DOES MONETARY POLICY MATTER?
THE NARRATIVE APPROACH AFTER 35 YEARS

ONLINE APPENDIX B:
ROBUSTNESS

Christina D. Romer
David H. Romer

University of California, Berkeley

March 2023

This appendix investigates the robustness of our empirical results. Section I examines some additional outcome measures. Section II considers a wide range of alternative specifications, as well as different approaches to computing the standard errors. Section III reports a very simple type of nonparametric evidence—plots of the errors in forecasting real GDP, inflation, and the unemployment rate after the dates of each of our monetary shocks, where the forecasts are from simple univariate forecasting equations. Finally, Section IV provides details on the additional data used in the robustness tests.

I. ADDITIONAL OUTCOME MEASURES

VARIABLES

In the paper, we consider two indicators of real activity and three measures of inflation. For real activity, they are the monthly unemployment rate and quarterly real GDP. For inflation, they are quarterly inflation as measured by the GDP price index, the price index for personal consumption expenditures (PCE), and the price index for PCE excluding food and energy (core PCE). Here, we examine two additional measures of real activity and two additional inflation measures. All four of these series are monthly.

For real activity, the first additional series is employment. We use the payroll series, in logarithms. This series is well measured, higher frequency than real GDP, and, in contrast to unemployment, a measure of activity rather than inactivity. The second series is industrial production (also in logarithms). Industrial production is a monthly rather than quarterly measure of production, and it was a major focus of our 1989 paper. At the same time, it omits large portions of output, and its importance in the economy has fallen steadily over time. Both series are available on a consistent basis back to the start of our sample, and each has been produced by a single agency throughout (the Bureau of Labor Statistics in the case of employment, the Federal Reserve Board in the case of industrial production).

For inflation, the two additional series are the monthly versions of inflation measured using the PCE price index and the core PCE price index. The obvious advantage of these series is that they are higher frequency than the corresponding quarterly series that we consider in the paper. But they have the drawback that they are noisier than the quarterly series. An additional disadvantage of the monthly PCE price index is that it is not available before 1959. Constructing a monthly series back to the start of our sample therefore requires splicing to other series; the details are described in Section IV. (The core PCE price index is not available at either frequency before 1959.)

RESULTS

We estimate equation (1) in the paper for various horizons for each of the additional outcome measures. Because the series are monthly, we set K (the number of lags) to 12, the sample period is 1947:10 to 2016:12, and we estimate the equation for h equal to 0 to 60.

Figures B1 and B2 show the results for the two real activity measures. For employment, the patterns and statistical significance are very similar to those for unemployment. There is little effect for the first six months, then a steady worsening over the next two years (falling employment and rising unemployment), with a peak t-statistic over 3.5, and then a gradual recovery. The largest difference is that the estimates for unemployment revert almost all the way to zero by month 60, while for those for employment revert only slightly more than halfway to zero. For both series, however, the 2-standard error confidence interval at 60 months includes both full reversion and reversion to half of the maximum estimated effect. In terms of magnitudes, the peak fall in employment is 3.75 percent (at month 31), while the peak rise in the unemployment rate is 1.60 percentage points (at month 27).

For industrial production, the time patterns and statistical significance are very similar to those for real GDP. There is a substantial fall starting around six months (or two quarters) after the shock, a peak effect at 27 months (or 9 quarters), a return roughly two-thirds of the way to zero after 60 months (or 20 quarters), and very high statistical significance, with multiple *t*-statistics over 4. The obvious difference between the results for the two series is that the magnitudes are larger for industrial production; for example, the peak effect is about twice as large. This fits with the greater cyclical volatility of industrial production.

Figure B3 shows the results for the two monthly inflation measures. The results are quite similar to the results for the corresponding quarterly measures shown in Figure 7 in the paper. There is a modest price puzzle over the first few quarters (that is, a rise in inflation following a contractionary monetary policy shock), which is slightly more pronounced for core than headline PCE inflation. Inflation then falls steadily until about two years after the shock, roughly leveling off at somewhat more than 1 percentage point below what it would have been otherwise. One difference from the quarterly results is that the standard errors are about 25 percent larger, which is a natural consequence of the fact that inflation is more volatile when measured at a monthly than a quarterly frequency.

II. ALTERNATIVE SPECIFICATIONS

We consider multiple variations on our baseline specification for the core outcome variables we consider in the main text (the unemployment rate, real GDP, and the three quarterly inflation measures).

SPECIFICATIONS

For concreteness, we describe our alternative specifications for the case of quarterly data. The implementation for the monthly data on the unemployment rate is analogous; for example, 2011:4 becomes 2011:12, 8 lags become 24, and so on.

The first two variations concern how we treat the end of our sample period. Because our shock series ends in 2016:4 (since that is the last date for which transcripts of FOMC meetings were available as of 2022), in our baseline specification we only use data through the end of 2016. This has the effect that the sample size of our regressions varies with the horizon. The last observation in the sample for $h = 0$ is 2016:4; the last observation for $h = 1$ is 2016:3; and so on through 2011:4 for $h = 20$. We therefore consider a variant where for all h, the last observation in the sample is 2011:4.

The other variant involving the end of the sample is to potentially use values of the dependent variable (Y_{t+h} in equation [1] in the paper) through the end of 2019. When we do this, the last observation in the sample is 2016:4 for the $h = 0$ through $h = 12$ regressions, 2016:3 for the $h = 13$ regression, and so on through 2014:4 for the $h = 20$ regressions. We do not use any data from 2020 or beyond so that our results are not influenced by the extraordinary behavior of the economy in the pandemic.

Our baseline specification includes the contemporaneous value and four lags of our shock variable, and four lags of the dependent variable. To investigate the effects of including more lags, we consider a specification with eight lags. Because our data begin in 1946:4, however, including the additional lags requires that our sample starts in 1948:4 rather than 1947:4, which causes the contemporaneous value of the 1947:4 shock to drop out of the sample period. To determine how much of any change in the results from adding the extra lags is due just to the change in the sample period, we also reestimate our baseline specification (with the four lags) with the sample restricted to start in 1948:4.

Our baseline specification allows for the possibility that monetary shocks affect the macroeconomy in the period when they occur. It is often argued, however, that it is reasonable to impose the restriction that this cannot occur—that is, to assume that changes in monetary policy cannot affect output or inflation in the quarter when they occur. Since this is fundamentally an empirical question, our preference is to let the data speak. However, we also consider a specification that imposes the restriction. To do this, we simply include not just the value of the dependent variable for periods $t-1$, ..., $t-4$, but also the period-t value.

Another variant we consider is to include four lags of both real GDP and inflation in both the real GDP and inflation regressions. Specifically, for the real GDP regressions, we add four lags of GDP price index inflation as additional controls; for each of the inflation regressions, we add four lags of real GDP as additional controls. This specification controls for the normal joint dynamics of both variables, and is closer to the spirit of conventional VARs.

For the inflation variables, we consider one more specification. Because we focus on inflation, our baseline specification does not allow for the possibility that the price *level* has predictive power for inflation. We therefore also consider a specification where we include not just four lags of inflation, but also one lag of the price level (in logarithms). Since inflation is proportional to the difference in the log price level, this is equivalent to including five lags of the log price level.

STANDARD ERRORS

In our baseline results, we report conventional standard errors. Our monetary shock series is not noticeably serially correlated—it consists of long strings of zeroes with an occasional one. As a result, serial correlation in the regression residuals should have little impact on the variance of the estimated effects of monetary policy shocks. And while there is some variation in macroeconomic volatility over the postwar period, it is not enormous and not obviously strongly correlated with our shock series. Moreover, when there are a small number of important observations (as in our case, with only ten monetary shocks), heteroskedasticity-robust standard errors are prone to volatility (Young 2019). For those reasons, we view conventional standard errors as the appropriate starting point.

Despite concerns about other approaches, we consider four other ways of computing the standard errors. Three are standard: heteroskedasticity-robust (or Huber-White), Newey-West with h lags, and Hansen-Hodrick with h lags. Our fourth approach is nonparametric. For concreteness, consider a given h. We randomly select 10 dates over the period 1947:4–2016:4 and construct a "pseudo" monetary shock series equal to 1 on only those 10 dates. We then estimate equation (1) in the paper and record the resulting estimate of the "effect" of the shock. We repeat this 100,000 times. Since the shocks are drawn completely at random, the *average* estimated effect is very close to zero. But the *variability* of the estimates provides a measure of how much the estimated effects of a monetary shock series like ours can vary due to chance correlation of the shocks with other forces affecting macroeconomic outcomes. We therefore use the standard deviation of the estimates as a nonparametric standard error. We view these nonparametric standard errors as the most useful alternative to the baseline ones.

RESULTS

Table B1 reports the results of these robustness exercises for real GDP and Table B2 reports results for the unemployment rate. For each specification, the tables report the effect after a year, the maximum estimated effect (negative for real GDP and positive for the unemployment rate) and the horizon at which it occurs, and the effect after 5 years. The changes relative to the baseline specification at those horizons provide a good sense of the changes at other horizons. For example, if the estimated impact is slightly larger than in the baseline at those three horizons, it is generally slightly larger at other horizons. (As we discuss below, however, the situation with regard to the alternative standard errors is more complicated.)

The basic message of the tables is that the results are very robust. For unemployment, the changes relative to the baseline are all minor. For real GDP, the biggest changes are that either restricting the contemporaneous effect to zero or including lags of inflation as well as real GDP reduces the estimated effects slightly. However, they remain similar to the baseline estimates and continue to be highly statistically significant.

The Huber-White, Newey-West, and Hansen-Hodrick standard errors show more variability—they are noticeably larger than the baseline standard errors at some horizons and noticeably lower at others. (The peak effects remain highly statistically significant regardless, however.) But we view this as more informative about practical problems with these standard errors in applications like ours than about the actual precision of our estimates. Since the variation in real GDP and unemployment relative to their current values is greater at longer horizons, it seems almost obvious that there is more uncertainty about the effects of monetary policy shocks at longer horizons. And indeed, the conventional standard errors for both real GDP and unemployment increase steadily with the horizon. But the Huber-White, Newey-West, and Hansen-Hodrick standard errors do not. For example, with the Huber-White (heteroskedasticity-robust) standard errors, the standard error of the estimated effect on real GDP is lower at 9 quarters than at 2, and lower at 20 quarters than at 12. The Huber-White standard errors for unemployment exhibit similar puzzling variability, as do the Newey-West and Hansen-Hodrick standard errors for both variables. The nonparametric standard errors, in contrast to these standard errors but like the conventional ones, rise steadily with the horizon. They are consistently slightly (about 5 to 10 percent) larger than the conventional standard errors.

The nonparametric approach also allows for a nonparametric test of statistical significance: we can ask how often choosing a set of 10 dates at random produces a peak estimated "effect" greater than the one we find with our series. This gives an estimate of how likely it would be to find effects as large as we do if monetary shocks actually had no impact. Recall that with our shock series, we find peak effects at 9 quarters for real GDP and 27 months for unemployment. With randomly chosen dates, for real GDP only 15 of the 100,000 draws (0.015 percent) yield a larger (negative) effect within the first 9 quarters, and only 402 (0.402 percent) yield one anytime over the full 20 quarters we consider. For unemployment, the fractions are 0.212 percent within the first 27 months, and 1.180 percent over the full 60 months. Thus, the effects we find are extremely unlikely to arise just by chance.

Table B3 presents the results for the three inflation measures. The one change from Tables B1 and B2 in what we report is that because we find a mild price puzzle, the table shows the maximum positive effect on inflation (and the horizon at which it occurs) in place of the effect after a year. As before, the changes relative to the baseline at the three horizons we report are a good guide to the changes at other horizons. As with the measures of real activity, most of the changes in specification have little effect. The biggest changes are in the direction of making the price puzzle weaker and the estimated negative impact on inflation larger. For all three inflation measures, including lags of real GDP in addition to inflation or controlling for the price level increases the estimated (negative) effect somewhat. For PCE price index inflation, these variations also attenuate the price puzzle somewhat. And for both PCE and core PCE price index inflation, restricting the contemporaneous impact on inflation to zero both attenuates the price puzzle and increases the estimated (negative) impacts at longer horizons.

The Huber-White, Newey-West, and Hansen-Hodrick standard errors again show patterns that suggest they are not highly reliable. The nonparametric standard errors are very close to the baseline ones for all three inflation measures. They are a few percent higher for GDP price index inflation and core PCE price index inflation, and generally very similar for PCE price index inflation. Not surprisingly, the chances of obtaining results as strong as our estimates by choosing dates at random are larger for inflation than for real activity. Nonetheless, they are still small. For both GDP and core PCE price index inflation, the chances of getting estimates as large and negative as our peak effects either anytime up to the time of the peak effects or over the full 20 quarters are about 2 percent; for PCE price index inflation, they are 5 percent through the time of the peak effect and 10 percent over the full 20 quarters.

III. FORECAST ERRORS

A simple way of visualizing how key outcome variable behave after our shocks is to plot errors from simple forecasting equations after the date of each shock. In particular, we estimate equation (1) from our paper with the shock variable (and its lags) omitted, for different values of h:

(B1) $$Y_{t+h} = \alpha^h + \sum_{k=1}^{K} \theta_k^h Y_{t-k} + e_t^h.$$

For each h, we then find the forecast errors (that is, the residuals) for the observations corresponding to the dates of our shocks (1947:4, 1955:3, and so on). Finally, we plot the sequence of errors (for h = 0, h = 1, ...) for each of those dates. For the 1972:1 shock, which is our only positive one, what we plot is −1 times the errors, so they can be more easily compared with those for the other shocks.

Figure B4 shows the results for real GDP. Not surprisingly, there is substantial variation across episodes. There is little systematic pattern in the quarter of the shock and the subsequent quarter, when output appears to be responding to factors other than the shocks. And there is wide variation in how the errors are changing at distant horizons, presumably reflecting the fact that there is wide variation in the shocks hitting the economy over a 5-year period. For example, the Korean War started 10 quarters after the 1947:4 shock, while the 1978:3 shock was followed by additional negative monetary shocks over the next three years. Nonetheless, the commonalities across the episodes are striking. From 6 to 10 quarters after the shocks, 90 percent of the forecast errors are negative (or positive for the 1972:1 shock). The only shock for which there is a substantial departure is the 1958:3 shock, where the forecast errors are large and positive for the first few quarters and then fall considerably but remain positive until quarter 9. Looking at the *changes* in the errors, in every case there is a substantial fall in the errors sometime in the first 10 quarters after the shock. The mildest is after the 1988:4 shock, where the swing is slightly under 3 percentage points (which occurs from $h = 5$ to $h = 9$).

Figure B5 shows the results for GDP price index inflation. On average the forecast errors are negative; that is, inflation generally fell relative to the univariate forecast after the negative monetary shocks (and rose after the positive one). But there is considerable variability. As with real GDP, there is no clear pattern at short horizons (in this case, for about the first year), and large variation at long horizons. For intermediate horizons, the forecast errors are generally negative, but there are exceptions. Most notably, inflation generally rose following the 1968, 1974, and (until year 5) 1978 shocks. And when it moved in the expected direction, the timing and magnitude varied substantially, ranging from quick, large movements after the 1947 and 1972 shocks to slow, small ones following the 1955 and 1988 shocks. All of this is consistent with our finding that the evidence points in the direction of contractionary monetary policy acting to reduce inflation, but with considerable uncertainty about the specifics of the effects.

Figure B6 shows the results for the monthly unemployment rate. After each shock, there are substantial positive forecast errors, suggesting that contractionary monetary policy shocks led to rises in unemployment. There is obviously substantial variation in timing—the positive forecast errors appeared fairly quickly in 1974, 1981, and (with a sign change) following the expansionary shock in 1972, and more slowly following the other stocks. Typically, the positive forecast errors are most noticeable between 18 and 36 months following the shock. The similarity of the forecast errors across the ten shocks is consistent with our finding that the estimated positive impact of a contractionary monetary shock on the unemployment rate is highly statistically significant.

Finally, we also consider another robustness exercise for the unemployment rate. As we do in the paper for quarterly real GDP and the GDP price index, we examine the impact of leaving out the shocks one-by-one. That is, we consider ten variants of our shock series, where each variant eliminates one of the shocks, but leaves the other shocks unchanged. Figure B7 shows the results. Leaving out the shocks in 1947 or 1972 results in somewhat larger estimates of the impact of a monetary shock; leaving out the 1979 Volcker shock results in somewhat smaller estimates. But overall, the results for the different variants are remarkably similar. This suggests that the estimated impact of a monetary shock on the unemployment rate is not being driven by any one shock.

IV. DATA SOURCES AND CONSTRUCTION

This section describes the data sources used in the robustness exercises. The sources for the series used in both the paper and this appendix are described in the data appendix to the paper.[1]

EMPLOYMENT

The monthly employment data for 1946:10 to 2016:12 are from the U.S. Bureau of Labor Statistics (BLS), series CES0000000001, seasonally adjusted, thousands, downloaded

[1] The data appendix to the paper describes the data sources through the end of 2016. For robustness exercises where we include data through the end of 2019, we continue with the same modern sources, downloaded on the same date as the data through 2016.

11/20/2022. The data come from the Current Employment Statistics program, and report all employees on nonfarm payrolls. All BLS data are from https://www.bls.gov/data.

INDUSTRIAL PRODUCTION

The monthly industrial production data for 1946:10 to 2016:12 are from the U.S. Board of Governors of the Federal Reserve System, https://www.federalreserve.gov/datadownload, G.17 Industrial Production and Capacity Utilization, series IP.B50001.S, seasonally adjusted, index, 2017=100, downloaded 11/20/2022. The series is the total index of industrial production.

PCE PRICE INDEX

The monthly data for the PCE price index for 1959:1 to 2016:12 are from the U.S. Bureau of Economic Analysis (BEA), Table 2.8.4, seasonally adjusted, index, 2012=100, downloaded 11/20/2022. All BEA data are from https://www.bea.gov/itable/national-gdp-and-personal-income.

To construct a proxy for the PCE price index for 1947:1 to 1958:12, we take the ratio of the PCE price index to the Consumer Price Index for All Urban Consumers (CPI-U) at a monthly frequency in 1959:1, and multiply it by the CPI-U for 1947:1 to 1958:12. The CPI-U data are from the BLS, series CUSR0000SA0, seasonally adjusted, index, 1982–84=100, downloaded 11/20/2022.

To continue the series back to 1946:9, we take the ratio of the extended PCE price index in 1947:1 to the non-seasonally-adjusted CPI-U in 1947:1, and multiply it by the unadjusted CPI-U for 1946:9 to 1946:12. The CPI-U data are from the BLS, series CUUR0000SA0, not seasonally adjusted, index, 1982–84=100, downloaded 11/20/2022.

Inflation at an annual rate is calculated as the difference in logarithms times 1200.

CORE PCE PRICE INDEX

The monthly data on the PCE price index excluding food and energy (Core PCE) for 1959:1 to 2016:12 are from the BEA, Table 2.8.4, seasonally adjusted, index, 2012=100, downloaded 11/20/2022.

To construct a proxy for the core PCE price index for 1957:1 to 1958:12, we take the ratio of the core PCE price index to the CPI-U less food and energy in 1959:1, and multiply it by the CPI-U less food and energy for 1957:1 to 1958:12. The CPI-U less food and energy data are from the BLS, series CUSR0000SA0L1E, seasonally adjusted, index, 1982–84=100, downloaded 11/20/2022.

To continue the series back to 1947:1, we take the ratio of the extended core PCE price index to the CPI-U less food in 1957:1, and multiply it by the CPI-U less food for 1947:1 to 1956:12. The CPI-U less food data are from the BLS, series CUSR0000SA0L1, seasonally adjusted, index, 1982–84=100, downloaded 11/20/2022.

To continue the series back to 1946:9, we take the ratio of the extended core PCE price index to the non-seasonally-adjusted CPI-U less food in 1947:1, and multiply it by the non-seasonally-adjusted CPI-U less good for 1946:9 to 1946:12. The CPI-U less food data are from the BLS, series CUUR0000SA0L1, not seasonally adjusted, index, 1982–84=100, downloaded 11/20/2022.

Inflation at an annual rate is calculated as the difference in logarithms times 1200.

REFERENCE

U.S. Board of Governors of the Federal Reserve System. 2022–2023. https://www.federalreserve.gov/datadownload. Various data series.

U.S. Bureau of Economic Analysis. 2022. https://www.bea.gov/itable/national-gdp-and-personal-income. Various data series.

U.S. Bureau of Labor Statistics. 2022–2023. https://www.bls.gov/data. Various data series.

Young, Alwyn. 2019. "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results." *Quarterly Journal of Economics* 134 (2): 557–98.

Table B1
Robustness Results for Quarterly Real GDP

	Specification	Effect at 4 Quarters	Maximum Negative Effect	Effect at 20 Quarters
(1)	Baseline	−2.68	−4.45	−1.42
		(0.82)	(1.09)	(1.51)
			[9]	
(2)	Equal observations at each horizon	−2.69	−4.45	−1.42
		(0.85)	(1.11)	(1.51)
			[9]	
(3)	Use data through 2019:4	−2.68	−4.45	−1.43
		(0.82)	(1.07)	(1.48)
			[9]	
(4)	Include 8 lags	−2.77	−4.51	−1.87
		(0.86)	(1.22)	(1.61)
			[10]	
(4')	Baseline, same sample as (4)	−3.02	−4.60	−2.05
		(0.86)	(1.22)	(1.59)
			[10]	
(5)	Contemporaneous effect restricted to zero	−2.35	−4.22	−1.15
		(0.73)	(1.06)	(1.48)
			[9]	
(6)	Include lags of both real GDP and GDP price index inflation	−2.19	−3.77	−0.77
		(0.85)	(1.12)	(1.56)
			[9]	
(7)	Alternative standard errors, baseline specification			
	Baseline (conventional)	(0.82)	(1.09)	(1.51)
	Huber-White	(0.85)	(0.73)	(1.20)
	Newey-West (*h* lags)	(0.86)	(0.99)	(1.33)
	Hansen-Hodrick (*h* lags)	(0.87)	(1.11)	(1.40)
	Nonparametric	(0.91)	(1.21)	(1.61)

Notes: Numbers in parentheses are standard errors. Numbers in brackets are the horizon. See text for explanation of the nonparametric standard errors.

Table B2
Robustness Results for the Monthly Unemployment Rate

	Specification	Effect at 12 Months	Maximum Positive Effect	Effect at 60 Months
(1)	Baseline	0.68	1.60	0.12
		(0.35)	(0.46)	(0.52)
			[27]	
(2)	Equal observations at each horizon	0.66	1.57	0.12
		(0.36)	(0.46)	(0.52)
			[27]	
(3)	Use data through 2019:12	0.69	1.64	0.19
		(0.35)	(0.45)	(0.52)
			[27]	
(4)	Include 24 lags	0.73	1.70	0.35
		(0.36)	(0.52)	(0.55)
			[32]	
(4')	Baseline, same sample as (4)	0.81	1.68	0.40
		(0.36)	(0.50)	(0.53)
			[32]	
(5)	Contemporaneous effect restricted to zero	0.62	1.57	0.11
		(0.33)	(0.45)	(0.52)
			[27]	
(6)	Alternative standard errors, baseline specification			
	Baseline (conventional)	(0.35)	(0.46)	(0.52)
	Huber-White	(0.45)	(0.32)	(0.52)
	Newey-West (*h* lags)	(0.44)	(0.38)	(0.63)
	Hansen-Hodrick (*h* lags)	(0.44)	(0.42)	(0.65)
	Nonparametric	(0.37)	(0.50)	(0.55)

Notes: Numbers in parentheses are standard errors. Numbers in brackets are the horizon. See text for explanation of the nonparametric standard errors.

Table B3
Robustness Results for Quarterly Measures of Inflation

	Specification	Maximum Positive Effect	Maximum Negative Effect	Effect at 20 Quarters
		A. GDP Price Index Inflation		
(1)	Baseline	1.09	−1.84	−1.37
		(0.59)	(0.64)	(0.67)
		[2]	[14]	
(2)	Equal observations at each horizon	1.07	−1.86	−1.37
		(0.61)	(0.65)	(0.67)
		[2]	[14]	
(3)	Use data through 2019:4	1.09	−1.84	−1.36
		(0.59)	(0.63)	(0.66)
		[2]	[14]	
(4)	Include 8 lags	1.36	−1.46	−1.02
		(0.60)	(0.63)	(0.70)
		[2]	[11]	
(4')	Baseline, same sample as (4)	1.18	−1.65	−1.16
		(0.61)	(0.66)	(0.70)
		[2]	[12]	
(5)	Contemporaneous effect restricted to zero	1.05	−1.86	−1.39
		(0.53)	(0.64)	(0.67)
		[2]	[14]	
(6)	Include lags of both real GDP and inflation	0.94	−2.06	−1.60
		(0.59)	(0.64)	(0.66)
		[2]	[14]	
(7)	Include lagged price level	1.04	−2.02	−1.59
		(0.59)	(0.62)	(0.64)
		[2]	[14]	
(8)	Alternative standard errors, baseline specification			
	Baseline (conventional)	(0.59)	(0.64)	(0.67)
	Huber-White	(0.62)	(0.58)	(0.88)
	Newey-West (*h* lags)	(0.61)	(0.64)	(0.50)
	Hansen-Hodrick (*h* lags)	(0.63)	(0.50)	(0.30)
	Nonparametric	(0.59)	(0.68)	(0.71)

Table B3 (continued)

	Specification	Maximum Positive Effect	Maximum Negative Effect	Effect at 20 Quarters
		B. PCE Price Index Inflation		
(1)	Baseline	0.80 (0.77) [3]	−1.83 (0.76) [9]	−1.30 (0.76)
(2)	Equal observations at each horizon	0.79 (0.79) [3]	−1.82 (0.76) [9]	−1.30 (0.76)
(3)	Use data through 2019:4	0.80 (0.77) [3]	−1.82 (0.74) [9]	−1.29 (0.75)
(4)	Include 8 lags	0.98 (0.71) [2]	−1.57 (0.74) [11]	−0.82 (0.78)
(4')	Baseline, same sample as (4)	0.90 (0.72) [2]	−1.62 (0.77) [12]	−0.95 (0.79)
(5)	Contemporaneous effect restricted to zero	0.58 (0.71) [3]	−1.85 (0.74) [14]	−1.36 (0.76)
(6)	Include lags of both real GDP and inflation	0.64 (0.77) [3]	−2.11 (0.75) [9]	−1.55 (0.76)
(7)	Include lagged price level	0.73 (0.77) [3]	−2.04 (0.74) [9]	−1.56 (0.74)
(8)	Alternative standard errors, baseline specification			
	Baseline (conventional)	(0.77)	(0.76)	(0.76)
	Huber-White	(0.68)	(1.09)	(1.12)
	Newey-West (*h* lags)	(0.69)	(1.24)	(0.76)
	Hansen-Hodrick (*h* lags)	(0.70)	(1.20)	(0.52)
	Nonparametric	(0.76)	(0.77)	(0.78)

Table B3 (continued)

	Specification	Maximum Positive Effect	Maximum Negative Effect	Effect at 20 Quarters
		C. Core PCE Price Index Inflation		
(1)	Baseline	1.39	−1.64	−1.13
		(0.45)	(0.56)	(0.60)
		[1]	[10]	
(2)	Equal observations at each horizon	1.38	−1.64	−1.13
		(0.46)	(0.57)	(0.60)
		[1]	[10]	
(3)	Use data through 2019:4	1.39	−1.63	−1.12
		(0.45)	(0.55)	(0.59)
		[1]	[10]	
(4)	Include 8 lags	1.20	−1.12	−0.66
		(0.47)	(0.58)	(0.61)
		[1]	[14]	
(4')	Baseline, same sample as (4)	1.20	−1.29	−0.75
		(0.47)	(0.61)	(0.62)
		[1]	[14]	
(5)	Contemporaneous effect restricted to zero	0.92	−1.87	−1.28
		(0.35)	(0.55)	(0.60)
		[1]	[10]	
(6)	Include lags of both real GDP and inflation	1.33	−1.81	−1.29
		(0.45)	(0.56)	(0.60)
		[1]	[10]	
(7)	Include lagged price level	1.37	−1.79	−1.32
		(0.45)	(0.55)	(0.58)
		[1]	[10]	
(8)	Alternative standard errors, baseline specification			
	Baseline (conventional)	(0.45)	(0.56)	(0.60)
	Huber-White	(0.58)	(1.05)	(0.81)
	Newey-West (*h* lags)	(0.56)	(1.13)	(0.50)
	Hansen-Hodrick (*h* lags)	(0.55)	(0.99)	(0.40)
	Nonparametric	(0.45)	(0.60)	(0.64)

Notes: Numbers in parentheses are standard errors. Numbers in brackets are the horizon. See text for explanation of the nonparametric standard errors.

Figure B1
Response of Employment to a Monetary Policy Shock

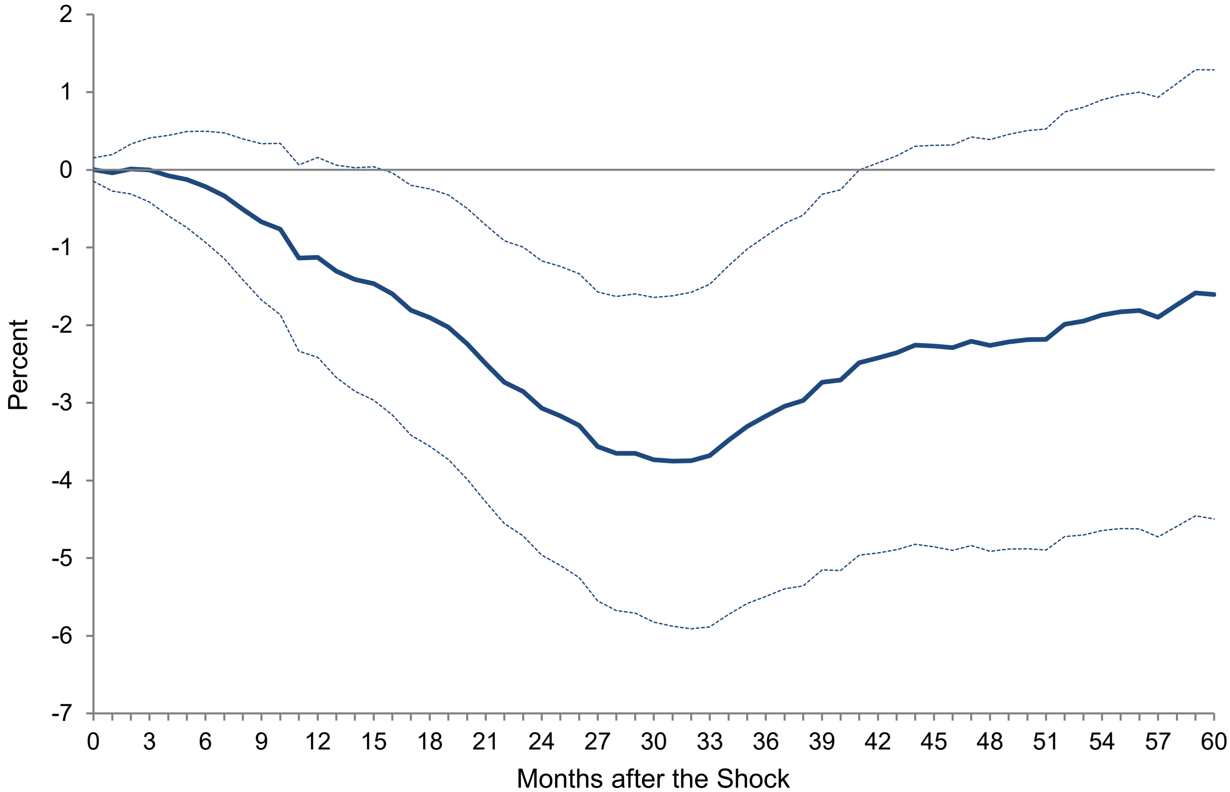

Notes: The figure shows the results of estimating equation (1) of the paper for horizons 0 to 60. The dependent variable is the log of payroll employment. The dotted lines show the two-standard-error confidence bands. The new shock series is given in Table 2 of the paper. See the text of the paper for details of the estimation and Section IV of this appendix for the source of the employment series.

Figure B2
Response of Industrial Production to a Monetary Policy Shock

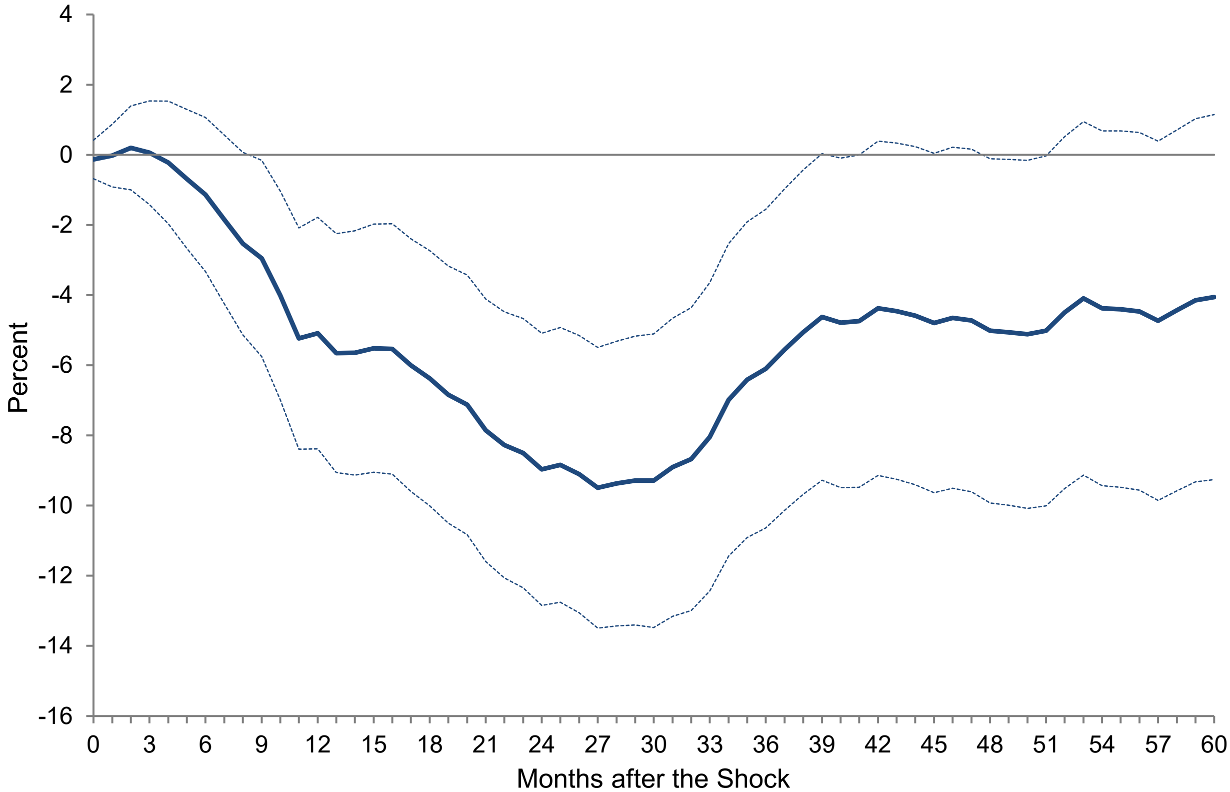

Notes: The figure shows the results of estimating equation (1) of the paper for horizons 0 to 60. The dependent variable is the log of industrial production. The dotted lines show the two-standard-error confidence bands. The new shock series is given in Table 2 of the paper. See the text of the paper for details of the estimation and Section IV of this appendix for the source of the industrial production series.

Figure B3
Response of Monthly Inflation to a Monetary Policy Shock

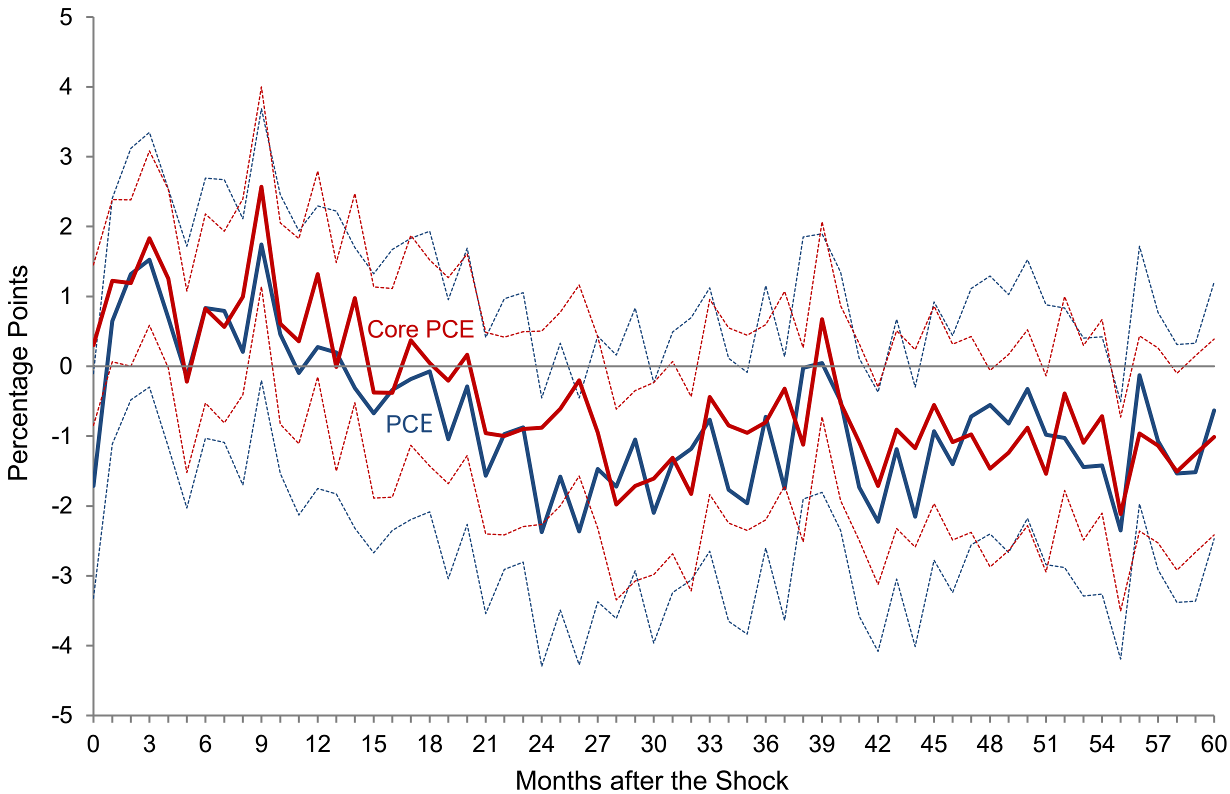

Notes: The figure shows the results of estimating equation (1) of the paper for horizons 0 to 60. The dependent variable is the inflation rate, measured using either the PCE price index or the PCE price index excluding food and energy (core PCE). The dotted lines show the two-standard-error confidence bands. The new shock series is given in Table 2 of the paper. See the text of the paper for details of the estimation and Section IV of this appendix for the sources of the monthly PCE and core PCE price index series.

Figure B4
Forecast Errors for Real GDP

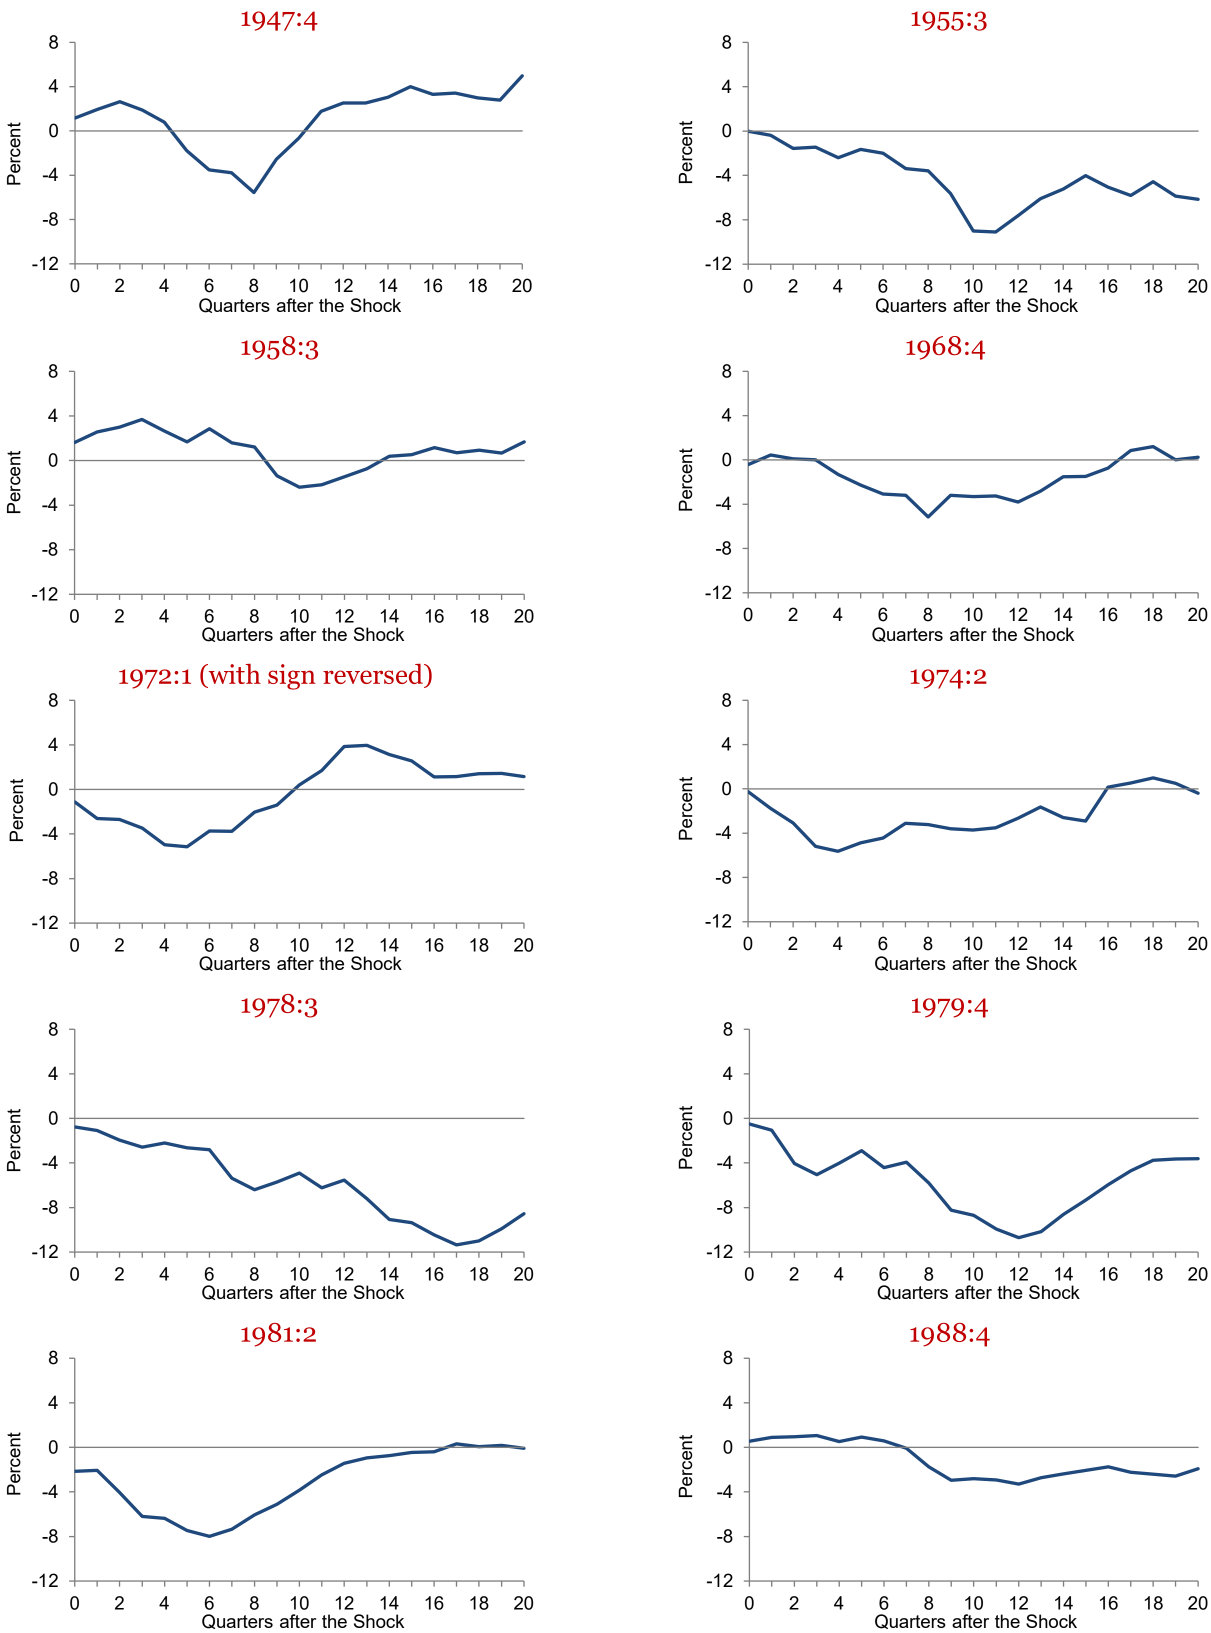

Notes: See the text of this appendix for details of the estimation, and the data appendix of the paper for the sources of the real GDP data.

Figure B5
Forecast Errors for GDP Price Index Inflation

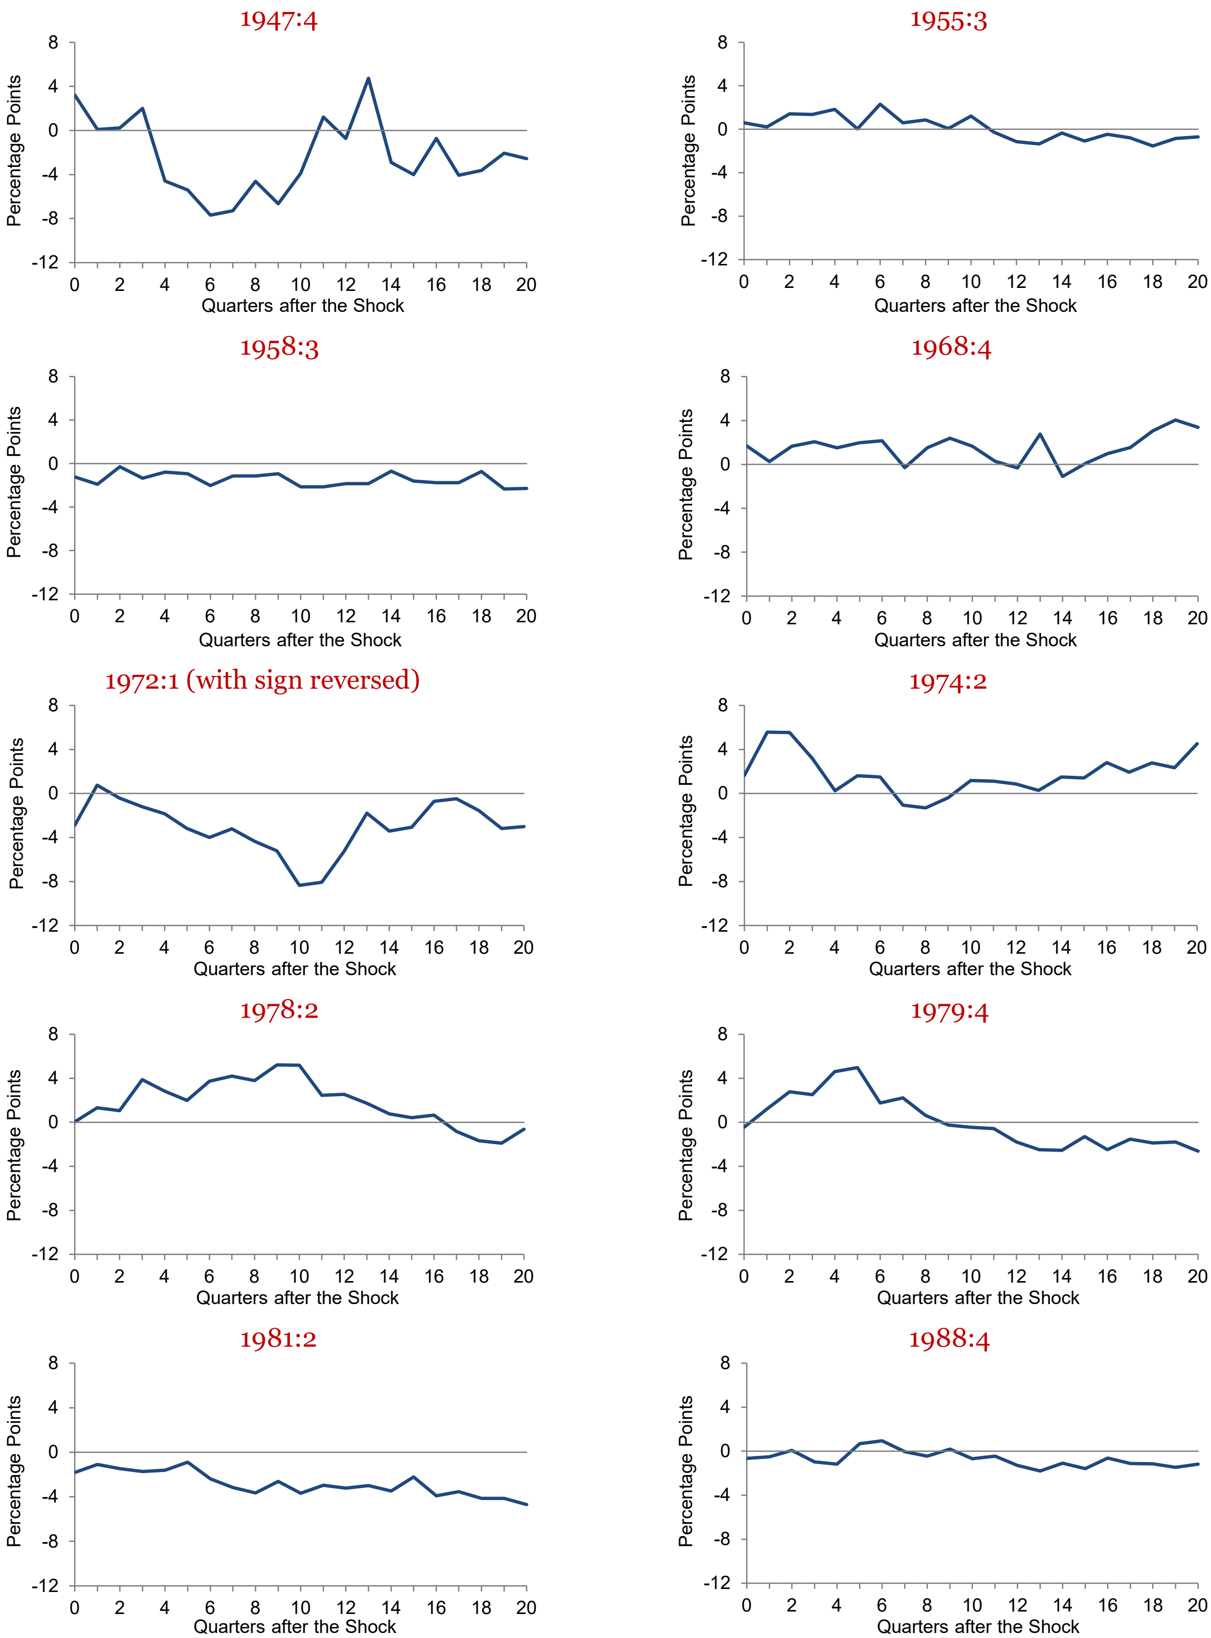

Notes: See the text of this appendix for details of the estimation, and the data appendix of the paper for the sources of the GDP price index inflation data.

Figure B6
Forecast Errors for the Unemployment Rate

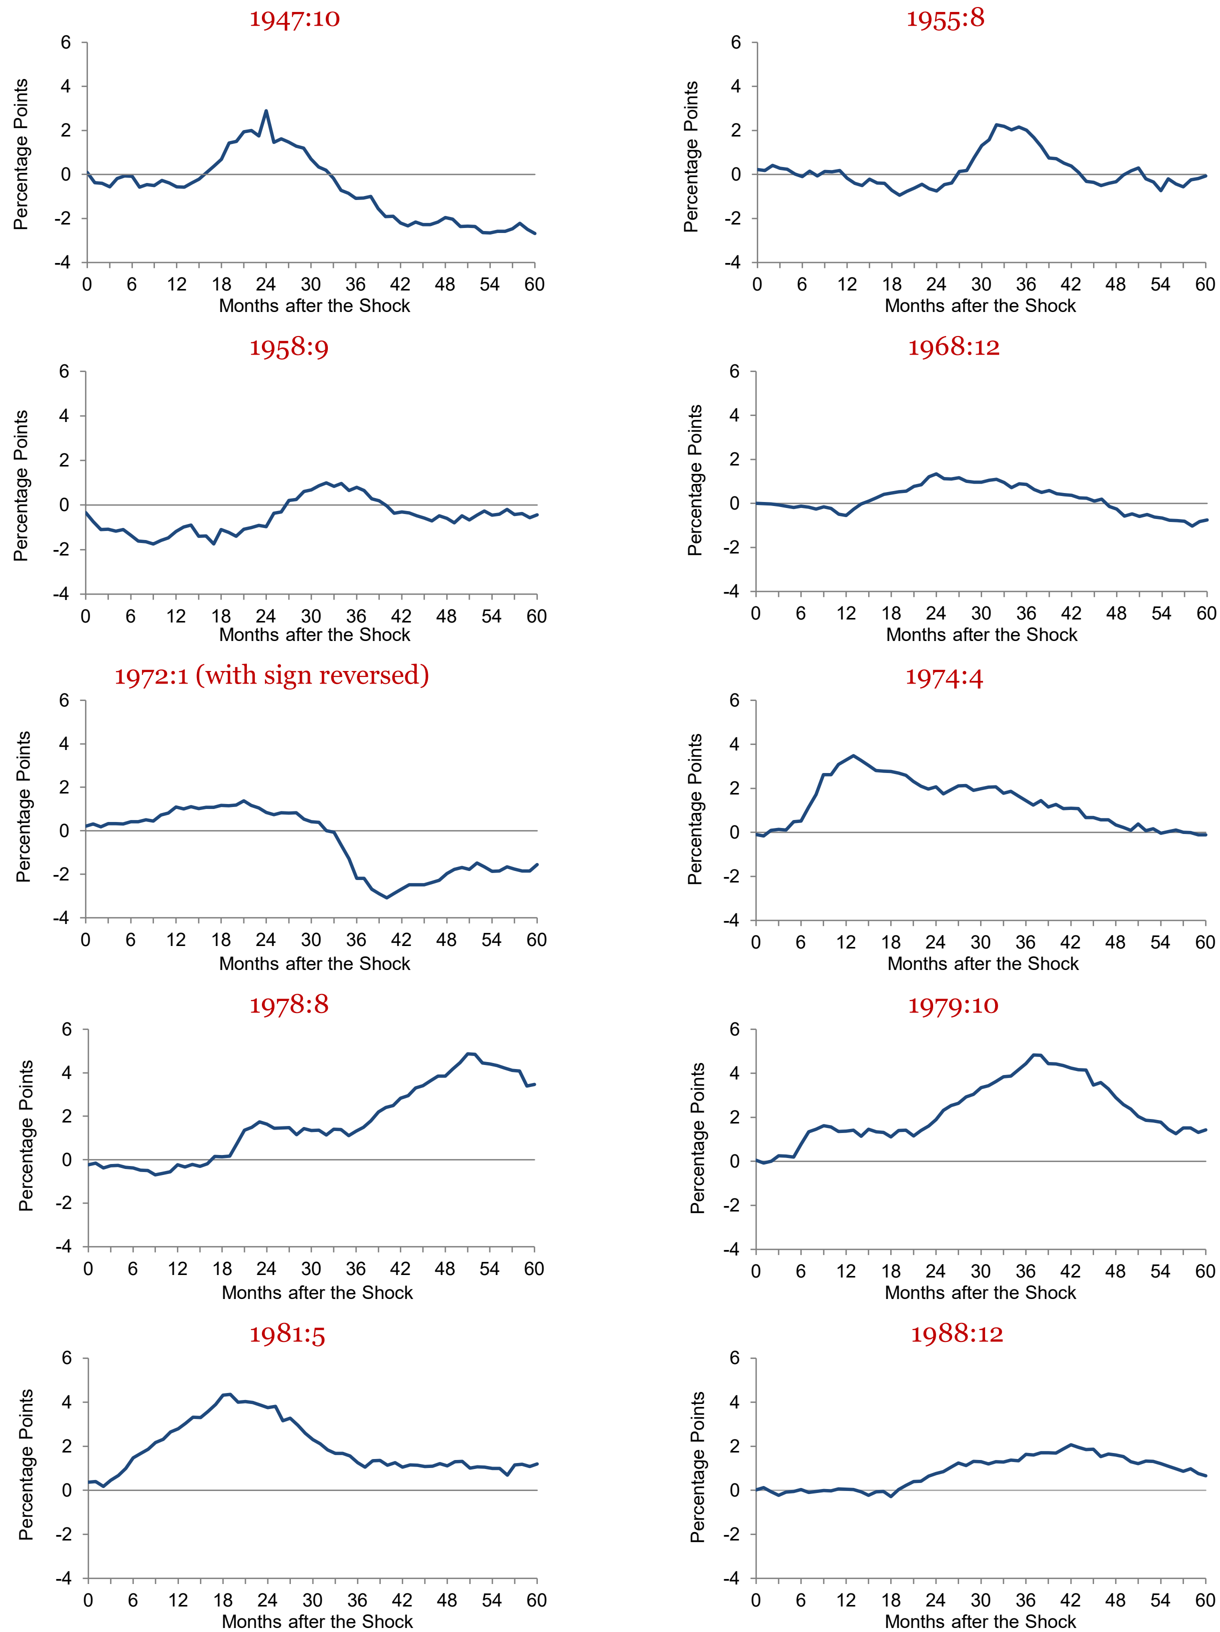

Notes: See the text of this appendix for details of the estimation, and the data appendix of the paper for the sources of the unemployment rate data.

Figure B7
Response of the Unemployment Rate to a Monetary Policy Shock,
Leaving out One Shock at a Time

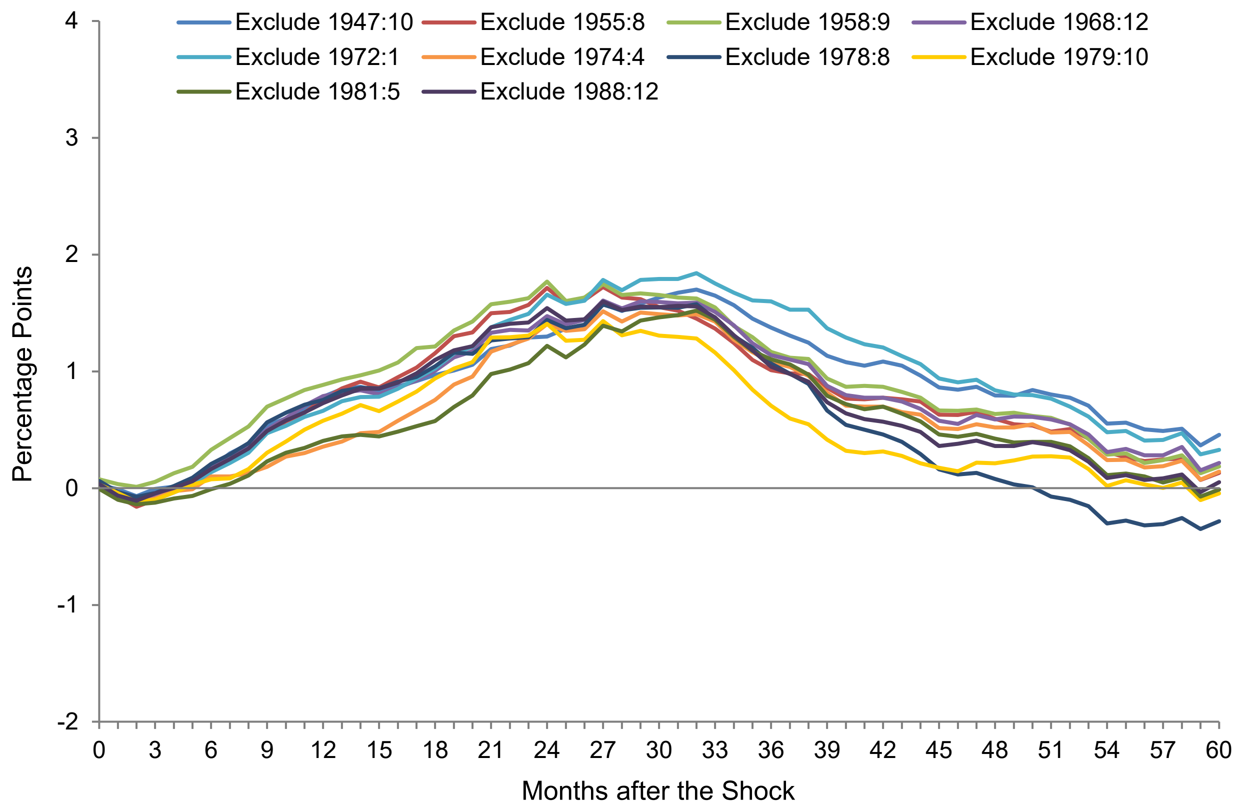

Notes: The figure shows the results of estimating equation (1) of the paper for horizons 0 to 20 for ten variants of the new shock series. The variants are formed by sequentially eliminating one of the shocks. The dependent variable is the level of the unemployment rate. The new shock series is given in Table 2 of the paper. See the text of the paper for details of the estimation and the data appendix of the paper for the sources of the unemployment rate series.