Understanding Changes In Aggregate Business Fixed Investment

BY AUBHIK KHAN

When economists talk about business fixed investment, they mean the expenditures by firms on equipment and structures. Business fixed investment is commonly held to be an important determinant of an economy's long-run growth. [1]

On average, higher levels of such investments raise production by increasing the productivity of the labor force. While the significance of short-term changes in business investment is less widely recognized, the importance of such changes for the business cycle has been known to economists since the beginning of the last century. For example, many believe that the current record expansion has been driven, at least in part, by strong investment in computers and related equipment. In this article, I attempt to explain some of what economists have learned about how investment changes over the business cycle.

[1] The definition of business fixed investment used throughout this paper does not include software expenditures by firms because these data are not available.

Aubhik Khan is an economist in the Research Department of the Philadelphia Fed.

INVESTMENT AT THE PLANT LEVEL

For individual plants, investment is simply the expenditure required to adjust its stock of capital. Capital includes all equipment and structures the plant uses. The plant combines capital with other inputs, such as labor and energy, to produce goods or services. When a mining company acquires diesel engines, it is investing in equipment. When an automobile manufacturer builds a new warehouse, it is investing in structures. Because it takes time to manufacture, deliver, and install new capital goods, investment expenditures today do not immediately raise the level of a plant's capital. So investment involves a planning decision that trades off present against future earnings. Investment expenditures today reduce current profit but increase a plant's future possible production and, as a result, future profit.

Since investment spending raises future capital and thus the quantity of goods and services that may be produced in the future, plants will tend to adjust their investment levels in response to forecasted changes in the market's demand for their own output. Changes in productivity — the efficiency with which inputs may be combined to produce output — will also tend to increase investment. For example, if productivity increases, the firm may be able to sell more of its product, since it can offer it at a more attractive price. The firm may then expand and more workers may be hired. These workers will need equipment, and, as a result, investment will rise.

AGGREGATE INVESTMENT OVER THE BUSINESS CYCLE

When plants anticipate increased demand for their output or higher productivity, they will generally raise their investment spending. For most sectors of the economy such increases in investment occur when GDP rises, for example, during economic expansions. In contrast, if plants expect a decline in demand, such as occurs for most plants when GDP falls, investment spending will fall. As a result, *aggregate investment* — the sum of all investments by all plants in the economy — is *procyclical:* it rises when output rises and falls when output falls over the business cycle.

Even a casual glance at the

data will reveal that investment is much more volatile than output (Figure 1). [2] During periods of above-trend growth, aggregate investment experiences a much larger percentage rise. Moreover, when growth rates are below trend, such as during recessions, aggregate investment falls far more sharply than does aggregate output. Indeed, if we use a standard measure of variability, quarterly investment is 3.4 times more volatile than quarterly output over 1956 - 1994.[3]

The reader will likely note another striking regularity between the two series: investment and output almost always move in the same direction. For example, in the sharp recession of the early 1980s, detrended output fell almost 5 percent, and concurrently, investment fell more than 10 percent.[4] Plants adjust investment in anticipation of changes in output demand, and consequently, investment moves similarly to output. It follows that to understand the business cycle, we must understand why aggregate investment changes over time. As Harvard economist Robert J. Barro has stated, "As a first approximation, explaining recessions amounts to explaining the sharp contractions in the investment components" (p. 245).

But to understand why aggregate investment fluctuates, economists are learning that they must understand the decisions of individual plants. Such emphasis on the role of an individual entity characterizes recent progress in many areas of macroeconomics.

[2] Episodes of negative growth rates in Figure 1 do not imply recessions, at least as they are commonly understood. For example, if the output trend is 3 percent, and actual output grows at 2 percent, Figure 1 will report –1 percent. Of course, actual recessions will be recorded when growth is negative.

[3] The percentage standard deviation of output is 1.4 while that of investment is 4.9, hence the ratio of 3.4.

[4] This co-movement in investment and output is captured by a correlation coefficient of 0.92 between the two series.

FIGURE 1

Investment and Output over the Business Cycle

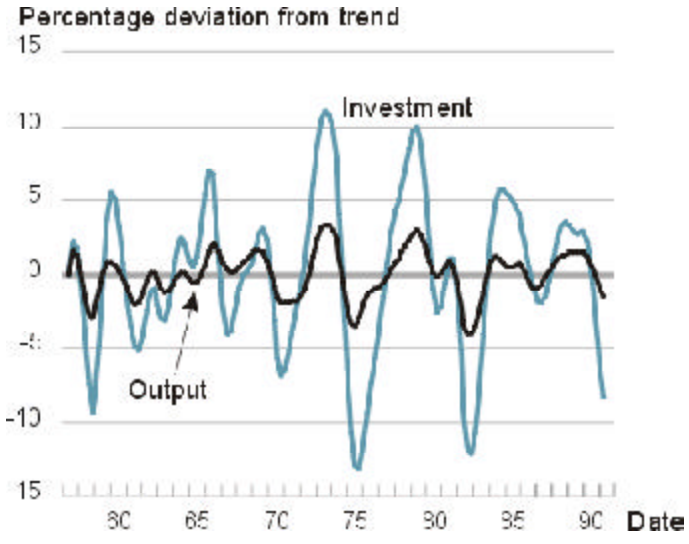

Figure 1 displays detrended quarterly total real business fixed investment and GDP in the United States in each quarter over the years 1956 – 1994. Since we want to concentrate on how these series move over the business cycle, we have detrended them. That is, the figure shows changes in output and investment from their longer-term trends. These trends were computed using the Band-Pass Filter developed by Marianne Baxter and Robert G. King in their 1999 paper. Note that the use of this filter eliminates several years of data at the beginning and end of our series.

THE PARTIAL ADJUSTMENT MODEL

As with all other forms of scientific progress, progress in economics relies on the development of theories. The success of these theories is determined by their ability to contribute to an explanation of observed phenomena. In the study of investment, this has led to a theory of how firms choose their levels of investment.

Traditionally, economists tried to understand aggregate investment using an approach that ignored possible differences across individual firms. This approach led to a theory that relied on the fiction of a representative firm that undertook all investment that actually occurred in the economy. In reality, many firms own and operate several plants, and investment decisions are made at both the firm and the plant level. But as long as we examine representative firms, there is no meaningful distinction between a firm (an ownership unit) and a plant (a production unit).

Let's consider a representative firm, BIGCAP. Even if BIGCAP sees no reason to change the level of its capital stock, it will nonetheless have to undertake some maintenance to sustain capital stock at current

levels, since capital depreciates over time. Investment beyond the level needed to offset depreciation will raise the stock of capital BIGCAP will have in the future. This higher level of capital will allow BIGCAP to raise production. So investment today will affect future earnings and, thus, future profits. Therefore, by undertaking investment today and building capital, BIGCAP can influence its future profits. The fundamental assumption of the standard theory of investment is entirely reasonable: A firm chooses its stock of capital in order to maximize its shareholder value. This is the firm's *target level* of capital.

Adjustment Costs. However, modern variants of this theory make another important assumption: there is a cost associated with changing a firm's capital, the cost of adjustment itself. In their 1996 review paper, Daniel Hammermesh and Gerard Pfann discuss some of the sources of these adjustment costs. Adding a new machine takes time. During installation, the firm must reallocate production across its other machines, a move that may overburden these other machines and may present machine operators with unfamiliar working conditions. As a result, production will fall during this first adjustment period. Next, after the new machine has been installed, workers must be trained to use it. Again, the firm will be operating at temporarily reduced levels of productivity during this second adjustment period.

Overall, when a firm installs new capital goods it incurs internal costs over and above the cost of the equipment itself. These costs reduce the firm's profits over the adjustment period.

Consider what happens if BIGCAP purchases a new computer to add to its existing stock. In addition to the price of the computer equipment, BIGCAP will incur additional costs of integrating the machine into its network and setting it up with the required software. The nature of these costs — how they change with respect to the quantity of investment undertaken by BIGCAP — is critical in determining their effect. Traditional investment theory assumes that it costs more, *per unit*, to install more capital. Thus, BIGCAP's cost of installing two new computers would be more than twice the cost of installing a single machine.

Rising costs of adjustment imply that adjusting capital rapidly would cost more than doing it gradually. So traditional theory said that firms adjusted to their target capital stock — that which maximized shareholder value — slowly in an effort to reduce adjustment costs. So this theory was called the *partial adjustment model.*

It is not at all obvious why the costs of adjustment should rise with the level of investment. We might well think that competent computer staff, learning from setting up the first computer, would install the second in much less time. However, when rising adjustment costs were ignored, the model performed very poorly, since it predicted too much volatility in aggregate investment. So by including rising adjustment costs, the model better matched the data for the economy as a whole.

Figure 2 shows how assuming rising adjustment costs leads to smoother aggregate investment. It displays two possible models of a firm's investment over time.[5] For each one, the vertical axis displays the firm's current level of investment over time, as a percentage of its long-run average level. Suppose the firm experiences a rise in productivity, lowering its costs, or, instead, a rise in expected demand

[5] Figures 2 and 3 were generated by solving economic models of a firm's behavior under different assumptions about the costs of capital adjustment.

FIGURE 2

Investment With and Without Adjustment Costs

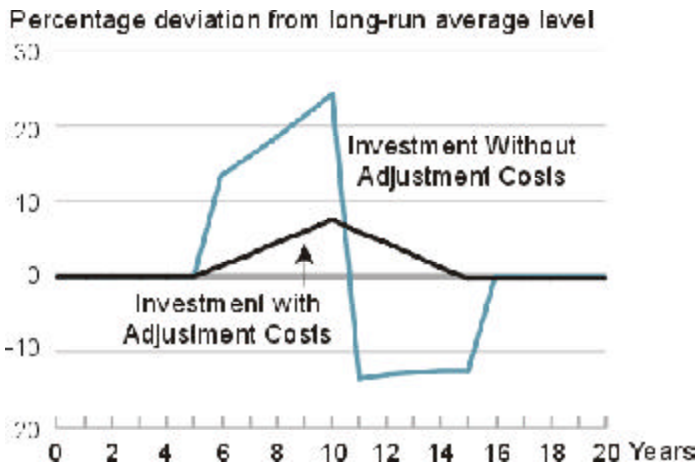

for its product. As a result, it chooses to increase its capital stock so that it can produce more. The blue line indicates the investment the firm will make if it faces no adjustment costs: there is a sharp rise in investment as the firm immediately adjusts its capital stock to allow it to efficiently increase production. Subsequently, when productivity or demand eventually returns to normal, there is an equally dramatic disinvestment episode, as the firm sells off its excess capital stock. In contrast, if adjustment costs rise with the level of investment, the change in capital is much more protracted. Capital partially adjusts in each period as investment slowly raises it toward its target value. As a result, when the change in productivity or demand ends, the plant has much less disinvestment to do. Investment is much more gradual under partial adjustment.

Since we are looking at a representative firm, total investment for the economy is the same as this firm's investment. Hence, more gradual investment at the firm level means that aggregate investment, that is, the total investment of all firms, shares the same properties.

When we compare Figures 1 and 2, we see that the model without adjustment costs generates an investment series that is too volatile when compared with the data. For example, in the model without adjustment costs, the largest deviation of investment from its trend is 25 percentage points, but in the data over 1956 – 94, the largest deviation was 10 percentage points. But when we examine the model with adjustment costs, we see that it exhibits much less variability in investment. As a result, the introduction of adjustment costs allows for a far better match with the aggregate data.[6]

Adjustment Costs Revisited. The partial adjustment model means gradual change in investment at the aggregate level, which matches the data, but it also means gradual adjustment in investment at each individual firm — and this does not seem to match the data! When researchers at the Bureau of the Census undertook an extensive study of how manufacturing plants adjusted their stock of capital, the story they uncovered was *inconsistent* with the predictions of the partial adjustment model. Instead of changing capital slowly and gradually, plants made capital adjustments that were lumpy, that is, they would invest a lot at one time, then refrain from investing for a while, then invest a lot again, and so on. Typically, plant capital remains roughly constant for long periods of time, with low levels of associated investment. These long episodes of relative inactivity are interrupted by sudden bursts of investment spending that drive large increases in plants' capital stock over short periods of time. The partial adjustment model with rising adjustment costs predicted plant-level investment that was too smooth. Given the limited success of the partial adjustment model, macroeconomists began to reconsider the plausibility of the assumption about rising adjustment costs. Indeed, much of the recent progress in our understanding of investment has arisen from replacing the unrealistic assumption of rising costs of adjustment with a better one. (See *How Do Plants Adjust Their Capital?*)

[6] It should be noted, however, that the match is still imperfect. Adjustment costs reduce variability too much (the largest deviation from trend in the model with adjustment costs is about 5 percentage points).

How Do Plants Adjust Their Capital?

In their 1998 paper, Mark Doms and Timothy Dunne examined capital adjustment at the plant level. Using the Longitudinal Research Datafile collected by the U.S. Bureau of the Census, they studied changes in the capital stock of 13,700 large U.S. manufacturing plants over 1972 – 1988. In terms of the total number of manufacturing plants, the sample is small: over this period, between 312,000 and 360,000 plants were operating in the manufacturing sector. However, the sample accounts for approximately 50 percent of total manufacturing production and 40 percent of employment. In addition to including relatively large plants, the sample is also unusual because all the plants present in the sample in 1972 were still in it through 1988.

In a typical year, over 80 percent of all plants in the sample undertook very little capital adjustment: their capital stocks changed less than 10 percent. But approximately 8 percent of plants adjusted capital by more than 30 percent, and more than half of the sample experienced capital growth of more than 37 percent in at least one year.

The partial adjustment model, which predicts gradual changes in investment due to the rising costs of undertaking too much capital adjustment at one time, cannot explain these sharp, sudden investment episodes followed by long periods of low adjustment.

As the mathematical sophistication of researchers in the field increased, they understood how to adapt the existing theory of investment to account for the new observations. The new theory assumes that the costs of capital adjustment are unrelated to the scale of the adjustment. Much of the adjustment cost borne by a plant would now be the same whether it was adding one, two, or even 10 computers to its network.

Such *fixed costs* (fixed because they are the same regardless of the amount of investment) lead to lumpy investment over time at the plant. Let's consider BIGCAP once again, assuming that BIGCAP is a firm that owns only one plant. BIGCAP determines its target level of capital, the level that maximizes shareholder value in the absence of adjustment costs. However, BIGCAP will adjust to this capital stock only if the rise in shareholder value from doing so is greater than the fixed cost associated with the capital adjustment. As explained by Ricardo Caballero in his 1999 paper, what this means is that a plant like BIGCAP will adjust its capital only when the current level of its capital stock is far enough away from its target level of capital stock.

If current and target capital levels are close, there's not much gain in shareholder value from adjustment; the fixed adjustment cost outweighs the benefits of adopting the target level of capital. But once it decides to adjust its capital stock, BIGCAP has no incentive to move gradually, since the adjustment cost is independent of the size of the adjustment. Notice that a simple modification of existing theory has led to a dramatic change in the model's predictions. Investment at the plant level is no longer slow and gradual but rather erratic and lumpy. Plants don't change their actual capital in response to small changes in their target capital. So there are typically long periods when plants don't undertake much investment. However, when target capital is sufficiently different from actual capital, there is sudden, sharp adjustment.

THE SUM OF INDIVIDUALS: THE IMPLICATION OF FIXED ADJUSTMENT COSTS FOR AGGREGATE INVESTMENT

Fixed adjustment costs seem to fit the plant-level data quite well. But how well do they match the aggregate data? Before we add up individual plants' behavior, we must understand how these fixed adjustment costs vary across plants and across time. Once this is accomplished, we will see that the new model actually fits the aggregate data better than the partial adjustment model.

How can a model of lumpy plant-level investment match the aggregate investment data, which show gradual changes in investment? The answer is that the fixed costs that are the foundation of the new theory are assumed to vary both across plants and over time, that is, fixed costs behave randomly.

Recall our example of installing new computers. Now, let's consider the installation of two new machines, on two separate occasions, at our hypothetical plant. For the first installation, managers may have available a very competent senior technician. He or she may be able to efficiently integrate the new machine into the plant's network. The cost of capital adjustment will be relatively small. However, at a later date, the senior technician may be unavailable, and managers may have to rely on a novice. This technician, new to the plant and unfamiliar with its computer systems, is likely to take far longer to install the new computer and will therefore incur a much larger adjustment cost. A simple way to introduce such variations into models of investment is to assume that adjustment costs are random.

What Do Random Adjustment Costs Mean for Aggregate Investment? If these costs differ randomly across plants and over time, then even two similar plants are likely to behave differently because they'll have different adjustment costs. Consider a world full of plants that all start out with the same level of capital. Over time, they'll face different adjustment costs, and thus, their capital adjustment behavior will differ.

The difference between a plant's actual and target capital stock will also differ across plants. Plants that had small fixed costs will have adjusted their capital stocks and be close to their targets. Plants that were less lucky and experienced several large adjustment costs in a row will have much larger capital imbalances. Generally, plant actions will not be synchronized. Plants with larger capital imbalances will see higher gains from adjusting capital, no matter what the adjustment cost; hence, they'll be more likely to undertake adjustments. Plants with low capital imbalances will not be willing to absorb even moderate adjustment costs and will be unlikely to adjust capital. At any time, someone studying the entire population of plants will find that some actively adjust their capital while others do not.

Plants with larger capital imbalances will see higher gains from adjusting capital, no matter what the adjustment cost

Changes in aggregate investment will arise for two reasons: changes in the level of investment undertaken by plants actively investing and changes in the number of these active plants. When there are many plants, small increases in productivity or demand that affect most plants will, generally, induce small changes in the number of plants actually investing. But by raising target capital a little, a few more plants will be induced to become active and adjust capital. As a result, while individual plants may exhibit lumpy investment, the number of plants investing will evolve more gradually, leading to slower changes in aggregate investment.

We see that the fixed cost model is able to preserve the success of the partial adjustment model in explaining changes in aggregate investment, while it improves the match with the microeconomic evidence on plant-level investment.[7]

SYNCHRONIZATION AND BUSINESS CYCLES

The fixed adjustment cost model and the partial adjustment model make different predictions about how investment should behave over the business cycle. While plants will typically not act together in the fixed adjustment cost model, at other times, plants will behave in a dramatically more synchronized manner in the model, mainly whenever there is a sharp change in some factor that affects all plants.

Economists agree that plants are subject to unforeseen events that can either increase or decrease their productivity. For example, a bank might be subject to new regulation, a farm might experience a drought, or a firm might adopt a new type of technology, for example, newer, faster computers.

The fixed adjustment cost model and the partial adjustment model make different predictions about how investment should behave over the business cycle.

Consider a large unforeseen rise in future productivity for all plants — what macroeconomists refer to as a large shock. Such a productivity shock, which might occur at the end of a recession, will yield a large change in the target capital of all plants. As a result, there will be few plants left with low capital imbalances, and most plants will adjust their capital. Their actions will, to a large extent, be synchronized. In their 1999 paper, Ricardo Caballero and Eduardo Engel show that such synchronization can lead to a sharp, unusual rise in aggregate investment.

The black series in Figure 3 represents the total investment of a group of plants when an extraordinary change in productivity results in a sudden synchronization of their investment. The blue series presents a hypothetical alternative case in which the number of plants allowed to adjust their capital is constrained to remain at ordinary levels. Notice the increased response in total investment due to the synchronization effect. Over the first 11 years, investment initially rises by a total of 17 percentage points more in the synchronized case. While this is partly offset by a steeper decline during years 12 – 14, the overall impact of synchronization is to raise investment spending by 4 percentage points.

This is the principal achievement of the new theory of investment. By allowing differences in capital imbalances across plants to evolve over the business cycle, the new investment theory allows the synchronization of investment activities during episodes involving large changes in the macroeconomy. It is through such episodes that the fixed cost model we have been examining overcomes the excessively low variability of investment in the partial adjustment model. The fixed cost model provides a considerably better match with both the aggregate and the plant-level data (these models are compared in Figure 4) and can explain the sharp increase in aggregate investment that follows a recession.[8]

CONCLUSION

The theory of investment has evolved into one that's now better able to explain the facts about investment at both the macro and micro levels. Traditional theory, known as the partial adjustment model, ignored differences across plants and firms. As a result, while it was reasonably

[7] In fact, the fixed cost model is actually better able to explain aggregate investment than the partial adjustment model because the partial adjustment model reduced the variability of investment too much. And while the fixed cost model typically behaves like the partial adjustment model, at other times it allows for much sharper changes in investment. This undoes much of the excess smoothness of the partial adjustment model.

[8] This is shown in the 1995 paper of Ricardo Caballero, Eduardo Engel, and John Haltiwanger and the 1999 paper of Russell Cooper, John Haltiwanger, and Laura Power.

successful at explaining aggregate investment, it did poorly at explaining lumpy plant-level investment. Newer theories that explicitly address plant-level investment resolve the problems not addressed by traditional theory. These new theories of investment emphasize the role of fixed costs of capital adjustment in inducing large but occasional plant-level investment. Moreover, once it was understood that these costs were likely to vary across plants and over time, the fixed cost theory has been able to explain not only plant-level investment but also aggregate investment. Indeed, by allowing for unusual synchronization of investment across plants, fixed cost theory is able to explain brisk recoveries following recessions, something traditional theory could not do.

Of course, even the new theory leaves something out. For example, recent work suggests that changes in interest rates, ignored in the new theory, may have powerful effects on firms' investment decisions.[9] Nevertheless, the new theory certainly represents progress — it provides an explanation of changes in aggregate investment that, in contrast to traditional theory, is consistent with our observations of plants' investment behavior.

[9] See the 2000 paper by Julia Thomas and the 2000 paper by both Julia Thomas and me.

FIGURE 3

Investment With and Without Synchronization

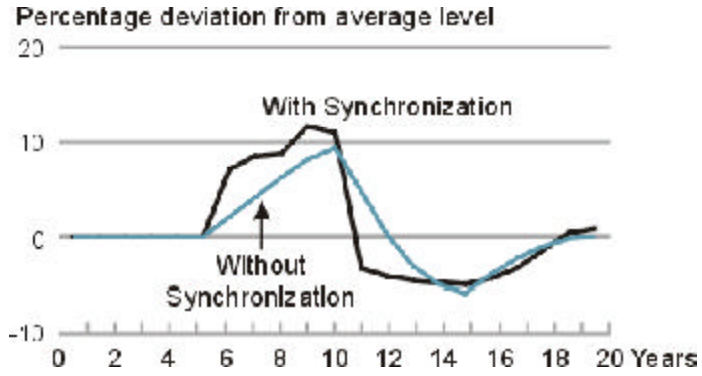

FIGURE 4

Investment Under Partial Adjustment and With Fixed Costs

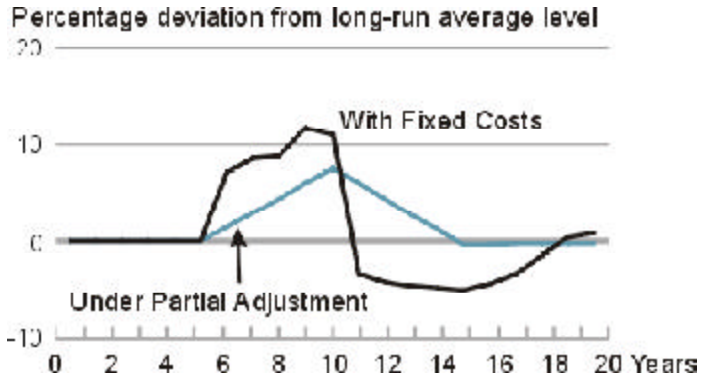

REFERENCES

Barro, Robert J. *Macroeconomics.* John Wiley & Sons, 1984.

Baxter, M., and R. G. King. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *Review of Economics and Statistics* 81, 1999, pp. 575-93.

Caballero, R. J. "Aggregate Investment," in M. Woodford and J. Taylor, eds., *Handbook of Macroeconomics.* Elsevier Science, 1999.

Caballero, R. J., and E. M. R. A. Engel. "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach," *Econometrica* 67, 1999, pp. 783-826.

Caballero, R. J., E. M. R. A. Engel, and J. C. Haltiwanger. "Plant-Level Adjustment and Aggregate Investment Dynamics," *Brookings Papers on Economic Activity* 2, 1995, pp. 1-54.

Cooper, R., J. Haltiwanger, and L. Power. "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review* 89, 1999, pp. 921–46.

Doms, M. and T. Dunne. "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics* 1, 1998, pp. 409-30.

Hammermesh, D., and G. A. Pfann. "Adjustment Costs in Factor Demand," *Journal of Economic Literature* 34, 1996, pp. 1264-92.

Khan, A., and J. Thomas. "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?" GSIA Working Paper No. 2000-E33, 2000.

Thomas, J. "Lumpy Investment, Partial Adjustment and the Business Cycle: A Reconciliation," GSIA Working Paper No. 1999-E250, 1999.