

Lumpy Investment, Lumpy Inventories

Rüdiger Bachmann Lin Ma*

February 3, 2016

Abstract

The link between the physical micro environment (frictions and heterogeneity) and the macroeconomic dynamics of general equilibrium macro models is influenced by the details of how exactly general equilibrium closes such a model. We make this general observation concrete in the context of the recent literature on how nonconvex capital adjustment costs influence aggregate investment dynamics. Specifically, we introduce inventories into a two-sector lumpy investment model. We find that with inventories nonconvex capital adjustment costs dampen and propagate the reaction of investment to shocks, and more so than without inventories: the initial response of fixed capital investment to productivity shocks is 46% higher with frictionless adjustment than in our preferred calibrated capital adjustment frictions model, once inventories are introduced, but only 26% higher without inventories. The reason for this result is that with two means of transferring consumption into the future, fixed capital and inventories, the tight link between aggregate saving and fixed capital investment is broken.

JEL Codes: E20, E22, E30, E32.

Keywords: general equilibrium, lumpy investment, inventories, heterogeneous firms, two-sector model.

*Respectively: University of Notre Dame, CEPR CESifo and ifo (e-mail: rbachman@nd.edu); National University of Singapore (e-mail: ecsml@nus.edu.sg). We thank Francois Gourio and Aubhik Khan for their discussions. We are grateful to participants at the ASSA (2012 meeting in Chicago), the 2012 CIREQ - Montreal conference, and the Royal Economic Society meeting (2014 in Manchester), as well as seminar participants at the University of Michigan and the Munich ifo Institute for their comments. We also thank two anonymous referees and the editor Pok-sang Lam for their comments and suggestions. Any remaining errors are our own.

1 Introduction

Researchers have now explored an ever more detailed and complex set of microeconomic frictions and heterogeneities in macroeconomic models. It has thus become an important question for macroeconomists, who on the one hand strive to build well-microfounded models, but are also, on the other hand, concerned about tractability and complexity of their models, how microeconomic frictions and heterogeneity affect macroeconomic dynamics. Caplin and Spulber (1987) present a striking example where any degree of nominal price stickiness at the micro level is consistent with the same aggregate outcome, money neutrality. In such a case, macroeconomic researchers arguably need not bother with the details of the microfoundation.

Conceptually, typical macroeconomic general equilibrium models can be split into a decision theoretic part where economic agents make often complex and dynamic decisions, which are, potentially, subject to a host of microeconomic frictions, e.g., physical adjustment frictions, informational frictions, etc. The second part of these models then consists of a formulation of aggregate resource and consistency constraints that will lead to the coordination of the individual decisions through prices (e.g., in Walrasian models) or aggregate quantities (e.g., in Non-Walrasian models, like search-and-matching models).

In this paper we argue that the answer to the question of how the microfoundations of decisions affect macroeconomic outcomes may depend on modeling choices in the second part, i.e., the details of how exactly general equilibrium closes a given physical environment, a perhaps obvious, but nevertheless underappreciated point. In other words, we will show – in a concrete, realistic and quantitative example – that there can be a cross effect between the general equilibrium part of a macroeconomic model and the mapping from microfoundations of decisions to macroeconomic outcomes.

Our example can simultaneously claim both realism with respect to a large body of microevidence (e.g., Doms and Dunne (1998) and Cooper and Haltiwanger (2006)) and also a certain notoriety in the literature: the debate about the aggregate effects of nonconvex capital adjustment costs. In a seminal paper, Caballero and Engel (1999) argue that nonconvex capital adjustment costs not only are powerful smoothers of aggregate investment, but also help explain certain nonlinearities in aggregate investment fluctuations. These results were produced in a macroeconomic model with essentially no general equilibrium elements, i.e., in a model with only a decision theoretic part that was aggregated by simple summation. In a series of papers, Thomas (2002), Khan and Thomas (2003) and Khan and Thomas (2008) argue, however, that once a general equilibrium part is added to the physical environment in Caballero and Engel (1999) not only do aggregate nonlinearities vanish, but also nonconvex capital adjustment costs have essentially no ability to smooth aggregate investment dynamics over and above what is done by general equilibrium price movements. Models with nonconvex capital adjustment costs thus deliver lumpy investment patterns at the micro level, but feature business cycle statistics that are very close to standard RBC models, once real wages and real interest rates adjust to clear markets.¹

¹Miao and Wang (2014a) provide an exact characterization of an economy with nonconvex fixed capital adjustment costs – namely, constant returns to scale production technology, nonconvex adjustment costs are proportional to the level of existing capital of a production unit, and shocks are small – where the aggregate dynamics are exactly identical to the same economy without nonconvex fixed capital adjustment costs. Conditional on the mentioned assumptions, this result is rather general and does not depend, for instance,

This perhaps somewhat striking irrelevance result in our view is best and most intuitively understood from the first-order conditions of the representative household, which are the same in a frictionless and a lumpy investment model, where the adjustment friction is on the firm side. With a representative household, the intratemporal and intertemporal first-order conditions govern the optimal paths of consumption and labor supply, which in turn govern the optimal paths of output/income and saving in the short run. Thus, the households in a lumpy investment model *would like* to follow the same consumption path as in the frictionless model. The question is, whether they *are able* to do so when adjusting the capital stock is costly. The answer turns out to be yes, as long as the economy can substitute between the extensive and intensive margins of investment (see Gourio and Kashyap (2007) and, ultimately, Caplin and Spulber (1987) for this insight).² To be concrete, after a positive aggregate productivity shock, the economy uses investment to increase consumption in the future. In a frictionless model this is entirely done through the intensive margin of investment: every firm invests a little more. With nonconvex capital adjustment costs this is no longer optimal, instead a few firms invest a lot. The desired amount of delayed consumption is concentrated into a few firms which really need to invest, and the same aggregate saving/investment path as in a frictionless model results. This intuition rests on the assumption that the economy provides only one means of transferring consumption into the future, fixed capital. This is the familiar dual role of fixed capital in standard models: factor of production on the one hand and the only means of saving on the other, which in turn implies the familiar equality between saving and (fixed capital) investment. For the economy as a whole, investment and consumption dynamics are thus tightly linked. However, it is important to realize that this is only one particular way of introducing general equilibrium in a lumpy investment physical environment. There are others conceivable, and in reality an economy may delay consumption through multiple channels. We show that once we introduce multiple channels of investment and thus break the tight link between aggregate consumption and aggregate fixed capital investment, nonconvex adjustment costs and their magnitude matter much more for fixed capital investment dynamics. As has been mentioned above, this paper is about a cross-derivative from how the aggregate resource constraint is formulated to the ability of nonconvex adjustment costs to impact aggregate dynamics.

on the distribution of fixed capital adjustment costs or the presence/absence of persistent idiosyncratic productivity shocks. On the other hand, Miao and Wang (2014b) show that nonconvex fixed capital adjustment costs in the presence of sufficiently large shocks, such as changes to corporate tax rates, do matter for the aggregate effects of such tax policies. In the more numerically focused literature, to which this paper belongs, a similar back-and-forth has been happening: Veracierto (2002) provides for kinked, but convex adjustment cost functions an approximate neutrality result similar to the Khan and Thomas series of papers. House (2014) argues that, as depreciation rates get small and capital goods are long-lived, nonconvex fixed capital adjustment costs become irrelevant for aggregate dynamics. On the other side of the debate are Gourio and Kashyap (2007) and Bachmann et al. (2013), who argue that these irrelevance results are a matter of degree, specific to the calibration strategy used, and inconsistent with some nonlinear aspects of the time series of the aggregate investment rate in the U.S. Recently, Cooper and Willis (2012) and Winberry (2016) have argued that the Khan and Thomas irrelevance results are inconsistent with real interest rate dynamics, and, once interest rate dynamics are modeled more realistically, the behavior of models with nonconvex fixed capital adjustment costs is much closer to the earlier results in, for instance, Caballero and Engel (1999). While different in the modelling strategy these papers reach conclusions similar to ours as far as the aggregate non-neutrality of nonconvex fixed capital adjustment frictions is concerned.

²We present a similar result in Section 5 in a robustness check to our baseline calibration.

The key intuition for this result is the substitution between different investment channels. Viewed from a social planners’ perspective,³ introducing more investment channels offers more margins to smooth households’ consumption, in addition to the extensive/intensive margin choice in fixed capital investment: if adjusting fixed capital is costly, the social planner can use other investment channels to optimally spread consumption over time. As a result, investment in fixed capital will be more sensitive to the level of frictions in capital adjustment.


To be concrete: we investigate the implications of multiple investment vehicles for the “neutrality question” in a quantitative DSGE model. Building on Khan and Thomas (2007), we study a two-sector setting with an intermediate goods sector and a final goods sector. The final goods sector has the opportunity to store, at a cost, the output from the intermediate goods sector as inventories. The incentive to hold inventories is generated by fixed ordering costs for shipments from the intermediate goods to the final goods sector. The intermediate goods sector uses fixed capital as a production factor, whose adjustment, and this is where we deviate from Khan and Thomas (2007), is subject to nonconvex costs. We choose inventories as the second capital type because, 1) it is a highly cyclical component in the national accounts and, 2) it is a natural means to buffer consumption against temporary shocks. Methodologically, our paper provides the first quantitative analysis of how nonconvex capital adjustment frictions impact aggregate dynamics in the presence of capital good heterogeneity.⁴

Figure 1 summarizes the point of the paper in a nutshell. It shows the impulse response functions of fixed capital investment to a one standard deviation productivity shock. The nonconvex fixed capital adjustment costs dampen the initial response of fixed capital investment to a productivity shock by 3.12 percentage points in the presence of inventories (‘Model I1’ versus ‘Model I2’). That is, the ‘no capital adjustment costs’-impact response is approximately 46% higher than the one with capital adjustment costs. In contrast, without inventories nonconvex fixed capital adjustment costs dampen the initial response of fixed capital investment to a productivity shock by only 2.39 percentage points (‘Model NI1’ versus ‘Model NI2’). That is, the ‘no capital adjustment costs’-impact response is only 26% higher than the one with capital adjustment costs. This highlights the aforementioned interaction effect or cross-derivative, namely, that the presence of inventories, a second capital good, will quantitatively affect the difference between a frictionless and a frictional model for fixed capital. In addition, with inventories the response of investment in the model with the baseline level of nonconvex fixed capital adjustment costs is flatter than that in the model without these capital adjustment frictions. This means that with inventories nonconvex capital adjustment costs stretch the propagation of the productivity shock by more than what capital adjustment frictions can do without inventories. We will cast this argument in more quantitative terms below, when we compare autocorrelation coefficients of aggregate fixed capital investment across these various models.

³We use in this paper a decentralized equilibrium model, where prices guarantee the social planners’ optimal allocations, but for the intuition a social planners’ perspective is useful.

⁴A paper related to ours is Fiori (2012), which also features lumpy capital adjustment in a two-sector model, though without inventories, so the focus there is on movements of the relative price of investment, which in our set up is constant by assumption. Another related paper is Berger and Vavra (2015) which features lumpy adjustment in durable goods in the presence of fixed capital.

Figure 1: Impulse Response Function of Fixed Investment



Notes: This figure shows the impulse response functions of fixed capital investment to a one standard deviation aggregate productivity shock in the intermediate goods sector. ‘Model I1’ has the baseline calibrated nonconvex fixed capital adjustment cost parameter and the baseline calibrated inventory order cost parameter. ‘Model I2’ has zero nonconvex fixed capital adjustment cost and the baseline inventory order cost parameter. ‘Model NI1’ has the baseline calibrated nonconvex fixed capital adjustment cost parameter and zero inventories. ‘Model NI2’ has zero nonconvex fixed capital adjustment cost and zero inventories. The difference between the IRFs of ‘Model I2’ and ‘Model I1’ is the effect of nonconvex fixed capital adjustment costs in the presence inventories. The difference between the IRFs of ‘Model NI2’ and ‘Model NI1’ is the effect of nonconvex fixed capital adjustment costs without inventories. There is no need to recalibrate the fixed capital adjustment cost parameter or the inventory order cost parameter, as our calibration targets, being long-run targets, are not sensitive across model specifications (see Table 2 below).

Figure 1 also shows that inventories dampen the impact response of fixed capital investment. With a positive productivity shock the higher demand for consumption transfer into the future can be partially satisfied by inventories, which are now relatively cheap to produce. And this is done the more so, the higher the nonconvex fixed capital adjustment is, i.e., the more costly the usage of fixed capital is: 11.68% impact response versus 9.95% impact response, a 1.7 percentage points difference, in the frictionless fixed capital adjustment model, yet 9.29% impact response versus 6.83% impact response, a 2.5 percentage points difference, in the model with the baseline calibrated nonconvex fixed capital adjustment cost parameter.

Another direct implication of our mechanism is, as we will show, that the households' ability to smooth consumption is enhanced when there are both inventories and fixed capital. In the end, inventories partially offset the hindering effect on consumption smoothing introduced by fixed capital adjustment frictions. As we will show, the impulse response functions of consumption to an aggregate productivity shock from the lumpy investment model and the frictionless adjustment model are very similar when inventories exist. Similarly, the volatility and persistence of aggregate consumption are much less sensitive to fixed capital adjustment frictions in models with inventories.

It is important to reiterate that the particular physical environment we chose – nonconvex capital adjustment costs as the friction and inventories as a way to modify the aggregate resource constraint – are not as important as the general insight here: when aggregate resource constraints and general equilibrium effects are important for aggregate dynamics, the precise details of *how* these general equilibrium effects are introduced into the physical environment, the precise details of *how* the model is closed matter. In the words of Caballero (2010): “But instead, the current core approach of macroeconomics preserves many of the original convenience-assumptions from the research on the periphery⁵ and then obsesses with closing the model by adding artificial factor supply constraints (note that the emphasis is on the word artificial, not on the word constraints).” This paper provides a quantitative analysis of the effects of closing the model in different ways for a specific, but prominent example. Put differently, unlike Khan and Thomas (2008) and Bachmann et al. (2013), who use the standard formulation for the aggregate resource constraint, this is not mainly a paper about the link between microfrictions and aggregate dynamics per se, but rather a paper about how this link is influenced by the formulation of the general equilibrium part of the model, i.e., a cross effect.

The rest of the paper proceeds as follows. Section 2 outlines the model. Section 3 discusses the calibration and model solution. Section 4 presents the results and is followed by a battery of robustness checks (Section 5). Section 6 concludes.

2 The Model

2.1 The Environment

There are three types of agents in the economy: final goods producers, intermediate goods producers and households. The final goods producers use the intermediate goods, of which

⁵Caballero's terminology for the first, decision theoretic part of macro models.

they hold inventories in equilibrium, and labor to produce the final goods.⁶ Final output can be either consumed or invested as fixed capital. The intermediate goods producers combine fixed capital and labor to produce the intermediate goods. Households consume final goods and provide homogeneous labor to both types of producers. They own all the firms. They receive wage and dividend payments from both types of firms and purchase their consumption goods from the final goods producers. All markets are competitive. This environment is identical to the one in Khan and Thomas (2007), arguably the state-of-the-art model in the general equilibrium inventory literature, with the exceptions that the adjustment of fixed capital in the intermediate goods sector is subject to nonconvex costs and the production function in the intermediate goods sector has decreasing returns to scale.

2.1.1 The Final Goods Producers

There is a continuum of final goods producers. They use intermediate goods, m , and labor, n , to produce the final output through a production function $G(m, n)$.⁷ The production function is strictly concave and has decreasing returns to scale. Whenever the final goods producers purchase intermediate goods, they face a fixed cost of ordering and delivery, denoted in units of labor, ϵ . To avoid incurring the fixed cost frequently, the final good producers optimally hold a stock of inventories of the intermediate goods. Denote the inventory level for an individual producer as $s \in \mathbb{R}_+$.

The final goods producers differ in their fixed cost parameter for ordering, $\epsilon \in [0, \bar{\epsilon}]$. In each period, this parameter is drawn independently for every firm from a time-invariant distribution $H(\epsilon)$. At the beginning of the period, a typical final goods firm starts with its stock of inventories, s , inherited from the previous period. It also learns its fixed cost parameter, ϵ . The firm decides whether to order intermediate goods. If the firm does so, it pays the fixed cost and chooses a new inventory level. Otherwise, the firm enters the production phase with the inherited intermediate goods inventory level s . We denote the quantity of adjustment by x_m . The inventory stock ready for production is $s_1 = s + x_m$, with $x_m = 0$ if the firm does not adjust.

After the inventory decision the firm determines its labor input, n , and the intermediate

⁶To be clear on terminology: inventories in this model are not a capital good in the sense that they enter directly a production function, as in some modeling approaches in the literature. Thus, in our model, they lack the dual role of fixed capital. But they are a capital good in the sense that they represent a means of transferring consumption into the future, just like fixed capital. In this sense, we follow the NIPA terminology and denote net inventory changes as investment and the corresponding stock variables as capital.

⁷Following Khan and Thomas (2007) we abstract from fixed capital in the final goods sector, which would complicate the model and its computation considerably. In addition, the BEA Fixed Assets Accounts by Industry tables (Table 3.1E and Table 3.1F) and the BEA Value Added by Industry table reveals that from 1997 to 2006 in the manufacturing sector, the sector arguably closest to the intermediate goods sector in the model, capital intensity was 1.13 versus 0.77 in the trade sector, the sector arguably closest to the final goods sector in the model. So, the importance of fixed capital is presumably lower in the trade than in the manufacturing sector. Moreover, two thirds of total capital in the trade sector is structures, whereas only 43% of total capital is structures in the manufacturing sector. While we do not take an explicit stand, the producers in our model economy are closer to individual establishments than firms. Our assumption of no fixed capital in the final goods sector is then essentially an assumption on how fixed and slowly-depreciating a production factor structures are at the business cycle frequency in the trade sector from the viewpoint of an individual producer unit, which, we believe, is a sensible first pass.

goods input, $m \in [0, s_1]$, for current production. Intermediate goods are used up in production. The remaining stock of intermediate goods, $s' = s_1 - m \geq 0$, is the starting stock of inventories for the next period. Stored inventories incur a unit cost of σ , denoted in units of final output. Inventory holding costs capture the idea that the storage technology that is used to partially circumvent the costly shipping technology is not free. Inventories require storage places, management and can lead to destruction of intermediate goods. The inventory management of the final good firms balances the trade-offs between costly shipping and costly storing optimally. In the end, the output of a typical final firm is $y = G(m, n) - \sigma s'$.

2.1.2 Intermediate Goods Producers

There is a continuum of intermediate goods producers. They are subject to an aggregate productivity shock, which, in the baseline model, is the sole source of aggregate uncertainty.⁸ Let z denote the aggregate productivity level. It follows a Markov chain, $z \in \{z_1, \dots, z_{N_z}\}$, where $P(z' = z_j | z = z_i) = \pi_{ij} \geq 0$ and $\sum_{j=1}^{N_z} \pi_{ij} = 1$ for all i .

Each firm produces with fixed capital and labor. Whenever the firm decides to adjust its capital stock, it has to pay a fixed cost, denoted in units of labor. In each period, the cost of adjusting capital is drawn independently for every firm from a time-invariant distribution $I(\zeta)$. A typical intermediate good producer is identified by its capital stock, k , and its cost of adjusting capital, $\zeta \in [0, \bar{\zeta}]$.

At the beginning of each period, the firm learns aggregate productivity, z , and its idiosyncratic cost of adjusting capital, ζ . It starts with a fixed capital stock, k , inherited from the previous period. First, it decides about the labor input, l . It combines l and k according to a production function $zF(k, l)$. The $F(\cdot)$ function is strictly concave and has decreasing returns to scale.⁹ After production, the firm chooses whether to adjust its capital stock. It can pay a fixed cost to adjust its capital stock by investing i . In this case, the new capital stock for the next period in efficiency units is $k' = [(1 - \delta)k + i]/\gamma$, where δ is the depreciation rate and γ is the steady state growth rate of the economy. Alternatively, the firm can avoid the adjustment cost and start the next period with the depreciated capital stock $k' = (1 - \delta)k/\gamma$.

⁸As pointed out in Khan and Thomas (2007), placing aggregate productivity in the intermediate sector is necessary in this physical environment to generate a countercyclical relative price of intermediate goods, a feature found in the U.S. data. In one of the robustness checks in Section 5 we do, however, consider the case where aggregate productivity shocks hit the final goods producers. We also abstract, in the baseline setting, for reasons of computability, from the persistent idiosyncratic productivity shocks that the recent literature (see Khan and Thomas (2008) and Bachmann et al. (2013)) has used to explain the observed micro-level heterogeneity in the data. Given the two firm problems, the computational burden in our model is already high. Nevertheless, we study in Section 5 the case with persistent idiosyncratic productivity shocks as a robustness check.

⁹As Miao and Wang (2014a) shows, fixed adjustment costs cannot be expected to have a large impact with a constant return to scale production technology. We follow the majority of the literature, e.g., Bachmann et al. (2013), Bloom (2009), Khan and Thomas (2008), Gourio and Kashyap (2007) as well as Cooper and Haltiwanger (2006), and use a decreasing returns to scale assumption.

2.1.3 Households

We assume a continuum of identical households who value consumption and leisure. They have access to a complete set of state-contingent claims. Households own all the firms. They provide labor to the firms and receive wage and dividend payments.

The households have the following felicity function:

$$u(c, n^h) = \log c - A^h n^h,$$

where n^h is the total hours devoted to market work.

2.2 Competitive Equilibrium

2.2.1 Aggregate State Variables

In addition to z , the aggregate productivity level, two endogenously determined distributions are aggregate state variables in this model: the distribution of the firm-specific inventory stocks, $\mu(S)$, and the distribution of firm-specific fixed capital stocks, $\lambda(K)$. Both S and K are subsets of a Borel algebra over \mathbb{R}_+ .

The aggregate state variables are summarized as (z, A) , where $A = (\mu, \lambda)$. The distribution of μ evolves according to a law of motion $\mu' = \Gamma_\mu(z, A)$, and similarly, the distribution of λ evolves according to $\lambda' = \Gamma_\lambda(z, A)$.

The final good is the numeraire. Workers are paid $\omega(z, A)$ per unit of labor input. The intermediate goods are traded at $q(z, A)$ per unit.

2.2.2 Problem of the Household

The households receive a total dividend payment $D(z, A)$ and labor income $n^h(z, A)\omega(z, A)$ from the firms. In each period the households determine how much to work and how much to consume. All we need from the household problem is an intertemporal and an intratemporal first-order condition.

We can express the dynamic programming problems for both types of firms with the marginal utility of consumption as the pricing kernel:

$$p(z, A) = \frac{1}{c(z, A)}.$$

Then every firm weighs its current profit by this pricing kernel and discounts its future expected earnings by β . This changes the unit of the firm's problems in both sectors to utils but leaves the policy functions unchanged.

The first-order conditions also imply that the real wage is given by:

$$\omega(z, A) = \frac{A^h}{p(z, A)}.$$

2.2.3 Problem of Final Goods Producers

Let V_0 be the value, in utils, of a final goods producer at the beginning of a period after the inventory adjustment cost parameter is realized and before any inventory adjustment and production decisions. Let V_1 be the expected value function after the adjustment decision but before the production decision. Given the aggregate laws of motion Γ_μ and Γ_λ , the firm's problem is characterized by the following three equations. For expositional ease, the arguments for functions other than the value functions are omitted.

$$V_0(s, \epsilon; z, A) = pqs + \max \left\{ -p\omega\epsilon + V_a(z, A), -pqs + V_1(s; z, A) \right\}, \quad (1)$$

$$V_a(z, A) = \max_{s_1 > 0} \{-pqs_1 + V_1(s_1; z, A)\}, \quad (2)$$

and:

$$V_1(s_1; z, A) = \max_{n \geq 0, s_1 \geq s' \geq 0} \left\{ p[G(s_1 - s', n) - \sigma s' - \omega n] + \beta E_z \left[\int_0^{\bar{\epsilon}} V_0(s', \epsilon; z', A') d(H(\epsilon)) \right] \right\}. \quad (3)$$

The expectation is taken over z' , next period's aggregate productivity.

Equation (1) describes the binary inventory adjustment decision of the firm. The firm adjusts if the value of entering the production phase with the optimally adjusted inventory level, described by $V_a(\cdot)$ in equation (2), minus the cost of adjustment, exceeds the value of directly entering the production phase with the inherited inventory level, $V_1(s; z, A)$.

The solution to equation (1) amounts to a cut-off rule in ϵ . The firm adjusts if:

$$-p\omega\epsilon + V_a(z, A) \geq -pqs + V_1(s; z, A).$$

Therefore the cut-off value is:

$$\tilde{\epsilon}(s; z, A) = \frac{V_a(z, A) - V_1(s; z, A) + pqs}{p\omega}.$$

Given the support of the adjustment cost distribution, this cut-off value is modified to:

$$\epsilon^* = \max(0, \min(\bar{\epsilon}, \tilde{\epsilon})).$$

The firm adjusts if its draw is smaller than or equal to $\epsilon^*(s; z, A)$.

Equation (2) describes the value of inventory adjustment. The solution to this equation is the optimal target level of inventory, s_1^* . Note that the optimization problem in equation (2) does not depend on any firm-specific characteristics. Therefore in any period, all the adjusting firms choose the same inventory target level, $s_1^*(z, A)$.

Equations (1) and (2) jointly determine the production-time inventory level, s_1 :

$$s_1(s, \epsilon; z, A) = \begin{cases} s_1^*(z, A) & \text{if } \epsilon \leq \epsilon^*(s; z, A) \\ s & \text{if } \epsilon > \epsilon^*(s; z, A) \end{cases}.$$

Notice that from $s_1(s, \epsilon; z, A)$ the quantity of intermediate goods ordered, the order flows, is given by $x_m(s, \epsilon; z, A) \equiv s_1(s, \epsilon; z, A) - s$.

Equation (3) describes the production phase. The firm finds the optimal inventory level for the next period and the optimal employment level for this period. The decision for next period's inventory level, s' , is equivalent to deciding about the amount of intermediate goods to be used up in current production.

The solution for employment does not depend on the continuation value function. Therefore, given s' , it is the analytical solution to:

$$\frac{\partial G(s_1 - s', n^*)}{\partial n} = \omega.$$

The optimal employment and inventory usage decision jointly imply the optimal output level:

$$y^*(s_1; z, A) = G(s_1 - s'^*(s_1; z, A), n^*(s_1; z, A)) - \sigma s'^*(s_1; z, A).$$

2.2.4 Problem of the Intermediate Goods Producers

Let W_0 be the value, in utils, of the intermediate good producers prior to the realization of the adjustment cost parameter ζ . Let W_1 be the value function after the realization of ζ . The intermediate good producer's problem can be summarized by the following equation:

$$W_1(k, \zeta; z, A) = \max_l \left\{ p \cdot [q \cdot zF(k, l) - l\omega] + \max \{ W_i(k; z, A), -p\zeta\omega + W_a(k; z, A) \} \right\}, \quad (4)$$

where:

$$W_a(k; z, A) = \max_{k'} \{ -(\gamma k' - (1 - \delta)k)p + \beta E_z [W_0((k'; z', A'))] \}, \quad (5)$$

$$W_i(k; z, A) = \beta E_z [W_0((1 - \delta)k/\gamma; z', A')], \quad (6)$$

$$W_0(k; z, A) = \int_0^{\bar{\zeta}} W_1(k, \zeta; z, A) d(I(\zeta)). \quad (7)$$

The expectation in equation (5) and (6) is taken over z' , next period's aggregate productivity.

In equation (4), the firm first solves for the optimal employment, given the fixed capital stock. The solution is:

$$\frac{\partial qzF(k, l^*)}{\partial l} = \omega.$$

After the production decision, the firm solves the binary fixed capital adjustment decision. The firm adjusts if the expected value from the optimally adjusted fixed capital stock, given in equation (5), minus the cost of adjustment, exceeds the expected value from the unadjusted fixed capital stock, given in equation (6).

The solution to the adjustment decision follows a cut-off rule for ζ . The firm adjusts if:

$$-p\omega\zeta + W_a(k; z, A) \geq W_i(k; z, A).$$

Therefore the cut-off value for ζ is:

$$\tilde{\zeta}(k; z, A) = \frac{W_a(k; z, A) - W_i(k; z, A)}{p\omega}.$$

The restriction from the support of the cost distribution applies, so that

$$\zeta^* = \max(0, \min(\bar{\zeta}, \tilde{\zeta})).$$

The firm adjusts to the target capital stock if its adjustment cost is smaller than or equal to $\zeta^*(k; z, A)$.

The optimal adjustment target for fixed capital is given by the solution to equation (5). Although the value function depends on the level of individual capital stocks, the resulting policy function, k^* , does not. After the binary adjustment decision, the capital stock for the next period is:

$$k'(k; z, A) = \begin{cases} k^*(z, A) & \text{if } \zeta \leq \zeta^*(k; z, A) \\ (1 - \delta)k/\gamma & \text{if } \zeta > \zeta^*(k; z, A) \end{cases}.$$

2.2.5 Recursive Equilibrium

A recursive competitive equilibrium for the economy defined by:

$$\left\{ u(c, n^h), \beta, F(k, l), G(m, n), \sigma, \delta, \gamma, H(\epsilon), I(\zeta), \{\pi_{ij}\} \right\},$$

is a set of functions:

$$\{V_0, V_1, W_0, W_1, x_m, n, s', k', l, i, c, n^h, p, q, \omega, D, \Gamma_\mu, \Gamma_\lambda\},$$

such that:

1. Given ω, q, p, Γ_μ and Γ_λ, V_0 and V_1 solve the final good firm's problem with policy functions x_m, n, s' .
2. Given ω, q, p, Γ_μ and Γ_λ, W_0 and W_1 solve the intermediate good firm's problem with policy functions k', l, i .
3. Given ω, D and p, c satisfies the household's first-order conditions.

4. The final goods market clears:

$$c(z, A) = \int_S \int_0^{\bar{\epsilon}} y(s, \epsilon; z, A) d(H(\epsilon)) d(\mu(s)) \\ - \int_K \int_0^{\bar{\zeta}} i(k, \zeta; z, A) d(I(\zeta)) d(\lambda(k)).$$

5. The intermediate goods market clears:

$$\int_S \int_0^{\bar{\epsilon}} x_m(s, \epsilon; z, A) d(H(\epsilon)) d(\mu(s)) = \\ \int_K \int_0^{\bar{\zeta}} zF(k, n(k, \zeta; z, A)) d(I(\zeta)) d(\lambda(k)).$$

6. The labor market clears:

$$n^h(z, A) = \int_S \int_0^{\bar{\epsilon}} (n(s; z, A) + \epsilon \cdot \mathbf{1}(x_m(s, \epsilon; z, A) \neq 0)) d(H(\epsilon)) d(\mu(s)) \\ + \int_K \int_0^{\bar{\zeta}} (l(k, n(k; z, A)) + \zeta \cdot \mathbf{1}(i(k, \zeta; z, A) \neq 0)) d(I(\zeta)) d(\lambda(k)).$$

7. The laws of motion for aggregate state variables are consistent with individual decisions and the stochastic processes governing z :

- (a) $\Gamma_\mu(z, A)$ defined by $s'(s, \epsilon; z, A)$ and $H(\epsilon)$;
- (b) $\Gamma_\lambda(z, A)$ defined by $k'(k, \zeta; z, A)$ and $I(\zeta)$.

2.2.6 Some Terminology

Final Sales (FS), is defined as the total output of the final goods sector. Intermediate goods demand, X, is the total amount of intermediate goods purchased by the final goods sector. Intermediate goods usage, M, is the total amount of intermediate goods used up in production by the final goods sector. The difference between the two evaluated at the relative price of intermediate goods is Net Inventory Investment (NII):

$$\text{NII} = q \times (\text{X} - \text{M}).$$

Finally, Gross Domestic Product (GDP) in this environment is defined as the sum of final sales and net inventory investment:

$$\text{GDP} = \text{FS} + \text{NII}.$$

3 Calibration and Computation

3.1 Baseline Parameters

The model period is a quarter. We choose the following functional forms for the production functions:

$$F(k, l) = k^{\theta_k} l^{\theta_l},$$

$$G(m, n) = m^{\theta_m} n^{\theta_n}.$$

We discretize the productivity process z into $N_z = 11$ points following Tauchen (1986). The underlying continuous productivity process follows an AR(1) in natural logarithms. We do not take a strong stance on the persistence of the auto-correlation parameter of this process, but instead choose two values, $\rho_z = 0.95$ as a lower bound and baseline case, and $\rho_z = 0.98$ as an upper bound in the robustness checks in Section 5.¹⁰ We then calibrate σ_z so as to make the model match the volatility of U.S. GDP, 1.66%.

We set the subjective discount factor, $\beta = 0.984$, the depreciation rate $\delta = 0.017$, and the steady state growth factor $\gamma = 1.004$. A^h is calibrated so that the aggregate labor input equals 0.33. $\theta_m = 0.5245$ is calibrated so as to make the model match the share of intermediate inputs in final output, 0.499, as reported by Khan and Thomas (2007). In the baseline case, we set $\theta_k = 0.25$ and $\theta_l = 0.5$, the values used in Bloom (2009), which amounts to a capital elasticity of the firms' revenue function, $\frac{\theta_k}{1-\theta_l}$, of 0.5.¹¹ We finally calibrate θ_n to match an aggregate labor share of 0.64. All these parameters are summarized in Table 1:

Table 1: Baseline Parameters

β	A^h	θ_m	θ_n	θ_k	θ_l	ρ_z	σ_z	δ	γ
0.9840	2.0720	0.5245	0.3530	0.2500	0.5000	0.9500	0.0167	0.0170	1.0040

Notes: β is the subjective discount factor of the households; A^h is the preference parameter for leisure; θ_m is the elasticity of materials in the final goods production function; θ_n is the elasticity of labor in the final goods production function; θ_k is the capital elasticity in the intermediate goods production function; θ_l is the labor elasticity in the intermediate goods production function; ρ_z is the auto-correlation for the aggregate productivity process; σ_z is the standard deviation for aggregate productivity innovations; δ is the depreciation rate; γ is the steady state growth rate.

3.2 Inventory and Adjustment Cost Parameters

We assume that the inventory adjustment costs, ϵ , are uniformly distributed, $H(\epsilon)$, on $[0, \bar{\epsilon}]$. $\bar{\epsilon}$ is set so that the average inventory-to-sales ratio in the model equals 0.8185, the average of the real private non-farm inventory-to-sales ratio in the United States between 1960:1 and

¹⁰We thus bracket not only the value used by Khan and Thomas (2007), $\rho_z = 0.956$, but also the canonical values used in business cycle research and published in seminal handbook articles, albeit in one-sector models; see Cooley and Prescott (1995), $\rho_z = 0.95$, and King and Rebelo (1999), $\rho_z = 0.979$.

¹¹Cooper and Haltiwanger (2006), using LRD manufacturing data, estimate this parameter to be 0.592; Hennessy and Whited (2005), using Compustat data, find 0.551. Gourio and Kashyap (2007) use a much lower capital elasticity, while Bachmann et al. (2013) calibrate also to 0.5, but use a different distribution onto θ_k and θ_l (in their inventory-focused paper Khan and Thomas (2007) assume constant returns to scale in the intermediate goods sector). We conduct a robustness check with respect to the calibration of θ_k and θ_l in Section 5 allowing for a production function that is closer to the linear case.

2006:4. The unit cost of holding inventories, σ , is chosen so that the annual storage cost for all inventories is 12% of aggregate final output in value (see Richardson (1995) for details). These two targets jointly determine $\bar{\epsilon} = 0.4580$ and $\sigma = 0.0128$.

We assume that $I(\zeta)$, the distribution of fixed capital adjustment costs, is uniform between $[0, \bar{\zeta}]$. The upper bound of the distribution is chosen so that the fraction of investors with lumpy fixed capital adjustments, defined as a gross investment rate larger than 20% in a given year, is 18%. This calibration target is taken from Cooper and Haltiwanger (2006)'s analysis of manufacturing firms in the Longitudinal Research Database (LRD). This yields $\bar{\zeta} = 0.1950$.¹²

3.3 Numerical Solution

The inherent nonlinearity of the model suggests global numerical solution methods. We use value function iterations from equation (1) to equation (3) to solve the problem of the final good producers. We use value function iterations from equation (4) to equation (7) to solve the intermediate good firm's problem. Howard policy function accelerations are used to speed up convergence (see Appendix A.1 for the details of the numerical implementation).

Our model gives rise to two endogenous distributions as state variables. We adopt the methods in Krusell and Smith (1997), Krusell and Smith (1998), Khan and Thomas (2003) as well as Khan and Thomas (2008) to compute the equilibrium. Denote the I th moment of distribution $\mu(S)$ and $\lambda(K)$ as $\mu_I(S)$ and $\lambda_I(K)$ respectively. We approximate each distribution function with its first moment. We find that a log-linear form for the $\Gamma(\cdot)$ functions approximates the law of motion rather well in terms of forecasting accuracy:

$$\log \mu'_1 = \alpha_\mu + \gamma_\mu \log(\mu_1) + \psi_\mu \log(z), \quad (8)$$

$$\log \lambda'_1 = \alpha_\lambda + \beta_\lambda \log(\lambda_1) + \psi_\lambda \log(z). \quad (9)$$

We adopt similar rules for the pricing kernel and the relative price of intermediate goods:¹³

$$\log p = \alpha_p + \beta_p \log(\lambda_1) + \psi_p \log(z), \quad (10)$$

$$\log q = \alpha_q + \beta_q \log(\lambda_1) + \psi_q \log(z), \quad (11)$$

¹²It should be clear that the exact numbers for $\bar{\epsilon}$ and $\bar{\zeta}$ have little direct economic meaning and cannot be compared to other calibrations for these parameters in the literature. They are essentially free parameters to hit observable calibration targets (which are what is common across papers), such as the inventory-to-sales ratio and the fraction of firms that have lumpy fixed capital investments. They will also lead to additional interpretable economic statistics like the average adjustment cost paid conditional on adjustment that we display below in Table 2. The precise values of these parameters are dependent on the entire model environment and its calibration.

¹³Since the cross-sectional averages of the capital stock and the inventory holdings distributions turn out to be highly correlated in the model's equilibrium, and thus including them both would lead to a multicollinearity problem, we use the average capital holdings in its own forecast equation and the two pricing equations, and the average inventory holdings only in its own forecast equation. We have also experimented with other functional forms for the forecasting rules such as adding interaction terms between aggregate productivity and the capital and inventory moments. This did not lead to significant improvements in goodness-of-fit and often jeopardized numerical stability. Our specifications perform very well as measured by the R^2 of the equilibrium OLS regressions, which exceeds 0.998 in all specifications.

where λ_1 is the first moment of the capital stock distribution, and μ_1 is the first moment of the inventory stock distribution.

Given an initial guess for $\{\alpha_{\cdot}, \beta_{\cdot}, \gamma_{\cdot}, \psi_{\cdot}\}$, we solve the value functions as described above. Then we simulate the model without imposing the pricing rules in equations (10) and (11). In each model simulation period we search for a pair of prices, (p, q) such that all the firms optimize and all the markets clear under the forecasting rules in equation (8) and (9). To guarantee numerical accuracy, we use the value functions to re-solve all the optimization problems period by period and for every guess of (p, q) . Given the market clearing prices, we update the capital and inventory stock distributions and proceed into the next period.

At the end of the simulation, we update the parameters $\{\alpha_{\cdot}, \beta_{\cdot}, \gamma_{\cdot}, \psi_{\cdot}\}$ using the simulated time series for the approximating moments and the market clearing prices. Then we repeat the algorithm with the updated parameters. Upon convergence of the parameters, we check the accuracy of the $\Gamma(\cdot)$ functions by the R^2 in the regression stage (see Appendix A.2 for additional accuracy checks on the numerical solution).

4 Results

We study the influence of nonconvex fixed capital adjustment costs on aggregate dynamics in our model by numerical simulation. The model with the calibration above is denoted by ‘Model I1’. In addition, we analyze three models that share all parameters with ‘Model I1’ and each other, except for $\bar{\epsilon}$ and $\bar{\zeta}$. ‘Model I2’ has the calibrated baseline equilibrium inventory holdings with $\bar{\epsilon} = 0.4580$, but features a frictionless technology for adjusting the fixed capital stock, $\bar{\zeta} = 0$. We also simulate two models without inventories, ‘Model NI1’ and ‘Model NI2’. In these models, we set $\bar{\epsilon} = 0$ to eliminate equilibrium inventory holdings.¹⁴ ‘Model NI1’ has the same level of $\bar{\zeta}$ as ‘Model I1’, while ‘Model NI2’ does not feature frictions in adjusting the fixed capital stock. The parameter specifications for the four models are summarized in Table 2. We do not recalibrate $\bar{\zeta}$ in ‘Model NI1’ as the calibration targets are largely insensitive to the changes in equilibrium inventory levels. To understand how the presence of inventories interacts with the effects of nonconvex fixed adjustment costs, we study the cross differences. That is, we contrast the differences between ‘Model I1’ and ‘Model I2’ with the differences between ‘Model NI1’ and ‘Model NI2’.

Table 2: Model Specifications

Model Name	$\bar{\zeta}$	$\bar{\epsilon}$	Average Adjustment Cost	Fraction of Investors with Lumpy Adjustments	Note
I1	0.1950	0.4580	0.9100%	18.00%	Baseline fixed capital adjustment cost with inventory
I2	0.0000	0.4580	0.0000%	0.000%	Frictionless fixed capital adjustment with inventory
NI1	0.1950	0.0000	0.8700%	18.18%	Baseline fixed capital adjustment cost without inventory
NI2	0.0000	0.0000	0.0000%	0.000%	Frictionless fixed capital adjustment without inventory

Notes: ‘Model I1’ has the baseline calibrated nonconvex fixed capital adjustment cost parameter and the baseline calibrated inventory order cost parameter. ‘Model I2’ has zero nonconvex fixed capital adjustment cost and the baseline inventory order cost parameter. ‘Model NI1’ has the baseline calibrated nonconvex fixed capital adjustment cost parameter and zero inventories. ‘Model NI2’ has zero nonconvex fixed capital adjustment cost and zero inventories. ‘Average Adjustment Cost’ is the average adjustment cost paid as a fraction of firms’ output, conditional on adjustment. ‘Fraction of Investors with Lumpy Adjustments’ is the share of firms that adjust more than 20% of their initial capital stocks in a given year.

¹⁴In theory, zero ordering costs are not inconsistent with positive inventory holdings as the firms might want to hedge against changes in the relative price of intermediate goods. However, in our simulations, given the inventory holding costs, no firm holds a positive level of inventories when $\bar{\epsilon} = 0$.

We present four sets of results on those four models. We first compare their unconditional business cycle moments. Second, we study the impulse response functions for fixed capital investment and consumption across the four models. Third, we plot the volatility and persistence for consumption, fixed capital investment and, for the models with inventories, net inventory investment for a wider range of $\bar{\zeta}$. And finally, we analyze the role of general equilibrium price movements in bringing about these results.

4.1 Unconditional Business Cycle Analysis

After computing the equilibrium, we simulate the model for 1,000 periods, of which we discard the first 100 to eliminate the influence of initial conditions. Except for net inventory investment and fixed capital investment, all the simulated time series are transformed by natural logarithms and then detrended by an HP filter with smoothing parameter 1600. We detrend fixed capital investment with the HP filter directly and then divide the deviations by the trend. We divide net inventory investment by GDP and then apply the HP filter to this ratio.

Table 3: Business Cycle Statistics

(a) Standard Deviation					
	GDP	Consumption	Fixed Investment	NII/GDP	Inventory Level
Model I1	1.664	0.720	10.299	0.411	1.343
Model I2	1.742	0.713	12.469	0.338	1.232
Model NI1	1.645	0.844	12.521	-	-
Model NI2	1.761	0.821	15.090	-	-
Data	1.663	0.901	4.890	0.422	1.655
(b) First Order Auto-correlation					
	GDP	Consumption	Fixed Investment	NII/GDP	Inventory Level
Model I1	0.687	0.760	0.746	0.612	0.930
Model I2	0.665	0.794	0.608	0.681	0.943
Model NI1	0.688	0.727	0.676	-	-
Model NI2	0.673	0.777	0.639	-	-
Data	0.842	0.883	0.901	0.370	0.891

Notes: "NII" denotes net inventory investment. GDP, consumption, and inventory levels are logged and detrended with an HP filter with a penalty parameter of 1600. We detrend fixed investment with the HP filter and then divide the deviations by the trend. We divide NII by GDP and then detrend this ratio with the HP filter. All the standard deviations reported in Panel (a) are percentage points. Time period for the data moments: 1960:1 - 2006:4.

The business cycle statistics in Panel (a) and (b) of Table 3 show several effects of inventories on aggregate dynamics.¹⁵ The first message is that nonconvex fixed capital adjustment costs matter for aggregate dynamics. Business cycle dynamics differ significantly between

¹⁵Bachmann et al. (2013) is explicitly about how nonconvex fixed capital adjustment costs shape the implied model investment dynamics in terms of higher than standard second moments. They argue that aggregate investment data exhibits conditional heteroskedasticity and that micro nonconvexities are a natural mechanism to explain this. In contrast, this paper is about micro nonconvexities and their role in shaping standard second moments and impulse response functions. Basically, this paper asks: is the fixed capital adjustment technology that is consistent with the micro data able to do what stand-in adjustment technologies do, namely, dampen and propagate aggregate investment.

‘Model I1’ and ‘Model I2’. For example, the percentage standard deviation of fixed capital investment decreases from 12.47 in the frictionless ‘Model I2’ to 10.30 in the lumpy investment ‘Model I1’. Persistence of fixed capital investment increases from 0.61 to 0.75. In contrast, consumption volatility and persistence move in the opposite way as investment volatility and persistence. Consumption is more volatile and less persistent with fixed capital adjustment costs, because the ability to use fixed capital as a means of consumption smoothing is hampered.¹⁶

Regarding the cross differences, the effects of nonconvex fixed capital adjustment costs change significantly in models where inventories are absent. Most notably, the persistence of fixed investment only increases by 0.04 between ‘Model NI2’ and ‘Model NI1’, while it increases by 0.14 between ‘Model I2’ and ‘Model I1’, bringing the model halfway, from 0.61 to 0.75, to the persistence of investment in the data, 0.90.

The unconditional volatility of consumption increases by 0.023 percentage points between ‘Model NI2’ and ‘Model NI1’, while it only increases by 0.007 percentage points between ‘Model I2’ and ‘Model I1’.¹⁷ The persistence of consumption decreases by 0.050 between ‘Model NI2’ and ‘Model NI1’, while it only decreases by 0.034 between ‘Model I2’ and ‘Model I1’. Thus, consumption dynamics are more insulated from variations in capital adjustment frictions in the presence of inventories.

These results suggest that inventories strengthen the propagation effect of fixed adjustment costs on fixed capital investments.¹⁸ At the same time, inventories enhance the households’ ability to smooth consumption, making fixed capital adjustment costs much less effective in affecting consumption volatility and persistence.

As for net inventory investment and the level of inventories, we see that they behave exactly the opposite way from fixed capital investment, when the latter is subject to adjustment frictions. Their volatility rises and their persistence falls, when capital adjustment frictions are introduced. This is due to the substitution towards inventories as a means of consumption smoothing, as fixed capital becomes more costly to use.

¹⁶The excessively high fixed investment volatility, as shown in the third column of Panel (a), is a common property of two-sector models where fixed capital is only used in intermediate goods production. Khan and Thomas (2007) find similar results. As fixed adjustment cost works to dampen investment volatility, this might point to our calibration of $\bar{\zeta}$ being conservative, especially in light of the insights of Bachmann et al. (2013), who argue that focusing only on the fraction of lumpy investment episodes when calibrating nonconvex adjustment costs might lead to a downward biased estimate.


¹⁷Looking at unconditional volatility is not ideal, however. The unconditional volatility numbers are a combination of changes in persistence and changes in conditional volatility, which is why we focus on the latter two in what follows.

¹⁸Note that already without inventories we have that nonconvex fixed capital adjustment costs matter somewhat for aggregate dynamics as, in line with the recent evidence in Bloom (2009) and Cooper and Haltiwanger (2006), our implied revenue elasticity of capital is closer to the calibration in Gourio and Kashyap (2007), where the substitution between the extensive and intensive margin of fixed capital investment is more difficult.

4.2 Conditional Business Cycle Analysis - Impulse Response Functions

The first two panels of Figure 2 show the impulse response functions of aggregate fixed capital investment and consumption to a positive, one standard deviation productivity shock in the intermediate goods sector (see Appendix A.3 for the details of how these impulse response functions are computed).

Figure 2: Impulse Response Functions



Notes: This figure shows the impulse response functions of fixed capital investment, consumption, net inventory investment (NII) over GDP and the relative price to a one standard deviation aggregate productivity shock in the intermediate goods sector. ‘Model I1’ has the baseline calibrated nonconvex fixed capital adjustment cost parameter and the baseline calibrated inventory level. ‘Model I2’ has zero nonconvex fixed capital adjustment cost and the baseline calibrated inventory level. ‘Model NI1’ has the baseline calibrated nonconvex fixed capital adjustment cost parameter and zero inventories. ‘Model NI2’ has zero nonconvex fixed capital adjustment cost and zero inventories. The impulse response of net inventory investment over GDP is reported in absolute values, instead of percentage points, as the steady state value of net inventory investment is zero.

Fixed Capital Investment Panel(a) of Figure 2 presents the four impulse response functions for fixed capital investment. Comparing the models with $\bar{\zeta} = 0.1950$ against the models

with $\bar{\zeta} = 0$ at the same level of inventories, we can see that nonconvex fixed capital adjustment costs dampen the initial responses both with and without inventories. However, at different levels of inventories, capital adjustment costs dampen these responses to a different degree. Without inventories, the initial response is dampened by 2.39 percentage points. In contrast, the initial response is dampened by 3.12 percentage points in models with inventories. Inventories also increase shock propagation. Comparing the impulse response function of ‘Model I1’ with that of ‘Model NII’ without inventories, we see that the impulse response function in the model with inventories is flatter.

Both the extra dampening effect and the increased propagation of the shocks come from the key mechanism in our model: the substitution between fixed capital investment and inventory investment as a means of consumption smoothing. When adjusting fixed capital is costly, the economy switches to inventories. As a result, fixed capital investments do not need to respond to productivity shocks as much as when inventories are absent. The responses are also more protracted because firms tend to wait for lower adjustment cost draws to invest.

The flip side of the substitution between the two investment means can be observed in Panel(c) of Figure 2, which shows the impulse response functions of net inventory investment (over GDP). As expected, the response of net inventory investment is stronger when adjusting fixed capital investment is costly. In ‘Model I1’, the impact response is roughly 0.0032, while in ‘Model I2’ it is only 0.0024.¹⁹

The same mechanism can also explain the other cross effect, namely, how lumpy fixed capital investment changes the effect of inventories on aggregate investment dynamics. For both levels of fixed capital adjustment costs, inventories dampen the positive response of fixed capital investment to a positive productivity shock, as the latter is no longer used as much to ensure consumption smoothing. This switching away from fixed capital investment as a means of transferring consumption into the future is stronger, the more costly it is to adjust fixed capital. This explains why inventories dampen the initial response of fixed capital investment by somewhat over 2.5 percentage points with fixed capital adjustment frictions, but only by 1.7 percentage point, when fixed capital can be freely adjusted.

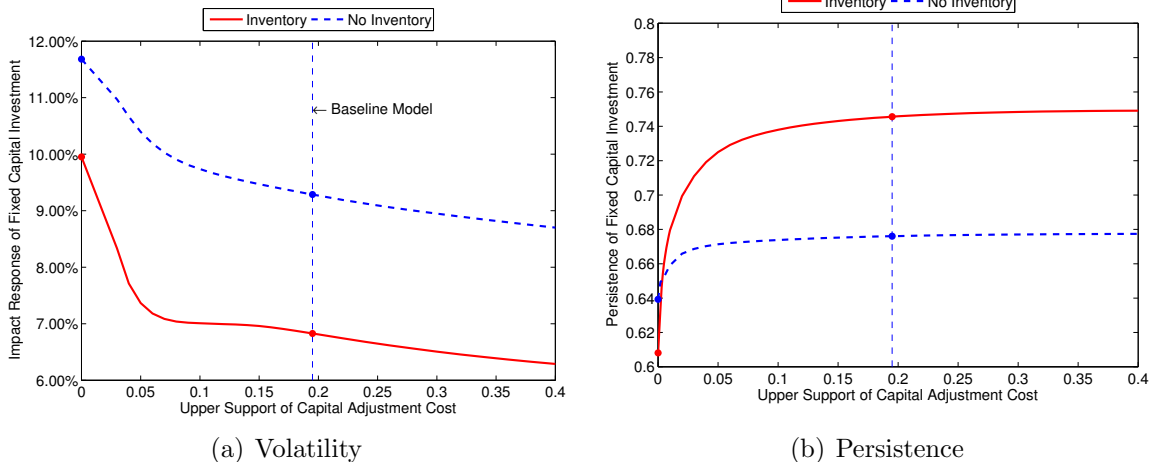
Consumption Another implication from the above mechanism is that consumers’ ability to smooth consumption is enhanced by inventories. We illustrate this with the impulse response functions for consumption in Panel(b) of Figure 2.

First, the impact response from the models with inventories is below those from the models without inventories, for every level of fixed capital adjustment costs. Secondly, the smoothing effectiveness of inventories is so good that consumers despite the presence of capital adjustment costs can almost exactly recreate their frictionless consumption path. Nonconvex fixed capital adjustment costs barely change the response of consumption after the initial impact, when there are inventories. In contrast, without inventories nonconvex fixed capital adjustment costs do interfere with consumption smoothing.

We interpret these response functions as evidence that inventories provide an effective smoothing device for the consumers. As a result, consumption dynamics are less volatile when productivity shocks hit and capital adjustment frictions are less relevant for consumption dynamics in the presence of inventories.

¹⁹The impulse responses for NII are reported in absolute changes as a fraction of GDP, not in percentage changes relative to the steady state. This is because the steady state value for NII is zero.

Figure 3: Conditional Volatility and Persistence of Fixed Capital Investment



Notes: This figure shows the impact response to an aggregate technology shock and the the first-order autocorrelation coefficient of fixed capital investment for models with $\bar{\zeta} \in [0, 0.4]$. The x-axis for both panels shows the upper bound of the capital adjustment cost distribution, $\bar{\zeta}$. In Panel(a), the y-axis shows the first element of the IRF of fixed capital investment to a one-standard deviation aggregate productivity shock in the intermediate good sector in percentage points. In Panel(b), the y-axis shows the first-order auto-correlation of fixed capital investment. For Panel(b) we detrend fixed capital investment with the HP(1600) filter and then divide the deviations by the trend.


4.3 Conditional Volatility and Persistence as a Function of Capital Adjustment Costs

In this section we illustrate the substitution mechanism between the two investment goods from a slightly different angle. We now simulate our model under our calibrated inventory level and the “No Inventory” setup over a wide range of $\bar{\zeta} \in [0, 0.4]$. The lower bound is frictionless adjustment, whereas the upper bound, 0.4, is approximately twice the baseline $\bar{\zeta} = 0.1950$.²⁰ We study how the conditional volatility, i.e., the impact response in the impulse-response function, and the persistence of fixed capital investment, consumption and net inventory investment change over this range of fixed capital adjustment costs.

Panel (a) of Figure 3 presents the conditional volatility of fixed capital investment over said $\bar{\zeta}$ -range for both the inventory model and the “No Inventory” model. Independently of the level of inventories, higher capital adjustment costs dampen the impact response of fixed capital investment to aggregate shocks, and they do this in a more pronounced way in the model with inventories. The interaction between inventories and nonconvex capital adjustment costs is also apparent in the behavior of the persistence of fixed capital investment in Panel (b) of Figure 3. With inventories, persistence increases from 0.61 to 0.75 when $\bar{\zeta}$ changes from 0 to 0.4. In contrast, without inventories persistence only increases from 0.64 to 0.68 over the same range of $\bar{\zeta}$. The agents rely less on fixed capital investment when inventories are available. As a result, the fluctuations in fixed capital investments are dampened and drawn out. It is important to emphasize again that the central message of the paper is depicted in the different slopes of the two lines in both panels of Figure 3,

²⁰At $\bar{\zeta} = 0.4$ the annual fraction of firms which have lumpy investments is 15.42%, and the annual average adjustment cost paid conditional on adjustment and measured as a fraction of the firm’s output is 1.56%.

Figure 4: Conditional Volatility and Persistence of Net Inventory Investment



Notes: See notes to Figure 3. This figure shows the impact response to an aggregate technology shock and the the first-order autocorrelation coefficient of net inventory investment (NII) divided by GDP for models with $\bar{\zeta} \in [0, 0.4]$.

which is precisely a graphical representation of the nontrivial cross effect between general equilibrium modeling and the impact of adjustment costs for fixed capital on aggregate statistics - conditional volatility and persistence.

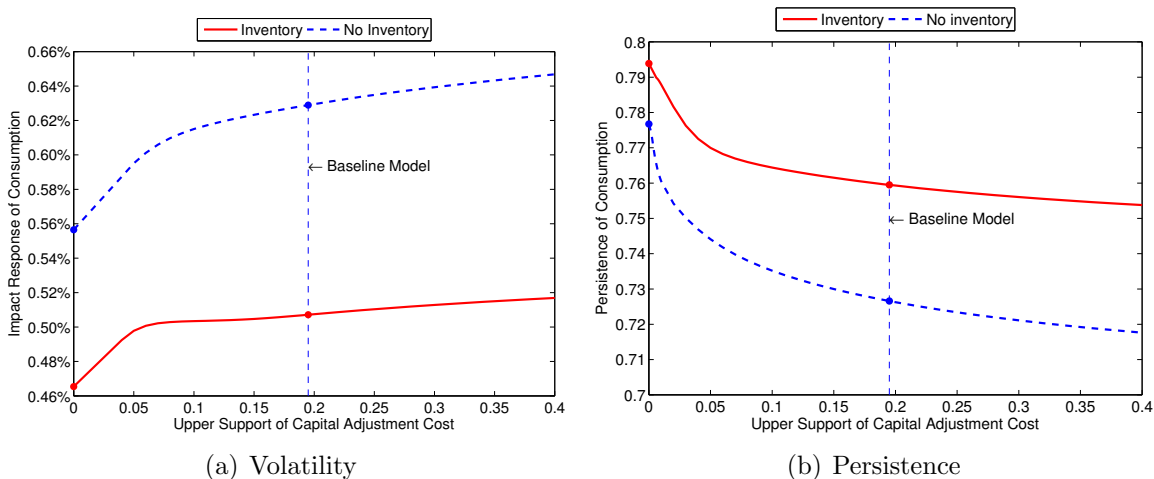
We can directly observe the substitution between different investment channels by contrasting the conditional volatility of fixed capital investment in Figure 3 to the conditional volatility of net inventory investment in Panel (a) of Figure 4. As fixed adjustment costs increase, the agents rely more on inventories and less on fixed capital for consumption smoothing. As a result, higher fixed adjustment costs lead to more volatile net inventory investment and less volatile fixed capital investment. Panel (b) of Figure 4 shows the opposite, albeit with a small nonmonotonicity, effect on persistence of net inventory investment.²¹

Also, we can see the implications of the investment substitution mechanism in the dynamics of consumption. Figure 5 shows that with inventories the conditional volatility of consumption is lower for every level of capital adjustment costs. More importantly, as the slopes of the two curves suggest, the rate at which fixed adjustment costs increases conditional consumption volatility is lower when inventories exist. In other words, the same increase in fixed adjustment cost makes conditional consumption volatility move up higher when inventories are absent from the economy, whereas it can barely increase conditional consumption volatility when inventories are present.

The change in consumption persistence reveals the same mechanism, as shown in Panel (b) of Figure 5. The existence of inventories changes the degree to which fixed capital adjustment costs affect consumption persistence. Over the same range of $\bar{\zeta}$, consumption

²¹The small nonmonotonicity is reflective of an overall persistence effect caused by adjustment costs on fixed capital investment and thus a more persistent output in the intermediate goods sector. The supply of intermediate goods becomes more persistent as fixed capital adjustment costs increase, while it is really the demand for intermediate goods as a means of smoothing consumption that causes the initial strong decrease of persistence in net inventory investment. Eventually, the effect from the supply side of intermediate goods slightly dominates.

Figure 5: Conditional Volatility and Persistence of Consumption



Notes: See notes to Figure 3. This figure shows the impact response to an aggregate technology shock and the the first-order autocorrelation coefficient of consumption for models with $\bar{\zeta} \in [0, 0.4]$. For Panel(b) consumption is logged and detrended with an HP filter with a smoothing parameter of 1600.

persistence decreases by much less in the inventory models compared to the “No Inventory” models. With inventories, consumption dynamics are more insulated from the effects of frictions in the adjustment of fixed capital.


4.4 The Effect of Market Clearing

The results on the effectiveness of fixed capital adjustment costs with or without inventories so far take into account all general equilibrium effects, i.e., adjustments of real interest rates and real wages, as well as the relative price of intermediate goods. In this section we isolate the effects of these price movements on how inventories impact the (ir)relevance of nonconvex fixed capital adjustment costs.

To this end, we solve three partial equilibrium versions of our model. In the first case, we fix the pricing kernel p (and thus the real wage), and the relative price q , at their long-run general equilibrium averages and simulate the model. In the second case, we fix the pricing kernel to its long-run general equilibrium average, but allow the relative price to adjust so that the intermediate goods market clears. In the last case we fix the relative price to its long-run general equilibrium average, but allow the pricing kernel (and the real wage) to adjust so that the final goods market clears.

The impulse response functions of fixed capital investment for all three cases are reported next to the full general equilibrium case – Panel (a) – in Figure 6. Panel(b) is the response from the first partial equilibrium case where both prices are fixed. Two messages emerge from this case. First, as is well known in the literature, nonconvex adjustment frictions matter a lot in partial equilibrium: the impact response drops substantively, and propagation arises only when fixed adjustment frictions are introduced. Second, inventories by and large do not change the effect of fixed adjustment frictions, as the differences between Model I1 and I2 are very similar to the differences between Model NI1 and NI2. Put differently, the effect

Figure 6: IRF for Fixed Capital Investments in Partial Equilibrium Models



Notes: These are the impulse response functions for fixed capital investments. Panel(a) is a reproduction of Figure 1. Panel(b) is based on models where both the pricing kernel and the relative price are fixed. Panel(c) is based on models where only the pricing kernel is fixed. Panel(d) is based on models where only the relative price is fixed.

of fixed capital adjustment frictions swamps the differential effect of inventories.

Panel(c) presents the response functions from the models where the pricing kernel is fixed but the relative price is not. The results in these models are very similar to those in the first case where both prices are fixed. Once again, nonconvex adjustment frictions matter a lot, but inventories do not interact with them significantly. Market clearing in the intermediate goods market only leads to slightly dampened fixed investment responses overall, as decreases in the relative price q (see Panel(d) of Figure 2) lead consumption smoothing activities away from fixed capital investment.

In other words, our exercise of comparing differences in differences really becomes only interesting, once real interest rate and real wage movements have been taken into account. The response functions in Panel(d) of Figure 6 come from the models where the pricing kernel and the real wage move freely to clear the final goods market, yet the relative price

of intermediate goods is fixed. These response functions resemble those from the general equilibrium case in that in models with inventories the impact response of fixed investment is 41% higher with frictionless fixed capital adjustment, whereas in models without inventories it is only 33% higher.²² Nevertheless, market clearing in the intermediate goods market does play a role in rendering fixed capital adjustment frictions more relevant. Recall that in full general equilibrium the difference in the initial fixed investment response between the frictionless model and the lumpy model was 46% (in the case with inventories) vs. 26% (in the case without inventories), that is the difference widens, when the intermediate goods market clear through a decline in q . The decline of the relative price q after an increase in aggregate productivity further facilitates the shifting of consumption smoothing through building up inventories and away from fixed capital investment. This substitution channel, for a given decline in q , is more valuable in an economy, when fixed capital adjustment is costly.²³

Taken together, the impulse response functions in Figure 6 illustrate two possible mechanisms to break the irrelevance result for nonconvex fixed capital adjustment costs (in addition to making the extensive and intensive margin of fixed capital investment difficult to substitute, see Gourio and Kashyap (2007)): decouple, with or without relative price movements for the second saving vehicle, the tight link between aggregate saving and fixed capital investment induced by the standard aggregate resource constraint with one saving vehicle, as in our paper (Panels (a) and (d) of Figure 6) and Berger and Vavra (2015); or dampen the real interest rate movements induced by this tight link (see Cooper and Willis (2012) and Winberry (2016)), an extreme version of which can be seen in Panels (b) and (c) of Figure 6. Both mechanisms, though distinct, are associated with smoother aggregate consumption dynamics.

5 Robustness

In the baseline model and its calibration we made several choices that we subject to four robustness tests in this section: (1) we set the persistence parameter of the aggregate productivity shock at $\rho_z = 0.95$; (2) we chose, following Bloom (2009), $\theta_k = 0.25$ and $\theta_l = 0.5$; (3) we let the aggregate productivity shock affect only the production function in the intermediate goods sector; and (4) we did not allow for persistent idiosyncratic productivity

²²The relative impact conditional on the same level of adjustment costs for fixed capital has changed between Panels (a) and (d) of Figure 6. For example, with no fixed capital adjustment costs, fixed capital reacts more to a productivity shock when there are inventories, but the price of intermediate goods is fixed, compared to the case where the price of intermediate goods adjusts downward, where the relative size of the reaction of fixed capital investment is reversed between the inventory and the ‘no inventory’-case. Of course, with frictionless fixed capital adjustment, positive inventory holding costs and a fixed price at which inventories can be stocked up, there is really not much reason to smooth consumption via inventories and thus fixed capital investment reacts more strongly. This changes, when the price of intermediate goods declines, which makes inventories as a smoothing device more attractive.

²³Notice the contrast to the mechanism in Fiori (2012), where the aggregate relevance of lumpy investment is generated entirely by movements of a relative price in a two-sector environment, namely, the relative price of investment. Here the relevant relative price, that is, of intermediate goods, strengthens, but does not drive our result.

shocks in either sector. In this section we conduct four robustness checks, all relative to the baseline scenario:

1. We set $\rho_z = 0.98$. All other parameters are calibrated using the same strategy as before, i.e., σ_z to match the volatility of U.S. GDP, A^h to match an aggregate labor input of 0.33, θ_m to match a material share in final output of 0.499, θ_n to match an aggregate labor share of 0.64, σ to match the annual storage costs of inventories as a fraction of final output of 12%, and $\bar{\zeta}$ and $\bar{\epsilon}$ to match, respectively, the fraction of investors with lumpy fixed capital adjustments of 18% and an average inventory-to-sales ratio of 0.8185.²⁴
2. We set $\theta_k = 0.256$ and $\theta_l = 0.64$, following Khan and Thomas (2008), and thus allowing for a much less curved production function. Put differently, the implied capital elasticity of the firms' revenue function, after labor is optimally chosen, is now $\frac{\theta_k}{1-\theta_l} = 0.71$, as opposed to the 0.5 of the baseline case. All other parameters are calibrated using the same strategy as before.
3. We assume the production function in the final goods sector to be $zG(m, n)$ and in the intermediate goods sector to be $F(k, l)$, i.e., there is now no aggregate productivity shock in the intermediate goods sector. ρ_z is set to 0.95 and σ_z is calibrated to match the volatility of U.S. GDP. All other parameters are calibrated using the same strategy as before.²⁵
4. We allow for persistent idiosyncratic productivity shocks in the intermediate goods sector, i.e., the production function in the intermediate goods sector is now $z\varepsilon F(k, l)$. We assume that $\log(\varepsilon)$ follows an AR(1) process and discretize it, using the method by Tauchen (1986), on a 15-state Markov chain. The underlying parameters are $\rho_\varepsilon = 0.95$, i.e., we set $\rho_\varepsilon = \rho_z$, and $\sigma_\varepsilon = 0.022$, both strategies following Khan and Thomas (2008). All other parameters are calibrated using the same strategy as before.

The results are concisely summarized in Table 4, which shows the reduction in the initial investment impulse response between the case with calibrated fixed capital adjustment costs and the frictionless capital adjustment case as a percentage of the case with calibrated fixed capital adjustment costs, where we compare the case with inventories (left column) with the case without inventories (right column); and in Figures 7 and 8, which are the equivalent of Figures 3 and 5 in Section 4.3, i.e., they show how in the various models the conditional volatility and persistence of aggregate investment and consumption change as a function of fixed capital adjustment costs, comparing the case with inventories to the case without them. The result is clear: in all cases inventories lead to a stronger dampening of the initial response of investment by fixed capital adjustment costs, and its persistence rises by much more, by and large consistent with the baseline case. The weakest relative dampening effect is generated in the case of the less curved production function, which is consistent with

²⁴Table 7 in Appendix B summarizes the parameters of all four robustness check models and juxtaposes them with the baseline parameters.

²⁵We should repeat that such an inventory model yields counterfactual comovements of the relative price of intermediate goods, which is procyclical here.


the intuition given in the introduction: the closer-to-linear the production function is, the easier the substitution between the extensive and intensive margin of investment, so that inventories are relatively less useful for smoothing consumption. As for the no-change result when we increase ρ_z to 0.98, while it is true that the value of ρ_z will change the absolute responsiveness of the model to shocks, the robustness check shows that the responsiveness of the model to shocks is altered in essentially the same way across the four models (I1, I2, NI1, NI2), so that the relative responsiveness of these models to shocks remains unaltered, at least in the range for the persistence parameter we study.

Table 4: Dampening of Initial Fixed Capital Investment Responses - Inventories versus No Inventories


	Inventories	No Inventories
Baseline	46%	26%
$\rho_z = 0.98$	50%	30%
$\theta_k = 0.256, \theta_k = 0.64$	28%	17%
Aggr. Prod. Shock in Final Goods Sector	61%	25%
Idiosyncratic Prod. Shocks	45%	25%

Notes: This table shows the reduction in the initial investment impulse response between the case with calibrated fixed capital adjustment costs and the frictionless capital adjustment case as a percentage of the case with calibrated fixed capital adjustment costs. We compare the case with inventories (left column) with the case without inventories (right column). We display these statistics for the baseline case from Section 4.2 as well as the four robustness checks described in this section.


Figure 7: Conditional Volatility and Persistence of Fixed Capital Investment




(a) Volatility - $\rho_z = 0.98$




(b) Persistence - $\rho_z = 0.98$




(c) Volatility - $\theta_k = 0.256, \theta_k = 0.64$




(d) Persistence - $\theta_k = 0.256, \theta_k = 0.64$




(e) Volatility - Aggr. Prod. Shock in Final Goods Sector



(f) Persistence - Aggr. Prod. Shock in Final Goods Sector




(g) Volatility - Idiosyncratic Prod. Shocks




(h) Persistence - Idiosyncratic Prod. Shocks

Notes: This figure shows the impact response to an aggregate technology shock and the the first-order autocorrelation coefficient of fixed capital investment for models with $\bar{\zeta} \in [0, 0.4]$. The x-axis for both panels shows the upper bound of the capital adjustment cost distribution, $\bar{\zeta}$. The left-hand side panels show on the y-axis the first element of the IRF of fixed capital investment to a one-standard deviation aggregate productivity shock in the intermediate good sector (except for the third row, where the productivity shock is to the final goods sector) in percentage points. In the right-hand side panels, the y-axis shows the first-order auto-correlation of fixed capital investment. For the right-hand side panels, we detrend fixed capital investment with the HP(1600) filter and then divide the deviations by the trend.


Figure 8: Conditional Volatility and Persistence of Consumption




(a) Volatility - $\rho_z = 0.98$




(b) Persistence - $\rho_z = 0.98$




(c) Volatility - $\theta_k = 0.256, \theta_k = 0.64$




(d) Persistence - $\theta_k = 0.256, \theta_k = 0.64$




(e) Volatility - Aggr. Prod. Shock in Final Goods Sector



(f) Persistence - Aggr. Prod. Shock in Final Goods Sector



(g) Volatility - Idiosyncratic Prod. Shocks



(h) Persistence - Idiosyncratic Prod. Shocks

Notes: See notes to Figure 7. This figure shows the impact response to an aggregate technology shock and the first-order autocorrelation coefficient of consumption for models with $\zeta \in [0, 0.4]$. For the right-hand side panels, consumption is logged and detrended with an HP filter with a smoothing parameter of 1600.

6 Conclusion

This paper shows that it matters for the aggregate implications of microfrictions *how* general equilibrium effects are introduced into the physical environment of dynamic stochastic general equilibrium models with these microfrictions. Specifically, we show that how relevant nonconvex fixed capital adjustment costs are for business cycle dynamics depends on how the aggregate resource constraint is modeled, depends on how the model is closed. Future research will explore the general insight in more general frameworks.

We develop a dynamic stochastic general equilibrium model to evaluate how the availability of multiple investment channels, here inventories in addition to fixed capital, affects the aggregate implications of nonconvex capital adjustment costs. We find that with more than one ways to invest, capital adjustment costs are more effective in dampening and propagating the response of fixed capital investment to an aggregate productivity shock.

References

- BACHMANN, R., E. ENGEL, AND R. CABALLERO (2013): “Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model,” *The American Economic Journal: Macroeconomics*, 5, 29–67.
- BERGER, D. AND J. VAVRA (2015): “Consumption Dynamics During Recessions,” *Econometrica*, 83, 101–154.
- BLOOM, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.
- CABALLERO, R. (2010): “Macroeconomics after the Crisis: Time to Deal with the Pretense-of-Knowledge Syndrome,” *Journal of Economic Perspectives*, 24, 85–102.
- CABALLERO, R. AND E. ENGEL (1999): “Explaining Investment Dynamics in US Manufacturing: a Generalized (S,s) Approach,” *Econometrica*, 67, 783–826.
- CAPLIN, A. S. AND D. F. SPULBER (1987): “Menu Costs and the Neutrality of Money,” *The Quarterly Journal of Economics*, 102, pp. 703–726.
- COOLEY, T. F. AND E. C. PRESCOTT (1995): “Economic Growth and Business Cycles,” in *Frontiers of business cycle research*, ed. by T. F. Cooley, Princeton University Press Princeton, N.J, 1 – 38.
- COOPER, R. AND J. HALTIWANGER (2006): “On the Nature of Capital Adjustment Costs,” *Review of Economic Studies*, 73, 611–633.
- COOPER, R. AND J. WILLIS (2012): “Do Interest Rates Smooth Investment?” *mimeo*, found: <http://www.economics.ku.edu/seminars/Friday/documents/oct4.pdf>.
- DE BOOR, C. (1978): *A practical guide to splines*, Springer.
- DEN HAAN, W. (2010): “Assessing the Accuracy of the Aggregate Law of Motion in Models with Heterogeneous Agents,” *Journal of Economic Dynamics and Control*, 34, 79–99.

- DOMS, M. AND T. DUNNE (1998): “Capital Adjustment Patterns in Manufacturing Plants,” *Review of Economic Dynamics*, 1, 409–429.
- FIORI, G. (2012): “Lumpiness, Capital Adjustment Costs and Investment Dynamics,” *Journal of Monetary Economics*, 59, 381–392.
- GOURIO, F. AND A. KASHYAP (2007): “Investment Spikes: New Facts and a General Equilibrium Exploration,” *Journal of Monetary Economics*, 54, 1–22.
- HENNESSY, C. A. AND T. M. WHITED (2005): “Debt Dynamics,” *The Journal of Finance*, 60, 1129–1165.
- HOUSE, C. L. (2014): “Fixed costs and long-lived investments,” *Journal of Monetary Economics*, 68, 86 – 100.
- KHAN, A. AND J. THOMAS (2003): “Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?” *Journal of Monetary Economics*, 50, 331–360.
- (2007): “Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies,” *The American Economic Review*, 97, 1165–1188.
- (2008): “Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics,” *Econometrica*, 76, 395–436.
- KING, R. G. AND S. T. REBELO (1999): “Chapter 14: Resuscitating real business cycles,” Elsevier, vol. 1, Part B of *Handbook of Macroeconomics*, 927 – 1007.
- KRUSELL, P. AND A. SMITH (1997): “Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns,” *Macroeconomic Dynamics*, 1, 387–422.
- (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *The Journal of Political Economy*, 106, 867–896.
- MIAO, J. AND P. WANG (2014a): “A Q-Theory Model with Lumpy Investment,” *Economic Theory*, 57, 133–159.
- (2014b): “Lumpy Investment and Corporate Tax Policy,” *Journal of Money, Credit and Banking*, 46, 1171–1203.
- RICHARDSON, H. (1995): “Control Your Costs Then Cut Them,” *Transportation & Distribution*, 36, 94–94.
- TAUCHEN, G. (1986): “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 20, 177–181.
- THOMAS, J. (2002): “Is Lumpy Investment Relevant for the Business Cycle?” *The Journal of Political Economy*, 110, 508–534.

VERACIERTO, M. L. (2002): “Plant-Level Irreversible Investment and Equilibrium Business Cycles,” *The American Economic Review*, 92, pp. 181–197.

WINBERRY, T. (2016): “Lumpy Investment, Business Cycles, and Stimulus Policy,” *mimeo*, *University of Chicago*, *found*: <http://faculty.chicagobooth.edu/thomas.winberry/research/winberryJMP.pdf>.

Appendix

A Numerics

A.1 Numerical Implementation

We start the algorithm to compute the recursive equilibrium (see Section 2.2.5) of the model outlined in Section 2 with a guess of the Krusell and Smith forecast rules on prices and future cross-sectional averages of the inventory and capital stocks, see equations (8) to (11) in the main text. With the conjectured forecast rules, we use value function iteration to solve the problems of the intermediate good producers and the final good producers. To approximate the value function, we use cubic spline interpolation with “not-a-knot” condition over fixed grids. The spline interpolation and evaluation algorithms are based on De Boor (1978) and implemented in the open-source library PPPACK. The grid points are equispaced over fixed intervals. The end points of the intervals and the number of grid points used for the benchmark model are reported in Table 5. In simulations of the equilibrium (see below) we check that the end points of the interval are never reached. We iterate the value functions until the relative difference between two consecutive iterations is smaller than 1.0E-7 under the l_∞ norm. We employ a 20-step Howard policy function acceleration module to speed up convergence. We discretize the aggregate productivity process using the method outlined in Tauchen (1986) with 11 grid points. We require the boundaries of the discretization to be 2 times the standard deviation of the underlying AR(1) process.

Table 5: Grid Points in Value Function Approximation

State Variable	Left End	Right End	Number of Grid Points
s	0.0	3.0	20
k	0.0	4.0	20
$\mu_1(S)$	0.25	0.75	10
$\lambda_1(K)$	0.5	1.5	10

Notes: s is the inventory holding at the firm level. k is the capital stock at the firm level. $\mu_1(S)$ is the cross-sectional mean of the inventory distribution. $\lambda_1(K)$ is the cross-sectional mean of the capital stock distribution.

With the approximated value functions, we simulate the model economy. In these simulations we track the distribution of firms with probability mass functions. For example, the intermediate goods firm size distribution over capital is tracked by two vectors: $\{m(k_i), k_i\}$, where $m(k_i)$ is the mass of firms with capital stock k_i . These mass points evolve endogenously in the simulation. The sum of weights over all grid points equals to one: $\sum_i^{n_k} m(k_i) = 1$, and n_k is the number of unique values of capital stock with non-zero mass. The firm size distribution over inventory stock is tracked using the same method. At the end of each period, we drop the mass points with a weight smaller than 1.0E-10 and reassign their weight uniformly across the other mass points to reduce the computational load. We need to track approximately 35 mass points for the final goods firm distribution, and approximately 85 mass points for the intermediate goods firm distribution.

We conduct the simulations without imposing the conjectured forecast rules used in the value function iteration part of the algorithm. Instead, we search for a pair of (p, q) to clear the final and intermediate good markets in each period of the simulation, and let the cross-sectional averages of the inventory and capital stocks evolve given these equilibrium prices. The excess demand functions for both markets are nonlinear, and thus we employ the Gauss-Jacobi algorithm to solve for the pair of (p, q) . Both markets are cleared with a relative error margin of 1.0E-8. For better accuracy, we re-solve the optimization problems of the intermediate good producers and the final good producers for every guess of (p, q) in the Gauss-Jacobi solver in every period.

We simulate the model over 1,000 periods and discard the first 100 periods when updating the forecast rules with OLS. After updating the forecast rules, we repeat the algorithm: solve the value functions again with the updated forecast rules and simulate the model without imposing the forecast rules. We loop over the entire algorithm until the forecast rules converge with a percentage difference smaller than 1.0E-4. The converged forecast rules for the benchmark model are reported in Table 6. Since the cross-sectional averages of the capital stock and the inventory holdings distributions turn out to be highly correlated in the model’s equilibrium, and thus including them both would lead to a multicollinearity problem, we use the average capital holdings in its own forecast equation and the two pricing equations ($\gamma_{\{.\}} = 0$), and the average inventory holdings only in its own forecast equation ($\beta_{\mu} = 0$).

Table 6: Forecast Rules for Benchmark Model

	α	β	γ	ψ	R^2
$\log \mu'_1$	-0.0892	-	0.8297	0.2262	0.9984
$\log \lambda'_1$	0.0025	0.9039	-	0.1048	0.9995
$\log p$	1.2521	-0.3355	-	-0.3027	0.9999
$\log q$	-0.8436	-0.1888	-	-0.3861	0.9999

Notes: This table reports the forecast rules for the benchmark model corresponding to equations (8) to (11). μ_1 is the cross-sectional mean of the inventory stock distribution, and λ_1 is the cross-sectional mean of the capital stock distribution.

A.2 Further Quality Checks on the Numerical Implementation

In addition to looking at the R^2 of the OLS estimation, we use the procedure outlined in Den Haan (2010) to check the accuracy of our numerical implementation. We start the exercise by drawing a new sequence of aggregate shocks z_t different from the sequence used in the solution algorithm. Using this new z_t -series, we first simulate the model with the method described in Appendix A.1 for 1,000 periods, using the forecast rules reported in Table 6. Denote the *simulated* time series of the mean of the capital stock distribution as \hat{k}_t , and the sequence of the mean of the inventory level distribution as \hat{s}_t . We then generate another set of *predicted* time paths of the same variables, denoted as \tilde{k}_t and \tilde{s}_t , by using only

the forecast rules reported in Table 6:


$$\begin{aligned}\log(\tilde{k}_1) &= \log(\hat{k}_1), \\ \log(\tilde{s}_1) &= \log(\hat{s}_1), \\ \log(\tilde{k}_{t+1}) &= \alpha_k + \beta_k \log(\tilde{k}_t) + \psi_k \log(z_t), \\ \log(\tilde{s}_{t+1}) &= \alpha_s + \gamma_s \log(\tilde{s}_t) + \psi_s \log(z_t).\end{aligned}$$

The accuracy of the forecast rules can be measured with the maximum and the mean of the relative deviation between the simulated and the predicted time series. For example, the Den Haan (2010) accuracy measures for the forecast rules of the capital stock distribution are:

$$\begin{aligned}d_k^{max} &= \max_{t>100} (|\log(\hat{k}_t) - \log(\tilde{k}_t)|), \\ d_k^{ave} &= \frac{\sum_{t=100}^{1000} (|\log(\hat{k}_t) - \log(\tilde{k}_t)|)}{900},\end{aligned}$$

and similar measures can be computed for the forecast rules of the inventory holdings distribution. We discard the first 100 periods of both sequences when computing the accuracy measures to avoid the influence of the initial conditions.

Figure 9: Simulated and Predicted k_t and s_t



Notes: This figure shows the simulated and predicted time paths for the mean of the capital stock distribution and the inventory distribution. The simulated time paths are generated with the algorithms described in Appendix A.1. The predicted time paths are generated using only the forecast rules reported in Table 6.

We plot the time paths of the simulated and the predicted sequences in Figure 9. The graphs suggest that differences between the simulated and the predicted sequences are small and not persistent, and do not accumulate over the course of the simulation. For k_t , $d_k^{ave} = 0.0027$ and $d_k^{max} = 0.0103$. For s_t , $d_s^{ave} = 0.0041$ and $d_s^{max} = 0.0139$. These numbers are well within the range of numbers that Den Haan (2010) reports in his accuracy exercises.

A.3 Computation of the Impulse Response Functions

To compute the impulse response functions in the main text, we first keep $z = 1$ for 100 periods thus letting the economy reach a steady state. All the aggregate variables converge to constants well before the 100th period, and we use the firm-individual capital stock and inventory holding distributions obtained at the end of this simulation as the initial condition to generate impulse response functions.

Starting from this initial condition, we next simulate the model for $m = 60$ periods for $m = 60$ times. For the i th simulation, we force the first i aggregate shocks to be one grid point above $z = 1$, and the rest of the shocks to be $z = 1$. Since we discretize z into 11 states, the median state corresponding to $z = 1$ is the 6th z -grid point, and the higher-productivity state is the 7th z -grid point.

We first compute the average time path of a certain variable χ as the weighted average of the time paths from all m simulations. Denote the time path of χ in the i th simulation by χ_i , then the average time path, $\bar{\chi}$, is given by:

$$\bar{\chi} = \sum_{i=1}^m \chi_i \cdot \phi_i,$$

where ϕ_i is the weight on the i th simulation:

$$\phi_i = \frac{(1 - \pi_s)\pi_s^{(i-1)}}{\sum_{j=1}^m (1 - \pi_s)\pi_s^{(j-1)}},$$

and π_s is the probability of staying in the higher-productivity state (the 7th z -grid point):

$$\pi_s = \sum_{j=\kappa}^{N_z} \pi_{\kappa j},$$

$$\kappa = \frac{N_z + 1}{2} + 1.$$

Recall that in the Markovian representation of the aggregate shocks, $\pi_{\kappa j}$ is the probability of transiting to state j from state κ , N_z is the total number of potential states, and κ is the state of higher productivity: one grid point above $z = 1$.

The IRF of χ is then computed as:

$$\text{IRF}_\chi = \left(\frac{\bar{\chi}}{\chi^s} - 1 \right) \cdot \left(\frac{\sigma_z}{z_\kappa - z_{\kappa-1}} \right),$$

where χ^s is the steady state value of χ . The last scaling factor serves to convert the size of the shock from one grid point to one standard deviation.

When implementing this procedure we set m high enough ($m = 60$) so that the weight on the last path of z is numerically negligible: $\phi_{60} \approx 1.3\text{E-}9$. In the graphs of the main text we report the first eight elements of the IRF.

B Parameters for Robustness Checks

Table 7: Parameters for Robustness Checks

Model	A^h	θ_m	θ_n	θ_k	θ_l	ρ_z	σ_z	σ	$\bar{\zeta}$	$\bar{\varepsilon}$
Baseline	2.0720	0.5245	0.3530	0.2500	0.5000	0.9500	0.0167	0.0128	0.1950	0.4580
$\rho_z = 0.98$	2.0700	0.5240	0.3530	0.2500	0.5000	0.9800	0.0173	0.0125	0.1950	0.4480
$\theta_k = 0.256, \theta_k = 0.64$	2.0800	0.5240	0.2820	0.2560	0.6400	0.9500	0.0150	0.0144	0.2697	0.1130
Aggr. Prod. Shock in Final Goods Sector	2.0700	0.5245	0.3530	0.2500	0.5000	0.9500	0.0088	0.0129	0.1900	0.4672
Idiosyncratic Prod. Shocks	2.0720	0.5240	0.3530	0.2500	0.5000	0.9500	0.0167	0.0128	0.1360	0.4650

Notes: A^h is the preference parameter for leisure; θ_m is the elasticity of materials in the final goods production function; θ_n is the elasticity of labor in the final goods production function; θ_k is the capital elasticity in the intermediate goods production function; θ_l is the labor elasticity in the intermediate goods production function; ρ_z is the auto-correlation for the aggregate productivity process; σ_z is the standard deviation for aggregate productivity innovations; $\bar{\zeta}$ is the upper bound of the fixed capital adjustment cost distribution; $\bar{\varepsilon}$ is the upper bound of the inventory adjustment cost distribution. For better readability, we have left out the parameters that stay the same across all models: β , δ and γ .