DRAFT

Using Structural Shocks to Identify Models of Investment *

John M. Roberts
Federal Reserve Board

August 2005

*The views expressed are those of the author and cannot be taken to represent those of the Board of Governors of the Federal Reserve System or any other member of its staff.

Abstract

This paper uses the response of investment to identified structural shocks to distinguish between some competing models of investment. One issue is the nature of adjustment costs: The classical treatments of adjustment costs have focused on costs to adjusting the level of the capital stock and a number of studies using firm-level data have found that capital stock adjustment costs are important. But a number of recent macroeconomic studies have considered only investment adjustment costs. Another issue concerns whether the response of investment to changes in interest rates is smaller than would be predicted by models with the conventional assumption of a unitary elasticity of substitution between labor and capital.

Investment is U.S. business spending on equipment other than information technology, a relatively homogeneous category that accounts for about half of business fixed investment. The empirical approach is to choose the parameters of the investment model to match as closely as possible the impulse responses from an identified VAR.

The results are sensitive to the impulse responses that are matched. In the preferred results, the estimated elasticity of substitution is estimated to be much smaller than one, both investment- and capital-stock adjustment costs are important, and the size of the capital-stock adjustment costs is in line with estimates from firm-level studies. But if the responses to an identified aggregate demand shock are included among those to be matched, the elasticity of substitution is not significantly different from one. Furthermore, only investment adjustment costs are found to be statistically important. This sensitivity turns on a large crowding out effect from the identified AD shock: The shock initially leads to an expansion in hours and output but a contraction in investment.

In this paper, I use the response of investment to structural shocks to distinguish between some of the major competing models of investment. One issue concerns the nature of adjustment costs. The classical treatments of adjustment costs have focused on costs to adjusting the level of the capital stock. A number of studies using firm-level data have found that capital stock adjustment costs provide an adequate characterization of investment dynamics (Gilchrist and Himmelberg, 1995; Cummins, Hassett, and Oliner, 1999). However, the sluggish response of investment to its determinants observed in macroeconomic data suggests that investment adjustment costs must be present. Recently, Christiano, Eichenbaum, and Evans (2005) have estimated a model in which only investment adjustment costs are present. In this paper, the firm is assumed to face costs of adjusting both investment and the capital stock.

Another key issue concerns whether the response of investment to movements in user cost is muted relative to its response to shocks arising from other sources. In his survey of the investment literature, Chirinko (1993) emphasized this muted response of investment to changes in user cost as one of the important stylized facts about investment. More recently, Tevlin and Whelan (2003) and Roberts (2003) have found such a muted response for U.S. non-high-tech equipment spending. Firm level studies have found that investment responds more strongly to movements in fundamental determinants of capital spending—such as sales and output—than to financial factors, such as the firm's share price (Gilchrist and Himmelberg, 1995; Cummins, Hassett, and Oliner, 1999).

The empirical approach adopted here will be to choose the parameters of the investment model to match as closely as possible the impulse responses from an identified VAR. This approach has been used in a number of recent macroeconomic studies, including Rotemberg and Woodford (1997), Amato and Laubach (2004), and Christiano, Eichenbaum, and Evans (2005). In an identified VAR, the shocks driving the movements in the variables in the VAR are given a structural interpretation by imposing various short-run and long-run restrictions on the effects of the shocks. Ideally, these restrictions have strong justification in economic theory. In the VAR examined here, investment will be included among the variables and so the response of investment to the structural shocks will be traced out. These responses provide a useful way of sorting out some of the competing hypotheses about investment. In particular, one conjecture about why investment responds less to movements in user cost than to shocks stemming from the real side of the economy is that of reverse causation (Chirinko, 1993). Under this hypothesis, interest rates are often low *because of* weak investment. Thus, there is a spurious correlation between low interest rates and low investment. Using structural shocks for estimation avoids the risk of reverse causation: For one thing, we can eliminate the response of shocks to investment from among the impulse responses we choose to match. Furthermore, if we can identify an exogenous shock to monetary policy, then the resulting movements in interest rates and user cost will allow us to trace out the responses of investment to user cost in a way that will presumably not be affected by reverse causation.

In identifying the structural shocks, I mostly rely on short-run restrictions; as Christiano, Eichenbaum, and Vigfusson (2005) have recently emphasized, short-run restrictions have better statistical properties in small samples than long-run restrictions. I do, however, use one long-run restriction in order to be able to identify a technology shock.

In the empirical work, I measure investment as business spending on equipment other than information technology (in the United States). This category excludes from overall business fixed investment spending on structures and the "high-tech" categories of computers, software, and communications equipment. I focus on this narrower category for two reasons. First, the excluded categories have important idiosyncrasies: For nonresidential structures, the lag in response to changes in its determinants appears to be considerably longer than for equipment spending. For information technology, changes in relative prices appear to be a much more important determinant than for other forms of equipment (Tevlin and Whelan, 2003). Second, the ratios of spending in the excluded categories to business-sector output appear to have shifted over time, with the share of structures falling and the share of information technology rising. By contrast, the share of non-high-tech equipment has been more stable. While non-high-tech equipment is a narrow category, it is not unimportant: Over the [1965-2002] period, non-high-tech equipment comprised over half of overall business fixed investment spending.

The structure of the paper is as follows. Section one reviews the canonical model of investment dynamics, under the hypothesis that the capital stock and investment are both costly to adjust. In addition, the elasticity of substitution between non-high-tech investment and labor is allowed to deviate from one. Section two presents the estimates of the structural VAR. Section three presents the estimates of the investment model. Section four presents additional discussion of the elasticity of substitution and section five, conclusions.

1 The Investment Model

The firm is assumed to choose investment and labor input to maximize profits, subject to a production function and a capital accumulation equation. The production function reflects costs to adjusting the level of the capital stock and the rate of investment. Thus, the adjustment costs are incurred by the firm and deducted directly from output.

$$
\begin{aligned}
Max_{\{H,K,I\}} \sum_{t=0}^{\infty} \prod_{s=0}^{t} (1+r_s)^{-1}\{Y_t - W_t H_t - I_t \\
-\Lambda_t[K_{t+1} - I_t - (1-\delta)K_t]\},
\end{aligned}
\tag{1}
$$

where,

$$
Y_t = F(K_t, A_t H_t) - \frac{\gamma_1}{2}\left(\frac{I_t}{K_{t-1}} - \delta - g\right)^2 K_t - \frac{\gamma_2}{2}\left(\frac{\Delta I_t}{I_{t-1}} - g\right)^2 I_t. \tag{2}
$$

Y is output, H is total worker hours, W is the wage, K is the capital stock, I is investment, δ is the depreciation rate of the capital stock, r is the required return on capital, A is an index of (labor-augmenting) technical progress, and g is the economy's average growth rate.

The first-order conditions for investment and the capital stock can be written as:

$$\begin{aligned}
\Lambda_t &= 1 + \gamma_1 \frac{K_t}{K_{t-1}} \left(\frac{I_t}{K_{t-1}} - \delta - g \right) &(3)\\
&+ \gamma_2 \ \frac{I_t}{I_{t-1}} \left(\frac{\Delta I_t}{I_{t-1}} - g \right) \\
&- \gamma_2 \ \frac{1}{1 + r_{t+1}} \frac{I_{t+1}}{I_t} \frac{\Delta I_{t+1}}{I_t} \left(\frac{\Delta I_{t+1}}{I_t} - g \right) \\
&+ \frac{\gamma_2}{2} \left(\frac{\Delta I_{t+1}}{I_t} - g \right)^2 \\
F_{Kt} &= \Lambda_{t-1} \left(1 + r_t \right) - (1 - \delta) \Lambda_t &(4)
\end{aligned}$$

Some comments on the model are worthwhile at this point. First, note that equation 4 is the familiar equation relating the shadow price of the capital stock, Λ, to the present- discounted value of future marginal product of capital. Second, note that when there are only capital-stock adjustment costs ($\gamma_2 = 0$), equation 3 becomes the familiar relationship between investment and the shadow price of the capital stock. Alternatively, when there are only investment adjustment costs ($\gamma_1 = 0$), the shadow price of capital continues to be crucial to investment dynamics, but the relationship now involves the change in investment rather than the level.

Looking ahead to the estimation, we now modify equations 3 and 4, by taking a log-linear approximations and by assuming that F has a constant-elasticity-of-substitution form.

$$\begin{aligned}
\lambda_t &= \gamma_1 \ (1 + g)(\Delta k_t - g) &(5)\\
&+ \gamma_2 \ (1 + g) \left[(\Delta i_t - g) - \frac{g}{1 + \bar{r}} (\Delta i_{t+1} - g) \right] \\
\sigma^{-1}(y_t - k_t) &= \frac{1 + \bar{r}}{\bar{r} + \delta} (\lambda_{t-1} + r_t) - \frac{1 - \delta}{\bar{r} + \delta} \lambda_t &(6)
\end{aligned}$$

where lower-case letters represent the logs of their upper-case counterparts, $\bar{r}$ is the long-run average value of r_t, σ is the elasticity of substitution between capital and labor input, and the approximations, $\Delta i_t \equiv (I_t - I_{t-1})/I_{t-1}$, and $\Delta k_t \equiv I_{t-1}/K_{t-1} - \delta$, have been used.

Some additional manipulations may make the structure of the model more transparent.

$$
\begin{aligned}
(\Delta i_t - g) &= \frac{1}{\gamma_2}\left[\frac{1}{1+g}\lambda_t - \gamma_1(\Delta k_t - g)\right] \\
&\quad + \frac{g}{1+\bar{r}}E_t(\Delta i_{t+1} - g) && (7)\\
\lambda_t &= \frac{1-\delta}{1+\bar{r}}E_t\lambda_{t+1} + \frac{\bar{r}+\delta}{1+\bar{r}}\sigma^{-1}(y_{t+1} - k_{t+1}) - r_{t+1}. && (8)
\end{aligned}
$$

Finally, the log-linearized capital accumulation equation is also needed to complete the model:

$$\Delta k_t = \frac{g+\delta}{1+g}(i_{t-1} - k_{t-1}). \tag{9}$$

2 The Structural VAR

2.1 The data

The structural VAR includes five variables: *Investment* is measured as the change in log of the stock of non-high-tech equipment, multiplied by 100 to make the units comparable to percent changes. The change in the log of the capital stock is, of course, approximately equal to net investment divided by the (lagged) capital stock. *Output* is for the business sector, and is also introduced as the log difference multiplied by 100. *Hours* are for the nonfarm business sector; they are detrended as described below and measured as 100 times the log deviation from trend. *Inflation* is 400 times the log difference in the personal consumption expenditures deflator. And the federal funds rate is introduced in percent.

Hours are detrended by running the Hodrick-Prescott filter through the log of hours relative to the working-age population, using an HP filter coefficient of 64,000. The hours trend is also taken out of output and the equipment stock prior to log differencing. The data are quarterly, consistent with the frequency of the investment, output, and hours data. The estimation period is

1965 to 2002; the starting date is limited by the availability of the investment data.

The reader may wonder why a stochastic trend is allowed to remain in productivity while hours are detrended prior to estimation. One reason based on economic logic is that while shocks to trend productivity are frequently posited as an important source of business-cycle variation, shocks to trend labor supply are a less frequent object of inquiry. Another reason is that, when per capita hours are allowed to follow a stochastic trend, the variance of that trend is estimated to be considerably smaller than that of productivity: For example, in an examination of stochastic trends in hours and productivity using state-space methods, Roberts (2001) finds that the variance of the productivity trend is about four times larger than the variance of the hours trend. This result suggests that the focus of the literature on the implications of shocks to trend productivity rather than to trend hours is well-founded.

2.2 The identification scheme

As noted in the introduction, the identification of the structural VAR is based mostly on short-run restrictions. First, it is assumed that investment is not affected by any current-period information. Given the likely existence of planning lags in investment, this seems a safe assumption (see Edge, 2000, for a discussion). On the other hand, investment is allowed to affect the other variables in the model contemporaneously; given that investment is, mechanically, a component of output, this seems a reasonable assumption. At the other extreme, the federal funds rate is assumed to be affected by all of the other variables in the model in current period, but is assumed not to affect any of the other variables. These assumptions are similar to those used by Christiano, Eichenbaum, and Evans (1999, 2005) to identify a shock to monetary policy. Briefly, the economic logic underlying this set of assumptions is that the central bank is well-informed about current-quarter economic activity and inflation, but that the private sector takes time to react to changes in monetary policy. The investment planning lags just discussed are one example of the behavior underlying such decision lags.

Inflation is allowed to be affected by current-period economic activity (hours and output) but not by monetary policy. While firms are likely to be aware of current market conditions for their own products, they have less incentive to monitor monetary-policy developments, leading to lags in the impact of monetary policy on inflation relative to the impact of hours and output. While inflation is assumed to be affected by economic activity contemporaneously, it is not allowed to affect activity in the current period.

So far, then, we have identified three shocks—to investment, inflation, and monetary policy—with contemporaneous restrictions. Yet to be disentangled are the shocks to the output and hours equations. There are likely many shocks that would affect both output and hours simultaneously: Shocks to aggregate demand would likely raise both, for example, and productivity shocks could also well affect both variables contemporaneously. One important productivity shock, however, can be identified by its long-run consequences. In particular, technology shocks can reasonably be assumed to be the only shocks that will affect output per hour in the long run. I therefore impose the restriction that of the two remaining structural shocks, only one affects output in the long run; the other is restricted to have only temporary effects on output.

The equations for investment, inflation, and the federal funds rate can be estimated with ordinary least squares. I estimate the equations for hours and output using the instrumental variables approach of Shapiro and Watson (1989)—modified to account for the fact that investment is exogenous to hours and output.

2.3 VAR results

Figure 1 shows impulse responses from the structural VAR, along with a 90 percent confidence band. I discuss the VAR results organized by the impulse responses to each shock, with the ordering roughly reflecting the degree to which each shock can claim to be "structural."

2.3.1 Monetary-policy shock

The effects of the identified monetary policy shock are similar to others found in earlier work: A shock to monetary policy that raises the federal funds rate leads to a reduction in output and in labor input within a few quarters. Output is below its baseline value for about three years, and the effect is statistically signficant for about half that time. Hours follow a similar pattern. Consistent with the well-known phenomenon of procyclical labor productivity, output falls more in percentage terms than do hours.

Inflation initially rises following the monetary-policy shock—the well-known “price puzzle.” Within four quarters, however, inflation is below its baseline level. Inflation remains low for an extended period thereafter.

The monetary-policy shock also leads to a drop in investment. According to the canonical model, the monetary-policy shock affects the shadow price of capital, and thus investment, through two possible channels. One channel is through higher interest rates, which reduce the shadow price of capital by causing dividends to be discounted more heavily. In the other channel, the drop in output implies a lower marginal product of capital both now and in the future and thus pushes down the shadow price of capital. One goal of the structural estimation will be to distinguish between these channels.

2.3.2 Technology shock

The technology (or trend productivity) shock has an immediate positive effect on output. Output then rises somewhat more before settling down to a level that is about the same as that immediately following the shock. Hours initially fall somewhat in response to the trend productivity shock. However, within four quarters, hours are above their baseline level. Investment also responds (with a lag) to the trend productivity shock. The response of investment is as expected given a shock that permanently raises output: If there is a permanent improvement in technology, there will be new investment opportunities.

One interpretation of the hump-shaped response of output, hours, and investment to the trend productivity shock is that the increase in trend productivity initially stimulates aggregate demand somewhat more than it raises aggregate supply. This interpretation is confirmed by the pattern of responses for inflation: The boost to productivity initially pushes down inflation, consistent with the reduction in costs. Eventually, however, inflation rises above its baseline level, consistent with the Phillips-curve notion that when output and hours exceed their equilibrium levels, inflation rises.

The initial effect of the trend productivity shock is to lower the federal funds rate. This response might be expected in a simple monetary-policy reaction function in which the central bank is responding to the initial drop in hours and inflation. Once hours and inflation exceed their baseline values, however, so does the funds rate.

One disappointing aspect of the effects of the identified productivity shock is that they are not estimated very precisely. While the initial effect of the shock on output is statistically significant at the 90 percent level, the confidence bands widen considerably thereafter, and within a few years after the initial shock, the effect is no longer statistically significant. None of the effects on the other variables ever exceeds 90 percent significance. Notably, beyond the first few quarters, the estimated effects of the trend productivity shock on investment have very wide confidence bands. Preliminary attempts to sharpen the estimates—for example, by reducing the number of lags included in the VAR from four to three—were unsuccessful; the issue of how to sharpen these estimates will be revisited in the conclusion. The imprecision of these estimates will limit the ability to make useful inferences based on the response of investment to productivity shocks.

2.3.3 Inflation shock

The shock to inflation has a persistent effect on inflation. The shock also leads to statistically significant declines in hours, output, and investment, with the effects reaching a peak about three or four years after the initial shock. The federal funds rate rises subsequent to the shock, consistent with the view that tight monetary policy may be responsible for the subsequent weakness in real economic activity.

The effect of the inflation shock on output is highly persistent. Indeed, the depressing effect is statistically significant at the 90 percent level even after ten years. This highly persistent effect was not ruled out in the estimation—it was only the IS shock that was restricted from having a permanent effect on output. Thus, one interpretation of the inflation shock is that it represents an additional supply shock that has a permanent (depressing) effect on output.

2.3.4 IS shock

Recall that, with the exception of the hours and output shocks, the model was identified using contemporaneous restrictions. Long-run restrictions were used to identify a shock to trend productivity shock. Conventional wisdom would suggest that the complement to the productivity shock ought to be an "aggregate demand" shock. The impulse responses to this shock bear out this prediction: The shock initially raises all three of hours, output, and inflation. Because the effects of monetary policy have been accounted for separately, it seems reasonable to associate this shock with the textbook "IS" shock.

The shock leads to an increase in the federal funds rate, consistent with the notion that countercyclical monetary policy should act to offset demand shocks. The increase in the funds rate is highly statistically significant, and durable: The funds rate remains significantly above its initial level in a statistical sense for several years, and the point estimate does not return to its initial level until about five years after the initial shock. Given this strong reaction of monetary policy to the shock, it is perhaps not surprising that both hours and output drop below their baseline levels within eight-to-ten quarters after the initial shock; with a lag, inflation follows.

By assumption, the IS shock has no permanent effect on output. As can be seen, however, after its decline, output returns to its baseline value only very gradually. Still, after about five years, the effect is well within the 90 percent confidence band.

After rising a bit initially in response to the shock, investment then drops sharply. Investment remains significantly below its baseline value for a considerable period following the initial shock. Such a drop in investment would make sense if there were an important effect of interest rates on investment. This interpretation suggests that when there is a surge in aggregate demand that originates outside the investment sector, it tends to have an important "crowding out" effect on investment. Furthermore, because this crowding-out effect is quite precisely estimated and so will have an important effect in the structural estimation.

2.3.5 Investment shock

Because of concerns over reverse causation, the shock to the investment equation will not be used in the estimation. Still, for completeness, it is worth noting some of the properties of this shock. The shock leads to a persistent increase in investment. The effect is initially slightly hump-shaped; investment then returns to baseline within three years after the initial shock. The increase in investment leads to an increase in output. That is perhaps not surprising since, mechanically, investment is one of the components that goes into the expenditure-side adding up of output. The investment shock also leads to an increase in hours that is strongly statistically significant and follows closely the movements in output and investment. The effect of the shock on hours suggests that the investment shock is indeed capturing an authentic effect and is not simply measurment error.

Inflation rises somewhat in the wake of the investment shock, although the effect is not statistically significant. The federal funds rate also increases—and the rise is statistically significant. An interpretation of this increase in terms of a monetary-policy reaction function is that monetary tightens as hours and output exceed their baseline values. This reaction of interest rates suggests that reverse causation could indeed be a risk.

3 Estimation

3.1 Estimation approach

I estimate the model by choosing the structural parameters—γ_1, γ_2, and σ—so as to make the impulse responses from the structural model as close as possible to those from the VAR. This approach has been used, for example, by Rotemberg and Woodford (1997), Amato and Laubach (2004), and Christiano, Eichenbaum, and Evans (2005). One important difference between the current application and these others, however, is that while they were choosing the parameters of all of the equations in the model, here, reduced-form specifications are used for the variables other than investment. In particular, I use the same VAR specifications that were used in the reduced-form VAR, holding the coefficients fixed. One advantage of this approach is that it is more robust to the possibility of misspecification in the non-investment parts of the model. As we saw in the previous section, however, one disadvantage is the large degree of uncertainty surrounding the VAR impulse responses.

Another advantage of the impulse-response-matching approach is that it allows estimation to focus on the effects of individual shocks. CEE, for example, estimate their model using only the responses of the economy to monetary-policy shocks. As they explain, one reason for this focus is that the assumptions underlying monetary-policy shocks may be especially credible, making inference based on these shocks of particular interest.

Following CEE's logic, I begin by estimating the model matching only the responses to the monetary-policy shock. I then broaden the analysis to consider other shocks as well. Technology shocks have received a great deal of focus as an important driving force in the economy, and so I next consider those shocks. On the other hand, shocks to investment are likely the least informative about the parameters underlying investment dynamics. As a consequence, I do not look at the responses to these shocks in the estimation.

The identified IS-curve and inflation shocks lie somewhere in between. These shocks have been less widely studied and there thus is less concensus on what their effects should be. In particular, the strong crowding-out effect of the IS shock has not, to my knowledge, been documented elsewhere. While we will examine the implications of this shock for the estimation of the model, it may be premature to draw strong conclusions from its effects.

3.2 Calibrating long-run averages

The variable r_t in the structural model is the required return on capital, whereas the VAR includes the federal funds rate, which is a risk-free interest rate. These variables will differ by a risk premium (as well as by expected inflation). If we assume a constant risk premium, then r_t moves one-for-one with changes in the real federal funds rate:

$$r_t = premium + rff_t - E_t dp_{t+1}, \qquad (10)$$

where *rff* is the federal funds rate and $E_t dp_{t+1}$ is the expected inflation rate. The parameter $\bar{r}$ is the average net return on capital (and not the average real fed funds rate). In the analysis, $\bar{r}$ is assumed to be 8 percent. Over the 1965-2002 period, the real federal funds rate averaged 3.2 percent, so the implicit premium is about 5 percent.

Two other parameters are calibrated: The average depreciation rate for non-high-tech equipment, δ, was 12 percent per year over the 1965-2002 period. And, g, the average growth rate of business-sector labor productivity, was 2 percent per year over this period.

3.3 Results

Table 1 presents results using only the monetary-policy shock to identify the model. In column 1, the elasticity of substitution between capital and labor input, σ, is constrained to be one; the adjustment cost parameters γ_1 and γ_2 are freely estimated. Both of the adjustment-cost parameters are statistically significant at the 10 percent confidence level (although not at the 5 percent level). The point estimates are quite reasonable: They imply that, in the absence of investment adjustment costs, the capital stock adjusts at about a 23 percent annual rate toward its target level. For investment, the estimates imply that investment adjusts about 28 percent *per quarter* toward its implicit target level.[1] These results are similar to those of Roberts (2003). As discussed in Roberts (2003), this pace of capital-stock adjustment is similar to that found in firm-level studies of investment dynamics. The implicit investment adjustment speed suggests that investment moves rapidly toward the path

[1] The capital-stock adjustment speed is computed by setting γ_2 to zero in equation 5, using this expression to substitute for λ into equation 6, and then solving for the roots of the resulting quadratic equation in k_t. For the investment adjustment speed, equation 9 is used to replace Δk_t in equation 5 with i_t. We then have a simple dynamic equation in invesment, with the implicit target investment rate proportional to λ_t.

dictated by the adjustment of the capital stock. As can be seen in the dashed line in figure 2, these estimates are consistent with the hump-shaped pattern of the response of investment to a monetary-policy shock reflected in the reduced form.

The estimates in column 2 impose the restriction that $\gamma_1 = 0$. This restriction was imposed, for example, by Christiano, Eichenbaum, and Evans (2005). As can be seen, with this restriction, investment is estimated to be much more costly to adjust than in the unconstrained case. γ_2 is now more precisely estimated, and is statistically significant at the 5 percent level. It is worth noting that although γ_1 was not highly statistically significant in column 1, imposing the restriction that $\gamma_1 = 0$ is strongly rejected, as the minimum distance criterion rises from 2.7 to 9.1.[2] The estimate of γ_2 in column 2 is about three times larger than that reported by Christiano, Eichenbaum, and Evans. There are a number of differences between this analysis and CEE's. A potentially important difference is that CEE look at a broad measure of investment while the analysis here uses non-high-tech equipment. Figure 2 shows estimates both with and without the restriction that $\gamma_1 = 0$ imposed. As can be seen, both well approximate the shape of the VAR impulse response for the first eight quarters following the shock. After that, however, the model that allows for capital-stock adjustment costs implies a faster return of investment to its baseline level and thus captures the VAR impulse response better.

In the final column, γ_1 is set equal to 30. This value is chosen to be consistent with a 25 percent annual capital-stock adjustment speed, in line with the firm-level evidence summarized in Roberts (2003). As can be seen, the overall fit of the model is little changed with this restriction.[3]

[2]Unfortunately, it is not possible to infer an investment adjustment speed based on the estimates in column 2. When $\gamma_1 = 0$, the capital stock is assumed to be costless to adjust. Hence, there is no well-defined target investment rate.

[3]Although the value of γ_1 in column 3 is 40 percent smaller than the value in column 1, the implicit capital-stock adjustment speed is only about 8 percent higher, reflecting the strongly nonlinear relationship between the adjustment-cost parameters and the adjustment speeds. This example suggests that we should be wary of putting too much weight on seemingly large changes in the adjustment-cost parameters, as the adjustment speeds are of greater economic relevance.

In table 2, the model is estimated by fitting only the impulse responses to the technology shock. As can be seen in column 1, here, the estimated value of γ_1 is pushed to its theoretical minimum of zero. The estimate of γ_2 is small relative to the values in table 1, but it is not precisely estimated. In column 2 of table 2, γ_1 is set equal to 30, consistent with moderate capital-stock adjustment costs. The fit of the model deteriorates somewhat; the point estimate of γ_2 is larger, but it remains statistically insignificant. Overall, the technology shock, by itself, gives little guidance to the parameters of the model.

Table 3 presents estimates based on fitting both the monetary-policy and technology impulse-repsonse functions. In the first column, σ is constrained to be one, as in the earlier estimation. Relative to the estimation based on monetary-policy shocks alone, the point estimates are little changed, although they are a bit more precise. In column 2, the assumption that $\sigma = 1$ is relaxed. As can be seen, the resulting point estimate of σ is very small—around 0.1. Allowing free estimation of σ improves the fit of the model notably; put another way, we can strongly reject the hypothesis that $\sigma = 1$. Looking at the other parameters, the point estimate of γ_1 drops while that of γ_2 rises. The estimate of γ_1 is now quite imprecise, although it is close to the moderate benchmark value of 30. In columns 3 and 4, γ_1 is constrained to be 30 and 0, respectively. As can be seen, neither restriction has much effect on the fit of the model, consistent with the impression given by the t-ratio on γ_1 in column 2. Figure 3 compares the impulse responses both with and without the restriction $\sigma = 1$ imposed. Both specifications capture the VAR response to a monetary-policy shock about equally well, but the specification with a smaller value of σ allows the model to do a better job approximating the response to the technology shock.

Table 4 adds the inflation shock to the mix. Relative to the estimates in column 2 of table 3, the estimate of γ_1 is more precise and is now statistically significant at the 5 percent level. The point estimate of σ remains well less than one, but it is higher than before and its difference from one is no longer statistically significant. Columns 2 and 3 impose two values of γ_1, 30 and 0. Given the precision of the γ_1 estimates in column 1, it is not surprising that imposing $\gamma_1 = 0$ leads to a significant reduction in fit. Note, however, that with this restriction, the point estimate of σ drops by somewhat more than one-half and is now significantly different from one. This result suggests that while it may be difficult to obtain precise estimates of both γ_1 and σ, it appears that the joint hypothesis of $\gamma_1 = 0$ and $\sigma = 1$ can be strongly rejected. The final two columns consider two alternative values of σ, 1.0 and 0.07. As can be seen, setting $\sigma = 1$ has little effect on the fit of the model; the estimates of the adjustment cost parameters are little changed and remain strongly statistically significant. Setting $\sigma = 0.07$ leads to a larger reduction in fit, and this restriction can be rejected at the 5 percent confidence level.

Figure 4 compares the impulse responses from the structural model estimates reported in column 1 of table 4 with the VAR impulse responses. With this set of parameter estimates, the structural model fits the technology shock IRF less well than the best-fitting model shown in figure 3. However, the structural model does a good job of capturing the effects of the monetary-policy and inflation shocks on investment.

Table 5 reports results based on matching the impulse responses of four shocks—to monetary policy, technology, inflation, and the IS curve. In table 5, the restriction $\sigma = 1$ is once again imposed: Unconstrained, σ tends toward very high values, but the value of the minimum-distance criterion changes little from its value with $\sigma = 1$ imposed. In column 1, the estimate of γ_1 is very small and not significantly different from zero. In column 2, the restriction $\gamma_1 = 0$ is imposed. Not surprisingly, neither the fit nor the coefficient estimates are much affected. In particular, γ_2 is 6.7, and quite precisely estimated. In column 3, moderate capital-stock adjustment costs ($\gamma_1 = 30$) are imposed. This restriction is strongly rejected.

Figure 5 shows the impulse responses from the model with the parameters shown in column 2 of table 5. Compared with the structural-model impulse responses shown in figure 4, the model with these parameters captures the responses to the monetary-policy and inflation shocks about equally well over the first eight to ten quarters. After that, however, the simulations shown in figure 5 predict considerably more persistence in investment than in the VAR IRFs or than in the model results shown in figure 4. This shift likely reflects the fact that $\gamma_1 = 0$ in figure 5; as we saw in figure 2, when capital-stock adjustment costs were not allowed, the investment-adjustment costs were high, implying a high degree of persistence in investment. In figure 5, the adjustment costs are estimated to be high because of the response of investment to the IS curve shock in the VAR, where there is a high degree of persistence. Overall, the model has a hard time fitting the responses to this shock, and the responses from the estimated structural model breach the VAR's 90 percent confidence interval in the first few quarters following the shock.

There is a possible economic explanation of the high value of σ that appears to be needed to take account of the effects of the IS shock. Recall from section 2 that while the IS shock leads initially to an increase in output and hours, it pushes down investment after a few quarters. Because the drop in investment is preceded by an increase in interest rates, one interpretation is that the drop represents an important crowding out effect, whereby the increase in interest rates leads to a reduction in investment despite elevated output. Crowding out would seem to suggest, then, that interest rates have important effects on investment and thus that σ is large. Furthermore, the crowding-out effect in the VAR was strongly statistically significant, so that taking account of it swamps the influence of the other shocks.

All that having been said, there are reasons to be suspicious of the results based on the IS shock. As noted earlier, while the crowding-out effect appears to be large and statistically significant in the structural VAR, it has not been widely remarked upon. Before drawing strong conclusions from the implications of this shock, we would want to confirm that it is a robust phenomenon. Moreover, it is worth noting that the overall fit of the model is much worse in table 4 than in the earlier tables. For example, the value of the minimum distance criterion rises from about 9 from the best-fitting models in table 3 to about 45 in column 1 of table 4. While a χ^2 statistic of 90 is far from indicating a statistically significant rejection of the model when there are 77 degrees of freedom, these results nonetheless suggest that the parameters needed to take account of the IS shock are very different from those that best fit the responses to the other three shocks: Recall that the joint restriction $\gamma_1 = 0$ and $\sigma = 1$, which seems to work best in table 5, was strongly rejected when only the responses to first three shocks shocks were considered in table 4.

4 Long-run and short-run evidence on the elasticity of substitution

The results using the responses to all of the shocks—including the response to the IS shock—suggested a value of σ equal to one. An elasticity of substitution equal to one is convenient, of course, because it makes balanced growth simple to calculate. But this result was sensitive to the set of impulses we chose to match, and dropping the IS responses—which were the most suspect—suggested a point estimate of σ that was well below one. It is therefore worth considering the possibility that $\sigma < 1$.

Looking outside the current setting, one piece of evidence that might point to a unit elasticity is the nominal share of non-high-tech investment, which has been fairly stable over time. [4] Whether this piece of evidence is informative about the elasticity of substitution, however, depends on the degree of trending behavior in the relative price of non-high-tech investment goods: If the relative price are trending (down) over time, then the stability of the nominal spending share is indicative of a unit elasticity. If, however, the relative price is stationary, then the stability of the share doesn't tell us much about the elasticity of substitution.

It turns out that the the trend in the relative price of non-high-tech investment goods is a matter of some controversy. In the official, national accounts data, there is little trend movement in these relative prices, and thus the stability of the nominal spending share is not very informative. However, Cummins and Violante (2002) have argued that the national accounts miss a good deal of quality improvement. They show that adjusting for overlooked quality improvements takes about 3 percentage points per year off the average rate of price increase. In that case, there would be an important trend in the relative price of non-high-tech equipment, in which case the stability of the nominal spending share would indeed indicate that the elasticity of substitution is one. [5]

[4] For example, the average share over the 1999-2002 period (6.4 percent of business-sector output) was about the same as the average level in the 1959-1962 period.

[5] It is also possible that there is upward bias in price indexes for other components of

There is thus some possibility that long-run relationships would suggest that the elasticity of substitution is one while short- and medium- run impulse responses suggest that it is much smaller. One possibility is that it is costly to change the capital-output ratio, perhaps because it is costly to reorganize production. This possibility has been considered by Roberts (2003). That paper found that such reorganization costs could provide a reasonable explanation for the attenuated response of investment to changes in user cost.

5 Conclusion

The results of this paper have a key sensitivity: When the responses to an identified IS shock are included among those to be matched, the response of investment to changes in user cost is estimated to be strong, implying an elasticity of substitution that is large and not signficantly different from one. In this case, only the investment adjustment cost is found to be estimated precisely; the point estimate of the capital stock adjustment cost parameter is estimated to be small and not significantly different from zero. On the other hand, when the IS responses are dropped from the set of impulse responses to be matched, the estimated elasticity of substitution is estimated to be much smaller than oneand the joint hypothesis that the elasticity is one and capital-stock adjustment costs are zero can be strongly rejected. Moreover, the results generally suggest that both investment- and capital-stock adjustment costs are important; the size of the capital-stock adjustment costs is in line with estimates from firm-level studies while the investment adjustment costs are relatively small and suggest that investment adjusts rapidly to the pace determined by the capital-stock adjustment process.

The argument that more information is better might argue in favor of preferring the results that include the IS impulse responses. There are reasons, however, to be suspicious of these results. For one, the overall fit of the model deteriorates when this shock is added to the mix: The responses to the other

spending. However, in their comprehensive survey, Lebow and Rudd (2003) estimate the bias in consumer prices at only 1 percent per year. Hence, at least relative to consumer prices, non-high-tech equipment prices would still have a significant downward trend.

three parameters can be jointly fit quite closely, whereas adding the IS curve leads to a compromise that is [less desirable]. Also, the strong crowding out effect of the IS shock—which is the likely source of the large elasticity estimates when this shock is used—has not been documented elsewhere and so prudence is called for in interpreting the results based on this shock.

Because they abstract from the effects of the shock to investment, all of the estimates in this paper are robust to reverse causation. Hence, the finding of a small elasticity of substitution in the set of results that exclude the IS shock suggests that reverse causation is not the source of this result.

An unfortunate aspect of the present estimation exercise has been the wide confidence interval around many of the impulse responses, esepcially for the technology shock. These wide confidence intervals limited the ability to obtain sharp estimates. One approach that might allow for more precise estimation would be to use structural specifications for the remaining equations of the model. Indeed, recent work by Del Negro and Schorfheide (2004) suggests that, once allowance is made for the number of parameters involved, structural models may actually dominate VARs as characterizations of the data. Because fewer parameters are involved, estimates from structural models are likely to be more precise. Structural estimation would also be useful for verifying the empirical relevance of the large crowding out effect found here.

Table 1
Estimates of Structural Investment Model
Based on matching monetary-policy IRFs only

	(1)	(2)	(3)
γ_1	49.8 (27.6)	0.0 --	30.0 --
γ_2	5.3 (2.7)	9.0 (3.6)	5.0 (2.2)
σ	1.0 --	1.0 --	1.0 --
Minimum distance criterion	2.7	9.1	3.3
Adjustment speeds: Capital stock Investment	 23% AR 28% QR	 n.a.	 25% AR 30% QR

Table 2
Estimates of Structural Investment Model
Based on matching technology IRFs only

	(1)	(2)
γ_1	0.0 --	30.0 --
γ_2	1.7 (1.5)	3.3 (3.4)
σ	1.0 --	1.0 --
Minimum distance criterion	1.8	3.3

Notes: Numbers in parentheses are standard errors.
-- indicates constrained parameter.
IRF—impulse response function.
AR—annual rate.
QR—quarterly rate.
n.a.—not applicable.

Table 3
Estimates of Structural Investment Model
Based on matching both monetary-policy and technology IRFs

	(1)	(2)	(3)	(4)
γ_1	41.6 (23.0)	24.5 (27.7)	30.0 --	0.0 --
γ_2	4.7 (2.1)	7.8 (4.3)	7.7 (4.2)	10.2 (4.8)
σ	1.0 --	.07 (.05)	.07 (.05)	.05 (.03)
Minimum distance criterion	6.6	2.9	3.0	3.8
Adjustment speeds: Capital stock Investment	 24% AR 44% QR	 26% AR 16% QR	 25% AR 19% QR	 n.a.

Table 4
Estimates of Structural Investment Model
Based on matching monetary-policy, technology, and inflation IRFs

	(1)	(2)	(3)	(4)	(5)
γ_1	39.0 (19.4)	30.0 --	0.0 --	39.5 (18.2)	59.2 (40.8)
γ_2	5.7 (2.9)	5.5 (2.6)	9.8 (4.0)	4.9 (2.0)	15.7 (5.2)
σ	.40 (.48)	.38 (.41)	.17 (.10)	1.0 --	.07 --
Minimum distance criterion	8.8	9.1	15.0	9.1	13.1
Adj't speeds: Capital stock Investment	 24% AR 34% QR	 25% AR 27% QR	 n.a.	 24% AR 40% QR	 22% AR 19% QR

Notes: See tables 1 and 2.

Table 5
Estimates of Structural Investment Model
Based on matching monetary-policy, technology, inflation, and IS-curve IRFs

	(1)	(2)	(3)
γ_1	1.5 (12.1)	0.0 --	30.0 --
γ_2	6.7 (1.7)	6.7 (1.7)	7.7 (2.3)
σ	1.0 --	1.0 --	1.0 --
Minimum distance criterion	45.5	45.5	51.5

Notes: See notes to tables 1 and 2.

Figure 1
Impulse Responses in Identified VAR

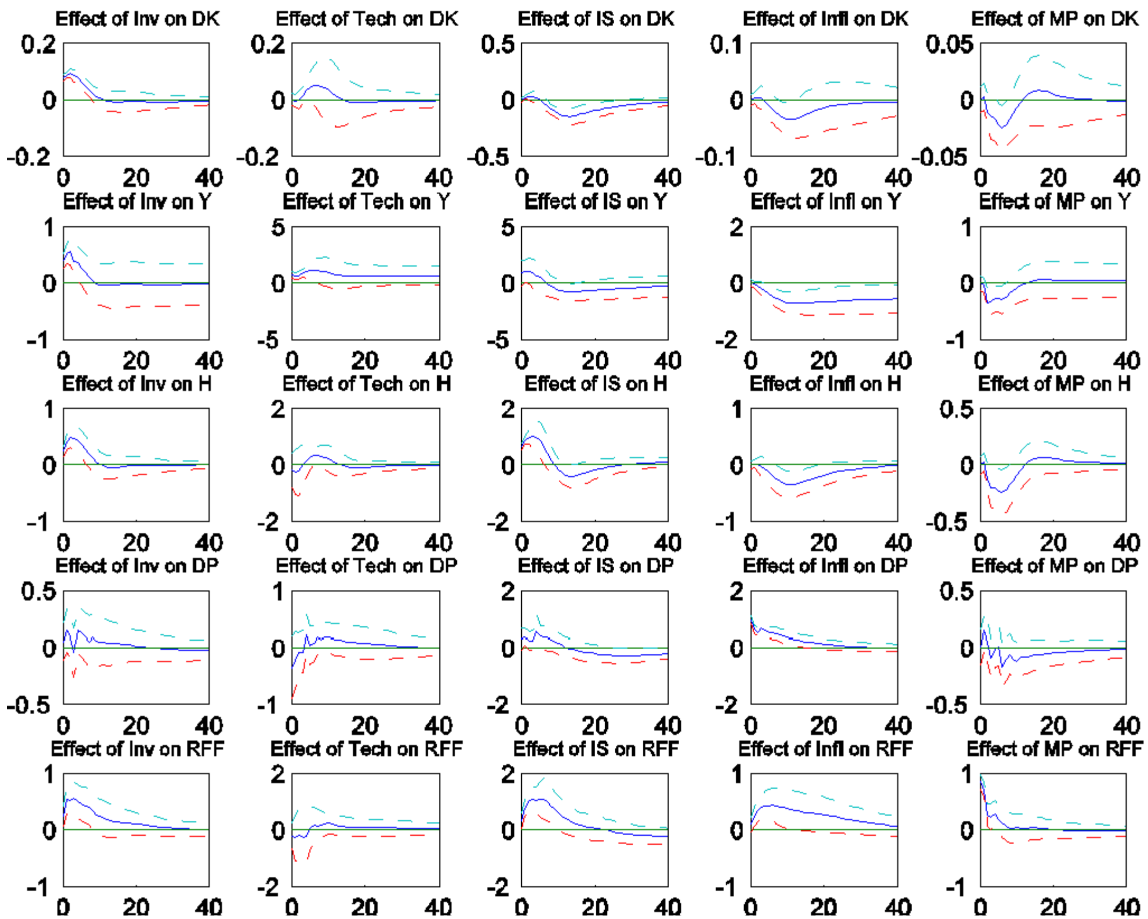

Note: Dotted lines depict 90 percent confidence interval.

Figure 2
Response of Investment to Monetary-Policy Shock
Comparison of VAR and Structural Model Responses

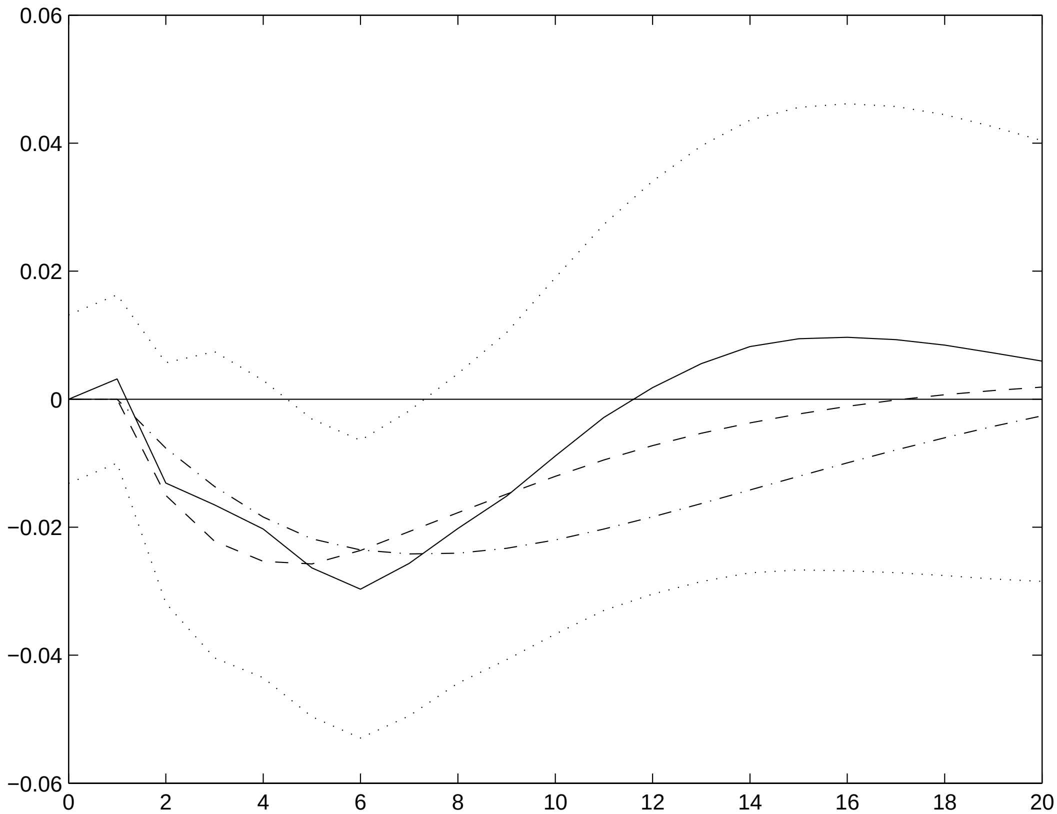

Notes: Responses to unit shocks
Solid: VAR impulse response
Dotted: 90 percent confidence interval around VAR impulse response
Dashed: Allows both investment and capital stock adjustment costs
Dot-dash: Investment adjustment costs only
Both impose $\sigma = 1$.

Figure 3
Response of Investment to Technology and Monetary-Policy Shocks
Comparison of VAR and Structural Model Responses

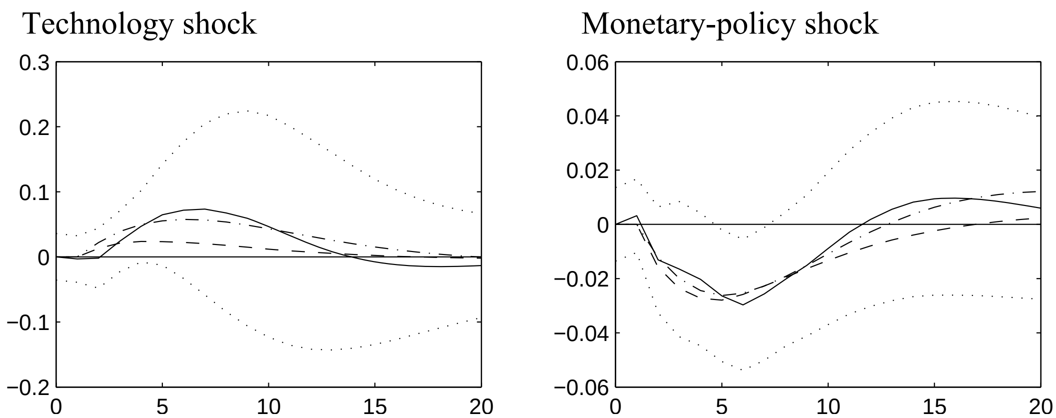

Notes: Responses to unit shocks.
Solid line: VAR impulse response
Dotted: 90 percent confidence interval around VAR impulse response
Dashed: σ = 1 imposed.
Dot-dash: σ is freely estimated
Both allow investment and capital stock adjustment costs

Figure 4
Response of Investment to Technology, Monetary-Policy, and Inflation Shocks
Comparison of VAR and Structural Model Responses

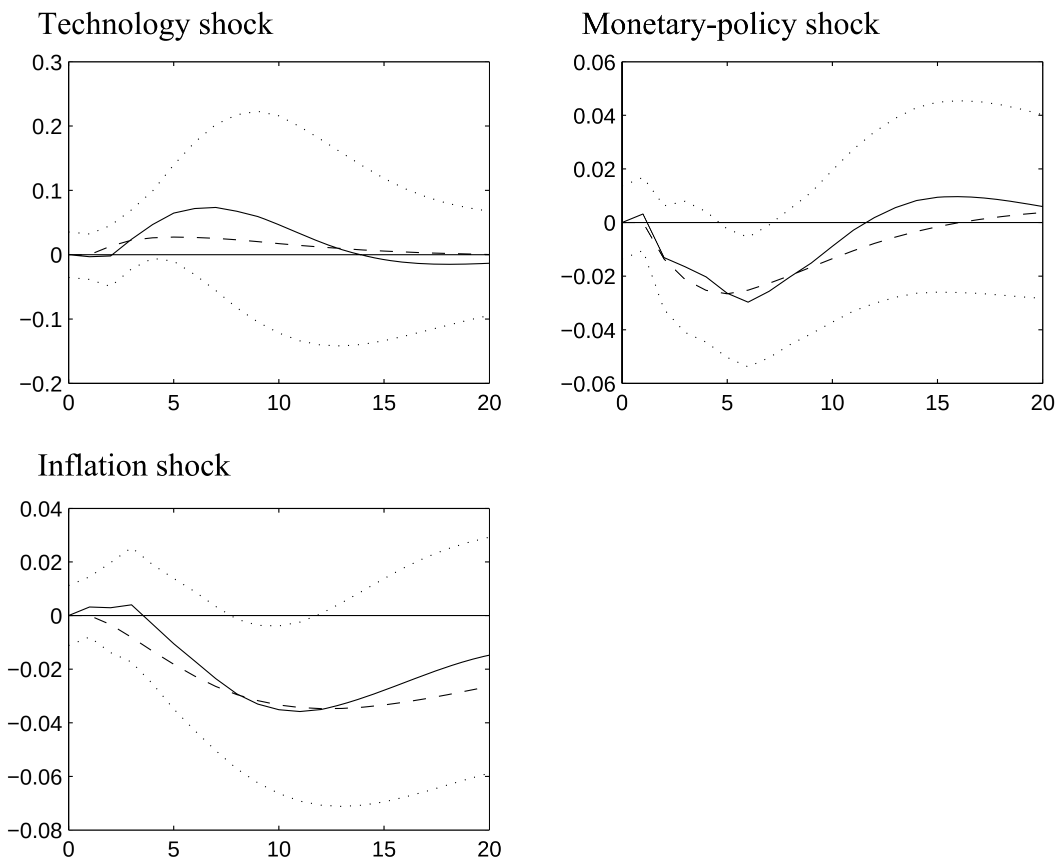

Notes: Responses to unit shocks.
Solid line: VAR impulse response
Dotted: 90 percent confidence interval around VAR impulse response
Dashed: Structural model; allows σ and investment and capital stock adjustment costs to be freely estimated.

Figure 5
Response of Investment to Technology, Monetary-Policy, Inflation, and IS Curve Shocks
Comparison of VAR and Structural Model Responses

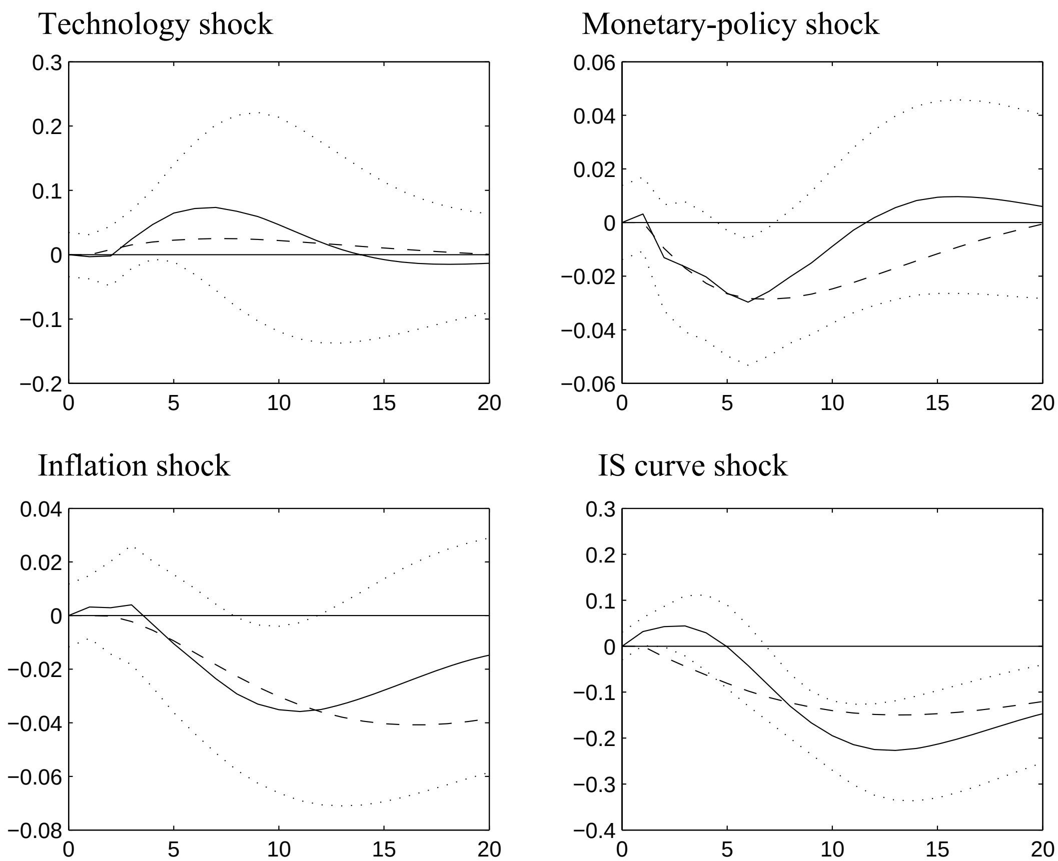

Notes: Responses to unit shocks.
Solid line: VAR impulse response
Dotted: 90 percent confidence interval around VAR impulse response
Dashed: Structural model; $\sigma = 1$, investment adjustment costs only

References

Amato, Jeffery D. and Thomas Laubach (2004) "Implications of Habit Formation for Optimal Monetary Policy," *Journal of Monetary Economics* 51, 305-25.

Chirinko, Robert S. (1993) "Business Fixed Investment Spending: Modeling Strategies, Empirical Results, and Policy Implications," *Journal of Economic Literature* 31, 1875-911.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005) "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy* 113, 1-45.

______, ______, and ______ (1999) "Monetary Policy Shocks: What Have We Learned and to What End?" in John B. Taylor and Michael Woodford, eds.. *Handbook of Macroeconomics, Volume 1a,* Elsevier Science, North-Holland, 65-148.

______, ______, and Robert Vigfusson (2005) “Assessing Structural VARs,” Board of Governors of the Federal Reserve System, manuscript (May).

Cummins, Jason G., Kevin A. Hassett, and Stephen D. Oliner (1999) "Investment Behavior, Observable Expectations, and Internal Funds," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, no. 1999-27.

______ and Giovanni L. Violante (2002) "Investment-Specific Technical Change in the United States (1947-2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics* 5, 243-84.

Del Negro, Marco and Frank Schorfheide (2004) "Priors from General Equilibrium Models for VARs," *International Economic Review* 45, 643-73.

Edge, Rochelle M. (2000) "Time-to-Build, Time-to-Plan, Habit-Persistence, and the Liquidity Effect," Board of Governors of the Federal Reserve System, International Finance Discussion Papers, no. 673.

Gilchrist, Simon and Charles P. Himmelberg (1995) "Evidence on the Role of Cash Flow for Investment," *Journal of Monetary Economics* 36, 541-72.

Lebow, David E. and Jeremy B. Rudd (2003) "Measurement Error in the Consumer Price Index: Where Do We Stand?" *Journal of Economic Literature* 41, 159-201.

Roberts, John M. (2003) "Modeling Aggregate Investment: A Fundamentalist Approach," Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, no. 2003-48.

______ (2001) "Estimates of the Productivity Trend Using Time-Varying Parameter Techniques," *Contributions to Macroeconomics* 1, issue 1.

Rotemberg, Julio J. and Michael Woodford (1997) "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual* 1997, 297-346.

Shapiro, Matthew D. and Mark W. Watson (1989) "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual* 1989, - .

Tevlin, Stacey and Karl Whelan (2003) "Explaining the Investment Boom of the 1990s," *Journal of Money, Credit, and Banking* 35, 1-22.