate Gamma prior distribution

$$p(\zeta_k) = Ga(c, d), \qquad k = 1, \ldots, K+1, \tag{7}$$

and c and d are the hyperparameters of ζ_k. These hyperparameters only determine the average regime duration since the expected regime durations are allowed to differ across breaks, with each individual regime having its own unique intensity parameter.

Our analysis calibrates the prior hyperparameters determining the regime duration so that breaks occur, on average, every twenty years. We achieve this by setting $d = 2$ and

structure.

[24] Our Bayesian approach has two features that help determine the number of groups. First, the marginal likelihood guiding the estimation prefers fewer groups and penalizes additional groups since these require estimating more parameters. Second, we use a prior that an average of five series comprise a group and apply a penalty to very small and very large groups although these are not ruled out. Still, our prior is towards not having groups with just a single member. In cases where we find empirically that some groups have just a single or very few members, the empirical evidence therefore strongly supports separating these units.

$c = 40$ ($c = 160$) for the annual (quarterly) data, respectively. Our priors are thus set to focus on rare, "secular", breaks in the Phillips curve.[25]

3.4.2. Priors on regression parameters

For regimes $k = 1, \dots, K+1$, we follow conventional practice and specify an inverse gamma prior distribution over the residual variances

$$p(\sigma_{ik}^2) \sim IG(a, b), \qquad i = 1, \dots, N, \tag{8}$$

while we assume a Gaussian prior on the regression coefficients

$$\begin{aligned} p(\lambda_k) &\sim N(0, \sigma_\lambda^2), \\ p(\rho_k) &\sim N(0, \sigma_\rho^2). \end{aligned} \tag{9}$$

Here σ_λ^2 and σ_ρ^2 are hyperparameters that control the degree to which λ_k and ρ_k are shrunk towards their prior means of zero.[26]

Our analysis sets $a = 2$ and $b = 1$, while σ_λ^2 and σ_ρ^2, which control the degree to which λ_k and ρ_k are shrunk towards their prior means (zero), are both set equal to 0.1. These are fairly uninformative priors which allow the autoregressive parameter and the slope of the Phillips curve to vary with the data.

3.4.3. Priors on heterogeneity in break dates

Our second specification allows for differences in the point in time when breaks affect the individual series within a break window, the length of which is estimated. Let τ_k denote the date at which the kth break window begins. The lag with which the ith series is hit by the kth break is denoted $\Delta_{ik} = \tau_{ik} - \tau_k$ which can be zero (hit immediately at the beginning of the break window) as well as positive (hit with a lag). We specify a Poisson prior over such break delays

$$p(\Delta_{ik} \mid \delta_k) \sim Po(\delta_k), \qquad k = 1, \dots, K, \qquad i = 1, \dots, N. \tag{10}$$

[25]Setting these priors to focus on breaks at higher frequencies (e.g., once every couple of years) tends to produce noisy regimes whose parameters and inflation dynamics are difficult to interpret economically.

[26]The grouped heterogeneity model specifies a normal prior over the coefficients λ_{g_k} and ρ_{g_k}.

We assume that the average expected lag with which the N series are hit by the kth break, δ_k, has a conjugate Gamma prior distribution

$$p(\delta_k) \sim Ga(e, f) \qquad k = 1, \ldots, K. \tag{11}$$

The hyper parameters e and f again control the average degree of heterogeneity in break dates across series and the lag in individual series' break dates from the beginning of the break window τ_k is allowed to vary across breaks. Some breaks might spread very rapidly across all series, while others may undergo a slower diffusion process.

For the quarterly data, we set $e = 8$ and $f = 1$ such that the prior expected break lag for each series is eight quarters (two years). Similarly, for the annual data we set $f = 1$ and $e = 2$.

3.4.4. Priors on heterogeneity and grouping structure

Our third specification introduces heterogeneity through an endogenous break clustering structure. We accomplish this by placing a Poisson prior over the number of series included in the g_kth group, N_{g_k}[27]

$$p(N_{g_k} \mid \psi_k) \sim Po(\psi_k), \qquad g_k = 1, \ldots, G_k + 1, \qquad k = 1, \ldots, K + 1, \tag{12}$$

where the expected number of series in every group, ψ_k, has a conjugate Gamma prior

$$p(\psi_k) \sim Ga(h, j). \tag{13}$$

The prior hyper parameters h, j control the average expected number of groups in the prior.

To determine group size for this specification, we set $h = 5$ and $j = 1$ to reflect our prior belief that there are, on average, five series in each group. This choice of prior on the groups thus leans towards not having a single series comprise a group.

3.5. Estimation

Each of our models is estimated using a multi-step reversible jump Markov chain Monte Carlo algorithm (Carlin and Chib 1995; Green 1995). Estimation of the baseline model

[27] This specification of multiple independent Poisson distributions is inferentially equivalent to a specification that uses a single Multinomial distribution.

consists of three steps. First, we estimate the regression coefficients from their full conditional distributions using a Gibbs step. Next, we estimate the break locations using a random-walk Metropolis-Hastings algorithm. Finally, the third step estimates the number of breaks using a reversible jump step. This latter step introduces the number of breaks K as a parameter and repeatedly attempts to "jump" to different values of K, with the proportion of iterations spent at each value of K approximating the posterior model probabilities.[28]

Estimation of the second, noncommon breaks, model proceeds in the same manner as for the baseline model, except it includes an additional Metropolis-Hastings step that estimates the exact break location for each series in the cross-section.

Finally, estimation of the third, grouping, model combines the first step of estimating the baseline model with a second reversible jump step that introduces the number of groups G_k as a parameter in the model and repeatedly attempts to 'jump' to different values of G_k, with the proportion of iterations spent at each value of G_k approximating the posterior model probabilities. The series are ordered with the first N_1 series in group 1 and so on. The ordering of the variables, and hence their group allocations, are further estimated using a random walk Metropolis-Hastings algorithm.

4. Empirical Results

Having introduced our data and estimation approach, we next turn to the empirical analysis. We begin with the industry-level data before turning to the MSA and EU country data.

4.1. Industry-level data

We separately analyze two panel data sets on industry-level inflation, namely 16 PCE inflation rates and 31 CPI series.

4.1.1. PCE inflation rates

We first estimate Phillips curves on quarterly sectoral data spanning the sample 1959-2022. Both the number of breaks and their location are very precisely estimated from the data: Our model assigns nearly 100% probability to the presence of two breaks with negligible

[28] For full details on how our three models are estimated, we refer the reader to the articles cited in the first paragraph of Section 3 and only provide a brief discussion here for completeness.

uncertainty as to the timing of these breaks.[29]

Table 2 displays the baseline results for the 16 PCE industry-level inflation rates. The first of the two breaks is a steepening in the Phillips curve around 1972. Prior to 1972, the estimated slope of the Phillips curve is -0.48. This slope estimate steepens notably in the 1972-2001 regime to -0.82.[30] Coupled with an AR coefficient of 0.38, this implies a dynamic slope of -1.33.[31] Inflation volatility, computed as the square root of an industry-weighted average of the individual σ_{ik}^2 estimates, is also notably higher in the 1972-2001 regime (2.91) than in the previous regime (1.58), consistent with major shocks to commodity prices and sharp shifts in inflation expectations accompanying the marked changes to the Federal Reserve's monetary policy during this period.

Our Bayesian panel model identifies a second break in the industry PCE data in 2001. After this break, the slope of the Phillips curve becomes insignificantly different from zero and inflation dynamics become notably less persistent with an AR(1) estimate of 0.16 compared with 0.38 in the regime prior to 2001.

The right-most column in Table 2 shows the equivalent panel estimate based on the full sample 1959-2022, i.e., for a conventional Phillips curve model with no breaks. At -0.42, the estimated full-sample slope shows that ignoring breaks results in a moderately steep Phillips curve. This estimate can be thought of as a weighted average of the slopes in the underlying regimes and so conceals the sharp differences in slope estimates across the more than six decades covered by our sample.

Food and energy prices are known to be more volatile than prices in other ("core") sectors. Focusing on core, rather than headline, inflation is also a simple way to remove some of the cost-push shocks that may affect headline inflation (McLeay and Tenreyro 2020). To examine the price dynamics in core industries, the second panel in Table 2 reports the Phillips curve slopes for a model estimated on all industries excluding food and energy. Excluding food and energy changes the slope during the 1972-2001 period from -0.82 to -0.49 which is notably flatter, but still quite steep. Moreover, this estimate continues to be steeper than that in both

[29] One might be concerned as to whether the identified breaks could be sensitive to the omission of time fixed effects for the sectoral data. While we cannot test this directly, we estimated models with and without time fixed effects for our MSA and EU country data examined below. For both data sets, we found that the number of breaks and their location were not affected by the presence of time fixed effects. While the precision of our estimated break location may partly hinge on the assumption of no cross-sectional dependence in the residuals, we show in our robustness analysis that such dependencies are in fact quite weak in our data.

[30] The high inflationary environment during the 1970s may have increased the speed of Calvo (1983) adjustment, consistent with the steepening of the Phillips curve that we find.

[31] The dynamic slope refers to the long-run effect of a sustained unit change in the unemployment gap.

the first regime (1959-1972), which equals -0.32, and in the last regime (2001-2022) which is -0.13. Excluding food and energy thus flattens the slope of the Phillips curve but the evidence of a steeper unemployment-inflation trade-off in the "middle regime" (1972-2001) continues to be strong.

To help interpret the underlying drivers of these breaks, we also report results separately after pre-assigning the individual price indices into goods and services groups. For both of these groups we obtain a similar pattern in slope coefficients with a steepening in 1972 and a flattening in 2001. However, the shifts in the estimated slopes is much sharper among the goods sectors (third panel in Table 2) as compared to the services sectors (bottom panel). Specifically, for the goods sectors the slope coefficient steepens from -0.59 in the first regime (1959-1972) to -1.09 between 1972 and 2001, only to flatten to a statistically insignificant value of -0.32 after 2001. For the services sectors, the corresponding slope estimates for the three regimes are -0.34, -0.57, and -0.19. In addition, the goods and services slope coefficients are significantly different from one another in the first two regimes (indicated through the bold font of the services slope), but not in the final regime.

Across all data sets examined in Table 2, the full-sample estimates (reported in the right-most column) imply a markedly flatter Phillips curve than the curve implied by the estimates in the first two regimes, 1959-1972 and 1972-2001. The reason for this is that the Phillips curve essentially becomes flat in the last period (2001-2022) which, when pooled with the earlier samples, flattens the curve. Ignoring breaks would therefore lead to the wrong conclusion of a rather flat Phillips curve and conceal the more complex story that, while quite flat during the last twenty years, the Phillips curve has, historically, been quite steep, especially during the nearly three decades 1972-2001. The full-sample estimates also show that ignoring breaks conceals the significant differences between the goods and services slopes that we find prior to 2001.

The disaggregate results in Table 2 assign industries to a set of pre-determined groups. Our unobserved grouped heterogeneity model in Equation (5) instead endogenously assigns industries to groups. Table 3 displays parameter estimates, along with the posterior mode group allocation, from applying this approach to the 16 industry PCE series. Within all three regimes, our approach identifies two groups with very different Phillips curve estimates. The really steep Phillips curve in the 1972-2001 subsample is concentrated in a group ("Group 1") that includes Gasoline and other energy goods along with Financial services and insurance and NPISH.

In fact, the behavior of the slope coefficient for Gasoline and other energy goods is so different from that of the other industries that this is the only sector to be included in Group 1 after 2001 and it is only grouped together with Financial services and insurance and NPISH in the 1972-2001 regime. This narrow, unbalanced grouping only happens when the behavior of a very small number of individual industries is truly different from that of the remaining industries. This point is further highlighted by the extremely high volatility estimate (39.79) for the Gasoline industry in the 2001-2022 regime which is more than twenty times higher than that of the other industries (1.95).

By averaging the group-specific Phillips curve slopes in each regime, using the number of series in each group as weights, we can back out the overall sectoral Phillips curve slope. This approach shows that the slope shifts from -0.41 to -0.79 across the breakpoint in the early-1970s. Comparing these values to the fully pooled estimates, we see that the fully pooled estimates over-estimate the steepness of the Phillips curve in the first two regimes.

Using our baseline panel break model, the black line in the top panel of Figure 1 plots the posterior mean of the Phillips curve slope within each of the three regimes with the blue bands denoting 95 percent posterior intervals. These bands are clearly narrower in the first regime and widest in the last regime after 2001. To illustrate the value of using cross-sectional information to estimate the Phillips curve, the red dotted lines plot industry-level estimates of the slope coefficients estimated separately for the three regimes identified by our panel breakpoint model and reported in Appendix Table A1. Two points stand out. First, consistent with the estimates in Table 3, we see strong evidence of variation in the industry-level Phillips curves both over time and across industries. The majority of industries have a significantly negative slope coefficient on the aggregate unemployment gap in the first regime: 9 of 16 slope estimates are negative and significant (at the ten percent level) for 1959-1972. In contrast, no more than five industries generate a significantly negative slope coefficient in either the final regime (2001-2022) or for the full sample (1959-2022). Second, we see that the individual industry PCE Phillips curves are imprecisely estimated with estimates covering a wide range of values that fall outside the 95% confidence band for our panel estimates. This demonstrates the value of using cross-sectional information to estimate the Phillips curve in a panel setting.

Figure 2 displays the posterior mode break dates for the 16 PCE industries obtained from the generalized version of the baseline model displayed in Equation (3). This model allows the break timing to vary across industries as described in Section 3.2. Industries that are hit first appear further to the left while industries hit later show up on the right in this figure.[32] The top and bottom panels show results for the 1972 and 2001 breaks, respectively.

The earliest industries to be hit by the 1972 break are Financial Services, and Food and Beverage. For these industries, the Phillips curve breaks in 1972Q3. Gasoline and NPISH follow suit in 1972Q4. Eight of the industries are affected by the break in 1973Q3 or 1973Q4, a full year later than the first-hit industries. Overall, the impact of the 1972 break to the Phillips curve took six quarters to percolate through the economy.

The 2001 break hits Gasoline very early (2001Q1) and the remaining industries are all hit in 2002. Thus, as for the first break an energy sector (Gasoline) is hit early, but there is less dispersion in the timing of the break across industries for the 2001 break compared to the 1972 break.

4.1.2. Results for individual inflation series

It is important to emphasize that our ability to detect breaks in the Phillips curve is closely linked to our use of panel data in conjunction with the assumption that both the timing of breaks and their impact on individual series is relatively homogeneous, i.e., there is a strong common component in the breaks.

To highlight this point, we undertook a set of Phillips curve regressions on the individual inflation series using the breakpoint methodology of Chib (1998). We fail to identify a single break in any of the PCE inflation series. Next, we dispensed with the assumption of homogeneous slope coefficients, imposing only that the timing of the break is identical across all variables in the panel. Once again, we fail to find evidence of breaks to the Phillips curves.

These results show that our ability to identify breaks in the Phillips curves hinges on the ability of our panel estimation approach to efficiently exploit multivariate information in a way that takes advantage of the relative homogeneity in both the timing and impact of the breaks across industries. This increases the power of the panel break tests compared with series-by-series approaches or approaches that rely on heterogeneous panels.

Next, we evaluate the ability of the frequentist breakpoint approach of Bai *et al.* (1998) to detect breaks in the Phillips curve in the settings we consider. Since their approach only permits a single break, the most direct comparison is with the data sets for which we identify

[32] The vertical ordering of industries is arbitrary.

just one break, namely, the MSA- and EU country data. Their approach, which assumes heterogeneous slope coefficients, does not detect a break in either data set, echoing what we find when applying our approach with heterogeneous slope coefficients and no pooling across variables. Indeed, the ability of our approach to exploit cross-sectional commonalities in the timing and impact of breaks by (fully or partially) pooling parameters accounts for the additional power our approach has to identify breaks.

4.1.3. CPI industry inflation

We next examine the results for the 31 CPI industry-level quarterly inflation rates (1954-2022). Once again, there is very little uncertainty about the number and timing of breaks and our model identifies two breaks–corresponding to three regimes–with posterior modes in 1971 and 2001.

Table 4 displays the baseline results for the 31 CPI industry-level quarterly inflation rates (1954-2022). The first of the three regimes (1954-1971) has a Phillips curve slope of -0.51. The second regime (1971-2001), has a very steep Phillips curve with an estimated slope of -1.46, while the third regime has a flat Phillips curve with an insignificant slope estimate of -0.27. Autoregressive dynamics are generally quite weak with estimates of 0.09 in the first regime, 0.23 in the middle regime and 0.07 in the last one.

Table 4 also reports results on the model that excludes food and energy prices (second panel). Here, we find a pattern of a Phillips curve that flattens across both breaks in 1971 and 2001. While the slopes of the CPI Phillips curves fitted to core and all prices are similar in the first and third regimes, the core CPI Phillips curve is noticeably flatter than the curve fitted to all prices in the middle regime (-0.49 versus -1.46). The third and fourth panels show inflation estimates generated separately for goods and services. In the two regimes prior to 2001, the Phillips curve is steeper for goods than for services. Conversely, the slope of the Phillips curve is insignificantly different from zero (at the 5 percent level) in the last regime for both goods and services.

Table 5 displays the results from the unobserved grouped heterogeneity model that uses the 31 CPI industries.[33] Our approach identifies a single group in the first regime but two groups in the second and third regimes. In the second regime (1971-2001), there is a group with an especially steep Phillips curve (slope estimate of -1.73 versus -0.26 for the

[33]Some industries do not have inflation data in the early parts of our sample and so cannot be allocated to a group. These show up as missing observations in the first two regimes in the table.

other group) which includes energy and some food components of CPI. In the third regime (2001-2022), two groups are again identified, both with flat Phillips curves and only the first group generates a significant slope estimate. While certain food and energy items are again overrepresented in the second (smaller) group of industries identified for this regime, others are included in the first group and, as a result, the group structure in the third regime is quite different from that in the second.[34]

The overall CPI sectoral Phillips curve slope implied by these grouped estimates is -1.09 in the second regime, considerably flatter than the fully pooled estimate. The fully pooled model therefore overstates the steepening that occurred in the early 1970s.

Appendix Table A2 examines the heterogeneity in further detail by estimating univariate Phillips curve time-series regressions separately for each of the three regimes identified by our panel break method as well as for the full sample. Many of the CPI series are not available in the first two regimes which limits the comparisons across time and industries. Nevertheless, for 18 of the 23 industry CPI series for which we have estimates for both the middle and last regime, the Phillips curve is steeper in the former (1971-2001) than in the latter (2001-2022) period. This again is strong evidence of a flattening of the Phillips curve at the industry level.

Figure 3 displays the posterior mode break dates for the 31 CPI industries based on the model that allows the break timing to vary across industries. Our findings are in line with what we found for the PCE industries: For the 1971 break (top panel), Food prices (Meats, Poultry, Fish and Eggs and Fruits and Vegetables) are the first categories to be affected in 1971:Q3, followed by food items and various energy sectors whose break date is estimated to occur in 1972. For the majority of industries, the break date is 1974:Q3, a full three years after the first sectors are affected, suggesting that it took a long time for this break to percolate throughout the economy.

The 2001 break (bottom panel) initially affects fuel sectors (Motor Fuel and Utility (piped) gas service) in 2001:Q1, followed by Fuel oil and other fuels and various food industries. Once again, the impact plays out over three years with the majority of industries

[34] Our approach allows all parameters to shift across regimes and clusters. Which cluster a particular industry gets assigned to will therefore depend on its persistence, slope, and volatility parameter. For example, in the third regime, group 1 has an AR slope of 0.30 with a t-statistic of 12.51 while Group 2 has a comparatively modest AR slope of 0.08 with a t-statistic of 1.92. Group 1 therefore tends to consist of industries with more persistent inflation dynamics. Similarly, among the food industries, those allocated to group 2 (the high-volatility cluster) have volatility estimates of 6.84 (Meats), 7.83 (Dairy), 7.80 (Fruits), and 25.65 (Food at employee sites). The remaining food sectors have much lower volatility estimates close to 2.5.

experiencing the break only in 2004:Q1.

4.2. MSA-level data

The top panel of Table 6 displays the baseline results for the 22 annual MSA-level inflation rates (1980-2022).[35] We identify a single break in 2000, with a marginal flattening of the Phillips curve which goes from a pre-break slope estimate of -0.26 to a post-break estimate of -0.21, with both being highly significant. The slope, scaled by the expenditure share on non-tradeables to give an implied aggregate Phillips curve slope (as discussed in subsection 3.1.1) flattens from -0.39 before the break to -0.31 after it. The persistence of the inflation process, measured through the autoregressive parameter, increases significantly from 0.23 before the break to 0.39 afterwards.

The MSA data suggests a much flatter slope of the Phillips curve in the pre-2000 period than that identified with either PCE or CPI sectoral data. There are a number of reasons for this. First, the sectoral Phillips curve is particularly steep in the period after the early seventies and the MSA data only starts in 1980. Consistent with this, sectoral Phillips curves are flatter if estimated only on data starting in 1980. Second, the MSA data are observed only at the annual frequency whereas sectoral data are quarterly, further attenuating the impact of the periods that experienced the steepest inflation-unemployment dynamics. Third, the MSA-level results apply to the slope of the regional Phillips curve and the implied national Phillips curve is steeper as we noted earlier.

To examine a possible source of breaks to the Phillips curve, the middle panel in Table 6 displays results when, conditional on the regimes identified by the baseline model, we estimate the Phillips curve regression separately for MSAs located in states with below and above median rates of import penetration from China based on the state-level import penetration rates calculated by Riker (2022).[36] We find that the flattening of the Phillips curve is concentrated in cities with above-median rates of import penetration. Specifically, whereas the slope of the Phillips curve changes only marginally from -0.16 to -0.15 for MSAs with below-median import penetration from China, it declines from -0.41 to -0.24 in cities with above-median import penetration from China. This finding lends credence to the role

[35] The number of breaks and break dates are, once again, very precisely estimated.

[36] Riker (2022) estimates these values using a structural econometric model that exploits data on the location of import entry, domestic shipments, and distances between states. The MSAs that comprise the below median group are Detroit-Warren-Dearborn, MI, Dallas-Fort Worth-Arlington, TX, Denver-Aurora-Lakewood, CO, Philadelphia-Camden-Wilmington, PA-NJ-DE-MD, and St Louis, MO-IL.

of international trade as an explanation for the flattening of the Phillips curve. Moreover, the slope coefficients for the two groups are significantly different from one another in both regimes and in the full-sample results.[37]

The overall regional Phillips curve slope implied by these grouped estimates are -0.35 in the first regime and -0.22 in the post-2000 regime. The fully pooled model therefore under-estimates how much the Phillips curve has flattened.

In Appendix Table A3 we report the results from a series of MSA-level Phillips curve regressions on the two regimes identified by our benchmark model, i.e., 1980-2000 and 2001-2022, as well as for the full sample, 1980-2022. Importantly, when we conduct the break-point estimation for individual MSAs, we fail to find significant evidence of breaks for any of the series. This reflects the weak power of break tests conducted on individual (univariate) time series which fail to exploit information in the cross-section to identify breaks. However, we can still use the breaks identified by our panel model to examine evidence of time-variation in Phillips curve slope estimates across time and cities. Only 7 MSAs have a significant negative Phillips curve over the full sample, compared with six in each of the two regimes, suggesting that Phillips curves are poorly identified using inflation series at the individual MSA-level. Possible explanations for this include non-stationarities in the data and the relatively small samples (at most 43 observations) available at the annual frequency.

The top panel of Figure 4 displays the Phillips curve slope coefficient estimated on MSA data throughout our sample using the model specification displayed in Equation (1), but estimating it as a panel no-break OLS regression with two-way fixed effects using a ten-year rolling window exponentially down-weighting older data according to a decay parameter equal to 0.8. We see that the Phillips curve reached its steepest point for windows ending in the early-2000s and flattened thereafter. We find a similar pattern for the EU data (bottom panel). This suggests that a shift at 1990 as other studies have implemented might be too early, and that the shift occurs a little later as our breakpoint model suggests. Note that the volatility in the rolling window estimates with a steepening of the Phillips curve followed by a flattening is caused by the fixed window length, leading to "base" effects as observations drop out of the window as it is recursively rolled forward. This clearly illustrates why a formal breakpoint approach like our method is required to accurately estimate time-variation in the Phillips curve. Finally, the rapid steepening of the Phillips curve suggests that a breakpoint

[37]Many MSAs have missing data from the 1980s. This makes it difficult for our endogenous clustering approach to identify groups of cities with distinctly different Phillips curve dynamics, so we do not apply our grouping approach to the MSA-level data.

approach may better capture variation in the data than a time-varying parameter model with slowly-evolving coefficients, aside from also being easier to interpret. We do a formal comparison of time-varying parameters and discrete break models in subsection 6.3 below, and find that discrete break models fit the data better.

Other studies have found evidence of a flattening Phillips curve. Using a simple subsample split on state-level data, Hazell *et al.* (2022) find that the Phillips curve flattens by about 50% post-1990, but that it was already quite flat and the change is not statistically significant. We estimate our breakpoint approach using their data and corroborate their findings with a single break, except the break date is estimated to occur in 2000, aligned with the date estimated from our other data sets.[38] In both regimes, we observe a flat Phillips curve with the slope estimate being marginally insignificant in the first regime (1978-2000) and clearly insignificant in the second (2001-2017). Note that the data of Hazell *et al.* (2022) excludes shelter and so it is not surprising that the slope is much flatter than when we estimate it from our MSA-level data since shelter is one of the most cyclically sensitive categories (Stock and Watson 2020) and comprises about 30 percent of CPI.

4.3. Nonlinear Phillips Curve

Tests for breaks are conducted in the context of, and conditional upon, the maintained model specification, in our case a linear Phillips curve model. It is possible that our findings on the presence of breaks to this model reflect omitted non-linearities in the inflation-unemployment trade-off. Previous studies such as Babb and Detmeister (2017) have in fact identified non-linearities in the Phillips curve.

To examine this possibility, while still allowing for the possibility of breaks, we generalize the linear Phillips curve specification in Equation (1) to allow the unemployment-inflation trade-off to have a kink at a pre-specified threshold, θ, so that for regimes $k = 1, \ldots, K+1$ and $t = \tau_{k-1}+1, \ldots, \tau_k$,

$$\pi_{it} = \alpha_i + \gamma_t + \rho_k \pi_{it-1} + \lambda_k URATE_{it} + \omega_k (URATE_{it} - \theta) 1_{URATE_{it} < \theta} + \epsilon_{it}. \qquad (14)$$

Note that there is no discontinuity at the threshold point (θ) but the degree of steepening in

[38] Specifically, using the baseline model in Equation (1) on the state-level core CPI data from Hazell *et al.* (2022), we regress the 34 state-level nontradeables inflation rates from 1978 through 2017 on the lagged state-level unemployment rates and an autoregressive term while allowing for two-way fixed effects. The four-quarter inflation rate in each year is regressed on the four-quarter inflation rate in the previous year and the average monthly unemployment rate computed over the 12 months of the previous year.

the Phillips curve can differ across regimes. This allows us to examine if nonlinearities were more or less important in regimes with a steep or flat Phillips curve.

Because we only have aggregate measures of slack for the sectoral (PCE and CPI) data, we cannot estimate the model in Equation (14) on these data. Conversely, the MSA-level data has city-level unemployment rate data and so can be used to examine nonlinear (threshold) effects.

We exploit regional variation in labor market tightness to identify nonlinearities in the Phillips curve, which is crucial since the US national labor market has experienced relatively few episodes of being very tight. For example, the black circles in the top panel of Figure 5 plot the annual aggregate headline CPI inflation rate against the lagged annual unemployment rate during our sample period from 1980 through 2022. There are few observations corresponding to very tight labor markets. On the other hand, there are many observations in which labor markets have been very tight in the MSA-level data (green circles).

The lower panel of Table 6 displays results from estimating Equation (14) with threshold values for the unemployment rate (θ) of 5% (top panel) and 4.2% (second panel), respectively. These are the thresholds considered by Babb and Detmeister (2017).[39] In the post-2000 regime, at an unemployment rate below 4.2 percent, the slope is -0.44 versus -0.20 at higher unemployment rates. The estimated size of the post-2000 kink is notably bigger than the kink estimated without allowing for any regime change (-0.24 versus -0.17). This is consistent with a time-varying non-linearity and shows that, at least for the MSA data, the Phillips curve has become notably steeper after 2000 at low levels of unemployment.

Table A4 in the Appendix, shows that the nonlinearity not only interacts with regimes, but also with clusters of MSAs. In the post-2000 regime, we document that in a hot labor market the group of MSAs with below median rates of import penetration has an additional Phillips curve slope that is approximately three times as steep and significantly different from that of the group of remaining MSAs. This finding is clearly relevant to current policy debates about costs and benefits of a hot labor market.[40]

[39] We condition on the two regimes identified by the Phillips curve fitted to the MSA-level inflation data, but the break dates remain the same if we estimate the model augmented to allow for the kink.

[40] There is mounting evidence that households experience different inflation rates (Kaplan and Schulhofer-Wohl 2017; Jaravel 2019; Cavallo 2020; Argente and Lee 2021). Meanwhile, the Federal Reserve and the European Central Bank recently committed to "inclusive" monetary policy, emphasizing lower income households (Powell *et al.* 2020; Schnabel 2021). If these MSAs groupings have different representations of lower income households, then the central bank may wish to put more weight on the trade-off from running a hot labor market in those MSAs.

4.4. Wage Phillips curves

The top panel of Table 7 displays results for the wage Phillips curve when using the 51 state-level (including the District of Columbia) quarterly wage inflation rates from 1980 through 2019. We identify a single break in 2000Q2. This break results in a flattening of the wage Phillips curve with the estimated slope falling from -0.47 to -0.38 which is less dramatic than for some of the price Phillips curves as noted by Rognlie (2019) and Hooper *et al.* (2020), perhaps because in the price Phillips curves markups can absorb some of the effects of wage changes, especially with considerable global competition in product markets.

Table 7 also considers the same kind of non-linearity in the wage Phillips curve that we earlier examined for prices at the MSA level in Table 6, i.e., using thresholds of 5% (middle panel) and 4.2% (bottom panel). For both threshold values, we find very strong evidence of a much steeper wage Phillips curve in tight labor markets. For example, the slope of the Phillips curve after 2000 is -0.36 when the unemployment rate exceeds 4.2% but, at -0.98, is nearly three times steeper when unemployment falls below this level. Moreover, for this data, the steep threshold effect holds in both the early and late regimes and is even slightly stronger in the earlier data.

4.5. EU Data

Table 8 displays the baseline results for the 28 EU country-level annual inflation rates (1986-2021). Our model identifies a single break in 2004. Before the break, the slope of the Phillips curve is -0.78 with an AR coefficient of 0.10, implying a dynamic slope coefficient of -0.87. After the break, the slope coefficient declines to -0.14 with an AR coefficient of 0.51, implying a dynamic slope coefficient of -0.29. Both estimates are significant, so the inflation-unemployment trade-off remains valid, but the Phillips curve becomes much flatter after the break. Scaling by the expenditure share on nontradeables, the implied euro-area Phillips curve slope flattens from -1.11 to -0.20. Grouping the countries into "rich" and "poor" nations, based on whether GDP per capita deflated by PPP is above or below average, we find that the poor countries had a significantly steeper Phillips curve in the first subsample (1986-2003), whereas the slopes of the Phillips curve are the same, and flat, for the two groups in the second subsample (2004-2021). This finding is consistent with the greater goods and labor market integration of countries in Southern and Eastern Europe seen in recent years. The overall Phillips curve slope implied by these rich versus poor grouped estimates are -0.48 in the first regime and -0.09 in the second regime. The fully pooled model therefore over-estimates how much the Phillips curve has flattened, primarily because it over-estimates the steepness of the Phillips curve in the first regime.

To further track how the heterogeneity in Phillips curve slopes evolves over time, the black line in the lower panel of Figure 1 graphs the evolution of the posterior mean of the EU Phillips curve slope over time with blue bands denoting the 95 percent posterior interval and red lines tracking the Phillips curves estimated country-by-country. Uncertainty about the panel estimate of the Phillips curve slope is much stronger in the first regime and considerably smaller after 2001. As for the PCE inflation series, we see that the individual country-level Philips curves are imprecisely estimated and often fall outside the 95% posterior interval.

The lower panel of Table 8 displays results for the full sample, as well as separately in the two regimes identified by the baseline model, for a Phillips curve that uses either country-level total goods inflation, or total services inflation as the dependent variable. The flattening of the Phillips curve after 2004 is apparent for both goods and services inflation. Interestingly, while flatter in absolute terms, the Phillips curve remains significantly steeper for services inflation in the second regime (slope estimate of -0.22 versus -0.09), consistent with what we found for the US. For the full sample, the Phillips curve estimated on services inflation is significantly steeper than the curve estimated on goods inflation (slope estimates of -0.24 versus -0.11).

Table 9 uses the group heterogeneity model that endogenously determines if there are differences in individual countries' Phillips curves and how these are affected by breaks. We identify two groups of countries in the first regime that ends in 2003.[41] One group (labeled group 1 in the table) has a relatively steep Phillips curve before 2004 with a slope of -0.42. This cluster mainly comprises countries on the European periphery, including Bulgaria, Estonia, Ireland and Portugal. The second group has a much flatter Phillips curve with an estimated slope of -0.08 which fails to be significant. In the post-2004 regime, we identify a single group, whose estimated slope is -0.14.

To further understand this heterogeneity, Appendix Table A5 reports country-level estimates of the Phillips curve coefficients estimated separately for the two regimes (1986-2003, 2004-2021) and for the full sample (1986-2021). Only six countries (the Netherlands, Finland, Lithuania, Latvia, Austria, and Cyprus) generate a significantly negative estimate over the full sample (1986-2021), versus five countries in the early period and seven countries for the

[41] Inflation in Romania contains extreme outliers during the post-Communist transition, so this country is in a group of its own during the early sample. We simply mark it as missing in the table.

regime that starts in 2004. This demonstrates two important points. First, Phillips curves are poorly identified using inflation series at the individual country-level. Second, break tests conducted at the univarate level tend to have weak power. As for the U.S. data, break tests conducted at the individual country level based on the break estimation methodology proposed by Chib (1998) fail to find significant evidence of breaks for any of the countries. This point is linked to the large estimation errors associated with the country-level Phillips curves and shows up in the form of quite large variation in coefficient estimates across the two regimes for individual countries.[42]

In summary, our estimates for the EU Phillips curve suggest that in addition to a flattening of the Phillips curves, there has been some "convergence" in Phillips curve slopes with the flattening being most pronounced in countries that previously had steep Phillips curves. European integration thus appears to have been associated with a convergence of the slopes of country-level Phillips curves, consistent with what we found for "rich" and "poor" countries in Table 8 above.

To examine possible nonlinearities in the EU Phillips curve, the results displayed in the final panel of Table 8 allow for a single kink at an unemployment gap threshold below -1.5% as in Equation (14).

Under normal labor market conditions, we find a significantly negative and very steep Phillips curve in the first regime (1986-2003) with a slope estimate of -1.15. Conversely, the Phillips curve in a tight labor market ($UGAP < -1.5\%$) is poorly identified in this subsample, likely because Europe had so few cases with very tight labor markets in this time period. Turning to the second regime (2004-2021), there is a significant steepening of the Phillips curve which goes from -0.07 (flat) to -0.63 (steep) in tight labor markets. The full sample kink (-0.17) is insignificant, underscoring again the insights from considering nonlinearity and structural stability jointly.

These findings are consistent with the US findings and support the presence of a Phillips curve trade-off over the last two decades but only in tight labor markets.

The top panel of Figure 6 displays the posterior mode break dates for the 28 EU countries based on the generalized version of the baseline model in Equation (2) that allows the break timing to vary across countries as described in Section 3.2. As in the earlier figures, circles to the left (right) indicate countries that are affected first (last) by a break. The vertical array of circles to the far left of the figure comprises all the advanced, early EU members which,

[42]We also fail to identify breaks in a panel break model with heterogeneous slope coefficients. Hence, it is exploiting cross-sectional information *and* pooling parameters that generates break detection power.

thus, are affected first by the break to the Phillips curve in 2003. Countries such as Estonia, Poland, and Latvia follow in 2004. With a five-year delay, Romania and Bulgaria are the last countries to exhibit flattening of their Phillips curves. The Bayesian algorithm has no knowledge of the timing of EU accession, but it is noteworthy that these two countries were the last to join the EU, in 2007. This points to EU membership, and the associated trade linkages and freedom of movement of labor, as possible factors associated with the observed flattening and convergence of the Phillips curves.

The lower panel displays how the corresponding Phillips curve slope coefficients for two countries – Germany (red line) and Estonia (black line) – evolve over time, incorporating uncertainty surrounding the break dates estimated from the noncommon breaks model. We see that the slope for Germany adjusts quite rapidly, while the slope for Estonia evolves much more gradually. An attractive feature of our Bayesian noncommon break methodology is that it not only allows the break dates to differ across series, but it also captures the uncertainty surrounding the estimated break date for each series. The latter feature allows series to differ in how quickly their parameters evolve.

5. Accounting for recent inflation experience

In this section we explore implications of our findings for the recent inflation experience, including evidence of missing disinflation and the recent inflationary surge.

5.1. Missing disinflation, reinflation and the recent inflationary surge

The top panel of Figure 7 displays the Phillips curve fit from our linear MSA breakpoint model (black line). Specifically, in each year this is our prevailing regime-specific MSA regional linear Phillips slope coefficient divided by the nontradeables share and multiplied by the lagged national unemployment rate gap. The red line graphs the annual national headline CPI inflation rate minus long term inflation expectations, which are 10-year ahead SPF CPI inflation expectations.[43]

If there were missing disinflation during the Great Recession we would expect to see the black line run below the red line. Likewise, we would expect the black line to run above the red line if there were a missing reinflation during the recovery years following the Great Recession. In fact, the black line tracks the red line closely and so there appears to be little

[43]Missing observations prior to 1991 Q4 are filled using linear interpolation.

evidence of missing disinflation and missing reinflation according to the fit of our Phillips curve model. This echoes the results of Ball and Mazumder (2019) and Hazell *et al.* (2022).

The dotted black line graphs the implied estimates from our nonlinear Phillips curve estimates in the second regime using MSA-level data and increasing the noncyclical rate of unemployment (NROU) in 2021 and 2022 to the estimates from Crump *et al.* (2022) who argue that NROU has risen temporarily since COVID-19.[44] The dotted black line increases just under half as much as the red line between 2020 and 2022, suggesting that a steeper Phillips curve in hot labor markets combined with a higher NROU can explain a bit less than half of the recent inflationary surge. If we were to instead use a linear Phllips curve, shown by the solid black line, we would explain almost none of the pandemic era inflation surge.

We repeat this exercise for the EU and display the results in the lower panel of Figure 7. The black line uses estimates from our linear breakpoint Phillips curve model. The red line uses the EU inflation rate and long term (five-year ahead) Eurozone inflation expectations from the ECB SPF which goes back to 2002 Q3. Prior to this, we use one-year ahead expectations, going back to 1999 Q1.[45] We average expectations across the four quarters in a given year. Once again, we see little evidence of missing disinflation during the Great Recession or missing reinflation during the subsequent recovery.

5.2. Break monitoring and the prior for regime duration

As we will discuss more fully in subsection 6.4 below, in most respects our results are not sensitive to the priors. However, an exception to this is the prior expected regime duration. So far we assumed a prior expected regime duration of 20 years which is suitable for a longer-term historical analysis aimed at detecting "secular" breaks. This choice of prior does, however, make it hard to detect any breaks near the end of the sample due to insufficient observations in the new regime (Bai and Perron 1998). This matters because monetary policy-makers and financial market participants must make decisions in real-time and so are more focused on monitoring for recent breaks.

Our methodology is flexible enough to also be used to detect breaks emerging towards the end of the sample.[46] This is particularly salient in light of the post-pandemic inflationary

[44]Specifically, we use their 5.9 percent estimate in 2021 and a value of 5.6 percent in 2022 which is about the middle of their range of forecasts.

[45]Eurozone expectations data are sourced from the ECB statistical data warehouse.

[46]Previous studies have focused on detecting recent structural changes due to, for example, the implosion

surge and the possibility that the Phillips curve has steepened in the COVID era, as documented in other studies (Cerrato and Gitti 2022; Inoue *et al.* 2022). Unfortunately, these studies either impose ad hoc subsample splits or use time-varying parameter models which allow for drifting coefficients but do not have a formal statistical procedure to determine whether changes in coefficients truly reflect a shift to a new mean. Conversely, our panel data approach leverages cross-sectional information to deal with the weaker power of tests targeting break detection towards the end of the sample.

The simplest way we can modify our baseline model to detect a possible break towards the end of our sample is by reducing our prior expected regime duration. Specifically, we reduce this from 20 to five years.

Table 10 shows the resulting estimates based on our sectoral CPI data. In all, five regimes are found, with the most recent break date in 2020:Q1. In this post-pandemic regime, the Phillips curve slope coefficient is estimated as a significant -1.6 after having been essentially flat for two decades. Excluding the volatile food and energy sectors which, as noted earlier, can help to remove some of the supply shocks that undoubtedly played a role during the pandemic era, we still uncover a steepening of the Phillips curve to a significant value of -1.16 in the COVID era. Splitting the sectors into goods and services, we find that the steepening of the Phillips curve is concentrated in the goods sector, with the services slope being insignificant. The steepening of the goods Phillips curve may have been driven by the rotation of consumption from services towards goods that accompanied lockdowns and the persistence of work from home policies. Of course, supply side disruptions such as the chip shortage that caused motor vehicle prices to skyrocket may be playing a role. Still, removing the transportation commodities less motor fuel category does not materially change our results.

Turning to the sectoral PCE data, our approach with the five-year regime duration prior identifies the same terminal break date as the CPI sectoral data, namely 2020:Q1. In Table 11, we report the estimated slope of the Phillips curve in this post 2020:Q1 period only for PCE sectoral data, with the corresponding results for the CPI data repeated as a memo item. For the PCE data, there is again a considerable steepening. And again this holds when we exclude food and energy sectors, and is concentrated in goods sectors. Excluding the motor vehicle and parts category has a greater impact here: the slope for all sectors excluding food and energy is a significant -1.45, quite a bit flatter but still a marked steepening relative

of financial bubbles or COVID-19 (Phillips and Shi 2018; Bardwell *et al.* 2019; Otto and Breitung 2022).

to the Phillips curve slope seen over the two decades that preceded the pandemic. Taken together, these estimates imply a considerable steepening of the Phillips curve in the US in the COVID era.

Turning to the EU data, since we are using annual data even our algorithm cannot detect a break with so few observations in the post-pandemic regime. We therefore impose the break date identified from our US Phillips curves, and find that the EU Phillips curve has steepened to a significant -0.51, as also shown in Table 11. Once again, this steepening is concentrated in goods rather than services sectors. We also find that the steepening is concentrated in poorer countries (as also found by Baba *et al.* (2023)), and these are the countries that experienced most of the flattening in the early-2000s. The strains from Covid on supply chains may have, at least for a while, reversed some of the impacts of European integration. Finally, when estimating our nonlinear Phillips curve on EU data in the final regime we find that the slope is considerably steeper in hot labor markets.[47] This nonlinear Phillips curve is similar in magnitude to the slope estimated for the US using the CPI sectoral data. This implies not only that the EU has moved onto the steeper part of the Phillips curve, but also that the Phillips curve in hot labor markets has got steeper.

An advantage of having a steep Phillips curve is that the increase in slack to bring inflation back to target does not need to be that big. A concern for policy makers going forward is that the Phillips curve could experience another break and revert to a flatter Phillips curve with lower persistence. In this scenario, the COVID era would have represented a short-lived regime related to the unusual shock that the pandemic had on the macroeconomy. Unfortunately, this would mean that inflation would prove harder to dislodge.

6. Robustness checks

In this section, we perform a number of robustness checks on our results. Specifically, we first consider the possible effect on our panel break estimates of serial correlation or cross-sectional error dependence in the residuals. Next, we evaluate whether our panel break model better fits the data than a time-varying parameter model with smoothly-evolving coefficients. Finally, we consider the effect of adopting alternative specifications for the priors.

[47] We are considering the nonlinear specification only for EU data because we need regional slack measures for strong identification of the nonlinear model. Recall that our state-level wage data only go through 2019 and so are not suited to analysis of the COVID period.

6.1. Serial correlation

Serial correlation in the residuals of our model could potentially result in misleading inference. Across all four data sets, the top panel of Appendix Table A6 shows that the p value of the Durbin and Watson (1950) test statistic fails to reject the null hypothesis of no serial correlation in the Phillips curve residuals within every regime across the four data sets (CPI and PCE sectoral, MSA-level, and EU country-level) we consider in our analysis.

6.2. Cross-sectional error dependence

Next, we consider the possibility of cross-sectional error dependence in the residuals from our model. In applications with reasonably large cross-sections, weak dependence or dependence that is confined to a relatively small number of series will not pose serious estimation and inferential problems and only pervasive cross-section dependence is problematic (Pesaran 2015). Moreover, if cross-sectional dependence is caused by unobserved common factors that are uncorrelated with the regressors, our estimator remains consistent, though some of the efficiency gains from pooling may be lost and the standard error estimates may be biased (Phillips and Sul 2003; Chudik and Pesaran 2013).

We test for cross-sectional error dependence using the test statistic proposed by Juodis and Reese (2022) which is a bias-corrected version of the original CD test statistic proposed by Pesaran (2021). Results are reported in the lower panel of Appendix Table A6. We cannot reject the null hypothesis of no cross-sectional error dependence in any regime across the four data sets, although we are on the borderline of rejecting the null in the third regime for CPI data. If we exclude the Motor Fuel category in this regime when computing the test statistic, however, we cannot reject the null, suggesting that the cross-sectional error dependence is not pervasive. Reassuringly, the estimates from the CPI sectoral data in the third regime follow the same basic pattern as the PCE sectoral data (which has no cross-sectional error dependence), namely a flattening curve in the final regime. We therefore conclude that any cross-sectional error dependence is insufficiently pervasive to cause serious inferential problems in our settings.

6.3. Time-varying parameters

Finally, we compute Bayes factors for the baseline panel model with discrete breaks versus the same model estimated using a time-varying parameter (TVP) specification. Bayes factors are constructed using the marginal likelihood of each model computed using the methodology of Chib (1995), for our four price Phillips curve data sets (at the PCE and CPI sectoral-level, the MSA-level, and the EU country-level). Bayes factor values between 1 and 3 are inconclusive, values between 3 and 20 indicate positive evidence in favor of the restricted model, while values between 20 and 150 indicate strong evidence in support of the restricted model (Kass and Raftery 1995). The TVP model can be viewed through the lens of our breakpoint model, but imposes that a (typically small) break occurs every period. We do not impose this assumption. Instead we estimate the number of breaks, specifying a prior on the regime duration that places relatively little weight on very short regimes and so our framework tends to reveal few (typically large) breaks.

Across all four data sets we find Bayes factors above 20, suggesting strong evidence in favor of modeling time variation in the Phillips curve as discrete breakpoints rather than smoothly-evolving changes.

6.4. Alternative prior specifications

Our analysis uses fairly uninformative priors on the key parameters of the Phillips curve such as ρ_k and λ_k, both of which are centered on zero. Effectively, this stacks the results against finding a steep Phillips curve, but we mitigate such effects by allowing for relatively large values of the prior variances σ^2_ρ and σ^2_λ. Because our priors are relatively uninformative, changing the centering of either ρ_k or λ_k has little impact on our results.

Priors can also be used to incorporate economic beliefs into the model. For example, truncating the prior on the slope of the unemployment rate at zero can be used to rule out positive values for the Phillips curve slope coefficient. Empirically, we find that truncating the prior has little impact on our baseline estimates of the price Phillips curve across all four data sets. Specifically, the truncation never binds for the MSA-level data and only binds for the CPI and PCE sectoral data sets in the final regime, causing their slope coefficients to become slightly steeper but not affecting their significance. The truncation binds in each regime for the EU data but only on a relatively small number of posterior draws, causing the magnitude of the Phillips curve to steepen slightly without altering our conclusions in any way. Overall, truncating the Phillips curve slope coefficient at zero has little impact on inference in our study.

To evaluate the impact of the priors on our posterior estimates of the Phillips curve slope, persistence, and volatility parameters, we next compute the formal prior informativeness (PI) measure of Müller (2012). The values are displayed in Table A7 for our PCE, CPI, MSA, and EU data sets. Across all four data sets, the value of this PI measure implies that the priors contribute less than 15 percent to the posterior estimates of these three parameters. We conclude that our choice of priors over the regression coefficients is sufficiently uninformative, contributing relatively little to the posterior estimates.

Our prior on the break frequency is more informative. We select it so a break is expected to occur every 20 years. This means that our results tend to select relatively rare shifts in the Phillips curve which are likely to be of a more secular nature, representing "trend breaks". Lengthening this prior regime duration to 40 years does not affect the results, but shortening it by enough uncovers additional breaks as stated in the earlier post-pandemic analysis. Our methodology therefore is flexible enough to focus on "trend" breaks or be used for break monitoring purposes.

7. Conclusions

In this paper, we have applied new Bayesian panel methods with breakpoints to panel data on inflation and unemployment from the U.S. and the European Union. Our approach brings us three key insights.

Firstly, we consider a "partial pooling" approach that endogenously forms groups or clusters of inflation series, allowing the Phillips curves to differ across clusters (but assuming homogeneity within clusters). This approach is more flexible than conventional panel data methods and yet more efficient than estimating separate time series regressions for each region or industry. Our results provide motivation for future research on partial pooling.

Secondly, we allow for breaks in the slope of the Phillips curve and can estimate the number of breaks, their location, as well as the magnitude of the shift in both slope and volatility parameters. These breaks cannot be estimated with meaningful precision with aggregate data or individual inflation series alone. We find evidence for up to two breaks; one in the early 1970s and the other around 2000. The Phillips curve steepened after the first break, and flattened after the second. The flattening around 2000 is greater for goods than for services, is greater for MSAs with above-median rates of imports from China than for MSAs with below-median rates, is greater in price Phillips curves than in wage Phillips curves, and is greater in poorer EU countries than in richer ones. We identify a distinct pattern of convergence in EU country Phillips curve slope coefficients, consistent with greater geographic mobility. We also consider an extension that allows breaks to affect individual inflation series at different points in time.

Thirdly, we consider a nonlinearity in the Phillips curve whereby hot labor markets are associated with a steeper Phillips curve and find that the evidence for this nonlinearity at the end of the sample is strengthened by allowing for regime breaks.

Our estimates imply essentially no missing disinflation during the Great Recession and no missing reinflation during the subsequent recovery. Combining nonlinearity in the Phillips curve with a higher estimate of the natural rate, we can explain about half of the recent surge in U.S. inflation.

References

ANDREWS, D. W. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica*, **61** (4), 821–856.

ARGENTE, D. and LEE, M. (2021). Cost of living inequality during the great recession. *Journal of the European Economic Association*, **19** (2), 913–952.

AUER, R. A., BORIO, C. B. and FILARDO, A. (2017). The globalisation of inflation: The growing importance of global value chains, Bank for International Settlements Working Paper 602.

BABA, C., DUVAL, R., LAN, T. and TOPALOVA, P. (2023). The 2020-2022 inflation surge across europe: A phillips-curve-based dissection, iMF Working Paper 23/30.

BABB, N. and DETMEISTER, A. K. (2017). Nonlinearities in the Phillips curve for the United States: Evidence using metropolitan data, FEDS Working Paper.

BAI, J., LUMSDAINE, R. L. and STOCK, J. H. (1998). Testing for and dating common breaks in multivariate time series. *The Review of Economic Studies*, **65** (3), 395–432.

— and PERRON, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, **66** (1), 47–78.

BALL, L., MANKIW, N. G. and ROMER, D. (1988). The New Keynesian economics and the output-inflation trade-off. *Brookings Papers on Economic Activity*, (1), 1–65.

— and MAZUMDER, S. (2019). A Phillips curve with anchored expectations and short-term unemployment. *Journal of Money, Credit and Banking*, **51** (1), 111–137.

BALL, L. M. and MAZUMDER, S. (2011). Inflation dynamics and the Great Recession. *Brookings Papers on Economic Activity*, **42**, 337–405.

BARDWELL, L., FEARNHEAD, P., ECKLEY, I. A., SMITH, S. and SPOTT, M. (2019). Most recent changepoint detection in panel data. *Technometrics*, **61** (1), 88–98.

BARNICHON, R. and MESTERS, G. (2021). The Phillips multiplier. *Journal of Monetary Economics*, **117**, 689–705.

BENA, J. and SIMINTZI, E. (2022). Machines could not compete with Chinese labor: Evidence from U.S. firms' innovation, Available at SSRN 2613248.

Blanchard, O. (2016). The Phillips curve: Back to the'60s? *American Economic Review*, **106** (5), 31–34.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, **12** (3), 383–398.

Canova, F. (forthcoming). Should we trust cross sectional multiplier estimates? *Journal of Applied Econometrics*.

Carlin, B. P. and Chib, S. (1995). Bayesian model choice via Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 473–484.

Cavallo, A. (2020). *Inflation with Covid consumption baskets*. Tech. rep., National Bureau of Economic Research.

Cerrato, A. and Gitti, G. (2022). Inflation Since COVID: Demand or Supply, Available at SSRN 4193594.

Chib, S. (1995). Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association*, **90** (432), 1313–1321.

— (1998). Estimation and comparison of multiple change-point models. *Journal of Econometrics*, **86** (2), 221–241.

Chudik, A. and Pesaran, M. H. (2013). Large panel data models with cross-sectional dependence: A survey, CESifo Working Paper Series.

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.

—, — and Koustas, D. (2013). Amerisclerosis? The puzzle of rising us unemployment persistence. *Brookings Papers on Economic Activity*, **Fall**, 193–241.

Crump, R. K., Eusepi, S., Giannoni, M. and Şahin, A. (2022). The unemployment-inflation trade-off revisited: The Phillips curve in COVID times, NBER Working Paper.

Del Negro, M., Lenza, M., Primiceri, G. E. and Tambalotti, A. (2020). What's up with the Phillips curve? *Brookings Papers on Economic Activity*, **1**, 301–357.

Durbin, J. and Watson, G. S. (1950). Testing for serial correlation in least squares regression: I. *Biometrika*, **37** (3/4), 409–428.

Firat, M. (2020). The China shock, market concentration and the U.S. Phillips curve, Working Paper.

Fitzgerald, T. J., Jones, C., Kulish, M., Nicolini, J. P. *et al.* (forthcoming). Is there a stable relationship between unemployment and future inflation? *AEJ: Macro*.

Galí, J. and Gambetti, L. (2019). Has the U.S. wage Phillips curve flattened? A semi-structural exploration, National Bureau of Economic Research.

Gilchrist, S. and Zakrajšek, E. (2019). Trade exposure and the evolution of inflation dynamics, Finance and Economics Discussion Series 2019-7.

Green, P. J. (1995). Reversible jump Markov Chain Monte Carlo computation and Bayesian model determination. *Biometrika*, **82** (4), 711–732.

Hall, R. E. (2023). A major shock makes prices more flexible and may result in a burst of inflation or deflation, NBER Working Paper.

Hazell, J., Herreno, J., Nakamura, E. and Steinsson, J. (2022). The slope of the Phillips Curve: Evidence from U.S. states. *Quarterly Journal of Economics*, **137** (3), 1299–1344.

Hooper, P., Mishkin, F. S. and Sufi, A. (2020). Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating? *Research in Economics*, **74** (1), 26–62.

Inoue, A., Rossi, B. and Wang, Y. (2022). Has the Phillips curve flattened and why?, unpublished manuscript.

Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the U.S. retail sector. *The Quarterly Journal of Economics*, **134** (2), 715–783.

Jones, C., Kulish, M. and Nicolini, J. P. (2021). Priors and the slope of the phillips curve, Unpublished working paper.

Juodis, A. and Reese, S. (2022). The incidental parameters problem in testing for remaining cross-section correlation. *Journal of Business & Economic Statistics*, **40** (3), 1191–1203.

Kaplan, G. and Schulhofer-Wohl, S. (2017). Inflation at the household level. *Journal of Monetary Economics*, **91**, 19–38.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, **90** (430), 773–795.

Koop, G. and Potter, S. M. (2007). Estimation and forecasting in models with multiple breaks. *The Review of Economic Studies*, **74** (3), 763–789.

Lagarde, C. (2023). Policymaking in an age of shifts and breaks. In *Proceedings-Economic Policy Symposium-Jackson Hole*, Federal Reserve Bank of Kansas City.

Leduc, S., Marti, C., Wilson, D. *et al.* (2019). Does ultra-low unemployment spur rapid wage growth? *FRBSF Economic Letter*.

—, Wilson, D. J. *et al.* (2017). Has the wage Phillips curve gone dormant? *FRBSF Economic Letter*, **30**, 16.

Matheson, T. and Stavrev, E. (2013). The Great Recession and the inflation puzzle. *Economics Letters*, **120** (3), 468–472.

McLeay, M. and Tenreyro, S. (2020). Optimal inflation and the identification of the Phillips curve. *NBER Macroeconomics Annual*, **34**, 199–255.

Müller, U. K. (2012). Measuring prior sensitivity and prior informativeness in large bayesian models. *Journal of Monetary Economics*, **59** (6), 581–597.

Otto, S. and Breitung, J. (2022). Backward cusum for testing and monitoring structural change with an application to covid-19 pandemic data. *Econometric Theory*, pp. 1–34.

Perron, P. and Yamamoto, Y. (2015). Using OLS to estimate and test for structural changes in models with endogenous regressors. *Journal of Applied Econometrics*, **30** (1), 119–144.

Pesaran, M. H. (2015). Testing weak cross-sectional dependence in large panels. *Econometric Reviews*, **34** (6-10), 1089–1117.

— (2021). General diagnostic tests for cross-sectional dependence in panels. *Empirical Economics*, **60** (1), 13–50.

Phillips, P. C. and Shi, S.-P. (2018). Financial bubble implosion and reverse regression. *Econometric Theory*, **34** (4), 705–753.

— and Sul, D. (2003). Dynamic panel estimation and homogeneity testing under cross section dependence. *The Econometrics Journal*, **6** (1), 217–259.

Powell, J. H. *et al.* (2020). Opening remarks: New economic challenges and the fedâs monetary policy review. In *Proceedings-Economic Policy Symposium-Jackson Hole*, Federal Reserve Bank of Kansas City, pp. 1–18.

Riker, D. (2022). State-level import penetration, unpublished manuscript.

Roberts, J. M. (2006). Monetary policy and inflation dynamics. *International Journal of Central Banking*, **2** (3), 193–230.

Rognlie, M. (2019). Comment on optimal inflation and the identification of the phillips curve. **34**, 267–279.

Schnabel, I. (2021). Monetary policy and inequality. In *speech delivered at a virtual conference on "Diversity and Inclusion in Economics, Finance, and Central Banking"*, vol. 9.

Smith, S. C. (2018). Noncommon breaks. *Unpublished Manuscript, University of Southern California*.

— (2023). Structural breaks in grouped heterogeneity. *Journal of Business & Economic Statistics*, **41** (3), 752–764.

— and Timmermann, A. (2021). Break risk. *Review of Financial Studies*, **34** (4), 2045–2100.

Stock, J. H. and Watson, M. W. (2009). Phillips curve inflation forecasts. In *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, MIT Press.

— and — (2020). Slack and cyclically sensitive inflation. *Journal of Money, Credit and Banking*, **52** (S2), 393–428.

Table 1: Summaries of existing papers on instability in the Phillips curve

Authors	Sample	Method	Finding	Notes
Ball and Mazumder (2011)	1960-2010	Random Walk parameter Slope coefficient linear function of level and variance	Steepening around 1970, flattening in 80s	Lower and more stable inflation both flatten curve. Paper uses median and core CPI
Ball and Mazumder (2019)	1985-2015	Sup Wald test	Flattening break in 1995	Break identified indirectly from expectations formation. Paper uses median CPI.
Perron and Yamamoto (2015)	1960-1997	Sup Wald test	Break in 1991	Uses GDP deflator.
Matheson & Stavrev (2013)	1961-2012	Random Walk parameter	Flattening in 80s	Uses headline CPI inflation.
Gali and Gambetti (2019)	1964-2017	Regimes with fixed dates	Flattening in 2007	Wage Phillips curve
Leduc and Wilson (2017)	1991-2015	Regimes with fixed dates	Flattening in 2009	Wage Phillips curve
Hooper et al. (2019)	1961-2018	Regimes with fixed dates	Flattening in 1988	Uses headline and core PCE and average hourly earnings and MSA panel data.
Coibion & Gorodnichenko (2015)	1961-2007	Regimes with fixed dates	Possible break in 1985; mixed evidence	No break if augmented with household expectations. Uses various aggregate inflation measures (CPI, core CPI...)
Coibion et al. (2013)	1968-2013	Regimes with fixed dates	Flattening break in 1985	Break in price Phillips curve not wage Phillips curve
Roberts (2006)	1960-2002	Regimes with fixed dates	Flattening break in 1983	Uses core PCE inflation.
Hazell et al. (2002)	1978-2018	Regimes with fixed dates	Break in 1990 but not significant	State level panel data
Cerrato and Gitti (2022)	1990-2022	Regimes with fixed dates	Flattening in pandemic; steepened after	MSA level panel data
Fitzgerald et al. (2023)	1977-2018	Regimes with fixed dates	No significant break	MSA level panel data
Williams (2006)	1980-2016	Recursive regressions	Flattening in the 90s	Core CPI and PCE
Del Negro et al. (2020)	1964-2019	Regimes with fixed dates	Break in 1990	Estimated in VAR
Barnichon & Mesters (2021)	1969-2007	Regimes with fixed dates	Break in 1990	Phillips multiplier not slope of curve. Uses headline PCE
Gilchrist & Zakrajsek (2019)	1962-2017	Sup-Wald test Interact gap with trade share	Mixed results; possible break in 80s	Panel and aggregated data (CPI and PPI)
Inoue et al. (2022)	1970-2021	IV estimation with random walk parameters	Flattening until early 2000s; then steepening	Uses core PCE
Blanchard (2016)	1960-2014	Random walk parameter	Flattening in the 1980s	Uses headline CPI

Table 2: Quarterly 16 PCE industry-level inflation rates (1959-2022)

	1959-1972	1972-2001	2001-2022	1959-2022
		All industries		
PC	-0.48***	-0.82***	-0.28*	-0.42***
AR	0.33***	0.38***	0.16***	0.25***
vol.	1.58	2.91	3.03	2.89
		All industries (ex. food and energy)		
PC	-0.32***	-0.49***	-0.13**	-0.24***
vol.	1.43	2.41	1.96	2.18
		Goods		
PC	-0.59***	-1.09***	-0.32	-0.52***
AR	0.05**	0.38***	0.15***	0.23***
vol.	1.9	3.44	4.01	3.57
		Services		
PC	**-0.34*****	**-0.57*****	-0.19***	-0.27***
AR	0.59***	0.37***	0.38***	0.43***
vol.	1.39	2.34	1.41	2.07

Note: The top panel of this table displays estimates of the slope coefficients on the aggregate unemployment gap (PC) and the autoregressive term (AR) from the baseline model displayed in Equation (3). Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. The reported volatility (vol.) is the weighted average of the sectoral-level volatility estimates, weighted using the 2022:Q1 expenditure weights. This model regresses the 16 PCE sector quarterly inflation rates from 1959 through 2022 on an autoregressive term, the aggregate unemployment gap, and long-term inflation expectations, including industry fixed effects. We display results for the three regimes identified by the model, and for the full sample (by estimating the model but precluding any breaks). The second panel displays results when food and energy sectors are excluded from the model. The third and fourth panels display results when estimating the same model separately for goods and services sectors, while precluding breaks and conditioning on either the regimes identified by the baseline model or on the full sample. The goods group consists of Motor vehicles and parts, Furnishings and durable household equipment, Recreational goods and vehicles, Other durable goods, Food and beverages purchased for off-premises consumption, Clothing and footwear, Gasoline and other energy goods, and Other nondurable goods. The services group consists of Housing and utilities, Health care, Transportation services, Recreation services, Food services and accommodations, Financial services and insurance, Other services, and NPISH. Values in bold font denote that the services PC is significantly different from the goods PC at the 95% confidence level.

Table 3: Grouped heterogeneity estimates: Quarterly 16 PCE industry-level inflation rates (1959-2022)

	1959-1972	1972-2001	2001-2022	1959-2022
Parameter Estimates				
Group 1				
PC	-0.51***	-2.25***	-2.32	-0.96*
vol.	3.13	9.31	39.79	11.09
Group 2				
PC	-0.37***	-0.45***	-0.12**	-0.23***
vol.	1.20	1.89	1.95	1.93
Weighted average				
PC	-0.41	-0.79	-0.26	-0.32
Group Allocation Estimates				
Motor vehicles and parts	1	2	2	2
Furnishings and durable household equipment	2	2	2	2
Recreational goods and vehicles	2	2	2	2
Other durable goods	1	2	2	2
Food and beverages purchased for off-premises consumption	1	2	2	2
Clothing and footwear	2	2	2	2
Gasoline and other energy goods	1	1	1	1
Other nondurable goods	2	2	2	2
Housing and utilities	2	2	2	2
Health care	2	2	2	2
Transportation services	2	2	2	2
Recreation services	2	2	2	2
Food services and accommodations	2	2	2	2
Financial services and insurance	2	1	2	1
Other services	2	2	2	2
NPISH	1	1	2	2

Note: The top panel of this table displays estimates of the slope coefficient on the aggregate unemployment gap (PC) from the model that estimates an unobserved grouping structure as described in Section 3.3. We also display the average slope across groups, weighted by the number of series in each group. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. This model regresses the 16 PCE industry-level quarterly inflation rates from 1959 through 2022 on an autoregressive term, the aggregate unemployment gap, and long-term inflation expectations, including industry fixed effects. We also report the industry-weighted volatility (vol.) estimate within each group, using the 2022:Q1 expenditure weights. The model is estimated within the three regimes identified by the baseline model displayed in Equation (3) that uses the 16 PCE sector inflation rates, and for the full sample. The lower panel displays the corresponding posterior mode group allocations.

Table 4: 31 CPI industry-level quarterly inflation rates (1954-2022)

	1954-1971	1971-2001	2001-2022	1954-2022
		All industries		
PC	-0.51**	-1.46***	-0.27	-0.43***
AR	0.09**	0.23***	0.07***	0.13***
vol.	1.98	3.69	5.20	5.24
		All industries (ex. food and energy)		
PC	-0.64***	-0.49***	-0.12**	-0.21**
AR	0.40***	0.36***	0.08***	0.23***
vol.	1.34	1.69	2.22	2.53
		Goods		
PC	-0.53*	-1.91***	-0.19	-0.54***
AR	0.07	0.19***	0.02	0.09***
vol.	1.94	4.68	6.74	5.41
		Services		
PC	-0.40***	**-0.43****	-0.32*	**-0.28*****
AR	0.36***	0.57***	0.25***	0.29***
vol.	1.04	1.69	3.85	3.07

Note: The top panel of this table displays estimates of the slope coefficients on the aggregate unemployment gap (PC) and the autoregressive term (AR) from the baseline model displayed in Equation (3). Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. The reported volatility (vol.) is the weighted average of the sectoral-level volatility estimates, weighted using the October 2022 relative importance weights. This model regresses the 31 CPI sector quarterly inflation rates from 1954 through 2022 on an AR term, the unemployment gap, and long-term inflation expectations, including industry fixed effects. We display results for the three regimes identified by the model, and for the full sample (estimating the model but precluding any breaks). The other panels display results when excluding food and energy sectors, and running the models separately for goods and services sectors, conditioning on either the regimes identified by the baseline model or on the full sample. The services group consists of Full Service Meals and Snacks, Limited Service Meals and Snacks, Food at employee sites and schools, Food from vending machines and mobile vendors, Other food away from home, Utility (piped) gas service, Shelter, Water and sewer and trash collection services, Household operations, Medical care services, Transportation services, Recreation services, Education and communication services, and Other personal services. The remaining sectors comprise the goods group. Values in bold font denote that the unemployment gap slope for services is significantly different from that of goods at the 95% confidence level.

Table 5: Grouped heterogeneity estimates: 31 CPI industry-level quarterly inflation rates (1954-2022)

	1954-1971	1971-2001	2001-2022	1954-2022
Parameter Estimates				
Group 1				
PC	-0.51^{***}	-1.73^{***}	-0.17^{**}	-0.27^{***}
vol.	1.63	10.79	3.48	5.80
Group 2				
PC		-0.26^{***}	-0.81	-0.63^{***}
vol.		0.77	16.87	13.12
Weighted average				
PC	-0.51	-1.09	-0.27	-0.44
Group Allocation Estimates				
Cereals and Bakery Products		2	1	1
Meats, Poultry, Fish and Eggs	1	1	2	2
Dairy and Related Products		1	2	1
Fruits and Vegetables	1	1	2	2
Nonalcoholic Beverages and Beverage Matls	1	1	1	2
Other Food At Home	1	1	1	2
Full Service Meals and Snacks		2	1	1
Limited Service Meals and Snacks	1	2	1	1
Food at employee sites and schools		2	2	1
Food from vending machines and mobile vendors		2	1	1
Other food away from home		2	1	1
Fuel oil and other fuels	1	1	2	2
Motor fuel	1	1	2	2
Electricity	1	1	1	2
Utility (piped) gas service	1	1	2	2
Household furnishings and supplies			1	1
Apparel	1	1	1	2
Transportation commodities less motor fuel			2	1
Medical care commodities	1	2	1	2
Recreation commodities			1	1
Education and communication commodities	1		1	1
Alcoholic beverages	1	1	1	2
Other goods			1	1
Shelter	1	1	1	2
Water and sewer and trash collection services		2	1	1
Household operations		2	1	1
Medical care services	1	2	1	2
Transportation services	1	1	1	2
Recreation services			1	1
Education and communication services			1	1
Other personal services			1	1

Note: The top panel of this table displays estimates of the slope coefficient on the aggregate unemployment gap (PC) from the model that estimates an unobserved grouping structure as described in Section 3.3. We also display the average slope across groups, weighted by the number of series in each group. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. The reported volatility (vol.) is the weighted average of the sectoral-level volatility estimates within each group, weighted using the October 2022 relative importance weights. This model regresses the 31 CPI industry-level quarterly inflation rates from 1954 through 2022 on an autoregressive term, the aggregate unemployment gap, and long-term inflation expectations, including industry fixed effects. The model is estimated within the three regimes identified by the baseline model displayed in Equation (3) that uses the 31 CPI sector inflation rates, and for the full sample. The lower panel displays the corresponding posterior mode group allocations. Missing group allocations indicate that the corresponding series had no inflation observations in the regime and so was not assigned to any group.

Table 6: Annual 22 CPI MSA-level inflation rates (1980-2022)

	1980-2000	2001-2022	1980-2022
	All MSAs		
PC	-0.26***	-0.21***	-0.21***
PC (scaled)	-0.39	-0.31	-0.32
AR	0.23***	0.39***	0.34***
vol.	0.63	0.65	0.65
Above and below median rate of import penetration from China			
PC (above)	-0.41***	-0.24***	-0.26***
PC (below)	**-0.16****	**-0.15***	**-0.15****
Weighted average	-0.35	-0.22	-0.24
	Kink at 5% or 4.2% U rate		
PC	-0.26***	-0.18***	-0.19***
Extra PC (U. rate <5%)	-0.22**	-0.28***	-0.19***
AR	0.21***	0.36***	0.33***
PC	-0.27***	-0.20***	-0.20***
Extra PC (U. rate <4.2%)	-0.29**	-0.24**	-0.17*
AR	0.21***	0.37***	0.33***

Note: The top panel of this table displays estimates of the slope coefficients on the MSA-level unemployment rates (PC) and the autoregressive term (AR) from the baseline model that includes two-way fixed effects displayed in Equation (1). Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. We also report the slope scaled by the expenditure share on nontradeables to map the regional PC slope into the national PC slope as suggested by Hazell *et al.* (2022). The reported volatility (vol.) is the equal-weighted average of the series-specfic volatility estimates. We display results for the two regimes identified by the model, and for the full sample (by estimating the model but precluding breaks). The middle panel displays corresponding results when, conditional on the regimes identified by the baseline model and for the full sample, we estimate the regression separately for those MSAs that correspond to states with above or below median rates of import penetration from China based on the state-level import penetration rates estimated by Riker (2022) who estimates these values using a structural econometric model that exploits data on the location of import entry, domestic shipments, and distances between states. The MSAs that comprise the below median group are Detroit-Warren-Dearborn, MI, Dallas-Fort Worth-Arlington, TX, Denver-Aurora-Lakewood, CO, Philadelphia-Camden-Wilmington, PA-NJ-DE-MD, and St Louis, MO-IL. Values in bold font denote that the PC slope for the below median rate of import penetration group is significantly different from that of the above median group. We also display the average slope across groups, weighted by the number of series in each group. The lower panel displays results, conditional on the regimes identified by the baseline model and for the full sample, from estimating the nonlinear Phillips curve in Equation (14), using a kink at unemployment rate values below either five or 4.2 percent.

Table 7: Wage Phillips curve: 51 state-level quarterly wage inflation rates (1980-2019)

	1980:1-2000:1	2000:2-2019:4	1980:1-2019:4
		Linear model	
PC	-0.47***	-0.38***	-0.40***
AR	0.01	0.07***	0.04***
		Nonlinear model	
		Kink at 5%	
PC	-0.41***	-0.32***	-0.35***
Extra PC (U < 5%)	-0.54***	-0.48***	-0.52***
AR	0.00	0.06***	0.04***
		Kink at 4.2%	
PC	-0.43***	-0.36***	-0.37***
Extra PC (U < 4.2%)	-0.74**	-0.62***	-0.69***
AR	0.00	0.06***	0.04***

Note: The top panel of this table displays estimates of the slope coefficient on the state-level unemployment rate (PC) and the autoregressive term (AR) when regressing the 51 (including the District of Columbia) state-level quarterly wage inflation rates (growth rates of average hourly earnings of production and nonsupervisory workers) from 1980 through 2019 on an autoregressive term and the state-level unemployment rates, including industry and time fixed effects. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. We display results for the two regimes identified by the model. Average Hourly Earnings of production and nonsupervisory workers are at the quarterly frequency beginning in 1980:Q1 and ending in 2019:Q4, sourced from the CEPR extract of the underlying CPS data. The middle and lower panels display estimates when including a kinkpoint for unemployment rate values below 5 or 4.2 percent, and conditioning on either the full sample or the two regimes identified by the baseline model.

Table 8: Annual 28 EU countries inflation (1986-2021)

	1986-2003	2004-2021	1986-2021
All countries			
PC	-0.78**	-0.14***	-0.23**
PC (scaled)	-1.11	-0.20	-0.33
AR	0.10	0.51***	0.53***
vol.	2.37	1.04	1.78
Rich vs poor			
PC (rich)	-0.19**	-0.07	-0.12**
PC (poor)	**-0.92**	-0.11***	-0.23
Weighted average	-0.48	-0.09	-0.16
Goods vs services			
PC (servs.)	-0.26***	-0.22***	-0.24***
PC (goods)	-0.29**	**-0.09***	**-0.11***
Kink at -1.5%			
PC	-1.15***	-0.07**	-0.21*
Extra PC (UGAP < -1.5%)	3.28	-0.56***	-0.17
AR	0.09	0.49***	0.53***

Note: The top panel of this table displays estimates of the slope coefficient on the country-level unemployment gaps (PC) and the autoregressive term (AR) from the baseline model displayed in Equation (2). Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. We also report the slope scaled by the expenditure share on nontradeables to map the country-level PC slope into the EU aggregate PC slope as suggested by Hazell *et al.* (2022). The reported volatility (vol.) is the weighted average of the country-level volatility estimates, using HICP country weights. This model regresses the 28 EU (including the UK) country-level annual inflation rates from 1986 through 2021 on an autoregressive term and the country-level unemployment gaps, including two-way fixed effects. We display results for the two regimes identified by the model, and for the full sample (by estimating the model but precluding any breaks). The second panel displays results when estimating the same model separately for rich and poor countries – while precluding breaks and conditioning on either the regimes identified by the baseline model or on the full sample. Rich countries are defined as countries with real GDP per capita deflated by PPP in 2019 above the EU average and poor countries are defined as the rest. We also display the average slope across rich vs poor groups, weighted by the number of series in each group. The third panel displays results when using either total services or total goods inflation, rather than total inflation for each country. Values in bold font denote that the PC for goods (poor countries) is significantly different from that of services (rich countries) at the 95% confidence level. The final panel displays results, conditional on the regimes identified by the baseline model and for the full sample, when including a kink at an unemployment gap below minus 1.5 percent as displayed in Equation (14).

Table 9: Grouped heterogeneity estimates: Annual 28 EU countries inflation (1986-2021)

	1986-2003	2004-2021
Parameter Estimates		
Group 1		
PC	-0.42**	-0.14***
vol.	1.81	1.04
Group 2		
PC	-0.08	
vol.	0.67	
Group Allocation Estimates		
Germany	2	1
Belgium	2	1
Bulgaria	1	1
Cyprus	1	1
Croatia	2	1
Czech Republic	1	1
Denmark	2	1
Estonia	1	1
Spain	2	1
Finland	2	1
France	2	1
Greece	2	1
Hungary	2	1
Ireland	1	1
Italy	2	1
Lithuania	2	1
Latvia	2	1
Luxembourg	2	1
Malta	2	1
Netherlands	2	1
Austria	2	1
Poland	1	1
Portugal	1	1
Romania		1
Sweden	1	1
Slovenia	2	1
Slovakia	1	1
United Kingdom	2	1

Note: The top panel of this table displays estimates of the slope coefficient on the EU country-level unemployment gaps (PC) from the model that estimates an unobserved grouping structure as described in Section 3.3. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. The reported volatility (vol.) is the weighted average of the country-level volatility estimates within each group, using HICP country weights. This model regresses the EU country-level annual inflation rates from 1986 through 2021 on an autoregressive term and the country-level unemployment gaps, and includes two-way fixed effects. The model is estimated within the two regimes identified by the baseline model displayed in Equation (2). The lower panel displays the corresponding posterior mode group allocations. Due to high volatility and extreme outliers, Romania was omitted from the analysis in the first regime.

Table 10: Phillips curve estimates with 31 CPI industry-level inflation rates and shorter prior regime duration

1954-1971	1971-2001	2001-2007	2007-2019	2020-2022
		All industries (CPI)		
-0.51**	-1.46***	0.11	-0.02	-1.60***
		All industries (ex. food and energy)		
-0.64***	-0.49***	-0.39*	-0.04	-1.16***
		Goods		
-0.53*	-1.91***	0.06	0.23	-1.73**
		Services		
-0.40***	**-0.43****	0.36	-0.29*	**-1.24**

Note: The top panel of this table displays estimates of the slope coefficients on the aggregate unemployment gap from the baseline model displayed in Equation (3), but using a prior expected regime duration of five years rather than 20 years, when regressing the 31 CPI sector quarterly inflation rates from 1954 through 2022 on an AR term, the unemployment gap, and long-term inflation expectations, including industry fixed effects. We display results for the five regimes identified by the model. Coefficients other than the Phillips curve slope are not reported. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. The other panels display results when excluding food and energy sectors, and running the models separately for goods and services sectors, conditioning on the regimes identified by the baseline model. The services group consists of Full Service Meals and Snacks, Limited Service Meals and Snacks, Food at employee sites and schools, Food from vending machines and mobile vendors, Other food away from home, Utility (piped) gas service, Shelter, Water and sewer and trash collection services, Household operations, Medical care services, Transportation services, Recreation services, Education and communication services, and Other personal services. The remaining sectors comprise the goods group. Values in bold font denote that the unemployment gap slope for services is significantly different from that of goods at the 95% confidence level.

Table 11: Post-pandemic Phillips curve slope estimates (beginning 2020:Q1)

	EU	PCE	CPI
All sectors and ex. food and energy			
All sectors	-0.51**	-2.39**	-1.60***
ex. food and energy	-	-2.38**	-1.16***
Goods and services sectors			
Goods	-0.59***	-3.69**	-1.73**
Services	-0.40	-0.29*	**-1.24**
Rich and poor countries			
Rich	**-0.02**	-	-
Poor	-0.50***	-	-
Nonlinear PC			
PC	-0.43**	-	-
extra PC (UGAP <1.5%)	-0.87**	-	-

Note: The final two columns of the top row of this table display estimates of the slope coefficients on the aggregate unemployment gap from the baseline model displayed in Equation (3) using all of the (PCE or CPI) sectoral US inflation data in the final regime identified when using a prior expected regime duration of five years rather than 20 years. The break date that marks the final regime is identified in 2020:Q1 in both cases. The results for CPI are the same as in Table 10 for the last regime, but are repeated as a memo item. The next three rows display results when excluding food and energy, services, and goods sectors. The first column displays corresponding results for the EU when imposing the break date identified from the US sectoral data sets, such that the regime begins in 2020. For the EU, we also display results when using only rich or poor countries. The bottom panel displays results for the EU in the pandemic regime using the nonlinear Phillips curve specification with a threshold that defines a hot labor market as one with an unemployment rate gap below -1.5 percent. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. Values in bold font denote that the unemployment gap slope for services sectors (rich countries) is significantly different from that of goods sectors (poor countries) at the 95% confidence level.

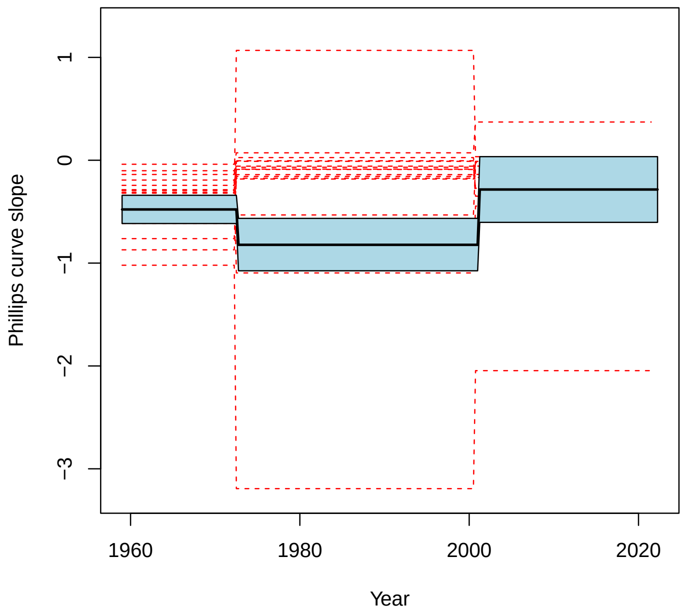

(a) PCE

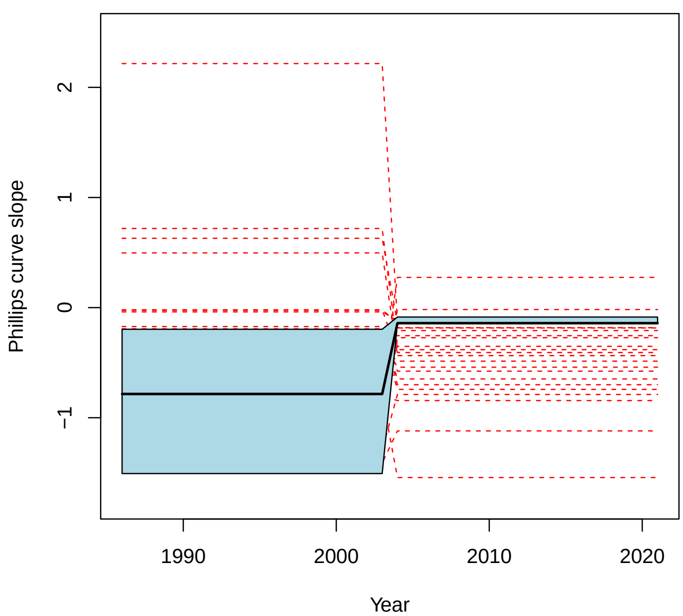

(b) EU

Figure 1: The black line in the top panel of this figure graphs the evolution of the posterior mean Phillips curve slope over time estimated from our baseline breakpoint model displayed in Equation (3) using the PCE sectoral data. The blue bands cover the corresponding 95 percent posterior interval of the estimates. The red dotted lines graph the OLS time series estimates for each individual sector, conditioning on each of the regimes identified by our breakpoint model. For illustrative clarity, the red dotted lines are not allowed to overlay the blue shaded area. The lower panel displays the same information but uses the EU data and the breakpoint model displayed in Equation (2).

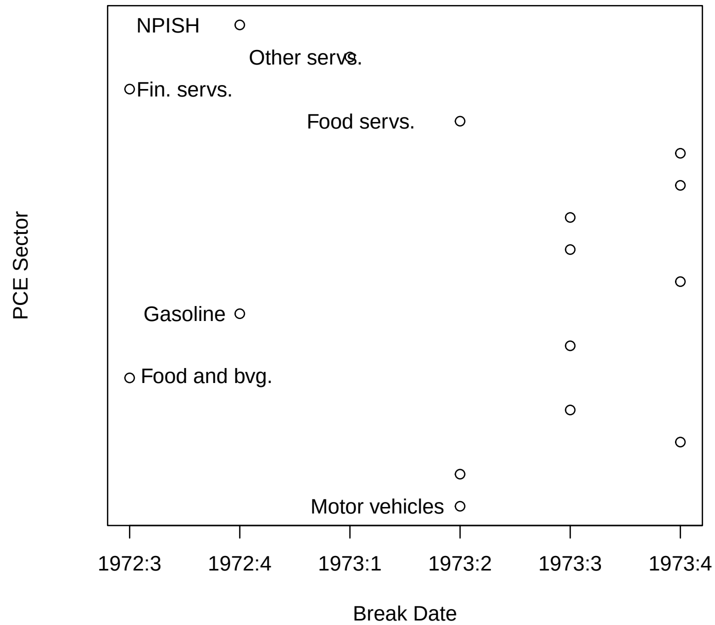

(a) 1972 Breakpoint

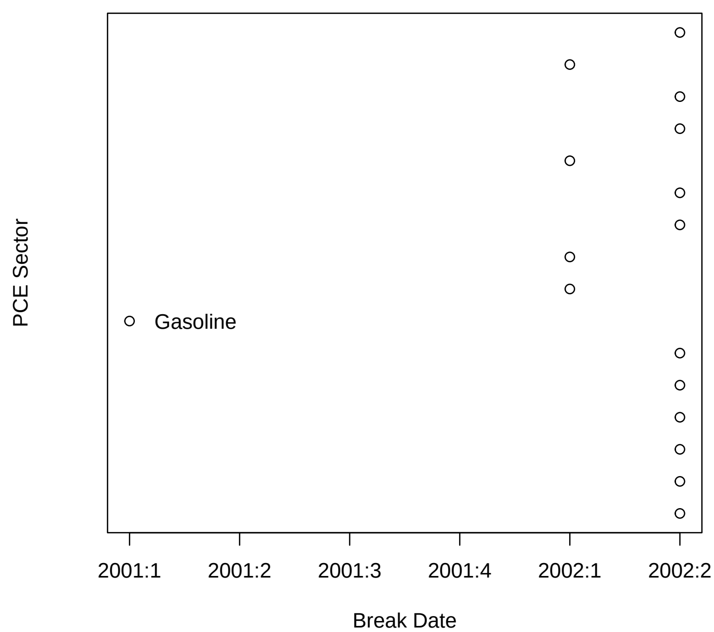

(b) 2001 Breakpoint

Figure 2: This figure displays the posterior mode break dates estimated from the model that regresses the 16 PCE industry-level quarterly inflation rates from 1959 through 2022 on an autoregressive term, the aggregate unemployment gap, and long-term inflation expectations, including industry fixed effects, and allowing the timing of the breaks to vary across industries as described in Section 3.2. The top panel displays results for the 1972 break, and the lower panel displays results for the 2001 break.

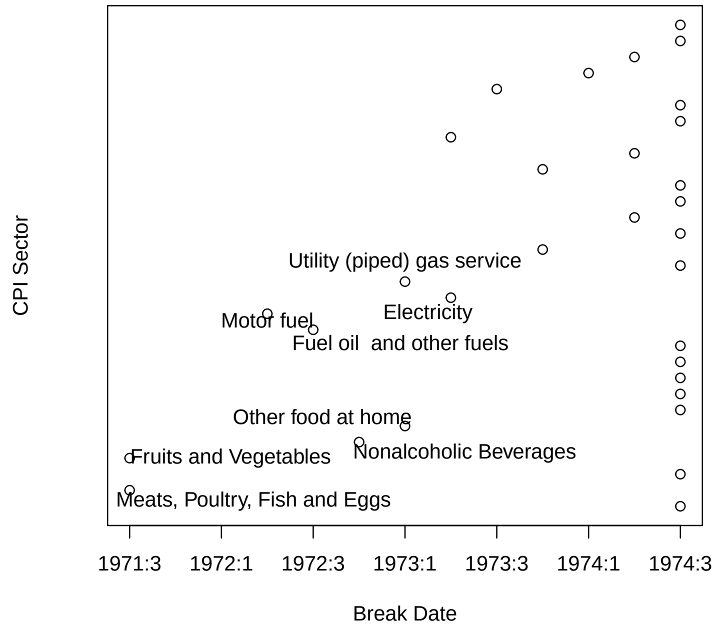

(a) 1971 Breakpoint

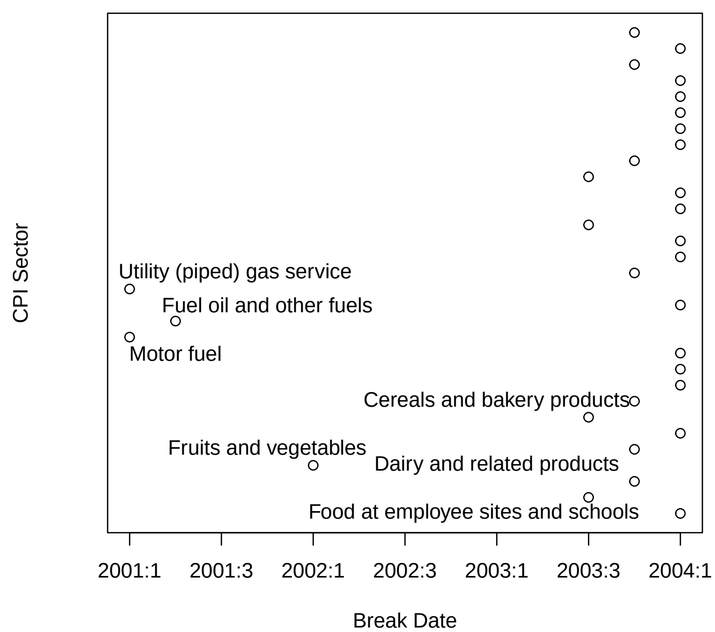

(b) 2001 Breakpoint

Figure 3: This figure displays the posterior mode break dates estimated from the model that regresses the 31 CPI industry-level quarterly inflation rates from 1954 through 2022 on an autoregressive term, the aggregate unemployment gap, and long-term inflation expectations, including industry fixed effects, and allowing the timing of the breaks to vary across industries as described in Section 3.2. The top panel displays results for the 1971 break, and the lower panel displays results for the 2001 break.

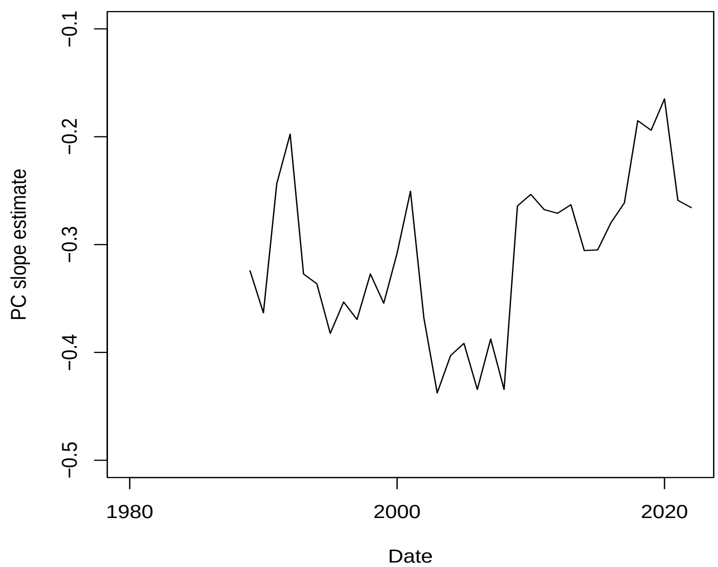

(a) MSA

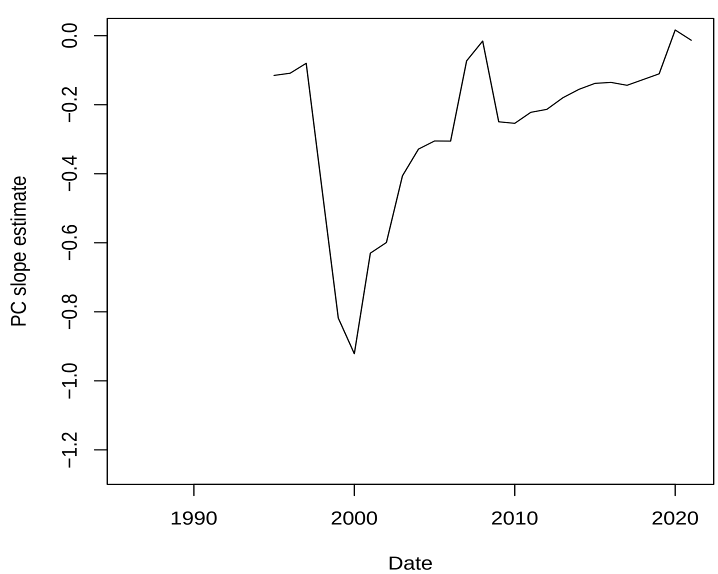

(b) EU

Figure 4: The top panel of this figure displays the Phillips curve slope coefficient estimated throughout our sample using the MSA-level data and the model specification displayed in Equation (1), but estimating it as a panel no-break OLS regression with two-way fixed effects with a ten-year rolling window that exponentially down-weights older data using a decay parameter of 0.8. Coefficient estimates are plotted against the end date of the rolling window. The bottom panel displays corresponding estimates using EU country-level data and the model specification displayed in Equation (2).

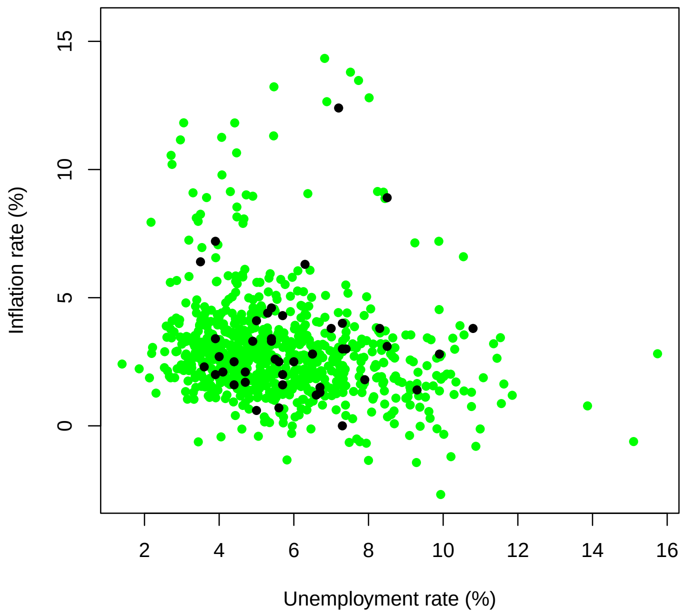

(a) MSA

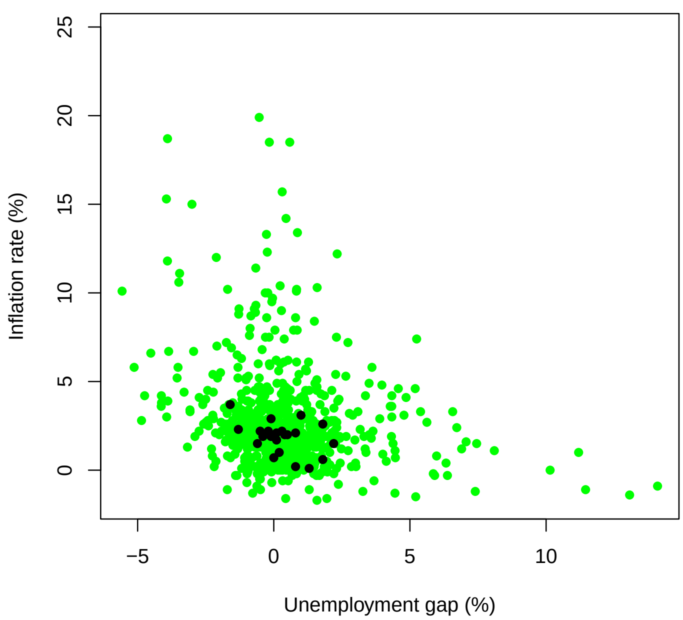

(b) EU

Figure 5: The green (black) circles in the top panel of this figure plot the annual headline CPI inflation rate against the annual unemployment rate for the 22 MSAs (national aggregate) during our sample period from 1980 through 2022. The green (black) circles in the bottom panel plot the annual inflation rate against the annual unemployment gap for each of the EU countries (the EU aggregate) in our sample from 1986 through 2021. Due to extreme outliers, Romania is excluded from the plot.

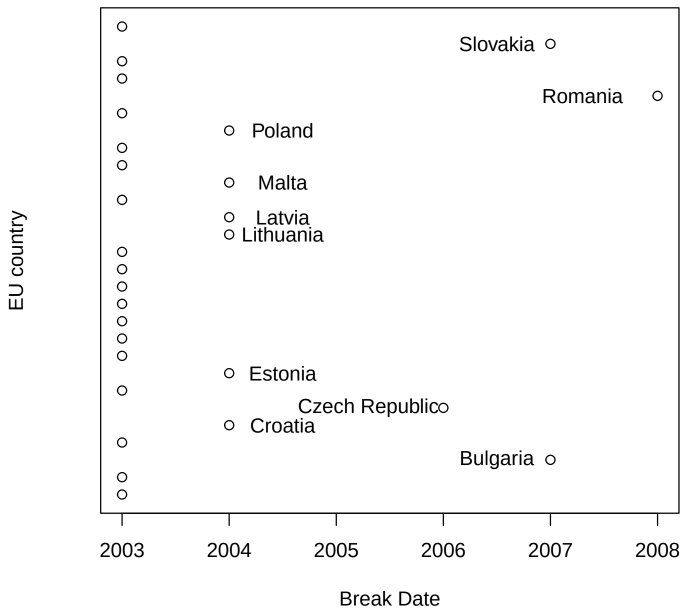

(a) Break date

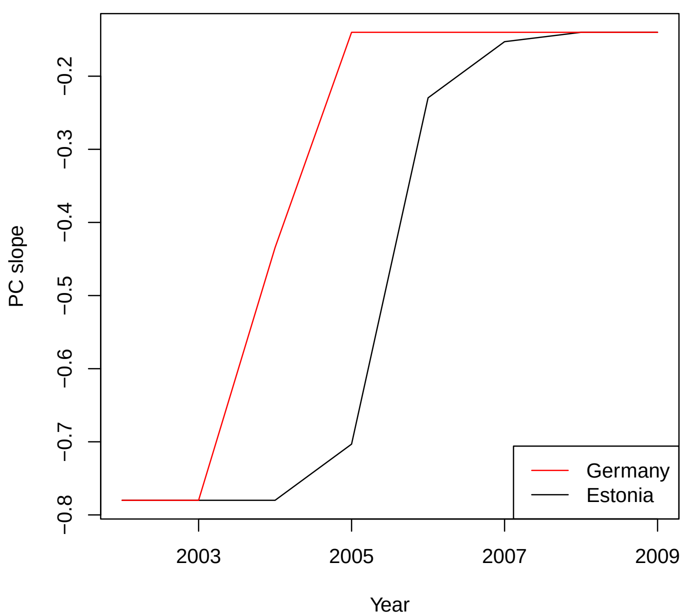

(b) Phillips curve slope

Figure 6: The top panel of this figure displays the posterior mode break dates estimated from the model that regresses the 28 EU country-level annual inflation rates from 1986 through 2021 on an autoregressive term and the country-level unemployment gaps, including industry and time fixed effects, and allowing the timing of the breaks to vary across countries as described in Section 3.2. The lower panel displays how the corresponding Phillips curve slope coefficients for two countries – Germany (red line) and Estonia (black line) – evolve over time, incorporating uncertainty surrounding the break date estimated from the noncommon breaks model.

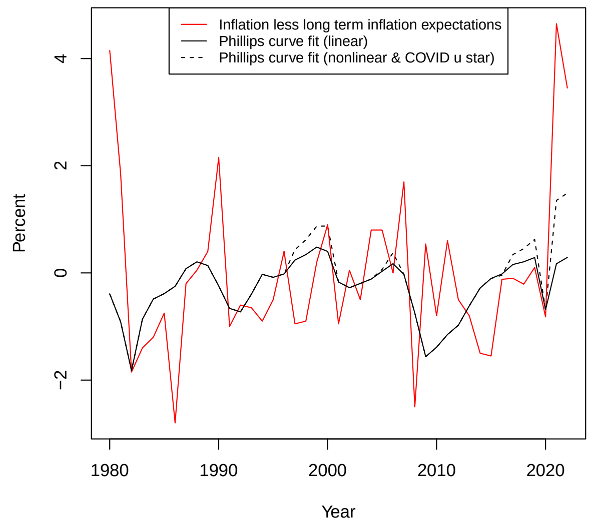

(a) Headline CPI

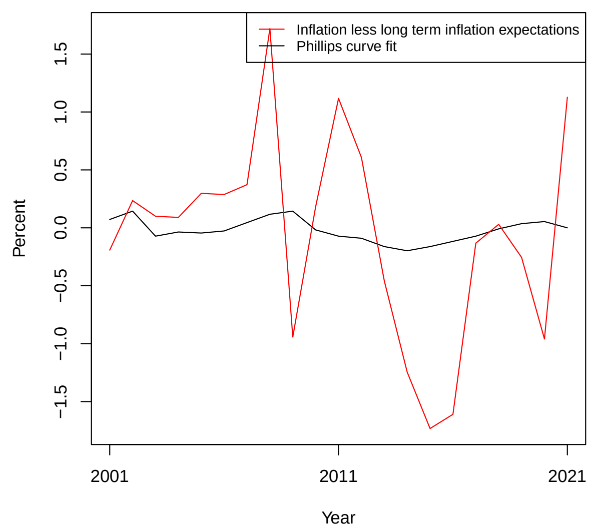

(b) EU

Figure 7: The top panel of this figure displays the linear Phillips curve fit from our MSA breakpoint model (solid black line). Specifically, in each year this is our prevailing regime-specific MSA regional Phillips slope coefficient divided by the nontradeables share multiplied by the national unemployment gap. The red line graphs the annual national headline CPI inflation rate minus long term inflation expectations, which are 10-year ahead SPF CPI inflation expectations back to 1991 Q4. Missing observations prior to 1991 Q4 are filled using linear interpolation. The dotted black line uses our implied national PC slopes from the nonlinear Phillips curve estimated in the second regime using MSA-level data and replacing the noncyclical rate of unemployment in 2021 and 2022 with the higher estimates from Crump *et al.* (2022). The lower panel plots the same information for the EU. Specifically, the black line is our estimated prevailing regime-specific EU linear PC slope coefficient multiplied by the EU unemployment gap. The red line uses the EU inflation rate and long term (five-year ahead) Eurozone inflation expectations from the ECB SPF which goes back to 2002 Q3. Prior to this, we use one-year ahead expectations, going back to 1999 Q1. Eurozone expectations data are sourced from the ECB statistical data warehouse. We average expectations across the four quarters in a given year.

Appendix A. Appendix Tables

Table A1: Time series regressions: 16 PCE industry-level quarterly inflation rates (1959-2022)

	1959-1972	1972-2001	2001-2022	1959-2022	1959-1972	1972-2001	2001-2022	1959-2022
	Motor vehicles and parts				Furnishings and durable household equipment			
PC	-0.04	0.07	0.37	0.22	-0.32***	-0.09	-0.17	-0.13
corr	-0.03	0.21	0.14	0.14	-0.61	0.19	-0.19	-0.11
	Recreational goods and vehicles				Other durable goods			
PC	-0.10	-0.07	-0.13	-0.09	-0.54*	-0.14	-0.04	-0.13
corr	-0.18	0.27	-0.25	-0.14	-0.33	0.04	-0.03	-0.08
	Food and beverages purchased for off-premises consumption				Clothing and footwear			
PC	-0.87***	-0.54**	-0.06	-0.23*	-0.61***	-0.01	-0.27	-0.39***
corr	-0.34	-0.29	-0.07	-0.17	-0.73	0.03	-0.11	-0.19
	Gasoline and other energy goods				Other nondurable goods			
PC	-1.02*	-3.19**	-2.05	-1.69	-0.29**	-0.18	-0.02	-0.09
corr	-0.20	-0.27	-0.11	-0.12	-0.52	0.21	0.01	0.04
	Housing and utilities							
PC	-0.14	0.03	-0.14**	-0.05				
corr	-0.22	0.33	-0.62	0.02				
	Health care				Transportation services			
PC	-0.76***	-0.16	-0.01	-0.09	-0.38	-0.09	-0.35	-0.17
corr	-0.58	0.39	-0.02	-0.00	-0.29	0.07	-0.21	-0.10
	Recreation services				Food services and accommodations			
PC	-0.24*	-0.06	-0.27***	-0.11*	-0.53***	-0.18	-0.14	-0.12**
corr	-0.29	0.06	-0.37	-0.16	-0.59	-0.06	-0.35	-0.21
	Financial services and insurance				Other services			
PC	-0.31	1.07**	0.03	0.06	-0.29	-0.00	-0.11	-0.09
corr	-0.35	0.18	-0.03	0.02	-0.36	0.18	-0.19	-0.04
	NPISH							
PC	-0.19	-1.09***	-0.45*	-0.50***				
corr	-0.26	-0.40	-0.33	-0.32				

Note: This table displays estimates of the slope coefficient on the aggregate unemployment gap (PC) when estimating OLS time series regressions of each of the 16 PCE sector quarterly inflation rates from 1959 through 2022 on an intercept, an autoregressive term, the aggregate unemployment gap, and long-term inflation expectations. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. We estimate the model conditioning on each of the three regimes identified by the baseline PCE model displayed in Equation (3), and for the full sample. Within each of the three regimes, and for the full sample, we also report the correlation between the industry's inflation rate and the aggregate unemployment gap.

Table A2: Time series regressions: 31 CPI industry-level quarterly inflation rates (1954-2022)

	1954-1971	1971-2001	2001-2022	1954-2022	1954-1971	1971-2001	2001-2022	1954-2022	1954-1971	1971-2001	2001-2022	1954-2022
	Cereals and Bakery Products				Meats, Poultry, Fish and Eggs				Dairy and Related Products			
PC		-0.16	-0.01	-0.06	-6.64***	-2.32**	0.71*	-0.45		-1.47	-0.30	-0.41
corr		0.22	-0.11	-0.09	-0.52	-0.15	0.12	-0.06		-0.15	-0.11	-0.13
	Fruits and Vegetables				Nonalcoholic Beverages and Beverage Matls				Other Food At Home			
PC	0.01	-2.93***	-0.38	-0.87**	-0.93	-0.15	-0.16	-0.25	-0.86	-1.33**	-0.10	-0.49**
corr	-0.01	-0.10	-0.06	-0.05	-0.09	0.03	-0.13	0.00	0.19	-0.14	-0.08	-0.12
	Full Service Meals and Snacks				Limited Service Meals and Snacks				Food at employee sites and schools			
PC		0.54	-0.07	-0.04		-1.42	-0.06	-0.04		5.19	-0.56	-0.45
corr		0.17	-0.36	-0.29		-0.24	-0.15	-0.08		0.33	-0.06	-0.06
	Food from vending machines and mobile vendors				Other food away from home				Fuel oil and other fuels			
PC		-0.89	-0.24	-0.10		-0.62	-0.24	-0.23*	-0.44	-6.53***	-0.58	-1.41
corr		-0.21	-0.12	-0.04		-0.11	-0.19	-0.20	-0.20	-0.29	-0.09	-0.12
	Motor fuel				Electricity				Utility (piped) gas service			
PC	-0.04	-5.74***	-1.22	-2.07	0.07	-1.18***	-0.43	-0.22	0.44	-0.20	-1.24	-0.43
corr	0.03	-0.24	-0.04	-0.08	0.04	0.07	-0.26	0.01	0.14	0.15	-0.15	-0.03
	Household furnishings and supplies				Apparel				Transportation commodities less motor fuel			
PC			0.28	0.28	-0.96***	-0.71***	0.03***	-0.42			0.49*	0.49*
corr			-0.11	-0.11	-0.66	0.02	-0.02	-0.15			0.05	0.05
	Medical care commodities				Recreation commodities				Education and communication commodities			
PC	-0.00	-0.14	0.08	0.08			-0.07	-0.07			0.29	0.29
corr	0.26	0.45	-0.01	0.19			-0.14	-0.14			0.13	0.13
	Alcoholic beverages				Other goods				Shelter			
PC	-0.25	-0.65***	-0.08**	-0.29***			-0.27	-0.27	-0.93***	-1.56***	-0.33**	-0.55***
corr	-0.01	-0.01	-0.13	-0.11			-0.28	-0.28	-0.55	-0.09	-0.73	-0.19
	Water and sewer and trash collection services				Household operations				Medical care services			
PC		-0.26	0.17*	0.21***		-3.46	-0.81***	-0.78***	-0.45**	-0.30*	-0.09	-0.21***
corr		0.14	0.26	0.37		-0.45	-0.41	-0.40	-0.43	0.35	-0.17	-0.01
	Transportation services				Recreation services				Education and communication services			
PC	-0.39	-0.31	-0.52	-0.39***			-0.28*	-0.28*			0.18	0.18
corr	-0.25	0.23	-0.19	-0.08			-0.23	-0.23			0.23	0.23
	Other personal services											
PC			-0.27*	-0.27*								
corr			-0.32	-0.32								

Note: This table displays estimates of the slope coefficient on the aggregate unemployment gap (PC) when estimating OLS time series regressions of each of the 31 CPI sector quarterly inflation rates from 1954 through 2022 on an intercept, an autoregressive term, the aggregate unemployment gap, and long-term inflation expectations. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. We estimate the model conditioning on each of the three regimes identified by the baseline CPI model displayed in Equation (3), and for the full sample. Within each of the three regimes, and for the full sample, we also report the correlation between the industry's inflation rate and the aggregate unemployment gap. Missing values indicate that the industry has insufficient inflation observations in the corresponding regime to either estimate the regression or compute the correlation.

Table A3: Time series regressions: 22 CPI MSA-level annual inflation rates (1980-2022)

	1980-2000	2001-2022	1980-2022	1980-2000	2001-2022	1980-2022
	Urban Alaska			Atlanta-Sandy Springs-Roswell, GA		
PC	-0.46	-0.89*	-0.41	-0.17	-0.48**	-0.34
corr	-0.08	-0.40	-0.18	0.39	-0.53	-0.28
	Boston-Cambridge-Newton, MA-NH			Baltimore-Columbia-Towson, MD		
PC	0.54	-0.51***	-0.24*	-0.33**	-0.47*	-0.40***
corr	0.75	-0.60	-0.39	0.44	-0.36	0.16
	Chicago-Naperville-Elgin, IL-IN-WI			Detroit-Warren-Dearborn, MI		
PC	-0.09	-0.29*	-0.19*	-0.03	-0.19*	-0.13*
corr	0.38	-0.39	0.05	0.11	-0.42	-0.24
	Denver-Aurora-Lakewood			Houston-The Woodlands-Sugar Land, TX		
PC	0.14	-0.52*	-0.36*	0.47**	-0.33*	-0.28*
corr	0.42	-0.42	-0.29	0.65	-0.37	-0.37
	Los Angeles-Long Beach-Anaheim, CA			Miami-Fort Lauderdale-West Palm Beach, FL		
PC	0.07	-0.24	-0.12	-0.33**	-0.29***	-0.25***
corr	0.33	-0.23	-0.12	-0.11	-0.57	-0.49
	Minneapolis-St Paul-Bloomington, MN-WI			Dallas-Fort Worth-Arlington, TX		
PC	0.07	-0.48	-0.36	-0.57	-0.45	-0.40
corr	0.26	-0.59	-0.46	0.32	-0.42	-0.06
	New York-Newark-Jersey City, NY-NJ-PA			Philadelphia-Camden-Wilmington, PA-NJ-DE-MD		
PC	-0.25	-0.15	-0.15	-0.39	-0.34*	-0.31**
corr	0.08	-0.29	-0.16	0.15	-0.41	-0.31
	Phoenix-Mesa-Scottsdale, AZ			Riverside-San Bernardino-Ontario, CA		
PC		-0.56**	-0.56**		-0.35	-0.35
corr		-0.57	-0.57		-0.33	-0.33
	San Diego-Carlsbad, CA			San Francisco-Oakland-Hayward, CA		
PC	-0.53***	-0.29**	-0.27***	-0.44***	-0.35***	-0.33***
corr	0.21	-0.55	-0.12	-0.56	-0.68	-0.66
	St Louis, MO-IL			Seattle-Tacoma-Bellevue WA		
PC	-0.10	-0.21	-0.17		-0.46	-0.44
corr	0.51	-0.27	0.22		-0.59	-0.55
	Tampa-St Petersburg-Clearwater, FL			Washington-Arlington-Alexandria, DC-VA-MD-WV		
PC		-0.26*	-0.23	-0.69***	-0.38	-0.36**
corr		-0.44	-0.39	0.20	-0.39	-0.13

Note: This table displays estimates of the slope coefficient on the MSA-level unemployment rate (PC) when estimating OLS time series regressions of each of the 22 MSA-level annual inflation rates from 1980 through 2022 on an intercept, an autoregressive term, and the MSA-level unemployment rate. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. We estimate the model conditioning on each of the two regimes identified by the baseline MSA model displayed in Equation (1), and for the full sample. Within each of the two regimes, and for the full sample, we also report the correlation between the MSA's inflation rate and its unemployment rate. Missing values indicate that the MSA has insufficient inflation observations in the corresponding regime to either estimate the regression or compute the correlation.

Table A4: Annual 22 CPI MSA-level inflation rates (1980-2022)

	1980-2000	2001-2022	1980-2022
Above median rate of import penetration from China			
PC	-0.39***	-0.21***	-0.23***
Extra PC (U. rate < 5%)	-0.21	-0.22**	-0.13*
PC	-0.39***	-0.23***	-0.25***
Extra PC (U. rate < 4.2%)	-0.26	-0.17	-0.11
Below median rate of import penetration from China			
PC	-0.16**	-0.15*	-0.15***
Extra PC (U. rate < 5%)	-0.09	**-0.51****	-0.41**
PC	-0.16***	-0.16*	-0.16***
Extra PC (U. rate < 4.2%)	-0.02	**-0.55***	-0.41*

Note: The top panel of this table displays estimates of the slope coefficients on the MSA-level unemployment rates from the nonlinear Phillips curve displayed in Equation (14) using – a kink at unemployment rate values below either five or 4.2 percent – and only those MSAs that correspond to states with above median rates of import penetration from China based on the state-level import penetration rates estimated by Riker (2022) who estimates these values using a structural econometric model that exploits data on the location of import entry, domestic shipments, and distances between states. We display results conditional on the two regimes identified by the baseline model and for the full sample. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. The lower panel displays the same results using only those MSAs with below median rates of import penetration from China. The MSAs that comprise the below median group are Detroit-Warren-Dearborn, MI, Dallas-Fort Worth-Arlington, TX, Denver-Aurora-Lakewood, CO, Philadelphia-Camden-Wilmington, PA-NJ-DE-MD, and St Louis, MO-IL. Values in bold font denote that in the post-2000 regime, the additional Phillips curve slope in a hot labor market for the below median rate of import penetration group is significantly different from that of the above median group.

Table A5: Time series regressions: 28 EU country-level annual inflation rates (1986-2021)

	1986-2003	2004-2021	1986-2021	1986-2003	2004-2021	1986-2021
		Germany			Belgium	
PC	-0.43*	-0.12	-0.31*	-0.29	-0.74	-0.46*
corr	-0.46	-0.12	-0.34	-0.47	-0.35	-0.35
		Bulgaria			Cyprus	
PC	0.49	-0.44	0.44	-0.54	-0.21*	-0.24**
corr	-0.31	-0.30	-0.17	-0.23	-0.49	-0.47
		Croatia			Czech Republic	
PC	2.22*	-0.13	0.01	-0.79	-1.54***	-1.15*
corr	0.91	-0.17	0.04	-0.34	-0.66	-0.45
		Denmark			Estonia	
PC	-0.17	-0.18	-0.17	-0.31	-0.54*	-0.48*
corr	-0.43	-0.19	-0.25	-0.52	-0.47	-0.39
		Spain			Finland	
PC	-0.04	-0.12	-0.07	-0.20***	-0.70*	-0.18***
corr	0.36	-0.39	-0.12	-0.69	-0.53	-0.47
		France			Greece	
PC	-0.19	-0.58	-0.25	-0.54	-0.18***	-0.09
corr	-0.29	-0.51	-0.30	-0.29	-0.79	-0.45
		Hungary			Ireland	
PC	0.62	-0.14	-0.19	-0.02	-0.41*	-0.21
corr	0.89	0.26	0.17	-0.16	-0.49	-0.28
		Italy			Lithuania	
PC	-0.25	-0.25	-0.09	0.72***	-0.38**	-0.34***
corr	-0.34	-0.48	-0.01	-0.59	-0.32	-0.39
		Latvia			Luxembourg	
PC	-0.20	-0.65***	-0.65***	-1.29	-0.79	-0.67
corr	-0.37	-0.54	-0.55	-0.51	-0.28	-0.29
		Malta			Netherlands	
PC	-1.41**	-1.12	-1.11	-0.65***	-0.35*	-0.51***
corr	-0.89	-0.28	-0.26	-0.53	-0.32	-0.45
		Austria			Poland	
PC	-0.80*	-0.84***	-0.61**	-1.08	-0.14	-0.22*
corr	-0.09	-0.54	-0.34	-0.89	-0.15	-0.35
		Portugal			Romania	
PC	-0.72	-0.02	-0.38	-3.47	-1.49***	-2.35
corr	-0.46	0.14	-0.18	-0.48	-0.03	-0.39
		Sweden			Slovenia	
PC	-0.36	-0.49***	-0.33*	-1.18	-0.27	-0.16
corr	-0.46	-0.39	-0.39	-0.33	-0.41	-0.23
		Slovakia			United Kingdom	
PC	-0.02	-0.13	-0.10	-0.67**	0.27	-0.25
corr	0.07	0.05	0.26	-0.27	0.51	-0.00

Note: This table displays estimates of the slope coefficient on the EU country-level unemployment gap (PC) when estimating OLS time series regressions of each of the 28 EU country-level annual inflation rates from 1986 through 2021 on an intercept, an autoregressive term, and the country-level unemployment gap. Significance at the 10, 5, and 1 percent levels are denoted by *, **, and ***. We estimate the model conditioning on each of the two regimes identified by the baseline EU model displayed in Equation (2), and for the full sample. Within each of the two regimes, and for the full sample, we also report the correlation between the country's inflation rate and its unemployment rate gap.

Table A6: Testing for error dependence

	PCE	CPI	EU	MSA
Durbin-Watson *p*-value				
Regime 1	0.43	0.68	0.90	0.48
Regime 2	0.62	0.72	0.31	0.15
Regime 3	0.31	0.58		
CD test statistic				
Regime 1	0.89	0.32	1.58	1.75
Regime 2	1.48	1.57	0.29	0.58
Regime 3	0.17	1.35		

Note: The top panel of this table reports the p-values from a Durbin-Watson test for serial correlation in the residuals from our baseline panel breakpoint models across every regime and all four data sets we consider for the price Phillips curve. Here, we exclude observations for Romania when computing the p-value of the DW test in the first regime due to extreme and volatile outliers. The lower panel displays the bias corrected CD test statistic, which has a standard Normal distribution, proposed by Juodis and Reese (2022) in each regime across the same four data sets. The first ten time periods of the first regime are excluded when computing the test statistic for the MSA data because more than half of the series have missing observations.

Table A7: Prior informativeness

	PCE	CPI	MSA	EU
λ	0.07	0.13	0.04	0.14
ρ	0.05	0.03	0.03	0.09
σ^2	0.06	0.14	0.03	0.06

Note: The first column of this table reports the prior informativeness (PI) measure of Müller (2012) for the Phillips curve slope (λ), persistence parameter (ρ), and error variance (σ^2) used in our baseline common breakpoint model applied to the PCE data displayed in Equation (2). The next three columns report the same information for the CPI, MSA, and EU data sets. The priors are the same across all four data sets.