Trade Exposure and the Evolution of Inflation Dynamics

Simon Gilchrist
New York University and NBER

Egon Zakrajšek
Federal Reserve Board

Pages 173-226 | Chapter printout

Changing Inflation Dynamics, Evolving Monetary Policy

Gonzalo Castex, Jordi Galí, and Diego Saravia
editors

27
Series on Central Banking, Analysis, and Economic Policies

Banco Central de Chile
Central Bank of Chile

The views and conclusions presented in the book are exclusively those of the authors and do not necessarily reflect the position of the Central Bank of Chile or its Board Members.

Copyright © Banco Central de Chile 2020
Agustinas 1180
Santiago, Chile
All rights reserved
Published in Santiago, Chile by the Central Bank of Chile
Manufactured in Chile

This book series is protected under Chilean Law 17336 on intellectual property. Hence, its contents may not be copied or distributed by any means without the express permission of the Central Bank of Chile. However, fragments may be reproduced, provided that a mention is made of the source, title, and author(s).

ISBN (print) 978-956-7421-67-1
ISBN (digital) 978-956-7421-68-8
Intellectual Property Registration 2020-A-2881

ISSN 0717-6686 (Series on Central Banking, Analysis, and Economic Policies)

Trade Exposure and the Evolution of Inflation Dynamics

Simon Gilchrist
New York University and NBER

Egon Zakrajšek
Federal Reserve Board

The Phillips curve—the relationship between price inflation and fluctuations in economic activity— is a central building block of economic models that allow for nominal rigidities and are relied upon by central banks around the world to gauge cyclical inflationary pressures and forecast inflation. The lack of deflationary pressures during the Great Recession and, more recently, the apparent lack of inflationary pressures during the recovery have brought into the forefront the question of whether this relationship still exists in the data.[1] More generally, the fact that inflation appears to have become less responsive to fluctuations in economic activity during the past couple of decades has been documented for the United States by Atkeson and Ohanian (2001), Roberts (2006), Mavroeidis and others (2001), and Blanchard (2016). This flattening of the Phillips curve appears to have occurred in other advanced economies as well;

We are grateful to Kimberly Bayard for her expert help with the industry-level industrial production data. We also thank Larry Ball, Mark Watson, Eduardo Zilberman (our discussant), Kurt Lewis, and conference participants for numerous helpful comments and suggestions. Tyler Pike and Gerardo Sanz-Maldonado provided excellent research assistance. The views expressed in this chapter are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

1. Recent work that studies the unusual inflation dynamics during the Great Recession and its aftermath in the United States and other advanced economies includes Stock and Watson (2010b), Ball and Mazumder (2011, 2018), Gordon (2013), Friedrich (2016), Berganza and others (2016), Miles and others (2017), Blanchard (2018), and Stock and Watson (2018).

Changing Inflation Dynamics, Evolving Monetary Policy edited by Gonzalo Castex, Jordi Galí, and Diego Saravia, Santiago, Chile. © 2020 Central Bank of Chile.

see Beaudry and Doyle (2000) for Canada, and Kuttner and Robinson (2010) for Australia, for example.

Reasons for the apparent attenuation of the relationship between inflation and resource utilization are often linked to the rise in globalization and an associated increase in the cross-border movement of goods, services, technology, labor, and capital since the 1990s.[2] The resulting greater openness of national economies implies that a greater share of an increase in domestic demand is satisfied through imports, rather than domestic production. In turn, this implies that changes in the domestic output gap will have a smaller effect on domestic marginal costs, thereby reducing the responsiveness of domestic inflation to fluctuations in domestic economic slack, while increasing the sensitivity of domestic inflation to foreign economic slack. Increased international trade also gives rise to a common component for inputs such as commodities, thus implying that local costs—and hence prices—become less sensitive to domestic economic conditions. Increased openness of labor markets is another factor that attenuates the link between inflation and fluctuations in economic activity at the local level.[3]

Although prominent in recent policy discussions, the evidence in favor of a weakening in the relationship between inflation and economic activity due to increased global economic integration is mixed. Ball (2006) and Ihrig and others (2010) argue that there is little evidence to suggest that increased international trade and other globalization factors have attenuated the relationship between inflation and economic slack in the United States. Borio and Filardo (2007), Auer and others (2017), and Zhang (2017), on the other hand, present evidence that globalization has indeed led to a decline in the sensitivity of inflation to domestic factors, arguing that the integration of China and other lower-cost producers in world production networks has increased competition, thereby inducing downward pressure on wages and import prices in the U.S. and other industrial countries.

2. Another hypothesis posits that the observed flattening of the Phillips curve over the past couple of decades is due to a lower frequency of price adjustment at the firm level, reflecting the significantly lower average inflation rate that has prevailed over that period (see Ball and others, 1988). Relatedly, some economists have hypothesized that firms and households have started to pay less attention to macroeconomic conditions when setting wages and prices because of a prolonged period of low and stable inflation—the so-called rational inattention hypothesis (see Sims, 2003; Pfajfar and Roberts, 2018).

3. See Bernanke (2007) for an overview of the various channels through which ongoing global economic integration can affect inflation dynamics.

Similarly, Forbes (2018) shows that global factors have played a more prominent role in determining U.S. inflation outcomes since the 1990s; these global factors, however, are primarily linked to the food and energy component of consumer prices and play a diminished, rather than an increased, role in explaining movements in the core measures of U.S. consumer price inflation.

In this chapter, we re-examine this "globalization" hypothesis by using both U.S. aggregate data on measures of inflation and economic slack and a rich panel data set containing producer prices, wages, output, and employment at a narrowly defined industry level. Industries in our data set are defined at the 6-digit North American Industry Classification System (NAICS) level, and the data on prices and output serve as the basis for the construction of the U.S. producer price index (PPI) published by the Bureau of Labor Statistics and the industrial production index (IPI) published by the Federal Reserve Board. We also measure international trade exposure at the industry level—albeit at a somewhat coarser level of aggregation (i.e., 4-digit NAICS)—by using information on exports, imports, and value-added output. Linking these trade exposures to industry-level prices, wages, employment, and production allows us to directly determine the extent to which the response of inflation to fluctuations in output differs systematically across industries that are more or less exposed to international trade.

We begin our analysis by examining the time-series relationships between inflation and fluctuations in economic activity. Specifically, we consider the extent to which the relationship between inflation and economic activity has evolved over time. We address this question by estimating the sensitivity of both producer and consumer price inflation to economic slack using 15-year rolling-window regressions, starting in the early 1960s. This evidence shows that this relationship has indeed weakened substantially over the past 30 years or so. Importantly, our findings are robust to using both headline inflation measures, as well as core measures of inflation that remove the direct influence of swings in the volatile food and energy prices. They are also robust to measuring economic slack using alternative concepts such as the output gap or the unemployment gap.

We next consider the responsiveness of inflation to economic activity at the industry level. In this analysis, we exploit the cross-sectional dimension of our data and can directly control for the common aggregate component driving both inflation and output. We again find that fluctuations in output are an important determinant of inflation—indeed, the estimated response of industry-level inflation to variation in industry-level output is very similar to that obtained from aggregate time-series data over comparable sample periods.

Using both the industry-level and aggregate time-series data, we then examine the extent to which an increase in trade exposure has altered the response of inflation to fluctuations in economic activity. Here again our findings are consistent across both aggregate and industry-level data. In the time-series dimension, the rising exposure of the U.S. economy to international trade can indeed explain a significant fraction of the overall decline in responsiveness of aggregate inflation to economic slack. This result is confirmed by our cross-sectional evidence, which shows that increased trade exposure significantly dampens the response of inflation to fluctuations in output across industries.

The results discussed above, however, do not directly determine the causal impact of fluctuations in economic activity on inflation. While demand shocks typically move inflation and output in the same direction, supply shocks have the opposite effect. Thus any attenuation in the observed response of inflation to output may be due to changes in the mixture of demand and supply shocks that the U.S. economy has experienced over our sample period. To address this issue, we examine the effect of identified aggregate shocks on industry-level outcomes. In this exercise, we are explicitly interested in the extent to which the intensity of trade exposure at the industry level alters the responsiveness of inflation to such aggregate shocks.

Given the high dimensionality of our industry-level data, we consider the dynamic effects of identified aggregate demand shocks using a Factor-Augmented Vector Autoregression (FAVAR) model, whereby the information contained in the large panel of industries is summarized by a small subset of common factors. By using this framework, we study how shocks to broad financial conditions—a specific form of aggregate demand shocks—affect the dynamics of price and wage inflation, output, and employment at the industry level. We focus on disturbances to the financial intermediation process because we view them as readily identified from economic and financial time-series data; moreover, there exists a large body of empirical evidence indicating that financial shocks account for a sizable fraction of the variability in output and inflation over the past 30 years.[4]

4. See Gilchrist and others, 2009; Gilchrist and Zakrajšek, 2012; Jermann and Quadrini, 2012; and Peersman and Wagner, 2014.

Using the FAVAR approach, we first document that an unanticipated tightening in broad financial conditions causes a significant decline in price and wage inflation, as well as in output and employment growth across all industries. Thus financial shocks deliver the positive comovement between inflation and output that is typically associated with shocks to aggregate demand. We then examine the extent to which responses of inflation and output to financial shocks differ across industries based on their trade exposure. Our results indicate that industries with a high trade exposure exhibit a substantially smaller response of inflation to movements in output induced by the unanticipated changes in financial conditions, relative to industries with a low trade exposure.

These differential dynamics occur despite the fact that the effect of such shocks on economic activity is virtually identical across these two industrial groupings. Translated into the movements of inflation—relative to output—our results imply that, in response to such shocks, inflation is about three times more responsive to changes in output for industries with a low trade exposure, compared with industries with a high trade exposure. These findings further support the argument that external trade exposure attenuates the link between inflation and fluctuations in economic activity and that increased international trade is indeed a likely reason behind the reduced responsiveness of aggregate inflation to economic slack that has been observed in the data since the early 1990s.

The remainder of the chapter is organized as follows. Section 1 considers the aggregate time-series relationship between inflation and economic activity and documents its evolution over time. Section 2 explores the relationship between inflation and economic activity using industry-level data and documents the extent to which differences in trade exposure across industries affect this relationship. Section 3 provides the FAVAR analysis, which shows how industry-level variables respond to financial shocks, as well as the extent to which these responses differ across industries depending on their exposure to international trade. Section 4 offers a brief conclusion.

1. Aggregate Phillips Curve

In this section, we establish some stylized facts about the relationship between inflation and economic slack by using aggregate time-series data, which serve as a useful benchmark for the subsequent industry-level analysis. While the vast literature on this topic has focused on consumer price inflation, we analyze inflation dynamics at both the producer and consumer levels; the focus on the former is especially important because movements in producer prices directly capture the price response of production units to changes in the underlying economic conditions.[5] In terms of data used in this analysis, the solid line in panel A of figure 1 shows the behavior of prices received by U.S. producers for their output, measured by the four-quarter percent change in the PPI for final demand, while the solid line in panel B shows the four-quarter percent change in the consumer price index (CPI), a measure of prices paid by urban consumers for a market basket of consumer goods and services. The slashed lines in each panel show the corresponding core inflation, which strips out items belonging to the food and energy categories from each headline price index.[6]

Clearly evident in the data are several distinct inflation regimes. First, the 1970s, a period of high and volatile inflation that was early on influenced importantly by the OPEC-induced increases in oil prices (Hamilton, 1983) and later by the Federal Reserve's overly optimistic view of the natural rate of unemployment (Orphanides and Williams, 2013). The early 1980s, in contrast, were marked by a gradual step-down in inflation reflecting the tightening of monetary policy under Chairman Volcker, who was determined to fight inflation and reverse the rise in inflation expectations (Lindsey and others, 2005). Since the mid-1980s, inflation—at both the producer and consumer levels—has stabilized in a narrow range around two percent, a pattern consistent with the well-anchored inflation expectations engendered by credible monetary policy, aimed at achieving the so-called dual mandate stipulated by the Full Employment and Balanced Growth Act of 1978.[7]

5. It is worth noting that the frequency of price changes in the narrow-item categories that are both in the consumer and producer micro-level price data sets collected by the Bureau of Labor Statistics are highly correlated (see Nakamura and Steinsson, 2008).

6. Each quarterly price index is constructed as a simple average of the monthly (seasonally adjusted) index values, and four-quarter percent changes are computed as 100 times the four-quarter log-difference of the specified series. In addition, while we use the CPI to measure inflation at the consumer level, all the results reported below are robust to using the personal consumption expenditure (PCE) price index to track the change in prices of goods and services purchased by the U.S. consumers throughout the economy.

7. More commonly known as the Humprey-Hawkins Act, the Full Employment and Balanced Growth Act established price stability and full employment as national economic policy objectives.

Figure 1. Producer and Consumer Price Inflation

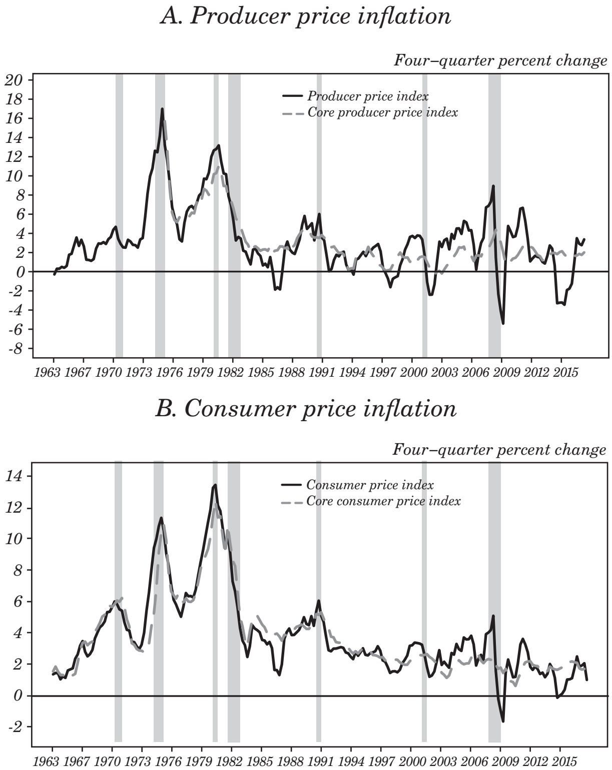

Source: Authors' calculations using data from the Bureau of Labor Statistics.
Note: All price indices are seasonally adjusted. The shaded vertical bars denote the NBER-dated recessions.

A striking way to illustrate how inflation is unresponsive to fluctuations in economic activity—in other words, how flat the Phillips curve is—is to focus on economic downturns. To that end, figure 2 examines the relationship between inflation and economic activity during the past five recessions, downturns in which supply-side disturbances—which cause inflation and economic activity to move in opposite directions—were arguably not the dominant factor. The first three panels of the figure depict the behavior of detrended prices two years before and after each NBER-dated cyclical peak since the early 1980s; the bottom right panel, by contrast, shows the corresponding dynamics of detrended real GDP, a simple measure of economic slack.

Figure 2. Inflation and Output in Recessions

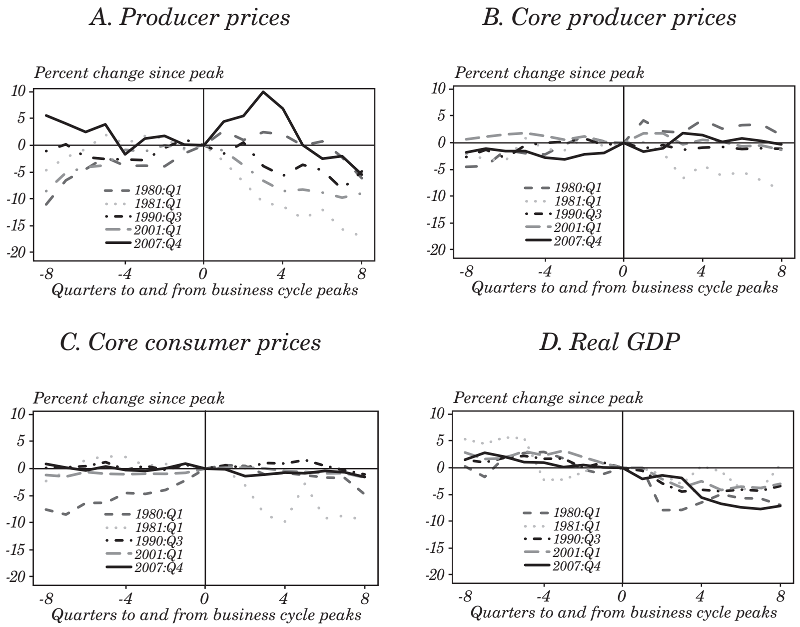

Source: Authors' calculations using data from the Bureau of Economic Analysis, the Bureau of Labor Statistics, and the Federal Reserve Board.
Note: The panels depict the behavior of various price measures and real GDP eight quarters before and eight quarters after the specified NBER-dated cyclical peak. All series are plotted as deviations from their respective stochastic trends, estimated using the Hamilton (2018) filter.

As shown in the top two panels, with the exception of the 2001 recession, producer prices—especially those that exclude the volatile food and energy components—showed virtually no deceleration during the past five economic downturns, relative to their trends. And even during the bursting of the tech bubble in 2001, the decline in both the headline and core PPI is due entirely to the plunge in producer prices in the immediate aftermath of the September 11 terrorist attacks—in October 2001, the Bureau of Labor Statistics reported that the PPI dropped almost 20 percent at an annual rate.[8] As shown in the bottom left panel, the resilience of inflation in response to the emergence of

8. It is also worth noting that the sharp increase in commodity prices prompted by the First Gulf War confounds the behavior of PPI inflation during the 1990 recession to some extent.

substantial economic slack is also evident at the consumer level. At the same time, as shown in the bottom right panel, real GDP declined markedly—relative to its trend—during these five episodes.

1.1 Baseline Estimates

To investigate more formally how the relationship between inflation and fluctuations in economic activity may have changed over time, we begin by estimating a standard Phillips curve specification, which expresses inflation as a linear function of expected inflation and a measure of economic slack. Specifically, letting lower-case variables denote variables in logarithms and defining $\Delta_h x_{t+h} = \frac{400}{h}(x_{t+h} - x_t)$, we estimate the following Phillips curve specification:

$$\Delta_{h+1} p_{t+h} = \mu + \lambda \mathrm{gap}_t + \sum\nolimits_{s=1}^{4} \phi_s \Delta p_{t-s} + \epsilon_{t+h}, \tag{1}$$

where p_t denotes the logarithm of a price index (i.e., PPI or CPI) and gap_t is a measure of economic slack, a degree of resource over- or under-utilization. Thus equation (1) posits a relationship between (annualized) inflation from quarter $t-1$ to quarter $t+h$ and a measure of economic slack in quarter t, while the lags of inflation $\Delta p_{t-s}, s = 1,\ldots,4$, are a proxy for expected inflation.[9,10] In this canonical formulation, the error term ϵ_{t+h} encompasses cost-push shocks—shock to commodity prices, for example—which are assumed to be uncorrelated with the contemporaneous economic slack.[11]

Figure 3 shows two measures of economic slack used in our analysis: the output gap and the unemployment gap. The output gap, denoted by $[y_t - y_t^*]$, is defined as (100 times) the logarithm of the ratio of real GDP to its estimate of potential, while the unemployment

9. See Gordon, 1982; Stock and Watson, 2009.

10. An alternative approach to using lagged values of inflation to capture expected inflation would be to use survey measures of expected inflation. However, as documented by Mankiw and others (2004), such survey measures do not appear to be consistent with either rational expectations or adaptive expectations used in specification (1).

11. It is worth noting that the presence of very low frequency variation in both the producer and consumer inflation rates (see figure 1) has the potential to confound the relationship between inflation and fluctuations in economic slack at the business cycle frequency, which is the primary interest of our analysis. To ensure that our baseline time-series results are not unduly affected by this low frequency variation, appendix A contains a robustness analysis in which all inflation series are expressed as deviations from their respective local means. As evidenced by those results, the main conclusions of this section are robust to this transformation of the data.

gap, denoted by $[U_t - U_t^*]$ corresponds to the unemployment rate less its estimate of the natural rate. The estimates of both the potential real GDP and the natural rate of unemployment are taken from the FRB/US model, a large-scale estimated general equilibrium model of the U.S. economy that has been in use at the Federal Reserve Board since 1996. While the definition of these two slack measures naturally produces series of the opposite sign, they paint a very similar picture of cyclical resource utilization over the last 50 years or so. One exception to this pattern has occurred during the past several years, a period in which the unemployment rate has moved below its natural rate, whereas the real GDP has yet to return to its potential.[12]

Table 1 present estimates of the coefficient λ for producer price inflation at horizons of one and four quarters (i.e., $h = 1, 4$), with panel A showing estimates of λ for headline PPI inflation and panel B showing estimates of λ for core PPI inflation; the corresponding estimates of λ for consumer price inflation—both headline and core—are shown in table 2.

Figure 3. Economic Slack

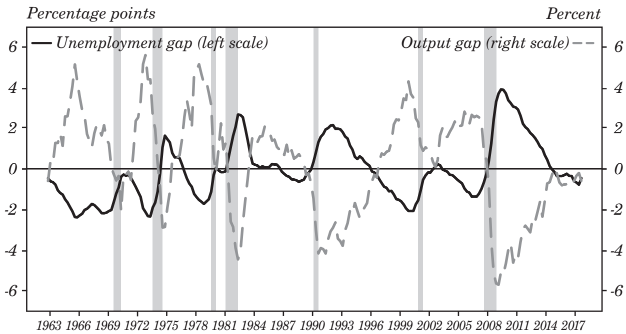

Source: Bureau of Economic Analysis, Bureau of Labor Statistics, and Federal Reserve Board.
Note: The output gap is defined as (100 times) the log-ratio of real GDP to its estimate of potential; the unemployment gap is defined as the civilian unemployment rate less its estimate of the natural rate. The shaded vertical bars denote the NBER-dated recessions.

12. Movements in the output gap can be interpreted as capturing fluctuations in real marginal cost, which microfounded models emphasize as a key determinant of inflation dynamics (see Roberts, 1995; Galí and Gertler, 2000; Galí and others, 2001; Sbordone, 2002; and Galí and others, 2007).

Table 1. Phillips Curve – Producer Price Inflation

Explanatory variables	h = 1		h = 4	
	(1)	(2)	(3)	(4)
A. Producer prices				
$[y_t - y_t^*]$	0.356**	-	0.414***	-
	(0.144)		(0.153)	
$[U_t - U_t^*]$	-	-0.396*	-	-0.469*
		(0.238)		(0.257)
Sum: inflation lags[a]	0.578***	0.600***	0.470***	0.495***
	(0.113)	(0.113)	(0.093)	(0.100)
sup W^b	15.185***	11.345**	33.370***	23.284***
	[81.Q2]	[91.Q4]	[80.Q3]	[93.Q4]
q_{LL}^c	-6.230	-5.375	-5.297	-4.325
Adj. R^2	0.360	0.333	0.392	0.343
B. Core producer prices				
$[y_t - y_t^*]$	0.186***	-	0.223***	-
	(0.056)		(0.067)	
$[U_t - U_t^*]$	-	-0.243**	-	-0.273**
		(0.105)		(0.131)
Sum: inflation lags[a]	0.776***	0.797***	0.730***	0.755***
	(0.071)	(0.076)	(0.071)	(0.081)
sup W^b	21.278***	18.00***	70.033***	39.261***
	[81.Q4]	[93.Q4]	[81.Q4]	[82.Q3]
q_{LL}^c	-9.554**	-7.550*	-6.304	-5.737
Adj. R^2	0.743	0.725	0.760	0.727

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4 for headline PPI (panel A), and 1974:Q1 to 2017:Q4 for core PPI (panel B). The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the annualized log-difference in the specified PPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap and $[U_t - U_t^*]$ = unemployment gap. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

[a] Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

[b] The Andrews (1993) sup-Wald statistic of the null hypothesis that there is no structural break in the coefficient on economic slack; the estimated break dates are reported in brackets below.

[c] The Elliott and Müller (2006) q_{LL} statistic of the null hypothesis that the coefficient on economic slack is constant over time.

Table 2. Phillips Curve – Consumer Price Inflation

Explanatory variables	*h* = 1		*h* = 4	
	(1)	*(2)*	*(3)*	*(4)*
A. Consumer prices				
$[y_t - y_t^*]$	0.258***	-	0.318***	-
	(0.075)		(0.084)	
$[U_t - U_t^*]$	-	-0.321***	-	-0.380***
		(0.120)		(0.128)
Sum: inflation lags[a]	0.779***	0.795***	0.690***	0.709***
	(0.066)	(0.070)	(0.068)	(0.077)
sup W^b	34.118***	28.008***	70.231***	44.548***
	[83.Q1]	[83.Q2]	[83.Q1]	[83.Q1]
$q_{LL}{}^c$	-8.199*	-6.892	-6.347	-4.986
Adj. R^2	0.657	0.635	0.676	0.632
B. Core consumer prices				
$[y_t - y_t^*]$	0.176***	-	0.265***	-
	(0.044)		(0.060)	
$[U_t - U_t^*]$	-	-0.263***	-	-0.364***
		(0.079)		(0.107)
Sum: inflation lags[a]	0.868***	0.875***	0.787***	0.797***
	(0.056)	(0.060)	(0.065)	(0.074)
sup W^b	38.828***	38.420***	112.255***	66.416***
	[83.Q1]	[83.Q1]	[83.Q1]	[83.Q1]
$q_{LL}{}^c$	-8.639***	-6.259	-7.278*	-6.056
Adj. R^2	0.802	0.794	0.778	0.750

Source: Authors' calculations.
Note: Sample: 1962:Q2 to 2017:Q4. The dependent variable in each Phillips curve specification is $\Delta_{h+1}p_{t+h}$, the annualized log-difference in the specified CPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, and $[U_t - U_t^*]$ = unemployment gap. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.
[a] Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.
[b] The Andrews (1993) sup-Wald statistic of the null hypothesis that there is no structural break in the coefficient on economic slack; the estimated break dates are reported in brackets below.
[c] The Elliott and Müller (2006) q_{LL} statistic of the null hypothesis that the coefficient on economic slack is constant over time.

According to the entries in panel A of table 1, fluctuations in economic slack have a significant effect on the subsequent behavior of producer prices. A decrease in resource utilization of one percentage point in quarter *t*—that is, a decline in the output gap or a rise in the unemployment gap of that magnitude—is estimated to reduce annualized headline producer price inflation over the next several quarters about 40 basis points. The corresponding estimates for core PPI inflation shown in panel B are about one-half as large as those reported in panel A, though the estimates are significant in both economic and statistical terms. As shown in table 2, economic slack is also a significant determinant of consumer price inflation. In that case, a decrease in resource utilization of one percentage point is estimated to shave off about 25 basis points from annualized CPI inflation over the subsequent few quarters.

As a first pass on the question of whether the relationship between economic slack and subsequent inflation may have changed over the past 50 years or so, we report results of two statistical tests. The first is the well-known Andrews (1993) test of a structural break—at an unknown date— in the coefficient λ. The second is the Elliott and Müller (2006) test of stability of the coefficient λ, which encompasses diverse forms of parameter instability—from relatively rare (including a single break) to frequent small breaks, persistent temporal parameter variation, and breaks occurring with a regular pattern.[13] This statistical analysis, however, yields a mixed picture. Turning first to producer prices (table 1), the Andrews (1993) test provides strong evidence of a structural break in λ, with the point estimate of a break date generally falling in the early 1980s, a result consistent with that of Roberts (2006). The evidence of parameter instability from the Elliott and Müller (2006) test, in contrast, is considerably weaker. A similar picture emerges when we look at consumer prices (table 2). Here again, the Andrews (1993) test strongly suggest a structural break in λ that occurred in the early 1980s, whereas the results from the Elliott and Müller (2006) test are far less conclusive.

All told, the results reported in tables 1 and 2 clearly indicate an important role—in both economic and statistical terms—for economic

13. In both tests, the null hypothesis is that the coefficient λ is stable over the sample period. The alternative in the Andrews (1993) test is that $\lambda = \lambda_1$ for $t = 1,2,...,\tau - 1$ and $\lambda = \lambda_2$ for $t = \tau, \tau + 1, ... T$, where τ is the unknown (single) break date. The alternative in the Elliott and Müller (2006) test is $\lambda = \lambda_t$, where the time variation in the parameter λ_t is unspecified and can take on a variety of forms.

slack as a determinant of cyclical inflation dynamics. Nevertheless, empirical Phillips curves of the type given by equation (1) predicted a significantly greater downward pressure on inflation—if not outright deflation—during the Great Recession than was actually realized. Economists have advanced a number of hypotheses to explain this case of "missing deflation." A prominent hypothesis that received a lot of attention in policy circles argues that the Federal Reserve's credibility has led businesses and households to discount inflation outcomes that fall outside the narrow range bracketing the Federal Open Market Committee's inflation target of two percent; this anchoring of agents' expectations has—through the standard expectational effects—prevented actual inflation from falling significantly below that level.[14]

Another frequently cited hypothesis posits that the relevant measure of economic slack in empirical Phillips curves is not the overall unemployment rate gap, but rather the short-term unemployment rate.[15] Compared with the former, this latter indicator of slack increased notably less during the Great Recession and has also returned more quickly to its pre-recession levels, thus providing substantially less deflationary impetus. And although it has proven difficult to identify structural changes in the economy that could account for the diminished sensitivity of inflation to the level of unemployment, a number of economists have singled out the apparent flattening of the Phillips curve as an important reason for the fact that the U.S. economy did not experience a Fisherian debt-deflation spiral during the 2008–2009 global financial crisis.[16, 17]

1.2 Time-Varying Estimates

In light of the above discussion and the results reported in tables 1 and 2, it seems clear that a further investigation in the time-varying nature of the relationship between inflation and economic slack is warranted. As a simple and relatively straightforward way to

14. See Bernanke, 2010; Yellen, 2013.

15. Underlying this argument is the idea that workers who have been unemployed for a relatively short time are the relevant margin for wage adjustment. The longer-term unemployed, by contrast, do not put much downward pressure on wages because these potential workers are disconnected from the labor market (see Stock and Watson, 2010b; Gordon, 2013; Krueger and others, 2014).

16. See Ball and Mazumder, 2011; Simon and others, 2013.

17. Gilchrist and Zakrajšek (2016); Gilchrist and others (2017), in contrast, emphasize how the interaction of financial distress and customer markets attenuated deflationary pressures during the Great Recession.

consider the possibility of time variation in the coefficient λ—as well as in other parameters of the standard Phillips curve—we re-estimate specification (1) by using a 15-year rolling window. We then plot the time-varying coefficient on the specified measure of economic slack, along with its 95-percent confidence interval. To conserve space, we focus on the Phillips curve specifications for inflation at the horizon of four quarters (i.e., $h=4$). The resulting time-varying estimates of the coefficient λ, for both the headline and core PPI inflation, are shown in figure 4, with panel A showing the time-varying sensitivity to the output gap and panel B showing the time-varying sensitivity to the unemployment gap; the comparable estimates for CPI inflation are shown in figure 5.[18]

The left chart in panel A of figure 4 shows the evolution of the response of headline PPI inflation to the output gap. In the early part of the sample, the estimates of λ are greater than one and significantly different from zero, according to the 95-percent confidence intervals. Starting in the mid-1980s, however, these estimated sensitivities begin to decline steadily before stabilizing in the late 1990s. From then onward, the estimates of λ fluctuate in a fairly narrow range between zero and 0.5, though for most of this latter sample period, one would not reject the hypothesis that the coefficient on the output gap is statistically different from zero.

The left chart in panel B shows the evolution of the response of headline PPI inflation to the unemployment gap. We observe roughly the same general pattern in this case. The estimates of λ start out negative and large in economic terms, as well statistically different from zero, according to the conventional significance levels. Once the late 1980s enter the sample period, however, the estimates begin to converge rapidly to zero. The estimated response of headline PPI inflation to the unemployment gap then remains around zero for the remainder of the sample period.

18. The convention is that the data point labeled "1994:Q4," for example, represents an estimate based on the 1980:Q1–1994:Q4 estimation window. For both the headline producer and consumer price inflation, as well as for the core consumer price inflation, our sample period—allowing for lags—starts in 1962:Q2, so that the rolling-window estimates begin in 1978:Q1 and run through 2017:Q4, the end of our sample period. Core producer prices, by contrast, start in 1974:Q1, which implies that the first rolling-window estimates—again allowing for lags—become available in 1990:Q1.

Figure 4. Time-Varying Coefficient on Economic Slack
(Phillips Curve – Producer price inflation)

A. Economic slack: output gap

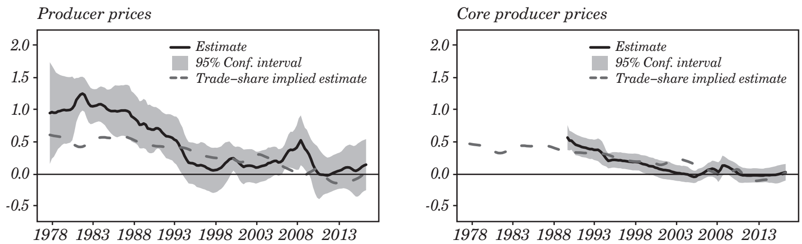

B. Economic slack: unemployment gap

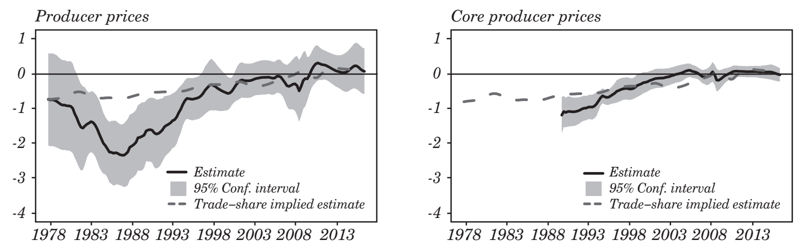

Source: Authors' calculations.
Note: The dependent variable in each Phillips curve specification is $\Delta_5 p_{t+4}$, the annualized log-difference in the specified PPI from date $t-1$ to date $t+4$. The solid line in each panel depicts the time-varying coefficient on the specified measure of economic slack estimated using a 60-quarter moving window; the dashed lines depict the corresponding time-varying coefficients implied by specifications (3) and (4) in table 3 (see notes to the table and the text for details).

The corresponding right charts of figure 4 trace out the estimated sensitivities of core PPI inflation to the output gap (panel A) and the unemployment gap (panel B). Although the sample begins later in this instance, the rolling-window estimates of the coefficient λ in the Phillips curve for core PPI inflation are much more precisely estimated than their counterparts for headline inflation. The estimates of λ for the output gap begin at about 0.5 for the sample that extends from the mid-1970s to the end of the 1980s and then decline monotonically to zero as the sample period moves forward; in fact, the estimate of λ based on the last 15 years of available data implies a sensitivity of core PPI inflation to the output gap that is economically and statistically indistinguishable from zero. The time-series pattern of coefficients on the unemployment gap is very similar: The estimates of λ start out negative, large in absolute value, and are precisely estimated and then converge to zero by the end of the 1990s.

Figure 5. Time-Varying Coefficient on Economic Slack
(Phillips Curve – Consumer price inflation)

A. Economic slack: output gap

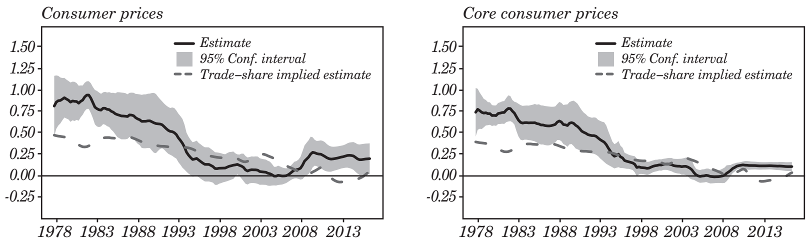

B. Economic slack: unemployment gap

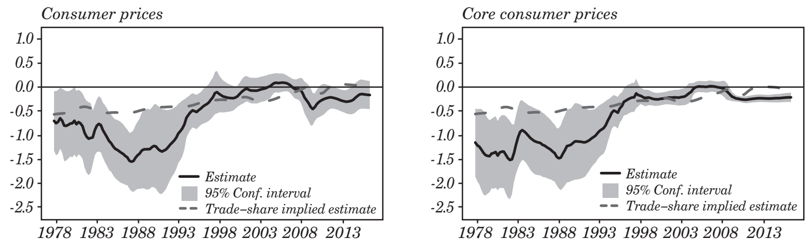

Source: Authors' calculations.
Note: The dependent variable in each Phillips curve specification is $\Delta_5 \, p_{t+4}$, the annualized log-difference in the specified CPI from date $t-1$ to date $t+4$. The solid line in each panel depicts the time-varying coefficients on the specified measure of economic slack estimated using a 60-quarter moving window; the dashed lines depict the corresponding time-varying coefficients implied by specifications (3) and (4) in table 4 (see notes to the table and the text for details).

Figure 5 shows the time-varying coefficient estimates on economic slack for both the headline and core measures of CPI inflation. As before, panel A shows coefficient estimates on the output gap, while the corresponding estimates for the unemployment gap are shown in panel B. Consistent with the full-sample estimates of λ reported in panel A of table 2, the time-varying coefficient estimates of the response of inflation to the output gap for headline CPI inflation are very similar to those for core inflation, both in terms of their magnitude and their evolution over time. They also show a pattern similar to that shown in figure 4: The estimates of λ are positive, economically and statistically significant in the early part of the sample, and then begin to decline sharply once the 1990s enter the estimation window. In contrast to the estimated response coefficients for PPI inflation shown in panel A of figure 4, the sensitivity of CPI inflation to the output gap is estimated to have increased notably at the end of our sample period. That said, these late-sample estimates of λ, though statistically different from zero, are only about one-fourth of those estimated during the early part of our sample period. According to panel B of figure 5, these patterns are robust to using the unemployment gap, rather than the output gap, as a measure of economic slack.

While there are a variety of phenomena that may help explain the declining sensitivity of aggregate inflation to fluctuations in economic activity, we are specifically interested in the extent to which increased globalization and trade may have contributed to the flattening of the Phillips curve. The notion that increased trade may help account for such changes is consistent with the rising trade intensity in the United States—defined as the sum of exports and imports relative to GDP—shown in figure 6. According to this metric, the trade intensity of the U.S. economy has risen by nearly a factor of three over the past 50 years or so.

To test the hypothesis that increased trade intensity of the U.S. economy may have contributed to the observed decline of the sensitivity of inflation to economic slack, we estimate the following variant of our baseline Phillips curve specification:

$$\Delta_{h+1}\mathrm{p}_{t+h} = \mu + \lambda_1 \mathrm{gap}_t + \lambda_2\left[\mathrm{gap}_t \times \mathrm{TrdShr}_{t-1}\right] + \sum_{s=1}^{4} \phi_s \Delta p_{t-s} + \epsilon_{t+h}, \tag{2}$$

where TrdShr_t denotes an eight-quarter trailing moving average of the U.S. trade share shown in figure 6.[19] The resulting coefficient estimates of λ_1 and λ_2 for PPI inflation are reported in table 3, while those for CPI inflation are reported in table 4.

19. The Phillips curve specification (2) is similar to that used by Ball (2006), except that it does not include the "smoothed" trade share, TrdShr_{t-1}, as a separate explanatory variable; the inclusion of this term, however, had no material effect on any of the results reported below. Note also that appendix A contains results from the estimation, which controls for the slow-moving changes in the average inflation rate over our sample period; again, those results are qualitatively and quantitatively similar to those reported in the main text.

Figure 6. U.S. Trade Share

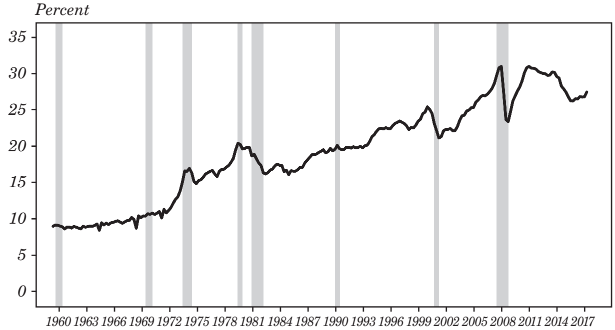

Source: Bureau of Economic Analysis.
The trade share is defined as the sum of the nominal value of U.S. imports and exports, expressed as a percent of nominal GDP. The shaded vertical bars denote the NBER-dated recessions.

According to the entries reported in panel A of table 3, the coefficient on the interaction term between the output gap and the trailing moving average of the U.S. trade share is negative— though not statistically different from zero—at the one-quarter horizon (column 1) and negative and marginally significant at the four-quarter horizon (column 3). Similarly, the interaction effect between the unemployment gap and trade share is positive and imprecisely estimated for $h = 1$, whereas the coefficient on the interaction term for $h = 4$ is positive and statistically different from zero at the 10-percent significance level. On balance, therefore, the evidence based on headline PPI inflation does not seem to support strongly the hypothesis that increased trade exposure of the U.S. economy can explain the decline in the sensitivity of inflation to fluctuations in economic activity.

As shown in panel B, however, the corresponding estimates for core PPI inflation paint a very different picture. The coefficients on the interaction terms between the output gap and trade share are negative and quite precisely estimated at both the one- and four-quarter horizons (columns 1 and 3). And similarly, the coefficients on the interaction terms between the unemployment gap and trade share are negative and statistically different from zero for both $h = 1$ and $h = 4$ (columns 2 and 4). Moreover, these estimates are economically meaningful. At the four-quarter horizon, they imply that when the trade share was at the 5th percentile of its distribution, the sensitivity of core PPI inflation to the output gap was 0.766 (std. error = 0.186) and −1.262 (std. error = 0.416) when the unemployment gap is used to gauge the degree of resource utilization in the economy; the corresponding estimates at the 95th percentile, in contrast, are −0.058 (std. error = 0.122) and 0.092 (std. error = 0.188), respectively. In other words, these results indicate a robust relationship between the rising trade share and the diminished sensitivity of core PPI inflation to fluctuations in economic activity.

In table 4, we report the estimates of coefficients λ_1 and λ_2 for Phillips curve specifications involving headline (panel A) and core (panel B) CPI inflation. These results again imply an economically large and statistically significant reduction in the responsiveness of inflation to economic slack as the trade share rises. This is true for both the headline and core measures of CPI inflation and holds at both the one- and four-quarter horizons. Moreover, the strong attenuation of the response of CPI inflation to economic slack is robust to using either the output or the unemployment gap as a gauge of cyclical resource utilization.

To summarize how the increasing exposure of the U.S. economy to international trade over the past 50 years affected the responsiveness of inflation to fluctuations in economic activity, we calculate the time-series evolution of the response coefficients associated with economic slack, as implied by the estimates of coefficients λ_1 and λ_2 reported in tables 3 and 4 and the trajectory of the U.S. trade share shown in figure 6. We then plot these estimates, as dashed lines in figures 4 and 5, next to their corresponding time-varying estimates based on the 15-year rolling window. The comparison of solid and dashed lines in the panels of these two figures shows that this specific parametrization of the time-varying slope of the Phillips curve—a simple interaction between the trade share and economic slack—can account for about one-half of the decline in the responsiveness of PPI and CPI inflation to economic slack observed over the past 50 years. In addition, this interaction effect captures remarkably well the attenuation in the response of core PPI inflation to changes in economic slack that we observe during the latter part of the sample period.

Table 3. Phillips Curve and the Trade Share – Producer Price Inflation

Explanatory variables	*h* = 1 (*1*)	*h* = 1 (*2*)	*h* = 4 (*3*)	*h* = 4 (*4*)
A. Producer prices				
$[y_t - y_t^*]$	0.968**	-	1.459**	-
	(0.449)		(0.584)	
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.030	-	-0.052*	-
	(0.023)		(0.07)	
$[U_t - U_t^*]$	-	-1.349**	-	-1.759**
		(0.674)		(0.822)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.047	-	0.063*
		(0.033)		(0.038)
Sum: inflation lags[a]	0.583***	0.604***	0.479***	0.500***
	(0.111)	(0.108)	(0.086)	(0.093)
Adj. R^2	0.369	0.342	0.440	0.371
B. Core producer prices				
$[y_t - y_t^*]$	0.903*	-	1.125***	-
	(0.309)		(0.305)	
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.031**	-	-0.040***	-
	(0.014)		(0.014)	
$[U_t - U_t^*]$	-	-1.591***	-	-1.851***
		(0.596)		(0.648)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.056**	-	0.065**
		(0.024)		(0.026)
Sum: inflation lags[a]	0.751***	0.790***	0.698***	0.747***
	(0.067)	(0.067)	(0.058)	(0.068)
Adj. R^2	0.762	0.742	0.794	0.754

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4 for headline PPI (panel A), and 1974:Q1 to 2017:Q4 for core PPI (panel B). The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the annualized log-difference in the specified PPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap; $[U_t - U_t^*]$ = unemployment gap, and TrdShr_{t-1} = eight-quarter (trailing) moving average of the trade share. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

[a] Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

Table 4. Phillips Curve and the Trade Share – Consumer Price Inflation

Explanatory variables	*h* = 1 (1)	*h* = 1 (2)	*h* = 4 (3)	*h* = 4 (4)
A. Consumer prices				
$[y_t - y_t^*]$	0.728***	-	1.093***	-
	(0.231)		(0.311)	
$[y_t - y_t^*] \times \mathrm{TrdShr}_{t-1}$	-0.023**	-	-0.038***	-
	(0.011)		(0.013)	
$[U_t - U_t^*]$	-	-0.997***	-	-1.282***
		(0.353)		(0.425)
$[U_t - U_t^*] \times \mathrm{TrdShr}_{t-1}$	-	0.033**	-	0.044**
		(0.016)		(0.018)
Sum: inflation lags[a]	0.801***	0.812***	0.727***	0.734***
	(0.063)	(0.067)	(0.063)	(0.075)
Adj. R^2	0.670	0.645	0.721	0.656
B. Core consumer prices				
$[y_t - y_t^*]$	0.526***	-	0.918***	-
	(0.133)		(0.200)	
$[y_t - y_t^*] \times \mathrm{TrdShr}_{t-1}$	-0.017***	-	-0.032***	-
	(0.005)		(0.008)	
$[U_t - U_t^*]$	-	-0.802***	-	-1.197***
		(0.237)		(0.340)
$[U_t - U_t^*] \times \mathrm{TrdShr}_{t-1}$	-	0.026***	-	0.040***
		(0.009)		(0.013)
Sum: inflation lags[a]	0.891***	0.896***	0.831***	0.851***
	(0.053)	(0.058)	(0.061)	(0.072)
Adj. R^2	0.811	0.803	0.815	0.874

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4. The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the annualized log-difference in the specified CPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, $[U_t - U_t^*]$ = unemployment gap, and TrdShr_{t-1} = eight-quarter (trailing) moving average of the trade share. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

[a] Sum of coefficients on Δp_{t-s}, $s = 1,\ldots,4$.

2. Industry-Level Phillips Curve

The combination of a rising trade share with the concomitant decline in the responsiveness of aggregate inflation to fluctuations in economic activity provides suggestive evidence that the observed flattening of the Phillips curve is at least partly due to increased trade intensity of the U.S. economy. The variation used to estimate this effect, however, relies solely on the secular increase in the U.S. trade share over the past 50 years or so and moreover does not fully explain the substantial reduction in the estimated slope of the aggregate Phillips curve. To provide a more thorough analysis of this phenomenon, we now turn to industry-level data, where we can exploit variation in trade shares across industries to test whether a differential trade exposure influences the sensitivity of inflation to economic slack.

2.1 Data Sources and Methods

To construct the panel data set used in this analysis, we utilize the most detailed (i.e., 6-digit NAICS) industry-level PPIs published by the Bureau of Labor Statistics, which we merge with the corresponding industry-level data on industrial production—a measure of output—constructed by the Federal Reserve.[20] The resulting data set covers all 6-digit NAICS industries—excluding those in the Utilities sector (i.e., 2-digit NAICS code 22)—that are used to produce both the producer price and industrial production indices for the U.S. economy. The industry-level price and production data are available at the monthly frequency, and we convert them to quarterly frequency by simply averaging the values of each index over the three months of each quarter.

The industry-level price and production data are available starting in the early 1970s. However, the data are not available for every industry from the beginning—that is, the panel is unbalanced—and there is an especially large expansion in the number of industries covered that occurred in the mid-1980s. To capture this broad array of industries, we thus begin our sample in 1984:Q1. All told, our

20. IPIs are not available for the full set of 6-digit NAICS industries. At such a fine level of disaggregation, there are in some cases an insufficient number of production units to construct a meaningful estimate of the index. In those instances, the staff at the Federal Reserve Board aggregates the underlying data across several of such closely related industries. In our matching algorithm, we assigned such industrial production data to all the 6-digit industries in the index.

unbalanced panel includes price and production data for 319 industries at the 6-digit NAICS level, covering the period from 1984:Q1 to 2017:Q4. We complement these industry-level data on output and prices with the corresponding data on wages and employment from the Quarterly Census of Employment and Wages (QCEW), a data collection program that publishes a quarterly count of employment, total wages, and average weekly wages per employee, reported by companies covering more than 95 percent of U.S. jobs. The QCEW data, however, are available only starting in 1990:Q1. We thus also consider a more in-depth analysis by using a balanced panel of 185 industries for which all of these variables are available over the 1990:Q1–2017:Q4 period.[21]

To measure trade exposure at the industry level, we rely on the annual (nominal) import and export data, which are made available by the Center for International Data at the University of California Davis and cover the period from 1972 to 2006.[22] The data provided are disaggregated by country (source for imports and destination for exports) and Schedule B number. These data were first aggregated to the total annual imports and exports at the industry level by using the 5-digit Standard Industrial Classification (SIC) codes. The annual (nominal) imports and exports for the 2007–2017 period were obtained from the U.S. Census Bureau's USA Trade Online database and are available at the 10-digit Harmonized System Code (HTS) level. By using various crosswalks, all of these data had to be first mapped to industries at the 6-digit NAICS level. At such a fine level of disaggregation, however, there are numerous missing industry/year observations. Accordingly, we aggregated trade date to the 4-digit NAICS level. The resulting panel data set was then merged with the annual 4-digit NAICS data on (nominal) value-added output provided by the Bureau of Economic Analysis; these data were then used to calculate trade exposure—the sum of imports and exports relative to output—for each 4-digit NAICS industry.

21. The industry-level data exhibit significant seasonal fluctuations. Accordingly, we filtered all industry-level variables by using the Census Bureau's X12 seasonal adjustment procedure—thus all of our growth rates (i.e., log differences) are constructed by using seasonally adjusted level series. To ensure that our results were not influenced by a small number of extreme observations, all quarterly growth rates were winsorized at the 0.5th and 99.5th percentiles.

22. These data were assembled by Robert Feenstra through the project funded by a grant from the National Science Foundation to the NBER; see http://cid.econ.ucdavis.edu/usix.html for further details.

Figure 7. Industry-Specific Producer Prices and Industrial Production

A. Producer prices

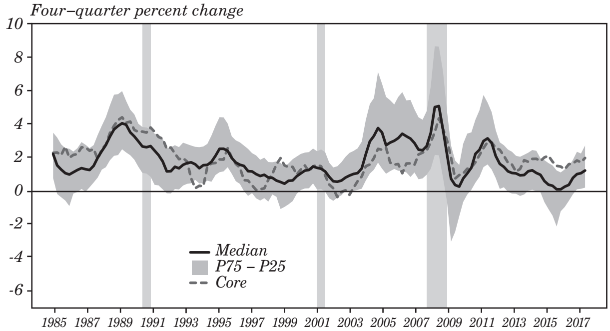

B. Industrial production

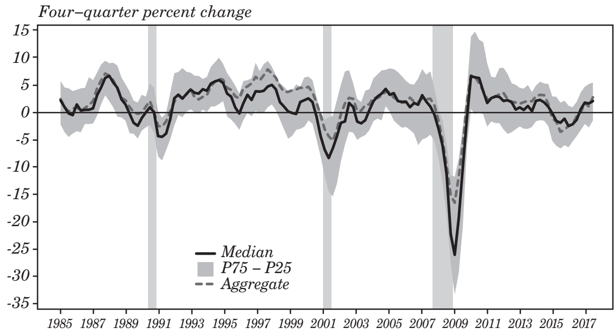

Source: Authors' calculations using data from the Bureau of Labor Statistics and the Federal Reserve Board.
Note: All industry-specific producer price and industrial production indices are seasonally adjusted. The solid lines depict the cross-sectional medians of the specified series, while the shaded bands depict the corresponding interquartile (P75–P25) ranges. For comparison purposes, the dashed line in panel A shows the four-quarter log-difference of the published core PPI, while the dashed line in panel B shows the four-quarter log-difference of the published IPI. The shaded vertical bars denote the NBER-dated recessions.

The solid line in panel A of figure 7 shows the time-series evolution of the (unweighted) cross-sectional median of the four-quarter percent change in PPI inflation across 319 industries in our unbalanced panel, while the shaded band depicts the corresponding (unweighted) interquartile range. The dashed line, in contrast, shows the four-quarter percent change based on the published core PPI. In panel B, the solid line and the shaded band depict the same moments of the four-quarter percent change in industrial production across the same set of industries, while the dashed line shows the corresponding growth rate of total industrial production. As evidenced by the shaded bands, the inflation rates and output growth vary significantly across industries. At the same time, the time-series fluctuations in the two medians closely match dynamics of their corresponding aggregates, an indication that our industry-level data are representative of the economy as a whole.

2.2 Baseline Estimates

To analyze the relationship between producer prices and economic activity at the industry level, we reformulate our baseline Phillips curve specification given by equation (1) above to accommodate the cross-sectional aspect of the industry-level data. Specifically, we estimate the following panel-data version of the Phillips curve:

$$\Delta_{h+1} p_{i,t+h} = \lambda \text{gap}_{it} + \sum_{s=1}^{4} \phi_s \Delta p_{i,t-s} + \mu_i + \eta_t + \epsilon_{i,t+h}, \tag{3}$$

where $p_{i,t}$ denotes the logarithm of the PPI for industry i in quarter t and gap_{it} is a measure of economic slack (or activity) in that industry. This specification also allows for an industry-specific intercept μ_i that is estimated using industry fixed effects and a full set of time dummies—denoted by η_t, $t = 1, 2, \ldots, T$—that capture variation in common factors across industries. To measure the extent of resource utilization within each industry, we compute the "industrial production" gaps for each industry—denoted by $[q_{it} - \tilde{q}_{it}]$—as (100 times) the log-deviation of IPI (q_{it}) from its stochastic trend ($\tilde{q}_{it}$), where the latter is estimated by using the Hamilton (2018) filter. As an alternative, we also consider a simple four-quarter log-difference of industrial production, denoted by $\Delta_4 q_{it}$.

Columns (1) and (2) of table 5 report estimates of the Phillips curve at the four-quarter horizon (i.e., $h = 4$) for the full sample of industries from 1984:Q1 to 2017:Q4. Columns (3) and (4), on the other hand, provide comparable estimates for a subsample based on the 1998:Q1–2017:Q4 period, which corresponds to the time period in which the slope of the aggregate Phillips curve for PPI inflation is estimated to have stabilized near zero (figure 4).[23] According to

23. Because our panel data set is unbalanced, the coefficient estimates are not strictly comparable across these two periods.

columns (1) and (2), fluctuations in economic activity—measured either as deviations of industrial output from its trend or as four-quarter growth in output—are important determinants of producer price inflation at the industry level. Although precisely estimated, the economic magnitudes of these coefficients are fairly small: An increase in the industrial production gap of 10 percentage points in quarter t—an increase of a bit less than one standard deviation—is estimated to boost annualized PPI inflation from quarter $t-1$ to $t+4$ a mere 15 basis points; the same-sized increase in the four-quarter growth of industrial output leads to a rise in PPI inflation of about a quarter of a percentage point over the same horizon.

While small in economic terms, these estimates are nonetheless broadly consistent with those based on the aggregate time-series data. For example, the coefficient on the output gap in the aggregate Phillips curve for core producer prices estimated over the 1984:Q1–2017:Q4 period is 0.063 (std. error = 0.054), while the corresponding coefficient estimate based on the 1998:Q1–2017:Q4 sample is –0.021 (std. error = 0.056).

Table 5. Industry-Level Phillips Curve

Explanatory variables	*Sample: 1984:Q1-2017:Q4*		*Sample: 1998:Q1-2017:Q4*	
	(1)	*(2)*	*(3)*	*(4)*
$[q_{it} - \tilde{q}_{it}]$	0.014**	-	0.020***	-
	(0.006)		(0.007)	
$\Delta_4 q_{it}$	-	0.027***	-	0.030***
		(0.008)		(0.008)
Sum: inflation lags[a]	-0.057*	-0.054*	-0.082**	-0.079**
	(0.031)	(0.030)	(0.037)	(0.037)
Adj. R^2	0.220	0.222	0.246	0.246
Panel dimensions				
No. of industries	319	319	319	319
Avg. T_i (quarters)	95.6	95.8	60.4	60.5
No. of observations	30,512	30,566	19,266	19,287

Source: Authors' calculations.

Note: The dependent variable in each Phillips curve specification is $\Delta_5 p_{i,t+4}$, the annualized log-difference in industry-specific PPI from date $t-1$ to date $t+4$. Explanatory variables: $[q_{it} - \tilde{q}_{it}]$ = industry-specific industrial production gap, and $\Delta_4 q_{it}$ = log-difference in industry-specific IPI from date $t-4$ to date t. All specifications include industry and time fixed effects and lags 1,...,4 of Δp_{it} (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are clustered across industries and time, according to Cameron and others (2011): * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

[a] Sum of coefficients on $\Delta p_{i,t-s}$, $s = 1,...,4$.

In other words, the slope of the aggregate Phillips curve for core PPI inflation is statistically indistinguishable from zero over this period. It is also worth noting that the estimates of coefficients on economic activity reported in table 5 are remarkably stable across the two sample periods. Thus, the industry-level estimates of the response of PPI inflation to fluctuations in industrial output do not show the same kind of attenuation pattern that we estimate by using the aggregate time-series data.

2.3 The Role of the Trade Share

With these results in hand, we now turn to the question of whether differences in external trade exposure across industries influence the sensitivity of PPI inflation to economic slack. A straightforward way to test this hypothesis would be to estimate our baseline industry-level Phillips curve given in equation (3) on a sample of "low" trade intensity industries and compare the results with those based on a sample of "high" trade intensity industries. However, to make a statement of whether differences in trade exposure across industries matter in the aggregate, we must specify some kind of a weighting scheme.[24] Unfortunately, the value of shipments, which would provide an economically most sensible weighting scheme for the industry-specific inflation rates, is not available at the 6-digit NAICS level. As an alternative, we rely on the QCEW employment data and aggregate the industry-specific PPI inflation rates using the industry-specific *average* employment shares as weights. Because the employment data are available only starting in 1990:Q1, we restrict the analysis to the balanced panel of 185 industries, which ensures that our aggregation scheme is not affected by changes in the composition of industries over time.

To gauge the reasonableness of our aggregation scheme, the solid line in figure 8 shows the time-series evolution of a weighted cross-sectional average of four-quarter PPI inflation rates across the 185 industries in our balanced panel, while the dashed and dashed-dotted lines show the corresponding behavior of the headline and core producer price inflation, respectively. As can be seen from the figure,

24. Note that in the above regression analysis, each industry received an equal weight. As such, the results in table 5 may not provide an accurate picture of the aggregate relationship between inflation and economic slack that is central to our analysis.

our employment-weighted aggregate inflation broadly tracks a mix of the headline and core PPI inflation. It is clearly more cyclical than the core inflation and somewhat less cyclical than the headline inflation. Importantly, this aggregation exercise gives us confidence that an employment-weighted version of the 6-digit industry data captures the cyclical variation that we see in other time-series aggregates and hence provides a meaningful laboratory from which one can infer aggregate phenomena from the industry-level estimates.

We use the balanced panel—with the associated average employment shares—to examine the extent to which the responsiveness of inflation to fluctuations in economic activity differs with the degree of trade intensity across industries. As noted above, we split our sample of 185 industries into two groups, based on whether their average trade share is above or below 5 percent. This cutoff corresponds to the median of the industry-specific average trade shares, weighted by the industry-specific average employment shares, and implies that the low and high trade intensity industry groups account for about one-half each of total employment in our balanced panel.

Figure 8. Industry vs. Aggregate Producer Price Inflation

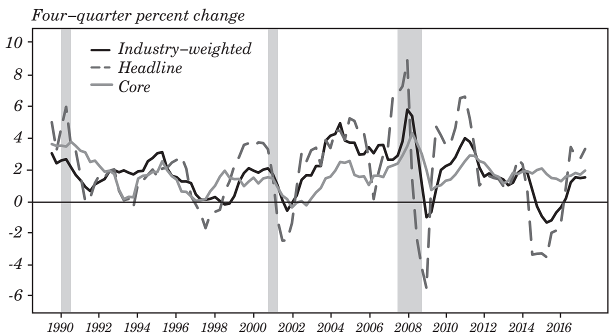

Source: Authors' calculations using data from the Bureau of Labor Statistics.
Note: The solid line depicts a cross-sectional weighted average of producer price inflation across 185 industries in the balanced panel, with weights equal to the corresponding average industry-specific employment shares. The dashed grey line depicts the headline (core) producer price inflation. The shaded vertical bars denote the NBER-dated recessions.

Table 6. Industry-Level Phillips Curve and the Trade Share
(weighted vs. unweighted estimates)

Explanatory variables	*Industry category*		
	All	*Low trade shr.*	*High trade shr.*
A. Weighted estimates			
$[q_{it} - \tilde{q}_{it}]$	0.015	0.029***	0.006
	(0.010)	(0.011)	(0.011)
Sum: inflation lags[d]	−0.060	−0.159***	0.044
	(0.041)	(0.043)	(0.043)
Adj. R^2	0.243	0.228	0.306
B. Unweighted estimates			
$[q_{it} - \tilde{q}_{it}]$	0.025***	0.035***	0.014**
	(0.007)	(0.013)	(0.006)
Sum: inflation lags[d]	−0.060	−0.091**	0.004
	(0.036)	(0.042)	(0.045)
Adj. R^2	0.198	0.198	0.227

Source: Authors' calculations.
Note: Sample: a balanced panel of 185 industries from 1990:Q1 to 2017:Q4 (Obs. = 19,239). The dependent variable in each Phillips curve specification is $\Delta_5 p_{i,t+4}$, the annualized log-difference in industry-specific PPI from date $t-1$ to date $t+4$. Explanatory variables: $[q_{it} - \tilde{q}_{it}]$ = industry-specific industrial production gap. All specifications include industry and time fixed effects and lags 1,...,4 of Δp_{it} (not reported). In panel A, the specifications are estimated by WLS—using average industry employment shares as weights—while in panel B, they are estimated by OLS. Asymptotic standard errors reported in parentheses are clustered across industries and time, according to Cameron and others (2011): * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.
[a] Sum of coefficients on $\Delta p_{i,t-s}$, $s = 1,\dots,4$.

Table 6 reports the results of this exercise for inflation at the four-quarter horizon (i.e., $h = 4$) and using the industrial production gap, $[q_{it} - \tilde{q}_{it}]$, to measure slack at the industry level. In the first column of panel A, we report the weighted least squares (WLS) estimates of the coefficient on the industrial production gap for all industries, while in the second and third column, we report the corresponding WLS estimates for low and high trade share industry groupings, respectively; for comparison purposes, panel B contains the corresponding ordinary least squares (OLS) estimates, which weight all industries equally.

The WLS estimate of the coefficient on economic slack for all industries is a bit smaller than its corresponding OLS estimate—0.015 vs. 0.025—and also less precisely estimated. More importantly, the WLS estimates of coefficients on economic slack show a clear difference across the two industry groupings: In low trade intensity industries, the coefficient on economic slack is positive and statistically highly significant, whereas in high trade intensity industries, the coefficient on economic slack is essentially zero, in both economic and statistical terms. These results provide further support for the argument that globalization and increased international trade may be responsible, at least in part, for the observed attenuation in the response of inflation to fluctuations in economic activity. However, swings in producer prices at the industry level are far more likely to reflect a confluence of demand shocks, which push prices and output in the same direction, and supply shocks, which push them in opposite directions. Thus one should be cautious in providing a structural interpretation to the coefficient estimates reported in table 6.

3. Trade Share and the Effects of Aggregate Shocks

In this section, we employ an alternative approach to investigate the role that international trade may play in determining domestic inflation outcomes. Specifically, we identify aggregate shocks that simultaneously influence inflation and output dynamics and trace out their effects on industry-level outcomes. We then examine the extent to which the industry-level responses of prices, wages, output, and employment to such aggregate shocks differ across industries with a differential exposure to international trade and thus to global factors.

3.1 Econometric Methodology

As in the previous section, we focus on a balanced panel of 185 industries for which all variables are available over the 1990:Q1–2017:Q4 sample period. Given the high dimensionality—in both the cross-sectional and time-series dimensions—of our industry-level data, we use the FAVAR methodology proposed by Bernanke and Boivin (2003) and Bernanke and others (2005) to identify aggregate shocks and trace out their effect on price and wage inflation and the growth of output and employment at the industry level. To identify an aggregate shock of interest, we study the response of industry-level variables to a sudden deterioration in broad domestic financial conditions. An adverse shock to financial conditions may be interpreted as a reduction in aggregate demand and such shocks have featured prominently in recent discussions regarding the source of business cycle fluctuations over the time period under our consideration.[25]

25. See Stock and Watson, 2012.

Our estimation and identification procedure broadly follows the empirical methodology outlined in Gilchrist and others (2009). In particular, we combine the industry-level data on price and wage inflation and on the growth of output and employment in an $(n_1 \times 1)$-dimensional vector X_{1t}.[26] We then consider a set of macro-level variables that summarize domestic financial conditions—these series are combined in an $(n_2 \times 1)$-dimensional vector X_{2t}. This data-rich environment can be succinctly represented by an $(n \times 1)$-dimensional vector $X_t = [X'_{1t}, X'_{2t}]'$, where $n = n_1 + n_2$ and $t = 1,2,\ldots,T$. We assume that X_t has a (linear) factor structure, whereby $X_{it} = \lambda'_i F_t + v_{it}$, $i = 1,\ldots,n$, where F_t is a $(k \times 1)$-dimensional vector of common latent factors (with $k << n$), λ_i is the corresponding vector of factor loadings, and v_{it} is an idiosyncratic random disturbance that is assumed to be uncorrelated across i and t.

When analyzing the dynamic effects of aggregate financial shocks, we assume that a subset of these common factors—denoted by a $(k_2 \times 1)$-dimensional vector F_{2t}—are factors that are specific to the aggregate financial variables contained in the vector X_{2t}. These factors do not contemporaneously influence the industry-level variables in the vector X_{1t}, but they do affect contemporaneously the variables in the vector X_{2t}. The rest of the factors—denoted by a $(k_1 \times 1)$-dimensional vector F_{1t}, where $k = k_1 + k_2$—are assumed to span the information contained in the entire data vector X_t. The relationship between the observed variables and the unobserved factors is assumed to be linear and is given by the following system of measurement equations:

$$\begin{bmatrix} X_{1t} \\ X_{2t} \end{bmatrix} = \begin{bmatrix} \Lambda_{11} & \Lambda_{12} \\ \Lambda_{21} & \Lambda_{22} \end{bmatrix} \begin{bmatrix} F_{1t} \\ F_{2t} \end{bmatrix} + \begin{bmatrix} v_{1t} \\ v_{2t} \end{bmatrix}, \tag{4}$$

where

$$\Lambda = \begin{bmatrix} \Lambda_{11} & \Lambda_{12} \\ \Lambda_{21} & \Lambda_{22} \end{bmatrix}$$

is an $(n \times k)$ matrix of factor loadings.

The latent factors are assumed to follow a vector autoregressive process of the form:

$$\begin{bmatrix} F_{1t} \\ F_{2t} \end{bmatrix} = \Phi(L) \begin{bmatrix} F_{1,t-1} \\ F_{2,t-1} \end{bmatrix} + \begin{bmatrix} \in_{1t} \\ \in_{2t} \end{bmatrix}, \tag{5}$$

26. Note that $n_1 = 4 \times 185 = 740$; that is, four series for each of the 185 industries. Wage inflation is measured as the log-difference in the average weekly earnings.

where $\Phi(L)$ is a matrix lag-polynomial of finite order p. As it is standard in these models, we assume that $E[v_{it}\in_{st}] = 0$, for all $i = 1,2,\ldots,n$ and $s = 1,2,\ldots,k$; and $E[\in_{it}\in_{jt}] = 0$, for all $i \neq j$. In this form, our model constitutes a static representation of a dynamic factor model;[27] it is static in the sense that factors enter only contemporaneously in the system of measurement equations (4).

To identify the aggregate factors F_{2t}, we impose the following restrictions on the system of measurement equations. First, we assume that the matrix $\Lambda_{12} = 0$. This restriction on the factor loading matrix Λ implies that, once we have conditioned on the factors F_{1t}, the remaining variation in the aggregate block X_{2t} has a systematic component that is reflected in its own factor structure. Although the aggregate factors F_{2t} have no contemporaneous effect on the vector X_{1t}, they affect the factors F_{1t} and, by extension, the variables in the industry block X_{1t} with a lag through the autoregressive dynamics of equation (5). The second identifying assumption is that the contemporaneous innovations associated with the factors F_{1t} and F_{2t} are orthogonal, an assumption that separates the residual information content in the aggregate block from the factors summarizing the state of the economy, as measured by the full set of industry-specific information contained in the vector X_{1t}.[28]

In implementing this identification strategy, we let the vector X_{2t} include a broad array of domestic financial indicators. Specifically, when considering how financial shocks affect industry-level outcomes, the vector X_{2t} consists of the following five financial indicators: the GZ corporate bond credit spread and the associated excess bond premium;[29] the Moody's Baa-Aaa corporate bond credit spread; the term spread measured as the difference in yields on the ten- and two-year U.S. Treasury coupon securities; and the option-implied volatility on the S&P 500 stock price index, the VIX. The GZ and Baa-Aaa credit spreads and the excess bond premium are widely used indicators of financial strains obtained from the corporate bond market. The VIX, on the other hand, is a measure of risk appetite in equity markets, while

27. See Stock and Watson, 2010a.

28. We can estimate the FAVAR model given by equations (4) and (5) by using a Gaussian maximum likelihood method and a Kalman filter to construct the likelihood function. However, in the presence of identifying assumptions with large n, this method is computationally demanding. We, therefore, follow the four-step procedure outlined in Gilchrist and others (2009), as it is straightforward to implement and directly imposes the necessary identification restrictions.

29. See Gilchrist and Zakrajšek, 2012.

the term spread primarily reflects investors' appetite for duration risk. These five indicators provide a broad summary of domestic financial conditions that are entirely reliant on market prices and, therefore, should capture changes in broad financial conditions in a timely manner. In the FAVAR specifications, we allow for four common factors in the industry-level block X_{1t}—that is, $k_1 = 4$— and for one factor in the aggregate block X_{2t}, that is $k_2 = 1$.[30]

3.2 The Impact of Financial Shocks

In this section, we present impulse responses of variables in the industry block X_{1t} to the identified aggregate financial shock. We begin by reporting these baseline results for all industries. Next, we examine how international trade exposure influences industry-level inflation dynamics by again dividing our sample of industries into those with a low external trade exposure and those with a high trade exposure.

Figure 9 plots the distribution of industry-level outcomes in response to an adverse financial shock of one standard deviation in quarter zero. Though not shown, this shock causes a broad-based tightening of domestic financial conditions, implying an increase in the excess bond premium of about 30 basis points upon impact.[31] The solid line in each panel shows the median industry response of the specified variable to such a shock, while the dark shaded bands denote the range of responses between the 75th and 25th percentiles (the $P75$–$P25$ range) and the light shaded bands denote the range of responses between the 95th and 5th percentiles (the $P95$–$P5$ range). Recall that the factor F_{2t} is, by assumption, contemporaneously orthogonal to the variables in the industry block and thus aggregate shocks have no effect on industry-level outcomes upon impact.

30. These choices were based on the information criteria proposed by Bai and Ng (2002); however, all of the results reported in this paper are robust to allowing a greater number of factors in either block.

31. Over the 1990:Q1–2017:Q4 period, the standard deviation of the excess bond premium is about 50 basis points. As a point of comparison, the excess bond premium shot up more than 300 basis points following the collapse of Lehman Brothers in September 2008.

Figure 9. Implications of an Adverse Financial Shock
(all industries)

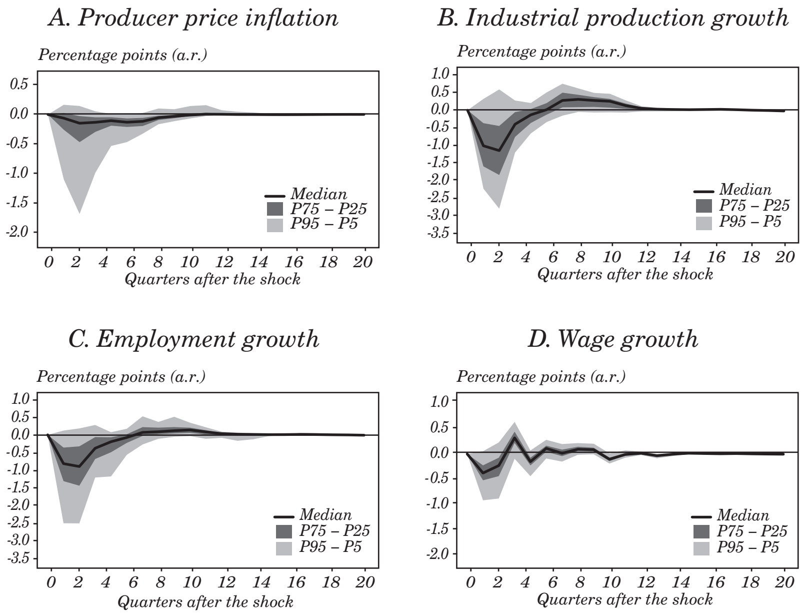

Source: Authors' calculations.
Note: The solid line in each panel depicts the median response of the specified variable to an adverse financial shock of one standard deviation across 185 industries; the shaded bands depict the corresponding *P75–P25* and *P95–P5* ranges. See the text for details.

The identified financial shock is clearly contractionary—it induces a substantial decline in the growth of industrial production and employment for the median industry. A couple of quarters after its impact, this shock is cutting 1.2 percentage points from the annualized growth of output and 0.8 percentage points from the annualized growth of employment at the median. It also causes a significant step-down in both price and wage inflation: For the median industry, annualized price inflation is lowered 0.1 percentage points, whereas the reduction in annualized wage inflation is on the order of 0.4 percentage points. Notably, the reduction in the rate of growth of economic activity, prices, and wages occurs relatively quickly, peaking a mere two quarters after the shock. Economic growth remains depressed for several more quarters before recovering slowly and returns to its long-run level only after about eight quarters. Judging by the shaded regions, the estimated range of industry responses implies that these effects are broad based. The combination of steep declines in the growth of output, employment, prices, and wages implies that the deterioration in broad domestic financial conditions delivers a response that is consistent with a reduction in aggregate demand within a New Keynesian framework.

With these baseline results in hand, we now analyze the extent to which differential trade exposure across industries changes the results reported above. As before, we sort our sample of industries based on their average trade exposure over the 1990:Q1–2017:Q4 period and group them into a low and high trade exposure categories. We then separately estimate our two FAVAR specifications for each of the two groupings, an approach that ensures that we do not artificially constrain the factor structure to be the same across industries with a differential trade exposure. As a reminder, recall that each category of industries accounts, on average, for about 50 percent of total employment in our sample.

Unlike our baseline exercise, this exercise is focused on the implications of the common financial shock for aggregate outcomes. Specifically, for each industry-level endogenous variable, we compute a weighted-average response across industries, where weights are equal to the industry-specific average employment shares within each group of industries (i.e., low vs. high trade exposure industry categories). In addition, we report the aggregate responses for all industries by weighting the industry-specific responses from figure 9 with their corresponding average employment shares; these results are shown in figure 10.[32]

As shown by the solid lines in figure 10, the aggregate responses of producer price and wage inflation and the growth of output and employment to an adverse financial shock follow closely the contours of the corresponding median industry-level responses shown in figure 9: Price and wage inflation, along with output and employment growth, all fall sharply, with peak responses occurring one to two quarters after the impact of the shock. Moreover, these aggregate responses remain persistently below their respective long-run values for six to eight quarters after the shock.

32. Figures B.1–B.2 in appendix B show the industry-level responses for the low and high trade share industry categories when the economy is perturbed by an aggregate financial shock.

Figure 10. Implications of an Adverse Financial Shock
(low vs. high trade share industries)

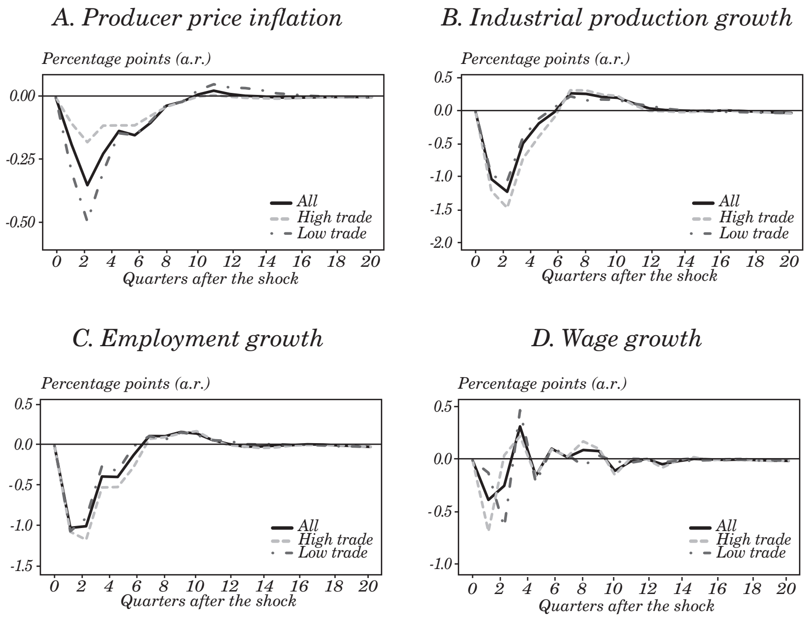

Source: Authors' calculations.
Note: The solid line in each panel depicts the employment-weighted-average response of the specified variable to an adverse financial shock of one standard deviation across 185 industries; the dashed (dashed-dotted) lines depict the corresponding employment-weighted-average responses for a subset of industries with a high (low) average trade share. See the text for details.

Note that the (absolute) magnitude of responses for the aggregates—as defined by the employment-weighted averages of industry-level responses—are somewhat larger than their corresponding unweighted median responses across industries. In particular, the annualized output and employment growth both fall by more than one percentage point, while the annualized producer price inflation declines about 30 basis points. The estimated decline in the growth of output in response to a financial shock is consistent with other studies that find that such disturbances lead to a significant contraction in economic activity.[33] That said, the estimated drop in producer price inflation is both larger and

33. See Gilchrist and others, 2009; Gilchrist and Zakrajšek, 2012; and Boivin and others, 2018.

occurs more quickly than the one estimated by VARs that use aggregate time-series data. Overall, these results indicate that producer price inflation is fairly sensitive to fluctuations in economic activity induced by changes in broad financial conditions—producer price inflation declines roughly 25 basis points when a tightening of financial conditions induces a one percentage point decline in the growth of industrial output.

Figure 10 also displays the aggregate responses to an adverse financial shock for high and low trade industries. As shown in the upper left panel, the dynamics of inflation differ markedly across industries with a differential trade exposure. Notably, the peak decline in producer price inflation of 0.5 percentage points for industries with low trade exposure is more than three times as large as that for industries with high trade exposure. Although the unanticipated tightening of financial conditions causes a somewhat greater contraction in economic activity among high trade industries, the responses of output and employment growth are broadly similar—in terms of both timing and their magnitudes—across the two industry groupings. Wage inflation also behaves in a similar manner across these two industry groupings, though in high trade industries, the deceleration in wages occurs more quickly.

A useful way to highlight the difference in inflation dynamics between low and high trade industries is to compute the cumulative responses of price inflation and output growth. The ratio of the resulting price response to the output response then provides an estimate of the decline in prices relative to output that occurs at different horizons in response to an adverse financial shock. As shown in figure 11, in low trade intensity industries, producer prices are estimated to decline about 0.3 percent for every one percent decline in output at very short horizons and about one percent for the same-sized reduction in output at the two-year horizon. In high trade intensity industries, by contrast, producer prices are estimated to decline about 0.1 percent for a one percent reduction in output at very short horizons and about 0.3 percent at the two-year horizon. In sum, these findings imply that the inflation-output tradeoff is—at every horizon—three times larger in low trade intensity industries than in their high trade intensity counterparts.

Figure 11. Inflation-Output Tradeoff

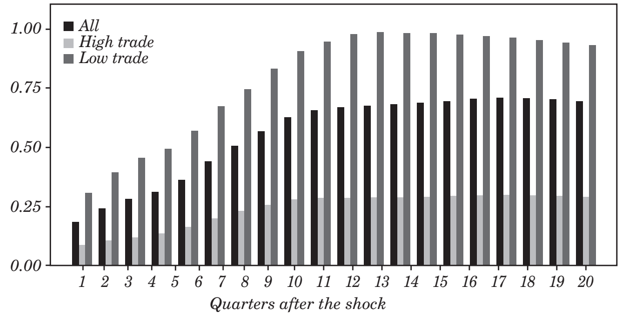

Source: Authors' calculations.
Note: The bars in the figure depict the estimated sensitivity of producer prices to fluctuations in output induced by aggregate financial shocks. See the text for details.

In summary, our FAVAR analysis implies that producer price inflation is three to four times more responsive to aggregate demand shocks in low trade intensity industries than their high trade intensity counterparts. Responses of wages, output, and employment, by contrast, are strikingly similar across the two industry groupings. These results are consistent with the notion that the Phillips curve is indeed much flatter in industries that are more exposed to international trade and are thus broadly consistent with our earlier findings, which show that the estimated flattening of the aggregate Phillips curves coincides to a substantial degree with the increased exposure of the U.S. economy to international trade.

4. Conclusions

In this chapter, we examine the extent to which the response of inflation to fluctuations in economic activity has weakened over time. Furthermore, we analyze the role of globalization and rising trade shares behind these structural shifts. Our evidence points to a significant flattening of the Phillips curve that occurred in the 1990s. Although there is some evidence of a recent rise in the responsiveness of CPI inflation to changes in economic slack, it remains the case that both PPI and CPI inflation are substantially less responsive to fluctuations in economic activity today, relative to estimates that rely on the pre-1990 data. To a significant degree, this reduced responsiveness of inflation to economic slack coincides with a rising U.S. trade share and a concomitant increase in global economic integration.

Industry-level data provide further evidence in favor of the notion that trade intensity attenuates the response of inflation to fluctuations in economic activity. Industry-level estimates of the Phillips curve imply a substantially lower sensitivity of PPI inflation to output in industries with a high trade share, relative to those with a low trade share. We confirm these results by examining the response of industry-level PPI inflation and output to identified aggregate financial shocks. This evidence implies that the inflation-output tradeoff is about three times larger for low trade intensity industries than for their high trade intensity counterparts. In this sense, increased international trade and globalization do indeed appear to help explain the observed flattening of the aggregate Phillips curve over the past several decades.

References

Andrews, D.W.K. 1993. "Tests for Parameter Instability and Structural Change With Unknown Change Point." *Econometrica* 61: 821–56.

Atkeson, A. and L.E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" *Federal Reserve Bank of Minneapolis Quarterly Review* (Winter): 2–11.

Auer, R., C. Borio, and A. Filardo. 2017. "The Globalisation of Inflation: The Growing Importance of Global Value Chains." Working Paper No. 602, Bank for International Settlements.

Bai, J. and S. Ng. 2002. "Determining the Number of Factors in Approximate Factor Models." *Econometrica* 70: 191–221.

Ball, L. 2006. "Has Globalization Changed Inflation?" NBER Working Paper No. 12687.

Ball, L.M. and S. Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 42: 337–405.

———. 2018. "A Phillips Curve with Anchored Expectations and Short-Term Unemployment." *Journal of Money, Credit and Banking*. Available at https://doi:10.1111/jmcb. 12502.

Ball, L.M., N.G. Mankiw, and D.H. Romer. 1988. "The New Keynesian Economics and the Output-Inflation Tradeoff." *Brookings Papers on Economic Activity* 19: 1–82.

Beaudry, P. and M. Doyle 2000. "What Happened to the Phillips Curve in the 1990s in Canada?" In *Proceedings of the Bank of Canada Conference on Price Stability and the Long-Run Target for Monetary Policy*, Bank of Canada.

Berganza, J.C., P. Del Río López, and F. Borrallo. 2016. "Determinants and Implications of Low Global Inflation Rates." Banco de España Occasional Paper No. 1608. Available at *http://dx.doi.org/10.2139/ssrn.2891619*.

Bernanke, B.S. 2007. "Globalization and Monetary Policy." Speech delivered at the Fourth Economic Summit, Stanford Institute for Economic Policy Research, Stanford, CA. Available at *https://www.federalreserve. gov/newsevents/speech/bernanke20070302a.htm*.

———. 2010. "The Economic Outlook and Monetary Policy." Speech delivered at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, 27 August. Available at *http://www.federalreserve. gov/newsevents/speech/bernanke20100827a.htm*.

Bernanke, B.S. and J. Boivin. 2003. "Monetary Policy in a Data-Rich Environment." *Journal of Monetary Economics* 50: 525–46.

Bernanke, B.S., J. Boivin, and P. Eliasz. 2005. "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach." *Quarterly Journal of Economics* 120: 387–422.

Blanchard, O.J. 2016. "The Phillips Curve: Back to the '60s?" *American Economic Review* 106: 31–4.

———. 2018. "Should We Reject the Natural Rate Hypothesis?" *Journal of Economic Perspectives* 32: 97–120.

Boivin, J., M.P. Giannoni, and D. Stevanović. 2018. "Dynamic Effects of Credit Shocks in a Data-Rich Environment." *Journal of Business and Economic Statistics* 25(1): 52–60. Available at *https://doi.org/ 10.1080/07350015.2018.1497507*.

Borio, C. and A. Filardo. 2007. "Globalisation and Inflation: New Cross-Country Evidence on the Global Determinants of Domestic Inflation." Working Paper No. 227, Bank for International Settlements.

Cameron, A.C., J.B. Gelbach, and D.L. Miller. 2011. "Robust Inference with Multi-Way Clustering." *Journal of Business and Economic Statistics* 29(2): 238–49.

Elliott, G. and U.K. Müller. 2006. "Efficient Tests for General Persistent Time Variation in Regression Coefficients." *Review of Economic Studies* 73: 907–40.

Forbes, K.J. 2018. "Has Globalization Changed the Inflation Process?" Working Paper, MIT Sloan School of Management.

Friedrich, C. 2016. "Global Inflation Dynamics in the Post-Crisis Period: What Explains the Twin Puzzle?" *Economic Letters* 142: 31–4.

Galí, J. and M. Gertler. 2000. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44: 195–222.

Galí, J., M. Gertler, and D. López-Salido. 2001. "European Inflation Dynamics." *European Economic Review* 45: 1237–70.

———. 2007. "Markups, Gaps, and the Welfare Costs of Business Fluctuations." *Review of Economics and Statistics* 89: 44–59.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajšek. 2017. "Inflation Dynamics During the Financial Crisis." *American Economic Review* 107; 785–823.

Gilchrist, S., V. Yankov, and E. Zakrajšek. 2009. “Credit Market Shocks and Economic Fluctuations: Evidence From Corporate Bond and Stock Markets.” *Journal of Monetary Economics* 56: 471–93.

Gilchrist, S. and E. Zakrajšek. 2012. “Credit Spreads and Business Cycle Fluctuations.” *American Economic Review* 102: 1692–720.

————. 2016. “Customer Markets and Financial Frictions: Implications for Inflation Dynamics.” Federal Reserve Bank of Kansas City, *Proceedings – Economic Policy Symposium – Jackson Hole*.

Gordon, R.J. 1982. “Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment.” In *Workers, Jobs, and Inflation*, edited by M. Baily. Washington D.C.: The Brookings Institution.

————. 2013. “The Phillips Curve is Alive and Well: Inflation and the NAIRU during the Slow Recovery.” NBER Working Paper No. 19390.

Hamilton, J.D. 1983. “Oil and the Macroeconomy Since World War II.” *Journal of Political Economy* 91: 228–48.

————. 2018. “Why You Should Never Use the Hodrick-Prescott Filter.” *Review of Economics and Statistics* 100: 831–43.

Ihrig, J.E., S.B. Kamin, D. Linder, and J.R. Marquez. 2010. “Some Simple Tests of the Globalization and Inflation Hypothesis.” *International Finance* 13; 343–75.

Jermann, U.J. and V. Quadrini. 2012. “Macroeconomic Effects of Financial Shocks.” *American Economic Review* 102: 238–71.

Krueger, A.B., J. Cramer, and D. Cho. 2014. “Are the Long-Term Unemployed on the Margins of the Labor Market?” *Brookings Papers on Economic Activity* 48: 229–98.

Kuttner, K.N. and T. Robinson. 2010. “Understanding the Flattening of the Phillips Curve,” *North American Journal of Economics and Finance* 21: 110–25.

Lindsey, D.E., A. Orphanides, and R.H. Rasche. 2005. “The Reform of October 1979: How It Happened and Why.” *Review, Federal Reserve Bank of St. Louis* 87: 187–235.

Mankiw, N.G., R. Reis, and J. Wolfers. 200.: “Disagreement about Inflation Expectations.” In *NBER Macroeconomics Annual*, edited by M. Gertler and K. Rogoff, Cambridge MA: MIT Press.

Mavroeidis, S., M. Plagborg-Møller, and J.H. Stock. 2001. “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve.” *Journal of Economic Literature* 52: 124–88.

Miles, D., U. Panizza, R. Reis, and A. Ubide. 2017. "And Yet It Moves: Inflation and the Great Recession." Geneva Reports on the World Economy 19. The International Center for Monetary and Banking Studies and the Centre for Economic Policy Research.

Nakamura, E. and J. Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123: 1415–64.

Newey, W.K. and K.D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation-Consistent Covariance Matrix." *Econometrica* 55: 703–08.

Orphanides, A. and J.C. Williams. 2013. "Monetary Policy Mistakes and the Evolution of Inflation Expectations." In *The Great Inflation: The Rebirth of Modern Central Banking*, edited by M. D. Bordo and A. Orphanides, Chicago, IL: University of Chicago Press.

Peersman, G. and W. Wagner. 2014. "Shocks to Bank Lending, Risk-Taking, Securitization, and their Role for U.S. Business Cycle Fluctuations." CEPR Discussion Paper No. 10547.

Pfajfar, D. and J.M. Roberts. 2018. "The Role of Expectations in Changed Inflation Dynamics." Finance and Economics Discussion Series Paper No. 2018-062, Federal Reserve Board.

Roberts, J.M. 1995. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit, and Banking* 25: 975–84.

————. 2006. "Monetary Policy and Inflation Dynamics," *International Journal of Central Banking* 2: 193–230.

Sbordone, A.M. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49: 265–92.

Simon, J., T. Matheson, and D. Sandri. 2013. "The Dog that Didn't Bark: Has Inflation Been Muzzled or Was It Just Sleeping?" In *World Economic Outlook: Hopes, Realities, and Risks*, Washington D.C.: International Monetary Fund.

Sims, C.A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50: 665–90.

Stock, J.H. and M.W. Watson. 2009. "Phillips Curve Inflation Forecasts." In *Understanding Inflation and the Implications for Monetary Policy*, edited by J. Fuhrer, Y. Kodrycki, J. Little, and G. Olivei, Cambridge, MA: MIT Press.

————. 2010a. "Dynamic Factor Models." In *Oxford Handbook of Economic Forecasting*, edited by M.P. Clements and D.F. Henry, Oxford, UK: Oxford University Press.

————. 2010b. "Modeling Inflation after the Crisis." Federal Reserve Bank of Kansas City, *Proceedings – Economic Policy Symposium – Jackson Hole*.

————. 2012. "Disentangling the Channels of the 2007–09 Recession." *Brookings Papers on Economic Activity* 44: 81–135.

————. 2018. "Slack and Cyclically Sensitive Inflation." Working Paper, Harvard University.

Yellen, J.L. 2013. "Panel Discussion on 'Monetary Policy: Many Targets, Many Instruments. Where Do We Stand?'" Speech delivered at the "Rethinking Macro Policy II," a conference sponsored by the International Monetary Fund, Washington, D.C., 16 April. Available at *http://www.federalreserve.gov/newsevents/speech/yellen10130416a.htm*.

Zhang, C. 2017. "The Great Globalization and Changing Inflation Dynamics." *International Journal of Central Banking* 13: 191–226.

Appendix

A. Controlling for Trend Inflation

As discussed in the main text, both producer and consumer price inflation exhibit significant low frequency variation over our sample period (figure 1). To ensure that this low frequency variation does not affect our baseline time-series estimates of the aggregate Phillips curves, this appendix reports a set of results in which all inflation series were "detrended" to eliminate very low frequency variation. Specifically, following Stock and Watson (2012), we calculated the deviations of each quarterly inflation series from a local mean, where the latter is estimated using a bi-weight kernel with a bandwidth of 100 quarters. As noted by Stock and Watson (2012), these local mean estimates are roughly the same as those computed using a centered moving-average window of ±30 quarters. This approach of eliminating low frequency variation in inflation rates has the desirable feature that it makes no assumption about reversion to the local mean.

As shown in figures A.1 and A.2, the values of these local means change substantially over our sample period. Tables A.1 and A.2 contain estimates of the baseline Phillips curve specifications for producer and consumer price inflation, respectively, which use the detrended inflation data; these estimates are directly comparable with those reported in tables 1 and 2 of the main text, which use the untransformed inflation series. Tables A.3 and A.4, in contrast, use the detrended inflation series to examine the role the trade share in influencing the slope of the Phillips curve, and the results in those tables are directly comparable to those reported in tables 3 and 4 of the main text.

Figure A1. Producer Price Inflation

A. Headline

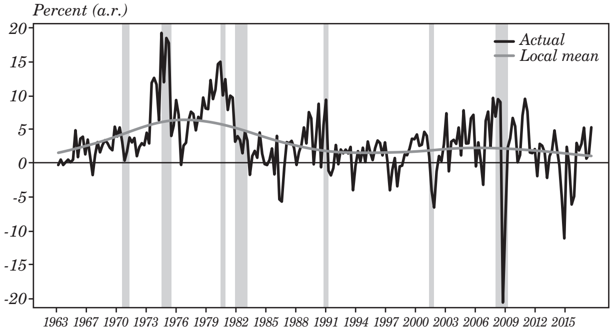

B. Core

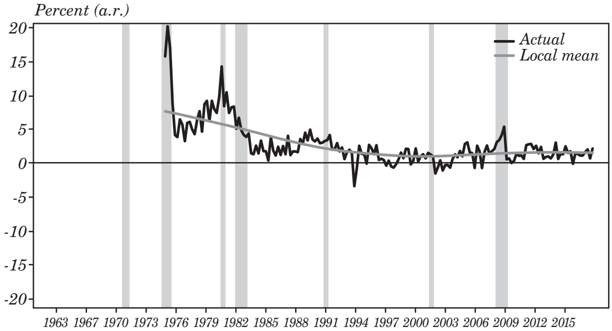

Source: Authors' calculations using data from the Bureau of Labor Statistics.
Note: Panel A depicts the annualized quarterly log-difference of headline PPI inflation and its estimated local mean, while panel B depicts the corresponding series for core PPI inflation (see the text for details). The shaded vertical bars denote the NBER-dated recessions. All price indices are seasonally adjusted.

Figure A2. Consumer Price Inflation

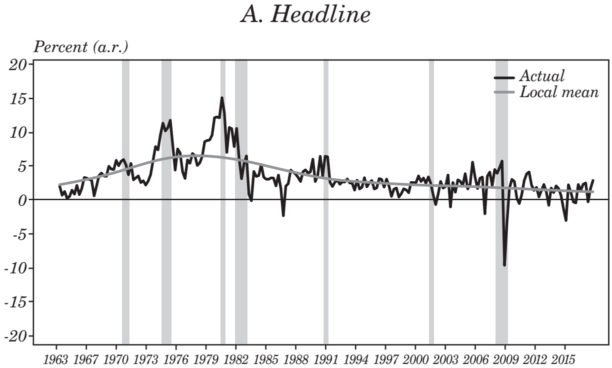

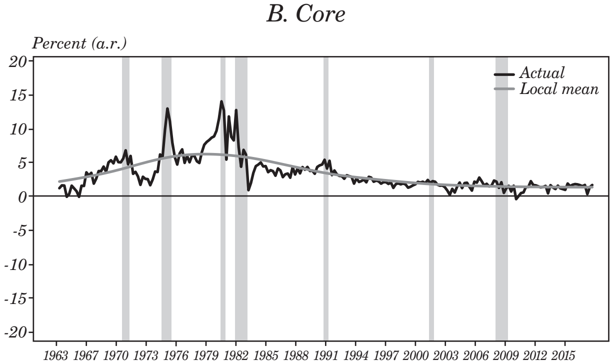

Source: Authors' calculations using data from the Bureau of Labor Statistics.
Note: Panel A depicts the annualized quarterly log-difference of headline CPI inflation and its estimated local mean, while panel B depicts the corresponding series for core CPI inflation (see the text for details). The shaded vertical bars denote the NBER-dated recessions. All price indices are seasonally adjusted.

Table A1. Phillips Curve – Detrended Producer Price Inflation

Explanatory variables	*h* = 1 (*1*)	*h* = 1 (*2*)	*h* = 4 (*3*)	*h* = 4 (*4*)
A. Producer prices				
$[y_t - y_t^*]$	0.263**	-	0.294**	-
	(0.129)	-	(0.135)	-
$[U_t - U_t^*]$	-	-0.306	-	-0.331
	-	(0.224)	-	(0.232)
Sum: inflation lags[a]	0.371**	0.369**	0.212*	0.211*
	(0.147)	(0.147)	(0.112)	(0.112)
sup W[b]	16.168***	115.353***	32.027***	34.338***
	[09:Q1]	[09:Q1]	[84:Q1]	[08:Q1]
q_{LL}[c]	-4.536	-4.975	-3.793	-3.178
Adj. R^2	0.161	0.141	0.160	0.115
B. Core producer prices				
$[y_t - y_t^*]$	0.136**	-	0.163**	-
	(0.052)	-	(0.062)	-
$[U_t - U_t^*]$	-	-0.240**	-	-0.271**
	-	(0.099)	-	(0.122)
Sum: inflation lags[a]	0.565***	0.538***	0.434***	0.403***
	(0.117)	(0.114)	(0.101)	(0.096)
sup W[b]	17.562***	25.350***	60.378***	78.601***
	[84:Q1]	[84:Q1]	[81:Q4]	[82:Q2]
q_{LL}[c]	-7.776*	-7.716*	-5.355	-5.395
Adj. R^2	0.378	0.373	0.409	0.389

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4 for headline PPI (panel A), and 1974:Q1 to 2017:Q4 for core PPI (panel B). The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the detrended annualized log-difference in the specified PPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, and $[U_t - U_t^*]$ = unemployment gap. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

a. Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

b. The Andrews (1993) sup-Wald statistic of the null hypothesis that there is no structural break in the coefficient on economic slack; the estimated break dates are reported in brackets below.

c. The Elliott and Müller (2006) q_{LL} statistic of the null hypothesis that the coefficient on economic slack is constant over time.

Table A2. Phillips Curve – Detrended Consumer Price Inflation

Explanatory variables	*h* = 1		*h* = 4	
	(1)	*(2)*	*(3)*	*(4)*
A. Consumer prices				
$[y_t - y_t^*]$	0.196***	-	0.232***	-
	(0.072)	-	(0.075)	-
$[U_t - U_t^*]$	-	-0.264**	-	-0.290**
	-	(0.114)	-	(0.113)
Sum: inflation lags[a]	0.571***	0.550***	0.392***	0.369***
	(0.114)	(0.118)	(0.101)	(0.108)
sup W[b]	22.185***	17.394***	52.617***	30.743***
	[83:Q2]	[83:Q2]	[83:Q1]	[91:Q3]
q_{LL}[c]	-6.173	-6.729	-4.367	-3.846
Adj. R^2	0.339	0.316	0.346	0.290
B. Core consumer prices				
$[y_t - y_t^*]$	0.124***	-	0.179***	-
	(0.045)	-	(0.053)	-
$[U_t - U_t^*]$	-	-0.200***	-	-0.257***
	-	(0.071)	-	(0.085)
Sum: inflation lags[a]	0.677***	0.652***	0.479***	0.441***
	(0.100)	(0.100)	(0.116)	(0.122)
sup W[b]	20.048***	19.449***	65.059***	38.216***
	[83:Q2]	[83:Q2]	[83:Q1]	[83:Q1]
q_{LL}[c]	-6.486	-6.384	-6.068	-5.870
Adj. R^2	0.499	0.495	0.427	0.394

Source: Authors' calculations.
Note: Sample: 1962:Q2 to 2017:Q4. The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the detrended annualized log-difference in the specified CPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, and $[U_t - U_t^*]$ = unemployment gap. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.
a. Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.
b. The Andrews (1993) sup-Wald statistic of the null hypothesis that there is no structural break in the coefficient on economic slack; the estimated break dates are reported in brackets below.
c. The Elliott and Müller (2006) q_{LL} statistic of the null hypothesis that the coefficient on economic slack is constant over time.

Table A3. Phillips Curve and the Trade Share – Detrended Producer Price Inflation

Explanatory variables	h = 1 (1)	h = 1 (2)	h = 4 (3)	h = 4 (4)
A. Producer prices				
$[y_t - y_t^*]$	0.702*	-	1.112**	-
	(0.397)	-	(0.494)	-
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.022	-	-0.040*	-
	(0.021)	-	(0.023)	-
$[U_t - U_t^*]$	-	-1.033	-	-1.310*
	-	(0.627)	-	(0.739)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.037	-	0.048
	-	(0.030)	-	(0.033)
Sum: inflation lags[a]	0.382**	0.374**	0.231**	0.217**
	(0.146)	(0.143)	(0.105)	(0.106)
Adj. R^2	0.166	0.147	0.210	0.142
B. Core producer prices				
$[y_t - y_t^*]$	0.636**	-	0.808***	-
	(0.270)	-	(0.281)	-
$[y_t - y_t^*] \times \text{TrdShr}_{t-1}$	-0.022*	-	-0.028**	-
	(0.011)	-	(0.018)	-
$[U_t - U_t^*]$	-	-1.908***	-	-2.239***
	-	(0.529)	-	(0.516)
$[U_t - U_t^*] \times \text{TrdShr}_{t-1}$	-	0.069***	-	0.081***
	-	(0.021)	-	(0.021)
Sum: inflation lags[a]	0.562***	0.490***	0.431***	0.346***
	(0.113)	(0.110)	(0.096)	(0.083)
Adj. R^2	0.402	0.443	0.466	0.519

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4 for headline PPI (panel A), and 1974:Q1 to 2017:Q4 for core PPI (panel B). The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the detrended annualized log-difference in the specified PPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, $[U_t - U_t^*]$ = unemployment gap, TrdShr_{t-1} = eight-quarter (trailing) moving average of the trade share. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

a. Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

Table A4. Phillips Curve and the Trade Share – Detrended Consumer Price Inflation

Explanatory variables	*h* = 1 (1)	*h* = 1 (2)	*h* = 4 (3)	*h* = 4 (4)
A. Consumer prices				
$[y_t - y_t^*]$	0.505**	-	0.772***	-
	(0.227)	-	(0.275)	-
$[y_t - y_t^*] \times \mathrm{TrdShr}_{t-1}$	-0.015	-	-0.028**	-
	(0.011)		(0.012)	
$[U_t - U_t^*]$	-	-0.717**	-	-0.867**
	-	(0.320)	-	(0.355)
$[U_t - U_t^*] \times \mathrm{TrdShr}_{t-1}$	-	0.022	-	0.028*
	-	(0.015)	-	(0.015)
Sum: inflation lags[a]	0.602***	0.566***	0.447***	0.390***
	(0.117)	(0.116)	(0.173)	(0.105)
Adj. R^2	0.348	0.324	0.398	0.313
B. Core consumer prices				
$[y_t - y_t^*]$	0.315**	-	0.576***	-
	(0.146)	-	(0.186)	-
$[y_t - y_t^*] \times \mathrm{TrdShr}_{t-1}$	-0.009*	-	-0.019***	-
	(0.006)	-	(0.007)	-
$[U_t - U_t^*]$	-	-0.503**	-	-0.698***
	-	(0.201)	-	(0.242)
$[U_t - U_t^*] \times \mathrm{TrdShr}_{t-1}$	-	0.015*	-	0.021**
	-	(0.008)	-	(0.009)
Sum: inflation lags[a]	0.711***	0.623***	0.549***	0.471***
	(0.102)	(0.099)	(0.112)	(0.118)
Adj. R^2	0.504	0.501	0.465	0.414

Source: Authors' calculations.

Note: Sample: 1962:Q2 to 2017:Q4. The dependent variable in each Phillips curve specification is $\Delta_{h+1} p_{t+h}$, the annualized log-difference in the specified CPI from date $t-1$ to date $t+h$. Explanatory variables: $[y_t - y_t^*]$ = output gap, $[U_t - U_t^*]$ = unemployment gap, and TrdShr_{t-1} = eight-quarter (trailing) moving average of the trade share. All specifications include a constant and lags 1,...,4 of Δp_t (not reported) and are estimated by OLS. Asymptotic standard errors reported in parentheses are computed according to Newey and West (1987) with the "lag-length" parameter equal to four: * $p < 0.10$; ** $p < 0.05$, and *** $p < 0.01$.

a. Sum of coefficients on Δp_{t-s}, $s = 1,...,4$.

B. Supplementary FAVAR Results

Figure B1. Implications of an Adverse Financial Shock
(industries with a low trade share)

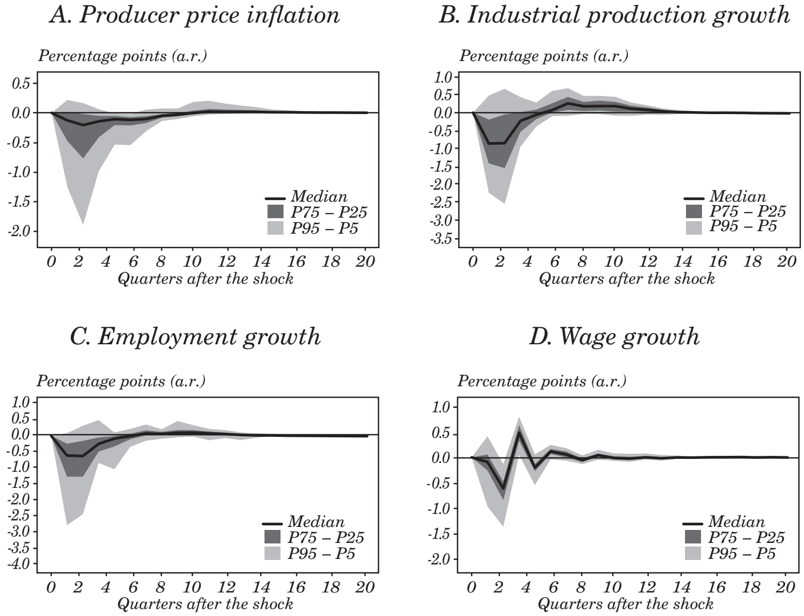

Source: Authors' calculations.
Note: The solid line in each panel depicts the median response of the specified variable to an adverse financial shock of one standard deviation across a subset of industries with a low average trade share; the shaded bands depict the corresponding *P75 – P25* and *P95 – P5* ranges. See the main text for details.

Figure B2. Implications of an Adverse Financial Shock
(industries with a high trade share)

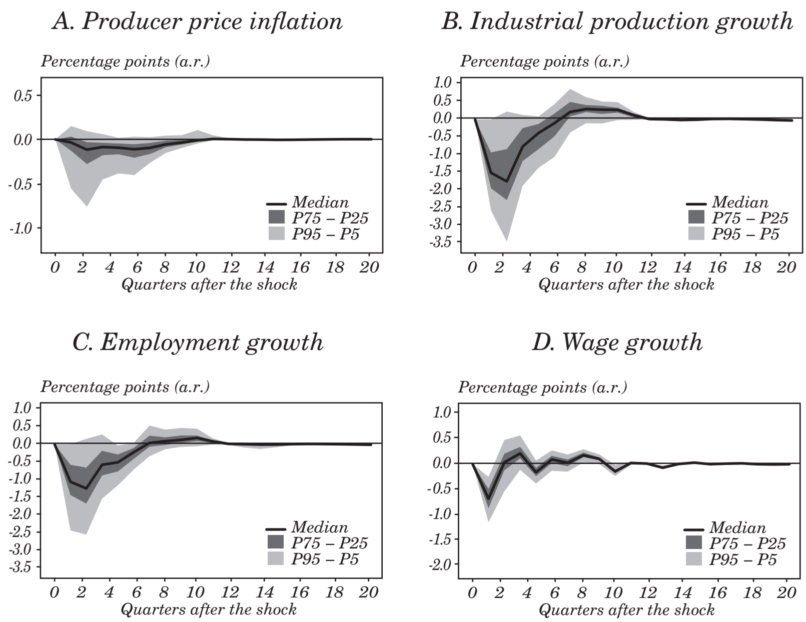

Source: Authors' calculations.
Note: The solid line in each panel depicts the median response of the specified variable to an adverse financial shock of one standard deviation across a subset of industries with a high average trade share; the shaded bands depict the corresponding *P75 – P25* and *P95 – P5* ranges. See the main text for detail.