Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

The Passthrough of Labor Costs to Price Inflation

Ekaterina V. Peneva and Jeremy B. Rudd

2015-042

Please cite this paper as:
Peneva, Ekaterina V. and Jeremy B. Rudd (2015). "The Passthrough of Labor Costs to Price Inflation," Finance and Economics Discussion Series 2015-042. Washington: Board of Governors of the Federal Reserve System, http://dx.doi.org/10.17016/FEDS.2015.042.

The Passthrough of Labor Costs to Price Inflation

Ekaterina V. Peneva
Federal Reserve Board*

Jeremy B. Rudd
Federal Reserve Board**

May 21, 2015

Abstract

We use a time-varying parameter/stochastic volatility VAR framework to assess how the passthrough of labor costs to price inflation has evolved over time in U.S. data. We find little evidence that changes in labor costs have had a material effect on price inflation in recent years, even for compensation measures where some degree of passthrough to prices still appears to be present. Our results cast doubt on explanations of recent inflation behavior that appeal to such mechanisms as downward nominal wage rigidity or a differential contribution of long-term and short-term unemployed workers to wage and price pressures.

*E-mail: ekaterina.v.peneva@frb.gov.

**Corresponding author. Mailing address: Mailstop 61, 20th and C Streets NW, Washington, DC 20551. E-mail: jeremy.b.rudd@frb.gov. We thank David Lebow, Deb Lindner, and seminar participants at Auburn University for helpful comments on earlier versions of this work. We also express our considerable gratitude to Todd Clark for providing us with his computer algorithms for estimating TVP/SV VAR models. (Any remaining errors are ours.) The analysis and conclusions set forth are our own and do not necessarily reflect the views of the Board of Governors or the staff of the Federal Reserve System.

I Introduction

Many formal and informal descriptions of inflation dynamics assign an important explicit or implicit role to labor costs. Intuitively, labor compensation should be a key determinant of firms' pricing behavior as, in the aggregate, it represents about two-thirds of firms' total costs of production. More formally, economic theory suggests that increases in labor costs in excess of productivity gains should put upward pressure on prices; hence, many older theoretical and empirical models (including the large-scale econometric models of the 1970s and 1980s) assumed that prices are determined as a markup over unit labor costs. Similarly, many empirical implementations of the new-Keynesian Phillips curve have used real unit labor costs as a proxy for real marginal costs, which are the theoretical driver of inflation in these models.

Wage-based explanations of inflation dynamics have seen increased prominence of late, as a number of observers have sought to use developments in the labor market to explain why price inflation did not decline by as much as conventional models would have predicted following the 2007–2009 recession (the so-called "missing disinflation" puzzle).[1] First, some analysts have argued that the presence of downward nominal wage rigidity has propped up aggregate wage inflation to an unusual degree in recent years, which has in turn led price inflation to decline by less than would be expected given the magnitude and persistence of the shortfall in real activity that resulted from the Great Recession.[2] Second, some researchers (for example, Gordon, 2013) have argued that

[1]See Ball and Mazumder (2011), Watson (2014), and Coibion and Gorodnichenko (2015) for some representative discussions of the missing disinflation puzzle.

[2]For example, Paul Krugman has made this type of argument in his popular writings. Formal modelling suggests that the effects of downward nominal wage rigidity could be more complicated: In Daly and Hobijn's (2014) theoretical analysis, downward nominal wage rigidity props up wage inflation in a recession; as the labor market recovers, however, the existence of "pent-up" wage cuts puts downward pressure on wages even as the unemployment rate is falling.

recent inflation behavior can be better explained if real activity is measured in terms of the short-term unemployment rate (that is, the share of the labor force unemployed for 26 weeks or less), on the grounds that the long-term unemployed seem to put less (or no) downward pressure on inflation.

Ultimately, these proposed explanations for the recent behavior of price inflation only make sense if there is an economically significant influence of compensation costs on prices. Regarding the first explanation, it is clear that downward nominal wage rigidity can have an important effect on inflation dynamics only if price setting is closely connected to labor costs. Regarding the second explanation, we would not expect a rise in long-term unemployment to have a smaller effect on aggregate demand than a rise in short-term unemployment: Presumably, the long-term unemployed—who have suffered a relatively larger and more persistent shock to their permanent income—would reduce their contribution to aggregate demand to a *greater* degree than would the short-term unemployed. Hence, it seems difficult to invoke the idea that the short-term unemployment rate provides a better gauge of the level of real activity that is relevant for price inflation without simultaneously arguing that the fundamental source of this relation is the differential effect that the short- and long-term unemployed have on *wage* inflation (and, again, that labor costs are an important determinant of prices).

In this paper, we explore whether there is a tight—and stable—link between labor costs and price inflation. Overall, we find it difficult to discern an important independent effect of changes in average labor costs on aggregate price inflation in recent years once we account for labor market slack. In particular, we find evidence either that the passthrough of labor costs to prices has fallen over the past several decades or—for compensation measures where there is still evidence of passthrough—that changes in labor costs have had essentially no material effect on price inflation in recent years.

A number of authors have examined whether movements in labor costs lead changes in price inflation. Although the results are often specific to various methodological choices and data definitions, the general conclusion that emerges from this literature is that there appears to be a break in the relation between labor costs and broad price measures, with changes in labor costs having little or no predictive power for price inflation after the early 1980s. For example, Mehra (2000) divides the postwar period into three subperiods and finds that wage inflation helps predict price inflation only in the middle (high-inflation) subperiod of 1966–1983; similarly, Emery and Chang (1996) find that labor costs are only useful in forecasting core consumer price inflation in the 1970s. Our work complements and extends this earlier research in two ways. First, the empirical framework that we use to gauge how the passthrough of labor costs to prices has evolved over time—a VAR model with time-varying parameters and stochastic volatility—has not, to our knowledge, been previously employed for this purpose.[3] This framework allows us to better identify the source of any changes in passthrough that we observe, as well as their implications for inflation dynamics. Second, our analysis covers a more-recent period, one that includes the Great Recession and a significant portion of the subsequent recovery.

While it casts doubt on explanations based on downward nominal wage rigidity or similar labor-market developments, our finding that the behavior of labor costs appears to have had little material influence on price inflation leaves unanswered the question of how to explain the evolution of inflation following the Great Recession. Based on our results, we conclude that the dynamics of prices and labor costs have changed significantly in recent decades, such that the stochastic trends for price inflation and labor cost growth

[3]Clark and Terry (2010) use this approach to analyze time variation in the passthrough of energy price changes to core consumer prices.

have both been essentially constant since the mid-1990s. As a result, price inflation now responds less persistently to changes in real activity or costs; at the same time, the joint dynamics of inflation and compensation no longer manifest the type of wage–price spiral that was evident in earlier decades. Hence, the recent behavior of inflation (and our inability to find an important independent role for labor costs in driving inflation movements) reflects a change to the inflation process that predates the 2007–2009 recession, not something specific to the Great Recession itself.

II Empirical framework and data

We use a time-varying parameter/stochastic volatility vector autoregression model (a TVP/SV VAR) to examine whether and to what degree the passthrough of labor costs to price inflation has changed over time. In general, an n-variable recursively identified VAR can be written as

$$\begin{aligned} y_t^1 &= a_0^1 + A^{11}(L)y_{t-1}^1 + A^{12}(L)y_{t-1}^2 + \cdots + A^{1n}(L)y_{t-1}^n + \varepsilon_t^1 \\ y_t^2 &= a_0^2 + A^{21}(L)y_{t-1}^1 + A^{22}(L)y_{t-1}^2 + \cdots + A^{2n}(L)y_{t-1}^n + a_1^2 y_t^1 + \varepsilon_t^2 \\ &\vdots \\ y_t^n &= a_0^n + A^{n1}(L)y_{t-1}^1 + A^{n2}(L)y_{t-1}^2 + \cdots + A^{nn}(L)y_{t-1}^n + a_1^n y_t^1 + \\ &\quad a_2^n y_t^2 + \cdots + a_{n-1}^n y_t^{n-1} + \varepsilon_t^n, \end{aligned} \tag{1}$$

where the $A^{ij}(L)$ terms denote lag polynomials and ε_t^i is the structural residual associated with equation i. (Note that as the system is written, the variables are ordered $y_t^1, y_t^2, \ldots, y_t^n$.) In the TVP/SV framework we consider, the values of $A^{ij}(L)$ and a_j^i and the standard deviations of the ε_t^i terms are allowed to drift over time (they are modelled as random walks). Hence, by using the relevant sets of parameter values we can examine impulse response functions at various points in time; similarly, we can use the VAR to decompose the historical movements in a given variable into the cumulative contributions of the various structural shocks. In addition, the model can be used to produce estimates of the variables' stochastic trends. Following Cogley, *et al.* (2010), write the VAR in its companion form as

$$z_{t+1} = \mu_t + B_t z_t + e_{t+1}, \tag{2}$$

where z_t stacks the current and lagged values of the variables y_t^i, μ_t contains the (time-varying) intercepts from each VAR equation, and B_t contains the VAR's autoregressive parameters (which are also time-varying). At time t, we can obtain estimates of the stochastic trends $\bar{z}_t$ from

$$\bar{z}_t = (I - B_t)^{-1}\mu_t, \tag{3}$$

where I denotes the identity matrix.[4]

The TVP/SV approach complements alternative approaches to evaluating changes over time in the passthrough of labor costs to inflation, such as examining models estimated over rolling samples or specified subperiods. In a rolling regression, coefficient estimates can fluctuate purely because of sampling variability; by explicitly modelling parameter drift and using information from the full sample, the TVP/SV approach can in principle provide a clearer picture of the amount and type of drift that is truly present. Likewise, in the TVP/SV approach the timing of any parameter shifts is determined by the data, rather than by the analyst's choice of estimation subperiods. Finally, the TVP/SV approach allows us to model changes in the volatility of shocks over time, which can be important in determining whether observed changes in passthrough actually reflect

[4]Note that the trend definition being used here is analogous to the Beveridge–Nelson concept.

parameter breaks.[5]

The baseline VAR system that we consider is a four-variable, two-lag, quarterly model consisting of weighted relative import price inflation, a measure of trend unit labor cost growth, core price inflation, and an unemployment gap, with that causal ordering.[6] (All growth rates are defined as annualized log differences.) We include a relative import price term to control for the effect of an important component of non-labor costs on price inflation. The unemployment gap, which we include to capture the degree of labor- and product-market slack in the economy, is defined as the difference between the total civilian unemployment rate and the Congressional Budget Office's (CBO's) estimate of the short-term natural rate of unemployment.[7] The core inflation measure that we use is the market-based component of the core PCE price index—that is, the chain price index for market-based personal consumption expenditures excluding prices for energy and food at home.[8] Finally, we estimate the model using Clark and Terry's (2010) implementation of

[5]Of course, these advantages come at a cost inasmuch as a TVP/SV model is more difficult to estimate; in addition, the model's dynamic structure will tend to be relatively parsimonious compared with other types of empirical inflation equations (such as a Phillips curve).

[6]Our choice of lag length is informed by applying the Schwarz criterion to a constant-coefficient VAR estimated over the full sample; for each of the measures of trend unit labor costs that we use in our analysis, this criterion was minimized by a two-lag system.

[7]This is the same gap definition used by Coibion and Gorodnichenko (2015) in documenting the presence of "missing disinflation" in the aftermath of the Great Recession. The CBO defines the natural rate as "the estimated rate of unemployment arising from all sources except fluctuations in aggregate demand," with the short-term variant defined to include structural factors that act to temporarily boost the natural rate relative to its long-term level. (For reference, note that the CBO's short-term natural rate estimate rises by a percentage point from 2008 to 2012, peaking at 6 percent.) The short-term natural rate is intended to be compared to the *total* unemployment rate to obtain a measure of aggregate labor market slack; it has no connection to the short-term *unemployment* rate—which, as noted above, is defined as the fraction of the labor force unemployed for 26 weeks or less.

[8]We confine our attention to market-based prices because several nonmarket components of consumption are priced using input cost indexes that are in turn derived from wage or compensation measures. (In total, core market-based prices account for nearly 90 percent of the overall core PCE price index.) Note that before the core inflation measure is used in the VAR, we subtract out Blinder and Rudd's (2013) estimates of the effects of the Nixon-era price controls.

the Metropolis-within-Gibbs posterior sampler.[9]

We consider two alternative measures of trend unit labor costs for our analysis; in each case, trend unit labor cost growth is defined by subtracting an estimate of the trend growth rate of average labor productivity for the nonfarm business sector from a measure of hourly compensation growth.[10] The first compensation measure that we use, hourly compensation for the nonfarm business sector, is taken from the Productivity and Costs (P&C) report constructed by the Bureau of Labor Statistics (BLS). The P&C series includes wage and salary payments to employees (derived from source data that are benchmarked to full-universe tax records), benefit costs, and an imputation for the portion of proprietors' income that is attributable to labor. This series therefore represents a relatively comprehensive measure of labor-related production costs. The second measure that we use, the Employment Cost Index (ECI) for private industry workers, also includes wage and salary payments and benefit costs, though its coverage excludes proprietors, self-employed workers, and those with substantial discretion over their own pay. More importantly, the ECI uses fixed weights for industry and occupational groups to control for the effect of changes in the mix of jobs on measured hourly compensation.[11]

It is not clear *a priori* which measure of hourly compensation—the P&C measure or the ECI—provides a better estimate of the compensation costs that are relevant for firms' pricing decisions. Although the ECI, by controlling for the effects of mix-shifts on

[9]See the Appendix for additional details regarding the estimation procedure and variable definitions.

[10]We obtain an estimate of trend productivity growth by applying a low-pass filter to actual productivity growth (see the Appendix for details). Our use of trend unit labor costs is informed by the fact that in other contexts, it is difficult to find an influence of *actual* unit labor cost growth on price inflation once we condition on trend unit labor cost growth, likely because actual productivity growth is extremely noisy at a quarterly frequency.

[11]Additional background on the P&C measure can be found in a March 11, 2008 BLS note entitled "Technical Information About the BLS Major Sector Productivity and Costs Measures" (available at `www.bls.gov/lpc/lpcmethods.pdf`); chapter 8 of the BLS *Handbook of Methods* discusses how the ECI is constructed.

compensation, might seem to smooth away important variation in labor costs, it is possible that such mix-shifts are not of first-order importance to firms' pricing decisions (in which case an ECI-based measure of unit labor costs would potentially provide a better gauge of the labor costs that are relevant for price setting). In any case, both the P&C and ECI hourly compensation series are commonly followed aggregate measures that are broadly representative of the compensation costs faced by a large set of private businesses, so it is of interest to examine how each measure influences observed price inflation.

The estimation period for our VARs ranges from 1965:Q1 to 2012:Q2 for the P&C-based models, and from 1982:Q1 to 2012:Q2 for the ECI-based specifications (the choice of a later starting date for the ECI-based models reflects that fact that the ECI for total compensation does not exist prior to the 1980s). Our use of a 2012:Q2 ending date for our sample is informed by several considerations. First, income shifting in advance of an anticipated tax increase resulted in a large, transitory swing in measured P&C compensation at the end of 2012; moreover, the implementation of federal budget sequestration provisions in early 2013 had large temporary effects on some of the medical services prices that enter the PCE price index. Hence, in order to prevent these unusual endpoint observations from unduly influencing the most-recent parameter estimates from our model, we stop our estimation period in mid-2012 (note, however, that this ending date still gives us three years' worth of data from the recovery that followed the 2007–2009 recession). In addition, at the time that we constructed the dataset for our study (in early 2014), 2012 was the last full year for which the national accounts data—from which the P&C compensation measure is derived—had undergone an annual revision; thus, the compensation data from 2012 and earlier should be somewhat less subject to measurement error than the most-recent available data.

III Time variation in the passthrough of labor costs to prices

To gauge how the passthrough of trend unit labor costs into core inflation has changed over time, we use the parameter estimates from our TVP/SV models to evaluate impulse response functions for core market-based PCE inflation at different dates.[12] Figure 1 plots the median response of core inflation following a 2.7 percentage point shock to the P&C-based measure of trend unit labor cost growth (expressed at an annual rate) at various times over the period 1975–2012, along with 70 percent credible sets. (A 2.7 percentage point shock is used because this is the standard deviation of this measure of trend unit labor cost growth over the full sample period.) As can be seen from the figure, the passthrough of unit labor costs to core inflation has diminished markedly over time; in particular, in the last year of the sample the point estimate for the response's peak value is only about one-fourth as large as in 1975 (and about one-third as large as in 1985), and is statistically indistinguishable from zero.

We obtain a somewhat different picture of how the passthrough of labor costs to price inflation has evolved if we measure labor costs with the ECI. Figure 2 plots the median response of core inflation at various dates following a one-standard-deviation shock to trend unit labor cost growth from the ECI-based models.[13] In this specification, the passthrough of labor cost changes into core inflation varies little over the sample period, with a peak response that remains statistically significant throughout.

A possible explanation for these findings can be found by comparing how the volatility

[12] Alternatively, we could examine the sum of the coefficients on the lagged trend unit labor cost growth terms in the core inflation equation of the VAR. However, because core inflation is ordered below unit labor cost growth, there can be a potentially important contemporaneous effect of the unit labor cost shock; in addition, time variation in the coefficients on lagged core inflation implies that the persistence of the effect of these shocks on inflation could change over time in ways that would not be captured by only considering their (full) impact effect. Hence, in this context we believe that it is more useful to look at the impulse response functions.

[13] A one-standard-deviation shock to this measure equals 0.8 percentage point at an annual rate.

of innovations to these two measures of unit labor cost growth has varied over time. We do this in figure 3, which plots the posterior medians for the standard deviation of the structural residuals from the unit labor cost equation of the P&C- and ECI-based VARs. Starting around 1985, the volatility of own innovations to P&C trend unit labor cost growth (the solid line in the figure) has moved steadily higher, reaching a level at the end of the sample that is more than twice as large as the level that prevailed over the first half of the sample. By contrast, the volatility of innovations to ECI trend unit labor cost growth (the dashed line) changes little over the sample; if anything, a modest *downward* trend is evident in the standard deviation of these shocks. This difference in volatility is also apparent in the raw data on trend unit labor costs, which we plot in figure 4: There is a clear increase in the variability of the P&C-based measure, both relative to its earlier history and relative to the ECI-based series (note that the figure shows four-quarter log differences).[14]

Of course, such an increase in volatility should only result in a reduction in the passthrough of labor costs to price inflation to the extent that it actually reflects a rise in the degree to which (measured) compensation movements are unimportant for price setting. It is difficult, though, to pinpoint specific changes in compensation practices (or measurement) that might explain both the observed reduction in passthrough and the similarly timed rise in volatility. One possibility is that the greater use over time of employee stock options could be driving both phenomena: The ECI does not capture stock options in any form, while the employee compensation data from the national accounts

[14]Because the same estimate of trend productivity growth is used to construct both unit labor cost measures, the relative volatility across the two measures reflects the relative volatility of the P&C and ECI measures of hourly compensation growth. Likewise, because the trend productivity growth series is reasonably smooth, the volatility of the trend unit labor cost measures is itself mostly attributable to the volatility of the hourly compensation series that are used to construct them.

that are used to construct the P&C measure include the value of options when they are exercised.[15] If, as seems consistent with economic theory, the *grant* value of a stock option is a better measure of the relevant cost to the firm, then including stock option exercises in compensation could both raise measured volatility (again because options have become more prevalent over time and because the value of exercises can be subject to large quarter-to-quarter swings) while at the same time reducing the passthrough of measured compensation changes to price inflation. However, the timing of the rise in importance of stock options in compensation, which appears to occur after the mid-1990s (see Moylan, 2008, p. 7), does not line up especially well with the corresponding decline in passthrough (which, according to our estimates, appears to have occurred somewhat earlier).

It is also possible that the decline in passthrough and increase in volatility for the P&C-based measure relative to the ECI-based measure reflects increased measurement error in the former, or a rise in the importance of changes in the mix of jobs (which the ECI controls for) in driving quarterly movements in compensation growth. Regarding measurement error, we are not aware of any evidence that the quality of the P&C hourly compensation series has deteriorated over time. Indeed, the measurement of the wage and salary component of the employee compensation data that enter the P&C measure has arguably improved (at least over the past decade): Starting in 2002, the Bureau of Economic Analysis (who are responsible for constructing the U.S. national accounts) began using full-universe tax records to measure employee compensation on a *quarterly* basis (previously, these tax records were only used to provide an annual benchmark, with quarterly estimates interpolated using a proxy measure of the quarterly wage bill). Regarding the

[15]See Moylan (2008) for a discussion of how employee stock options are treated in the U.S. national accounts.

second possibility, it is certainly plausible that a compensation measure such as the ECI, which controls for changes in the mix of jobs, might provide a better read of the labor-related costs that are relevant for firms' price-setting decisions. That said, we are also not aware of any evidence that these sorts of mix-shifts have made an increasingly important contribution in recent years to the volatility of compensation growth (nor do we know of any plausible explanation as to why an increase in the importance of mix-shifts might have occurred).[16]

IV Stochastic trends in inflation and labor cost growth

Another interesting feature of the inflation process is revealed by considering the stochastic trends in price inflation and unit labor cost growth that we obtain from our model, which we plot (together with the four-quarter log differences of the actual data) in panels A and B of figure 5.[17]

As is evident from panel A of the figure, trend price inflation (the thick dashed line) rises steadily over the 1960s and 1970s, peaking at 6½ percent at the end of 1979. Trend inflation then drops sharply following the back-to-back recessions of 1980–1982, after which it stays roughly flat at around 4 percent for the rest of that decade. The 1990–1991 recession results in another—though much smaller—decline in trend inflation (to around 2 percent) that is essentially complete by the mid-1990s. From then on, however, there are no persistent movements in the trend—in particular, the 2001 recession leaves

[16] In principle, it would be possible to assess time variation in the importance of mix-shifts by comparing ECI-based hourly compensation with the hourly compensation measure from the Employer Costs for Employee Compensation (ECEC) report (very roughly, the ECEC is computed from the "raw" data used to estimate the ECI, without any correction for changes in the mix of jobs). Unfortunately, quarterly ECEC data are only available starting in 2002, which is not early enough for this purpose.

[17] We focus on the estimates from the model with the P&C-based unit labor cost measure because it allows us to consider a longer time period; the stochastic trends from the ECI-based model are similar.

no discernable imprint on the trend inflation rate, nor does the much more severe recession of 2007–2009. This behavior of inflation's stochastic trend is largely mirrored by the stochastic trend in our measure of trend unit labor cost growth, which is shown in panel B of the figure. Interestingly, the broad contour—and recent stability—of these stochastic trends is also apparent in survey measures of longer-run expected inflation, such as the expected five-to-ten-year price change from the Michigan survey (the dotted line in panel A).[18]

When inflation dynamics are characterized by a stable long-run trend, certain types of empirical inflation specifications will tend to fit poorly in periods that see persistent changes in the other determinants of inflation. For example, under this characterization of inflation dynamics, a persistent widening of the unemployment gap (such as that seen in the 2007–2009 recession and subsequent slow recovery) will tend to push actual inflation below its trend for as long as the gap persists. As the economy recovers and the gap closes, however, actual inflation will move back to its (unchanged) trend, with no persistent effect on its level. This sort of behavior will be at odds with the predictions of a traditional "accelerationist" model of inflation of the form

$$\pi_t = A(L)\pi_{t-1} + \gamma X_t + \varepsilon_t, \tag{4}$$

in which X_t captures other influences on inflation (for example, the unemployment gap or supply shocks) and where the accelerationist restriction $A(1) = 1$ is imposed when the equation is estimated. In this model, the presence of a persistent unemployment gap

[18] This correspondence between inflation's stochastic trend and survey measures of longer-term expected inflation was noted by Clark and Davig (2008); see Faust and Wright (2013) for a related discussion in the context of inflation forecasting. (Note that to obtain a longer time series for expected inflation in figure 5, prior to 1990:Q2 we splice the Michigan survey measure to the longer-term expected inflation series from the Hoey survey—see the Appendix for additional details.)

causes predicted inflation to drift lower and lower over time; when the gap finally does close, the model predicts that inflation will bottom out at some new, lower value (and will show no tendency to return to its pre-recession level).

In practice, fitting an accelerationist specification to a period where there is a large widening of the unemployment gap—and, again, where inflation dynamics are actually characterized by a stable long-run trend—will tend to attenuate the coefficient on the gap, thereby suggesting the presence of nonlinearities in the inflation–unemployment relation (this result can obtain even if the full sample period is relatively long, since a rise in the unemployment gap similar to that seen over the previous recession will represent a large and influential outlier). Relatedly, any modification to the baseline accelerationist specification that reduces the size or persistence of the measured unemployment gap—for example, defining the unemployment gap in terms of the short-term (as opposed to the total) unemployment rate or allowing for an increase in the natural rate—will tend to improve the model's performance over the past several years.[19] Of course, an alternative explanation is simply that the inflation process has changed in a manner that makes an accelerationist-style model a poor description of current actual inflation dynamics.[20]

Finally, the joint behavior of the stochastic trends for inflation and unit labor cost growth shown in figure 5 also suggests why it is that reduced-form models would tend to

[19]Because the total unemployment rate and the short-term unemployment rate behave similarly in most previous postwar U.S. business cycles—it is really only in the most-recent recovery that a noticeable difference is present—this particular modification need not compromise the accelerationist model's ability to fit inflation in earlier periods.

[20]It is noteworthy, therefore, that observers who have pointed to nonlinearities in the relationship between unemployment and inflation (such as those induced by downward nominal wage rigidity) in order to explain recent price behavior appear to have the predictions from an accelerationist model in mind as a benchmark. Likewise, many of the studies that have pointed to a differential role for short-term unemployment in determining price inflation—such as Gordon (2013) and the results presented in chapter 2 (pp. 82–83) of the 2014 *Economic Report of the President*—have made the case in the context of an accelerationist framework (this statement also applies to the analysis of Watson, 2014, inasmuch as his specification for inflation can be thought of as an accelerationist model with a long distributed lag).

find a smaller role for labor costs in driving price inflation over more recent periods: Since the early 1980s, there have been no instances of a significant wage–price spiral of the sort that resulted in a persistent and roughly contemporaneous increase in the stochastic trends of inflation and labor cost growth over the 1960s and 1970s. As a result, in recent decades movements in labor costs have not tended to carry much information about persistent movements in price inflation (and *vice-versa*).

V The role of labor costs in explaining recent inflation behavior

The results in section III suggest that the passthrough of trend unit labor costs—defined using the P&C-based measure of hourly compensation—to core inflation has declined markedly in recent years, to the point where the response of core inflation to a shock to this measure of labor costs is statistically indistinguishable from zero. Nevertheless, given that the volatility of these shocks has risen sharply over time (recall figure 3), it is still possible for this measure of unit labor costs to have a nontrivial effect on inflation. To assess the extent to which recent movements in inflation are driven by changes in unit labor cost growth, we use the estimated VAR system to decompose actual movements in core inflation into the VAR's baseline forecast (that is, the projected path of inflation absent any structural shocks but given any time-variation in the model's coefficients) and the cumulative contribution of the model's estimated structural shocks. In the figures that follow, we focus on the shocks to trend unit labor cost growth and the unemployment gap; note that for a given variable, the contribution of *all* of the model's structural shocks (*i.e.*, the shocks to core inflation, unit labor cost growth, relative import price inflation, and the unemployment gap), combined with the VAR's baseline forecast, will by construction exactly sum to the variable's actual value. The specific period over which we perform this decomposition extends from 2001:Q1 to 2012:Q2, and therefore includes both the 2001 and 2007–2009 recessions.

Panel 1 of figure 6 gives the results from this exercise for the VAR specification that uses the P&C-based measure of labor costs. (To improve readability, the figure shows actual core inflation as a four-quarter log difference, and plots the baseline forecast and innovation contributions as four-quarter moving averages.) According to the model, very little of the movement in core inflation over this period can be attributed to innovations to trend unit labor cost growth (compare the dashed line, which gives the baseline forecast, with the dotted line, which adds in the contribution of the unit labor cost shocks). Importantly, this result obtains even though these innovations account for much of the actual variation in unit labor costs themselves (see panel 3 of the figure, which repeats this calculation for trend unit labor cost growth), and even though unit labor cost growth is ordered before core inflation in the VAR. By contrast, if we instead consider the effect of shocks to the unemployment gap (which is ordered last in the VAR), the resulting decomposition suggests that the widening of the gap in both the 2001 and 2007–2009 recessions made an important contribution to pushing both price inflation (panel 2) and unit labor cost growth (panel 4) below their respective baselines.[21]

What about the VAR system that uses an ECI-based measure of trend unit labor costs? Here, the passthrough of labor costs to core inflation appeared to be essentially stable over time, which suggests that we might be able to find a more important role for labor costs in explaining recent inflation behavior if we instead use this specification. In addition, the stability of the dynamic responses of inflation to unit labor cost shocks in this model raises

[21]Note that a corresponding historical decomposition of the unemployment gap (not shown) indicates that the VAR attributes almost all of the widening of the unemployment gap in each recession to own-innovations to the gap.

the possibility that we might be able to employ a *constant-coefficient* specification, so long as the other dynamic reponses implied by the VAR manifest a similar degree of stability. As it turns out, they do: In figure 7, we plot the impulse response functions for core inflation, the ECI-based measure of trend unit labor cost growth, and the unemployment gap that we obtain from the time-varying parameter VAR at various dates (the first set of panels gives the responses following an unemployment gap shock, and the second set shows the responses following a shock to core inflation). None of the responses from this model shows any significant variation over time, mirroring the results that we obtained for the response of inflation to a unit labor cost shock (figure 2).[22]

One advantage of using a constant-coefficient VAR is that it allows us to indirectly test for the presence of downward nominal wage rigidity of the sort invoked by Daly and Hobijn (2014) in their interpretation of recent U.S. wage dynamics. Recall that in the Daly–Hobijn model, downward nominal wage rigidity induces a nonlinearity in the relationship between wage growth and labor market slack that causes wage inflation to decline by less than it otherwise would following an increase in unemployment; later, as the labor market recovers, the existence of "pent-up" wage cuts puts downward pressure on wages. Hence, if we use a constant-coefficient VAR model—in which a constant linear relationship between the unemployment gap and labor costs is imposed—to describe the evolution of labor costs in the wake of the 2007–2009 recession, the presence of downward nominal wage rigidity should result in our seeing a sequence of positive innovations to labor cost growth as the recession proceeds and the unemployment gap widens (that is, growth in labor costs should be higher than expected over the recession period). Afterwards, as the labor market starts to recover, we should expect to see a sequence of *negative*

[22]The responses of the other model variables following a shock to unit labor costs (not shown) display a similar degree of stability, as do the responses following a shock to import price inflation (also not shown).

innovations to labor costs as compensation growth is held down by pent-up wage cuts.[23] More broadly (and outside of this specific model), to the extent that downward nominal wage rigidity has had an important effect on recent *price* inflation dynamics, it should be possible to find an economically significant influence of labor costs on core inflation during and after the 2007–2009 recession.

In figure 8, we repeat our historical decompositions of core inflation and trend unit labor cost growth using a constant-coefficient VAR in which unit labor costs are defined using the ECI for hourly compensation.[24] In contrast to the results from the P&C-based VARs, we find that shocks to labor cost growth have had a less-trivial effect on core inflation over this period (panel 1), though quantitatively their contribution remains small (generally no greater than ¼ percentage point) and accounts for very little of the overall movement in core inflation from 2007 to 2012. In addition, the shocks to labor cost growth have exactly the *opposite* pattern to what we would expect if downward nominal wage rigidity were playing an important role: Over the course of the 2007–2009 recession, a sequence of negative own-innovations pushes *down* the rate of growth of labor costs (see panel 3), leaving the pace of labor cost growth lower than would be expected given the innovations to the unemployment gap alone (panel 4). The labor market starts to recover after 2009 (the unemployment gap reaches its widest point in 2009:Q4); however, instead of seeing negative innovations to labor costs (as would be expected were pent-up wage cuts holding down compensation growth), over the next couple of years a sequence of *positive* own-innovations puts net *upward* pressure on labor costs.

[23]This discussion is loose in that it does not distinguish between compensation growth and growth in trend unit labor costs (recall that the latter is the former adjusted for trend productivity). However, over the 2007–2012 period that is our main focus, virtually all of the movement in our measure of trend unit labor cost growth reflects changes in nominal compensation growth.

[24]The specification and estimation period for the VAR are otherwise identical to the corresponding time-varying parameter VAR described in section II.

Whatever conclusion one draws about the presence or absence of downward nominal wage rigidity from this pattern of labor-cost innovations, the fact remains that on average these innovations make a small *negative* contribution to core price inflation over this period, not the large *positive* contribution that would be needed were wage-related developments to be a good candidate explanation for the existence of significant "missing disinflation."

VI Additional implications and caveats

In this paper, we have documented that shocks to labor costs have made a relatively small contribution to the observed behavior of price inflation in recent years. Our findings therefore cast doubt on explanations of recent inflation behavior that appeal to such mechanisms as downward nominal wage rigidity or a differential contribution of long-term and short-term unemployed workers to wage pressures. We have also proposed an alternative way to understand the recent behavior of price inflation that does not rely on wage-based explanations—specifically, price inflation is currently tied down by a stable stochastic trend, to which it ultimately returns once resource utilization rates return to normal levels and the influences of any other shocks dissipate. If correct, this alternative view of the inflation process implies that most of the "missing disinflation" puzzle that has been discussed by previous analysts simply reflects the use of a model of inflation (an accelerationist specification) that no longer provides an especially accurate characterization of U.S. inflation dynamics, and that therefore generates a misleading benchmark for how we would have expected inflation to behave following the 2007 business cycle peak. In addition, our results suggest that wage developments are unlikely to be an important independent driver of (or an especially good guide to) future price developments.

We would emphasize that our results do not *necessarily* imply that labor costs are unimportant for pricing. Instead, a more-nuanced interpretation is that as long as the stochastic trends for inflation and labor costs remain stable—in particular, so long as the sort of wage–price spiral that characterized earlier decades does not emerge—observed year-to-year movements in price inflation are likely to mostly reflect a mix of changes in resource utilization, supply shocks, and idiosyncratic variation, not independent movements in the growth of labor costs. Indeed, it is quite possible that the greater observed stability of inflation's stochastic trend is itself directly attributable to the greater stability that we observe in the stochastic trend for labor cost growth (even if the cause of this latter phenomenon ultimately lies elsewhere).

This last point highlights an important question that is left unanswered here—namely, what has caused the processes for inflation and labor costs to change in such as way as to make their long-run levels essentially constant? While a simple explanation—that the public's expectations of longer-term inflation have become better anchored over time—certainly seems plausible (recall our figure 5), such an answer itself begs the further question of *why* this greater anchoring has taken place. (The obvious reply, that the improved conduct or credibility of monetary policy has played a key role in anchoring expectations, turns out to have surprisingly little hard evidence to support it—though see Clark and Davig, 2011, for some suggestive circumstantial evidence.) Unfortunately, until we come to a much better understanding of what determines the expectations of wage- and price-setters, we are unlikely to be able to claim much certainty regarding how the inflation process will evolve in the future.

VII References

Ball, Laurence, and Sandeep Mazumder, "Inflation Dynamics and the Great Recession," International Monetary Fund Working Paper no. WP/11/121 (2011).

Blinder, Alan S., and Jeremy B. Rudd, "The Supply-Shock Explanation of the Great Stagflation Revisited," in Michael D. Bordo and Athanasios Orphanides (Eds.), *The Great Inflation: The Rebirth of Modern Central Banking* (Chicago: University of Chicago Press, 2013).

Clark, Todd E., and Troy Davig, "An Empirical Assessment of the Relationships Among Inflation and Short- and Long-Term Expectations," Federal Reserve Bank of Kansas City Working Paper no. RWP 08-05 (2008).

Clark, Todd E., and Troy Davig, "Decomposing the Declining Volatility of Long-Term Inflation Expectations," *Journal of Economic Dynamics and Control* 35 (2011), 981–999.

Clark, Todd E., and Stephen J. Terry, "Time Variation in the Inflation Passthrough of Energy Prices," *Journal of Money, Credit and Banking* 42 (2010), 1419–1433.

Cogley, Timothy, and Thomas J. Sargent, "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US," *Review of Economic Dynamics*, 8 (2005), 262–302.

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

Daly, Mary C., and Bart Hobijn, "Downward Nominal Wage Rigidities Bend the Phillips Curve," *Journal of Money, Credit and Banking* 46:S2 (2014), 51–93.

Emery, Kenneth M., and Chih-Ping Chang, "Do Wages Help Predict Inflation?" Federal Reserve Bank of Dallas *Economic Review*, First Quarter (1996), 2–9.

Faust, Jon, and Jonathan H. Wright, "Forecasting Inflation," in Graham Elliott and Allan Timmermann (Eds.), *Handbook of Economic Forecasting*, vol. 2A (Amsterdam: Elsevier–North Holland, 2013).

Gordon, Robert J., "The Phillips Curve is Alive and Well: Inflation and the NAIRU during the Slow Recovery," NBER Working Paper no. 19390 (2013).

Koop, Gary, and Simon M. Potter, "Time Varying VARs with Inequality Restrictions," *Journal of Economic Dynamics and Control* 35 (2011), 1126–1138.

Mehra, Yash, “Wage-Price Dynamics: Are They Consistent with Cost Push?” Federal Reserve Bank of Richmond *Economic Quarterly* 86:3 (2000), 27–43.

Moylan, Carol E., “Employee Stock Options and the National Accounts,” *Survey of Current Business* 88:2 (2008), 7–13.

Staiger, Douglas, James H. Stock, and Mark W. Watson, “Prices, Wages, and the U.S. NAIRU in the 1990s,” in Alan Krueger and Robert Solow (Eds.), *The Roaring Nineties: Can Full Employment Be Sustained?* (New York: Russell Sage Foundation and Century Foundation Press, 2001).

Watson, Mark W., “Inflation Persistence, the NAIRU, and the Great Recession,” *American Economic Review* 104 (2014), 31–36.

A Appendix

This Appendix provides additional details regarding the data and estimation procedures that we use for our study.

A Data documentation

All standard data from the National Income and Product Accounts (NIPAs) were downloaded from the Bureau of Economic Analysis (BEA) website; data on unemployment, productivity, and compensation were downloaded from the Bureau of Labor Statistics website. (All data were current as of February 12, 2014.) Finally, the CBO short-term natural rate series is taken from the February 2014 edition of *The Budget and Economic Outlook: 2014 to 2024*.

Market-based PCE price index: Official data for the core market-based PCE price index are published from 1987 to the present. To extend back the market-based series before 1987, we use detailed PCE data and a Fisher aggregation procedure routine that replicates the procedure followed by the BEA in constructing the NIPAs to strip out the prices of core nonmarket PCE components from the published overall core PCE price index, where our definition of "nonmarket" mimics the BEA's.[25]

Relative import price term: We define import price inflation as the annualized log difference of the price index for imports of nonpetroleum goods excluding natural gas, computers, peripherals, and parts, which we compute using detailed NIPA series. (As the data required to construct this series only extend back to 1967:Q1, we use the annualized log difference of total goods imports prior to that date.) The relative import price inflation term that we use in our VARs is equal to the difference between this series and core market-based price inflation (lagged one period), weighted by the two-quarter moving average of the share of nominal imports (defined consistently with the import price measure) in nominal core PCE.[26]

Long-run expected inflation: We splice the median response to the Michigan survey's question on expected 5-to-10-year inflation to long-run expected CPI inflation from the Hoey survey. Specifically, we use the Hoey data from 1980:Q3 to 1989:Q4 (their first and last available dates), and the Michigan survey data starting in 1990:Q2 (the first date in which a continuous quarterly series is available). (For the 1990:Q1 observation, we extrapolate the Hoey data using the change from 1989:Q4 to 1990:Q1 in the median long-term CPI inflation forecast from the Survey of Professional Forecasters.)

[25] As noted in the text, the core inflation series that we use in our VAR models subtracts out Blinder and Rudd's (2013) estimates of the effects of the Nixon-era price controls.

[26] In constructing the relative import price term, we use the actual core market-based PCE price series (that is, we do not adjust the series for the effect of price controls), and scale the nominal import share by its sample mean.

Trend productivity growth: Trend productivity growth is defined as the low-frequency component of the annualized log difference of nonfarm business output per hour, which we obtain from a band-pass filter with the filter width and cutoffs set equal to the values used by Staiger, Stock, and Watson (2001). We use an ARIMA(4,1,0) model to pad the actual productivity growth series prior to its 1947:Q2 starting point; to pad the series after its 2013:Q3 endpoint, we set the series equal to the CBO's February 2014 forecast of average trend labor productivity growth from 2013 to 2024 (which equals 1.96 in log differences), and to the 2024 value of the CBO forecast (which equals 1.76) thereafter. (Note that the padded series is only used in the trend extraction routine, not to construct any of the unit labor cost series that we use in our VAR models.)

B Additional estimation details

We use Clark and Terry's (2010) implementation of the Metropolis-within-Gibbs posterior sampler, which in turn follows Cogley and Sargent (2005).[27] We set the number of burn-in draws equal to 50,000 and then run 50,000 additional draws, keeping every tenth draw. The priors for the initial values are computed by estimating the VAR over a training sample that runs from 1950:Q2 to 1964:Q4 (for the P&C-based models) or from 1967:Q2 to 1981:Q4 (for the ECI-based models).[28] Following Clark and Terry (2010), we use an uninformative prior for the degree of time variation in the VAR coefficients (specifically, we set the prior equal to 0.001 times the variance-covariance matrix of the VAR coefficients estimated over the training sample, with degrees of freedom set equal to the number of coefficients in the system plus one).

[27]In contrast to Cogley and Sargent (2005), we do not set the sampler to truncate explosive draws with a reflecting barrier or "backstep" algorithm; in line with the recommendation of Koop and Potter (2011), therefore, we generally report median values and use relatively interior percentiles (the 15th and 85th) to bound the credible sets. (An exception is in the historical decompositions, where we use mean values in order to ensure that the sum of the baseline forecast and the contributions of all shocks will exactly equal the actual value of the variable whose decomposition we are describing.)

[28]In two cases, it was necessary to extend the data back to 1950:Q2 for use in the training sample; specifically, we extended the market-based core PCE inflation series (which starts in 1959:Q2) with the log difference of a price index for total PCE less prices for food and energy goods, and extended our measure of trend productivity growth prior to 1955:Q1 by setting it equal to its 1955:Q1 value of 2.375. In addition, prior to 1980:Q2 (the starting date for our ECI inflation series), we used P&C hourly compensation growth to compute trend unit labor cost growth in the training sample that we constructed for the ECI-based VARs. (Again, all of these extended series were used for the training sample only.)

Figure 1
Response of core inflation to P&C trend unit labor cost growth shock

A. 1975

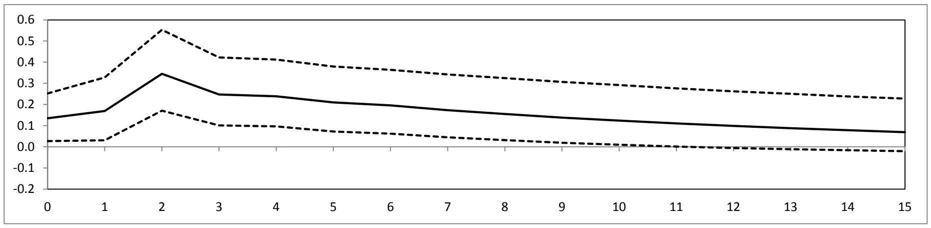

B. 1985

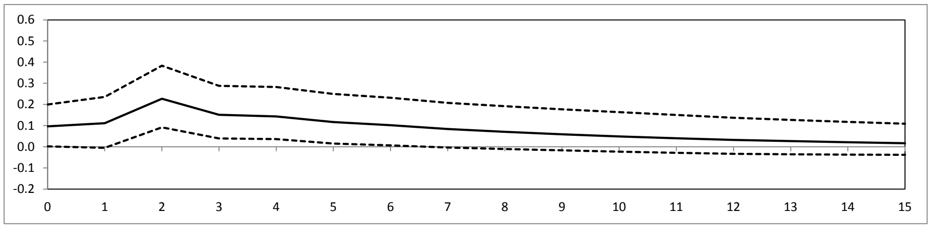

C. 1995

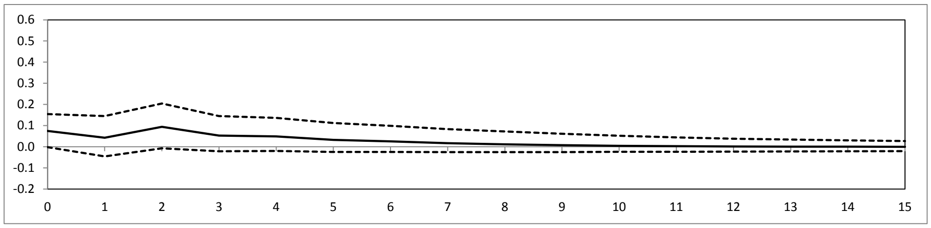

D. 2005

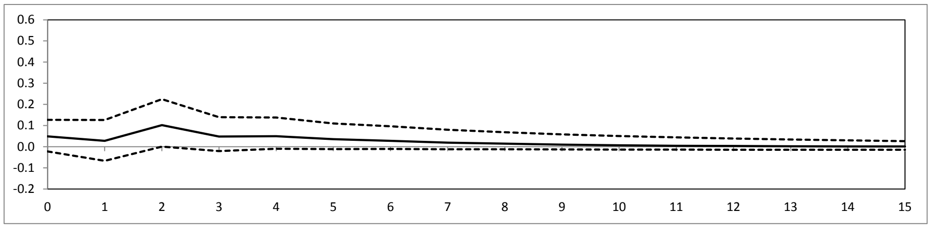

E. 2012

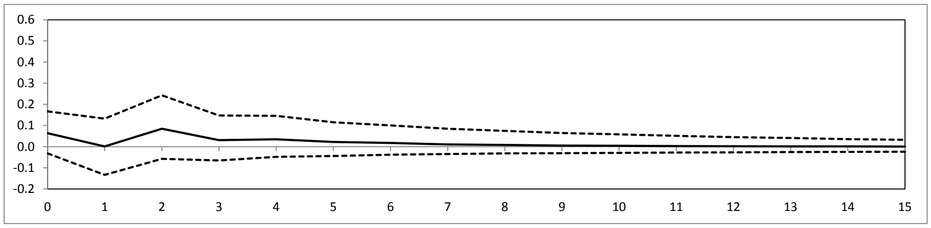

Note: Core inflation defined as log difference of core market-based PCE price index. Dashed lines denote 70 percent credible set.

Figure 2
Response of core inflation to ECI trend unit labor cost growth shock

A. 1985

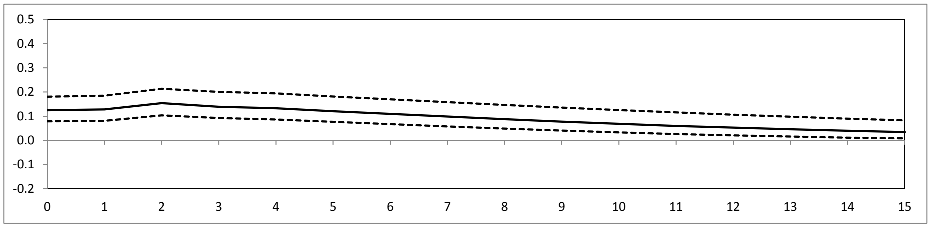

B. 1995

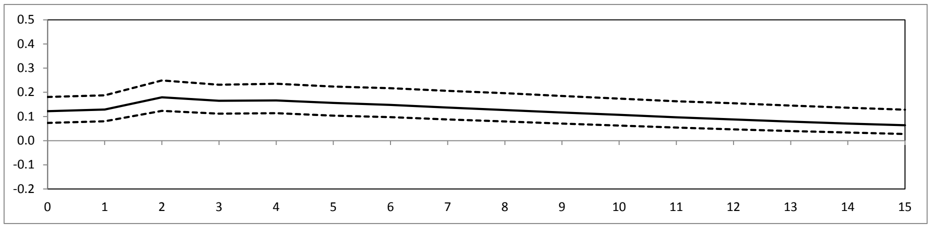

C. 2005

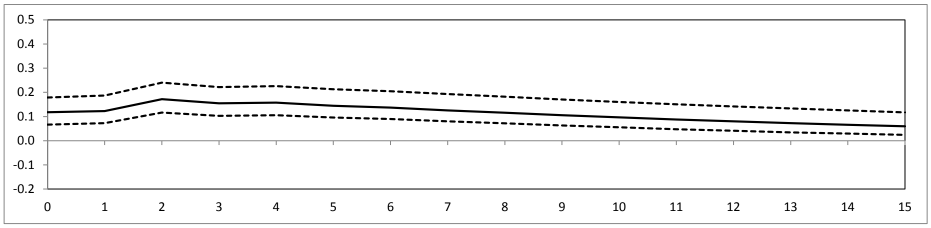

D. 2012

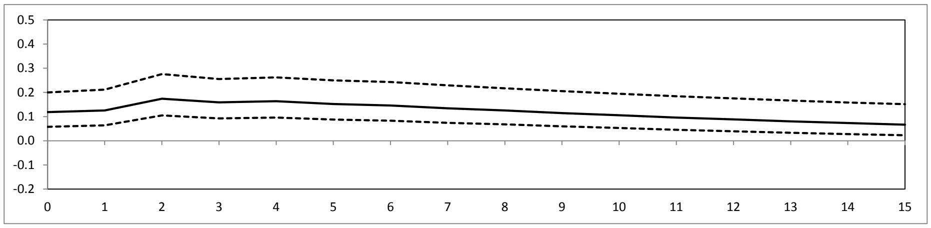

Note: Core inflation defined as log difference of core market-based PCE price index. Dashed lines denote 70 percent credible set.

Figure 3
Standard deviation of structural innovations
(Annualized log differences)

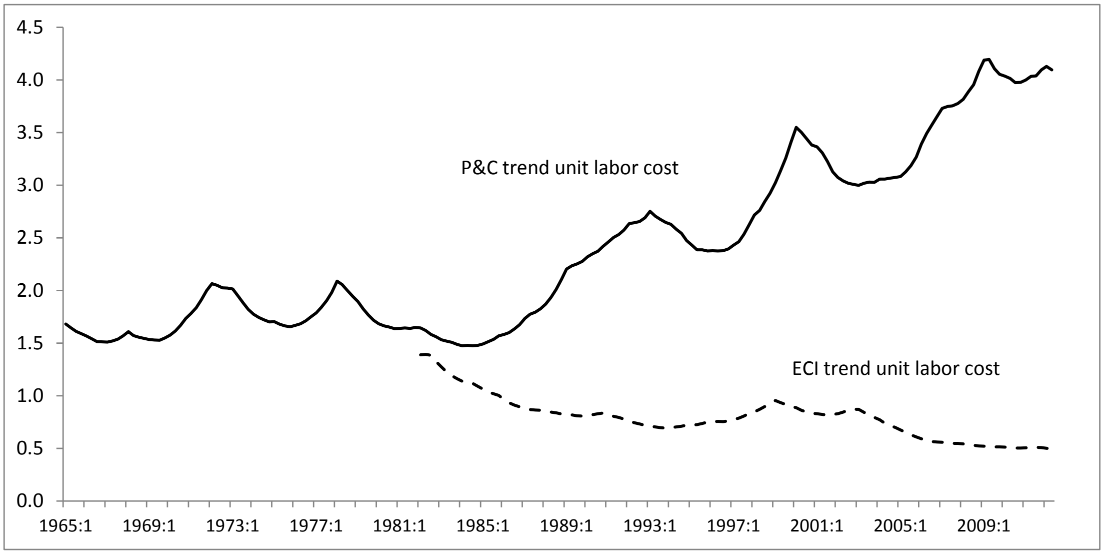

Note: Estimated from VAR systems with relative import price inflation, labor cost growth, core inflation, and unemployment gap.

Figure 4
Trend unit labor cost measures
(Four-quarter log differences)

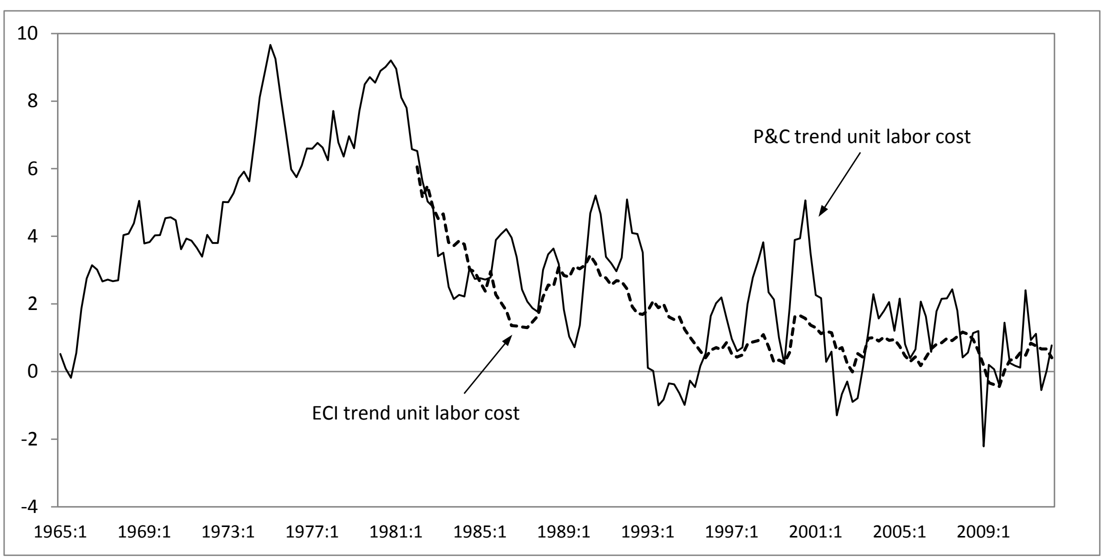

Figure 5
Measures of trend inflation

A. Core market-based PCE price index

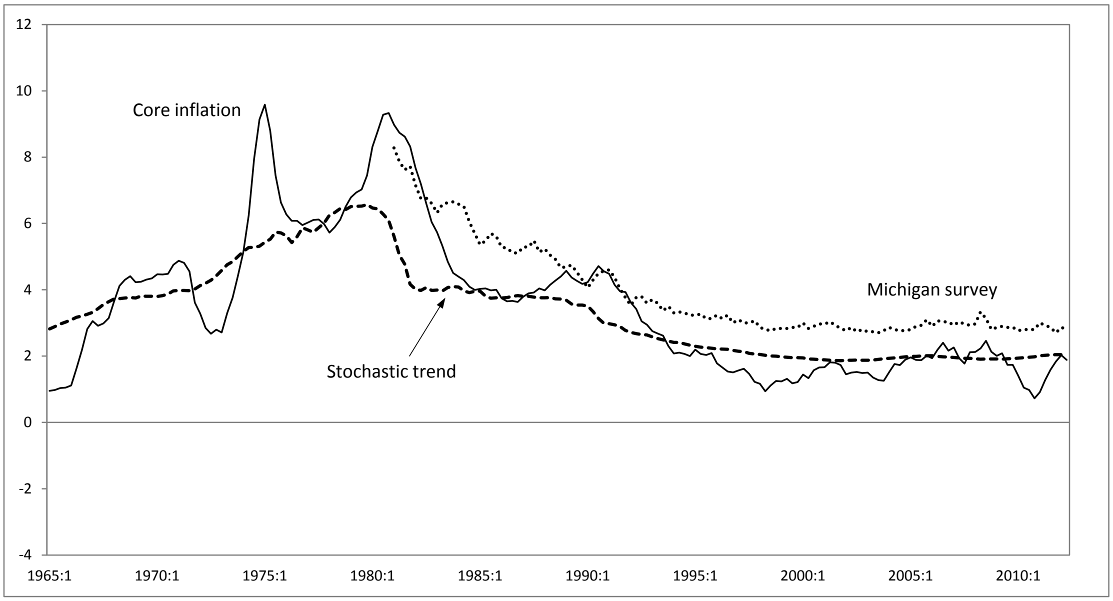

Note: Inflation measured as four-quarter log difference. Michigan survey is median response, spliced to Hoey data prior to 1990:Q2.

B. Trend unit labor costs (productivity and costs measure)

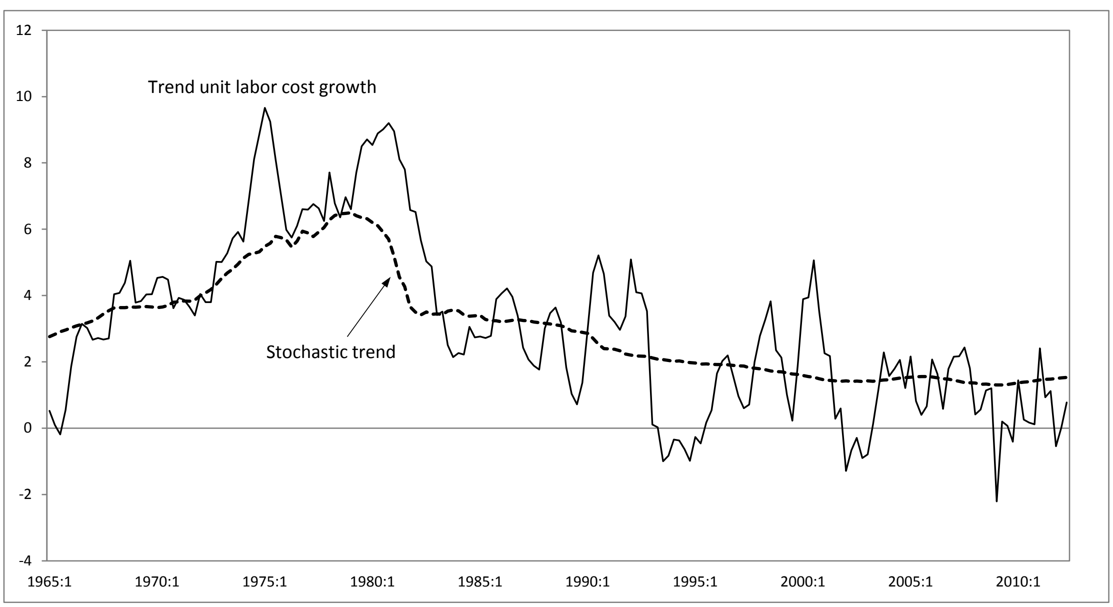

Note: Trend unit labor cost growth measured as four-quarter log difference.

Figure 6
Effect of structural shocks from P&C trend unit labor cost model
(Four-quarter log differences)

A. Effect of structural shocks on core inflation

1. P&C trend unit labor cost shocks

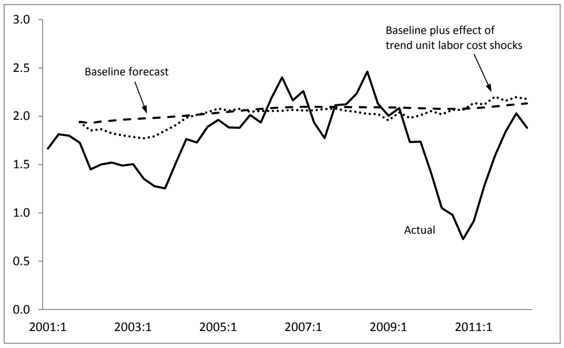

2. Unemployment gap shocks

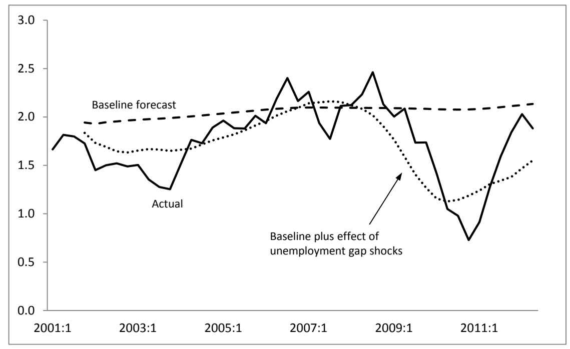

B. Effect of structural shocks on P&C trend unit labor cost growth

3. P&C trend unit labor cost shocks

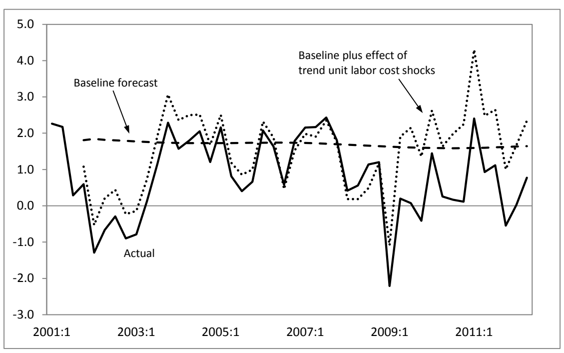

4. Unemployment gap shocks

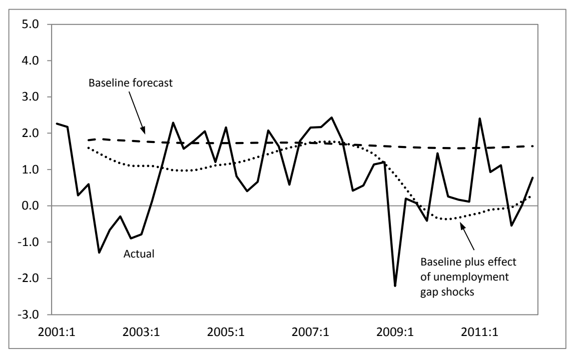

Note: Core inflation defined as four-quarter log difference of core market-based PCE price index.

Figure 7
Impulse responses at different points in time

1. Effect of an unemployment gap shock

Inflation response

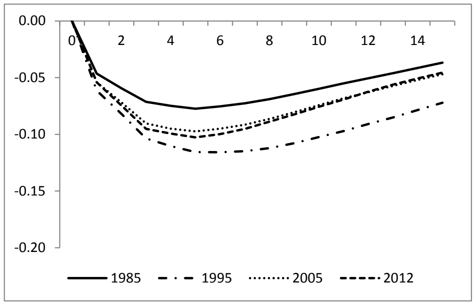

ECI trend unit labor cost growth response

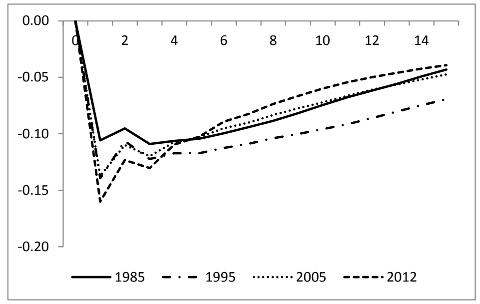

Unemployment gap response

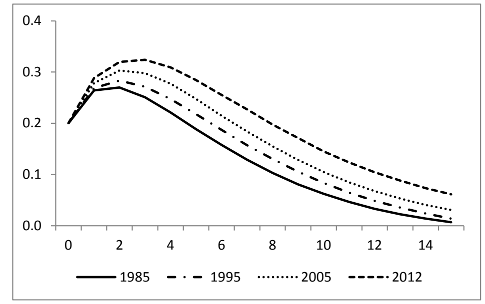

2. Effect of an inflation shock

Inflation response

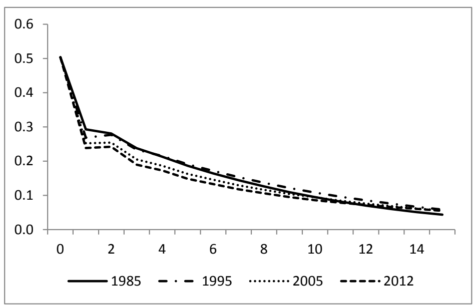

ECI trend unit labor cost growth response

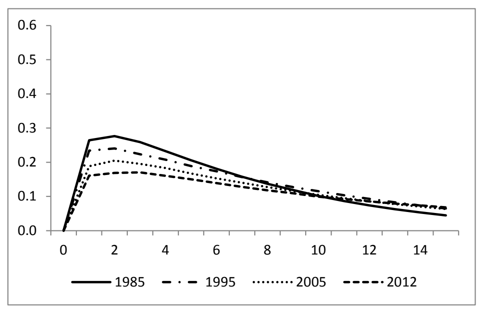

Unemployment gap response

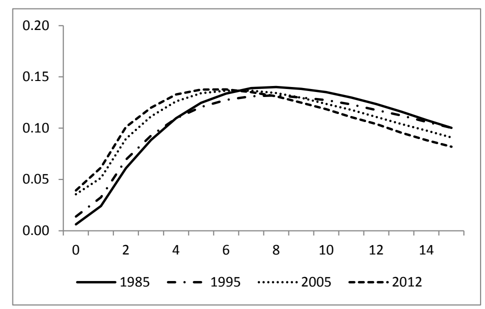

Note: VAR system includes relative import price inflation, ECI trend unit labor cost growth, core market-based PCE price inflation, and unemployment gap.

Figure 8
Effect of structural shocks from ECI trend unit labor cost model
(Four-quarter log differences)

A. Effect of structural shocks on core inflation

1. ECI trend unit labor cost shocks

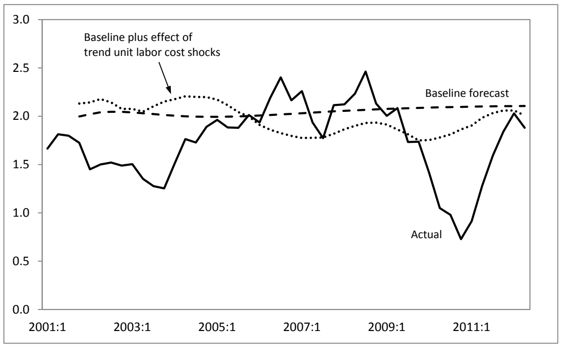

2. Unemployment gap shocks

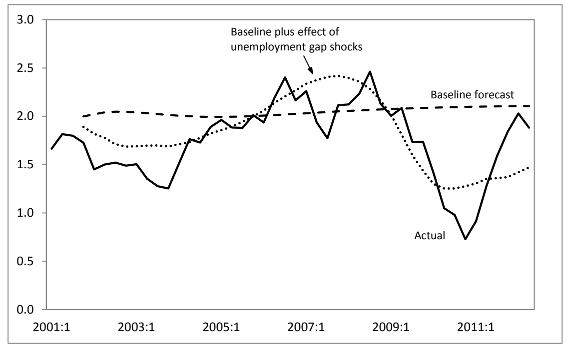

B. Effect of structural shocks on ECI trend unit labor cost growth

3. ECI trend unit labor cost shocks

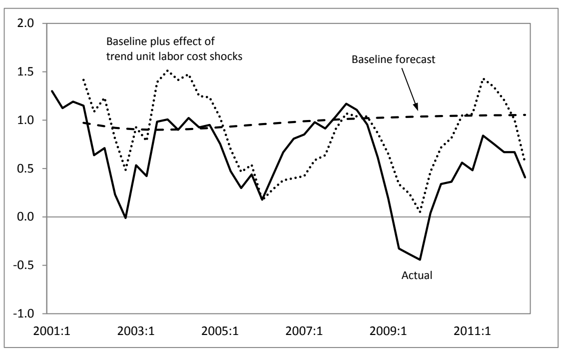

4. Unemployment gap shocks

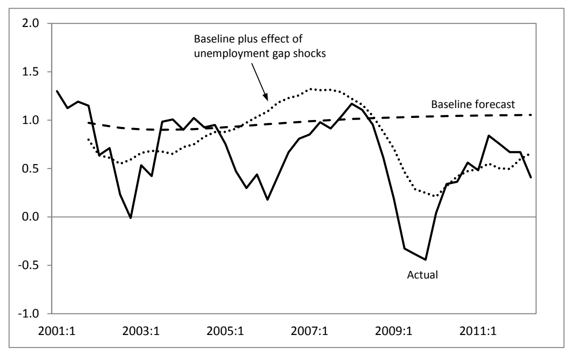

Note: Core inflation defined as four-quarter log difference of core market-based PCE price index.