Investment Spikes:
New Facts and a General Equilibrium Exploration

François Gourio† and Anil K Kashyap††

PRELIMINARY Pre-Conference Draft

September 2006

Abstract: We analyze plant-level data from Chile and the U.S. to confirm and extend previous findings about the lumpiness of investment. Investment spikes are highly pro-cyclical, so much that changes in the number of firms undergoing investment spikes (the "extensive margin") account for the bulk of variation in aggregate investment. Moreover, the number of firms undergoing investment spikes has independent predictive power for aggregate investment, even controlling for past investment and sales. These facts suggest that investment lumpiness is related to the business cycle. We use the information on the importance of the extensive margin of adjustment to evaluate calibrated versions of a DSGE model of investment with fixed costs of adjusting capital due to Thomas (2002). We show that a modified version of the Thomas (2002) model can match the importance of the extensive margin. With our preferred calibration, however, the model yields fairly different implications for the dynamic properties of aggregate investment than the original Thomas model or the standard RBC model with quadratic adjustment costs. Thus, we conclude that while general equilibrium attenuates the effects of fixed costs, it does not eliminate them.

Key words: adjustment costs, investment, fixed costs and intensive and extensive adjustment.

† Department of Economics, Boston University, †† Graduate School of Business University of Chicago, Federal Reserve Bank of Chicago, and National Bureau of Economic Research. Kashyap thanks the Center for Research in Securities Prices and the Stigler Center both at the University of Chicago Graduate School of Business for research support. We thank Shawn D. Klimek from the US Census, and Olga Fuentes and Simon Gilchrist from Boston University for very kind help in providing data. We thank participants in the BU macro lunch, and in the GSB Macro and Finance lunches for helpful suggestions, Ricardo Caballero and Janice Eberly for helpful conversations, and Robert King for many constructive comments. We also thank Dorianna Ruffino, Sandra Shao, and Jin Xu for excellent research assistance. The data, a technical appendix and Matlab programs to run the model are posted on the following web page: http://people.bu.edu/fgourio/extintpaper.html. All errors here are our own.

1. Introduction

In this paper, we add to the ongoing debate over the aggregate significance of the "lumpiness" of plant-level investment. By "lumpiness" we mean two basic facts on which there is broad agreement. First, Caballero, Engel and Haltiwanger (1995), Doms and Dunne (1998), Cooper, Haltiwanger and Power (1999), Cooper and Haltiwanger (2002), and Becker et al (2006) all document that the typical manufacturing plant is likely to have at least one year when their capital stock surges. For instance, Doms and Dunne find that over half the plants that they study have a year in which their capital stock increases by at least 37 percent, while Cooper and Haltiwanger (2002) note that 18 percent of the observations in their data set covering 7,000 U.S. manufacturing plants over 16 years report investment rates of 20 percent (relative to capital) or higher. Following others in this literature, we refer to these high investment episodes as "spikes".

Second, these same studies also show that many establishments forgo investment in some years: the prevalence depends on whether one looks only at cases where exactly zero investment is reported or whether "near zeros" are counted too. For instance, Becker et al (2006) find that between 28 percent and 9 percent (depending on the year) of the plants in their sample have exactly zero total investment.

These facts have convinced economists studying plant-level data to abandon the traditional quadratic adjustment cost model because there is no reason why many firms would bunch at a zero investment, and spikes, which are heavily penalized, should be rare. In contrast, models with fixed costs to adjusting the capital stock can generate both facts.

Yet, there is little agreement as to whether modeling this lumpiness is necessary for understanding aggregate investment. For instance, Caballero (1999), in his survey for the Handbook of Macroeconomics, argues that it is, stating "it turns out the changes in the degree of coordination of lumpy actions play an important role in shaping the dynamic behavior of aggregate investment." On the other hand, Thomas (2002) argues that "in contrast to previous partial equilibrium analyses, [my] model results reveal that the aggregate effects of lumpy investment are negligible. In general equilibrium, households' preference for relatively smooth consumption profiles offsets changes in aggregate investment demand implied by the introduction of lumpy plant-level investment." This "irrelevance result" inspired Prescott (2003) to argue "partial equilibrium reasoning to an inherently general equilibrium question cannot be trusted."

We contribute to this debate in three ways. First, we introduce several new facts about the lumpiness of investment, considering both spikes and very low rates of investment. As Becker et al (2006) stress, (1) the fraction of firms with near-zero investment is countercyclical, and (2) the fraction of firms with spikes is strongly pro-cyclical. We build on these findings by noting the spikes are sufficiently important that most of the variation in the total investment rate is due to variation in investment of firms undergoing spikes.

We go on to show that this approximation derives its explanatory power from changes in the number of firms making large investments, and not changes in the average size of the spikes. Moreover, information on prevalence of spikes in one year has predictive power for forecasting aggregate investment in the next year (even controlling for the past level of investment or sales): years with relatively more spikes are followed by less investment in the subsequent years.

These patterns suggest that lumpy investment is relevant for the business cycle, and that a suitable model should generate not only the *average* lumpy behavior, but also *variation* in lumpy behavior over the business cycle. Our second contribution is to revisit the model developed by Thomas (2002) to study the determinants of the variation of the number of zeros and spikes over the business cycle. We use this model because (to our knowledge) it is the only tractable dynamic stochastic general equilibrium (DSGE) model which incorporates lumpy investment. We use this model as a laboratory to explore the determinants of the cyclicality of spikes and zeros. We find that the exact model she developed has trouble fitting the facts about cyclical patterns in lumpiness. But by changing the calibration we can match better these facts.

The third contribution is to examine the model's predictions for the effect of productivity shocks on investment. Thomas found that in her version of the model, the fixed costs that she introduced to generate spikes were essentially "irrelevant" for aggregate dynamics. In particular, she found the aggregate dynamics to be the same as the standard real business cycle (RBC) model, which has no adjustment costs of any kind. In our calibration, the qualitative response of investment to a productivity shock is quite different from the standard RBC model. So we find that although general equilibrium attenuates the differences between the fixed cost model and the RBC model (i.e. the model without fixed costs or any other adjustment costs), it does not eliminate these differences. In other words, the irrelevance result is not a generic finding that comes from the general equilibrium, but rather a result that depends on the details of how the model is calibrated, especially regarding the production side.

The remainder of the paper is organized into three sections, each corresponding to the three main contributions. Section 2 documents the extent to which lumpiness varies over the business cycle. The data we analyze comes from a census survey of establishments in manufacturing in the U.S. and a similar survey for Chilean plants. One comforting finding is that the results mentioned above hold for both data sets. We propose a decomposition of aggregate investment into an "extensive margin" (the number of firms with positive investment) and an "intensive margin" (the average investment per firm that invests). By varying the threshold for positive investment to consider all non-zero investment, as well as large spikes we are able to provide a concise summary of the cross-sectional patterns in the data. In particular, these decompositions prove useful in calibrating the model that we study in the remainder of the paper.

In Section 3, we introduce the aforementioned Thomas (2002) model and compare that model to the patterns documented in section 2. The Thomas model as originally proposed includes a fixed cost for firms that invest. But these fixed costs are extremely (and in our view implausibly) low. Once we require firms to spend a non-trivial amount to adjust their capital it appears that three other critical changes are needed to bring the model closer to the data.

First, Thomas, following Caballero and Engel (1999), used a random fixed cost drawn from a uniform distribution, so that there is considerable variation in the actual costs firms wind up paying. In our specification, the distribution is more compressed, and the fixed cost is less random. This reduction in the heterogeneity of fixed costs across firms proves important to match the fact that the extensive margin dominates for spikes. We discuss its interpretation in terms of idiosyncratic shocks or heterogeneity.

Secondly, we also introduce machine failures in the model. Without breakdowns, a spike is effectively guaranteed to be followed by many firms choosing not to invest. Therefore the zero and the spikes have virtually the same correlations (with some lag). These exogenous breakdowns force some firms to invest sooner than normal and partially breaks the otherwise tight connection between zeros and spikes that would otherwise be present in the model (but not in the data).

Finally, several recent papers, e.g. Cooper and Haltiwanger (2005), Fuentes, Gilchrist and Rysman (2006) and Hennessy and Whited (2005), estimate that curvature of the profit function is greater than Thomas had presumed.[1] We use these estimates to guide our calibration and find that doing so improves the model's fit by increasing the benefit to adjusting and thus penalizing inaction.

In section 4, we study the augmented model's predictions regarding the response of aggregate investment with respect to aggregate productivity shocks. With our preferred calibration the model no longer generates the "irrelevance result" of Thomas. The model's impulse responses exhibit the kind of "echo effects" (i.e., non-monotonicity) stressed in the partial equilibrium literature (e.g., Caballero and Engel (1999)). Another noticeable effect of fixed costs is also (in general) simply to make aggregate investment smoother. These effects are starker after a succession of shocks in the same direction, or when a shock reshapes directly the cross-sectional distribution of firms' capital so that it departs noticeably from the steady-state distribution. Importantly, the smoother responses that we do recover from our model are not very similar to those that would emerge from an RBC model with quadratic adjustment costs.

Our concluding section briefly summarizes and suggests a couple of directions for future work.

[1] In the Thomas set up all the curvature comes because of decreasing returns to scale, but imperfect competition in the product market would also imply curvature.

2. Empirical Evidence on Lumpiness over the Business Cycle

To analyze lumpiness we study two establishment-level data sets covering manufacturing plants in Chile and the U.S.[2] We briefly describe the data construction in the data appendix and in what follows we will refer to the U.S. data as the census sample and the other sample as the Chilean data. We start by reviewing three measurement issues before reporting our main results.

The first issue is how to handle very small rates of investment, for example where investment is not exactly zero, but less than one or two percent of capital. If fixed costs of investing are present, then we would expect to find few cases of this sort. Yet, it appears empirically many plants report making these tiny investments. We suspect that these cases represent some sort of maintenance or replacement investment (for which the fixed cost presumably does not apply). So in what follows, we will typically aggregate the plants with near zero investment with those that report exactly zero investment.

On the other side, there is no clear definition of what constitutes an investment spike. The papers by Cooper, Haltiwanger and Power (1999), Cooper and Haltiwanger (2005), and Becker et al (2006) all define spikes to be cases where investment relative to the beginning of period capital is greater than 20 percent. To maintain comparability with these papers we use this threshold as a primary definition. But we will also note the results for the case where we set the threshold to be 35 percent; the results are very similar, and we note the few cases where we find a difference.

The last conceptual issue that arises relates to aggregation. To summarize the distribution of firms or establishments we must take a stand on whether each observation will be equally weighted or weighed by some other characteristic such as the plant's capital. We see equal-weighting as problematic (or at least inferior to capital-weighting) for several reasons.

First, at a sufficiently fine level of aggregation every decision is lumpy; no one disputes that integer constraints and the like are relevant for truly tiny firms.[3] Conversely for the entire economy there are never any zeros and spikes are rare. So as firm sizes vary, so will all measures of lumpiness. This means that as the size distribution of firms in a sample changes, either because of changes in the underlying population or because of changes in the sample coverage, the statistics on zeros and spikes will change. We would like our measurement to reflect more than just the mechanical effects that derive from the composition of the sample.

Conversely, one way to partially offset the attenuation of the zeros and spikes that results from aggregation is to weight the data based on firm size. Loosely speaking by giving firms with large capital stocks more weight when they report a zero or a spike we make

[2] See Hsieh and Parker (2006) for a concise summary of macroeconomic developments in Chile over our sample period.

[3] The model that we introduce below, like most models, abstracts from differences in size. So even if we wanted to account for this effect it would be difficult.

up for the zeros or spikes that might be occurring within some of these organizations that are otherwise obscured. As a bonus, a capital-weighted average of investment rates delivers a measure that equals the aggregate investment in the sample divided by the total capital in the sample.

A final consideration comes from our specific interest in the general equilibrium effects of lumpiness. Intuitively we expect general equilibrium effects that operate through prices to depend on aggregate indicators of lumpiness. This suggests another potential reason that actions of larger firms (or establishments) are more important than smaller firms. For all these reasons we will emphasize our findings that pertain to the capital-weighted data.

Figures 1 and 2 show the effects of varying the definitions of spikes and zeros, and of changing the aggregation schemes. In each figure we report four panels; the two panels on the left show the time series patterns for the prevalence of zero (dashed lines) and near zero investment (defined to include establishments with I/K less than two percent). The bottom panel shows the data when the observations are aggregated according to the capital stock for each establishment, while the top panel treats all plants identically. The right hand panels graph spikes, with the solid lines showing the percentages based on I/K greater than 20 percent and the dashed lines showing the 35 percent spikes; again the top and bottom differ based on the weighing scheme. Figure 1 gives the results for the census sample, while Figure 2 shows the data for our Chilean sample.

For both the U.S. and Chilean samples it is clear that many establishments are not investing in any given year, whereas at the same time there are other establishments where investment is spiking. The full distribution of the investment rates for each sample is shown in Table 1. Fuentes, Gilchrist and Rysman (2006) stress the fact that emerging markets such as Chile tend to have more plants that are not investing than in developed economies such as the U.S. Comparing the dashed lines in the upper left panel in each figure shows that the (unweighted) percentage of establishments with exactly zero investment is two to three times higher in the Chilean sample.

The figures also show that the level of zeros is sensitive to the weighting schemes used. As would be expected, fewer large firms report literally zero investment, so the reported percentages of zeros drops precipitously in the capital-weighted figures compared to the equally-weighted figures. The capital weighting makes less of a difference for level estimates for the near zeros and even less difference for the spikes.[4] For the rest of this section, we concentrate on capital weighted series.

The figures show that regardless of the weighting scheme that is used, the 20 and 35 percent thresholds for spikes are extremely highly correlated. For example, the correlation between the capital weighted series for census data spikes of 20 percent and 35 percent (the lower right panel in Figure 1) is 0.95. For the remainder of the section,

[4] The pairwise correlations for the weighted and capital weighted series for Chile (US) are as follows: exact zeros = 0.82 (0.60), near zeros = 0.94 (0.23), 20 percent spikes = 0.92 (0.90), and 35 percent spikes = 0.90 (0.87).

we concentrate on the 20 percent spikes and mention the 35 percent spikes only cases where there are differences.

One final noteworthy feature of Figures 1 and 2 is that several of the lumpiness proxies show trends; below when we look at the aggregate investment rates we will also find trends for those series too. These low frequency changes are outside of the scope of our investigation and in most of our analysis we will remove them by regressing the series on a linear time trend (although using a Hodrick-Prescott filter delivers very similar results for all of our findings).[5]

The general facts that we have mentioned thus far about the prevalence of zeros and spikes have been documented in a number of other studies (including all of the ones mentioned in the opening paragraph of the paper.) Some of these studies also describe the cyclicality of the lumpiness. Figure 3 shows the (de-trended) capital weighted shares of establishments with either spikes or with near zeros, along with the (de-trended) aggregate investment rate for each sample; the aggregate rate is calculated by taking the capital weighted average of the establishment level rates and we denote this as Itot/K. (The weighting scheme also means that Itot/K is the ratio of aggregate investment to aggregate capital in our sample.) In each country, the spikes are strongly pro-cyclical and near-zeros are strongly counter-cyclical. The correlation between the capital-weighted spikes and the aggregate investment rate (both detrended) is 0.87 for the US sample and 0.96 for the Chile sample; and the correlation between the capital-weighted near zeros and the aggregate investment rate (both detrended) is -0.94 for the U.S. sample and -0.56 for the Chilean sample. Thus, Figures 1, 2 and 3 show all the standard characteristics of plant-level investment.

In the remainder of this section we document several new facts regarding spikes. The first of these facts is documented in Figure 4. Each panel in the chart shows a pair of aggregate investment rates. The solid lines show Itot/K, the aggregate investment rates from our samples. The lines with circles show the total investment done by those establishments where investment is large (i.e. $I/K > 20$ percent), divided by the total stock of capital for all the firms in the sample; we label this series I20/K.

The relative levels of I20/K and Itot/K indicate that the spikes account for about half of total investment in each country; in other words, I20/Itot is about 0.5. More importantly, the investment rate constructed for the spiking firms tracks the movements in the aggregate investment rate closely; the correlations between the de-trended series is 0.99 for each sample. Clearly, the bulk of the variation in the aggregate Itot/K is accounted for by changes in I20/K. The share of variance of Itot/K accounted to by I20/K (as opposed to the residual (investment of firms with investment rates between 0 and 20 percent over total capital)) is 97 percent for the US sample and 86 percent for the Chile sample.[6] The converse of these observations is that there is little variation in total investment explained by the firms investing between zero and 20 percent. The investment for these firms is

[5] In future work we will investigate the role of changes in average firm size and in industry composition on these trends.
[6] This is measured as Cov (I20/K,I/K) / Var (I/K).

labeled I(0-20)/K in the graph. Thus, for the purposes of modeling investment fluctuations it is critical to understand the timing of the investment spikes.

To go further and better describe the spikes we start from the following identity:

$$\frac{I20}{K} \equiv \frac{I20}{K20} \bullet \frac{K20}{K} \quad \equiv \quad IPA20 \bullet ADJ20 \qquad (1)$$

$$\rightarrow Log(\frac{I20}{K}) \equiv \quad \log(IPA20) + \log(ADJ20)$$

where $\mathrm{I20} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}} > 0.20} I_{i,t}$, $\quad \mathrm{K20} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}} > 0.20} K_{i,t-1}$, $\quad \mathrm{K} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}} \geq 0} K_{i,t-1}$

In words, equation (1) simply says that the total investment done by the plants experiencing spikes can vary either because of a change in the investment per adjuster (IPA20, the intensive margin) or because of a change in the (capital-weighted) number of firms adjusting (the extensive margin). This approach is analogous to the one proposed by Klenow and Kryvstov (2005) for studying price dynamics, where they decompose inflation into changes in the number of firms resetting their prices and changes in the average size of changes for those firms resetting their price.

Figure 5 shows a graph of Log(I20/K), along with Log(IPA20) and Log(ADJ20) (after each series has had a linear time trend removed) for the census and Chilean sample. The striking conclusion is that the extensive margin, ADJ20, drives variation in spikes.

One way to conveniently summarize the information in the picture is to compute the following pair of statistics:

$$\mathrm{ShareADJ20} \equiv \frac{\mathrm{covariance}(\log(\mathrm{ADJ20}),\log(\frac{\mathrm{I20}}{\mathrm{K}}))}{\mathrm{variance}(\log(\frac{\mathrm{I20}}{\mathrm{K}}))} \text{ and } \mathrm{ShareIPA20} \equiv \frac{\mathrm{covariance}(\log(\mathrm{IPA20}),\log(\frac{\mathrm{I20}}{\mathrm{K}}))}{\mathrm{variance}(\log(\frac{\mathrm{I20}}{\mathrm{K}}))}$$

These shares (by construction) must sum to one. If the proportion of firms with spikes ADJ20 is constant, they would be zero and one, and if the average investment rate of firms with spikes is constant, they would be one and zero. For the census sample ShareADJ20 is 0.87, while for the Chilean sample it is 0.925. The dominant role of the extensive margin also appears when the threshold for identifying spikes is 35 percent (instead of 20), and for other de-trending procedures.[7]

[7] The only case where the ShareADJ is not above 0.85 is the equal weighted data are considered for Chile, when it drops to 0.5.

There is nothing mechanical that guarantees that the extensive margin has to account for the bulk of movements in the investment rate. Figure 6 shows Log(I2/K), Log(IPA2), and log(ADJ2) (where these variables are defined as in equation 1 except that the thresholds for what counts as a "spike" is all investment over two percent.) These pictures are interesting because they help gauge the importance the variation in the number of exact zeros and near zeros in aggregate investment fluctuations.

In contrast to the patterns of Figure 5, the intensive margin is now much more important. ShareIPA2 is 0.665 for the census sample and 0.605 for the Chilean sample. The dominant role of IPA for low investment thresholds is also a recurring pattern that is insensitive to the procedure for de-trending.[8]

Our last fact about spikes is to note that they seem to contain additional predictive content beyond just information that they convey about the past level of investment. The spirit of many models of lumpiness (e.g. Caballero and Engel (1999)) is that the cross-sectional distribution of firms' capital stock relative to the level that would prevail absent any adjustment costs should be an important determinant of aggregate investment. It is empirically difficult to construct this cross-sectional distribution, but there is a simple way to test for this possibility. We estimate regressions of the form:

$$\frac{Itot_t}{K_{t-1}} = \alpha + \beta Trend_t + \gamma \frac{Itot_{t-1}}{K_{t-2}} + \phi \frac{Sales_{t-1}}{\mathrm{K}_{t-2}} + \sum_{h=1}^{H} \omega_h ShareADJ20_{t-h} \qquad (2)$$

The novelty is that we add the share of adjusters to an otherwise standard accelerator type investment equation.[9] This type of accelerator style equation has repeatedly been shown to be an effective forecasting equation in horse-races of different specifications (Bernanke, Bohn and Reiss (1988) and Oliner, Rudebusch and Sichel (1995).)

Table 2 shows estimates of equation (2). The first six rows show the estimates for the U.S. data, while the last six rows show the estimates for the Chilean sample. For the census data the lagged dependent variable is always estimated to have a positive and highly significant coefficient. The sales proxy is positively related to investment, but not always significant. Conversely in the Chilean sample the sales variable is always estimated to have a positive and very significant effect on investment, but the lagged dependent variable does not systematically influence investment.

Our main coefficients of interest are the ω's that measure the effects of past spikes on current investment. For the census sample, the coefficients on both the first and second lags of ShareADJ20 are significant, whereas in the Chilean data, only the second lag is consistently significant.[10] Importantly, the estimated signs of the ω's are all *negative*, suggesting that investment is depressed in the period after an investment surge – a kind of

[8] When the threshold is set to zero the IPAshares are generally above 0.90 (with the exceptions coming for the equal weighted Chilean sample).
[9] For the census sample, we have shipments data which correspond to sales for establishment data.
[10] When the spikes are measured with the 35 percent threshold then both lags one and two are significant in both samples.

"echo" effect. This correlation is to be expected based on fixed costs models (and would be of the opposite sign if the past ShareADJ20 variable was standing in for productivity shocks or other factors that raise investment demand).

Taken literally, the coefficients suggest that the echoes from the spikes have a quantitatively important effect on investment. For the Census sample (Chile) the standard deviation of the spike variable is 0.046 (0.093), compared to the standard deviation of the investment rate of 0.017 (0.054). Taking the specifications where h=1, (shown in rows 5 and 11), the estimates for the census (Chile) sample imply that a one standard deviation increase in ShareADJ20 predicts an increase of the investment rate of 0.7 (0.57) of a standard deviation.

Collectively, we read our findings as implying six important facts that we think models should replicate. The first two are simply that there are both many small and many large rates of investment at a given point in time; so we will compute the percentage of near zeros and spikes for the models that we study and compare them to the averages in Table 1. Moreover, the share of spikes is highly procyclical and the share of zeros is countercyclical. The fourth fact is that aggregate investment is largely driven by investment spikes; so the model should have the property that I20/Itot is substantial and that variations of I/K are accounted for by variation in I20/K. Fifth, the spikes matter because of adjustment along the extensive margin, i.e. a change in the number of firms making large investments; these spikes are sufficiently important that they have independent predictive power for aggregate investment, even controlling for past investment and sales. We quantify this by looking at the model's predictions for ShareADJ20 and seeing if it is large. Finally, if one looks at overall investment done by all plants with non-zero investment, then that quantity fluctuates more because of changes in the amount of investment per plant (rather than changes in the number of plants investing). Hence, ShareADJ2 should be much lower than ShareADJ20. Overall, these facts suggest to us that lumpy investment is important, and that it is related to the business cycles.[11]

3. A DSGE model with fixed costs of adjusting capital

A. A brief review of the Thomas model

Thomas (2002) offers an elegant and compact model for analyzing the importance of fixed costs of adjusting capital on aggregate investment in a dynamic, stochastic general equilibrium model.[12] Using her notation, we begin with a review of the key ingredients of the model.

The economy has a fixed number of plants (normalized to be of measure one). In what follows, we refer to these as "plants" or "firms" interchangeably. Each plant has the production function: $y=Ak^{\psi}n^{\upsilon}$, where y is output, A is an aggregate productivity shock,

[11] We agree however that there may be competing explanations. For instance, some of our facts could be rationalized if the cross-sectional distribution of idiosyncratic shocks exhibited more variance in upturns.

[12] The setup is similar to the sticky price model of Dotsey, King and Wolman (1999).

k is capital, and n is labor. There are decreasing returns to scale so that $\psi + \upsilon < 1$ and there is no entry or exit.

Each period, each plant has the opportunity to adjust its factor usage. Labor can be freely varied, but adjusting capital can only be done if the firm pays a fixed cost. The fixed cost, ξ, is a random variable that is independently and identically distributed across time and plants and comes from the cumulative distribution G. This distribution has finite support and the maximum fixed cost that any firm would ever have to pay is called B. The firms that choose to pay the fixed cost, which we call "adjusters", bear no marginal adjustment costs: they can buy or sell capital at price 1. The fixed cost is measured in units of labor. Owing to the fixed cost, firms will not always adjust capital.

Much of the model's tractability derives from its inherent symmetry that leads all firms choosing to invest at a given point to pick the same new level of capital, $k_{0,t+1}$; since there is no heterogeneity except in the fixed cost drawn today and the current capital, all firms choose the same new level of capital, conditional on investing. So firms are distinguished by the time since their last investment. Regardless of whether a firm invests, its capital depreciates at rate δ. Therefore,

$$\mathrm{k}_{0,t+1} = (1-\delta)\mathrm{k}_{j,t} + \mathrm{i}_{j,t} \quad \text{and} \quad \mathrm{k}_{j+1,t+1} = (1-\delta)\mathrm{k}_{j,t}\,.$$

A firm that last adjusted capital j periods ago, henceforth a vintage j firm, will operate with capital k_j (and labor n_j). This implies the following maximization problem for a plant:

$$\max_{i_{jt},n_{jt}} \mathrm{E}_0\left(\sum_{t\geq 0} m_t (A_t k_{jt}^{\psi} n_{jt}^{\upsilon} - w_t n_{jt} - i_{jt} - \xi_t w_t 1_{i_{jt}\neq 0})\right)$$

subject to the capital accumulation laws above, where m_t is the stochastic discount factor (the ratio of marginal utilities in period t to period 0).

The TFP process, A_t, evolves according a first-order autoregressive process around a deterministic trend:

$$A_t = \Theta_A^t z_t\,,\ \ \log\ \mathrm{z}_t = \rho \log z_{t-1} + \varepsilon_t\,,\ \ \varepsilon_t \text{ is distributed independently } N(o,\sigma^2).$$

The combination of the fixed depreciation rate and the finite upper bound on the fixed cost guarantees that all firms will eventually find it optimal to invest; in other words, this structure delivers a maximum vintage J by which time all firms will invest. The solution to the problem involves finding that maximum vintage (J), along with the capital stock for each of the intervening vintages (k_j), and the percentage of total firms in each vintage (θ_j).

Thomas shows that firm's investment rules are such that for firms of a given vintage the decision of whether to invest and upgrade their capital is determined by a cutoff rule related to the fixed cost that they face. A proportion α_j will draw sufficiently low fixed costs that given the economy-wide wage it pays to adjust and the others will wait. In her simulations she chooses the uniform distribution function for the fixed costs, so these fixed costs are uniformly distributed between 0 and B. The level of fixed costs B is chosen to match two facts reported by Doms and Dunne (1998): i) in the average year, 8 percent of plants raise their real capital stocks by 30 percent or more; ii) these plants account for 25 percent of aggregate investment.

The rest of the model is intentionally chosen to follow the real business cycle (RBC) literature. So, for instance, Thomas adopts a utility function with indivisible labor of the form $U_t = \log c_t - \zeta n_t$. Thus, aside from the fixed costs and the mild decreasing returns, the calibrated parameters she uses are very standard.[13] (These parameter values are displayed in Table 3 below.) Indeed, when the upper bound of fixed costs, B, is set to 0, all firms adjust their capital each period, and equate their marginal product of capital and labor; in this case, there is a representative firm, and the model collapses to a standard RBC model with decreasing return to scale.

This model is solved numerically by a standard log-linearization around the steady-state. First, one finds the optimal J, the maximum time-since-last-adjustment such that all firms want to invest. Second, one solves the system of non-linear equations that define the non-stochastic steady-state. Finally, one computes the log-linear approximation itself. The log-linear method is advantageous here since the state space of the model is large: it includes the TFP shock, and the cross-sectional distribution of capital (the θ_j's and the k_j's).[14]

B. The determinants of the extensive-intensive decomposition

Given our interest in matching the intensive and extensive investment margins for different thresholds, it is instructive to briefly review the factors that govern this split in the model. To gain some intuition, consider the response to a positive technology shock in this model. An increase in TFP increases the marginal product of capital leading firms to want to accumulate more capital. As a result, more firms find it worth paying the fixed cost. This increases the number of firms "adjusting". (Of course, both spikes and small investments become more likely.) Moreover, a second margin is present: the capital stock of firms that do adjust, denoted by $k_{0,t+1}$, increases too. This means that the typical investment per adjuster will increase. Hence the model will generate a mix of extensive and intensive margins: both the number of adjusters and the investment per adjuster will typically be procyclical.

---

[13] Also, the model is calibrated to annual rather than quarterly data, because the plant-level evidence is based on annual surveys.

[14] For more details on the solution, we refer the reader to our separate technical appendix (available on http://people.bu.edu/fgourio). Khan and Thomas (2003) examine whether a nonlinear solution method gives different results than the usual log-linear approximation, and find that while some nonlinearities arise in partial equilibrium, they are not present in general equilibrium. In light of our results of Section 4, it would be interesting to solve the model with a nonlinear method as they do, and test for the presence of non-linearities for the GE version of our calibration.

Our empirical work in Section 2 examined how the *variation* of aggregate investment was split between the "extensive" and "intensive" margins. To build some intuition for the results, it is important to distinguish this from the decomposition of *average* investment between the two margins. In the model, this is the distinction between fluctuations and nonstochastic steady-state.[15]

Regarding the steady-state (average) decomposition, the magnitude of the fixed cost is the key determinant of the mix between the intensive and extensive margins: for instance, if the fixed cost is small enough, firms will find it worth to adjust their capital in every period. This yields the RBC model since there is exact aggregation of the production functions in this case. On the other hand, if the fixed cost is large, there will be a long inaction between any two adjustments at the firm level. This implies that firms will have a capital stock that is very different from the (static) optimal level most of the time: initially too high, it will become too low after a while. Clearly then, the main deterrent of long inaction is the curvature of the profit function (which in this model comes from the decreasing returns to scale but could also have been introduced by assuming monopolistic competition in the product market). To summarize, the steady-state solution will optimally balance the cost of having more firms adjusting in any given period – i.e., having more fixed costs to pay - with the inefficiency of having the size of adjusting firms $k_{0,t+1}$ depart too far from the level that would be chosen if capital could be freely continuously adjusted.

This steady-state tradeoff between extensive and intensive margins is intuitive and well understood. In this paper we are interested in how this trade-off affect is reflected in the business cycle fluctuations. We now use the model to see if it can replicate the facts from section 2 that summarize the lumpiness of the plant-level investment. The first two rows of Table 4 show the percentages of near zero investment rates (i.e. $I/K < 0.02$) and investment spikes ($I/K > 0.20$), along with the mean of I20/Itot, the percentage of the variance of I/K due to I20/K and ShareADJ2 and ShareADJ20 for our two samples.[16] The third row in the table shows the analogous statistics for the Thomas model.[17]

We find that for the Thomas calibration, ShareADJ2 (which measures the importance of the extensive margin for all non-zero investment) is 61.1 percent, while ShareADJ20 (which measures the extensive margin for investment spikes) is 30.6 percent. This is at odds with our empirical findings, which pointed to the opposite pattern where the extensive margin was dominant for the spikes but not for all investment. The model has too many zeros, which can be attributed to the absence of maintenance in the model. We also find that the share of variance of I/K accounted for by I20/K ($\text{cov}(I20/K, Itot/K)/\text{var}(Itot/K)$) is too low in Thomas model. Clearly there is a

---

[15] By "nonstochastic" we mean the model without aggregate productivity shocks; we maintain the random idiosyncratic shocks to the level of the fixed cost.

[16] Note that in the model, I/K>2 percent is equivalent to I/K>0. Also, we do not display the correlation between the share of zeros or spikes and aggregate investment, but they are high in all the models.

[17] We do not detrend the model output since it has only cyclical variation in the number of adjusters, and no trend. The results are in general insensitive to this.

large response by firms which are doing small investment: more firms choose to do small investment, and the average size of their investment also increases, so that both margins make I(0-20)/K react significantly to a shock. The model matches the data well in terms of share of spikes, and predicts that I20/Itot is very high.

We view the model as having three fundamentally free parameters: the curvature of the profit function, the level of fixed costs, and the shape of the adjustment cost function. We start by exploring how varying these characteristics affects the model's ability to match our five moments; in doing so we move one characteristic at a time, keeping the remaining parameters fixed at the baseline values in Table 3 (unless otherwise mentioned). We also extend the model to allow for breakdowns, which generate small positive investment rates, a robust feature of the data. Finally, we consider the effect of varying the persistence of aggregate shocks which helps to show further how the model operates. Looking across these experiments we propose a new calibration of this DSGE model which replicates better the extensive-intensive decomposition, i.e. the dominant role of spikes.

The effect of the level of fixed costs

It is natural to ask how the Thomas model changes as the level of the fixed cost changes. While there is substantial debate about the magnitude of adjustment costs, it seems clear that the calibration that Thomas chose has very low fixed costs. There are several ways to measure them in the model. The share of adjustment costs in total investment is 0.21 percent, and this number is 0.38 percent on average for the plants which are just indifferent between adjusting and not (i.e., plants which are the ones paying the highest adjustment costs). This cost seems small on an anecdotal basis, if we think of the costs of the planning, budgeting, and committee work that accompany most investments. There are obvious cases when adjustment costs are much larger: think of the disorganization of a factory floor, or the temporary closure of a retail store.

One recent study that computes adjustment costs is by Cooper and Haltiwanger (2005). They study a host of specifications that include convex and non-convex adjustment costs, including fixed costs, quadratic costs, gaps between the buying and selling price of capital, and productivity distortions created by capital adjustment. Using the Census data, they find statistically significant costs of each type, either when estimated in isolation or when several costs are simultaneously present. The total implied adjustment costs in this model and all the others (e.g. the one including just fixed costs) are substantial. For instance, their preferred estimates suggest that profits are reduced 20 percent during investment spikes. They simulate the model and find that on average spending on adjustment costs is equal to 0.91 percent of capital. Given that investment for their sample is about 12.2 percent of capital, this implies that adjustment costs average roughly 7.5 percent of investment; in other words, they find adjustment costs roughly 20 times the size assumed by Thomas. Abel and Eberly (2002) in their study of listed firms find a similar magnitude of adjustment costs (between 1.1 and 9.7 percent of investment).

Why, then, does Thomas choose such a small number? She calibrates the uniform distribution of fixed costs to match the Doms-Dunn statistics. However, since the model abstracts from any heterogeneity, the benefits to adjusting are small: since firms have the same productivity, it makes little sense to pay the fixed cost unless the capital is very much out of line. Hence, she naturally obtains small fixed costs to equilibrate these small benefits of adjustment and generate realistic lumpiness statistics. In our opinion, a better description of the data is that costs and benefits to adjustment are both large.

When we increase the maximum fixed cost, B, to 0.02 firms adjust less often in steady-state, so that the maximum vintage J rises to 20 years. (However, most of the distribution has a far smaller time-since-last-adjustment.) The number of zeros rises, and the number of spikes falls; almost all investment is done in spikes now as small deviations are not costly enough to justify paying the large fixed cost. The share of adjusters increases now for both the low and high thresholds. More generally, we found by varying only B it was very difficult to raise the ShareADJ20 above 60 percent in a robust way.[18]

The effect of curvature

In this model, the profit function is not a linear function of the level of capital because firms operate with a decreasing returns to scale technology; the curvature of the profit function could also have been motivated by assuming imperfect competition in the product market. Regardless of its origin, the extent of curvature is important because it determines costs that firms bear from not continuously adjusting the level of capital.

Subsequent to Thomas' paper a large empirical literature has estimated this curvature to be between 0.5 and 0.7, markedly lower than one (see e.g., Cooper and Haltiwanger (2005), Fuentes, Gilchrist and Rysman (2006), and Hennessy and Whited (2005)). Unsurprisingly, when we lower return to scales to 0.6, we obtain that firms adjust more often: the number of firms with zero investment falls to 67.4 percent from 72.5 percent in Thomas' calibration, and J falls from 5 to 4. Hence, there is an intuitive connection between the returns to scale and the steady-state mix between the intensive and extensive margins.

What about the business cycle decomposition? We see that the share of adjusters becomes smaller for both thresholds. While lower return to scale increase the benefit to

---

[18] One might have expected that a higher fixed cost would lead firms to substitute intensive for extensive investment, leading the share of adjusters to fall. However, there seems to be another effect. Our interpretation (related to Klenow and Kryvstov (2005)) is as follows: when there is a higher B, the maximum vintage rises, so that a firm will typically wait longer to adjust. This makes the incentive to adjust in response to a shock larger than in the low B economy. In the low B economy, no matter if you choose to change your adjustment threshold in response to a shock, the probability of adjusting in the next two or three years is very large. In contrast, this probability is very small with high B. Hence the incentive to change the timing of adjustment and adjust precisely when there is a good positive shock makes the number of adjusters vary more. Dotsey, King, and Wolman (1999) and Klenow and Kryvstov (2005) mention similar "counterintuitive" results when discussing the effect of trend inflation rate on business cycle dynamics. This also seems similar to Caballero's "Fallacy of composition" (1992). Because of these two opposing effects we even found cases where. depending on the setting of the other parameters, raising B lowered ShareADJ20.

adjusting and thus lead to more adjustment, it has also an indirect effect by changing the steady-state distribution and lowering J, thus making firms more likely to wait in response to a shock. We found that these comparative statics were thus generally not very robust, but in any event changing only the curvature of the production function did not help the model get closer to the empirical estimates of ShareADJ20 .

The effect of the shape of the distribution of fixed costs

Given the uncertain effect of the level of fixed costs or the curvature on the extensive/intensive decomposition, how, then, can we make the model match the importance of extensive adjustment which was so striking in Section 2? We find that the only feature of the model which affects robustly the decomposition is the shape of the distribution from which the random fixed costs are drawn. The distribution chosen by many authors (e.g. Caballero and Engel (1999) or Thomas (2002)) in their baseline calibration is uniform. This implies that there is a lot of randomness in the fixed cost that firms draw. But there are several other interesting possibilities besides the uniform, and as we show the choice of the CDF can be economically meaningful.

We first experiment by making this distribution more convex. Figure 7 shows the CDF we used, and clearly making this distribution more convex is akin to making the fixed cost less random.[19] In the limit, if $G(x) = 0$ for $x<B$ and $G(B)=1$, every plants draws B in every period. This is the standard fixed cost model without randomness. Rows 7 and 8 of Table 4 compare the results for a convex and uniform distribution of fixed costs with the same mean.

When the fixed cost is less random, we find that firms wait less before investing, because the option value of waiting for a low fixed cost diminishes. As a result, J decreases and the number of people with zero investment rises. The convexity leads the share of adjusters to rise (for both the 2 percent and 20 percent threshold). The interpretation is straightforward. To have more plants investing requires shifting marginal plants from inaction to action. The marginal cost of doing this depends on the shape of the CDF. When the CDF is flat, increasing the number of plants investing runs quickly into higher fixed costs (there is a lot of heterogeneity at the margin in terms of fixed cost). When the CDF is curved, increasing the number of plants investing is not very costly, because the inactive plants draw nearly identical fixed costs at the margin. With the convex CDF, most firms at the margin have the same fixed cost, which is close to B. This implies that increasing the number of firms investing is cheap and consequently the extensive margin dominates. Since most of the investment is done by firms which are in the curved part of the CDF, the extensive margin dominates also for the low threshold.[20]

---

[19] The exact function form that we choose (following Dotsey, King, and Wolman (1999)) is $G(x) = H(x/B)$ with $H(x) = c1 + c2*\tan(c3*x-c4)$ where c1 and c4 are chosen so that $H(0) = 0$ and $H(1) = 1$ and c2 and c3 govern the shape of the distribution. We set $c2 = 0.02$ and $c3 = 1.4$.

[20] Note that for extreme cases of convexity, the extensive margin may become less important. This is because in this case, almost only plants from the last vintage adjust; and by definition the plants of the last vintage all adjust in every period, so the number of adjusting plants is constant. Interestingly, we also found that when the distribution is very convex, comparative statics generally accord with the standard intuition: an increase in B reduces the share of adjusters. Another important difference is that when the distribution of

The choice of the convex CDF thus implies that there are a lot of nearly-identical firms which are close to investing, i.e. there is "marginal homogeneity". In this sense, the model becomes closer to the first generation of Ss models [21] rather than the ones studied by Caballero and Engel (1999) and Thomas (2002). Because the adjustment costs are the only source of heterogeneity in the model, there are a variety of ways to interpret what it means to have the model set up with this property. It could be that adjustment cost shocks are large, but that there are a lot of firms with nearly identical fixed costs which are almost indifferent to adjusting (i.e. there is "marginal homogeneity"): this depends on the precise shape of the distribution of fixed costs. So long as there is some part of the CDF that is concentrated the marginal homogeneity condition will be satisfied. In our set up, we have single mode, but a bimodal distribution (and many other distributions) would also have the same property. Exploring the impact of other distributions is certainly a topic that merits further work, but it seems clear that if the extensive margin is going to drive the variation in investment spikes the model needs this property.

In contrast to the convex CDF, we also consider the possibility that a fraction d of all plants draw a zero fixed cost. If they draw this zero fixed cost, it is always optimal to adjust. We see this element as capturing large idiosyncratic shocks which dwarf the other variables determining the benefits and costs of adjustment. The formula for G is now G(x) = d + (1-d)*x/B. When we increase d starting from zero, we find that the share of adjusters in either decompositions decreases quickly towards zero. Unsurprisingly, only plants which draw the zero fixed cost choose to adjust; the other plants prefer to wait until they draw it. As a result, the number of adjusters is roughly constant (Row 9). For d = 0.20, i.e. a 20 percent chance of drawing a zero cost, ShareADJ2 and ShareADJ20 drop from 61.1 percent and 30.6 percent respectively to 38.7 percent and 14.1 percent.

Introducing breakdowns

The introduction of a convex G however makes the extensive margin dominant not only for spikes but also for all investment. But as we now show, this is probably because the model abstracts from "mandatory" maintenance investment, which must be modeled differently. Indeed, a criticism of our decomposition for the low threshold is that it is not robust to small investments which firms make for maintenance and which may not be subject to a fixed cost. Think of replacing some broken machine: little planning or consulting needs to be done, as the managers may know already exactly what type of machine to buy, where to get it, and how to install it. We amend the model in the following way to capture this. At the beginning of each period, each plant faces a probability λ of a breakdown. When a breakdown occurs, the plant must immediately replace a fraction χ of its capital stock which has been destroyed. Once the breakdown has (or not) occurred, the usual sequence of events takes place, with each plant drawing a

---

convex, the older-vintages plants are the ones which react most to the aggregate shock. This attractive feature is not true with the uniform distribution. We suspect that these features are all related.

[21] E.g. Sheshinksi and Weiss (1977)and (1983), Caplin and Spulber (1987), Caplin and Leahy (1991).

fixed cost and then deciding to adjust or to wait. This is a simple way to introduce small investment rates in the model.

We first start with the case of "mandatory maintenance" i.e. the breakdown occurs with probability one. We assume χ = 10 percent. In this case (Row 10), there is literally no-one with zero or near-zero investment.[22] Maintenance remedies the defect of the model of attributing too much of investment to spikes: the share due to spikes falls from 85.9 percent to 57.0 percent. In terms of the decompositions, the number of firms with investment greater than 2 percent of the capital stock is now constant equal to one, and the variance of adjusters is thus zero. Maintenance thus allows us to get a large difference between the decomposition for the low threshold and the decomposition for the high threshold.

However, the lack of zeros is unrealistic, and the 0/100 percent decomposition between extensive and intensive margins is excessive. A simple way to remedy this is to make the breakdowns random. For instance, in row 11, we use a 50 percent probability of a breakdown. Each breakdown still requires a 10 percent investment. We find that this creates some zeros and some low amounts of investment, which corresponds to the data. The share of adjusters is now low but not nil for low threshold. Allowing for these random machine failures thus allows us to match the volatility of the number of adjusters for the low threshold.

C. Our Calibration

Based on these various experiments, we select a "preferred calibration" which is presented in Table 5. The moments we obtain from this calibration are reported in the last line of Table 4. In the next section we investigate the aggregate dynamics that this calibration generates, but before doing so we briefly review the motivation behind the main differences between our parameters and the ones chosen by Thomas.

We want to have large fixed costs because the irrelevance result of Thomas (2002) is hardly surprising if fixed costs are very low, and is interesting only for large fixed costs. So we set B = 0.05. For this calibration, the fixed costs represent 3.5 percent of total investment (that includes maintenance in the measure of investment), and for the marginal plants they are on average 10.7 percent. This is substantially higher than in Thomas, but these numbers are reasonable and in line with the estimates of Cooper and Haltiwanger (2005).

In keeping with the econometric estimates cited above we set returns to scale to be 0.65. As noted earlier, this not only seems more empirically relevant but also helps to partially eliminate the excessive amount of inaction that is present when the Thomas estimate of 0.905 is used. Likewise, we also choose a convex cumulative distribution function for

[22] The smallest investment rate, apart from the one resulting from the breakdowns (10 percent), is from firms which adjusted last period, and readjust their capital for depreciation and trend growth, which is about 7.6 percent (6% of depreciation and 1.6% of growth). Of course this is along the balanced growth path, and there are deviations from this trend, but they are not very large and they average out.

fixed costs which makes the extensive margin dominate for spikes. Finally we also have breakdowns which allow us to match the fact that the intensive margin is large for a low threshold. We also match the importance of I20/K in accounting for changes of I/K, both because most of the investment in our calibration is spikes, and the non-spike investment is mostly maintenance, which by definition is constant over time.

In addition to the critical changes, we also make some minor modifications that have only a limited effect on the intensive-extensive decomposition. Our persistence parameter for TFP is slightly smaller (0.82 rather than 0.92); in this we follow Khan and Thomas (2005) rather than Thomas (2002). Similarly, our estimate of 0.12 for depreciation is perhaps on the high side of typical estimates, but like the higher curvature it penalizes excessive inaction, and thus helps to reduce the number of zeros.

This calibration is not fully optimized, i.e. it is likely that we can match the moments more closely. But, this calibration does roughly as well as the others in Table 4 regarding the steady-state, and it matches the two extensive-intensive decompositions and the variance share of I20/K well. The main shortcoming of this parameterization is that the model does not generate enough spikes, because with large fixed costs there is a lot of inaction. Future work will tackle this issue, which we believe can be corrected without changing the aggregate properties of the model.

4. Aggregate Dynamics and the Irrelevance Result

A. The Thomas result

We conclude our analysis by revisiting the Thomas (2002) "irrelevance result" using our new calibration of the fixed cost model. Thomas compared the effect that aggregate shocks have on investment dynamics when the fixed cost is positive and when the fixed cost is zero. In the later case, the model simplifies to the standard RBC model (with decreasing returns to scale) without any adjustment cost. Figure 8 plots the impulse response of the two models to the productivity shock.[23] The striking result is that the two models are virtually indistinguishable, with the two lines sitting on top of each other. The response on impact of the fixed cost model is about 99.8 percent of the response of the RBC model.

Thomas was careful to check that this result holds for many variations of parameter values. For instance, changing the elasticity of labor supply or the source of shocks does not affect the result. Increasing the level of fixed costs (B), while maintaining a uniform distribution, also makes little difference: for instance, when B is multiplied by a factor of 10, i.e. B = 0.02, so that the maximum vintage is J=20, the impact response of the fixed cost model is 98 percent of the response of the Thomas model. That is, larger fixed costs lead to a slightly smaller response of investment, but the difference between the two models remains negligible.

[23] In a one-shock linear model, the impulse response function summarizes the full dynamics of the system. Hence, models which have the same IRF have exactly the same dynamics in all respects.

Also, in contrast to partial equilibrium analysis, where fixed cost models can generate oscillatory dynamics ("echo effects"), these are completely absent in Figure 8. Thomas argued that the general equilibrium nature of the model was responsible for the inconsequential impact of the micro lumpiness.

While general equilibrium undoubtedly has important smoothing effects, and in particular makes non-monotonicity of the investment impulse response difficult to achieve, we show below that it is not the whole story. Depending on microeconomic assumptions, features typical of the partial equilibrium responses with fixed costs may still arise in general equilibrium. Bachman, Caballero and Engel (2006) present a different set of calculations, but reach a similar conclusion. In particular, like ours, their model presumes higher curvature, higher fixed costs, and some forced "maintenance" investment. They calibrate the preference parameters differently, to reproduce "sectoral level" volatility. With these features, they obtain like us differences between the impulse responses of the two models. They also obtain that the elasticity of aggregate investment with respect to shock is time-varying. (This feature is absent from our model because it is log-linear.)

We see two main differences between our paper and the study by Bachman, Caballero and Engel (2006). First, we keep the same preferences as Thomas (2002), i.e. log utility of consumption and linear disutility of leisure (as in Hansen (1985) and Rogerson (1988)). Since the dispute is about whether general equilibrium offsets are central to this debate, we believe this is the appropriate place to start. Second, we focus on the shape of the distribution of fixed costs.[24]

In contrast to our findings, Khan and Thomas (2005) show that the "irrelevance result" is robust to the introduction of idiosyncratic productivity shocks to the Thomas (2002) setup. Kahn and Thomas maintain the assumption that the distribution of fixed costs is uniform. Since idiosyncratic shocks are not very persistent, the heterogeneity that results increases benefits of adjusting in the current period. Intuitively, it seems possible that if the shocks were more persistent then the shocks could generate something akin to the marginal homogeneity condition. Khan and Thomas also emphasize that general equilibrium feedbacks affects plant-level investment dynamics, which would imply that the panel data estimates from partial equilibrium models that we use may be misleading.

B. Impulse response: Single shocks

We start by displaying in Figure 9 the impulse response function of aggregate investment to a productivity shock for our preferred calibration from Section 3, along with the RBC model with has the same parameters but zero fixed costs. The two models have noticeably different dynamics in two respects. First, the response is initially smaller in the

[24] Another recent paper on the topic is Svenn and Weinke (2005). In contrast to Thomas (2002) or Caballero and Engel (1999), they use a Calvo-style time-dependent adjustment rule for capital. Interestingly, they find that given this rule, the irrelevance result holds in the RBC model but not in a New Keynesian model.

fixed cost model: on impact the response of the fixed cost model is only 83 percent of the response of the RBC model. Second and more interestingly, the fixed cost model exhibits a substantial hump nine periods after the shock. We call this hump an "echo effect" because it is caused by the initial surge in investment: as many firms adjust, the distribution shifts toward more recent vintages, which are less likely to invest. This makes the investment response smaller than the RBC model, until the cycle ends (after eight years when everyone is forced to invest given this calibration) and these units adjust again. (Of course, adjustment is random, and probabilities of adjustment move over time, but on average the length of the cycle still plays an important role.) Figure 10 shows the evolution of the cross-sectional distribution. The substantial hump in number of units adjusting initially moves through time and is preserved until $t = 9$ when it translates into a higher investment. Clearly, this result relies on the hazard rate (the probability of adjusting as a function of vintage, i.e. alpha), which for our calibration is a steeply convex function: the alphas are almost negligible for all vintages except for the penultimate vintage and the oldest vintage (when alpha equals one).

C. Impulse responses with more complicated shocks

The differences shown after a single shock often become more pronounced when we consider more complicated disturbances. To illustrate this, consider the model's predictions when it is fed a succession of five positive productivity shocks in a row.[25] This is presented in Figure 11. In this case, the two differences we pointed out for the single shock case become starker: the investment response while the shocks are arriving is substantially smoother in the fixed cost model, but then multiple "echoes" appear over the course of several consecutive periods (whereas the RBC model responses exhibit a monotonic decay).

A second case in which the differences between the two models are accentuated is when there is a large change to the cross-sectional distribution of capital. Of course, the evolution of the cross-sectional distribution is endogenous, but it is interesting to see the effect of starting from a cross-sectional distribution that is not in steady-state and watching how aggregate investment evolves as the distribution returns to its steady-state.[26] To make the initial disequilibrium position realistic, we consider the cross-sectional distribution that arises after a series of five positive technology shocks.

Our experiment amounts to computing the aggregate, general equilibrium predictions for investment when the economy simply begins with this initial cross-sectional distribution (but is not hit by any future productivity shocks). Because productivity is positively serially correlated, normally if the cross-sectional distribution were to reach this position, firms would continue to have productivity above average for many periods (even without any new shocks) as a result of the past disturbances. This experiment can, therefore, be thought of as simulating the effects of a set of beliefs about future productivity that

[25] We borrow this device from Bachmann, Caballero and Engel (2006).

[26] One reason, for instance, why the cross-sectional distribution may shift independently, is when there is a rise in uncertainty , so that firms delay investment (Bloom 2006).

ultimately are not realized.[27] The initial series of positive innovations will mean that a large number of firms will have invested, so that when the firms revise their beliefs and recognize productivity will not improve further, they will be doing so while having more capital than they would have chosen given their current expectations.

Figure 12 plots the effect in the RBC model and in the fixed cost model. The differences are now extremely pronounced. While the RBC model displays the usual, monotonic, smooth convergence to the steady-state given a high starting initial capital, the fixed cost model exhibits oscillations and great differences in magnitudes. We emphasize that these results are obtained with log utility; as a point of reference Figure 13 shows the same experiment in the baseline Thomas model. With her calibration the RBC model and the fixed cost model yield essentially identical predictions even for this particular experiment.

This for us is proof that general equilibrium effects are not the only reason why Thomas found no aggregate effect of fixed costs. Depending on microeconomic assumptions, the equivalence result need not hold.

D. Comparing the fixed cost model to a quadratic adjustment cost model

There is a widely held conjecture that even when fixed costs matter, their aggregate effects will be similar to those of a standard representative firm quadratic adjustment cost model. For instance, Hall (2004) citing an earlier version of Cooper and Haltiwanger (2005) and Caballero and Engel (1999), concludes that "the quadratic specification provides a reasonably accurate approximation [at the industry level]". To explore this conjecture, we add a quadratic adjustment cost to the RBC model and choose the adjustment cost parameter η so that following a productivity shock the augmented RBC model impact response is identical to the one in our preferred calibration of the fixed cost model. We then compare the subsequent responses.

This experiment is shown in Figure 14. The parameter of the quadratic adjustment cost that we require to match the impact response of the fixed cost model is $\eta = 12.107$ (this is the elasticity of the investment rate with respect to Tobin's Q). Clearly, the dynamics implied by the fixed cost model and the quadratic adjustment cost model differ, because the quadratic adjustment cost omits any echo effect.[28] Thus, we conclude that appending a simple quadratic cost of adjusting capital to the standard RBC model is not necessarily a good way to approximate a fixed cost model either.

[27] Note given the linear structure of the solution, the full response to five consecutive shock (Figure 10) equals the dynamic response shown in Figure 11 *plus* the response due to productivity being higher than usual. Hence, Figure 11 is really the first part of a decomposition of Figure 10.
[28] When we change some parameters, we also sometimes find that investment in the fixed cost model is "paradoxically" more volatile than in the RBC model without adjustment costs.

5. Conclusions

We make three contributions to the debate over the aggregate significance of plant-level investment lumpiness. Remarkably, the basic plant-level facts on the lumpiness of investment are fairly similar in Chile and the U.S. In each country, we show that investment spikes drive total investment. The spikes draw their predictive power from changes in number of plants making large investments, rather than changes in the size of average investment per plant. When we look at all investment (not just the part associated with spikes), changes in the size of investment per plant is more important. We use these statistics regarding the decomposition between the intensive and extensive margins of adjustment to summarize the microeconomic facts about lumpiness that we ask a model to match.

We use the Thomas (2002) model to examine these facts. This model augments a relatively standard RBC model by assuming that firms must pay a fixed cost (that is randomly drawn each period) in order to adjust its capital. As originally calibrated, however, the model fails to account for the facts regarding the intensive and extensive margins and the dominant role of investment spikes in explaining total investment. We argue (appealing to recent econometric work by others) that the original calibration has an average level of fixed costs which is too low and a profit function that has too little curvature. We adjust the calibration to reflect these estimates and find that these changes alone do not fundamentally change the model's properties, particularly regarding its ability to fit the intensive and extensive margins.

We also consider a third change whereby the distribution of fixed costs from which firms sample is much more compressed (than the distribution considered by Thomas). This change raises the prominence of extensive adjustment (but does so for both spikes and regular investment). When we also add the necessity of some maintenance investment, then we finally arrive at a calibration that approximately accounts for the importance of investment spikes in explaining the variation of total investment and the differential importance of intensive and extensive adjustment for the different investment thresholds.

Our final contribution is to study the properties of the model (using our preferred calibration) regarding various shocks. In the original Thomas model the aggregate dynamics for investment following a productivity shock were indistinguishable from an RBC model with no adjustment costs. In our model this type of shock plays out differently in two respects. First, the impact response of total investment is partially muted (relative to the RBC model). Second, there is an "echo" whereby the response of firms who adjust immediately is repeated when they replace their initial investment. These differences are even more pronounced for some other shocks (such as a sequence of productivity shocks). Moreover, our preferred calibration of the model is not well approximated by a RBC model with quadratic costs of adjusting the capital stock.

Our conclusion from the last exercise is that there is nothing generically related to DSGE models that guarantee plant-level investment lumpiness is smoothed away. Rather we agree with Thomas that there can be substantial differences between the importance of lumpiness in a GE models and partial equilibrium models. However, many have gone farther and concluded that GE makes fixed costs to investment completely irrelevant for the business cycle. Both our empirical and theoretical work show this conclusion is premature. Instead, we see the answer to this question as being sensitive to details of how the model is set up. Given the currently available information, we think our calibration is reasonable, but we recognize much more work needs to be done in this respect to determine how these models should be estimated and calibrated. Therefore, it is too early to conclude that fixed costs are irrelevant for aggregate investment dynamics.

Appendix: Data

The purpose of this appendix is to briefly describe the data that we analyze in section 2.

US Census data:

One of our data sets relates to U.S. establishment-level data between 1972 and 1998. These data were kindly provided by Shawn Klimek of the Census Bureau. The capital expenditure data are taken from the Census Bureau's Annual Survey of Manufactures (ASM) and the details of the data construction are given in Becker et al (2006). Their core calculation involves building up a capital stock series using a perpetual inventory method.

The data provided to us were sorted establishments into different categories according to the ratio of investment to beginning of period capital ($I_{i,t}/K_{i,t-1}$). Totals for investment and (beginning of period) capital were computed by summing across all establishments in a given category for each year; for example, data for total investment for firms with $I_{i,t}/K_{i,t-1} > 0.2$ would be one entry in the spreadsheet that we received. By summing across categories we get total investment (or total capital) for the year.

Chilean data:

Our second data set is a plant-level census of manufacturing plants with ten or more employees from Chile. This data is collected by the National Statistics Institute of Chile and the series we exploit were provided to us by Olga Fuentes and Simon Gilchrist, who constructed real capital stocks and real investment series from a perpetual inventory equation, with industry-specific investment prices and depreciation rates. The data we use is an unbalanced panel which has on average 1780 plants per year, from 1981 to 1999. We delete firms with missing observations. We sort firms based on their investment-capital ratio using the same procedure we use for the Census data.

The spreadsheets with these data are available on the following web page: http://people.bu.edu/fgourio/extintpaper.html

References

Abel, Andrew B., and Janice C. Eberly, 2002, "Investment and *q* with Fixed Costs: An Empirical Analysis", (working paper).

Bachmann, Ruediger, Caballero Ricardo J. and Eduardo M.R.A. Engel, 2006, "Lumpy Investment in Dynamic General Equilibrium", working paper MIT.

Becker, Randy, Haltiwanger John, Jarmin, Ron, Klimek, Shawn and Dan Wilson, 2006, "Micro and Macro Data Integration: The Case of Capital", in A New Architecture for the U.S. National Accounts, edited by Dale Jorenson, Steve Landefeld and William Nordhaus, University of Chicago Press, pp. 541-609.

Bernanke, Benjamin S. Bohn, Henning and Peter C. Reiss, "Alternative Nonnested Specification Tests of Time Series Investment Models" in Journal of Econometrics, Vol. 37, No. 2, March 1988, pp. 293-326.

Bloom, Nick, 2006, "The Impact of Uncertainty Shocks: Firm Level Estimation and a 9/11 Simulation", mimeo Stanford University.

Caballero, Ricardo, 1992, "A Fallacy of Composition", American Economic Review, 82(5):1279-1292.

Caballero, Ricardo J. 1999. "Aggregate Investment", Handbook of Macroeconomics, Vol. 1B, Elsevier Science.

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

Caballero, Ricardo J., Engel, Eduardo M. R. A., and John C. Haltiwanger, 1995, "Plant-Level Adjustment and Aggregate Investment Dynamics", Brookings Papers on Economic Activity, no. 2, pp. 1-54.

Cooper, Russell W., and John C. Haltiwanger, 2005, "On the Nature of Capital Adjustment Costs", Review of Economic Studies, forthcoming.

Cooper, Russell W., Haltiwanger, John C. and Laura Power, 1999, "Machine Replacement and the Business Cycle: Lumps and Bumps", American Economic Review, 89(4), pp. 921-946.

Doms, Mark, and Timothy Dunne, 1998, "Capital Adjustment Patterns in Manufacturing Plants", Review of Economic Dynamics, 1 pp. 409-429, (April).

Dotsey, Michael, King, Robert G. and Alexander L. Wolman, 1999, "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output", Quarterly Journal of Economics, 114 pp. 655-690, (May).

Fuentes Olga, Simon Gilchrist and Marc Rysman, 2006, "Irreversibility and Investment Dynamics for Chilean Manufacturing Plants", mimeo Boston University.

Hall, Robert E., 2004, "Measuring Factor Adjustment Costs", Quarterly Journal of Economics, CIXX (August), pp. 899-927.

Hansen, Gary D., 1985, "Indivisible Labor and the Business Cycle", Journal of Monetary Economics, 16, pp. 309–327.

Hennessy, Christopher A. and Toni M. Whited, 2005, "Debt Dynamics", Journal of Finance, LX(3), pp. 1129-1165.

Hsieh, Chang-Tai and Jonathan A. Parker, 2006, "Taxes and Growth in a Financially Underdeveloped Country: Evidence from the Chilean Investment Boom", NBER Working Paper 12104.

Kahn, Aubhik, and Julia K. Thomas, 2003, "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?", Journal of Monetary Economics, 50, pp. 331-360.

Kahn, Aubhik, and Julia K. Thomas, 2005, "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics", mimeo.

Klenow, Peter J., and Oleksiy Kryvtsov, 2005, "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?", NBER Working Paper 11043.

Oliner, Stephen .D., Rudebusch Glenn D. and Daniel E. Sichel, 1995, "New and Old Models of Business Investment: A Comparison of Forecasting Performance", Journal of Money, Credit, and Banking, vol. 27 (August 1995), pp. 806-32.

Prescott, Edward C., 2003, "Non-Convexities in Quantitative General Equilibrium Studies of Business Cycles", (working paper).

Rogerson, Richard, 1988, "Indivisible Labor, Lotteries, and Equilibrium", Journal of Monetary Economics, 21, pp. 3-16.

Sheshinski, Eytan, and Yoram Weiss, 1977, "Inflation and Costs of Price Adjustment", Review of Economic Studies, vol. 44(2), pp. 287-303.

Sheshinski, Eytan, and Yoram Weiss, 1983, "Optimum Pricing under Stochastic Inflation", Review of Economic Studies, vol. 50(3), pp. 513-529.

Svenn, Tommy and Lutz Weinke, 2005, "Is Lumpy Investment Irrelevant for the Business Cycle?", mimeo.

Thomas, Julia K., 2002, "Is Lumpy Investment Relevant for the Business Cycle," Journal of Political Economy, CX, pp. 508-534.

Table 1: Distribution of Investment Rates for U.S. and Chilean Plants

	Census Equal Weighted (Percent)	Census Capital Weighted (Percent)	Chile Equal Weighted (Percent)	Chile Capital Weighted (Percent)
I/K=0	15.8	3.4	41.31	18.55
0<I/K<2	15.1	12.1	11.04	17.24
8>I/K>2	29.7	33.3	17.78	26.32
12>I/K>8	11.5	14.4	6.76	9.44
20>I/K>12	11.6	16	8.44	11.84
35>I/K>20	8	10.8	7.36	8.83
I/K>35	8.3	10	7.31	7.78

Table 2: Effect of Investment Spikes on Aggregate Investment

Dependent variable is $Itot_t/K_{t-1}$, rows of the table show regressions with different right hand side variables that are defined in the text. A time trend is always included (but not shown) to save space. For the census sample time period is 1974 to 1998. For the Chilean sample the time period is 1981 to 1999. The standard errors are computing using the Newey-West correction with three lags.

			Coefficient estimates (standard errors)			
Row	Sample	$\bar{R}^2$	$Itot_{t-1}/K_{t-2}$	$Sales_{t-1}/K_{t-2}$	$ShareADJ20_{t-1}$	$ShareADJ20_{t-2}$
1	Census	0.748	0.743 (0.101)			
2	Census	0.738	0.690 (0.094)	0.0078 (0.0098)		
3	Census	0.776	1.255 (0.180)		-0.204 (0.044)	
4	Census	0.893	1.553 (0.165)		-0.228 (0.035)	-0.161 (0.048)
5	Census	0.786	1.257 (0.153)	0.0199 (0.009)	-0.258 (0.039)	
6	Census	0.866	1.531 (0.167)	0.010 (0.008)	-0.250 (0.033)	-0.157 (0.055)
7	Chile	0.809	0.353 (0.292)			
8	Chile	0.848	0.151 (0.257)	0.055 (0.017)		
9	Chile	0.802	0.999 (0.804)		-0.331 (0.341)	
10	Chile	0.847	1.152 (0.753)		-0.454 (0.272)	-0.405 (0.061)
11	Chile	0.839	0.462 (0.764)	0.054 (0.018)	-0.156 (0.339)	
12	Chile	0.856	0.790 (0.629)	0.034 (0.12)	-0.323 (0.264)	-0.331 (0.075)

Table 3: Baseline Parameters in the Thomas (2002) Calibration

Parameter	Value
Depreciation rate (δ)	0.06
Persistence of TFP shock (ρ)	0.9225
Returns to scale ($\psi + \upsilon$)	0.905
Share of capital in Production Function ψ	0.325
B (maximum fixed cost)	0.002
Discount factor (β)	0.954

Table 4: Steady-State and Business Cycle Lumpiness Statistics for various calibrations.

			Capital-Weighted Investment Shares					
		J	Mean % Plants I/K<0.02	Mean % Plants I/K>0.20	Mean I20/Itot	% Variance of Itot/K due to I20/K	Share ADJ2	Share ADJ20
1	Data US	NA	15.5	20.8	49.9	97.0	33.5	87.0
2	Data Chile	NA	35.8	16.6	57.3	86.0	39.5	92.5
3	Thomas (2002) Calibration	5	72.5	19.7	85.9	62.7	61.1	30.6
4	Thomas with Higher B (B=0.02)	20	86.4	12.5	98.1	93.4	64.0	53.1
5	Thomas with Lower Returns to Scale (0.6)	4	67.4	20.0	77.4	44.8	62.9	33.5
6	Thomas with Lower Returns to Scale (0.6) and Higher B (0.02)	13	83.7	14.5	96.9	90.0	63.1	49.5
7	Thomas with Convex G (B=0.01)	9	91.8	8.1	99.9	99.9	78.0	77.8
8	Thomas with Uniform G (B= 0.019 i.e. same mean as row 7)	20	86.2	12.7	98.0	93.1	64.0	52.9
9	Thomas with 20% of plants getting a zero fixed cost.	5	62.5	16.7	67.9	45.0	38.7	14.1
10	Thomas with Breakdowns (10%)	5	0.00	26.0	57.0	109.9	0	61.1
11	Thomas with Random Breakdowns (50% chance of a 10% breakdown)	5	36.0	22.8	65.5	85.8	22.6	48.5
12	Preferred Calibration (See Table 5)	8	65.1	7.0	83.2	101.0	12.5	72.0

Note: See the text for the full characteristics of the alternative calibrations. The definitions of I20, Itot, ShareADJ2 and ShareADJ20 are:

$$\text{I2} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}}>0.02} I_{i,t}, \quad \text{I20} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}}>0.20} I_{i,t}, \quad \text{Itot} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}}\geq 0.0} I_{i,t},$$

$$\text{K2} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}}>0.02} K_{i,t-1}, \quad \text{K20} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}}>0.20} K_{i,t-1}, \quad \text{K} \equiv \sum_{\frac{I_{i,t}}{K_{i,t-1}}\geq 0.0} K_{i,t-1}$$

$$\text{ShareADJ2} \equiv \frac{\text{covariance}(\log(\frac{\text{K2}}{\text{K}}),\log(\frac{\text{I2}}{\text{K}}))}{\text{variance}(\log(\frac{\text{I2}}{\text{K}}))} \text{ and ShareADJ20} \equiv \frac{\text{covariance}(\log(\frac{\text{K20}}{\text{K}}),\log(\frac{\text{I20}}{\text{K}}))}{\text{variance}(\log(\frac{\text{I20}}{\text{K}}))}$$

% Variance of Itot/K due to I20/K = Cov(I20/K, Itot/K)/Var(Itot/K).

Table 5: Preferred Calibration of the DSGE Model.

Parameter	Value
Depreciation rate (δ)	0.12
Persistence of TFP shock (ρ)	0.82
Returns to scale ($\psi + \upsilon$)	0.65
Share of capital in Production Function ψ	0.233
B (maximum fixed cost)	0.05
Discount factor (β)	0.954
Convex CDF with G(x) = H(x/B) with H(x) = c1 + c2*tan(c3x-c4) with H(0)=0, H(1)=1	
C2	0.02
C3	1.4
Lamba (Probability of a breakdown)	0.3
Psi (Investment Rate required if breakdown)	0.10

Figure 1: Investment Lumpiness in U.S. Manufacturing Plants

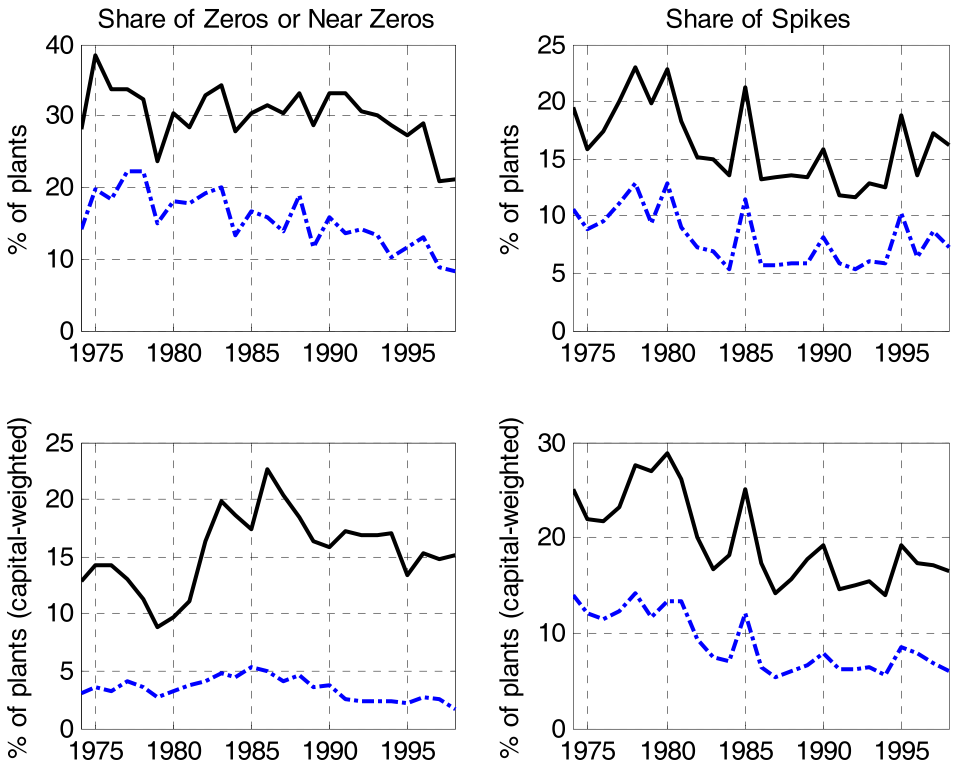

Notes:
Left Panels: "Near Zeros" are defined as plants with 0< I/K < 0.02 and are shown in the solid line. Plants with I/K =0 are shown with dashed line.

Right panels: Solid line is plants with I/K > 0.2. Dashed line is I/K > 0.35.

Figure 2: Investment Lumpiness in Chilean Manufacturing Plants

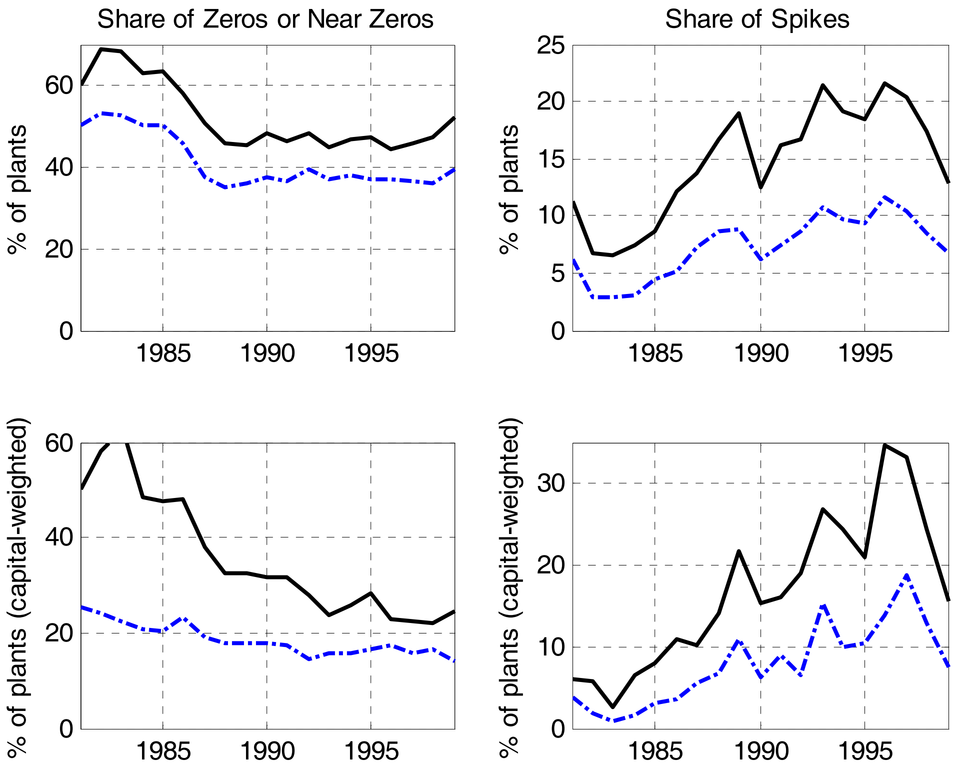

Notes:
Left Panels: "Near Zeros" are defined as plants with 0< I/K < 0.02 and are shown in the solid line. Plants with I/K =0 are shown with dashed line.

Right panels: Solid line is plants with I/K > 0.2. Dashed line is I/K > 0.35.

Figure 3: Cyclicality of Near Zero Investment and Investment Spikes in U.S. and Chilean Manufacturing Plants

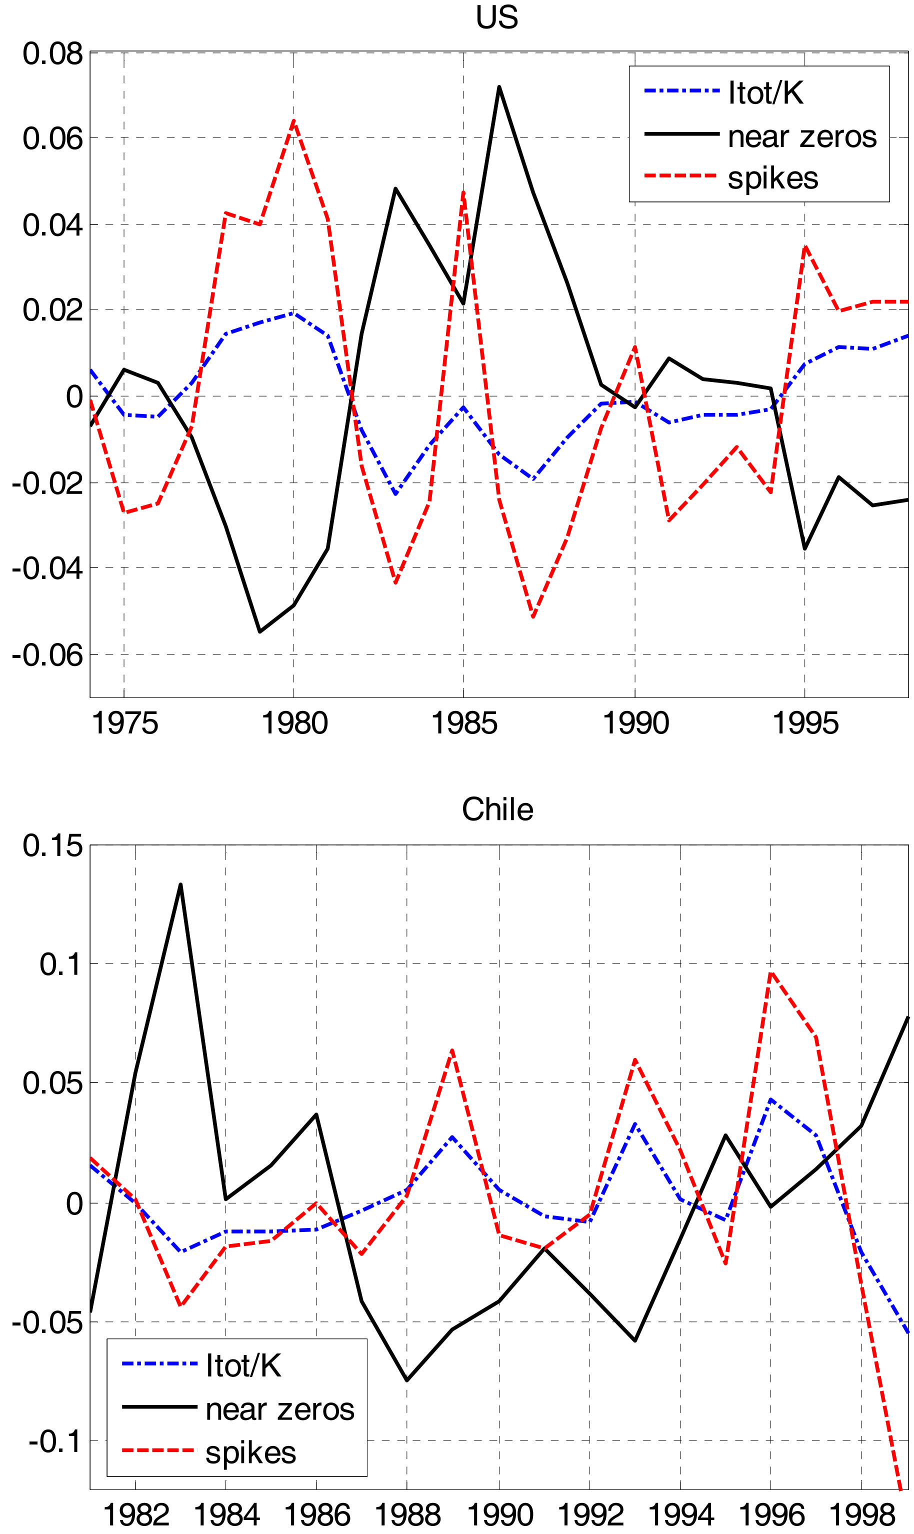

Note: Data are de-trended as described in the text.

Figure 4: Investment Spikes and Investment Non-Spikes Relative to Total Investment for U.S. and Chilean Manufacturing Plant.

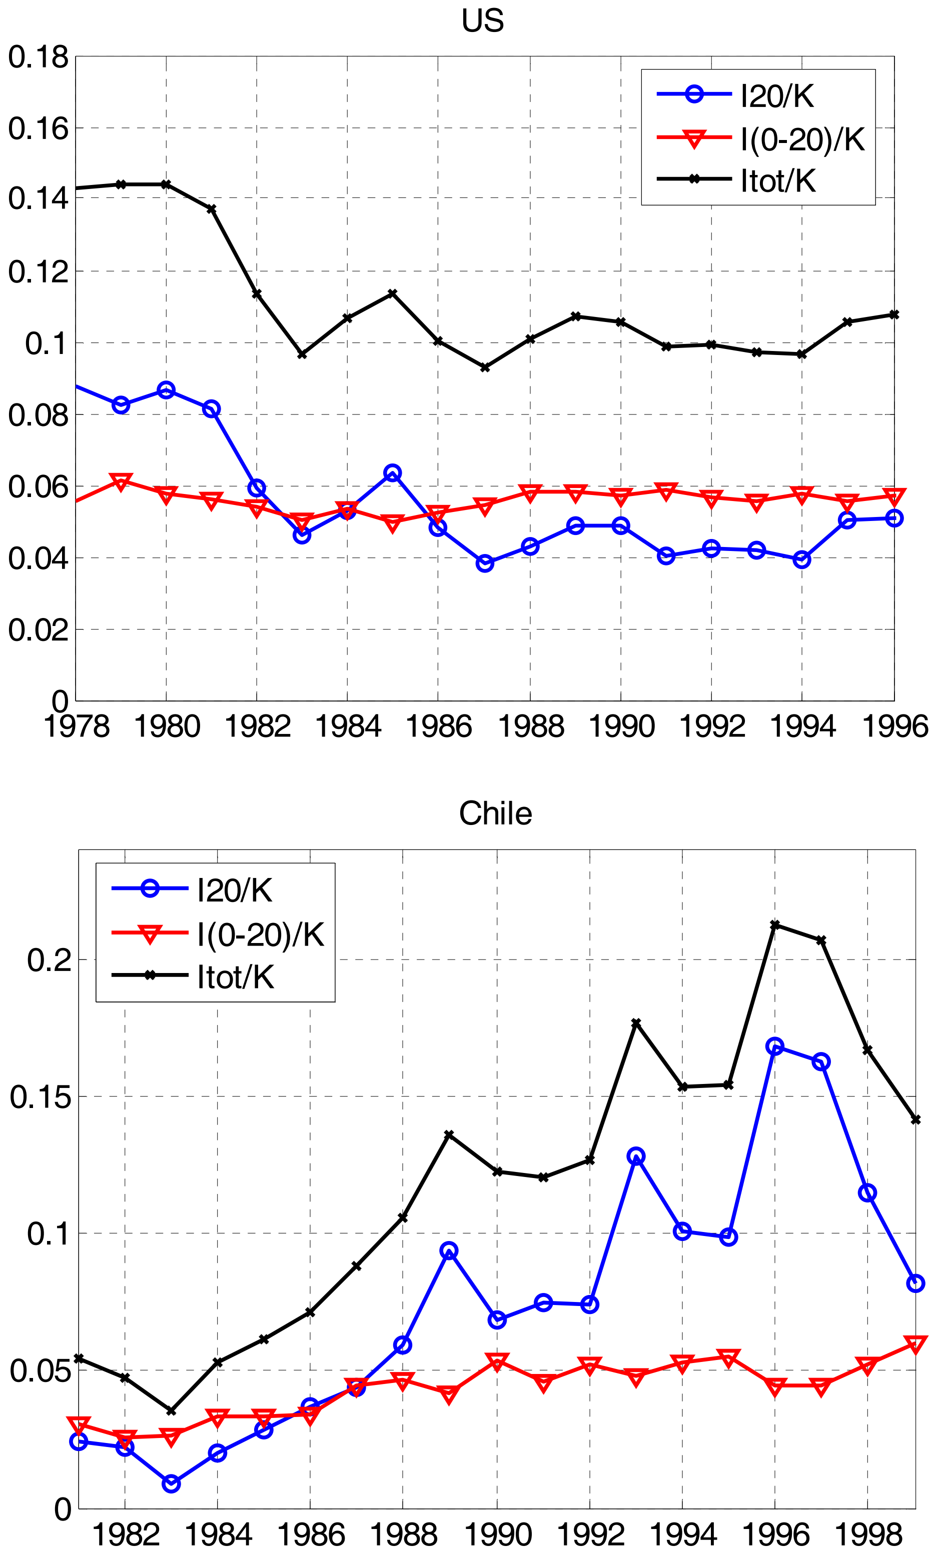

Figure 5: Decomposition of Investment in Intensive and Extensive Adjustment for U.S. and Chilean Manufacturing Plant

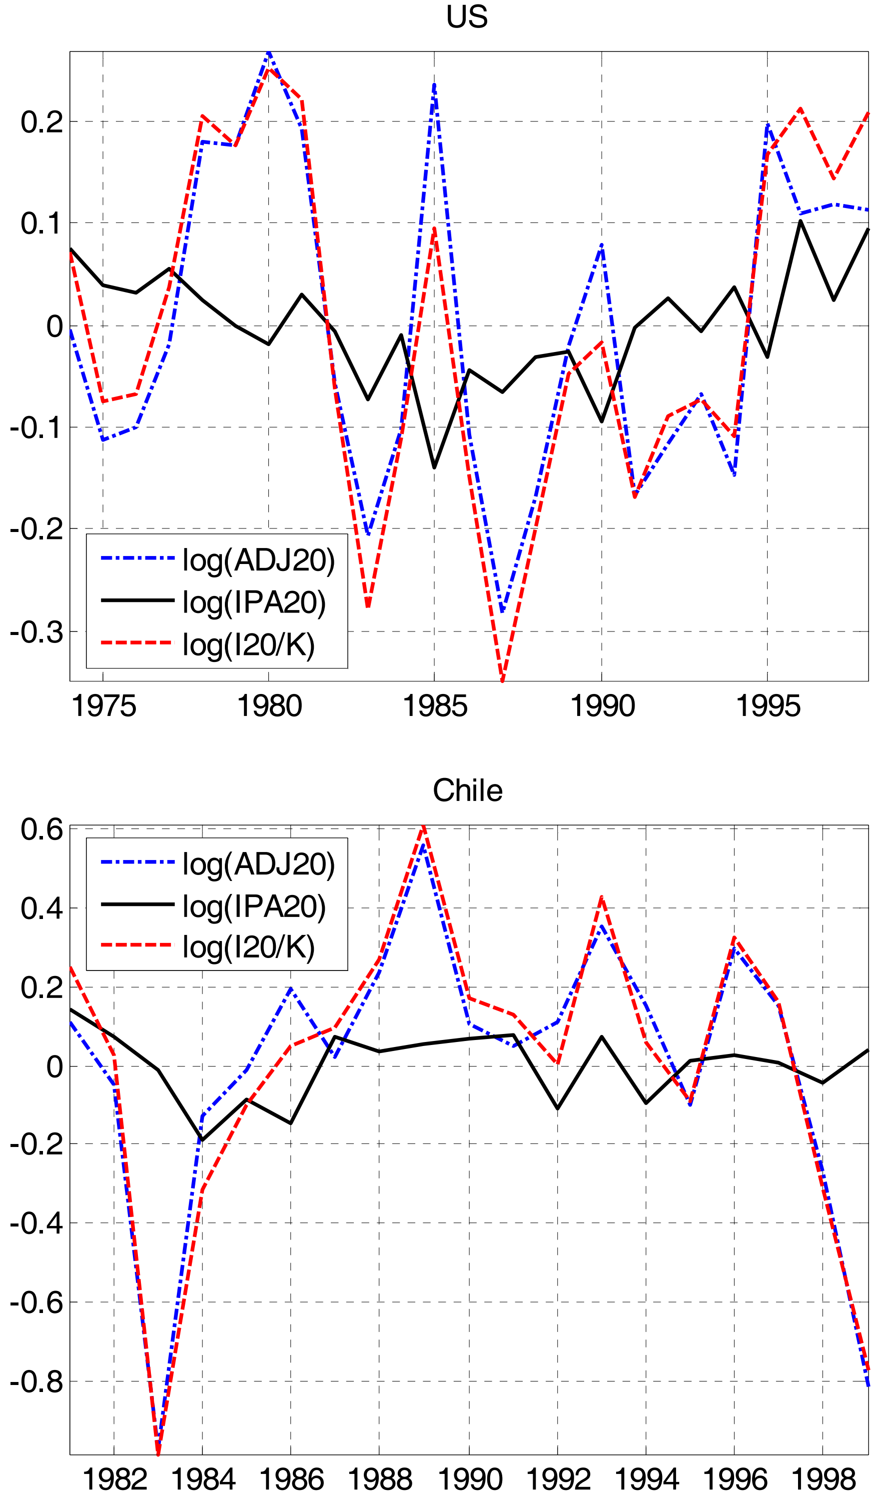

Note: Data are de-trended as described in the text.

Figure 6: Decomposition of Investment in Intensive and Extensive Adjustment for U.S. and Chilean Manufacturing Plant

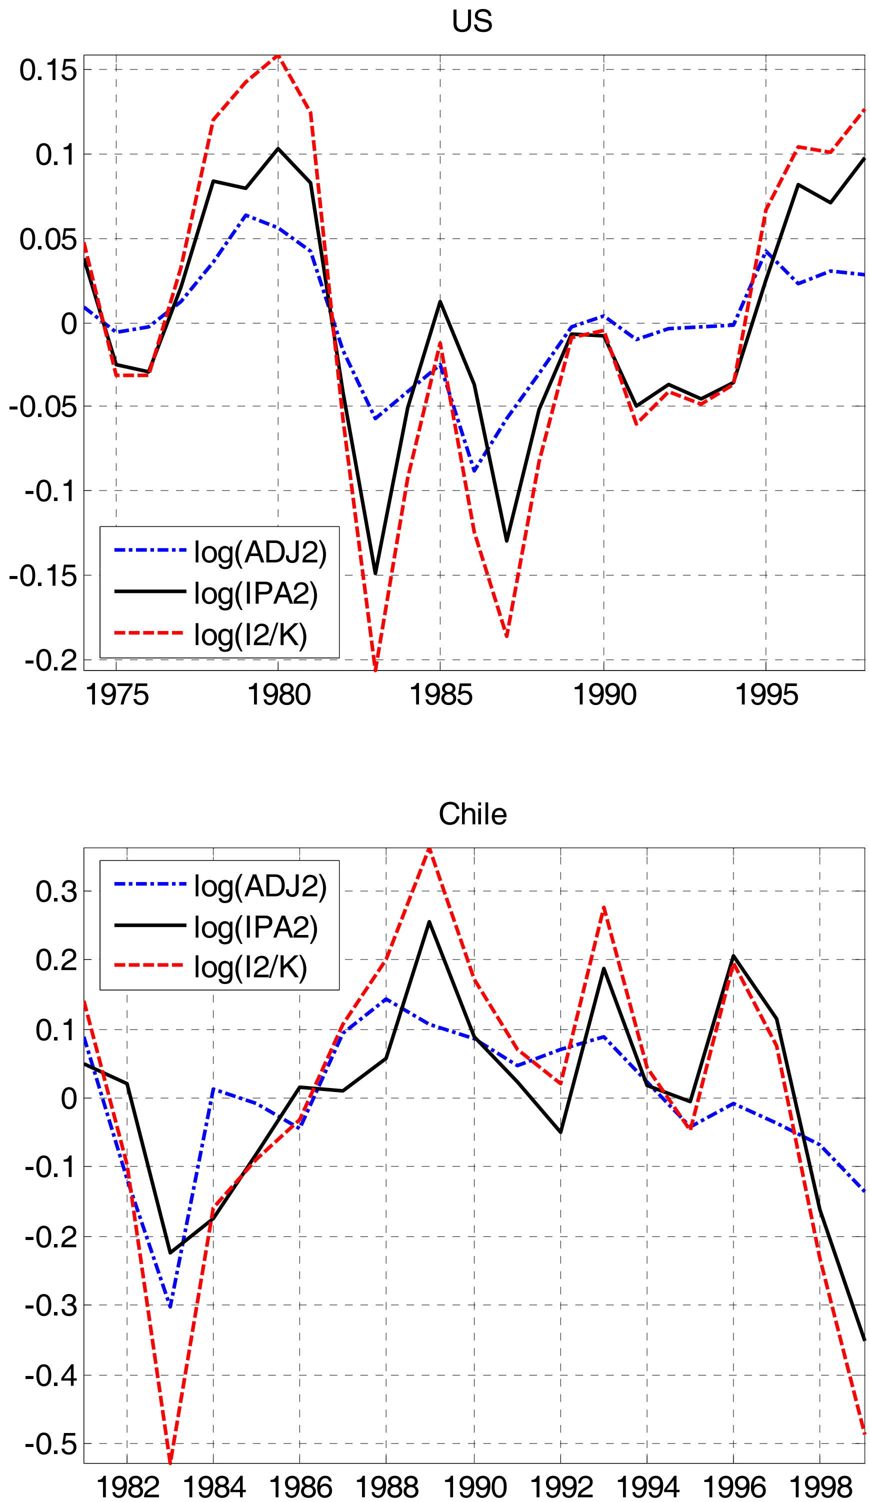

Note: Data are de-trended as described in the text.

Figure 7: Cumulative Distribution Function G of Fixed Costs used in our preferred calibration.

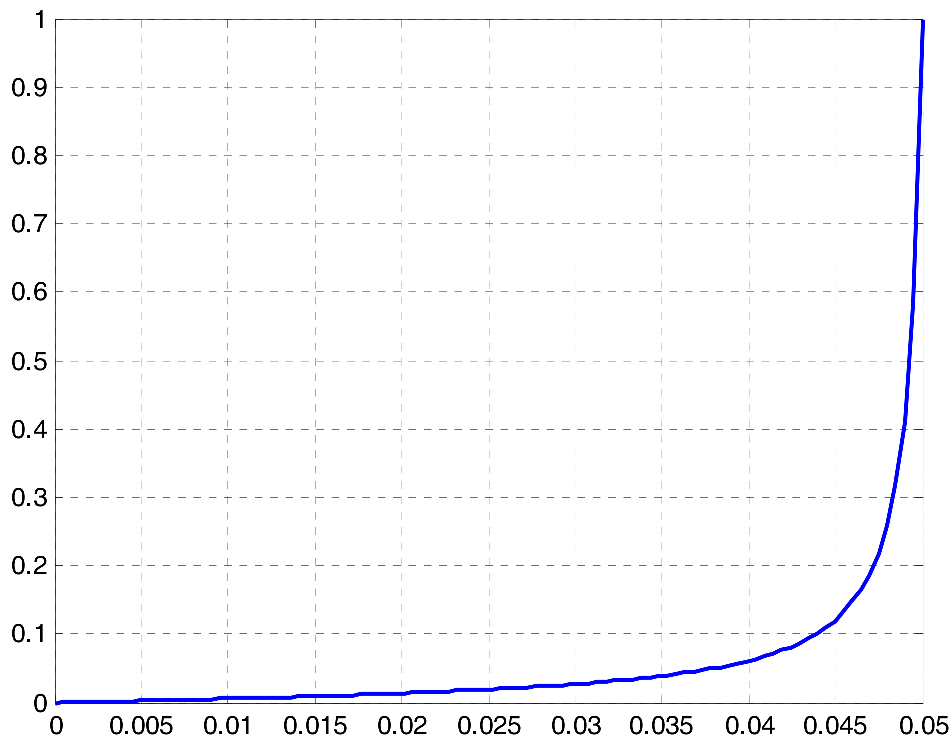

Figure 8: Impulse Response of Aggregate Investment to an Aggregate Productivity Shock for the Original Thomas Calibration of the DSGE Model with Fixed Costs.

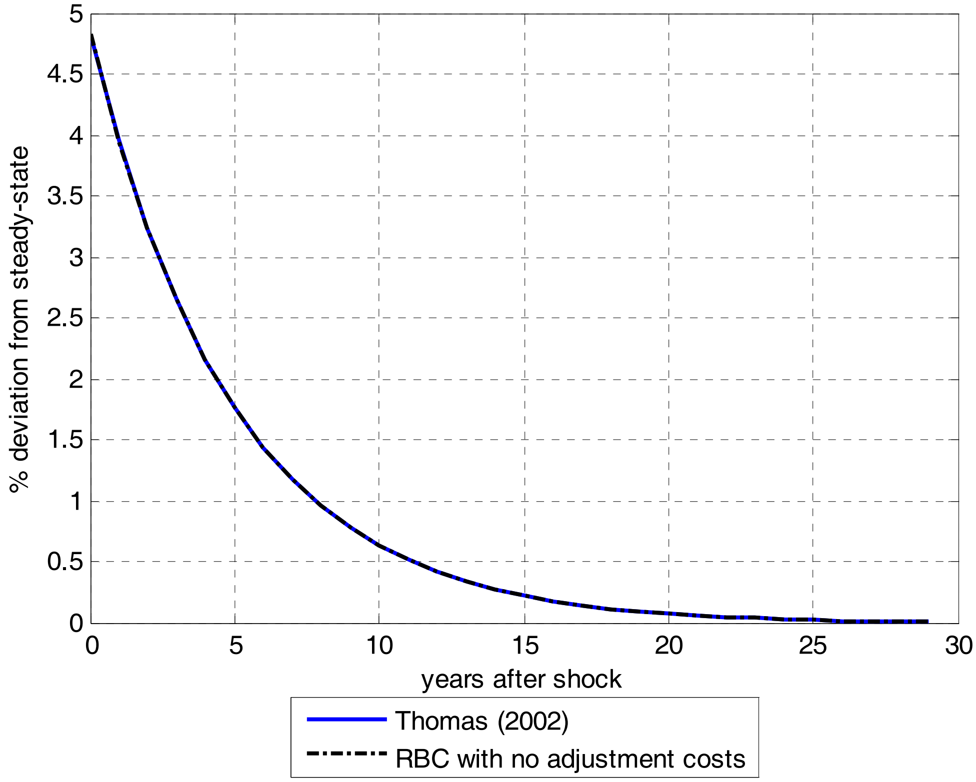

Figure 9: Impulse Response of Aggregate Investment to an Aggregate Productivity Shock for Our Preferred Calibration of the DSGE Model with Fixed Costs.

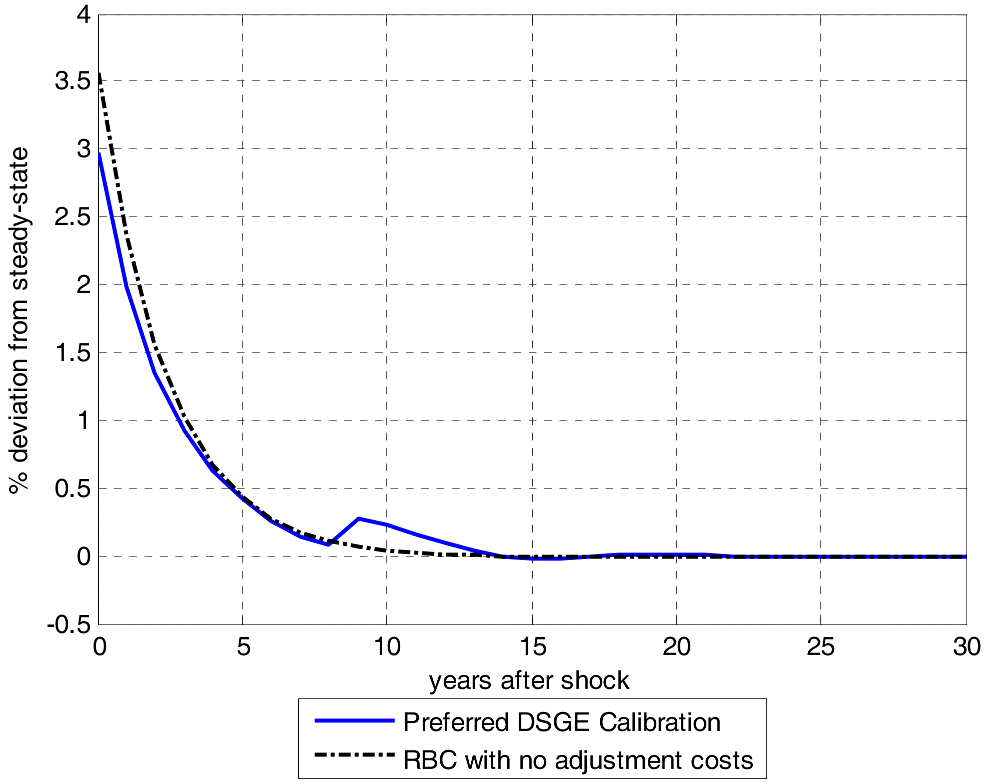

Figure 10: Evolution of Distribution across Vintages following an Aggregate Productivity Shock at Different Points in Time For Our Preferred Calibration of the DSGE Model.

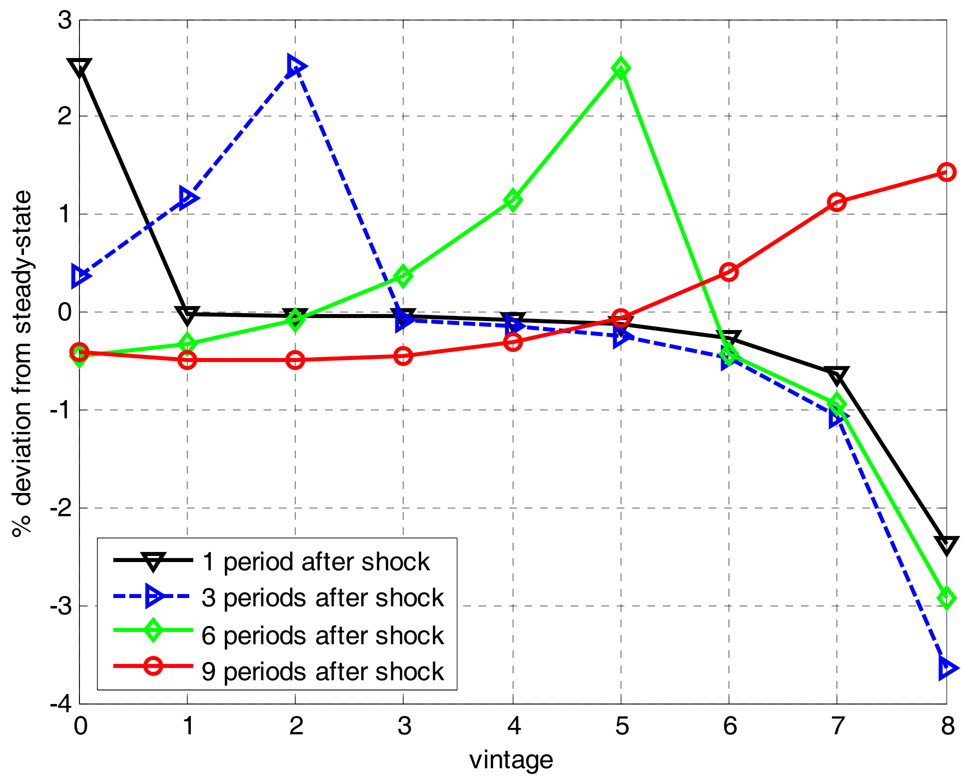

Figure 11: Dynamic Response of Aggregate Investment to Five Consecutive Positive Aggregate Productivity Shocks for Our Preferred Calibration of the DSGE Model with Fixed Costs.

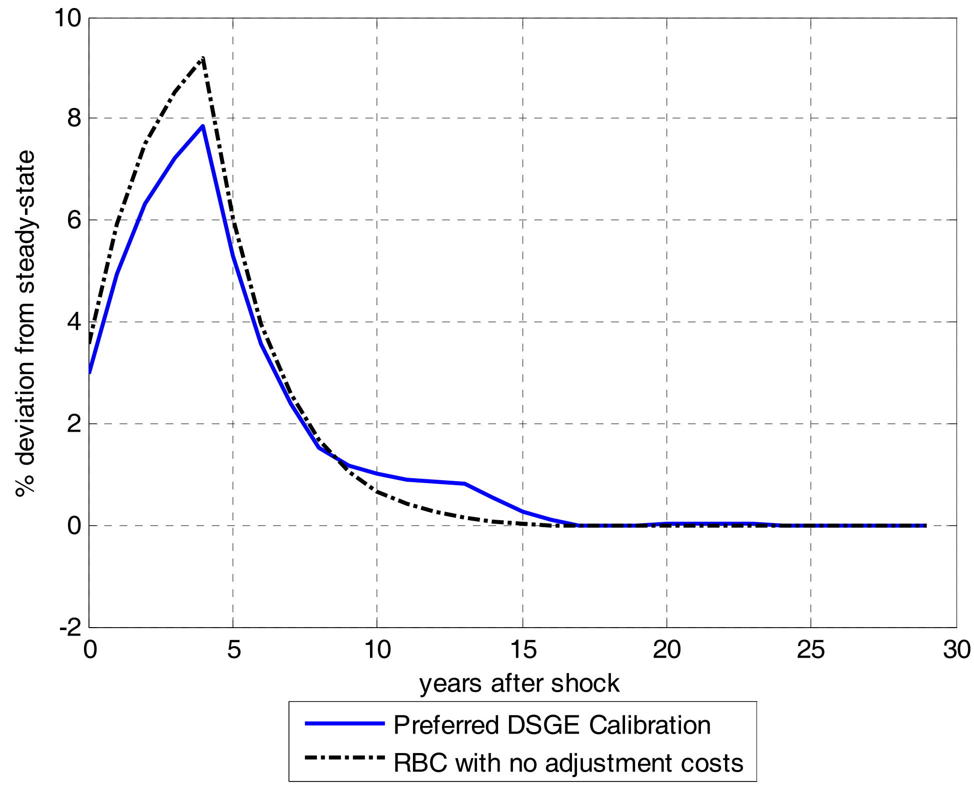

Figure 12: Dynamic Path for Aggregate Investment When the Initial Distribution of Capital is Distorted in Our Preferred Calibration of the DSGE Model with Fixed Costs.

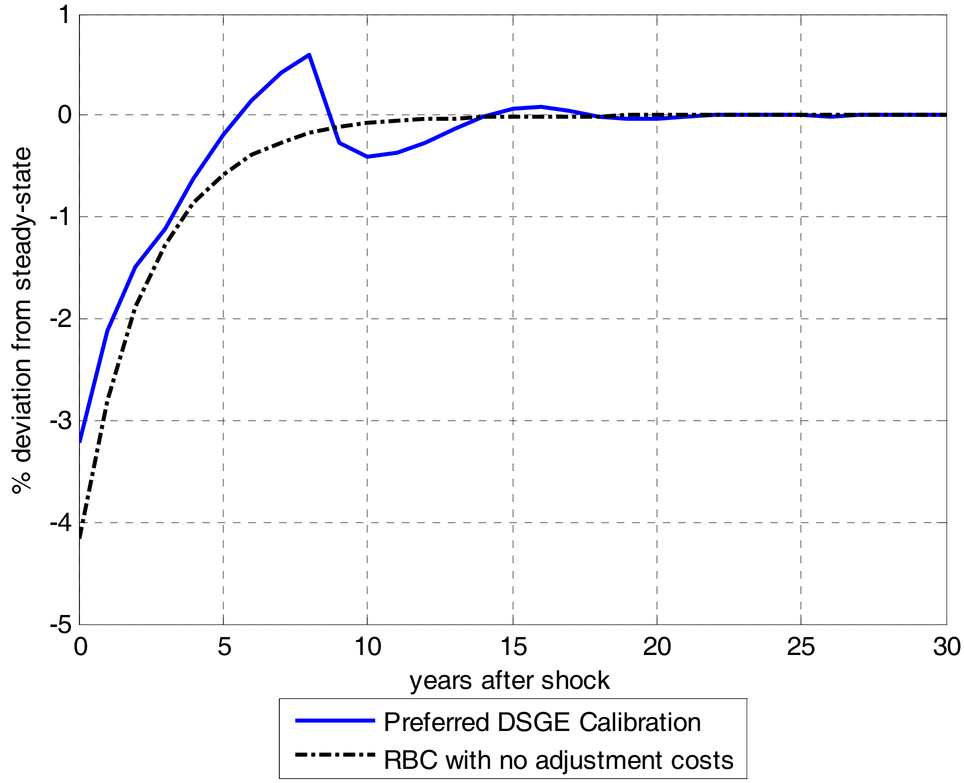

Figure 13: Dynamic Path for Aggregate Investment When the Initial Distribution of Capital is Distorted in the Original Thomas Calibration of the DSGE Model with Fixed Costs.

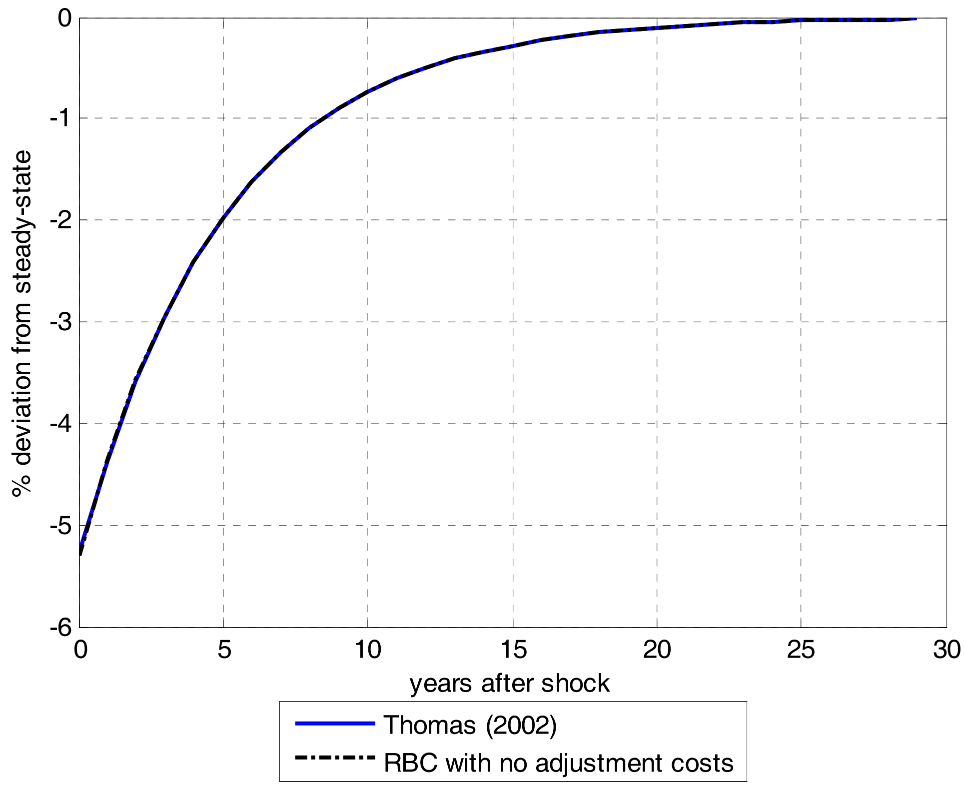

Figure 14: Impulse Response of Aggregate Investment to an Aggregate Productivity Shock for RBC Model with Quadratic Adjustment Costs and for our Preferred Calibration of the DSGE Model with Fixed Costs.