Idiosyncratic Shocks, Lumpy Investment and the Monetary Transmission Mechanism*

Michael Reiter[a,b], Tommy Sveen[c], Lutz Weinke[d]

[a]Institute for Advanced Studies, Vienna, [b]NYU Abu Dhabi

[c]BI Norwegian Business School, [d]Humboldt-Universität zu Berlin

May 27, 2020

Abstract

Standard (S,s) models of lumpy investment allow us to match many aspects of the micro data, but it is well known that the implied interest rate sensitivity of investment is unrealistically large. The monetary transmission mechanism is therefore a natural starting point to assess the macroeconomic relevance of any investment theory. Our results show that lumpy investment can coexist with a realistic monetary transmission mechanism, but that we are nevertheless still a step away from a micro-founded theory of monetary policy.

Keywords: Lumpy Investment, Sticky Prices.

JEL Classification: E22, E31, E32

*Thanks to Jordi Galí, Adrian Ochs, Christian Wolf, and to seminar participants at several venues. Patrick Brock, Clara Hoffmann, Dennis Zander and Lara Zarges provided excellent research assistance. This research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 402884221, and the Fonds zur Förderung der wissenschaftlichen Forschung (FWF Austrian Science Fund) – I3840-G27. Their financial support is gratefully acknowledged. The usual disclaimer applies.

1 Introduction

What explains the short-run effects of monetary policy on real variables of interest? This question takes center-stage in much of the literature in macroeconomics. In the words of Galí (2015): "Over the past two decades, monetary economics has been among the most fruitful research areas within macroeconomics. The efforts of many researchers to understand the relationship among monetary policy, inflation, and the business cycle have led to the development of a framework - the so called New Keynesian model - that is widely used for monetary policy analysis." Our motivation to reconsider this question originates in a well-known micro-macro puzzle in investment theory. In fact, (S,s)[1] models of lumpy investment allow us to match many aspects of the micro data, but the implied interest rate sensitivity of investment is unrealistically large (see, e.g., Thomas 2002 and Khan and Thomas 2008). It therefore seems to us that the monetary transmission mechanism is a natural starting point to assess the macroeconomic relevance of any investment theory. For instance, Reiter et al. (2013) have shown that once an otherwise conventional NK model is augmented with a lumpy investment decision à la Thomas (2002), the implied monetary transmission mechanism becomes counterfactual. Specifically, the impact responses of investment and output to a change in the nominal interest rate become very large and the dynamic consequences of that shock are only short-lived.[2]

A drawback of our work in Reiter et al. (2013) is, however, that the micro data on investment could not be fitted in a satisfactory way by just relying on a fixed adjustment cost for capital. The present paper therefore develops a HANK[3]

[1] The nature of optimal microeconomic decisions implied by fixed adjustment costs is typically referred to as (S,s), this way highlighting the range of inaction, which is a general feature of those decisions (see, e.g., Dotsey et al. 1999).

[2] McKay and Wieland (2019) show that this mechanism is also an important channel of monetary policy transmission in a fixed-cost model of durable consumption demand.

[3] HANK stands for Heterogeneous Agent New Keynesian. This term has been popularized by Kaplan et al. (2018).

model consistent with the cross-sectional distribution of establishment investment rates.[4] To this end, we combine the investment model by Khan and Thomas (2008) with a convex capital adjustment cost and integrate the resulting framework into an otherwise standard NK model. More concretely, each investor is assumed to face both a fixed cost and a convex cost of adjusting the capital stock, but low-level investments are exempt from the fixed cost. There is also an idiosyncratic shock to plant-level productivity. This set of assumptions allows us to match the micro data on investment that have been established in the seminal work by Cooper and Haltiwanger (2006). In this context, we also point at a problematic aspect of the calibration in Khan and Thomas (2008). They target most of the micro-facts on lumpy investment reported by Cooper and Haltiwanger (2006), but Khan and Thomas (2008) ignore the serial correlation in investment rates. Making the model consistent with the small and positive correlation reported by Cooper and Haltiwanger (2006) requires, however, a substantial extension of the model proposed by Khan and Thomas (2008), namely a combination of fixed and convex costs of adjusting the capital stock. Our model solution relies on the methods developed in Reiter (2009, 2010 and 2019).

But what does this imply for the monetary transmission mechanism? Under our baseline calibration a quantitatively relevant monetary transmission mechanism can coexist with lumpiness in investment at the micro level. In a nutshell, the intuition is as follows. A convex capital adjustment cost incentivizes firms to smooth investment. In the context of our quarterly model of the monetary transmission mechanism the combined size of investment over a year can, however, still be substantial enough to be consistent with both the investment spikes that are observed in the yearly

[4] That lumpiness is reported by, e.g., Doms and Dunne (1998), and Cooper and Haltiwanger (2006). In the context of our theory there is no distinction between a plant and a firm and we therefore use those terms interchangeably.

data and the lack of persistence in annual investment rates at the firm level. When we recalibrate our model in the spirit of Khan and Thomas (2008), i.e., assuming much smaller capital adjustment costs (in our case, of course, both fixed and convex costs) then our model can also generate negative investments (as well as investment spikes) in the stationary distribution. However, the implied interest rate sensitivity of investment becomes so large as to imply a counterfactual monetary transmission mechanism. It is therefore fair to say that we are still a step away from a micro-founded theory of monetary policy.

Let us relate our results to the literature. NK models often abstract from capital accumulation,[5] and if capital accumulation is taken into account in the context of NK theory then it is common practice to postulate convex adjustment costs in the investment block of the framework.[6] But the existence of those adjustment costs makes NK models inconsistent with the observed lumpiness in plant-level investment. An early attempt to make a NK model consistent with the lumpy nature of investment at the micro level is the work in Sveen and Weinke (2007). In this paper infrequent pricing and investment decisions are made in a Calvo (1983) fashion, and this framework is shown to be observationally equivalent in the aggregate to a model of convex capital adjustment costs at the firm-level, as in Woodford (2005). One drawback of our 2007 paper is, however, that the step taken in the direction of having an empirically relevant investment decision is relatively small compared to the standard approach in modern investment theory. The theory of lumpy investment is an active field of research, and the work in Winberry (2020) is an interesting recent contribution to it. He studies, however, the dynamic consequences of technological shocks in an RBC framework, whereas our paper is concerned with the monetary

[5] See, e.g., Galí (2015), among many others.

[6] For instance, Christiano et al. (2005) assume a convex investment adjustment cost, whereas Woodford (2005) postulates a convex capital adjustment cost.

transmission mechanism. In the investment block of his model, Winberry (2020) extends the model in Khan and Thomas (2008) by combining it with convex capital adjustment costs. He also assumes habit formation in the preferences of the representative household.[7] In one of their robustness checks, Reiter et al. (2013) had also combined fixed and convex costs of adjusting the capital stock. In that paper, we pointed at a tension associated with having a substantial convex portion of the capital adjustment cost. On the one hand, this gives rise to a realistic monetary transmission mechanism. On the other hand, it makes the model inconsistent with investment spikes at the quarterly frequency. In the present paper, we show that a plausible model of the monetary transmission mechanism can coexist with the lumpiness in yearly investment data that is documented in Cooper and Haltiwanger (2006). The recently emerging HANK literature has mostly studied the aggregate consequences of heterogeneity at the household level. A notable exception is the paper by Ottonello and Winberry (2019). They analyze the monetary transmission mechanism in the presence of financial heterogeneity, but abstracting from the lumpy nature of investment at the micro level. In the part of their paper that is most related to our work, Koby and Wolf (2020) embed a rich heterogeneous-firm block with lumpy firm investment into an otherwise standard medium-scale New Keynesian model. Their heterogeneous-firm block is calibrated to be jointly consistent with firm-level investment lumpiness and their novel evidence on investment price elasticities. They study the response of aggregate investment to expansionary monetary policy shocks over the business cycle, as a function of the underlying cross-sectional distribution of capital, and they show that lumpiness of investment can dampen the effectiveness of monetary policy in classical TFP recessions.

[7] In Winberry (2020), the consumption habit is important in order to generate a plausible degree of volatility of the real interest rate. This is, of course, not an issue in the context of a New Keynesian model.

The remainder of the paper is organized as follows. Section 2 outlines the model. Section 3 presents the dynamic analysis, and section 4 concludes.

2 The Model

Our model integrates lumpy investment into an otherwise standard New Keynesian model of the monetary transmission mechanism. There are households, intermediate goods firms, retail firms and a central bank in charge of conducting monetary policy.

2.1 Households

Households are assumed to have access to a complete set of financial markets. The representative household has the following period utility function

$$U\left(C_t, L_t\right) = \ln C_t - \frac{\varphi}{1+1/\phi} L_t^{1+1/\phi}, \tag{1}$$

which is separable in its two arguments C_t and L_t. The former denotes a Dixit-Stiglitz consumption aggregate while the latter is meant to indicate hours worked. A household's time endowment is normalized to one per period, and throughout the analysis the subscript t denotes the time period. The steady state labor supply elasticity is given by ϕ, and parameter φ is used to make sure that the representative household spends one third of time working in the labor market. The consumption aggregate reads

$$C_t \equiv \left(\int_0^1 C_t\left(i\right)^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}}, \tag{2}$$

where ϵ is the elasticity of substitution between different varieties of goods $C_t(i)$. The associated price index is defined as follows

$$P_t \equiv \left(\int_0^1 P_t(i)^{1-\epsilon} di \right)^{\frac{1}{1-\epsilon}}, \tag{3}$$

where $P_t(i)$ is the price of good i. Requiring optimal allocation of any spending on the available goods implies that consumption expenditure can be written as $P_t C_t$. Households are assumed to maximize expected discounted utility

$$E_t \sum_{k=0}^{\infty} \beta^k U(C_{t+k}, L_{t+k}),$$

where β is the subjective discount factor. The maximization is subject to a sequence of budget constraints of the form

$$P_t C_t + E_t \{Q_{t,t+1} D_{t+1}\} \leq D_t + P_t w_t L_t + T_t, \tag{4}$$

where $Q_{t,t+1}$ denotes the stochastic discount factor for random nominal payments and D_{t+1} gives the nominal payoff associated with the portfolio held at the end of period t. We have also used the notation w_t for the real wage and T_t is nominal dividend income resulting from ownership of firms.

The labor supply equation implied by this structure takes the standard form

$$\varphi \; C_t L_t^{1/\phi} = w_t, \tag{5}$$

and the consumer Euler equation is given by

$$Q_{t,t+1}^R = \beta \; \left(\frac{C_{t+1}}{C_t} \right)^{-1}, \tag{6}$$

where $Q^R_{t,t+1} \equiv Q_{t,t+1}\Pi_{t+1}$ is the real stochastic discount factor, and $\Pi_{t+1} \equiv \frac{P_{t+1}}{P_t}$ is the gross rate of inflation between periods t and $t+1$. We also note that $E_t\{Q_{t,t+1}\} = R_t^{-1}$, where R_t is the gross risk free nominal interest rate.

2.2 Intermediate Good Firms

There is a continuum of intermediate good firms indexed on the unit interval. They produce with capital and labor, and they face idiosyncratic shocks to their productivity. Let us note already that the relative intermediate good price is also the real marginal cost for retail firms. A key difference with respect to the model proposed in Khan and Thomas (2008) is that intermediate good firms are assumed to face not only a fixed cost but also a convex cost of adjusting the capital stock. In each period the time-line is as follows:

1. The idiosyncratic productivity shock realizes.
2. The firm chooses its current level of labor input, production takes place, and workers are paid.
3. The fixed cost of adjusting the capital stock realizes.
4. The firm invests (or not).

Each period, an intermediate good firm therefore solves a problem of the form[8]

$$\max E_t \sum_{k=0}^{\infty} Q_{t,t+k}^{R} \; D_{t+k}$$

s.t.

$$D_t = q_t^M x_t z_0 e^{\gamma t} \widetilde{L}_t^{\nu} K_t^{\alpha} - w_t \widetilde{L}_t - \Psi\left(K_t, K_{t+1}\right),$$

with

$$\Psi\left(K_t, K_{t+1}\right) = \begin{cases} i_t + \epsilon_\psi K_t \left(\frac{i_t}{K_t}\right)^2 \text{ if } i_t \in [aK_t, bK_t] \\ i_t + \epsilon_\psi K_t \left(\frac{i_t}{K_t}\right)^2 + f_t w_t \text{ if } i_t \notin [aK_t, bK_t], \end{cases} \tag{7}$$

and

$$i_t = \gamma K_{t+1} - (1-\delta) K_t. \tag{8}$$

All variables measured in units of output are defined as a Dixit-Stiglitz aggregate of the same form as the consumption aggregate. An intermediate good firm's capital stock, K_t, evolves according to (8), where i_t is its current investment, and $\delta \in (0,1)$ is the rate of capital depreciation. The growth rate of labor-augmenting technological progress is $\gamma - 1$, and all variables measured in units of output are deflated by the level of labor-augmenting technological progress. Equation (7) reflects the restrictions on an intermediate good firm's capital adjustment. Specifically, investments that are sufficiently minor relative to the existing capital are only subject to a convex adjustment cost. The latter is measured in terms of the aggregate good, and it is given by $\epsilon_\psi K_t \left(\frac{i_t}{K_t}\right)^2$, with parameter $\epsilon_\psi \geq 0$. The range of exemption is defined by parameters a and b, with $a \leq 0 \leq b$. Otherwise, an intermediate good firm also needs to pay a fixed adjustment cost, f_t, measured in units of labor and drawn from a time-invariant uniform distribution $U : [0, f] \to [0, 1]$. Adjustment

[8] In order to lighten the notation in this place of the text, we omit a j-index to refer to the intermediate good firm being modeled, one among the continuum of intermediate good firms in our model.

cost shocks are *iid* across firms and over time. Labor used in the production of intermediate goods is denoted by $\widetilde{L}_t$, and D_t is meant to indicate dividends, measured in terms of the aggregate good. Each intermediate good firm produces its output by combining labor, $\widetilde{L}_t$, with its predetermined capital stock, K_t. The corresponding parameters in the production function are ν and α. Total factor productivity is common across intermediate good firms and evolves according to $z_0 e^{\gamma t}$.[9] Finally, x_t is an intermediate good firm's idiosyncratic productivity, which is assumed to follow a Markov chain.

2.3 Retail Firms

Retail firms introduce the New Keynesian (NK) elements into our model. Since the details of the NK model have been discussed elsewhere (see, e.g., Woodford 2003 or Galí 2015 for textbook treatments) we turn directly to the implied set of optimality conditions. A standard representation reads

$$\Pi_t = \left[\theta_p + (1-\theta_p)\left(p_t^*\right)^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}}, \tag{9}$$

$$q_t^M = \frac{1}{M_t}, \tag{10}$$

$$\Phi_t P_t^* = \mu_p \Upsilon_t P_t, \tag{11}$$

where θ_p is the Calvo parameter, i.e., the probability according to which a firm is not allowed to change price in a given period. We have also used the notation $p_t^* \equiv \frac{P_t^*}{P_{t-1}}$ for the optimal newly set price, P_t^*, that is chosen by all time t price-setters in our model, relative to the price of the consumption good one period earlier. The average price markup in period t is M_t, and $\mu_p \equiv \frac{\varepsilon}{\varepsilon-1}$ denotes the desired frictionless

[9] In Khan and Thomas (2008) total factor productivity is stochastic. This difference is explained by our research question. In fact, we restrict our attention to the monetary transmission mechanism.

markup. Finally, Φ_t and Υ_t are functions of the form

$$\begin{aligned}\Phi_t &= Y_t + \theta_p E_t \left\{ \Pi_{t,t+1}^{\varepsilon} Q_{t,t+1} \Phi_{t+1} \right\}, \\ \Upsilon_t &= Y_t + \theta_p E_t \left\{ \Pi_{t,t+1}^{\varepsilon+1} Q_{t,t+1} \Upsilon_{t+1} \right\},\end{aligned}$$

where Y_t denotes aggregate output, defined as a Dixit-Stiglitz aggregate of the same form as the consumption aggregate.

2.4 To Close the Model

All markets are assumed to clear. Specifically, the aggregate goods market clearing condition reads

$$Y_t = C_t + \int_0^1 \Psi \left(K_t \left(j \right), K_{t+1} \left(j \right) \right) dj, \tag{12}$$

where $K_t \left(j \right)$ is meant to indicate intermediate good firm j's time t capital stock. The labor market clearing condition is of the form

$$L_t = \int_0^1 \widetilde{L}_t \left(j \right) dj + \int_0^1 f_t \left(j \right) J \left(\frac{i_t \left(j \right)}{K_t \left(j \right)} \right) dj,$$

where $J \left(x \right) = 0$, if $x \in [a, b]$, and $J \left(x \right) = 1$ otherwise. Finally, we follow Walsh (2005) in assuming a monetary policy rule of the form

$$R_t = \left(R_{t-1} \right)^{\rho_r} \left[\frac{\Pi}{\beta} \left(\frac{\Pi_t}{\Pi} \right)^{\gamma_\pi} \left(\frac{Y_t}{Y} \right)^{\gamma_y} \right]^{1-\rho_r} e^{e_{r,t}}. \tag{13}$$

Parameters γ_π and γ_y indicate the long-run responsiveness of the nominal interest rate to changes in current inflation and output,[10] respectively, and parameter ρ_r

[10] Usually, the output gap, i.e., the ratio between equilibrium output and natural output (defined as the equilibrium output under flexible prices) enters the specification of monetary policy. Notice, however, that natural output does not change in response to a monetary disturbance.

measures interest rate smoothing. We adopt the convention that a variable without time subscript indicates its steady state value. The shock, $e_{r,t}$, is $i.i.d.$ with zero mean.

2.5 Baseline Calibration

We consider a quarterly model. There are three sets of parameters. For the parameters in the first set we assign values that are standard in the NK literature. They are shown in table 1.

[Table 1 about here]

We wish to make our model consistent with the micro facts on lumpy investment reported in Cooper and Haltiwanger (2006). This means that we add the serial correlation in investment rates to the targets that are also considered in Khan and Thomas (2008). Nevertheless, some of their modeling choices are still an excellent starting point for our purposes. However, since the length of a period corresponds to one year in their model, we had to adjust some of the parameter values that are taken from Khan and Thomas (2008) in an appropriate way. We also set the rate of depreciation, δ, to a value that makes our model consistent with the conventional 10% annual rate of investment in the stationary distribution. Those parameter values belong to the second set, and they are shown in table 2.

[Table 2 about here]

As Khan and Thomas (2008) do, we model idiosyncratic productivity shocks and the Markov chain determining their evolution by discretizing a log-normal process

$$\log \epsilon_{t+1} = \rho_\varepsilon \log \epsilon_t + \eta_t,$$

where η_t is iid with standard deviation σ_η. We also follow Khan and Thomas (2008) in assuming $|a| = b$. We then choose the parameter values in the third set. They measure, respectively, the range of exemption from capital adjustment costs (b), the standard deviation of the idiosyncratic shock to productivity (σ_η), the upper bound of the fixed cost distribution (f), and the convex (ϵ_ψ) portion of the capital adjustment cost. At this point, we find it useful to consider two alternative calibrations. They are shown in table 3. Those parameter values are used in order to target the objects that are stated in table 4.

[Tables 3 and 4 about here]

The first one is our baseline. In this case, we set ϵ_ψ to 1.5, combined with $f = 0.9$ and $\sigma_\eta = 0.08$. This makes our model consistent with both a plausible number of investment spikes and a realistic persistence in annual investment rates at the firm level. We also choose $b = 0.011/4$. This implies that there is no inaction nor any negative investment in the stationary distribution of our model. Intuitively, negative investment is unattractive to firms in the presence of substantial costs of adjusting the capital stock, and in this high costs environment variations of the range of exemption give rise to abrupt changes in the frequency of inactive firms. In order not to have a knife-edge result for this frequency, we simply choose a range of exemption for which all investment is positive. This is different in the calibration that is called "Low Capital Adjustment Costs" in table 4. In this case, we set $f = 0.008$, a value that is close to the one chosen in Khan and Thomas (2008). The convex portion of the capital adjustment cost is correspondingly also much smaller than in the baseline. Concretely, we set $\epsilon_\psi = 0.0063$. When combined with $\sigma_\eta = 0.037$ and $b = 0.0096$ this calibration makes our model reasonably consistent with the micro-facts reported in Cooper and Haltiwanger (2006).

For future reference, let us also mention two additional calibrations. The term "Khan and Thomas (2008)" in table 4 is meant to indicate a version of "Low Capital Adjustment Costs" that takes away the convex portion of the capital adjustment cost from that specification. All the remaining parameter values are held constant, and also in this version of the model the representative household spends one third of time working in the labor market. The problematic aspect of the "Khan and Thomas (2008)" model is that it does not allow us to target the positive serial correlation in investment rates, which is one of the micro-facts that have been established by Cooper and Haltiwanger (2006). The label "Traditional" model in table 4 refers to a version of our baseline without any idiosyncratic productivity shocks and without any range of exemption from the fixed capital adjustment cost. The models in Thomas (2002) and Reiter et al. (2013) would also fall into this category. The (well know) problematic aspects of "Traditional" models are manyfold. For instance, they imply that a very large portion of positive investment takes the form of an investment spike.

2.6 Solution Method

We solve the model by linearization around the stationary state without aggregate shocks (see Reiter 2009), using almost-exact state aggregation (see Reiter 2010). The details of how to handle the non-convexity of the firm problem are explained in Reiter (2019). We solve the firm problem on a discrete grid. For the value function, we use 400 grid points in the capital dimension, and 51 grid points for idiosyncratic productivity. We approximate the cross-sectional distribution with 1000 grid points in the capital dimension, and again 51 grid points for idiosyncratic productivity. This implies an aggregate state space of about 51000 variables. The loss-less state reduction shrinks the state space to 367 variables. The loss-less value function

reduction writes the 20400 elements of the value function as a linear combination of 157 basis functions. As a robustness check, we add an i.i.d. productivity shock on top of the Markov productivity shock. The results are almost identical. Similarly, changing grid sizes has no relevant effect on the results.

3 The Monetary Transmission Mechanism

We wish to isolate the role of a realistic degree of lumpiness in plant-level investment for the monetary transmission mechanism. It is natural to start by comparing our baseline calibration to a standard textbook treatment of this mechanism.

3.1 Baseline

Figures 1 and 2 illustrate the dynamic consequences of a 100 basis point decrease in the annualized nominal interest rate. The rate of inflation is also annualized. All other variables are measured as the respective log deviation of the original variable from its steady state value.

[Figure1 1 and 2 about here.]

Figure 1 shows the monetary transmission mechanism under our baseline calibration, whereas figure 2 displays the corresponding outcome under a standard textbook calibration of our model. In the latter case, the convex adjustment cost parameter, ϵ_ψ, is set to 9, there is no fixed cost, no range of exemption and no idiosyncratic shocks to intermediate goods firms' productivity. The results shown in figures 1 and 2 are similar, and they also resemble the corresponding outcomes in Galí (2015, p.69).[11] He observes that the dynamic consequences of monetary policy shocks, as implied by a Calvo pricing model, are (at least qualitatively) con-

[11] His model does not feature endogenous capital accumulation though.

sistent with the empirical evidence that has been obtained using structural vector autoregressive (SVAR) methods. A similar observation can be made for the results shown in figures 1 and 2. In fact, those calibrations predict that monetary policy shocks have strong and persistent consequences for real variables. For instance, the estimates reported by Christiano et al. (2005) indicate that the maximum output response to an identified monetary policy shock is about 0.5 percent (with 95 percent confidence interval around this point estimate of about $\pm$ 0.2).[12] After that, output is estimated to take about one and a half years to revert to its original level which is in line with the model's prediction. Christiano et al. (2005) also estimate a maximum investment response of about one percent (with 95 percent confidence interval around this point estimate of about $\pm$ 0.5). The estimated maximum consumption response is roughly 0.2 percent (with 95 percent confidence interval around this point estimate of about $\pm$ 0.1). By and large, the results shown in figures 1 and 2 are consistent with that evidence. Moreover, both specifications are also able to capture the observed inertial behavior of inflation, but the maximum inflation response lies outside the empirically plausible range. In fact, Christiano et al. (2005) estimate a maximum inflation response of roughly 0.2 percent (with 95 percent confidence interval around this point estimate of about $\pm$ 0.15).[13] The reason is, of course, that price-setting and investment decisions take place in two different sectors of our model. Price-setters therefore do not internalize the consequences of their price-setting decisions for the marginal costs that they are expecting to face over the life-time of a newly chosen price. Assuming firm-specific capital would allow us to deal out of this problem, as analyzed in Sveen and Weinke (2005).[14] In the

[12] The maximum response is estimated to occur about six quarters after the shock. This is one reason why additional real and nominal frictions are typically added to New Keynesian models in order to increase their empirical realism. See, e.g., Christiano et al. (2005).

[13] The estimated maximum inflation response occurs about two years after the shock.

[14] The basic intuition has been developed in Galí et al. (2001) and Sbordone (2002) in the context of models featuring decreasing returns to labor resulting from a fixed capital stock at the firm level.

present paper, however, we wish to focus on the role of investment behavior for the monetary transmission mechanism. What are the economic mechanisms at work?

3.2 Inspecting the Mechanism

In order to uncover the economic mechanism behind our results we find it useful to consider the "Low Adjustment Costs" calibration shown in table 4. Figure 3 illustrates the result.

[Figure 3 about here.]

The dynamic consequences of the monetary policy shock under consideration are out of line with their empirical SVAR counterpart. In fact, as figure 3 makes clear, the impact responses of investment and output to a change in the nominal interest rate become very large and the dynamic consequences of that shock are only short-lived. For instance, in the period when the monetary shock hits the economy investment deviates by about seven percent from its steady state value. In other words, the impact investment response is about seven times larger than the size of the response that appears to be plausible based on the above mentioned SVAR evidence. The (S,s) nature of investment decisions is crucial to understand this result. To show this, one simply needs to follow well-trodden paths. In response to an expansionary monetary policy shock firms choose to undertake some of the investment activity that they would have otherwise done later. This is crucially different in the presence of capital adjustment costs of a size that allow us to entertain a plausible monetary transmission mechanism, as we have seen in figure 2.

Let us further inspect the economic mechanisms at work. To this end, we compare the monetary transmission mechanism for three versions of our model: (i) "Low

Sveen and Weinke (2005) have shown that this simple intuition also helps understand the large degree of endogenous price stickiness that is implied by the assumption of firm-specific capital.

Adjustment Costs", (ii) "Khan and Thomas (2008)", i.e., a version of "Low Adjustment Costs" where only the convex portion of the capital adjustment cost is taken away from that specification, (iii) "Traditional", i.e., a version of our baseline where not only the convex portion of the capital adjustment cost but also the idiosyncratic productivity shocks as well as the range of exemption from the fixed capital adjustment cost are taken away. As it turns out, the results are very similar to the outcome that is shown in figure 3. We saw in table 4 that the small convex portion of the capital adjustment plays an important role for the "Low Adjustment Costs" calibration. In fact, as discussed there, without that portion we are unable to match the positive serial correlation in investment rates that is reported by Cooper and Haltiwanger (2006). As far as the dynamics are concerned, this version of the model implies, however, a monetary transmission mechanism that is very similar to its counterpart under the "Low Adjustment Costs" calibration. The reason is, of course, that the positive serial correlation in investment rates at the micro level is relatively small. This puts empirical discipline on the convex adjustment cost parameter. We also saw in table 4 that the idiosyncratic productivity shocks combined with a range of exemption from the fixed capital adjustment cost play an important role for the "Low Adjustment Costs" calibration. In fact, as documented there, without those features we are unable to match many aspects of the micro data on investment. As far as the dynamics are concerned, the "Traditional" version of our model implies, however, a monetary transmission mechanism that is very similar to its counterpart under the "Low Adjustment Costs" calibration. The reason why we find this interesting is as follows. If idiosyncratic factors are relevant for investment decisions at the micro level, firms will respond differently to aggregate shocks. For a realistic size of the idiosyncratic shocks, there are always firms just at the margin between investing and not investing, which will then change behavior in response to

a change in the interest rate. With larger idiosyncratic shocks, the distribution is more spread out, and the density at those margins is smaller. This limits the extent to which firms choose to undertake some of the investment activity that they would have otherwise done later. The quantitative relevance of this effect is, however, very small. We also find it interesting to compare our baseline results to a flexible capital case. The latter is obtained by doing away with the fixed capital adjustment cost in the context of the "Traditional" model of the monetary transmission mechanism. Also in this case, the results are very similar to their counterpart under the "Low Adjustment Costs" calibration. The last result is reminiscent of the irrelevance results in Thomas (2002) and Khan and Thomas (2008).

4 Conclusion

We introduce lumpy investment into an otherwise standard New Keynesian framework, and our main result shows that a quantitatively relevant monetary transmission mechanism can be entertained in the context of the resulting model. In the investment block, we extend the (S,s) model in Khan and Thomas (2008) by allowing for a combination of convex and non-convex capital adjustment costs. The key insight is that the smooth investment pattern generated by the convex capital adjustment cost can coexist with investment spikes in yearly investment rates. At the same time, idiosyncratic shocks to firm-level productivity can generate the observed small and positive autocorrelation in yearly investment rates. Our results also show that a calibration with much smaller capital adjustment costs is even better able to match the micro data on investment at the micro level. In the latter case, however, the monetary transmission mechanism becomes counterfactual due to the extremely large interest rate sensitivity of investment that is implied by this calibration.

Estimated impulse responses to identified monetary policy shocks have many other properties that are left out of the focus of this paper, and we have noted that this is one reason why additional real and nominal frictions are typically added to New Keynesian models (see, e.g., Christiano et al. 2005). It seems to us that also the modeling of those additional frictions should be disciplined by the micro data, and this is another reason why our work is just one more step towards a micro-founded theory of the monetary transmission mechanism. In this context it would also be interesting to take up a question raised by Eberly et al. (2012): How does the right model of investment vary with the level of aggregation?

References

Calvo, Guillermo (1983), "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics*, 12(3), 383-398.

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

Cooper, Russell W., and John C. Haltiwanger (2006), "On the Nature of Capital Adjustment Costs," *Review of Economic Studies* 73(3), 611-633.

Doms, Mark, and Timothy Dunne (1998), "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics* 1(2), 409-429.

Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999), "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *The Quarterly Journal of Economics* 114(2), 655-690.

Eberly, Janice, Sergio Rebelo, and NicolasVincent "What Explains the Lagged-Investment Effect?," *Journal of Monetary Economics*, 59(3), 370-380.

Galí, Jordi (2015), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications,* Second Edition, Princeton University Press.

Galí, Jordi, Mark Gertler, and David López-Salido (2001), "European Inflation Dynamics," *Euroean Economic Review* 45(7),1237-1270.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018), "Monetary Policy According to HANK," *American Economic Review* 108(3), 697-743.

Khan, Aubhik, and Julia K. Thomas (2008), "Idiosyncratic Shocks and the Role of Noncenvixities in Plant and Aggregate Investment Dynamics," *Econometrica* 76(2),

395-436.

Koby, Yann, and Christian K. Wolf (2020), "Aggregation in Heterogenous-Firm Models: Theory and Measurement," mimeo.

McKay, Alisdair, and Johannes F. Wieland (2019), "Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy," NBER Working Paper No. 26175.

Ottonello, Pablo, and Thomas Winberry (2019), "Financial Heterogeneity and the Investment Channel of Monetary Policy," mimeo.

Reiter, Michael (2009), "Solving Heterogeneous-Agent Models by Projection and Perturbation," *Journal of Economic Dynamics and Control* 33(3), 649-665.

Reiter, Michael (2010), "Approximate Aggregation in Heterogeneous-Agent Models," IAS Economics Working Paper 258.

Reiter, Michael (2019), "Solving Heterogeneous Agent Models with Non-convex Optimization Problems: Linearization and Beyond," mimeo.

Reiter, Michael, Tommy Sveen, and Lutz Weinke (2013), "Lumpy Investment and the Monetary Transmission Mechanism," *Journal of Monetary Economics* 60(7), 821-834.

Sbordone, Argia M. (2002), "Prices and Unit Labor Costs: A New Test of Price Stickiness,", *Journal of Monetary Economics* 49(2), 265-292.

Sveen, Tommy, and Lutz Weinke (2007), "Lumpy Investment, Sticky Prices, and the Monetary Transmission Mechanism," *Journal of Monetary Economics* 54S, 23-36.

Thomas, Julia K. (2002), "Is Lumpy Investment Relevant for the Business Cycle?," *Journal of Political Economy* 110(3), 508-534.

Winberry, Thomas (2020), "Lumpy Investment, Business Cycles, and Stimulus Policy," mimeo.

Woodford, Michael (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

Woodford, Michael (2005), "Firm-Specific Capital and the New Keynesian Phillips Curve," *International Journal of Central Banking* 1(2), 1-46.

Table 1: New Keynesian Parameters

ϵ	γ_π	γ_y	ρ_r	θ_p	ϕ
7	1.5	0.125	0.7	0.75	0.5

Table 2: Parameters in the Spirit of Khan and Thomas (2008)

γ	β	δ	α	ν	ρ_ε
$1.016^{\frac{1}{4}}$	$0.977^{\frac{1}{4}}$	0.0187	0.256	0.640	$0.859^{\frac{1}{4}}$

Table 3: "Baseline" and "Low Capital Adjustment Costs" (LCAC)

	σ_η	b	f	ϵ_ψ
Baseline	0.08	0.00275	0.9	1.5
LCAC	0.037	0.0096	0.008	0.0063

Table 4: Distribution of Plant Investment Rates

	Inaction	Positive Spike	Negative Spike	Positive Invest.	Negative Invest.	Invest. Autocorr.
Data*	*0.081*	*0.186*	*0.018*	*0.815*	*0.104*	*0.058*
Baseline	0.000	0.151	0.000	1.000	0.000	0.054
LCAC	0.085	0.234	0.008	0.758	0.156	0.066
KT (2008)	0.103	0.262	0.024	0.690	0.207	-0.071
Traditional	0.785	0.159	0.000	0.214	0.000	-0.121

*Establishment data are from Cooper and Haltiwanger (2006).

Inaction $|i/k| < 0.01$; positive spike, $i/k > 0.20$; negative spike, $i/k < -0.20$;

positive investment $i/k \geq 0.01$; negative investment, $i/k \leq -0.01$; serial correlation of i/k

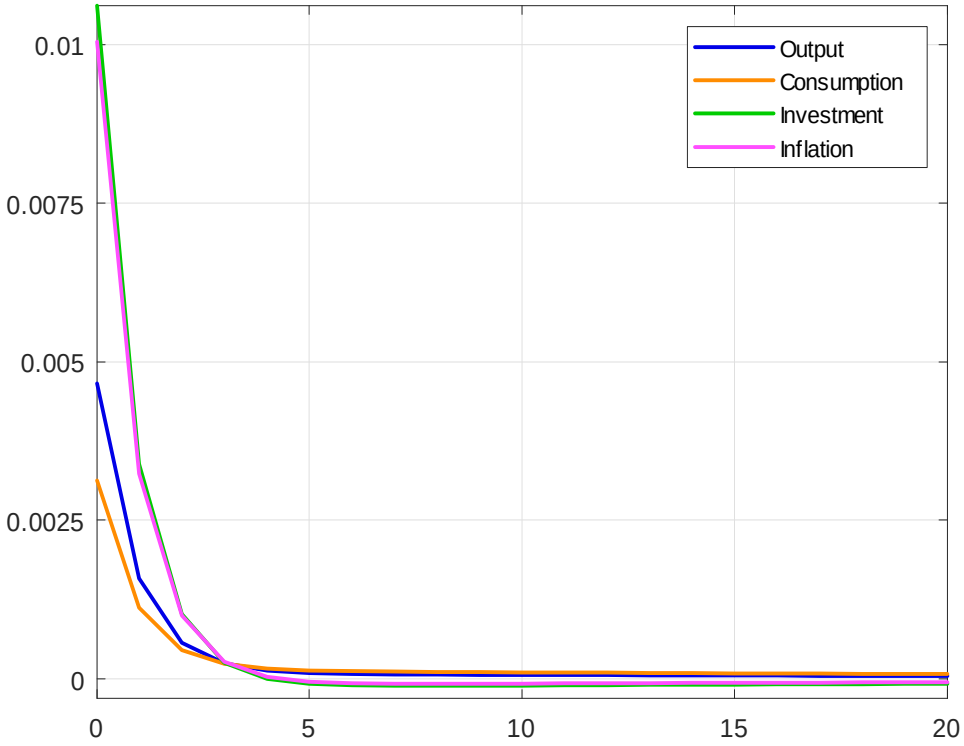

Figure 1: Monetary Transmission Mechanism (Baseline)

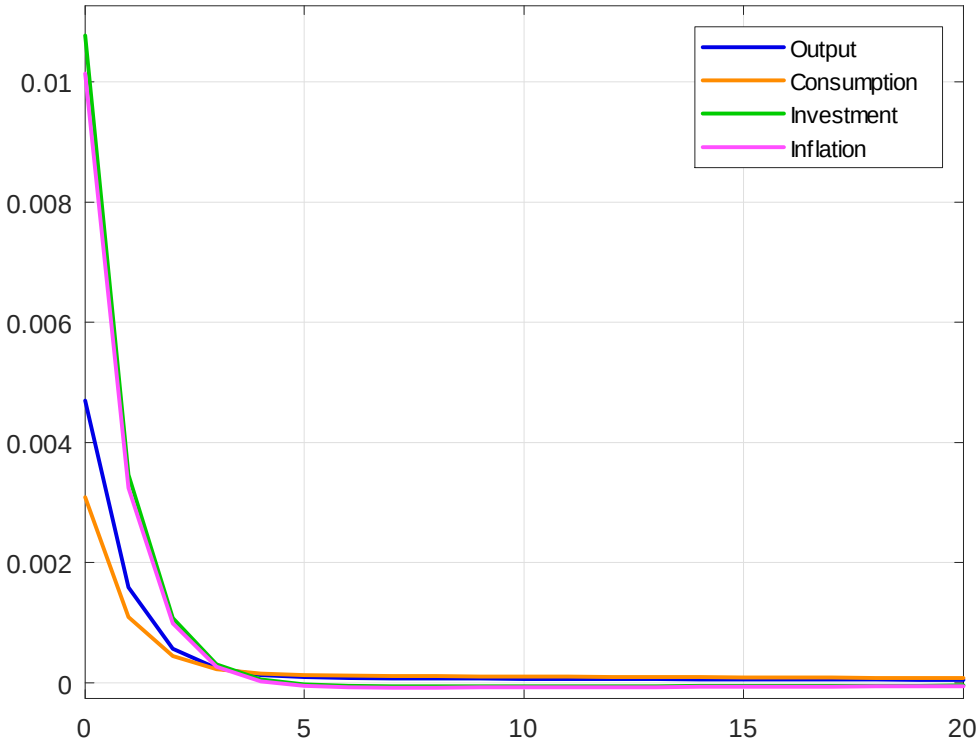

Figure 2: Monetary Transmission Mechanism (Standard Textbook Case)

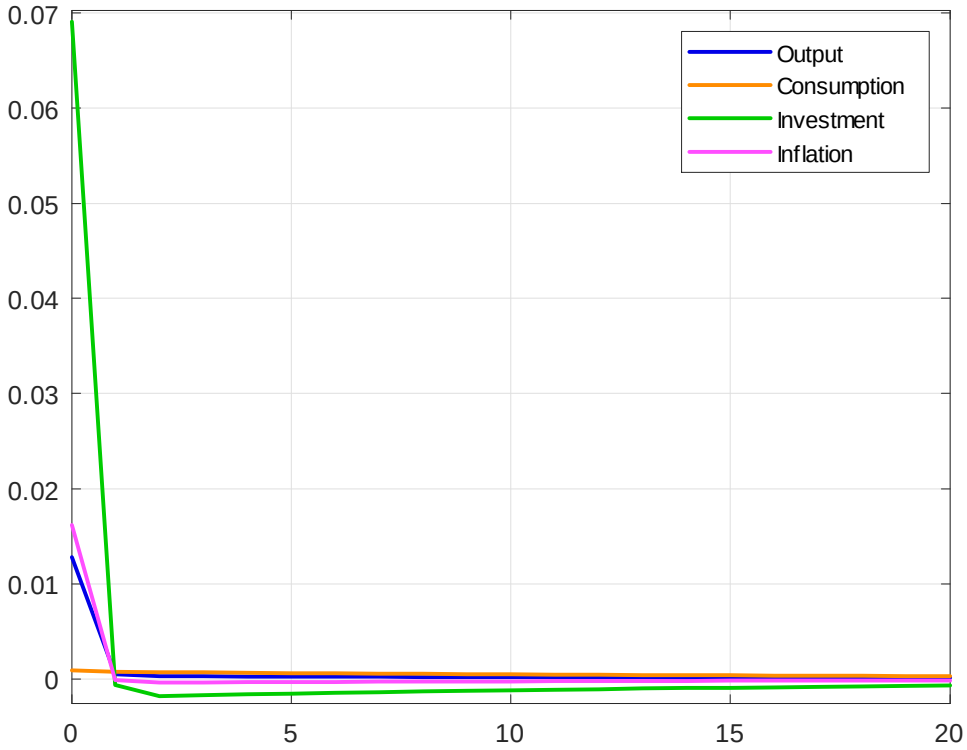

Figure 3: Monetary Transmission Mechanism (LCAC)