

Modeling Aggregate Investment: A Fundamentalist Approach

John M. Roberts
Board of Governors of the Federal Reserve System
Stop 61
Washington, D.C. 20551
jroberts@frb.gov

August 2003
Revised version

Abstract: This paper applies some lessons from recent estimation of investment models with firm-level data to the aggregate data with an eye to rehabilitating convex costs of adjusting the capital stock. In recent firm-level work, the response of investment to output and other “fundamental” variables is interpreted in terms of the traditional convex-adjustment-cost model, implying annual capital-stock adjustment speeds on the order of 15 to 35 percent. In aggregate data, I find that this “fundamentalist” model can account for the reduced-form effect of output on investment and the estimated capital-stock adjustment speed is similar to those from firm-level studies – around 25 percent per year. To account for the slower adjustment to changes in the cost of capital, I consider a model in which the capital-intensity of production is also costly to adjust. I find that this model can account for the reduced-form effects of both output and the cost of capital on investment.

JEL classification codes: E22, D92

Keywords: Investment, adjustment costs, putty-clay

I am grateful to Flint Brayton, Chris Carroll, Darrel Cohen, Jason Cummins, Rochelle Edge, Stacey Tevlin, John Williams, and seminar participants at Johns Hopkins University and the Federal Reserve Board for helpful discussions, and to Sarah Alves for her help running the model simulations in section 5. The views expressed in this paper are those of the author and are not to be construed as those of the Board of Governors of the Federal Reserve Board or other members of its staff.

In this paper, I examine whether some recent results from the empirical literature on investment at the firm level might help with the interpretation of aggregate investment data through the prism of convex costs to adjusting the capital stock. Using firm-level data, Gilchrist and Himmelberg (1995, 1998) and Cummins, Hassett, and Oliner (1999) found that business fixed investment could be explained with a model assuming convex costs of adjusting the capital stock and only modest costs of adjusting the capital stock – implying, for broad samples of firms, capital-stock adjustment speeds that range from 14 percent per year to 36 percent per year (Gilchrist and Himmelberg, 1995, 1998; Cummins, Hassett, and Oliner, 1999, obtain an intermediate pace of 27 percent per year).¹ Furthermore, in many cases, these models are able to capture fully the predictable variation in investment at the firm level.

By contrast, studies that attempt to explain *aggregate* business fixed investment using models with convex capital-stock adjustment costs have found implausibly slow adjustment of the capital stock: For example, Summers (1981) reported annual adjustment speeds of 1-1/2 to 4 percent per year, and Abel and Blanchard (1986), about 1-1/2 percent per year. And these studies found that the structural model left much of the predictable movement in investment unexplained.² The inability of the convex-capital-stock-adjustment-cost model to explain movements in aggregate investment has led some builders of empirical macroeconomic models

¹ Appendix A describes the relationship between the adjustment cost estimates presented in these papers and the estimates of capital-stock adjustment speed discussed here.

² Shapiro (1986a, b) found plausible estimates of capital-stock adjustment speeds using limited-information estimation methods. However, the underlying parameter estimates were not statistically significant, and Oliner, Rudebusch, and Sichel (1995) found that even when limited-information estimates of investment models pass standard specification tests, the estimated models could not adequately account for predictable movements in investment.

(Kiley, 2001; Christiano, Eichenbaum, and Evans, 2001) to abandon costly capital-stock adjustment in favor of models in which only investment, and not the capital stock, is costly to adjust.

In the recent firm-level studies, the estimation exploits the response of investment to quantity variables such as firm output and sales. Many of these studies have also examined the response of investment to the firm's share price. When they do so, they obtain results similar to those of Summers and Abel and Blanchard, with capital-stock adjustment speeds on the order of 2 to 5 percent per year. These results are consistent with the conclusions of Chirinko (1993), who characterizes the investment literature as finding much smaller movements of investment in response to changes in the cost of capital than to changes in output. Gilchrist and Himmelberg (1995, 1998) and Cummins, Hassett, and Oliner (1999) argue that the strong response of investment to quantity variables should be considered "fundamental," in contrast to the response of investment to movements in the firm's share price. These authors argue that the response of investment to quantity variables reflects the working of the convex-adjustment-cost mechanism, whereas the response of investment to financial-market information is attenuated by other processes.

I bring the insights of the firm-level work to the explanation of aggregate investment by first using only the information contained in the response of investment to movements in output to inform estimation of the model's parameters. I find that estimates of the speed of adjustment of the capital stock are similar to the estimates obtained in the firm-level studies. In particular, I find that the capital stock adjusts at an annual rate of about 25 percent in response to a

hypothetical permanent output shock. Moreover, the model is able to fit the reduced-form responses of investment to an output shock very closely.³

These “fundamentalist” results demonstrate that the capital-stock-adjustment-cost model can account for the relationship between output and investment in aggregate data. As will be illustrated in the empirical work below, a drawback of the model with these estimates is that it predicts that shocks to the cost of capital will have a much larger effect on investment than is consistent with the reduced-form evidence. To account for the more-sluggish response of investment to innovations in the cost of capital, I consider a model with the additional features (1) that the firm’s desired capital intensity of production may be costly to adjust and (2) that utilization of the capital stock can vary.⁴ Costly adjustment of the desired capital intensity of production means that the model can, in principle, account for the sluggish adjustment of investment in the aftermath of a change in the cost of capital; adding variable capital utilization as well allows output to vary freely despite costs to adjusting the capital intensity of production.

This expanded model captures the responses of investment to both output and the cost of capital that are evident in reduced-form estimates quite closely. At the same time, the estimates of the speed of capital-stock adjustment are similar to the estimates I found using “fundamentals” alone and are thus consistent with the pace found with firm-level data. The other key adjustment in this model – the pace at which the capital intensity of production adjusts – occurs at an annual rate of about 20 percent. It is this additional sluggish adjustment that accounts for the attenuated effect of the cost of capital on investment in this model.

³ As explained in detail in section 2 below, I examine the stock of equipment excluding high-tech categories such as computers, software, and communications equipment for the business sector. This category accounted for about half of overall U.S. business fixed investment in the 1990s.

⁴ Edge and Rudd (2003) were the first to propose convex costs of adjusting the capital intensity of production.

The idea that firms might face costs of adjusting the capital intensity of their production goes back at least to Bischoff (1971), who argued that firms were subject to a “putty-clay” technology in which capital goods, once installed, have a fixed capital-labor ratio. A number of studies pointed to putty-clay technology as a motivation for a slower rate of adjustment of investment to movements in the cost of capital in empirical work on investment, including Bischoff (1971), Clark (1979), Brayton and Mauskopf (1985), and Oliner, Rudebusch, and Sichel (1995). However, these studies (quite explicitly) relied on approximations to the structural model. One reason the link between theory and model was not tighter is that structural analysis of the putty-clay model is computationally demanding, requiring a model in which distinct “vintages” of installed capital – which vary, for example, according to the required labor input and level of embedded technology – must be kept account of.⁵

An advantage of the model of convex capital-intensity adjustment costs is thus its tractability. This greater tractability comes from treating currently installed and newly purchased capital goods similarly: In the putty-clay model, there is a sharp asymmetry in the costs of adjusting the capital-intensity of production between new and existing capital: It is infinitely costly to change the capital-output ratio of installed capital, whereas it is costless to change the capital intensity for new capital. In the convex model, the capital-intensity of existing capital can be changed, if a cost is incurred. However, this cost is just as large for newly installed capital as for existing capital. This difference in the treatment of new and existing capital can be related to different sources of costs of changing the capital-intensity of production: In the original putty-clay model, the main motivation for the rigidity of the capital

⁵ Recent work by Gilchrist and Williams (2000) illustrates the difficulties involved: Although they provide statistical evidence that putty-clay mechanisms may be important, they do not provide empirical estimates of the key parameters of their model.

intensity of production was the notion that the physical properties of, for example, a particular machine would prohibit adjustment of the capital-labor ratio. Such a rigidity would impinge asymmetrically on new and existing capital. But changing the capital intensity of production entails other costs, notably the need to reorganize production processes, and these costs apply equally to new and existing capital.

Investment plays a prominent role in cyclical fluctuations. To assess how this role may be affected by the sluggish adjustment of the capital intensity of production, I include my estimated investment model as part of a “New Keynesian” macroeconomic model. Besides the model of investment just discussed, this macroeconomic model also includes an optimization-based model of consumer behavior; a New Keynesian price adjustment model; and a monetary-policy reaction function. As might perhaps be expected, I find that slowing the speed at which the cost of capital affects investment has important implications for the effects of monetary policy on the economy: Based on the empirical estimates of this paper, the maximum effect of a monetary policy shock on output is reduced by about a third.

* * *

Section 1 presents the formal model, and section 2 discusses empirical implementation. Sections 3 and 4 present empirical estimates, first of the “standard” model without costly adjustment of the capital intensity of production and then adding this additional source of rigidity. Section 5 presents the model simulations, and section 6 concludes.

1. The model

The firm makes its investment decisions conditional on processes for output and the cost of capital, subject to the following production technology: The firm’s “gross output” is produced using a Cobb-Douglas production function:

$$X_t \equiv (U_t K_t)^\alpha (A_t L_t)^{(1-\alpha)}, \quad (1)$$

where K is the capital stock, U is the capital utilization rate, L is labor input, and A is the efficiency of labor. The capital stock is subject to the capital-accumulation identity,

$$K_t = (1-\delta)K_{t-1} + I_t, \quad (2)$$

where I is investment and δ is the depreciation rate of capital.

The net output of the firm, Y , is the output available after a number of adjustment costs have been incurred. These costs include:

- A convex cost to adjusting the capital stock.
- A convex cost to changing the investment rate. Costly adjustment of the investment rate was first suggested by Topel and Rosen (1988). It is currently a part of the Federal Reserve's FRB/US model (Kiley, 2001) and of the empirical macroeconomic model developed by Christiano, Eichenbaum, and Evans (2001).
- A convex cost of changing the capital-output ratio, where the capital stock has been adjusted for utilization, along the lines proposed by Edge and Rudd (2003).
- A convex cost of deviations in the utilization rate of capital from its normal value. Utilization costs of this sort have been used, for example, by Christiano, Eichenbaum, and Evans (2001).

For simplicity, I assume that the firm faces a log-linear-quadratic cost minimization problem, in which output and the cost of capital are taken as given.⁶ In particular, I assume that the firm minimizes the following dynamic cost function:

$$\text{Min}_{\{k, \rho\}} E_t \sum \beta^t \{ (\rho_t + c)^2 + \theta_1 (k_t - k_{t-1})^2 + \theta_2 (\Delta k_t - \Delta k_{t-1})^2 + \theta_3 (\rho_t - \rho_{t-1})^2 + \theta_4 (\rho_t + x_t - k_t)^2 \}. \quad (3)$$

⁶ Appendix B shows that a formulation of the problem in terms of profit maximization leads to similar first-order conditions.

Here, constant terms have been ignored; lower-case letters represent logs of the corresponding upper-case variables; c is the log of the user cost of capital, β is a constant discount factor, and ρ is the log of the desired capital-output ratio – that is, the actual capital-output ratio, adjusted for utilization:

$$\rho_t \equiv u_t + k_t - x_t, \quad (4)$$

The first term in problem 3 is the cost to the firm of a capital-output ratio that differs from the value indicated by the user cost of capital: Given the assumption of a Cobb-Douglas production function, the firm would want to set $\rho_t = -c_t$ in the absence of adjustment costs (recall that constant terms are being ignored). The second term indicates the cost to the firm of changing the level of the capital stock, while the third indicates the cost of changing the firm's investment rate. The term $(\rho_t - \rho_{t-1})^2$ represents the cost to the firm of changing its capital-intensity of production. The final term captures the cost to the firm of capacity utilization that differs from its equilibrium level (recall that $u_t = \rho_t + k_t - x_t$).

The first-order conditions from this problem are:

Capital:
$$E_t [\theta_1 (1 - \beta L^{-1}) \Delta k_t - \theta_4 u_t + \theta_2 (1 - \beta L^{-1})^2 \Delta^2 k_t] = 0 \quad (5)$$

Capital-output ratio:
$$E_t [(\rho_t + c_t) + \theta_3 (1 - \beta L^{-1}) \Delta \rho_t + \theta_4 u_t] = 0 \quad (6)$$

To help interpret the first-order condition for the capital stock, note that in the absence of costs to changing investment ($\theta_2 = 0$), the capital stock would adjust according to:

$$\Delta k_t = \beta E_t \Delta k_{t+1} + (\theta_4/\theta_1) u_t. \quad (7)$$

Thus, in this model, the capital stock adjusts to close the deviation in capital-stock utilization from its long-run value. This relationship can be used to define a variable Δk^* , the rate of investment that would occur in the absence of investment adjustment costs:

$$\Delta k_t^* \equiv E_t (\theta_4/\theta_1) u_t / (1 - \beta L^{-1}). \quad (8)$$

The first-order condition for the capital stock can then be rewritten as:

$$\Delta^2 k_t = \beta E_t \Delta^2 k_{t+1} + (\theta_1/\theta_2) (\Delta k_t^* - \Delta k_t). \quad (9)$$

Written in this way, the first-order condition for the capital stock can be interpreted as capturing the dynamic adjustment of investment toward a target rate of investment, where that target rate of investment, in turn, reflects the adjustment of the capital stock to deviations in the utilization rate from its long-run value.

The first-order condition for the utilization-adjusted capital-output ratio can be rewritten:

$$\Delta \rho_t = \beta E_t \Delta \rho_{t+1} - (1/\theta_3)[(\rho_t + c_t) + \theta_4 u_t] \quad (10)$$

This equation indicates that ρ adjusts gradually to close the gap between ρ_t and $-c_t$ and between the log of the utilization rate and zero. Recalling the definition of ρ in equation 4, note that the choice of ρ also determines u , the (log of the) utilization rate. Thus, the firm's problem can be thought of as choosing both the capital stock and the utilization rate, conditional on output and the cost of capital.

To get a heuristic understanding of how costly adjustment of the capital-intensity of production can help account for the different responses of investment to movements in output and user cost, it is useful to compare these first-order conditions to those from the "standard" model – that is, the model with no cost of adjusting the capital-intensity of production. In the standard model, the parameter $\theta_3 = 0$ and as a result, $u = 0$, $\rho = -c$, and equations 6 and 7 imply:

$$\Delta k_t = \beta E_t \Delta k_{t+1} + (1/\theta_1)(x_t - c_t - k_t). \quad (11)$$

Equation 7 can be written as:

$$\Delta k_t = \beta E_t \Delta k_{t+1} + (\theta_4/\theta_1)(x_t + \rho_t - k_t). \quad (12)$$

As can be seen, in the standard model (equation 11), the role of the cost of capital is similar to that of output. By contrast, in equation 12, the impact of the cost of capital on investment is

mediated through the variable ρ , which is itself subject to adjustment costs. Thus, in the model with costly adjustment of the capital intensity of production, the impact of movements in the cost of capital on investment, can, in principle, be attenuated relative to the impact of output. Of course, this is a broad-brush characterization, and the exact predictions of the models will depend, among other things, on the time-series properties of output and the cost of capital and the exact parameter estimates of the model. But, as we will see, these “heuristic” implications will be borne out.

2. Empirical implementation

2.1 – Data

I estimate the model using data on non-high-tech equipment spending in the U.S. business sector, where the excluded “high-tech” categories include computers, software, and communications equipment. It is sensible to consider equipment separately from structures because the lead time for structures is considerably longer than for equipment. It is appropriate to exclude the “high-tech” categories from equipment investment because, as emphasized by Tevlin and Whelan (2003), the effects of the cost of capital on high-tech investment are very different than on other equipment spending, mostly because user cost for high-tech spending is dominated by a strong downward trend in the price of these capital goods, which is not an important factor for user cost in other areas. Notwithstanding these exclusions, non-high-tech equipment spending was about half of business fixed investment in the 1990's.

My main source of data on output and investment is the U.S. National Income and Product Accounts (NIPAs). As an adjunct to the national accounts, the Bureau of Economic Analysis publishes capital stock data. Unfortunately, these data are published at an annual frequency. To construct quarterly capital-stock data, quarterly investment data were used to

interpolate finely disaggregated capital stocks to the quarterly frequency, and these stocks were then aggregated to form the capital-stock aggregate used in this paper. Kiley (2001) provides more details on how these quarterly data were constructed.

For output, the available data don't quite match the model's concepts: I use published business sector output, whereas the variable that appears in the first-order conditions is x , the measure of output that is gross of the firm's internal adjustment costs. Estimates by Lichtenberg (1988) suggest that the difference between net and gross output is likely to be small, and so this approximation likely has little effect on the results.

The cost of capital is unfortunately not available in the national accounts. Here, I measure the user cost of capital as the sum of the depreciation rate, the real rate of interest as measured by the three-month commercial paper rate less a measure of inflation expectations, and a constant that raises the average level of the cost of capital to the level implicit in the average return on capital. This sum is multiplied by the relative price of non-high-tech investment goods. The additive constant can be thought of as reflecting an additional internal "hurdle rate" that the firm applies before deciding to undertake a capital investment project. The depreciation rate is the time-varying rate consistent with the capital-stock measure. The results are little affected if the time-varying depreciation rate is replaced by its sample average of the rate.

In preliminary work, I also considered a more elaborate measure of the cost of capital that is used in the Federal Reserve's FRB/US model, which also takes into account the effects of taxes and an costs of equity and bond finance. I found that the results were not very sensitive to the choice of the measure of user cost and I therefore focus on the simpler measure in the empirical work presented below. An advantage of using the measure of user cost based on a short-term interest rate that I focus on is that it is closer in spirit to what the model demands: In

this model, the user cost of capital captures how the firm would choose its capital stock in the absence of adjustment costs. In that case, the firm would adjust its capital stock each period in line with the opportunity cost in that period, which is best captured by a short-term interest rate. One disadvantage of the commercial paper rate, however, is that it may not reflect the financing cost for the typical firm because only large firms with good credit ratings participate in the commercial paper market. While the adjustments to user cost that are made to be consistent with the average cost of capital ensure that the measure is correct on average, it may miss some variation in the cost of capital for the typical firm. Because the more-elaborate measure of user cost takes account of the costs of equity and bond finance, which apply to a broader range of firms, it will better capture this additional variation – albeit at the cost of an inappropriate duration. But as already noted, the results were little affected by the choice of user cost measure.

As the discussion in the introduction suggests, in principle, estimation of the model would benefit from inclusion of data on the utilization rate of capital. Unfortunately, appropriate aggregate data on the utilization rate of capital are not available. The closest available proxy is the Federal Reserve's "capacity utilization" measure, which is available for several sectors of the economy, including manufacturing, mining, and utilities. For a number of reasons, capacity utilization is not an ideal measure of the utilization of capital. First of all, it covers all types of capital for a narrow sector of the economy, whereas the ideal measure for this model would be the utilization rate for a specific type of capital (non-high-tech equipment) for the entire business sector. In addition to the problems of coverage, the Fed's capacity utilization measure is not intended as a measure of the utilization of the capital stock. Rather, it is intended as a measure of the utilization of all of the resources of the firm, including, for example, labor.

Because this measure is not ideal, I first examine estimates of the parameters of the model using only data on investment, output, and the cost of capital. Limiting the data set to these series has the advantage of making the estimates of this paper more comparable to most other studies of the determinants of investment, which have not jointly consider the behavior of investment and utilization. I then present estimates of the model using a measure of the utilization of capital that is based on the Fed's measure of manufacturing capacity utilization. The measure makes a partial adjustment for coverage to account for the greater volatility of the manufacturing sector relative to the overall business sector: The standard deviation of quarterly output growth in the business sector is only 0.58 times as large as the standard deviation of the quarterly growth rate of industrial production in the manufacturing sector. As a consequence, the log of the manufacturing capacity utilization rate is multiplied by 0.58 to obtain the measure of the utilization rate of capital.

2.2 – Estimation approach

In the empirical work, I present estimates using both full-information maximum likelihood and an approach that chooses the model parameters to match as closely as possible the impulse responses predicted by the model to output and cost of capital shocks. The latter method is similar to the approach of Christiano, Eichenbaum, and Evans (2001) and Altig, Christiano, Eichenbaum, and Linde (2002): It seeks to minimize a weighted sum of squared differences between the impulse responses from the model and impulse responses from a reduced-form model, such as a VAR, where the weights are the square of the estimated standard error around the impulse responses of the unconstrained model. The minimum-distance-based parameter estimates are asymptotically consistent; they will be less econometrically efficient than the

estimates from a correctly specified model estimated with FIML, but, as noted below, they are less subject to specification bias than are FIML estimates.

An alternative approach to the estimation of forward-looking models is the generalized method of moments (GMM). GMM was used by Shapiro (1986a, b) and Oliner, Rudebusch, and Sichel (ORS; 1995). Fuhrer, Moore, and Schuh (1995) have examined the relative merits of FIML and GMM in the context of forward-looking adjustment-cost models such as the present one. They found that GMM estimates can have poor small-sample properties and FIML-based estimates are typically superior. As noted in the introduction, ORS – looking specifically at investment – found that even when GMM estimates pass standard diagnostic tests, forecasts based on the structural GMM estimates are inferior to those from reduced-form models. One advantage of the moment-matching approach is that it directly addresses the forecasting concern raised by ORS by choosing the parameters specifically to match the ability of a reduced-form model to fit the responses of investment to output and cost of capital shocks. Hence, the estimates from the moment-matching approach should perform better than GMM estimates by this criterion.

The textbook critique of FIML is that it depends on correct specification of the entire model, and so mis-specification in any part of the model – even in an area of secondary concern – can lead to biased estimates throughout the model. One way to avoid this problem is to use a relatively unconstrained model to characterize the evolution of the parts of the model that are not of central concern, as in Fuhrer (2000). Here, I use a bivariate VAR (with only a few restrictions) to specify the process for output and user cost and thus impose little prior structure on this part of the model. I use this “core” VAR to determine the evolution of output and user cost in both the FIML-based and moment-matching approaches. An additional advantage of the

moment-matching approach, however, is that it is not necessary to complete the model by, for example, specifying a process for shocks to investment. The moment-matching approach allows consistent estimation without having to give a (possibly incorrect) structural interpretation to this shock.

Another advantage of the moment-matching approach is that it allows comparison with the “fundamentalist” firm-level studies – Cummins, Hassett, and Oliner (1999) and Gilchrist and Himmelberg (1995, 1998). In particular, with the moment-matching approach, it is possible to look at the response of investment to output only, which is similar to the approach in the firm-level studies in which the response of investment to “fundamentals” is used to estimate the structural parameters of the model.

2.3 – Specification of the reduced-form model

The reduced-form model that will be the basis of comparison in the moment-matching estimation procedure consists of two parts. The first is the “core” bivariate VAR in output growth and (the log of) the user cost of capital, with four lags of each variable; the second is a reduced-form model that relates investment to lags of itself, output, and user cost. As noted in the previous section, the core VAR is also the specification for output and the cost of capital in the FIML estimation. In the core VAR, output is differenced because it clearly has a unit root. User cost is not differenced, because while movements in the cost of capital are highly persistent, the evidence that it may contain a unit root process is not as clear-cut as for output, and by leaving user cost in level form, the data are allowed to determine the persistence of the process. The residuals from the two equations were virtually uncorrelated (correlation coefficient of -0.03) and so I assume zero contemporaneous correlation. In addition, I impose

the restriction that only the shock to the output equation is allowed to have a permanent effect on output; this restriction is not violated by the data.

Figure 1 shows the impulse responses in the core bivariate VAR. In the VAR, the output shock has a large permanent effect on output, with the ultimate effect of the shock about one-and-a-half times larger than the initial effect. The output shock has only a small, transitory effect on the cost of capital, with the cost of capital significantly higher than the baseline a few quarters after the initial shock, but then returning to baseline shortly thereafter. The user cost shock has a persistent, but transitory, effect on user cost, with user cost returning about half-way to its initial value about six quarters after the initial shock. It has a depressing effect on output that is (barely) statistically significant after two quarters; the effect becomes a bit larger in subsequent quarters before fading away gradually. The depressing effect of user cost on output might occur if, for example, monetary policy shocks have been an important source of variation in user cost.

The second part of the reduced-form model used in the minimum-distance estimation consists of a regression of the change in the log of the capital stock on lagged values of itself, output, and user cost.⁷ I impose several restrictions that are suggested by the theory, leaving the lagged coefficients of the model otherwise unconstrained. In particular, I assume that, of the two shocks from the output-and-user-cost VAR, only the shock that has long-run effects on output can have long-run effects on the capital stock. I further restrict the long-run effect on the capital stock to be proportional to the effect on output. In addition, I assume that neither of the shocks

⁷ Relative to a standard VAR in output, the cost of capital, and investment, the present approach eliminates any feedback from investment to output and the cost of capital. An advantage of this approach is that it is more transparent than the approach with feedback, because the focus is on effects only in one direction. This advantage likely comes at little cost, as the effects of the investment shock on output and the cost of capital in an unconstrained three-variable VAR are economically small.

has any contemporaneous effect on investment. This assumption is consistent with a one-quarter “planning lag” before equipment can be purchased. The model is thus overidentified; however, conditional on the current-period zero restrictions, the long-run restrictions are not rejected, suggesting that the overidentifying restrictions are not rejected. Furthermore, because the restrictions that are imposed on this model are also imposed on the structural model, only the implications of the model for the dynamics of investment in response to output and user-cost shocks are tested. As with the core VAR, the reduced-form investment model includes four lags of each variable.

Figure 2 shows the effects of the output and user-cost shocks on investment in the reduced-form model. These will be the “moments” that the moment-matching approach will attempt to match. As in Altig, Christiano, Eichenbaum, and Linde (2002), the goal will be to match the impulse responses over the first twenty periods. This choice seems reasonable since the first five years will cover most of the “business cycle” variation that is of interest.

The effects of output and the cost of capital on investment are as might be expected: The output shock leads to an increase in investment, whereas the user-cost shock depresses investment. In both cases, the effects mount gradually, consistent with the notion that it is costly for the firm to adjust investment in response to shocks. The effect of the output shock on investment is strongly statistically significant. The effect of the user-cost shock is also statistically significant, a result that raises the prospect that we may be able to obtain precise estimates of the effects of the cost of capital on investment.

In addition to the output and user-cost responses that the method will attempt to match, the top panel of figure 2 also shows the effects of a shock to the investment equation itself. While the responses of investment to output and the cost of capital gradually mount before

reaching a peak after several quarters, the peak impact of the investment shock is immediate. The effect then fades gradually. One specification that could capture this pattern would be to add a serially correlated error term to either equation 8 or 9. An error term added to equation 8 would have the interpretation of being a transitory shock to the firm's desired level of investment. I will make such an assumption in my FIML estimation, allowing the error term to follow a first-order autoregressive process.

3. Results for the standard model

This section presents estimates of the “standard” model – that is, the model with no cost of adjusting the capital-intensity of production. Recall that in this case, the parameter $\theta_3 = 0$ and as a result, $u = 0$, $\rho = c$, and the model reduces to:

$$\Delta k_t^* = \beta E_t \Delta k_{t+1}^* + \gamma_1 (x_t - c_t - k_t) \quad (13)$$

and
$$\Delta^2 k_t = \beta E_t \Delta^2 k_{t+1} + \gamma_2 (\Delta k_t^* - \Delta k_t), \quad (14)$$

where $\gamma_1 \equiv (1/\theta_1)$ and $\gamma_2 \equiv (\theta_1/\theta_2)$. In the estimation, I assume there is a one-period planning lag, so that investment decisions are made on the basis of period $t-1$ information.

Recall that in the minimum-distance approach, the parameters are chosen to best fit the impulse responses of investment to output and user-cost shocks that are shown in figure 2. To illustrate the different implications of fitting these two impulse responses, I first fit the model to each impulse response separately before turning to joint estimation. As noted in the introduction, an advantage of this approach is that it allows comparison with firm-level studies that estimate capital-stock adjustment costs based on the response of investment to “fundamentals.” As in that literature, these estimates focus on how investment responds to its major nonfinancial determinants.

The first column of table 1 shows estimates when the parameters are chosen to match as closely as possible the effect of the output shock on investment over the first twenty periods following the shock. As can be seen, both γ_1 and γ_2 are precisely estimated, with t -ratios of 6.4 and 3.0, respectively. The top panel of figure 3 shows the model's predicted impact of an output shock on investment along with the reduced-form impulse response of investment to an output shock. With these parameters, the structural model can fit the response of investment to an output shock very closely. The bottom panel of figure 3 shows the implications of this set of parameters for the response of investment to a cost-of-capital shock. Here, the fit is much poorer: The structural model implies a much larger response than the reduced-form model after the increase in user cost and in the initial six quarters, the response lies well outside the 90 percent confidence interval around the reduced-form impulse response. Hence, while these parameters lead to a very tight fit for the effect of *output* on investment, they imply a response of investment to user cost that is much faster than is seen in the reduced-form impulse response.

These conclusions are also clear from the squared distance parameters reported in table 1: The sum of squared distances, normalized by the reduced-form standard errors, is quite small for the response of investment to output and very large for the response to a user-cost shock. If the reduced-form impulse responses were statistically independent, these "minimum-distance statistics" for the output and user-cost impulse responses would be distributed χ^2 , with eighteen degrees of freedom in each case. The differences between the responses to an output shock would then not be statistically significant, whereas the differences between the responses to a user cost shock would be strongly statistically significant. However, the reduced-form impulse responses are almost certainly correlated, and so the χ^2 distribution provides only a rough guide

to the joint statistical significance of the deviations of the model from the reduced-form impulse responses.

The parameter estimates in column 1 imply a *capital-stock* adjustment speed of 25 percent at an annual rate and an *investment* adjustment speed that implies that, after four quarters, investment adjusts four-fifths of the way to its unconstrained level.⁸ The capital-stock adjustment speed is very similar to estimates from the firm-level research cited earlier. Hence, when adjustment speeds are estimated with respect to the response of investment to movements in “the fundamentals,” there appears to be little tension between estimates based on firm-level and aggregate data. An implication of this finding is that, at least between the level of the firm and that of the aggregate economy, aggregation does not appear to affect estimates of investment dynamics.

I next turn to choosing the parameters to fit the reduced-form response of investment to a cost of capital shock. In this case, it was not possible to obtain precise estimates of the parameters γ_1 and γ_2 separately; unconstrained, the parameter estimates would imply that there were very large costs of changing *investment* – and thus a very small value of γ_2 – and very small costs of changing the capital *stock* – and thus a very large value of γ_1 . Columns 2 and 3 illustrate the point by presenting estimates with the parameter γ_1 constrained to two specific values: In column 2, γ_1 is constrained to be 0.0064, the same value as was estimated in column 1, whereas in column 3, γ_1 is allowed to take on the larger value of 0.02, a value that implies a very rapid pace of capital-stock adjustment – 40 percent at an annual rate. In both

⁸ The root λ that determines the adjustment speed is related to γ through the formula:

$$\lambda = \{(1 + \gamma + \beta) - [(1 + \gamma + \beta)^2 - 4\beta]^{0.5}\} / (2\beta).$$

The annualized adjustment speed is equal to $1 - \lambda^4$.

cases, the parameter estimates imply much slower adjustment of investment than was the case in column 1.

As the minimum-distance statistics indicate, the ability of the model to fit the reduced-form impact of user-cost on investment is greater in column 3 than in column 2. And as shown in the bottom panel of figure 4, the model in column 3 matches the first twenty reduced-form impulse responses to a user-cost shock very closely with these parameters. However, as can be seen in the top panel, the adjustment is too sluggish to capture the response of investment to output, and for the first six quarters after the shock, the model-based response of investment to output lies below the 90 percent confidence interval.

These results suggest that the response of investment to the cost of capital does not provide sufficient econometric power to identify more than one parameter. This lack of power is a problem that will return when the model is extended to include the parameters of the sticky capital-intensity model.⁹

Column 4 presents parameter estimates that are chosen to fit as closely as possible the responses of investment to *both* output and cost-of-capital shocks. The two adjustment-cost parameters are now individually strongly statistically significant. The estimated parameters indicate a capital-stock adjustment speed of 24 percent at an annual rate and an investment adjustment speed that implies that investment adjusts about half way to its desired level in a year. As the “distance” statistics indicate, this pair of parameters doesn’t fit the individual

⁹ These results may also explain why Christiano, Eichenbaum and Evans (2001) adopt a model with costly adjustment of investment but no capital-stock adjustment cost: CEE estimate their model using only the responses of investment (and other variables) to monetary-policy shocks, which would presumably have their largest effect through the channel of the user cost of capital.

impulse responses to output or the cost of capital as well as the models in columns 1 or 3, respectively, but it fits the sum better than either.

The implications of these parameters for the impulse responses to output and user-cost shocks is shown in figure 5. As might be expected, the resulting impulse responses are a compromise: They are too slow to match the output responses and too fast to match the user-cost responses, and the predicted impulse responses lie on or outside the 90 percent confidence bands in several quarters. To the extent that the minimum distance criterion can be thought of as having a χ^2 distribution, it would imply rejection of the model at the 10 percent, but not at the 5 percent, confidence level.

Column 5 shows the results of FIML estimation, under the assumption that target investment is subject to a first-order autoregressive shock. These estimates suggest a somewhat more-sluggish pace of capital-stock adjustment, although the estimate remains within the range suggested by firm-level studies. Not surprisingly, by the minimum-distance criterion, these impulse responses are not as close to the reduced-form responses as are those based on the estimates in column 4. As might be expected based on figure 2, the investment shock is estimated to be highly serially correlated, with an autoregressive parameter of 0.96.

Tevlin and Whelan (2003) argue that one reason investment may respond less to shocks to the cost of capital than to shocks to output is that movements in output that follow an initial shock are more persistent than the movements in the cost of capital following a shock to that variable. The results presented in this section, however, suggest that the different time-series properties of output and the cost of capital provide only a partial explanation for the smaller response of investment to the cost of capital. For example, in figure 5 – where the estimates are a compromise between the response of investment to output and user-cost impulse responses –

the predicted response of investment to output shocks is both too slow and too small relative to the response in the reduced-form model, whereas the predicted response of investment to the cost of capital shocks is too fast and too large. But these predicted responses already take account of the different time-series properties of output and user cost shocks. Hence, time-series properties alone cannot fully account for the failure of the standard model to account for the responses of investment to both kinds of shocks.

4. Results for model with costly adjustment of the capital-intensity of production

This section presents estimates of the model represented by equations 8, 9, and 10, which includes costly adjustment of the capital-intensity of production. The model is reproduced for convenience:

$$\Delta k_t^* = \beta E_t \Delta k_{t+1}^* + \gamma_1 \gamma_4 u_t, \quad (15)$$

$$\Delta^2 k_t = \beta E_t \Delta^2 k_{t+1} + \gamma_2 (\Delta k_t^* - \Delta k_t), \quad (16)$$

and
$$\Delta \rho_t = \beta E_t \Delta \rho_{t+1} - \gamma_3 [(\rho_t + c_t) + \gamma_4 u_t], \quad (17)$$

where $\gamma_1 \equiv (1/\theta_1)$ and $\gamma_2 \equiv (\theta_1/\theta_2)$, as before, and $\gamma_3 \equiv (1/\theta_3)$ and $\gamma_4 \equiv \theta_4$. Recalling the identity,

$$\rho_t \equiv u_t + k_t - x_t, \quad (18)$$

it is clear that this model jointly determines both investment and the utilization rate of capital. In the estimation that follows, I first estimate the model using only investment data (and not the utilization rate). I then turn to estimation using both investment and the proxy for the utilization rate of capital described in section 2. As noted earlier, the first set of results is of interest because most other work on investment does not also take into account the joint relationship with utilization, and so it is useful to see how well the present model can be estimated using similar information.

The results in section 3 suggested that the moment-matching approach may face a problem of econometric power: The information from the response of investment to user cost appeared to allow precise estimation of only one parameter. Hence, the introduction of two additional parameters, as in the model with costly adjustment of the capital intensity of production, may be problematic: Once again, it was necessary to constrain one of the four parameters of the model in implementing the moment-matching method. In the following, I will examine the implications of some alternative parameter settings.

In column 1 of table 2, the parameter γ_2 is set to its value in column 4 of table 1, where the parameters were estimated using the responses of investment to both output and user cost shocks. This choice of γ_2 implies that investment adjusts half-way to its target in a year. A motivation for this setting is that it takes from the previous estimates only a measure of the high-frequency dynamics of investment and thus asks to what extent the lower-frequency dynamics can be refined by the additional aspects of the extended model.

With this setting for γ_2 , both of the other underlying adjustment-cost parameters, γ_1 and γ_3 , are highly statistically significant. The estimate of γ_1 implies that the “fundamental” pace of capital-stock adjustment is quite rapid, at 37 percent at an annual rate. The capital-output ratio is estimated to adjust toward its equilibrium level at a 21 percent annual rate, a pace that is slower than the rate of adjustment of the capital stock itself. The parameter γ_4 , which captures the differential between the impact of the cost of capital and utilization on the adjustment of the rate of capital intensity, is estimated to be 1.5, but without a great deal of precision.

The capital-stock adjustment speed estimated in column 1 is on the high side of estimates using firm-level data. In column 2, I constrain γ_1 to 0.0064, which implies a capital-stock adjustment speed of 25 percent at an annual rate, a pace that is more consistent with estimates

from firm-level studies as well as with the estimates in columns 1 and 4 of table 1. Under this constraint, the implicit investment adjustment speed is slightly faster than in column 1, while the adjustment speed of the capital-intensity of production is about the same. The estimate of γ_4 is considerably larger than in column 1, and implies that utilization puts about four-and-a-half times more pressure on adjustment of the capital-output ratio than does the user-cost gap. On balance, however, the fit of the model in column 1 is somewhat better.

How much does costly adjustment of the capital-intensity of production improve the fit of the model? Figure 6 shows the impulse responses of the model using parameters from column 1. Relative to the standard model in figure 5, this model comes closer to the reduced-form impulse responses, with the impulse responses from the model no longer breaching the 90 percent confidence bands around the reduced-form impulse responses in the first twenty quarters following the shock. Based on comparison of the minimum-distance criteria, the improvement in fit is statistically significant, with the addition of two degrees of freedom reducing the criterion from 52.0 to 19.5: If this difference were distributed χ^2 with two degrees of freedom, it would be strongly statistically significant, although it bears repeating that this test may not be appropriate, since it does not account for possible correlation among the impulse responses.

Column 3 presents FIML estimates of the model. As before, I add a first-order autoregressive error to the investment equation. In this case, all of the parameters of the model can be estimated without constraints. As with the previous FIML estimates, the capital-stock adjustment speed is on the sluggish side of the minimum-distance estimates, but it nonetheless implies a 21 percent annual pace of adjustment. The investment adjustment speed is actually a bit faster than in the minimum-distance estimates in columns 1 and 2, and implies that

investment adjusts about two-thirds of the way to the desired level in a year. The capital-intensity adjustment speed is similar to, but slightly on the slow side of, the estimates using the minimum-distance approach. As in the table 1 FIML estimate, the autoregressive term is estimated to be 0.96.

As is often the case with nonlinear FIML estimation, the numerical estimates of the parameter standard errors are sometimes too small to be credible, notably for γ_1 and γ_4 . One advantage of the FIML approach, however, is that we can compare the likelihood of this model with that of the standard model estimates in table 1. As can be seen, the likelihood increases by eight, implying that the parameters γ_3 and γ_4 are jointly highly statistically significant. Thus, it appears that the increase in the fit of the model that is permitted by costly adjustment of the capital-intensity of production is large.

In columns 4 and 5, I estimate the model using information from the proxy for the utilization rate of capital that was described in section 2. In deriving the minimum-distance estimates in column 4, the basis of comparison were the impulse responses from a model that includes the core VAR as well as reduced-form models for investment and the capital utilization rate that include four lags of output growth, user cost, capital-stock growth, and utilization. The reduced-form capital utilization equation also included contemporaneous output growth, on the grounds that the firm can vary utilization instantaneously; other contemporaneous correlations were small. As before, the minimum-distance estimator chooses the parameters so as to minimize the weighted squared deviations between the model impulse responses and the reduced-form impulse responses; in this case, the responses of utilization to output and user-cost shocks are also part of the minimum-distance objective.

As can be seen in column 4, with the additional information from utilization, all of the parameters of the model are precisely estimated. The estimated parameters imply reasonable adjustment speeds: The capital stock adjusts at a 29 percent annual rate, the capital-output ratio adjusts at about a 20 percent annual rate toward the rate indicated by the cost of capital, and investment adjusts about three-quarters of the way to its desired level in a year. This model doesn't do quite as well at fitting the reduced-form investment impulse responses as the model in column 1, but as can be seen in figure 7, the fit isn't bad, and it continues to be better than that of the model in figure 5.

Column 5 presents FIML estimates of the model. In the FIML estimation, first-order autoregressive error processes are added to equations 15 and 17. Most of the parameters are in the range seen in earlier estimates, with the exception of the adjustment speed of the capital-output ratio, which, at 11 percent at an annual rate, is the smallest of the estimates. Once again, the numerical standard error estimates are implausible for γ_1 and γ_4 . By the criterion of fitting the investment impulse responses, the fit is about as good as that in column 3. The serial persistence of the shock to the investment equation is once again estimated to be 0.96; however, the serial correlation of the error to the capital-output ratio is much smaller, at only 0.14.

Overall, the results suggest, first, that if "the fundamentals" are allowed to manifest themselves, either because they are the sole piece of information used to estimate the model, as in column 1 of table 1, or because investment is allowed to respond at different rates to output and user cost shocks as in table 2, the estimates suggest that the capital stock adjusts rapidly, and in line with the results from studies with firm-level data. The results also suggest that introducing costly adjustment of the capital-intensity of production significantly improves the ability of the model to fit the responses of investment to both output and user cost shocks.

5. Simulations

The estimation in the preceding sections has been conditional on a reduced-form specification for portions of the model that are not central to the main topic of concern in this paper. In this section, I switch gears and assess the importance of the sluggish adjustment of the capital-intensity of production for the evolution of the economy in response to structural shocks. For this purpose, I embed my investment model in a standard structural macroeconomic model.

Two structural shocks that have long been of interest in macroeconomics are shocks to fiscal and to monetary policy. To explore the implications of these shocks, I complete the model by adding a production function, a model of consumer behavior, a model of “sticky” price adjustment, a monetary-policy reaction function, and a simple fiscal-policy reaction function. The model is described in detail in appendix C. Briefly, the model of consumer behavior assumes that most spending is consistent with intertemporal optimizers but that some spending is directly related to income each period. The behavior of the optimizers includes habit persistence. The model of price adjustment is a New Keynesian specification that assumes “sticky inflation” (see Roberts, 1997, for a discussion of sticky inflation). I assume that a central bank sets short-term interest rates in response to current observations on output and inflation, with coefficients chosen to be consistent with the post-1987 U.S. experience. For fiscal policy, I assume that government spending is a fixed share of (trend) output and that there are lump-sum taxes that are adjusted in order to gradually return the ratio of government debt to output to a target level. As discussed in the appendix, most of the model coefficients are based on empirical estimates from other research.

Figure 8 shows the effects of a 100 basis-point downward shock to the monetary-policy reaction function. The figure compares the reactions in the model with costly adjustment of the

capital-output ratio – where the adjustment speed is 20 percent at an annual rate, a typical estimate from table 2 – with the reactions in the standard model where $\theta_3 = 0$ – that is, in which firms adjust to cost-of-capital shocks at the same rate as output shocks. In both cases, the capital-stock and investment adjustment speeds are assumed to be 25 and 50 percent at an annual rate, respectively.

In both models, the effects of the cut in interest rates on output and inflation are “conventional,” with the cut in interest rates leading to a small, temporary increase in output and inflation. Also in both models, because of costly adjustment of investment, the peak impact on investment is delayed. Output has returned to near its base level after about two years in both models, largely because the monetary policy shock does not have a very persistent effect on interest rates.

Although there are broad similarities in the predictions of the two models, there are also some noticeable differences. In particular, in the model with sticky capital intensity, the peak impact of the shock on investment is considerably smaller and is delayed by several quarters relative to the model without this feature. By contrast, the initial movements in consumer spending are very similar in the two models. Because of the smaller increase in investment, the peak increase in output is about a third smaller in the model with sticky capital intensity.

To explore the implications of shocks that imply more persistent movements in the cost of capital, in figure 9, I examine the effects of a simple fiscal policy shock. Here, the shock is a temporary but persistent tax cut, with no change in government spending. In this model, tax cuts have important effects on interest rates because some households are assumed to consume all of their income each period. Also, because taxes in this model are lump sum, the shock has no

first-order implications for the cost of capital. As a consequence, the shock can be thought of as a generic aggregate demand shock.

As can be seen in the figure, the tax cut has traditionally “Keynesian” consequences in this model: Consumer spending rises, as the “liquidity-constrained” households spend their higher after-tax incomes. Given the assumption of sticky prices, this higher spending leads to greater output, and interest rates are higher as the central bank raises rates to offset the increase in output. With the increase in interest rates, investment falls, leading to the conventional prediction of “crowding out” of capital investment in response to deficit-increasing changes in fiscal policy. Indeed, the increase in interest rates is sharp enough that output is below its baseline level in both models within six quarters following the shock. Output remains below its baseline level thereafter largely because the capital stock is smaller.

The persistence in fiscal policy leads to an increase in interest rates that is also persistent – more so than in the case of the monetary-policy shock. As a consequence, the peak impact on investment is much greater (in absolute value) than in response to the increase in interest rates following the monetary-policy shock, even though the increase in (short-term) interest rates is smaller.

As was the case for the monetary-policy shock, the impact of the increase in the cost of capital on investment is much greater in the standard model than in the model with sticky capital intensity, even though the cost of capital rises by more in the latter model. However, because the increase in interest rates is more persistent in this case, the difference is not quite as great as in the monetary-policy case. Nonetheless, the movements in overall output are about the same in the two models, in large part because monetary policy acts quickly to keep output near its baseline level.

6. Conclusions

The results of this paper demonstrate that the “fundamentalist” findings of several studies of investment using firm-level data carry over to aggregate investment data. In particular, the results indicate that the stock of (non-high-tech) equipment adjusts to fundamental shocks to the desired stock at an annual rate of about 25 percent. Furthermore, the model is able to match closely the reduced-form effect of an output shock on investment. When the standard model is generalized to permit costly adjustment of the capital intensity of production, it can match the reduced-form effects of both output and the cost of capital on investment.

The results of this paper go a long way toward addressing the criticism, leveled by Oliner, Rudebusch, and Sichel (1995), that structural investment models cannot match the predictions of reduced-form forecasting models: Once the structural model has been generalized to allow the capital stock to respond at different speeds to output and cost of capital innovations, it can fit reduced-form impulse responses quite well. Simulation results also suggest that costly adjustment of the capital intensity of production can have important business-cycle implications, with the model that includes such costs predicting a considerably smaller effect of a monetary-policy shock on output than the effect the standard model predicts.

Appendix A – What do firm-level studies suggest for the adjustment speed of the capital stock?

In the introduction, I cited the implications of several firm-level studies for the adjustment speed of the capital stock (Gilchrist and Himmelberg, GH, 1995, 1998; Cummins, Hassett, and Oliner, CHO, 1999). The results these studies report are for equations of the form:

$$(I/K)_t = qcoef Q_t + \text{an error term}, \quad (\text{A.1})$$

where I/K is the ratio of (gross) investment to the capital stock and Q is an estimate of the present discounted value of expected future marginal product of capital. I infer estimates of the adjustment speed of the capital stock from the following relationship between the structural adjustment cost parameters $qcoef$ and γ , taken from Kiyotaki and West (1996, p. 296):

$$\gamma = (r + \delta) qcoef, \quad (\text{A.2})$$

where δ is the rate of depreciation of the capital stock and r is the (average) cost of funds from financial markets. With an estimate of γ , the adjustment speed can be inferred using the formula,

$$\lambda = \{(1 + \gamma + \beta) - [(1 + \gamma + \beta)^2 - 4\beta]^{0.5}\} / (2\beta), \quad (\text{A.3})$$

where $\beta = 1 + r$. The annualized adjustment speed is equal to $1 - \lambda^4$. Thus, conditional on assumptions about r and δ , the estimates of $qcoef$ from the firm-level studies can be used to infer

Table A.1 Implications of Fundamentals-Based Investment Models for the Capital-Stock Adjustment Rate				
<i>Study</i>	<i>qcoef</i>	<i>r</i>	<i>δ</i>	<i>Adjustment speed (1-λ⁴)</i>
GH (1995, table 1, line 1)	.18	.08	.10	.14
GH (1998, table 5, column 3)	1.3	.08	.10	.36
CHO (1999, table 5, column 2) ^a	.52	.10	.15	.27

Notes:

a. Middle panel.

the speed of adjustment of the capital stock. Table A.1 combines estimates of *qcoef* with estimates of r and δ presented in these studies to obtain estimates of adjustment speeds.¹⁰

Appendix B – A profit-maximizing approach

B.1 – The firm’s problem

In section 1, I considered a log-linear cost-minimization problem for the firm and derived the resulting first-order conditions for the capital stock and the capital-output ratio. In this appendix, I consider the profit-maximization problem for the firm and compare the log-linearized first-order conditions with those from the cost-minimization problem in the text. Here, the firm is assumed to maximize profits subject to a number of constraints, as indicated in the following equations:

$$\text{Max}_{\{L,K,I,U,\rho\}} \sum \beta_t \{Y_t - w_t L_t - I_t - \lambda_{1t} [K_t - (1-\delta)K_{t-1} - I_t] - \lambda_{2t} [\kappa_t - (U_t K_t)/(A_t L_t)]\}, \quad (\text{B.1})$$

where,

$$Y_t = (U_t K_t)^\theta (A_t L_t)^{(1-\theta)} - (\phi_1/2)(I_t/K_{t-1} - \xi)^2 K_{t-1} - (\phi_2/2)(U_t - 1)^2 K_{t-1} - (\phi_3/2)(\kappa_t - \kappa_{t-1})^2 K_{t-1} \quad (\text{B.2})$$

Here, w is the wage; as before, Y is output, L is labor input, and I is investment, which is assumed to have the same price as output. The term multiplying λ_1 is the capital accumulation identity; K is the capital stock and δ is the depreciation rate. The term multiplying λ_2 is the identity for the adjusted capital-labor ratio, κ , where the capital-labor ratio is adjusted for the utilization of capital – U – and for the efficiency of labor – A . β_t is a (time-varying) discount factor that reflects the firm’s financing costs. As before, “gross” production is Cobb-Douglas, but various adjustment costs reduce the amount of product that is available for sale outside the firm, including a cost of changing the capital stock; a cost when capacity utilization deviates from one; and a cost to changing the adjusted capital-labor ratio, ρ .

¹⁰ Each of these studies includes extensive sensitivity analysis and so there are multiple estimates of *qcoef* from which to choose. I selected estimates that (a) were based on a broad sample of firms and (b) appeared to be consistent with the “bottom-line” conclusions that the authors emphasized.

B.2 – The first-order conditions

The first-order conditions for the firm's problem are shown in the box below.

First-Order Conditions	
Labor:	$(1-\theta)(U_t K_t / L_t)^\theta A_t^{(1-\theta)} - w_t - \lambda_{2t} [(U_t K_t)/(A_t L_t^2)] = 0$
Investment:	$-\phi_1 (I_t / K_{t-1} - \xi) - I + \lambda_{1t} = 0$
Utilization:	$\theta K_t^\theta (A_t L_t / U_t)^{(1-\theta)} - \phi_2 (U_t - 1) K_{t-1} + \lambda_{2t} [K_t / (A_t L_t)] = 0$
Capital:	$\beta_t \{ \theta U_t^\theta (A_t L_t / K_t)^{(1-\theta)} - \lambda_{1t} + [U_t / (A_t L_t)] \lambda_{2t} \}$ $+ \beta_{t+1} \{ \phi_1 (I_{t+1} / K_t - \xi) (I_{t+1} / K_t) + (1-\delta) \lambda_{1t+1}$ $- (\phi_1 / 2) (I_{t+1} / K_t - \xi)^2 - (\phi_2 / 2) (U_{t+1} - 1)^2 - (\phi_3 / 2) (\kappa_{t+1} - \kappa_t)^2 \} = 0$
Capital-labor ratio:	$\beta_t \{ \phi_3 (\kappa_t - \kappa_{t-1}) K_{t-1} - \lambda_{2t} \} - \beta_{t+1} \{ \phi_3 (\kappa_{t+1} - \kappa_t) K_t \} = 0$

The first-order conditions for utilization and the capital-labor ratio can be combined as:

$$\Delta \kappa_t = (\beta_{t+1} / \beta_t) (K_t / K_{t-1}) \Delta \kappa_{t+1} + (1/\phi_3) (A_t L_t / K_t) [\phi_2 (U_t - 1) - \theta \kappa_t^{(\theta-1)} (K_t / K_{t-1})] \quad (\text{B.3})$$

And the conditions for investment and utilization can be inserted into the first-order condition for capital to give:

$$\begin{aligned} & (I_t / K_{t-1} - \xi) - (\beta_{t+1} / \beta_t) (1 - \delta + I_{t+1} / K_t) (I_{t+1} / K_t - \xi) \\ & = (1/\phi_1) \{ [(1-\delta)(\beta_{t+1} / \beta_t) - 1 + U_t \phi_2 (U_t - 1) (K_{t-1} / K_t)] \\ & \quad - (\beta_{t+1} / \beta_t) [(\phi_1 / 2) (I_{t+1} / K_t - \xi)^2 + (\phi_2 / 2) (U_{t+1} - 1)^2 + (\phi_3 / 2) (\kappa_{t+1} - \kappa_t)^2] \} \quad (\text{B.4}) \end{aligned}$$

These two equations determine the joint evolution of investment and the capital utilization rate.

B.3 – Linearizing the model

The model in the text is log-linear. To facilitate comparison, it is useful to log-linearize the profit-maximization-based model as well. The box below summarizes some definitions and approximations that are useful in linearizing the model.

Some Definitions and Approximations Used in Deriving the Linearized Model

Definitions:

- “Gross” output, $X_t \equiv (U_t K_t)^\theta (A_t L_t)^{(1-\theta)}$.
- The user cost of capital, $C_t \equiv 1 - (1-\delta)(\beta_{t+1}/\beta_t)$.
- The log of the capital-output ratio, $\rho_t \equiv \ln (U_t K_t / X_t)$.
 - ▶ ρ_t is related to κ_t according to: $\rho_t = (1-\theta) \ln \kappa_t$.

Approximations:

- ▶ The average value of β_{t+1}/β_t is $(1-r)$.
- ▶ The average value of K_t/K_{t-1} is $(1+g)$.
- ▶ The average value of $\exp(\rho_t)$ is *rhobar*.
- ▶ $(I_t/K_{t-1} - \xi) \approx \Delta k_t + \delta - \xi$, where the lower-case letter indicates the log of the upper-case letter.
- ▶ $[(1-\delta)(\beta_{t+1}/\beta_t) - 1 + \theta U_t^\theta (A_t L_t / K_t)^{(1-\theta)}] \approx \text{cbar} (\ln \theta + x_t - k_t - c_{t+1})$, where *cbar* indicates the average value of C .
- ▶ $\Delta \kappa_t \approx \text{rhobar} \Delta \rho_t$.
- ▶ $\kappa_t^{(1-\theta)} \approx \text{rhobar} [1 - (\ln \text{rhobar}) / (1-\theta)] + \text{rhobar} \rho_t$.

Using these definitions, and ignoring constants and higher-order terms, we can derive the following two equations:

$$\Delta \rho_t = (1+g)(1-r) \Delta \rho_{t+1} + [(1+g)/(\phi_3 (1-\theta)\text{rhobar}^2)] \{ [\phi_2/(1+g)] u_t - \theta \text{rhobar} \rho_t \} \quad (\text{B.5})$$

$$\Delta k_t = (1+g)(1-r) \Delta k_{t+1} + (1/\phi_1) \{ -\text{cbar} c_t + [\phi_2/(1+g)] u_t \} \quad (\text{B.6})$$

These equations are similar to equations 7 and 10 in the main text. In both sets of equations, both the change in the log of the capital stock and the change in the log of the capital-output ratio are affected by their own discounted future value and by the capital utilization rate. A difference between the cost-minimization- and profit-maximization-based models is that the term $(\rho + c)$

enters the text equation for the capital-output ratio, but here, ρ and c are separated, with ρ entering the equation for the capital-output ratio and c entering the equation for capital-stock growth. An advantage of the specification in the text is that $(\rho + c)$ has a straightforward interpretation as an “error-correction” term, with the capital-output ratio proportional to user cost in the long-run.

Appendix C – A structural macroeconomic model

In this appendix, I describe the macroeconomic model that I used to simulate the effects of monetary and fiscal shocks in section 5. To facilitate solution, I use a “log-linearized” version of the model; Campbell (1994) has a useful discussion of log-linearization in the context of a simple stochastic growth model. Note that throughout this derivation, constant terms will be dropped, because, as a Taylor approximation, the log-linearized model applies only to deviations from some steady-state level.

The structural investment model (equations 8, 9, and 10) is already log-linear. However, looking to eventual closure of the model, it is helpful to express the model in terms of the real interest rate rather than the cost of capital. The log-linear approximation that is used is:

$$c_t = [(1 + rbar)/(rbar + \delta)] rr_t, \quad (C.1)$$

where $rbar$ is the average real interest rate and δ is the depreciation rate. The specific values chosen for $rbar$ and δ are discussed later in the appendix.

The standard capital-stock accumulation equation (2) also needs to be approximated in a log-linear model:

$$k_t = b_1 k_{t-1} + b_2 bfi_t \quad (C.2)$$

where $b_1 = (1 - \delta)/(1 + growth)$, $b_2 = (\delta + growth)/(1 + growth)$, and $growth$ is the average rate of growth of labor-augmenting technical progress, which I assume to be 2 percent per year, its average pace over the past fifty years in the United States.

Turning to the other structural relationships of the model, I assume that part of consumer spending is determined by a standard forward-looking Euler equation with “external” habit

formation and the remainder is proportional to (after-tax) income. Thus, overall consumer spending is:

$$pce_t = (1-\theta) pce_t^* + \theta (y_t - tbar \times tax_t), \quad (C.3)$$

where pce_t is (the log of) overall consumer spending, pce_t^* is the log of the forward-looking portion of consumer spending, tax_t is the log of taxes, and $tbar$ is the average ratio of taxes to income. The forward-looking portion of consumer spending is determined by:

$$pce_t^* = [E_t pce_{t+1}^* + h (1-\sigma) pce_{t-1}^* - \sigma rr_t] / [1 + h (1-\sigma)]. \quad (C.4)$$

I take my estimates of the share of consumer spending that is proportional to current income (θ), the intertemporal elasticity of substitution (σ), and the degree of habit persistence (h), from Fuhrer (2000, table 1); the values are 0.25, 0.16, and 0.80, respectively.

I assume that prices are sticky. In particular, I assume inflation is determined by a “hybrid New Keynesian Phillips curve,”

$$\Delta p_t = 0.5 \Delta p_{t-1} + 0.5 E_t \Delta p_{t+1} + \alpha (y_t - y_t^*) + \epsilon_t, \quad (C.5)$$

where y_t^* is the level of output that would exist in the absence of nominal rigidities (defined below). The microeconomic underpinnings of such a model are discussed in various places; see, for example, Roberts (2001) or Christiano, Eichenbaum, and Evans (2001). For the slope coefficient α , I use 0.02, consistent with the estimates of Roberts (2001).

Monetary policy is set according to a “dynamic Taylor rule” for the nominal interest rate,

$$(rr_t + \Delta p_t) = \mu_r (rr_{t-1} + \Delta p_{t-1}) + (1-\mu_r)[\mu_y (y_t - y_t^*) + \mu_p \Delta p_t]. \quad (C.6)$$

Monetary policy rules of this type have been estimated by many authors, including Clarida, Gali, and Gertler (2000) and English, Nelson, and Sack (2002). I use parameters of $\mu_p = 1.5$, $\mu_y = 1.0$, and $\mu_r = 0.72$, which are consistent with the estimates of English, Nelson, and Sack (2002), table 1, column 1; these estimates are for the 1987-2000 period.

Government consumption spending (gov_t) is assumed to be a fixed fraction of output (16 percent, the average over the 1960-to-2000 period). Taxes (tax_t) are lump-sum and set according to the fiscal-policy reaction function,

$$tax_t - y_t = 0.9(tax_{t-1} - y_{t-1}) + 0.03(debt_{t-1} - y_{t-1}) \quad (C.7)$$

where *debt* is government debt. The parameters ensure (a) that government debt will eventually converge to a fixed fraction of GDP (assumed to be 40 percent in the simulations) and that (b) any deviation of taxes from the level dictated by debt stability will be moderately persistent. These parameters are consistent with those for the fiscal reaction function in the Federal Reserve Board's FRB/US model.

The production function can be used to define "potential output," the level of output determined by the level of labor-augmenting technology and the capital stock:

$$x_t^* = (1-\theta) a_t + \theta k_t. \quad (C.8)$$

This level of output abstracts from the influence of a number of model "rigidities" on output, most notably, the assumption of sticky prices. Implicit in this definition is the assumption that labor supply, in the absence of sticky prices, is inelastic. The elasticity of output with respect to the capital stock, θ , is 1/3, as in Campbell (1994).

Finally, the (log-linearized) expenditure identity is:

$$y_t = cbar \times pce_t + ibar \times bfi_t + gbar \times gov_t, \quad (C.9)$$

where *cbar*, *ibar*, and *gbar* are the shares of consumer spending, investment, and government consumption spending in overall spending. As noted earlier, I set *gbar* = 0.16, its average share over the 1960-to-2000 period. I set the average real interest rate equal to 8 percent and the depreciation rate equal to 7 percent, each in annual terms. These parameter choices yield *ibar* = 0.20 and *cbar* = 0.64, the average shares of investment and consumer spending over the 1960-to-2000 period. Note that this investment share includes all categories of business fixed investment as well as residential and government investment. I am thus implicitly assumes that all investment behaves like non-high-tech equipment.

Table 1 Estimates of Standard Model (Convex Costs of Adjusting the Stock and Rate of Investment in Capital) Estimation period: 1965:Q1-2001:Q4					
	(1)	(2)	(3)	(4)	(5)
	Minimum distance, output	Minimum distance, user cost; γ_1 constr'd	Minimum distance, user cost; γ_1 constr'd	Minimum distance, both output & user cost	FIML
γ_1	.0064 (.0010)	.0064 –	.020 –	.0059 (.0009)	.0038 (.0013)
γ_2	.195 (.065)	.0158 (.0025)	.0040 (.0005)	.036 (.010)	.038 (.027)
<i>Minimum-distance statistics</i>					
<i>Output impulse responses</i>	1.8	57.8	70.0	33.0	49.0
<i>User cost impulse responses</i>	165.6	7.3	2.0	19.0	15.7
<i>Both</i>	167.4	61.0	72.0	52.0	64.7
<i>Annualized adjustment speed:</i>					
Capital	.25	.25	.41	.24	.19
Investment	.82	.37	.19	.51	.52
<i>Likelihood</i>	–	–	–	–	2182.3

Notes: Numbers in parentheses are standard errors.

Table 2 Estimates of Model with Convex Costs of Adjusting the Capital Intensity (In Addition to the Stock and Rate of Investment in Capital) Estimation period: 1965:Q1-2001:Q4					
	Without utilization			With utilization	
	Minimum distance		FIML	Minimum distance	FIML
	(1)	(2)	(3)	(4)	(5)
γ_1	.0154 (.0061)	.0064 –	.0045 <i>(.00003)</i>	.0091 (.0022)	.0052 <i>(.00002)</i>
γ_2	.036 –	.048 (.012)	.075 (.017)	.149 (.038)	.102 (.034)
γ_3	.0046 (.0008)	.0047 (.0012)	.0037 (.0008)	.0047 (.0006)	.0017 (.0007)
γ_4	1.50 (.88)	4.54 (1.57)	2.97 <i>(.08)</i>	.95 (.21)	2.18 <i>(.01)</i>
<i>Minimum-distance statistics</i>					
<i>Output impulse responses</i>	5.3	3.2	9.4	7.8	6.8
<i>User cost impulse responses</i>	14.2	22.6	30.1	17.1	34.5
<i>Both</i>	19.5	25.8	39.5	24.9	41.3
<i>Annualized adjustment speed:</i>					
Capital	.37	.25	.21	.29	.22
Investment	.51	.57	.65	.78	.71
Capital-output ratio	.21	.21	.19	.21	.11
<i>Minimum-distance statistics, investment and utilization impulse responses</i>	–	–	–	56.3	91.3
<i>Likelihood</i>	–	–	2190.3	–	2948.9

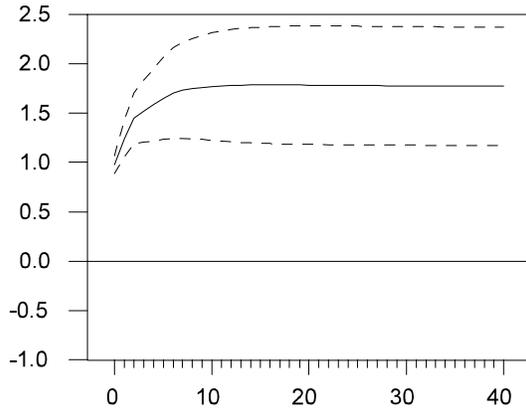
Notes: Numbers in parentheses are standard errors.

Numbers in italics indicate FIML standard errors that may manifest numerical estimation problems.

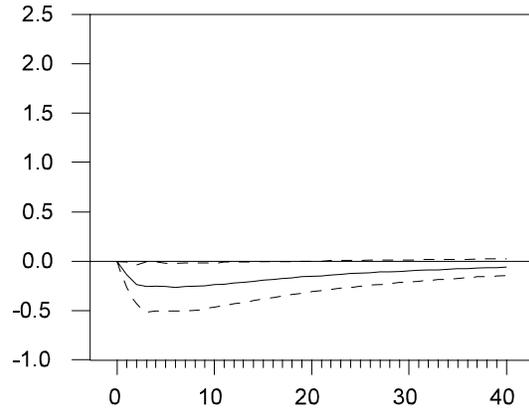
Figure 1

Impulse Responses in the "Core" VAR

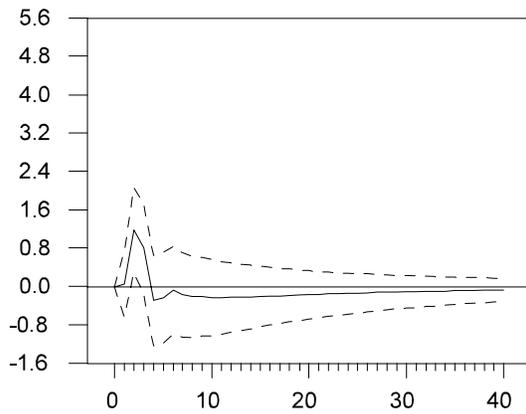
Effect of Output Shock on Output



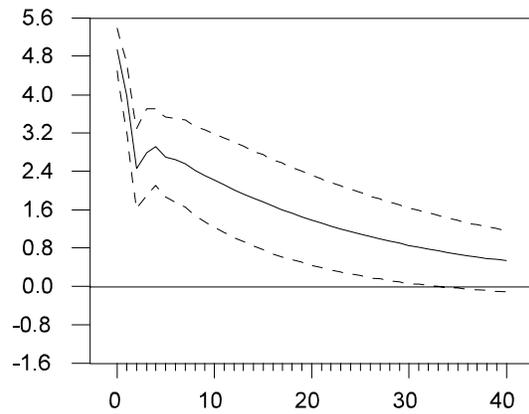
Effect of UserCost Shock on Output



Effect of Output Shock on UserCost



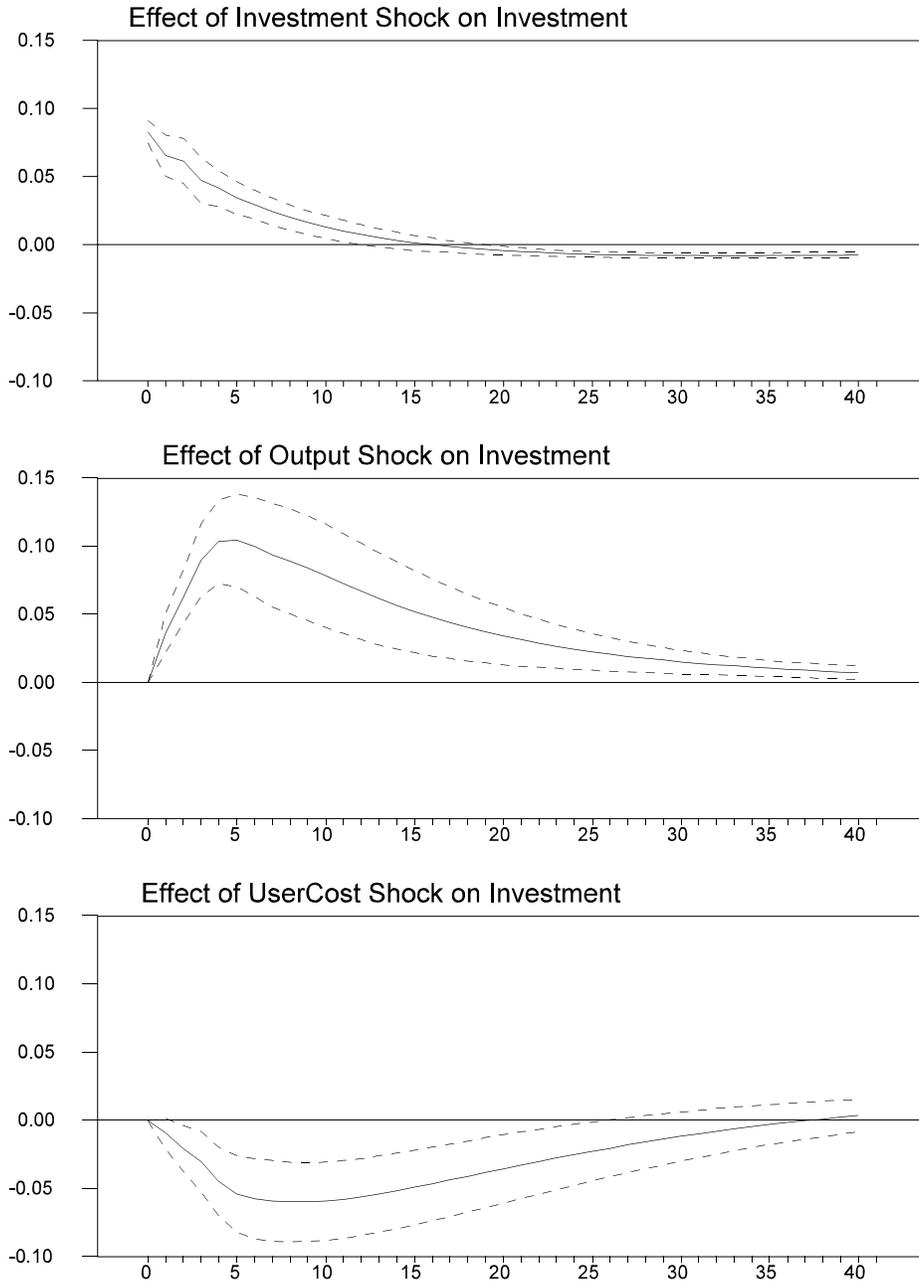
Effect of UserCost Shock on UserCost



Note: Dotted lines are upper and lower bounds of a centered 90 percent confidence interval.

Figure 2

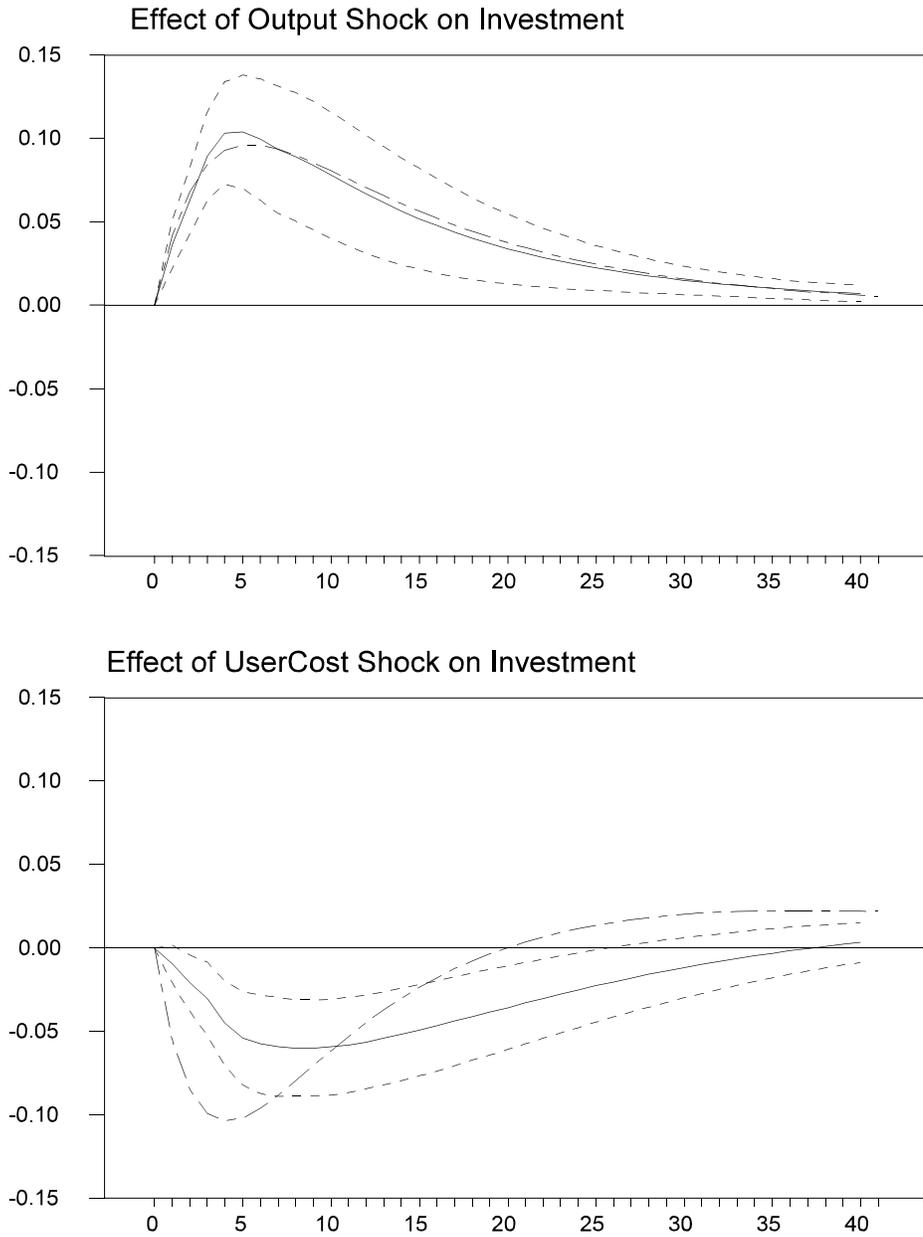
Responses of Investment to One-standard-deviation Shocks



Note: Dotted lines are upper and lower bounds of a centered 90 percent confidence interval.

Figure 3

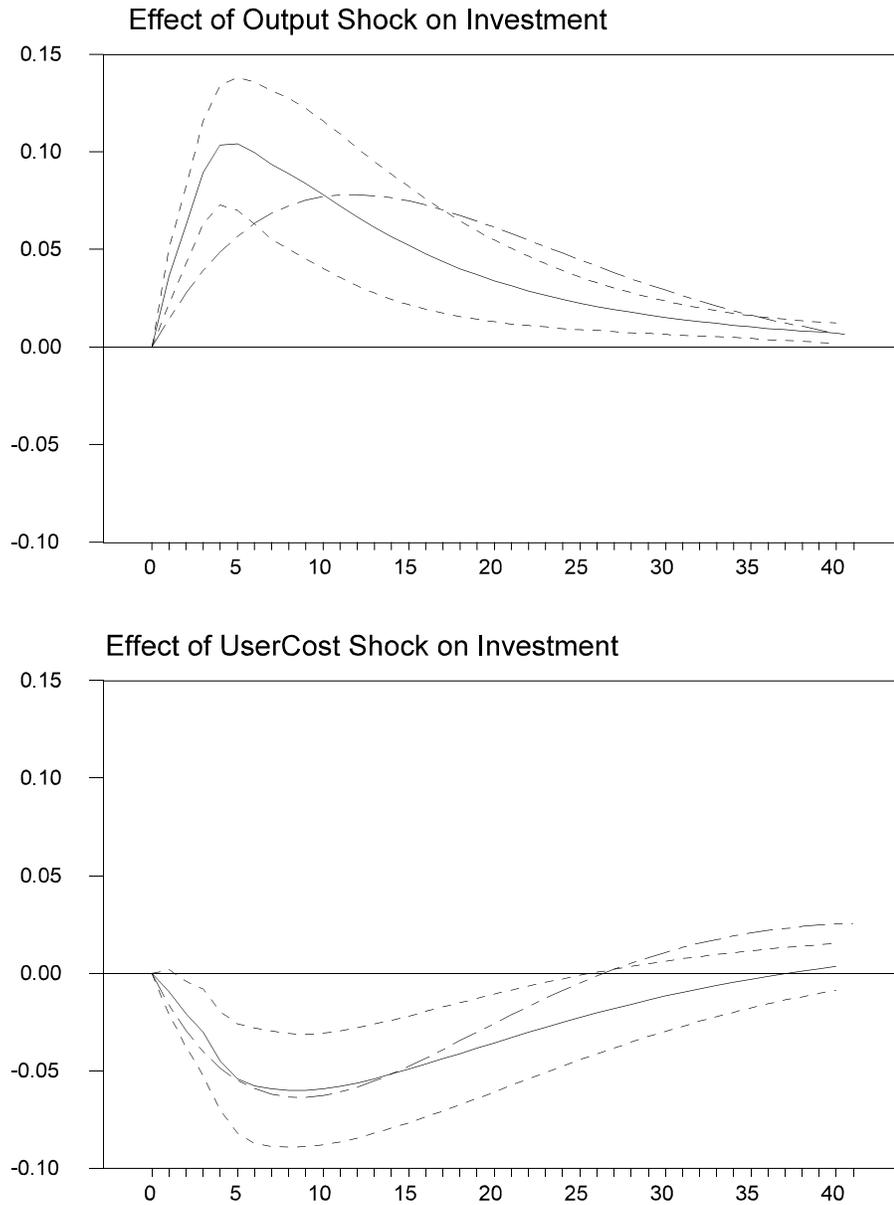
**Impulse Responses of “Standard” Model,
Parameters Chosen to Fit Response of Investment to Output Shock**



Note: Solid line is reduced-form impulse response; dotted lines are upper and lower bounds of a centered 90 percent confidence interval; dash-dot line is structural model impulse response.

Figure 4

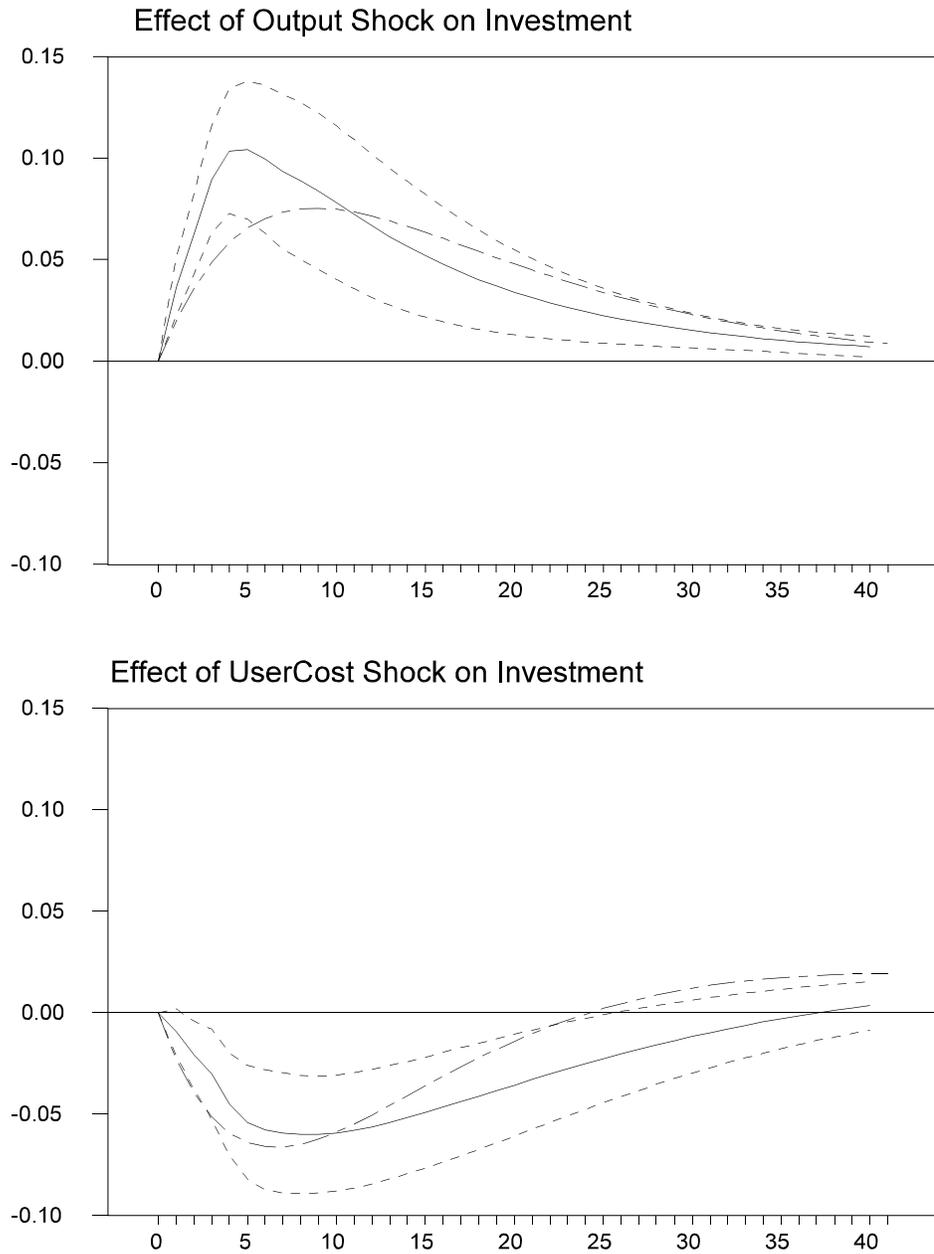
**Impulse Responses of "Standard" Model,
Parameters Chosen to Fit Response of Investment to User Cost Shock**



Note: Solid line is reduced-form impulse response; dotted lines are upper and lower bounds of a centered 90 percent confidence interval; dash-dot line is structural model impulse response.

Figure 5

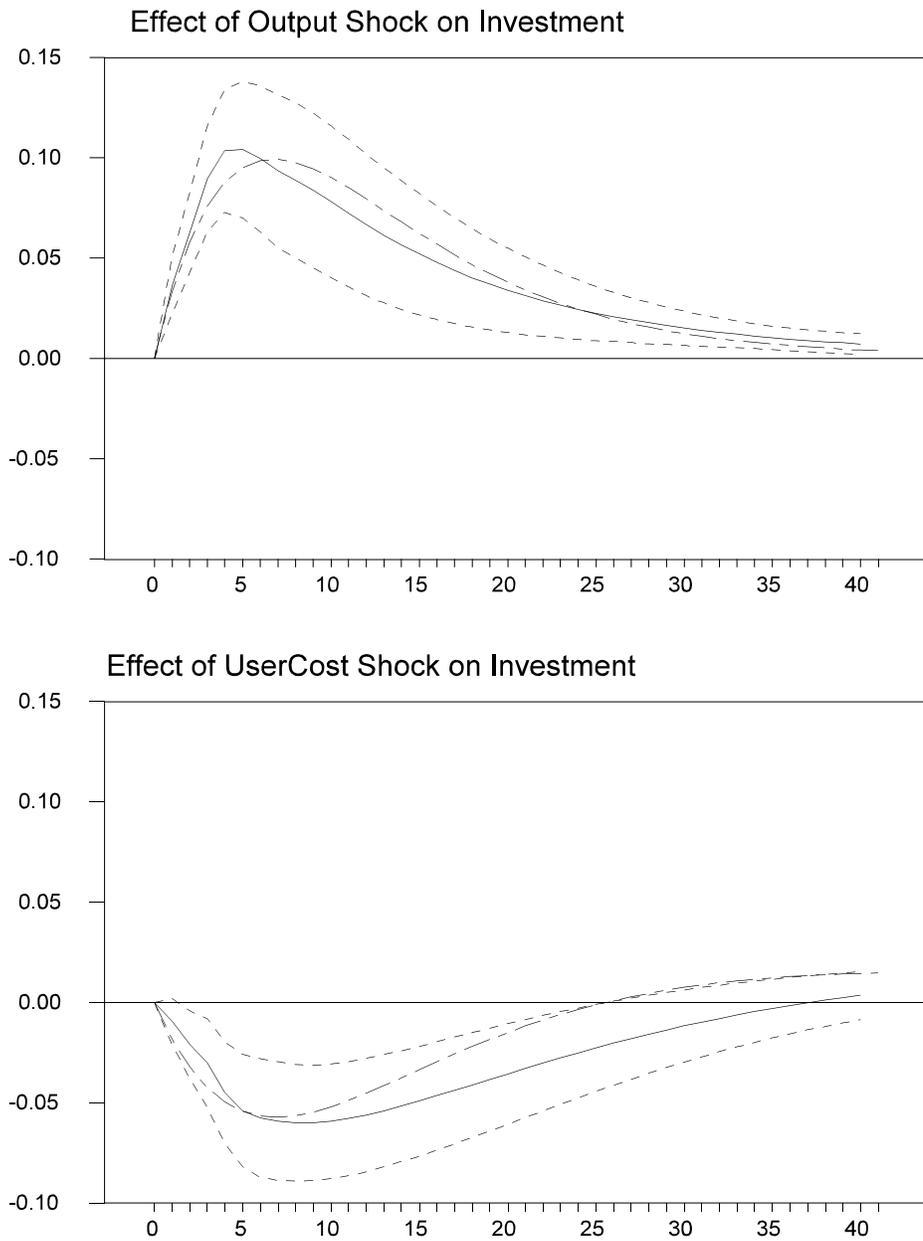
Impulse Responses of “Standard” Model, Parameters Chosen To Fit Response of Investment to Both Output and User Cost Shocks



Note: Solid line is reduced-form impulse response; dotted lines are upper and lower bounds of a centered 90 percent confidence interval; dash-dot line is structural model impulse response.

Figure 6

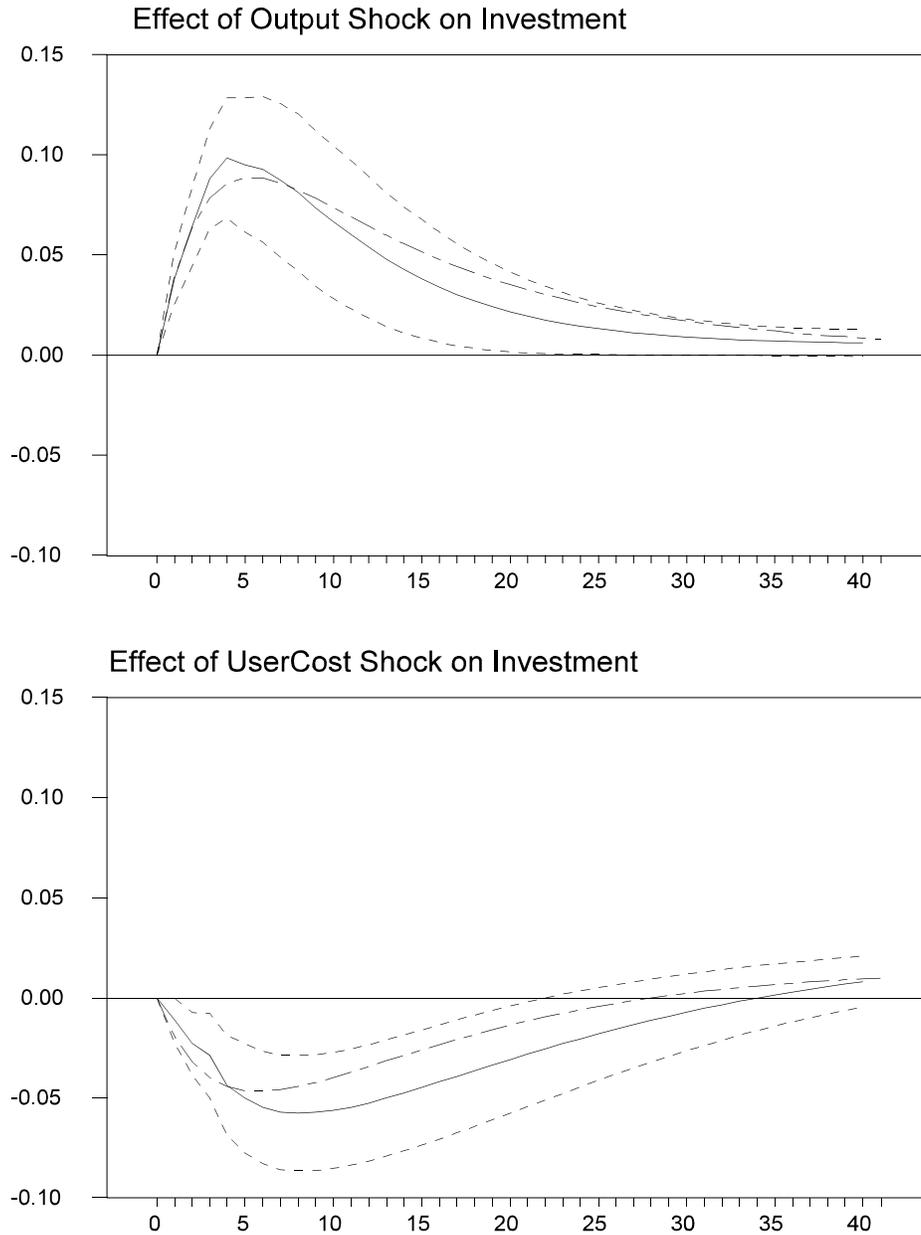
Impulse Responses of Model with Costly Adjustment of the Capital-Output Ratio, Matching Investment Impulse Responses Only



Note: Solid line is reduced-form impulse response; dotted lines are upper and lower bounds of a centered 90 percent confidence interval; dash-dot line is structural model impulse response.

Figure 7

Impulse Responses of Model with Costly Adjustment of the Capital-Output Ratio, Matching Both Investment and Utilization Impulse Responses



Note: Solid line is reduced-form impulse response; dotted lines are upper and lower bounds of a centered 90 percent confidence interval; dash-dot line is structural model impulse response.

Figure 8

Impulse Responses to a 1 Percentage Point Monetary Easing

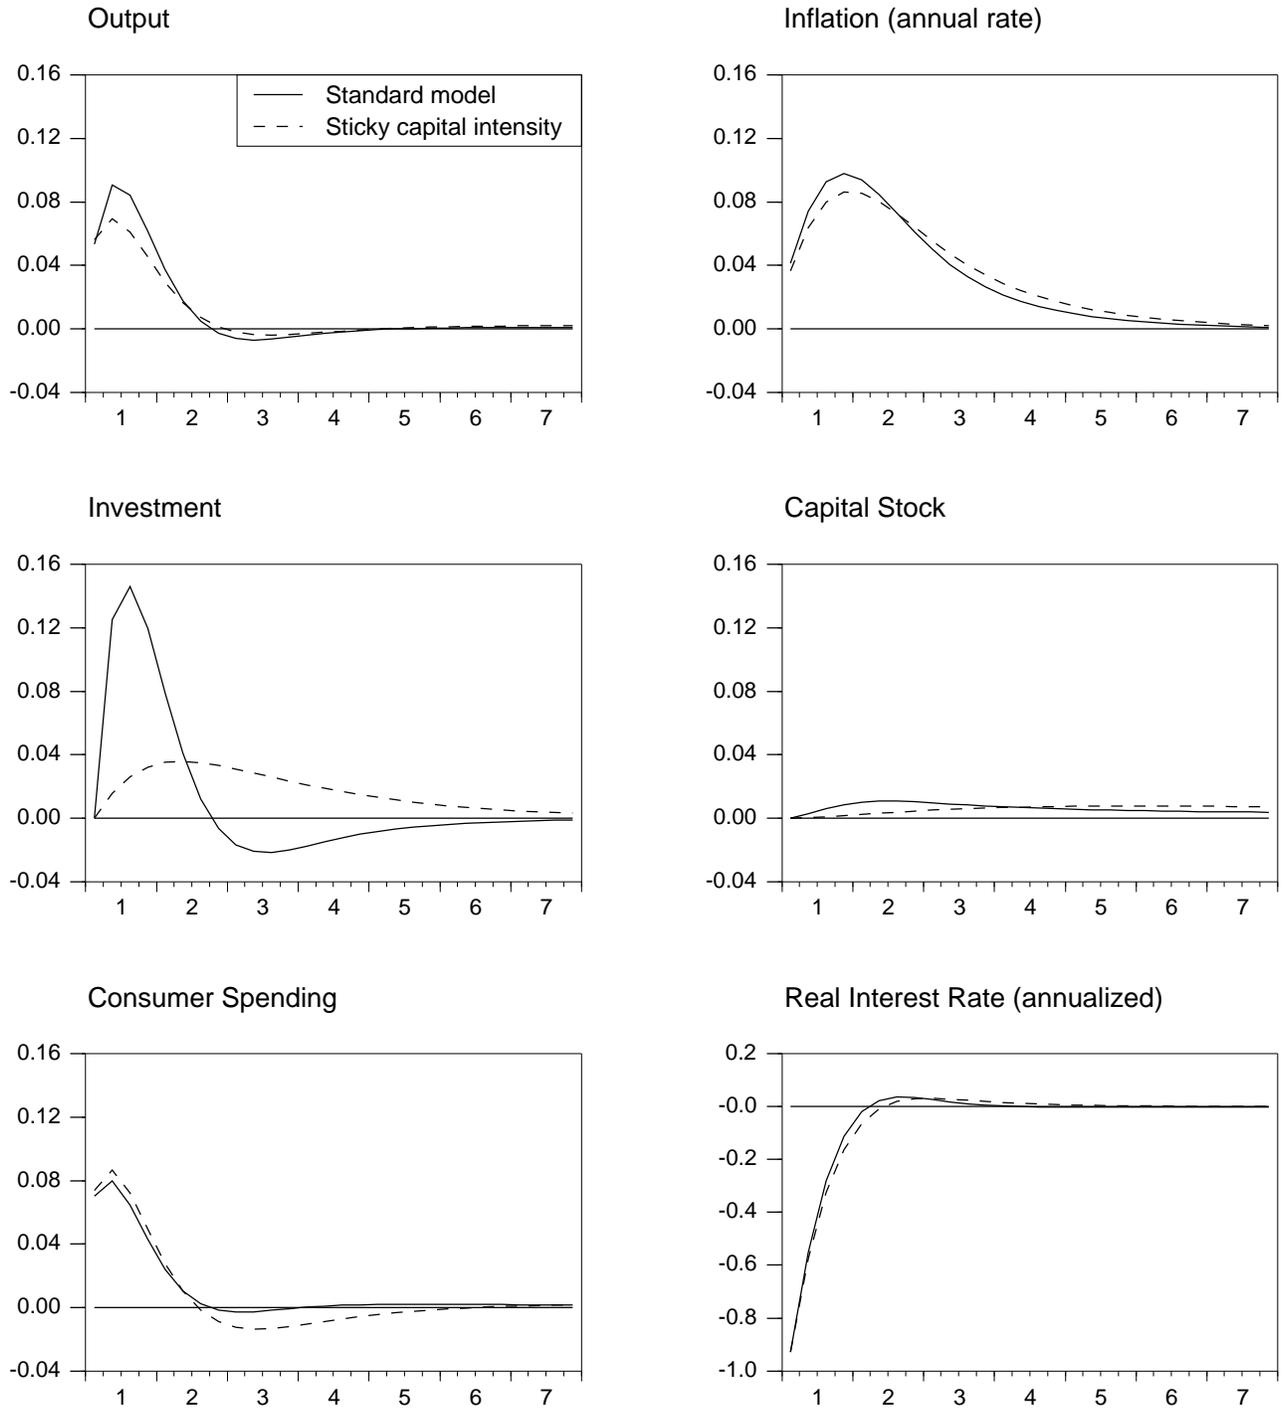
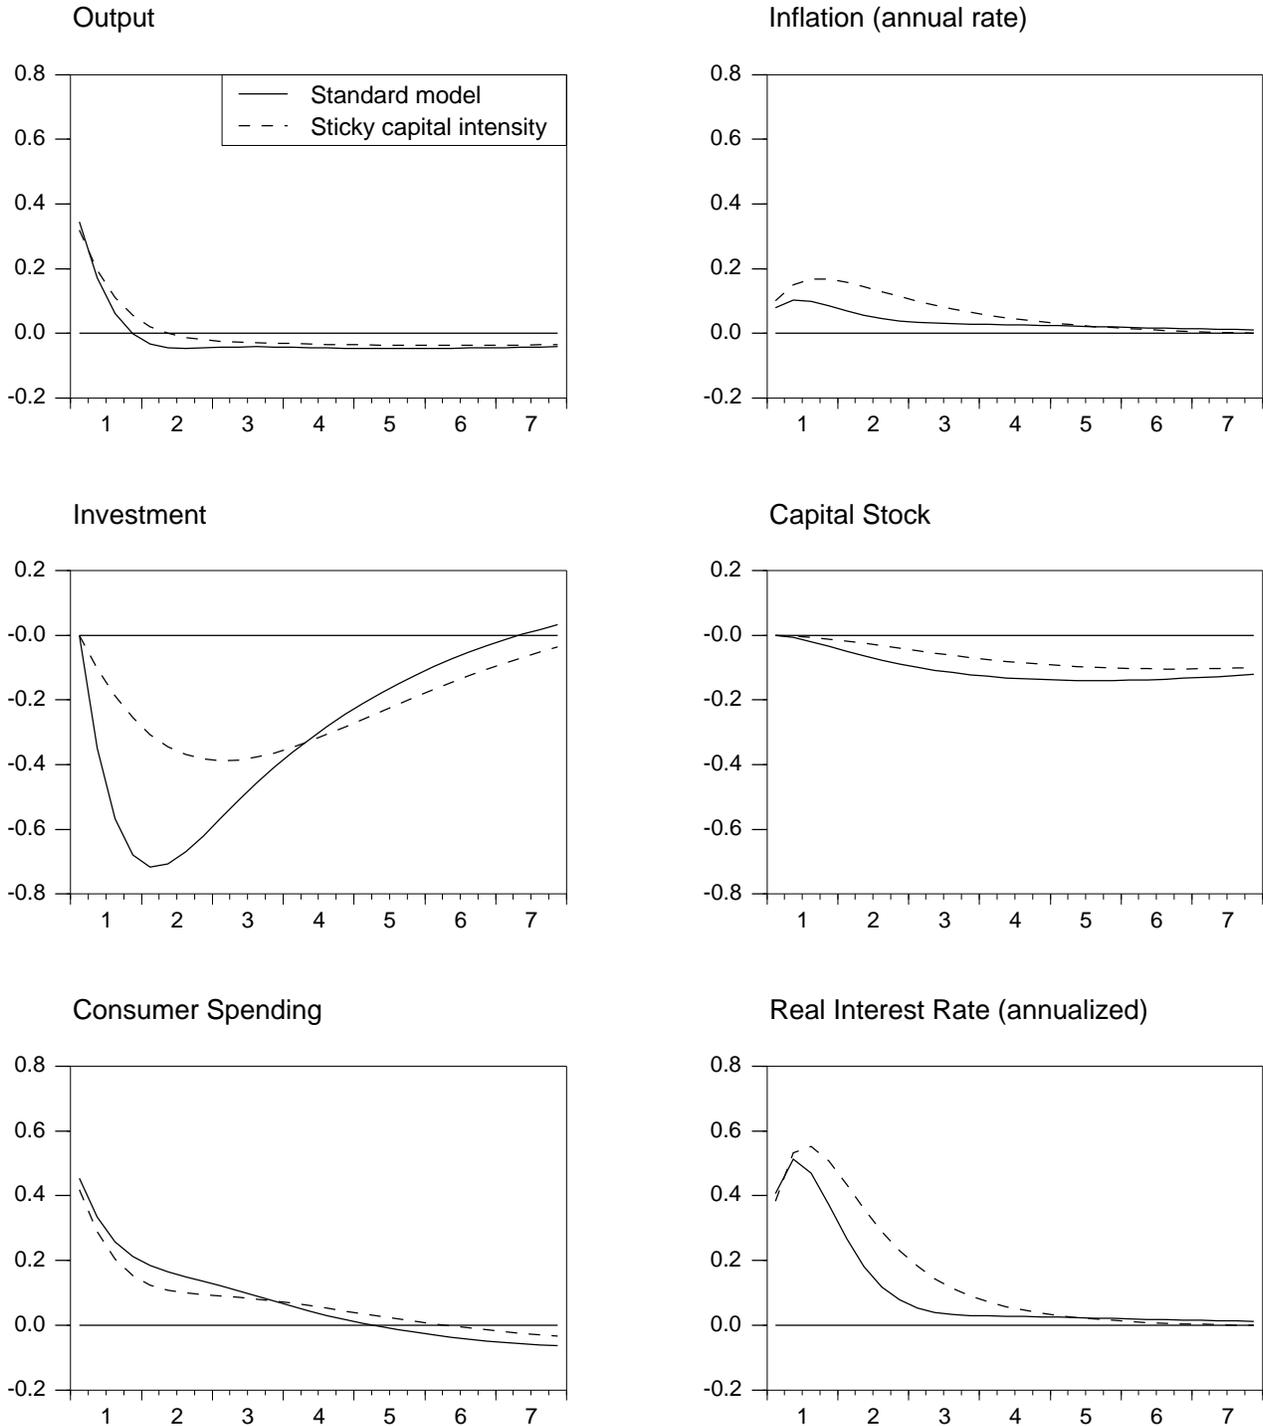


Figure 9
Impulse Responses to a 10% Tax Cut



References

- Abel, Andrew B., and Olivier J. Blanchard, 1986. The Present Value of Profits and Cyclical Movements in Investment, *Econometrica* 54, 249-73.
- Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde, 2002. Technology Shocks and Aggregate Fluctuations, working paper (June).
- Bischoff, Charles W., 1971. The Effect of Alternative Lag Distributions, in Gary Fromm, ed., *Tax Incentives and Capital Spending* (Brookings Institution) 61-130.
- Brayton, Flint, and Eileen Mauskopf, 1985. The MPS Model of the United States Economy, Federal Reserve Board, Washington, D.C.
- Campbell, John C., 1994. Inspecting the Mechanism: An Analytical Approach to the Stochastic Growth Model, *Journal of Monetary Economics* 33, 463-506.
- 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, 2001. Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy, PDF file (May).
<http://faculty.econ.nwu.edu/faculty/christiano/research/modeling/papermay10-2001.pdf>
- Clarida, Richard, Jordi Gali, and Mark Gertler, 2000. Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory, *Quarterly Journal of Economics* 115, 147-80.
- Clark, Peter K., 1979. Investment in the 1970s: Theory, Performance, and Prediction, *Brookings Papers on Economic Activity* 1979:1, 73-113
- Cummins, Jason G., Kevin A. Hassett, and Stephen D. Oliner, 1999. Investment Behavior, Observable Expectations, and Internal Funds, Federal Reserve Board *FEDS* working paper no. 1999-27 (May).
- Edge, Rochelle M. and Jeremy B. Rudd, 2003. Temporary Partial Expensing: A Sticky-Price DGE Analysis, working paper (July).
- English, William B., William R. Nelson, and Brian P. Sack, 2002. Interpreting the Significance of the Lagged Interest Rate in Estimated Monetary Policy Rules, Federal Reserve Board *FEDS* working paper no. 2002-24 (May).
- Fuhrer, Jeffrey C., 2000. Habit Formation in Consumption and Its Implications for Monetary-Policy Models, *American Economic Review* 90, 367-90.

- _____, George Moore, and Scott Schuh, 1995. Estimating the Linear-quadratic Inventory Model: Maximum Likelihood versus Generalized Method of Moments, *Journal of Monetary Economics* 35, 115-185.
- Gilchrist, Simon, and Charles P. Himmelberg, 1995. Evidence on the Role of Cash Flow for Investment, *Journal of Monetary Economics* 38, 541-72.
- _____ and _____, 1998. Investment: Fundamentals and Finance, *NBER Macroeconomics Annual 1998*, 223-73.
- _____ and John C. Williams, 2000. Putty-Clay and Investment: A Business-Cycle Analysis, *Journal of Political Economy* 108, 928-60.
- Kiley, Michael, 2001. Business Investment in the Federal Reserve's U.S. Model (FRB/US): Specification and Implications, Federal Reserve Board (April).
- Kiyotaki, Nobuhiro, and Kenneth D. West, 1996. Business Fixed Investment and the Recent Business Cycle in Japan, *NBER Macroeconomics Annual 1996*, 277-344.
- Lichtenberg, Frank, 1988. Estimation of the Internal Adjustment Costs Model Using Longitudinal Establishment Data, *Review of Economics and Statistics* 70, 421-430.
- Oliner, Stephen, Glenn Rudebusch, and Daniel Sichel, 1995. New and Old Models of Business Investment: A Comparison of Forecasting Performance, *Journal of Money, Credit, and Banking* 27, 806-26.
- Roberts, John M., 1997. Is Inflation Sticky? *Journal of Monetary Economics* 39, 173-96.
- _____, 2001. How Well Does the New Keynesian Sticky-Price Model Fit the Data? Federal Reserve Board *FEDS* working paper no. 2001-13 (February).
- Shapiro, Matthew D., 1986a. The Dynamic Demand for Capital and Labor, *Quarterly Journal of Economics* 103, 513-42.
- _____, 1986b. Capital Utilization and Capital Accumulation: Theory and Evidence, *Journal of Applied Economics* 1, 211-34.
- Summers, Lawrence H., 1981. Taxation and Corporate Investment: A q -Theory Approach, *Brookings Papers on Economic Activity* (1981:1), 67-127.
- Tevlin, Stacey, and Karl Whelan, 2003. Explaining the Investment Boom of the 1990s, *Journal of Money, Credit, and Banking* 35, 1-22.
- Topel, Robert, and Sherwin Rosen, 1988. Housing Investment in the United States, *Journal of Political Economy* 96, 718-40.