Economic Quarterly—Volume 97, Number 1—First Quarter 2011—Pages 45–66

Housing and the Great Recession: A VAR Accounting Exercise

Samuel E. Henly and Alexander L. Wolman

Measured in terms of either output or employment, the United States economy has gone through a period of extreme weakness since early 2008. Real gross domestic product (GDP) fell by 3.83 percent from the second quarter of 2008 to the second quarter of 2009. If it had instead grown at its average post-1982 rate, GDP would have risen by 3.2 percent over the same interval. Nonfarm payroll employment has fallen by 6.1 percent from December 2007–February 2010; growth at its average post-1982 rate would have meant an increase of 3.6 percent over the same interval.

Comparisons to some past episodes, presented in Figure 1, are useful for evaluating the severity of the current one. The decline in real GDP in 2008 and 2009 is larger than any that the United States has experienced since the immediate post-World War II period.[1] The next largest postwar decline was 3.73 percent, from 1957:Q3–1958:Q1. Even in the early 1980s, GDP never fell below its previous peak by more than 2.87 percent. Employment behavior also looks extreme when compared to other episodes (see Figure 1, Panel B). Since the end of World War II, the previous greatest peak-to-trough decline in employment was 4.37 percent from April 1957–June 1958, compared to the 6.1 percent decline in the Great Recession. And in the early 1980s the peak-to-trough decline in employment was never greater than 3.1 percent.

The views expressed in this article are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. The authors acknowledge helpful comments from Brian Gaines, Thomas Lubik, Yash Mehra, Pierre Sarte, and John Weinberg. E-mail: alexander.wolman@rich.frb.org.

[1] We are defining a decline in real GDP as the difference between the level at its peak and the level at its subsequent trough. If there are multiple local troughs before real GDP again reaches its initial peak, we choose the lowest trough to measure the decline.

Figure 1 Large Postwar Output and Employment Declines

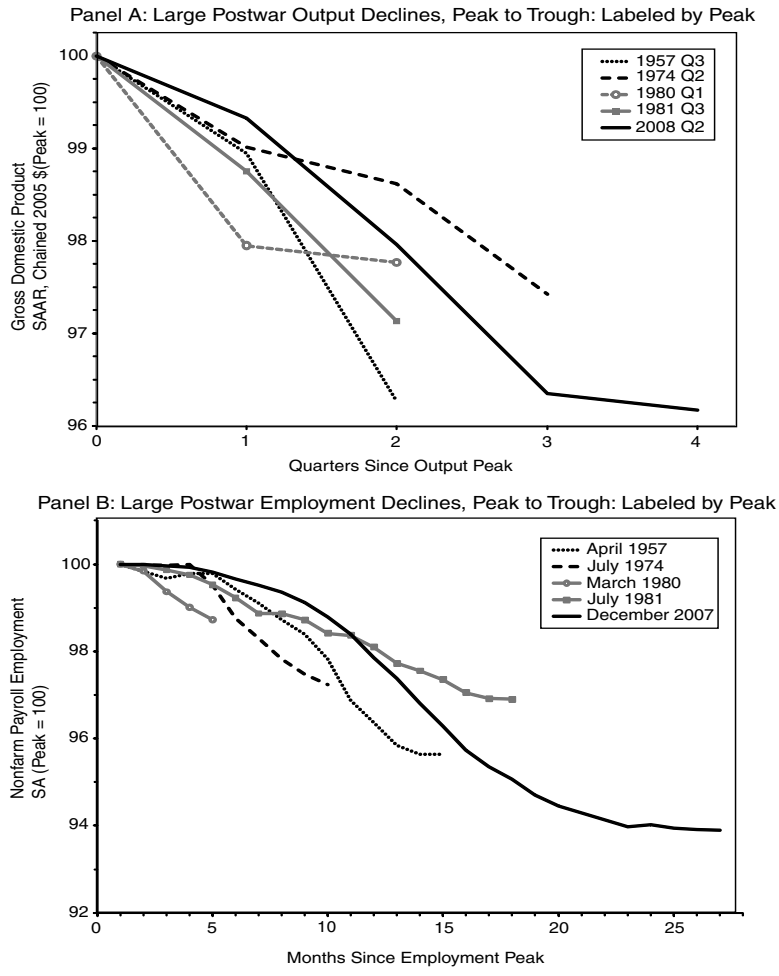

There is thus no doubt about the severity of the decline in economic activity since late 2007. Understanding why the decline occurred is more difficult. We aim to take some initial steps toward such an understanding by studying the behavior of the major components of output and, in less detail, employment. We will focus especially on the behavior of the housing components of output and employment. Much popular commentary on the recession has emphasized the role of the housing bust and the subsequent (and related) financial crisis. After documenting that the decline in the housing component of GDP was of roughly the same magnitude as the decline in GDP itself, we will use vector autoregressions (VARs) to investigate the extent to which housing

Figure 2 Real Residential Fixed Investment and Real GDP

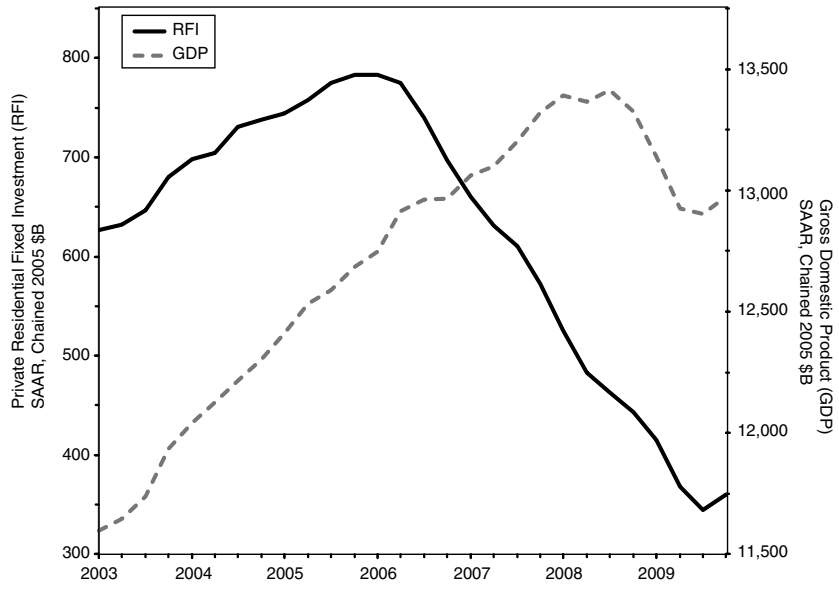

can be identified as a cause of the decline in aggregate activity. While we do not propose any economic model to explain the recession, our statistical analysis may be suggestive for subsequent modeling that does attempt to assign causality.

1. FACTS ABOUT THE RECESSION

Figure 2 displays the annualized quarterly levels of real GDP (right scale) and real residential investment (left scale) from 2002 to the present. Residential investment peaked in the fourth quarter of 2005, and fell $439 billion in real terms through the second quarter of 2009, a decline of more than 56 percent. As reported above, real GDP peaked in the second quarter of 2008 and bottomed out in the second quarter of 2009. The 3.8 percent fall in real GDP represents about $514 billion in chained 2005 dollars. Thus, the cumulative decrease in residential investment so far has been approximately the same magnitude as the decrease in real GDP that ended in the second quarter of this year. Just as striking as the similarity in magnitudes is the difference in timing of the declines in residential investment and GDP: Residential investment fell steadily for two years before GDP began to fall. Evidently other components were supporting GDP growth during 2006 and 2007 before weakening in 2008. Figures 3 and 4 show that consumption and nonresidential investment fit this

Figure 3 Real Nonresidential Fixed Investment and Real GDP

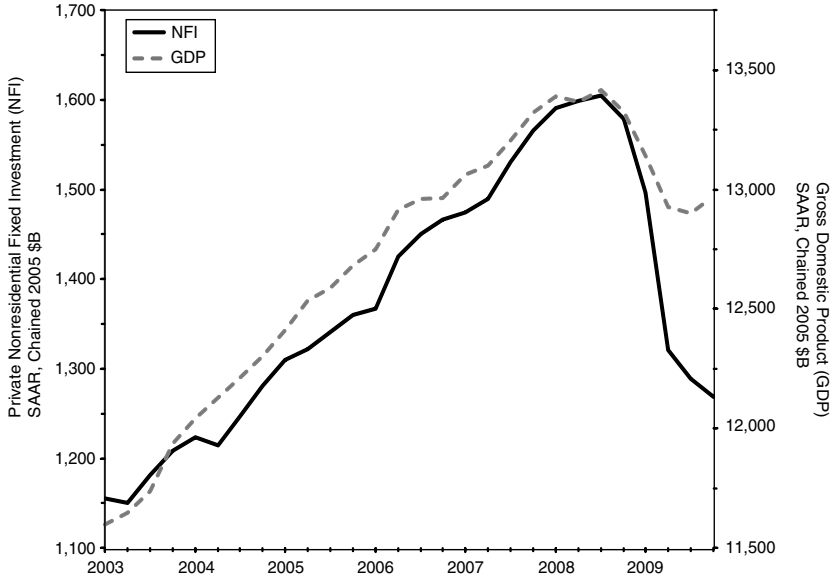

description; decreases in those two components—which together account for 80 percent of GDP—were roughly contemporaneous with the fall in real GDP.

Like output, aggregate employment lagged housing-related employment in the current recession. Figure 5 displays total nonfarm payroll employment, along with employment in residential building construction. The latter peaked in April 2006 at 1.02 million, and has since fallen by 43 percent, to 586,000 in November 2009. Aggregate employment began to fall 20 months after residential construction employment, in December 2007, and thus far has fallen by 6.1 percent, from 138 million to 130 million. Unlike output, the fall in housing-related employment has been much smaller in absolute terms than the fall in overall employment.[2] Even if we add residential specialty trades to residential building construction, the decline in housing-construction related employment amounts to only about one-sixth of the total job losses since December 2007.

The data portrayed in Figures 2–5 raise many questions. The remainder of this article will focus on just one of them: To what extent did the fall in aggregate output represent the economy's typical lagged response to a shock

[2] In percentage terms, whether one looks at employment or output, the decline in housing construction dwarfs the decline in the aggregate economy.

Figure 4 Real Personal Consumption Expenditures and Real GDP

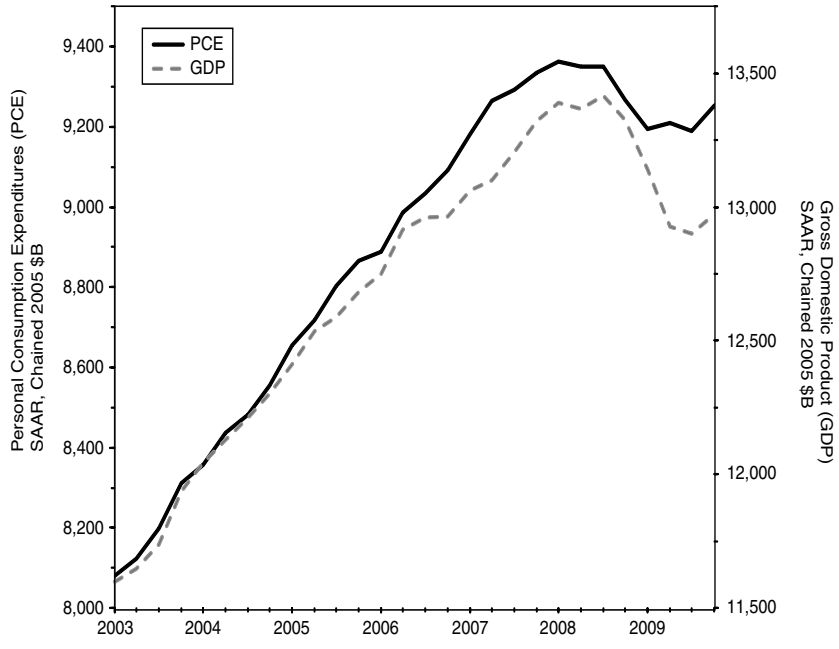

to the housing sector? Although overall employment fell much more than employment in home construction, we can and will ask the same question about employment. That is: Did the large fall in overall employment reflect the normal propagation of a shock to housing? We will determine this normal statistical pattern using vector autoregressions for the components of output and employment.

2. BACKGROUND ON VAR METHODOLOGY

A VAR is a statistical model of the behavior across time of a set of variables. A VAR specifies that the value of each variable in a given time period is a linear function of (1) the lagged values of all the variables and (2) one or more exogenous random variables. For our purposes, a VAR for the components of real GDP (or employment) is useful because it provides a summary of the relationship between those components over the sample with which the VAR is estimated. With a VAR, we will be able to assess whether the recent behavior of real GDP and employment can be interpreted as reflecting the normal response to a large shock to the housing sector.

Figure 5 Residential Construction Employment and Total Nonfarm Employment

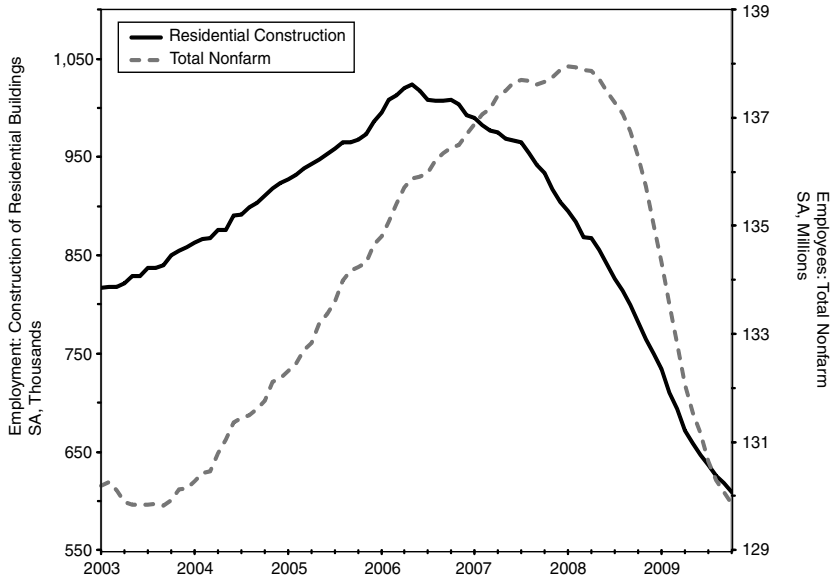

A VAR for N variables observed over T periods, $Y_t = \left\{y_{1,t}, y_{2,t}, ..., y_{N,t}\right\}'$, $t = 1, 2, ..., T$, would be written as

$$\underset{Nx1}{Y_t} = \underset{Nx1}{c} + \sum_{k=1}^{K} \underset{NxN}{\Phi_k}\, \underset{Nx1}{Y_{t-k}} + \underset{Nx1}{u_t}\,. \tag{1}$$

In (1), c is a vector of constant terms, Φ_k are coefficient matrices for each of K lags, and u_t is an error vector. The coefficient matrices, Φ_k, can be estimated using ordinary least squares.[3]

We are interested in isolating the aggregate effect of a shock to the housing component of GDP or employment. Consider the case of GDP and allow residential investment to be the first element of Y_t; an obvious approach is to treat the first element of u_t as the shock to residential investment, and use the estimated Φ_k matrices to determine the effect of this shock on all the components of GDP and thus on GDP itself. Unfortunately, the elements of u_t are generally correlated within the period. That is, a high value for the first element of u_t provides information about other elements of u_t, and that

[3] There are many good introductory treatments of VAR analysis; two examples are Sims (1986) and Hamilton (1994).

information needs to be taken into account in determining the effect of a shock to residential investment.

The process of accounting for correlation of elements of u_t within a period is called *orthogonalization*. Let us denote the covariance matrix of u_t by Σ,

$$Euu' = \Sigma. \tag{2}$$

To orthogonalize the N errors u_t, we decompose them into N errors v_t, which are a linear combination H of u_t,

$$v_t = Hu_t. \tag{3}$$

For the matrix H to orthogonalize the errors, it must be the case that the elements of the resulting vector v_t are uncorrelated within the period:

$$E\left(vv'\right) = I. \tag{4}$$

From (2)–(4), it follows that the linear transformation H must be related to the covariance matrix Σ as follows:

$$H^{-1}\left(H^{-1}\right)' = \Sigma. \tag{5}$$

For ease of notation, define $G \equiv H^{-1}$. Then we have

$$GG' = \Sigma. \tag{6}$$

Thus, if we can find a matrix G that satisfies (6), we can meaningfully trace out the effects of the N different shocks v_t on each of the variables Y_t. Unfortunately, there are generally many matrices G that satisfy (6). We will restrict G to be lower triangular, meaning that it has only zeros above the diagonal. There is a unique lower triangular G satisfying (6), and it is known as the Cholesky matrix.

With the Cholesky approach to orthogonalization, it is natural to think about the *order* of variables in the VAR. To make this point clear, we use a two-variable example:

$$\begin{bmatrix} Y_{1,t} \\ Y_{2,t} \end{bmatrix} = \begin{bmatrix} c_1 \\ c_2 \end{bmatrix} + \sum_{k=1}^{K} \begin{bmatrix} \Phi_{11} & \Phi_{12} \\ \Phi_{21} & \Phi_{22} \end{bmatrix} \begin{bmatrix} Y_{1,t-k} \\ Y_{2,t-k} \end{bmatrix} + \begin{bmatrix} G_{11} & 0 \\ G_{21} & G_{22} \end{bmatrix} \begin{bmatrix} v_{1,t} \\ v_{2,t} \end{bmatrix}.$$

Note that we have replaced u_t in the VAR of (1) with Gv_t. We say that $Y_{1,t}$ is ordered first in the VAR because only the first element of v_t affects $Y_{1,t}$ within the period. With the Cholesky decomposition there is some justification for referring to the random variable $v_{1,t}$ as a shock to $Y_{1,t}$. In what follows, we will order the housing component of GDP or employment first in our VARs and investigate the role of "housing shocks." By ordering residential investment first, we unambiguously identify housing shocks—there is only one shock that affects residential investment independently of the VAR's dynamics (the Φ matrices). Of course, one may wonder whether our results are robust to different orderings—perhaps there is more than one shock that affects residential investment contemporaneously. This concern will be addressed.

Table 1 VAR Output Components

Category	Share in 2002
Personal Consumption Expenditures (C)	69.4 percent
Government Expenditures (G)	19.7 percent
Nonresidential Fixed Investment (NFI)	10.2 percent
Residential Fixed Investment (RFI)	5.3 percent
Change in Inventories (dI)	0.1 percent
Net Exports (X)	−4.8 percent

3. VAR FOR COMPONENTS OF GDP

Table 1 describes the breakdown of GDP into six components for which we estimate the VAR, along with the 2002 GDP shares of those components.[4] The VAR will be specified in log levels for most of the variables.[5] Net exports and inventory investment need to be treated differently because they can be negative; we include those variables in the VAR as shares of GDP. The VAR contains five lags and is estimated over the period 1985:Q1–2009:Q4.[6]

As described above, the results that we emphasize have residential investment ordered first in the VAR. The only shocks we study are shocks to residential investment, so the ordering of the other variables is inconsequential. Because our primary interest is the relationship between the decline in housing-related economic activity and the decline in overall activity, most of our analysis will be based on economic conditions as of the end of 2005—the peak of housing activity. We will look at the VAR's forecasts conditional on data through the end of 2005, and then ask how much shocks to residential investment can account for the deviation of outcomes from the VAR's forecast. The variables in the VAR represent components of GDP, but we can generate the VAR's forecasts for GDP itself by appropriately transforming the component forecasts.[7]

[4] We use 2002 to describe the shares because it occurred neither during the height of the housing boom nor during the depths of the bust.

[5] See Sims, Stock, and Watson (1990) on levels estimation of VARs with nonstationary variables.

[6] We use five lags because we want to capture any residual seasonality while conserving on parameters, given the relatively short sample.

[7] If r, n, c, g, x, i represent, respectively, VAR forecasts for the logs of residential investment, nonresidential investment, consumption, and government spending, and the GDP shares of net exports and inventory investment, we generate the forecast for GDP (Y) as follows:

$$Y = \frac{\exp^r + \exp^n + \exp^c + \exp^g}{1 - x - i}.$$

Figure 6 Impulse Response Functions for Baseline VAR

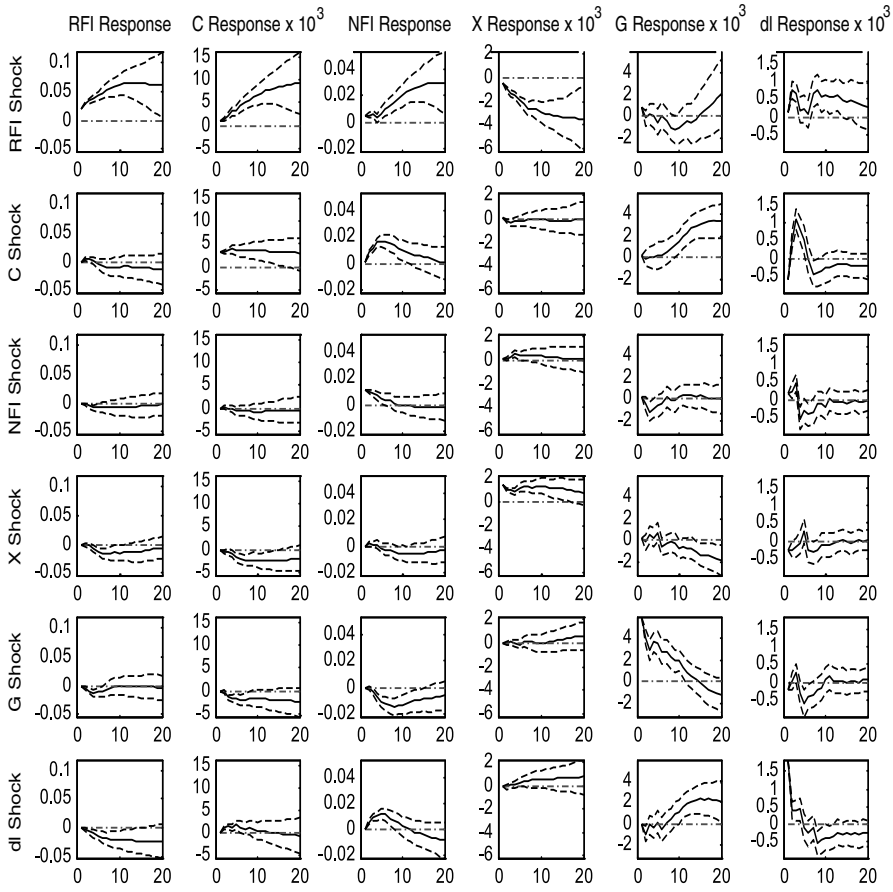

Baseline Results

Figure 6 displays a full set of impulse response functions for our baseline specification; the variables are ordered as they appear in the figure, with residential investment first. Note the relatively large and persistent effects that shocks to residential investment have on RFI, C, and NFI, which together account for about 85 percent of GDP. Turning to the Great Recession, we illustrate some of our main results in Figure 7. Focus first on the two solid lines, which represent the actual path of GDP (black) and the path forecasted by the VAR conditional on data through the end of 2005 (gray). If we define "trend" as the level of GDP forecasted by the VAR, then GDP was at or above trend until the third quarter of 2008. In the fourth quarter of 2008 and the first quarter of

Figure 7 U.S. GDP: Sample Period 1985–2009

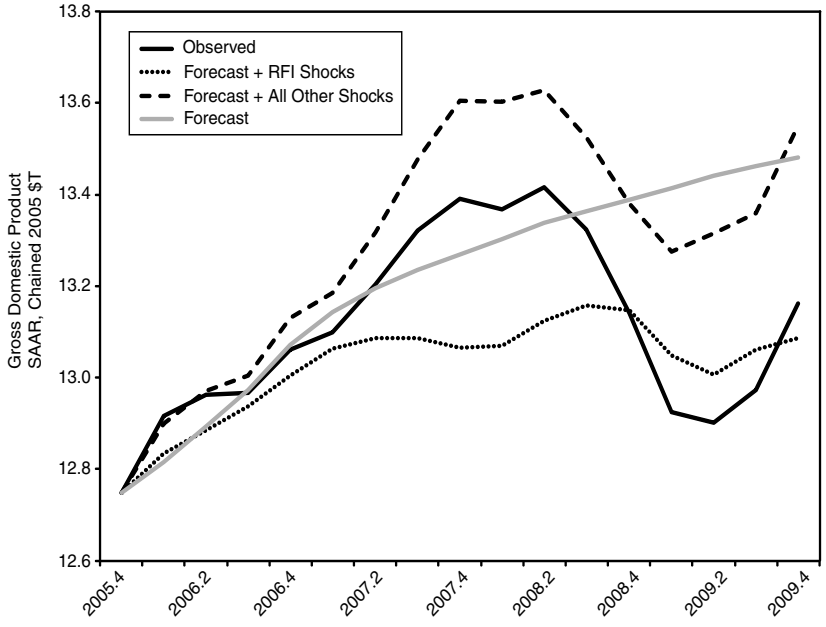

2009 there was a steep decline in GDP. As of the fourth quarter of 2009, GDP remained approximately 3 percent below trend.[8]

Next, turn to the dotted line, which represents the VAR forecast conditional not only on data through 2005 but also conditional on the estimated path for the residential investment shock. Until late 2006 the contribution of the residential investment shock was relatively unimportant in that the forecast conditional on the residential investment shock is close to the trend line. Starting in 2007, the residential investment shock accounts for an ever larger shortfall of GDP relative to trend. Yet until late 2008 GDP remains at or above trend because other shocks are accounting for an ever larger *surplus* of GDP relative to trend. The contribution of these other shocks (the other elements of v_t) is indicated by the gap between the dashed and the solid gray lines in Figure

[8] It is important to emphasize that our version of "trend" is itself influenced by the Great Recession, since the VAR was estimated over the full sample.

7.[9,10] Finally, the effects of other shocks disappear in importance late in 2008, and the cumulative effect of shocks to residential investment leaves output well below trend in mid-2009. In the second half of 2009 other shocks contributed to a significant increase in GDP, allowing it to come back toward trend.

The most notable aspect of Figure 7 is that, from the standpoint of late 2005, the level of output in mid-2009—the trough—can be explained almost entirely by a sequence of shocks to residential investment. Paraphrasing Section 1, we asked to what extent the unusual decline in output could be understood as the usual response of output to an unusual decline in residential investment. Based on the analysis presented in Figure 7, we have a two-part answer: The severe decline in output from 2008:Q4–2009:Q2 is accounted for primarily by shocks other than residential investment. However, in the preceding two years those other shocks had worked in the opposite direction, counteracting the negative effect of residential investment shocks. Thus, although the sudden decline in output late in 2008 cannot be explained by the economy's response to residential investment, the low *level* of output reached in mid-2009 can indeed be explained in this way. When filtered by the VAR, the severe decline in output looks like a delayed response to residential investment shocks that had been accumulating for years.

As a basic reality check on our story that shocks to residential investment have been important, it is useful to look at a version of Figure 7 that substitutes residential investment for GDP. The shock that we are referring to as a residential investment shock ought to be important for the behavior of residential investment. Figure 8 shows this to be the case. Residential investment persistently deviates from trend, and that deviation is overwhelmingly accounted for by the residential investment shock.

Our results are remarkably robust to the way variables are ordered in the VAR. For an ordering with residential investment in the n^{th} position (our baseline has $n = 1$), we call the n^{th} shock the RFI shock. We compare impulse response functions for our baseline to those for $n = 6$ and there is little change in the responses to the RFI shock. We also generate the analogue to Figure 7, showing the contribution of the RFI shocks to GDP in the current recession, for every Cholesky ordering. For every ordering, the difference between GDP in 2009:Q4 and the level explained by RFI shocks only (i.e., the distance between the solid black line and the dotted line) is less than 1.5 percent in absolute value, and for a majority of the orderings the difference is less than

[9] Because the VAR is specified in logs, in this and subsequent figures that plot levels, the distances between (i) the dashed line and the gray line and (ii) the dotted line and the gray line do not add up exactly to the distance between the black line and the gray line. In most cases the discrepancy is small, though it is large in Figure 9.

[10] From Figures 3 and 4 we can see that these offsetting shocks were reflected in the behavior of nonresidential investment and consumption. It would of course be interesting to pursue an economic interpretation of these offsetting shocks, but we leave that for future work.

Figure 8 Residential Fixed Investment: Sample Period 1985–2009

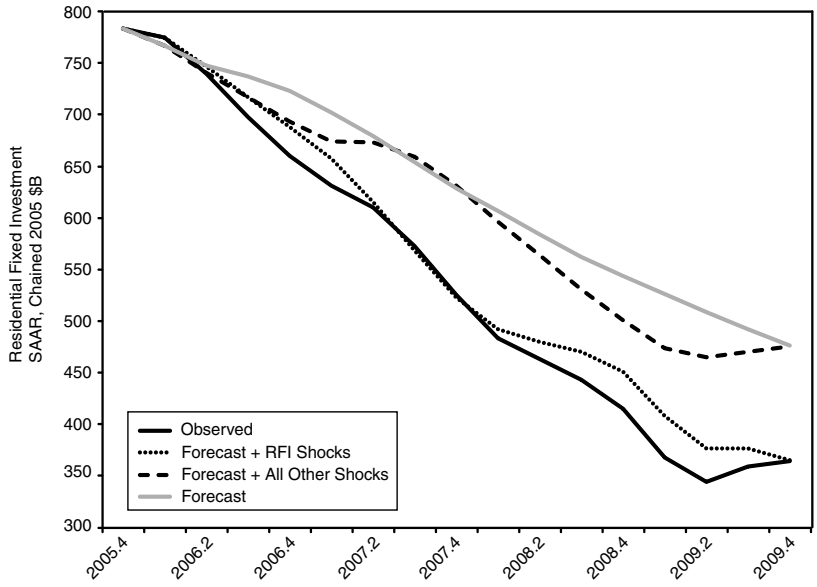

0.75 percent. Each ordering also preserves the property that weakness in GDP coming from RFI shocks was offset by other shocks in 2007 and early 2008.

Out-of-Sample Approach

Figures 7 and 8 were generated by a VAR for the (log) levels of GDP components, estimated from 1985:Q1–2009:Q4. Here we look at the implications of estimating the VAR only through 2005:Q4. Figures 9 and 10 are analogues to Figures 7 and 8, where the estimation period ends in 2005:Q4—these figures display out-of-sample forecasts. In Figure 9, the deviation of output from trend at the end of the sample rises to almost 6 percent from 3 percent with the full sample estimates. It is not surprising that output is further from the perceived trend as of 2005:Q4 than it is from the currently perceived trend. The most striking aspect of Figure 9, however, is the tremendous "overshooting" generated by the residential investment shock. The residential investment shock alone accounts for a shortfall of output from trend that is more than three times as large as the observed shortfall, meaning that the other shocks account for a large positive deviation of output from trend (the dashed line). To understand what is going on here, it helps to look at Figure 10, which plots residential

Figure 9 U.S. GDP: Sample Period 1985–2005

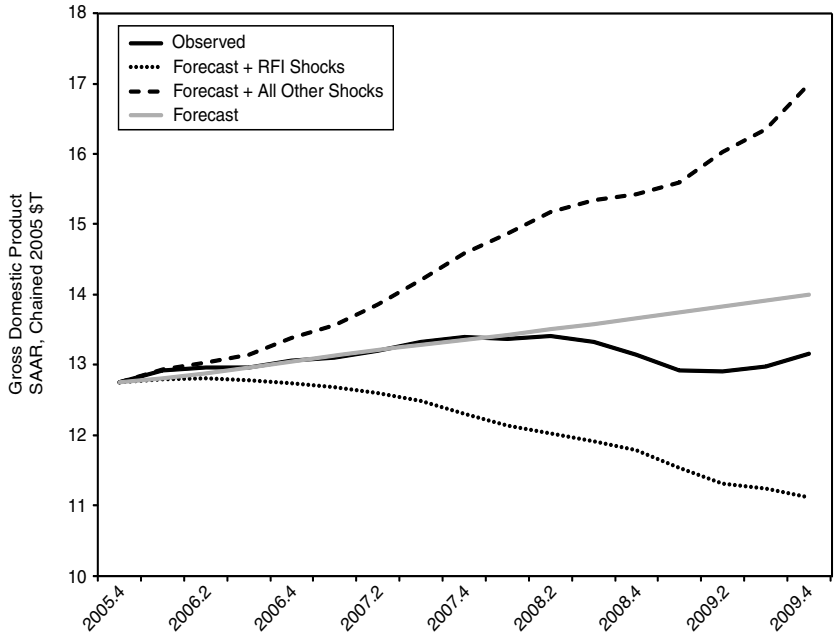

investment for this same case. When we estimate only through 2005, the housing bust is not in the sample, so the post-2005 trend for residential investment does not involve a large decline (compare the solid gray lines in Figures 8 and 10). Of course the large decline in residential investment did occur, and the VAR accounts for that decline primarily with shocks to residential investment (the dotted line in Figure 10). Those shocks were extreme by historical standards so, according to the VAR, they should generate an extreme decline in output—on the order of 21 percent! Such an extreme decline did not occur, so the other shocks must account for an extreme increase in output (the dashed line in Figure 9).

The Housing Boom

Our discussion thus far has centered on the housing bust and the subsequent decline in real GDP. We can also use VAR methodology to investigate the housing boom. Figures 11 and 12 provide the same information as Figures 7 and 8, except that the time scale is shifted back more than five years, to

Figure 10 Residential Fixed Investment: Sample Period 1985–2005

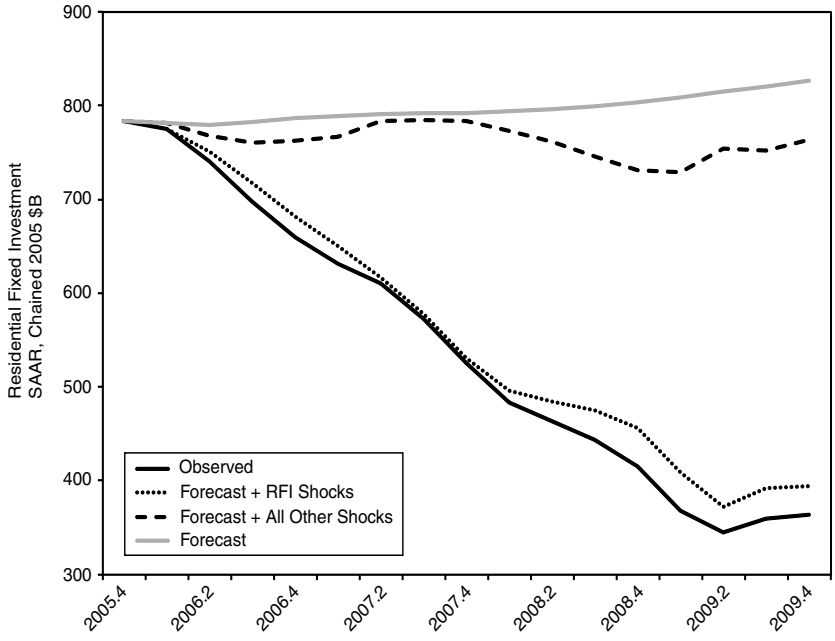

condition on data as of 2000:Q3 and then project ahead to 2006:Q1.[11] It is apparent in Figure 11 that, until late 2003, output was below trend and residential investment shocks accounted for little of the deviation from trend. For the next two-and-a-half years, however, output increasingly exceeded trend, and the deviation is more than accounted for by shocks to residential investment. Figure 12 shows that residential investment deviated persistently from trend for the entire period, but it was only after the middle of 2003 that the residential investment shocks began to account for an increasing share of the deviation. By the end of the boom period, shocks to residential investment accounted for most of the deviation of residential investment from trend. Figures 11 and 12 paint a picture of the housing boom changing shape in mid-2003. It would be interesting to investigate whether there is corroborating evidence for this view. For example, did underwriting standards for home mortgages change around this time?

[11] To generate Figures 11 and 12 we use the same full sample estimates that generated Figures 7 and 8.

Figure 11 U.S. GDP, Housing Boom: Sample Period 1985–2009

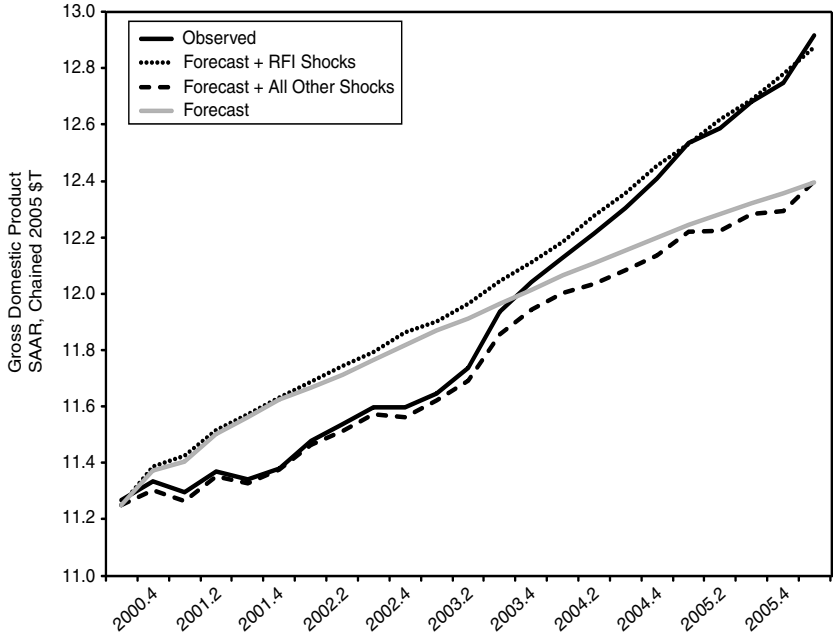

A Broader Perspective on Housing and Output

Although the VAR results are influenced by data as far back as 1985, we have focused exclusively on interpreting the most recent business cycle. We showed that a shock to residential investment plays a central role in accounting for the current shortfall in GDP relative to trend. The current recession is clearly special in terms of its magnitude; is it also special in terms of the temporal relationship between residential investment and output? There is an extensive literature arguing that housing fluctuations do have predictive power for future GDP fluctuations, with Leamer (2007) as perhaps the most forceful proponent. Here we provide some findings consistent with that view.

Figure 13 displays real residential investment for the eight quarters on either side of the peaks of the eight recessions since 1960. For each recession, the level of real residential investment is normalized to be 100 in the quarter of the National Bureau of Economic Research (NBER) business cycle peak. The figure shows that the peak in residential investment preceded the business cycle peak for each recession except for 2001. The current recession is unusual, however, in the length of time during which residential investment fell before the business cycle peak, and in the depth of the decline.

Figure 12 Residential Fixed Investment, Housing Boom: Sample Period 1985–2009

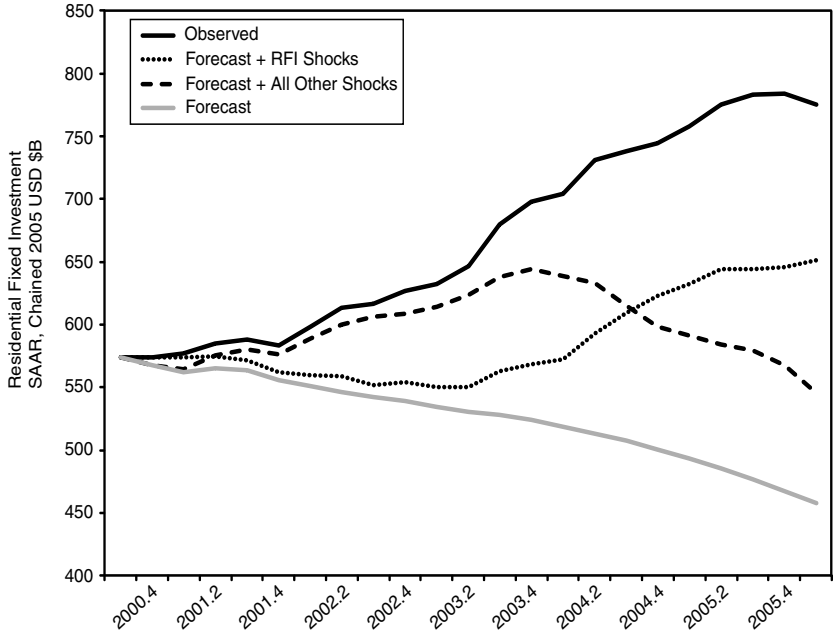

Next, we estimated a two-variable VAR in (log) residential investment and GDP, over the period 1959–2009. The top panel of Figure 14 displays the impulse response of GDP to a residential investment shock when residential investment is ordered first in the VAR. Indeed, a positive shock to residential investment generates a positive, hump-shaped response of GDP, with the peak response of GDP occurring seven quarters after the shock. The same result holds if the ordering of variables in the VAR is reversed. If we estimate the VAR over the more recent sample, the same qualitative relationship holds. However, the elapsed time between the arrival of the residential investment shock and the peak in GDP increases from seven quarters to 14 quarters (bottom panel of Figure 14). The 14-quarter lead time may seem high, given that the most recent peak in residential investment occurred just 10 quarters before the peak in GDP. Note, however, that the peak in residential investment most likely was not a period in which there was a large positive residential investment shock. In fact, the two-variable VAR shows that the peak own response of residential investment to a shock occurs nine periods after the shock, using the estimates for the recent sample.

Figure 13 RFI Across Recessions, Labeled by Peak

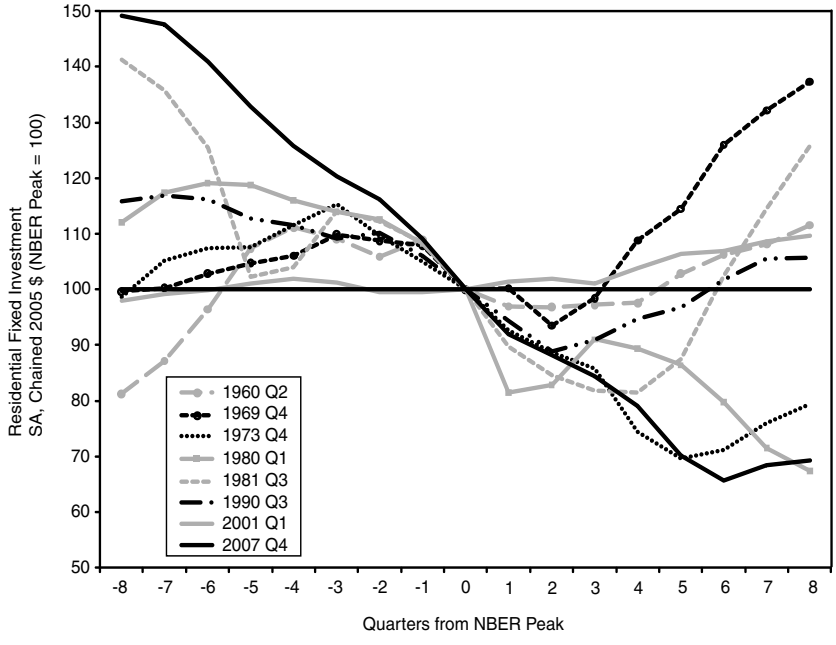

4. VAR FOR COMPONENTS OF EMPLOYMENT

In Section 1 we saw that viewing the economy in terms of employment rather than output led to a rather different picture of the relationship between housing activity and aggregate activity in the Great Recession. According to both metrics, the decline in aggregate activity lagged the decline in housing. However, the raw magnitudes of the decline in housing and aggregate activity were roughly equal in the case of output, whereas overall employment fell by more than 8 million, against a decline in housing employment of just 1.3 million![12] Based on the raw data then, it seems unlikely that identified shocks to housing employment could succeed in explaining the decline in overall employment in the way that residential investment shocks succeeded in explaining the decline in GDP. VAR analysis will allow us to investigate formally whether that intuition holds, or whether there is a large "housing employment multiplier" that amplifies the aggregate effects of shocks to housing employment.

[12] We define housing employment here as the sum of employment in residential building construction and residential specialty trade contractors.

Figure 14 Impulse Response of GDP to RFI Shock: VAR with RFI and GDP, RFI Ordered First

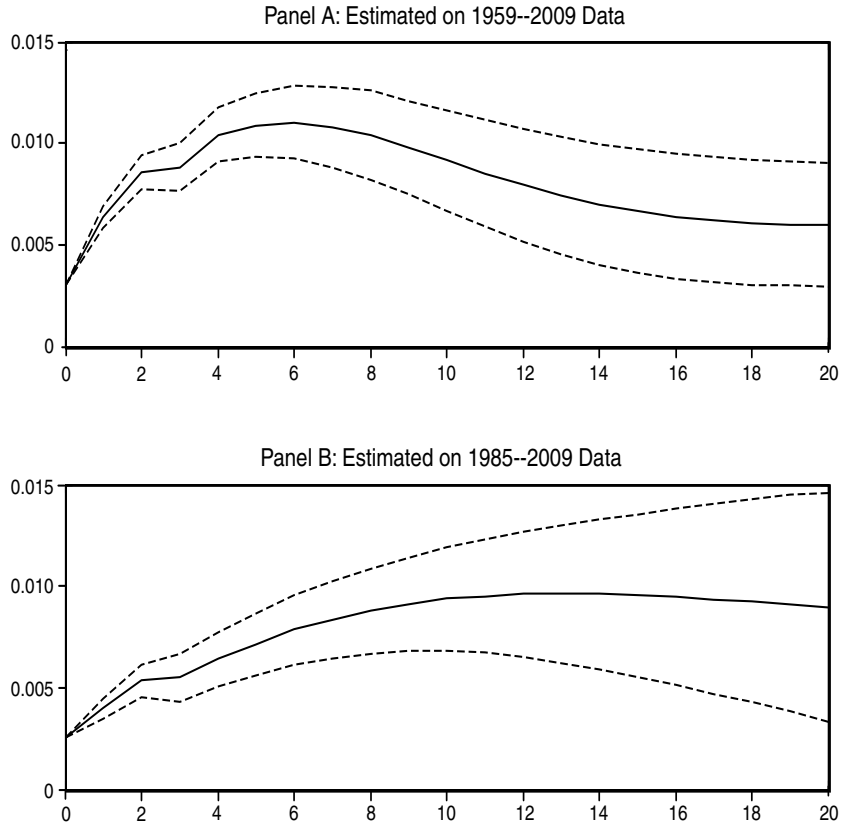

We use data from the Bureau of Labor Statistics's establishment survey to estimate a VAR for the components of employment. Employment is categorized differently than output, so the variables in the VAR will not match up well with the variables in the output VAR. Table 2 lists the components of employment in the VAR, along with their shares of total nonfarm employment in 2002. In specifying a component of employment that represents "housing," one has to choose how to classify *specialty trades* (e.g., plumbing, painting). Specialty trades comprise a sizable fraction of employment in the housing sector: In 2009:Q3 this number was approximately 70 percent. However, prior to 2001, data for specialty trades employees are not broken down into residential and nonresidential components. Thus, the choice is between including or omitting all specialty trades from our housing employment number. We choose to omit specialty trades so as to be confident that the employment category

Table 2 Employment VAR Components

Category	Share in 2002
Services	66.1 percent
Government	16.5 percent
Manufacturing	11.7 percent
Nonresidential Construction	4.5 percent
Residential Construction	0.6 percent
Mining and Logging	0.4 percent

we call "housing" is not contaminated by other areas of economic activity, such as commercial construction. Note that in 2002 residential construction accounted for just 0.6 percent of total nonfarm payroll employment.

Figure 15 displays the same basic set of results for employment that Figure 7 displays for output. The level of employment was on trend, if at all, only through about late 2006, then climbed well above trend before beginning to plummet—both absolutely and relative to trend—in the second half of 2008. At the end of the sample, employment lay less than 2 percent below trend.[13] Unlike what we saw for output, the innovation to employment in the housing sector for much of the sample contributes to a *positive* deviation of employment from trend, although by the end of the sample that shock explains roughly half the deviation of employment below trend. This is consistent with our intuition based on the raw data: The decline in overall employment was so much greater than the decline in housing-related employment that it seemed unlikely that the decline in overall employment could be explained as the usual response to a housing shock. Surprisingly, the housing shock is relatively unimportant in accounting for the weakness in employment in housing (Figure 16).

5. CONCLUSION

The NBER-defined recession that began in December 2007 has been referred to by many as the Great Recession. Indeed, the facts presented in Figure 1 confirm that the current recession stands out as particularly severe among post-1950 recessions. Why such a severe recession occurred is a question that will likely never be answered definitively. However, researchers in academia, government (including central banks), and the private sector have argued convincingly for the importance of various factors in the severity of the recession. Some of the factors that have been discussed are the financial crisis of September 2008, the dramatic increase in oil prices from January 2007–July

[13] We define trend as the level of employment predicted (in sample) by the VAR conditional on data through December 2005. Although employment has fallen much more in percentage terms than GDP, according to the VAR output is actually further from trend than employment.

Figure 15 Payroll Employment: Sample Period 1985–2010

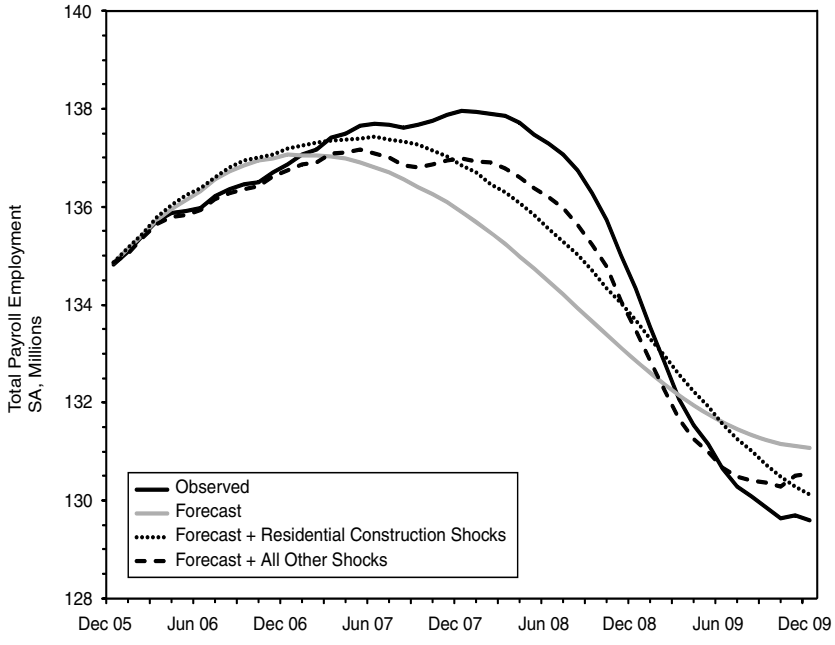

2008 (Hamilton 2009), and inappropriate monetary policy in mid-to-late 2008 (Hetzel 2009, Sumner 2009).

Almost all discussions of the Great Recession, however, include some role for the housing boom and bust. Residential investment declined before and during the Great Recession by about the same amount as GDP, although the decline in GDP occurred with a lag. This observation led us to investigate whether the severity of the recession could be understood as the typical response to a shock emanating from the housing sector. With respect to output, the answer is a qualified "yes": Viewed from the peak of the housing boom, subsequent shocks to residential investment can account for the level of GDP late in 2009. The qualification is that these shocks to residential investment account for approximately zero GDP growth over 2007–2009—they do not account for the sharp decline in late 2008 and early 2009. Similar analysis conducted on employment data attributes much of the shortfall in employment from trend to a housing shock by the end of the sample. However, over much of the sample shocks emanating in home construction push employment *above* trend.

Because there is no economics in our VAR model, our results cannot be used to talk about policies that should have or could have been used to

Figure 16 Residential Construction Employment: Sample Period 1985–2010

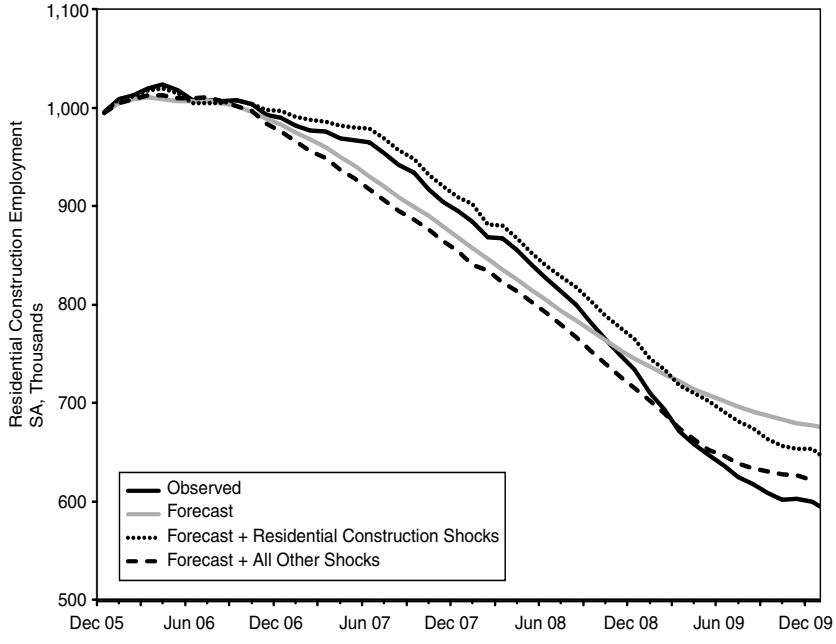

lessen the recession's severity. Further, our results cannot rule out that shocks emanating from monetary or financial policy may have played a role in the Great Recession. However, those results may be useful as an input to future economic modelling that can be used to discuss policy. It is clear that any macroeconomic model used to address the Great Recession ought to have a housing sector. Less trivially, such a model ought to be consistent with (1) the *fact* that housing's contribution to GDP fell by roughly the same amount as the subsequent fall in GDP, and (2) our finding that a shock originating in, or initially reflected in, the housing sector can broadly account for the behavior of GDP during the Great Recession. Any modelling along these lines also faces the challenge that the sectoral behavior of employment followed a very different pattern than the sectoral behavior of output.

REFERENCES

Hamilton, James D. 1994. *Time Series Analysis*. Princeton, N.J.: Princeton University Press.

Hamilton, James D. 2009. "Causes and Consequences of the Oil Shock of 2007–2008." Working Paper No. 15002. Cambridge, Mass.: National Bureau of Economic Research (May).

Hetzel, Robert L. 2009. "Monetary Policy in the 2008–2009 Recession." Federal Reserve Bank of Richmond *Economic Quarterly* 95 (Spring): 201–33.

Leamer, Edward. 2007. "Housing is the Business Cycle." Federal Reserve Bank of Kansas City *Proceedings* 2007: 149–233.

Sims, Christopher. 1986. "Are Forecasting Models Usable for Policy Analysis?" Federal Reserve Bank of Minneapolis *Quarterly Review* (Winter): 2–16.

Sims, Christopher A., James H. Stock, and Mark W. Watson. 1990. "Inference in Linear Time Series Models with Some Unit Roots." *Econometrica* 58 (January): 113–44.

Sumner, Scott. 2009. "The Real Problem was Nominal." www.cato-unbound.org/2009/09/14/scott-sumner/the-real-problem-was-nominal/ (September 14).