

Lumpy Investment, Business Cycles, and Stimulus Policy[†]

By THOMAS WINBERRY*

I study the aggregate implications of micro-level lumpy investment in a model consistent with the empirical dynamics of the real interest rate. The elasticity of aggregate investment with respect to shocks is procyclical because more firms are likely to make an extensive margin investment in expansions than in recessions. Matching the dynamics of the real interest rate is key to generating this result because it disciplines the interest-elasticity of investment and avoids counterfactual behavior of the model that would otherwise eliminate most of the procyclical responsiveness. Therefore, data on interest rates place important discipline in aggregating micro-level investment behavior. (JEL D25, E13, E22, E23, E43, G31, H25)

Aggregate investment is one of the most volatile components of GDP over the business cycle, accounting for a large share of the decline in GDP during recessions. These large swings in aggregate investment are primarily driven by changes in the number of firms undertaking an investment project (the extensive margin) rather than changes in the size of investment projects (the intensive margin).¹ However, most dynamic stochastic general equilibrium (DSGE) macro models make the simplifying assumption that all changes in aggregate investment are the result of a representative firm operating along the intensive margin. Therefore, a key question for business cycle modeling is: does this abstraction matter for understanding aggregate investment dynamics? In a benchmark real business cycle (RBC) model, the answer to this question is no; general equilibrium changes in the real interest rate bring the aggregate investment series in line with the preferences of the representative household, regardless of whether that investment occurs along the extensive or intensive margin at the micro level.²

In this paper, I argue that accounting for the extensive margin at the micro level is important for understanding aggregate investment dynamics. My argument has two main components. First, I show that the dynamics of the real interest rate which drive

*University of Chicago, Booth School of Business (email: Thomas.Winberry@chicagobooth.edu). Mark Aguiar was the coeditor for this article. I am extremely grateful to Richard Rogerson, Pablo Ottonello, Greg Kaplan, and Esteban Rossi-Hansberg for many conversations that shaped this paper. I also thank Diane Alexander, Adrien Auclert, Gabe Chodorow-Reich, Jesús Fernández-Villaverde, Oleg Itskhoki, Nobu Kiyotaki, Yann Koby, Ben Moll, Ezra Oberfield, Chris Sims, Christian vom Lehn, Christian Wolf, Eric Zwick, and four anonymous referees. Financial support from Chicago Booth is gratefully acknowledged.

[†]Go to <https://doi.org/10.1257/aer.20161723> to visit the article page for additional materials and author disclosure statement.

¹See, for example, the evidence in Doms and Dunne (1998) or Gourio and Kashyap (2007).

²This irrelevance result was established in an important series of papers by Thomas (2002), Khan and Thomas (2003, 2008), and further elaborated upon by House (2014).

the irrelevance results in the RBC environment are at odds with the data. Second, I build a heterogeneous firm model consistent with both the importance of the extensive margin of investment and the observed dynamics of the real interest rate, and find that the behavior of aggregate investment is substantially different than in the representative firm model. In particular, the elasticity of aggregate investment with respect to shocks is procyclical; in expansions, more firms are close to making an extensive margin investment, so an additional shock generates more total investment than it would otherwise. These results illustrate how data on interest rates place discipline on the role of general equilibrium in determining the aggregate implications of micro-level investment behavior.

In the benchmark RBC environment, changes in the real interest rate eliminate the procyclical responsiveness to shocks because investment is extremely interest-sensitive; therefore, small but procyclical movements in the real interest rate are able to restrain the large movements in the extensive margin of investment. However, I show that these movements in the interest rate are inconsistent with two key features of the data. First, the interest rate is negatively correlated with aggregate output and productivity, suggesting that it does not dampen cyclical movements in investment demand. Second, the interest rate is an order of magnitude more volatile in the data than in the RBC model, suggesting that investment is not as interest-sensitive as it is in the model.

Motivated by this evidence, I extend a simple heterogeneous firm model with extensive margin investment decisions to capture the empirical dynamics of the real interest rate. In the model, there is a fixed mass of firms who make investment decisions subject to both fixed and convex capital adjustment costs. The presence of fixed costs generates the extensive margin choice of whether to invest or not. There is a representative household whose preferences feature habit formation and aggregate dynamics are driven by shocks to aggregate productivity. The equilibrium dynamics of the real interest rate are determined by the strength of habit formation, which controls the sensitivity of investment supply to aggregate shocks, and the strength of adjustment costs, which control the sensitivity of investment demand to shocks. I calibrate these features of the model to match both the dynamics of the real interest rate and the behavior of micro-level investment.

Quantitatively, my calibrated model predicts that aggregate investment is nearly 30 percent more responsive to an aggregate shock starting from a brisk expansion than it is starting from a deep recession. As described above, this procyclical responsiveness reflects the fact that more firms are close to making an extensive margin investment in expansions, so an additional shock induces more firms to invest. However, as Khan and Thomas (2008) show, embedding this mechanism into an otherwise standard RBC model (without habit formation or convex adjustment costs) generates little variation in the responsiveness to shocks. In my version of their calibration, aggregate investment is only 10 percent more responsive to shocks in the expansion compared to recession.

My model generates substantial variation in responsiveness for two related reasons. First, matching the negative comovement between the real interest rate and aggregate productivity implies that the interest rate does not directly dampen the effect of a productivity shock. Second, matching both the dynamics of the interest rate and the behavior of investment indirectly disciplines the interest-sensitivity of

investment, which Koby and Wolf (2020) show is key to determining aggregation in this class of models. In my version of Khan and Thomas's (2008) calibration, the semi-elasticity of aggregate investment with respect to the real interest rate is approximately $-1,075$, so small changes in the interest rate have a strong influence on aggregate investment. In my model, the corresponding semi-elasticity is only -6.73 , giving general equilibrium a much smaller influence over aggregate dynamics.

I illustrate two implications of the model for investment stimulus policies, such as investment tax credits or the bonus depreciation allowance. First, the aggregate effect of investment stimulus is also state dependent and falls in recessions; therefore, predictions based on linear models overstate the effectiveness of these policies in recessions. Second, I develop a simple size-dependent stimulus policy that increases cost effectiveness by 30 percent compared to existing size-independent policies. The main insight of this alternative policy is to avoid subsidizing investment that would have been done even without the policy. Because investment primarily occurs along the extensive margin, most of this inframarginal waste is accounted for by subsidizing firms that would have made an extensive margin investment without the policy. In my model, as in the data, small firms grow faster than average and are therefore more likely to be inframarginal to the policy.

A key challenge throughout the analysis is efficiently computing the equilibrium of the model, which involves approximating the entire cross-sectional distribution of firms. I use the method developed concurrently in Winberry (2018), which approximates the distribution with a flexible but finite-dimensional parametric family. I find that this approach captures how changes in the shape of the distribution affect the dynamics of aggregate variables more accurately than simply using the mean capital stock.

Related Literature.—This paper contributes to three main strands of literature. First, it addresses the long-standing question of how the extensive margin of investment impacts aggregate dynamics. Early work, analyzing firms' decision rules with fixed prices, finds that the extensive margin generates procyclical responsiveness to shocks (as in my model).³ However, Thomas (2002) and Khan and Thomas (2003, 2008) show that most of this time-varying elasticity disappears when prices are endogenized in an otherwise standard RBC framework, rendering the extensive margin irrelevant for aggregate dynamics. House (2014) suggests that these irrelevance results are driven by the extreme sensitivity of investment to changes in the relative price of investment goods in a stylized partial equilibrium model. In recent work, Koby and Wolf (2020) analytically show that the elasticity of investment with respect to its user cost is a sufficient statistic for the aggregation properties of a wide class of models with firm heterogeneity. They argue against an extreme user cost elasticity by matching the empirical response of investment to tax changes measured in Zwick and Mahon (2017). I show that matching the dynamics of the real interest

³ See, for example, Caballero, Engel, and Haltiwanger (1995) or Caballero and Engel (1999).

rate also requires breaking this extreme sensitivity and, therefore, a key source of the irrelevance results.⁴

To match the dynamics of the real interest rate, I follow Beaudry and Guay (1996) in using habit formation and capital adjustment costs. Boldrin, Christiano, and Fisher (2001) also use this approach to match interest rate dynamics and the level of the equity premium. These papers work in representative agent environments; my results show that many of their lessons carry over to a heterogeneous firm environment in which adjustment costs are disciplined with micro-level investment data.

Finally, this paper contributes to a large literature which studies investment stimulus policy. Many papers estimate the effect of stimulus policy through linear regression models, most recently House and Shapiro (2008) and Zwick and Mahon (2017). Edge and Rudd (2011) introduce the Bonus Depreciation Allowance into a linearized New Keynesian model, which rules out state dependence by construction. I focus on the effect of stimulus policy over the business cycle and how micro-level targeting can increase its cost effectiveness.⁵

Road Map.—The rest of this paper is organized as follows. Section I describes the role of the real interest rate in driving the existing irrelevance results in the literature and argues that the implied interest rate behavior is counterfactual. Section II develops my quantitative heterogeneous firm model, which Section III parameterizes to jointly match micro-level investment behavior and macro-level interest rate dynamics. Section IV shows that the existence of the extensive margin implies that aggregate investment is more responsive to shocks in expansions than in recessions, and argues that matching the dynamics of the interest rate is key to generating this result in general equilibrium. Section V introduces stimulus policy into the model, shows that the effectiveness of these policies falls in recessions, and develops the alternative size-dependent stimulus to increase cost effectiveness. Finally, Section VI concludes.

I. Role of Real Interest Rate Dynamics

This section motivates the features of real interest rate dynamics on which I will focus for the rest of the paper. Section IA uses a simplified RBC model to illustrate the role of the interest rate in rendering the extensive margin of investment irrelevant

⁴Other papers challenge the irrelevance results on other grounds. Gourio and Kashyap (2007) and Miao and Wang (2014) show that the results are sensitive to the distribution of fixed adjustment costs and the degree of returns to scale. Bachmann, Caballero, and Engel (2013) argue that Khan and Thomas's (2008) calibrated fixed costs are implausibly small and that increasing them to empirically reasonable levels breaks the irrelevance results. Bachmann and Ma (2016) and Bayer and Tjaden (2016) argue that the irrelevance results are sensitive to the precise form of general equilibrium; Bachmann and Ma (2016) allow for two savings vehicles (physical capital and inventories) and Bayer and Tjaden (2016) allow for multiple countries. Finally, and most closely related to this paper, Cooper and Willis (2014) parameterize an interest rate process from the data and solve firms' decision problems given this process. My paper produces such an interest rate process endogenously in general equilibrium.

⁵Berger and Vavra (2015) analyze a related class of consumer durable stimulus policies in a model of lumpy durable investment. They find that stimulus policies are less effective in recessions for similar reasons as here; however, they focus on detailed features of the microdata while I focus on the role of real interest rate dynamics in aggregation and on designing more cost-effective policies.

for aggregate dynamics. Section IB shows that the key features of the interest rate which drive this result are inconsistent with the data.

A. Irrelevance of Extensive Margin in a Simple RBC Model

I use a stylized model to simplify the exposition. In this model, I take returns to scale to be nearly constant, which implies that investment is nearly infinitely sensitive to changes in the real interest rate. Therefore, small but procyclical movements in the real interest rate are able to bring aggregate investment in line with the representative household's desired path of consumption, regardless of the existence of fixed costs.

This analysis builds heavily on related work in the literature. Miao and Wang (2014) also show that fixed costs are irrelevant for aggregate dynamics under constant returns to scale, though they do not relate that mechanism to interest-sensitivity or allow for idiosyncratic productivity shocks. House (2014) shows that low depreciation rates can also generate extreme price-sensitivity of investment and render fixed costs irrelevant for aggregates. Finally, Koby and Wolf (2020) show that the user-cost elasticity of investment is a sufficient statistic for characterizing the aggregation properties of a general class of heterogeneous firm models. The main value added of my analysis here is to summarize the mechanism in a transparent way and to study its empirical implications for interest rate dynamics.

Simple RBC Model with Fixed Costs.—Consider a discrete time environment with heterogeneous firms indexed by $j \in [0, 1]$. Firm j produces output y_{jt} using the production function

$$(1) \quad y_{jt} = z_t \varepsilon_{jt} k_{jt}^\alpha,$$

where z_t is aggregate productivity, ε_{jt} is idiosyncratic productivity, k_{jt} is the firm's capital stock, and the parameter $\alpha \leq 1$ controls the returns to scale. Idiosyncratic productivity ε_{jt} follows a first-order Markov process with finite support $\varepsilon \in \{\varepsilon_1, \dots, \varepsilon_{n_\varepsilon}\}$. Firms have perfect foresight over the path of aggregate productivity z_t ; since firms are owned by the representative household (as described below), the absence of aggregate uncertainty implies that firms use the risk-free rate r_t to discount future profits. The capital stock k_{jt} is predetermined at time t and each period the firm chooses next period's capital $k_{j,t+1} = (1 - \delta)k_{jt} + i_{jt}$, where i_{jt} is gross investment and δ is the depreciation rate. Nonzero gross investment incurs a fixed resource adjustment cost $\bar{\xi}$ which is rebated lump sum to the representative household. The initial distribution of idiosyncratic productivity and capital across firms is stationary absent changes in aggregate TFP z_t .

There is a representative household with preferences over consumption C_t represented by the utility function

$$\sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\sigma} - 1}{1-\sigma},$$

where β is the discount factor and $1/\sigma$ is the elasticity of intertemporal substitution. The household owns all firms in the economy. Total output can be used for consumption or investment, which implies the aggregate resource constraint

$$Y_t = C_t + K_{t+1} - (1 - \delta)K_t,$$

where $Y_t = \int y_{jt} dj$ and $K_t = \int k_{jt} dj$.

PROPOSITION 1: *Let the returns to scale parameter $\alpha < 1$. Let $k_t^*(\alpha)$ be the optimal capital accumulation policy at date t for a firm with idiosyncratic productivity $\tilde{\varepsilon} = \max_i E[\varepsilon'_i | \varepsilon_i]$, conditional on paying the fixed cost. Let $\pi_t^*(\alpha)$ denote the flow profits associated with that choice. If $\tilde{\xi} \leq \pi_t^*(\alpha)$, then*

$$(2) \quad r_t + \delta \rightarrow z_{t+1} \tilde{\varepsilon} \quad \text{as } \alpha \rightarrow 1.$$

Furthermore, aggregate output Y_t , investment I_t , and consumption C_t approach the outcomes of a representative firm model with aggregate productivity $\tilde{Z}_t = Z_t \tilde{\varepsilon}$ and initial capital stock K_0 equal to the aggregate capital stock in the invariant distribution.

PROOF:

See online Appendix Section I.

In the limit of Proposition 1, the fixed cost $\tilde{\xi}$ is irrelevant in the sense that the aggregate outcomes can be derived from a representative firm without the fixed cost. As the returns to scale approaches 1, firms become infinitely elastic with respect to changes in the real interest rate because their profit function becomes linear with respect to capital. Since general equilibrium requires consumption to be positive and finite at all dates t , and therefore investment be finite at all dates t , the real interest rate adjusts to ensure that firms with the highest value of expected future productivity earn zero profits. At this real interest rate, only firms with $E[\varepsilon'_i | \varepsilon_i] = \tilde{\varepsilon}$ accumulate capital and all other firms do not.

The irrelevance of fixed costs is ensured by the real interest rate adjusting to (2) in order to generate an equilibrium which satisfies the household's preference for positive consumption. Although quantitative models do not exactly satisfy the conditions in Proposition 1, they are often close to that limit, giving the representative household's preferences for smooth consumption a powerful influence over aggregate investment. This occurs in quantitative models for two reasons. First, allowing for modest decreasing returns $\alpha < 1$ does not break the extreme sensitivity of investment with respect to the real interest rates. The semi-elasticity of investment with respect to the real interest rate for adjusting firms is

$$(3) \quad \frac{\partial i_{jt} / i_{jt}}{\partial r_t} = -\frac{1}{\delta} \frac{1}{1 - \alpha} \left(\frac{1 + r_t}{r_t + \delta} \right).$$

As $\alpha \rightarrow 1$, this semi-elasticity (3) becomes infinite. However, even under a more typical calibration, for example, with $\alpha = 0.7$, $\delta = 0.025$, and $r_t = 0.01$, the semi-elasticity is approximately $-3,847$.⁶ Second, the requirement of small fixed costs ξ is often guaranteed by assuming that fixed costs are a random draw from a $U[0, \xi]$ distribution, which ensures that there is always a positive mass of firms with arbitrarily small fixed costs.⁷

B. Comparing Interest Rate Dynamics to the Data

The simple framework described in Section IA implies that the real interest rate must move one-for-one with productivity to ensure that the zero variable profit condition (2) holds, generating exact aggregation to a representative firm. In more general models, the real interest rate adjusts to ensure that the zero profit approximately holds, generating approximate aggregation. In this subsection, I show that these real interest rate dynamics are counterfactual.

Measurement.—I study the joint dynamics of the real interest rate, aggregate total factor productivity (TFP), and aggregate output in the US data 1954:I–2016:IV. I measure the real interest rate r_t as the nominal return on 90-day Treasury bills adjusted for realized CPI inflation. I measure aggregate productivity Z_t as the aggregate Solow residual. Finally, I measure output Y_t as real GDP. Details of the data construction are contained in online Appendix Section II.A. I compare the data to the benchmark RBC model, which is quantitatively close to models in which fixed costs are irrelevant for aggregate dynamics (as discussed in the introduction). The detailed model specification and calibration, which follows standard practice in the business cycle literature, is contained in online Appendix Section III.

Descriptive Results.—Table 1 shows that the RBC model is counterfactual along two key dimensions. First, the real interest rate is negatively correlated with aggregate TFP (-0.20), while in the model the two are extremely positively correlated (0.97). Second, the standard deviation of the interest rate is an order of magnitude higher in the data (1.73 percent) than in the model (0.17 percent). Table 1 also shows that these two findings hold in three different subsamples of the data: pre-1983, post-1983, and excluding the years near the Volcker recession.

Figure 1 further illustrates the stability of these two statistics by plotting eight-year rolling windows of each statistic over the postwar sample. The correlation of the real interest rate with aggregate productivity is negative for nearly the entire sample. The correlation of the real interest rate with GDP is negative before 1983 but positive after 1983, consistent with the fact that productivity is less procyclical in the later period. Nevertheless, both correlations are consistently below the level implied by

⁶Gourio and Kashyap (2007) calibrate a strong degree of decreasing returns and show that helps break this irrelevance result.

⁷Even in models in which fixed costs are nonrandom and strictly positive, aggregate investment may be extremely interest-sensitive. For example, House (2014) shows that when $\delta \rightarrow 0$, the elasticity of the timing of investment episodes with respect to the relative price of capital is infinite, playing a similar role to the infinite interest elasticity here.

TABLE 1—CYCLICAL DYNAMICS OF RISK-FREE RATE

	$\sigma(r_t)$ (percent)	$\rho(r_t, y_t)$	$\rho(r_t, z_t)$
Whole sample	1.73	−0.11	−0.20
(<i>p</i> -value)		(0.09)	(0.00)
No Volcker	1.61	−0.03	−0.17
		(0.66)	(0.01)
Pre-1983	1.57	−0.39	−0.17
		(0.00)	(0.07)
Post-1983	1.87	0.21	−0.24
		(0.01)	(0.01)
RBC	0.17	0.96	0.97

Notes: Real interest rate measured as the return on 90-day Treasury bills adjusted for realized CPI inflation, expressed in annual percentage points. Output measured as real GDP. TFP measured as the aggregate Solow residual. All variables have been HP-filtered and expressed as percentage deviation from an HP trend. *Whole sample* refers to the 1954:I–2016:IV time series. *No Volcker* excludes 1979:I–1983:IV. *Pre-1983* refers to the 1954:I–1982:IV sample. *Post-1983* refers to the 1983:I–2016:IV sample. *RBC* refers to the benchmark RBC model described in online Appendix Section III.




FIGURE 1. STABILITY OF CYCLICAL DYNAMICS OF RISK-FREE RATE

Notes: Statistics computed in forward-looking 8-year rolling windows over the sample. The left panel plots the standard deviation of the real interest rate. The right panel plots the correlation of the real interest rate with output and TFP. Dashed lines correspond to the population statistics in the benchmark RBC model described in online Appendix Section III. Real interest rate measured as the return on 90-day Treasury bills adjusted for realized CPI inflation, expressed in annual percentage points. Output measured as real GDP. TFP measured as the aggregate Solow residual. All variables have been logged and HP-filtered.

the RBC model. The standard deviation of the real interest rate is consistently above the prediction of the RBC model over the entire sample.

Impulse Response to TFP Shock.—In order to make a consistent comparison between the data and the RBC model, I estimate the impulse response of the real interest rate to a TFP shock using a simple bivariate vector autoregression (VAR):

(4)

$$X_t = \sum_{j=1}^p \Gamma_j X_{t-j} + e_t,$$




FIGURE 2. IMPULSE RESPONSE OF THE REAL INTEREST RATE TO TFP SHOCK

Notes: Impulse response of the real interest rate to a TFP shock identified from a bivariate VAR with TFP ordered first. TFP is adjusted for changes in utilization following Fernald (2014). Lag length of 3 chosen by the AIC criteria. *RBC theoretical* refers to the theoretical impulse response from the benchmark RBC model described in online Appendix Section III. *RBC measured* refers to the impulse response identified using the VAR estimation on simulated data from the model. *Empirical (90% CI)* refers to the empirical impulse response and 90 percent error bands.

where $X_t = (Z_t, r_t)^T$, p is the lag length, Γ_j are coefficient matrices, and e_t are residuals. I choose the lag length $p = 3$ following the AIC. I identify TFP shocks by assuming that shocks to the interest rate equation do not affect TFP upon impact.⁸ I use the adjusted TFP measure from Fernald (2014) in order to isolate changes in productivity that are not driven by endogenous changes in factor utilization.

Figure 2 shows that the RBC model fails along the same two dimensions highlighted in the descriptive analysis above. First, the empirical response is negative while the model's is positive, consistent with the differences in correlations presented in Table 1. Second, the magnitude of the empirical response is larger than in the model, consistent with the differences in volatilities in Table 1. The figure also shows that the model's theoretical impulse response is nearly identical to the response estimated using the VAR (4) on model-simulated data, validating the specification within that model.

Robustness.—The key mechanism driving the irrelevance result in Section IA is that the user cost of capital $r_t + \delta$ moves one-for-one with aggregate productivity. While the empirical results above indicate that is not the case, one may be concerned that a richer specification of the user cost does move one-for-one with productivity, reestablishing the zero profit condition. For instance, in a model in which consumption goods cannot be transformed into investment goods one-for-one, the user cost is $q_t(1 + r_t) - (1 - \delta)q_{t+1}$ where q_t is the relative price of capital. Table 2 shows that this generalized user cost also does not move one-for-one with productivity; in

⁸Note that this identifying assumption is weaker than the assumption that TFP is exogenous with respect to the real interest rate at all dates.

TABLE 2—CYCLICAL DYNAMICS OF USER COST u_t

	$\sigma(u_t)$ (percent)	$\rho(u_t, y_t)$	$\rho(u_t, z_t)$
Whole sample	1.98	−0.06	−0.09
(<i>p</i> -value)		0.33	0.15
No Volcker	1.93	0.05	−0.05
		0.43	(0.48)
Pre-1983	2.39	−0.13	0.14
		0.18	(0.14)
Post-1983	1.56	0.07	−0.50
		0.39	(0.00)
RBC	0.17	0.96	0.97

Notes: User cost defined as $u_t = q_t(1 + r_t) - (1 - \delta)q_{t+1}$, where r_t is the real interest rate (measured as in Table 1) and q_t is the price of investment goods relative to consumption goods (constructed by Riccardo DiCecio), log-linearly detrended. User cost is expressed in annualized percentage points. Output measured as real GDP. TFP measured as the aggregate Solow residual. All variables have been HP-filtered and expressed as percentage deviation from an HP trend. *Whole sample* refers to the 1954:I–2016:IV time series. *No Volcker* excludes 1979:I–1983:IV. *Pre-1983* refers to the 1954:I–1982:IV sample. *Post-1983* refers to the 1983:I–2016:IV sample. *RBC* refers to the benchmark RBC model described in online Appendix Section III.

fact, its time-series behavior is similar to that of the real interest rate r_t .⁹ Therefore, I focus on the simpler case in which $q_t = 1$ for the rest of the paper.¹⁰

Online Appendix Section II.B contains a number of additional robustness checks of these empirical results. The online Appendix shows that the results continue to hold if one uses a VAR to construct expected inflation, are robust to using core CPI rather than headline CPI to deflate the nominal interest rate, and are robust to different choices of business cycle filters. It also displays the impulse response of the ex ante real interest rate to a TFP shock, which is targeted in the model calibration in Section III.

II. Model

I now develop a quantitative heterogeneous firm model to be consistent with both the behavior of investment at the micro level and the dynamics of the real interest rate at the macro level.

A. Environment

The model is set in discrete time.

Firms.—The firm side of the model builds heavily on Khan and Thomas (2008), extended to include convex adjustment costs and the corporate tax code.¹¹ There is a fixed mass of firms $j \in [0, 1]$ that produce output y_{jt} using the production function

$$y_{jt} = z_t \varepsilon_{jt} k_{jt}^\theta n_{jt}^\nu,$$

⁹One could further extend the user cost to include capital taxes, but the cyclical variation in taxes is limited.

¹⁰Most empirical estimates of the user cost elasticity of investment are low and therefore also provide evidence against the extreme interest-sensitivity driving the irrelevance results. In fact, Koby and Wolf (2020) show that the user cost elasticity from Zwick and Mahon's (2017) estimates implies that fixed costs are quantitatively important for aggregate dynamics.

¹¹I include the tax code in order to study investment stimulus policy in Section V.

where z_t is an aggregate productivity shock, ε_{jt} is an idiosyncratic productivity shock, k_{jt} is the firm's capital stock, n_{jt} is its labor input, and θ and ν are parameters satisfying $\theta + \nu < 1$. The aggregate shock z_t is common to all firms and follows the AR(1) process

$$\log z_{t+1} = \rho_z \log z_t + \omega_{t+1}^z, \quad \text{where} \quad \omega_{t+1}^z \sim N(0, \sigma_z^2).$$

The idiosyncratic shock ε_{jt} is independent across firms but within firm follows the AR(1) process

$$\log \varepsilon_{jt+1} = \rho_\varepsilon \log \varepsilon_{jt} + \omega_{t+1}^\varepsilon, \quad \text{where} \quad \omega_{t+1}^\varepsilon \sim N(0, \sigma_\varepsilon^2).$$

Each period, a firm j observes these two shocks, uses its preexisting capital stock, hires labor from a competitive labor market at wage w_t , and produces output y_{jt} .

After production, the firm decides how much capital in which to invest for the next period. Gross investment i_{jt} yields $k_{jt+1} = (1 - \delta)k_{jt} + i_{jt}$ units of capital in period $t + 1$. This investment is subject to two capital adjustment costs. First, if $i_{jt} \neq 0$, then the firm must pay ξ_{jt} units of labor.¹² The fixed cost ξ_{jt} is a uniform random variable with support $[0, \bar{\xi}]$, distributed independently across firms and time. Second, any nonzero amount of investment incurs the quadratic adjustment cost $-(\varphi/2)(i_{jt}/k_{jt})^2 k_{jt}$ units of output.

After production and investment, the firm pays a linear tax rate τ on its revenue y_{jt} net of two deductions. First, the firm deducts its labor costs $w_t n_{jt}$. Second, it deducts capital depreciation costs according to the following geometric schedule. The firm enters the period with a stock of depreciation allowances d_{jt} , of which it writes off $\hat{\delta} d_{jt}$ from its tax bill. The firm also writes off the same fraction $\hat{\delta}$ of new investment i_{jt} from its tax bill. The remaining portion is then carried into the next period, so that $d_{jt+1} = (1 - \hat{\delta})(d_{jt} + i_{jt})$.¹³ In total, the tax bill in period t is

$$\tau(y_{jt} - w_t n_{jt} - \hat{\delta}(d_{jt} + i_{jt})).$$

Households.—There is a representative household with preferences represented by the expected utility function

$$E \sum_{t=0}^{\infty} \beta^t \log \left(C_t - \chi \frac{N_t^{1+\eta}}{1+\eta} - X_t \right),$$

where C_t is consumption, N_t is labor supplied to the market, and X_t is habit stock. The habit stock X_t is

$$(5) \quad X_t = \lambda \hat{C}_{t-1},$$

¹²Khan and Thomas (2008) allow for any investment in $[-ak_{jt}, ak_{jt}]$ to be free of the fixed costs. I set $a = 0$ for simplicity, but results are robust to allowing for empirically reasonable values of a .

¹³In reality, the US tax code follows an annual straight line depreciation schedule with a half-year purchase convention rather than this simple quarterly geometric schedule. Section IIB shows that only the present value of this schedule per unit of investment affects firm's decisions in my model. See Xu and Zwick (2020) for an analysis of the implications of the details of the tax code using a model of firm-level tax management.

where $\hat{C}_t = C_t - \chi(N_t^{1+\eta}/(1+\eta))$ is consumption net of the disutility of work and λ controls the sensitivity of habit with respect to \hat{C}_t .¹⁴ I assume that the household does not take into account the fact that their current choices impact the future habit stock X_{t+1} . The total time endowment per period is 1, so that $N_t \in [0, 1]$. The household owns all firms in the economy and markets are complete.

Government.—The government collects the corporate tax and transfers the proceeds lump sum to the household. In period t , this transfer is

$$(6) \quad T_t = \tau(Y_t - w_t N_t - \hat{\delta}(D_t + I_t)),$$

where Y_t is aggregate output, N_t the aggregate labor input, D_t the aggregate stock of depreciation allowances, and I_t is aggregate investment.

B. Firm Optimization

I characterize the firm's optimization problem recursively. The firm's individual state variables are ε_{jt} , its current draw of the idiosyncratic productivity shock; k_{jt} , its preexisting stock of capital; d_{jt} , its preexisting stock of depreciation allowances; and ξ_{jt} , its current draw of the fixed cost. I denote the aggregate state vector \mathbf{s}_t and postpone discussion of its elements until I define the recursive competitive equilibrium in Section IID.

The firm's value function $v(\varepsilon, k, d, \xi; \mathbf{s})$ solves the Bellman equation,

$$(7) \quad v(\varepsilon, k, d, \xi; \mathbf{s}) = \tau \hat{\delta} d + \max_n \left\{ (1 - \tau)(z\varepsilon k^\theta n^\nu - w(\mathbf{s})n) \right\} \\ + \max \left\{ v^a(\varepsilon, k, d; \mathbf{s}) - \xi w(\mathbf{s}), v^n(\varepsilon, k, d; \mathbf{s}) \right\}.$$

The first max operator represents the optimal choice of labor and the second max operator represents the extensive margin choice of investment. These two choices are independent because the choice of labor is a purely static problem.

If the firm chooses to pay the fixed cost, then it achieves the choice-specific value function $v^a(\varepsilon, k, d; \mathbf{s})$ defined by the Bellman equation,

$$(8) \quad v^a(\varepsilon, k, d, \mathbf{s}) = \max_i \left\{ (1 - \tau \hat{\delta})i - \frac{\varphi}{2} \left(\frac{i}{k} \right)^2 k + E \left[\Lambda(z'; \mathbf{s}) v(\varepsilon', k', d', \xi'; \mathbf{s}') \mid \varepsilon, k, d, \mathbf{s} \right] \right\},$$

subject to

$$k' = (1 - \delta)k + i \quad \text{and} \quad d' = (1 - \hat{\delta})(d + i),$$

where $\Lambda(z'; \mathbf{s})$ is the stochastic discount factor. I denote the implied "target" capital stock $k^a(\varepsilon, k, d; \mathbf{s}) = (1 - \delta)k + i^a(\varepsilon, k, d; \mathbf{s})$.

¹⁴I assume this particular form of preferences for two reasons. First, following Greenwood, Hercowitz, and Huffman (1988), they eliminate the wealth effect on labor supply and allow the model to generate procyclical hours worked with a countercyclical real interest rate. With standard KPR preferences, the fact that the real interest rate falls in expansions potentially induces households to intertemporally substitute future leisure for current leisure in expansions, leading to a fall in hours worked. Second, assuming habit formation over the consumption bundle \hat{C}_t simplifies the analysis of the stochastic discount factor. The main results of the paper also hold if habit is defined over actual consumption C_t only.

If the firm chooses not to pay its fixed cost then it achieves the choice-specific value function $v^n(\varepsilon, k, d; \mathbf{s})$ defined by the Bellman equation,

$$(9) \quad v^n(\varepsilon, k, d; \mathbf{s}) = E[\Lambda(z'; \mathbf{s})v(\varepsilon', k', d', \xi'; \mathbf{s}') | \varepsilon, k, d, \mathbf{s}],$$

subject to

$$k' = (1 - \delta)k \quad \text{and} \quad d' = (1 - \hat{\delta})d.$$

The only difference from the unconstrained Bellman equation (8) is that investment is constrained to be $i = 0$. I call the implied “constrained” capital stock $k^n(\varepsilon, k, d; \mathbf{s}) = (1 - \delta)k$.

The firm will choose to pay the fixed cost if and only if the value from doing so is higher than the value of not paying the fixed cost, i.e., if and only if $v^a(\varepsilon, k, d; \mathbf{s}) - \xi w(\mathbf{s}) \geq v^n(\varepsilon, k, d; \mathbf{s})$. For each tuple $(\varepsilon, k, d; \mathbf{s})$, there is a unique threshold $\hat{\xi}(\varepsilon, k, d; \mathbf{s})$ which makes the firm indifferent between these two options:

$$(10) \quad \hat{\xi}(\varepsilon, k, d; \mathbf{s}) = \frac{v^a(\varepsilon, k, d; \mathbf{s}) - v^n(\varepsilon, k, d; \mathbf{s})}{w(\mathbf{s})}.$$

For draws of the fixed cost ξ below $\hat{\xi}(\varepsilon, k, d; \mathbf{s})$, the firm pays the fixed cost; for draws of the fixed cost above $\hat{\xi}(\varepsilon, k, d; \mathbf{s})$, it does not. This threshold is increasing in the “capital imbalance” $|k^a(\varepsilon, k, d; \mathbf{s}) - k^n(\varepsilon, k, d; \mathbf{s})|$ since the value from adjusting is higher when the target capital stock is further away from the constrained capital stock. The firm will only find it optimal to pay the fixed cost infrequently, generating the lumpy investment patterns observed in the micro data.

It is possible to simplify the problem by eliminating the tax depreciation allowances d from the firm’s state vector. In this model, firms only care about the present value of the depreciation allowances generated by their investment because these allowances enter the firm’s flow profits separately from the other terms. Therefore, the tax depreciation schedule simply decreases the implicit price of new investment by the present value of depreciation allowances which that investment generates. I formalize this logic in Proposition 2.

PROPOSITION 2: *The firm’s value function is of the form $v(\varepsilon, k, d, \xi; \mathbf{s}) = v^1(\varepsilon, k, \xi; \mathbf{s}) + \tau PV(\mathbf{s})d$ where $PV(\mathbf{s})$ is defined by the recursion $PV(\mathbf{s}) = \hat{\delta} + (1 - \hat{\delta})E[\Lambda(z'; \mathbf{s})PV(\mathbf{s}')]$. Furthermore, $v^1(\varepsilon, k, \xi; \mathbf{s})$ is defined by the Bellman equation,*

$$(11) \quad v^1(\varepsilon, k, \xi; \mathbf{s}) = \pi(\varepsilon, k; \mathbf{s}) + \max_i \left\{ -q(\mathbf{s})i - \frac{\varphi}{2} \left(\frac{i}{k} \right)^2 k - \xi w(\mathbf{s}) \mathbf{1}\{i \neq 0\} \right. \\ \left. + E[\Lambda(z'; \mathbf{s})v^1(\varepsilon', (1 - \delta)k + i, \xi'; \mathbf{s}') | \varepsilon, k, \xi, \mathbf{s}] \right\},$$

where $q(\mathbf{s}) = 1 - \tau PV(\mathbf{s})$ is the tax-adjusted relative price of investment.

PROOF:

See online Appendix Section IV.

In the online Appendix, I show that this result can be leveraged to simplify the model's equilibrium following the strategy developed in Khan and Thomas (2003).

C. Household Optimization

Since investment is chosen by firms, there are no dynamic linkages in the household's choices and their decision problem is equivalent to the following static problem state by state:

$$(12) \quad \max_{C,N} \log \left(C - \chi \frac{N^{1+\eta}}{1+\eta} - X(\mathbf{s}) \right), \quad \text{subject to} \quad C \leq w(\mathbf{s})N + \Pi(\mathbf{s}) + T(\mathbf{s}),$$

where $\Pi(\mathbf{s})$ are profits from the firms and $T(\mathbf{s})$ is government transfers. Markets are complete with respect to aggregate risk, so the stochastic discount factor used by firms is equal to the household's intertemporal marginal rate of substitution state by state:

$$(13) \quad \Lambda(z'; \mathbf{s}) = \beta \frac{\hat{C}(\mathbf{s}) - X(\mathbf{s})}{\hat{C}(\mathbf{s}') - X(\mathbf{s}')}.$$

D. Definition of Equilibrium

The aggregate state vector is $\mathbf{s} = (z, X, \mu)$, where z is the aggregate productivity shock, X is the household's habit stock, and μ is the distribution of firms over their individual state vector (ε, k, ξ, d) .

DEFINITION 1: A Recursive Competitive Equilibrium for this economy is a list of functions $v(\varepsilon, k, d, \xi; \mathbf{s})$, $n(\varepsilon, k; \mathbf{s})$, $i^a(\varepsilon, k; \mathbf{s})$, $\hat{\xi}(\varepsilon, k; \mathbf{s})$, $C(\mathbf{s})$, $N(\mathbf{s})$, $T(\mathbf{s})$, $w(\mathbf{s})$, $\Pi(\mathbf{s})$, $\Lambda(z'; \mathbf{s})$, $X'(\mathbf{s})$, and $\mu'(\mathbf{s})$, such that

- (i) (Household Optimization) Taking $w(\mathbf{s})$, $\Pi(\mathbf{s})$, and $T(\mathbf{s})$ as given, $C(\mathbf{s})$ and $N(\mathbf{s})$ solve the utility maximization problem (12).
- (ii) (Firm Optimization) Taking $w(\mathbf{s})$, $\Lambda(z'; \mathbf{s})$, $X'(\mathbf{s})$, and $\mu'(\mathbf{s})$ as given, $v(\varepsilon, k, d, \xi; \mathbf{s})$, $n(\varepsilon, k; \mathbf{s})$, $i^a(\varepsilon, k; \mathbf{s})$, and $\hat{\xi}(\varepsilon, k; \mathbf{s})$ solve the firm's maximization problem (7)–(10).
- (iii) (Government) For all \mathbf{s} , $T(\mathbf{s})$ is given by (6).
- (iv) (Consistency) For all \mathbf{s} ,

$$(a) \quad \Pi(\mathbf{s}) = \int \left[(1 - \tau)(z\varepsilon k^\theta n(\varepsilon, k; \mathbf{s})^\nu - w(\mathbf{s})n(\varepsilon, k; \mathbf{s})) + \tau\hat{\delta}d \right. \\ \left. - (1 - \tau\hat{\delta})i(\varepsilon, k, \xi; \mathbf{s}) - \frac{\varphi}{2} \left(\frac{i(\varepsilon, k, \xi; \mathbf{s})}{k} \right)^2 k \right. \\ \left. - \xi w(\mathbf{s}) \mathbf{1} \left\{ \frac{i(\varepsilon, k, \xi; \mathbf{s})}{k} \neq 0 \right\} \right] \mu(d\varepsilon, dk, dd, d\xi),$$

where $i(\varepsilon, k, \xi; \mathbf{s}) = i^a(\varepsilon, k, \xi; \mathbf{s})$ if $\xi \leq \hat{\xi}(\varepsilon, k; \mathbf{s})$ and $i(\varepsilon, k, \xi; \mathbf{s}) = 0$ otherwise.

(b) $\Lambda(z'; \mathbf{s})$ is given by (13).

(c) $X'(\mathbf{s})$ follows (5).

(d) For all measurable sets $\Delta_\varepsilon \times \Delta_k \times \Delta_d \times \Delta_\xi$, $\mu'(\Delta_\varepsilon \times \Delta_k \times \Delta_d \times \Delta_\xi) = \int p(\varepsilon' \in \Delta_\varepsilon | \varepsilon) d\varepsilon' \times \mathbf{1}\{i(\varepsilon, k, \xi; \mathbf{s}) + (1 - \delta)k \in \Delta_k\} \times \mathbf{1}\{(1 - \hat{\delta}) \times (i(\varepsilon, k, \xi; \mathbf{s}) + d) \in \Delta_d\} \times G(\Delta_\xi) \times \mu(d\varepsilon, dk, dd, d\xi)$, where $G(\xi)$ is the c.d.f. of ξ and $p(\varepsilon' | \varepsilon)$ is the p.d.f. of idiosyncratic shocks ε .

(v) (Market Clearing) For all \mathbf{s} , $N(\mathbf{s}) = \int \left(n(\varepsilon, k, \xi; \mathbf{s}) + \left(\hat{\xi}(\varepsilon, k, d; \mathbf{s})^2 / (2\bar{\xi}) \right) \right) \times \mu(d\varepsilon, dk, dd, d\xi)$.

E. Solution Method

The key challenge to solving the model is that the aggregate state vector \mathbf{s} contains the cross-sectional distribution of firms, which is an infinite-dimensional object. I overcome this challenge using the computational method concurrently developed in Winberry (2018) (which itself builds on Campbell 1998; Reiter 2009; and Algan, Allais, and Den Haan 2008). The method approximates the distribution at any point in time using a flexible but finite-dimensional parametric family; the parameters of that family are then endogenous aggregate state variables of the approximated model. A good approximation of the distribution requires 5–10 endogenous parameters, leaving globally accurate approximation methods infeasible due to the curse of dimensionality. Therefore, I solve for the aggregate dynamics of the model using a second-order perturbation. See online Appendix Section V for details of the implementation.

I have found that Winberry's (2018) method has two advantages over the usual approach of approximating the distribution with moments, as in Krusell and Smith (1998). First, forecasts of key variables based only on the aggregate capital stock are inaccurate (which I show in online Appendix Section V). This fact indicates that higher-order features of the distribution are relevant in determining aggregate dynamics. Second, the method is computationally efficient due to the use of perturbation methods with respect to the aggregate state vector. A local approximation with respect to the aggregate state is appropriate because aggregate shocks are small relative to idiosyncratic shocks. The main advantage of using Krusell and Smith's (1998) method would be that the solution is globally accurate with respect to aggregate shocks.

III. Model Parameterization and Validation

I parameterize the model to jointly match the behavior of investment at the micro level and the dynamics of the real interest rate at the macro level.

TABLE 3—FIXED PARAMETER VALUES

Parameter	Description	Value
β	Discount factor	0.99
η	Inverse Frisch	0.5
θ	Labor share	0.64
ν	Capital share	0.21
δ	Capital depreciation	0.025
ρ_z	Aggregate TFP AR(1)	0.97
σ_z	Aggregate TFP AR(1)	0.0078
τ	Tax rate	0.35
$\hat{\delta}$	Tax depreciation	0.119

Note: Parameters fixed in calibration.

A. Parameterization

I parameterize the model in two steps. First, I fix a set of parameters to match standard macroeconomic targets in steady state. Second, given the values of those parameters, I choose the remaining parameters to match targets in the data. A model period corresponds to one quarter.

Fixed Parameters.—Table 3 lists the parameters that I fix. I set the discount factor $\beta = 0.99$. I set the Frisch elasticity of labor supply to 2, within the range of macro elasticities identified by Chetty et al. (2011). I set the labor share $\theta = 0.64$ and choose the capital share so that the total returns to scale is 85 percent. The returns to scale lies within the range considered in the literature, from 60 percent in Gourio and Kashyap (2007) to 92 percent in Khan and Thomas (2008). I set $\delta = 0.025$ so that the steady state aggregate investment rate is 10 percent, in line with the average in the postwar data. I set the stochastic process for TFP to $\rho_z = 0.97$ and $\sigma_z = 0.0078$, the values of the empirical process (which was used to estimate the impulse response of the real interest rate to TFP shocks in Section IB).

I set the tax rate $\tau = 0.35$ to match the top marginal income tax rate in the US federal corporate income tax code over most of the sample period. Given this value of the tax rate, I choose the slope of the tax depreciation schedule $\hat{\delta}$ to match the average present value of tax depreciation allowances per unit of investment from Zwick and Mahon (2017). Proposition 2 shows that this present value summarizes how the tax depreciation schedule affects firms' decisions in my model.

Fitted Parameters.—I choose the remaining parameters, listed in Table 5, in order to match the empirical targets in Table 4.¹⁵ The micro-level investment targets are computed from annual IRS corporate income tax returns, reported in Zwick and Mahon (2017).¹⁶ The IRS sample features significant micro-level

¹⁵ I exogenously fix the persistence of the productivity shocks $\rho_\varepsilon = 0.9$. Clementi and Palazzo (2015) show the persistence parameter is only weakly identified when using investment data alone.

¹⁶ Much of the literature with firm heterogeneity and investment calibrates models to match investment behavior from the Census of manufacturing firms, reported by Doms and Dunne (1998) or Cooper and Haltiwanger (2006). Zwick and Mahon's (2017) data have three important advantages over Census data in the context of this paper. First, they cover all sectors of the economy rather than just manufacturing, and therefore allow for a more representative

TABLE 4—EMPIRICAL TARGETS

Target	Data	Model
<i>Micro investment (annual)</i>		
Average investment rate (percent)	10.4	10.8
Standard deviation of investment rates	0.16	0.14
Spike rate (percent)	14.4	18.9
Positive investment rates (percent)	85.6	81.1
<i>Interest rate dynamics (quarterly)</i>		
Cumulative impulse response	−0.41	−0.29
$\sigma(I_t)/\sigma(Y_t)$	2.87	2.73

Notes: Micro investment moments from annual firm-level IRS data, 1998–2010, as reported in Zwick and Mahon (2017, online Appendix Table B.1). Statistics drawn from distribution of investment rates pooled over firms and time. Spike rate is fraction of observations with investment rate greater than 20 percent. Positive investment is fraction of observations less than 20 percent. *Cumulative impulse response* refers to the cumulated response of the ex ante real interest rate to a TFP shock over the first year, identified from the empirical VAR (4). $\sigma(I_t)/\sigma(Y_t)$ is the standard deviation of HP-filtered aggregate investment relative to the standard deviation of HP-filtered aggregate output 1947:I–2016:IV.

lumpiness, in line with previous findings in Census data (see, for example, Cooper and Haltiwanger 2006). About one-fourth of firm-years in the sample feature essentially zero investment while simultaneously one-sixth of firm-years have investment rate spikes greater than 20 percent.

I target two statistics related to the dynamics of the real interest rate. First, I target the one-year cumulative response of the expected real interest rate to a TFP shock identified using the VAR (4).¹⁷ Second, I target the volatility of aggregate investment relative to the volatility of aggregate output. As I discuss in Section IIIB, increasing the strength of either habit formation or adjustment costs has similar implications for the dynamics of the real interest rate. However, they have opposite implications for the volatility of investment; adjustment costs make investment less volatile while habit formation makes it more volatile. Targeting the volatility of investment therefore places discipline on the relative strength of these two forces.¹⁸

Although the model is nonlinear and over-identified, with six moments determining four parameters, it nonetheless fits the targets in Table 4 fairly well.¹⁹ The model captures the frequency of spikes relative to the frequency of non-spike observations,

sample of the economy than previous studies. Second, they cover a more recent sample period (1998–2010) than the Census data (1972–1988). Third, the IRS data are at the firm level, which is the appropriate unit of analysis for studying tax policy in Section V.

However, Zwick and Mahon's (2017) data also have two disadvantages relative to the Census data. First, the tax data only record investment expenditures, while the Census data also record retirement and sales of capital. Second, measured investment in the Zwick and Mahon (2017) sample mainly includes equipment goods while the measured capital includes both equipment and structures. On net, I prefer to work with the Zwick and Mahon (2017) data because of its greater sample coverage across sectors and time. I conjecture that my results are robust to calibrating the model to the Census data since both datasets indicate substantial lumpiness of investment in terms of the frequency of inaction and the frequency of spikes.

¹⁷The impulse response of the expected real interest rate is plotted in online Appendix Figure 1.

¹⁸Habit formation has also been shown to improve the empirical performance of DSGE models along two other dimensions. First, habit improves the internal propagation of shocks onto aggregate consumption; for example, Christiano, Eichenbaum, and Evans (2005) show that habit formation allows their model to match the hump-shaped response of consumption to a monetary policy shock. Second, habit helps match various features of asset prices (see, for example, Boldrin, Christiano, and Fisher 2001).

¹⁹I time-aggregate the model's quarterly observations to the annual frequency by summing over flow investment done within a year and using the end-of-year capital stock.

TABLE 5—FITTED PARAMETER VALUES

Parameter	Description	Value
<i>Micro heterogeneity</i>		
$\bar{\xi}$	Upper bound on fixed costs	0.703
φ	Convex adjustment cost	2.950
ρ_ε	Idiosyncratic productivity AR(1) (fixed)	0.900
σ_ε	Idiosyncratic productivity AR(1)	0.053
<i>Habit formation</i>		
λ	Sensitivity of habit with respect to consumption bundle	0.746

Notes: Parameters chosen to match moments in Table 4. I exogenously fix the persistence of idiosyncratic productivity following the discussion in footnote 15.

which is informative about the strength of fixed costs. The model also captures the dispersion of investment rates across firms, which is informative about the size of idiosyncratic shocks and strength of the convex adjustment costs.²⁰ While the model matches the negative response of the real interest rate to a TFP shock, it only captures around two-thirds of the overall decline. This failure is primarily due to the hump-shaped nature of the empirical response, which the simple habit formation process (5) does not generate.

Table 5 shows that the calibrated parameter values are broadly comparable to previous findings in the literature. The upper bound on the fixed cost $\bar{\xi}$ is within the (admittedly wide) range of 0.0083 in Khan and Thomas (2008) and 4.4 in Bachmann, Caballero, and Engel (2013). The calibrated value implies that the average fixed cost paid conditional on adjusting is 9.3 percent of firms' average quarterly output. The dispersion of idiosyncratic TFP shocks is in line with direct measures surveyed in Syverson (2011); for example, the 90-10 ratio of log productivity is 0.31 in my model versus 0.65 in the data. The average size of the habit stock is 75 percent of the households consumption bundle \hat{C}_t , close to the 65 percent in Christiano, Eichenbaum, and Evans (2005).

B. Identification

The identification of these parameters can be understood in two broad steps. First, the dynamics of the real interest rate pin down the overall strength of habit formation and adjustment costs. This occurs because the model's equilibrium real interest rate is determined by the interaction of investment supply, which is influenced by the degree of habit formation, and investment demand, which is influenced by the degree of adjustment costs. Second, given the overall strength of adjustment costs, the micro-level investment data pin down the split between fixed and convex adjustment costs and the dispersion of idiosyncratic shocks.

²⁰ A natural target which I omit is the occurrence of investment inaction, often defined as the frequency of observations with investment rates less than 1 percent annually. I do not target inaction for two reasons. First, as discussed by Cooper and Haltiwanger (2006), the precise definition of inaction in the data is rather arbitrary given heterogeneity in depreciation rates or in the types of investment episodes (e.g., maintenance versus large new projects). Second, since the IRS data only report investment expenditure, observed inaction may reflect firms which do not purchase capital but which nonetheless sell or retire capital. It is straightforward to allow for some degree of inaction in the model, as in Khan and Thomas (2008).

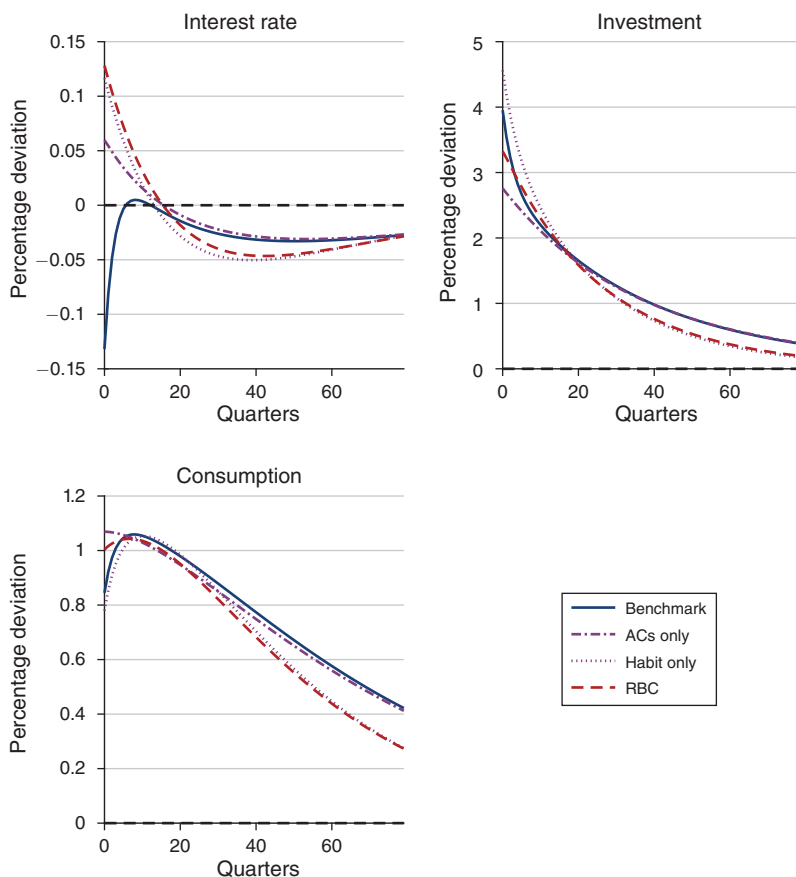


FIGURE 3. IDENTIFICATION OF HABIT FORMATION AND ADJUSTMENT COSTS

Notes: Impulse responses to a TFP shock in different parameterizations of the model. *Benchmark* refers to calibrated model. *ACs only* refers to model in which habit formation has been turned off (by setting $\lambda = 0$). *Habit only* refers to model in which adjustment costs have been turned off (by setting ξ and $\varphi =$ to low values). *No habit or ACs* refers to model in which both habit formation and adjustment costs have been turned off. In this case, the firm side of the model approximately aggregates to a representative firm.

In order to more formally understand the first step of this process, Figure 3 plots the impulse response of key model variables under four different parameter configurations: no habit formation or adjustment costs; only habit formation; only adjustment costs; and the calibrated model with both habit formation and adjustment costs.

First consider the cases without habit formation. The household's Euler equation is

$$1 + r_t = \frac{1}{\beta} E_t \left[\frac{\hat{C}_t}{\hat{C}_{t+1}} \right]^{-1},$$

which relates the real interest rate to expected consumption growth.²¹ While a positive TFP shock unambiguously increases both consumption C_t and investment I_t , its effect on consumption growth C_{t+1}/C_t , and therefore the real interest rate, depends on the responsiveness of investment because that determines the size of the capital stock in period $t + 1$. Without adjustment costs, the capital stock K_{t+1} rises by enough to allow C_{t+1} to rise relative to C_t , therefore causing the real interest rate to rise. Adjustment costs dampen the response of investment, which dampens the rise of consumption growth and therefore dampens the rise of the real interest rate. However, adjustment costs alone are not quantitatively strong enough to fully account for the response of the real interest rate in the data.

Habit formation further brings the model in line with the data by breaking the tight link between consumption growth and the real interest rate. With habit formation, the household's Euler equation is

$$1 + r_t = \frac{1}{\beta} E_t \left[\frac{\hat{C}_t - X_t}{\hat{C}_{t+1} - X_{t+1}} \right]^{-1},$$

where X_t is the stock of habit at time t . This equation implies that the real interest rate will fall following a positive TFP shock if the growth in the habit-adjusted consumption bundle $\hat{C}_t - X_t$ falls. And indeed, holding the path of the consumption bundle \hat{C}_t fixed, stronger habit formation decreases the response of habit-adjusted consumption growth; with strong enough habit formation, $\hat{C}_t - X_t$ increases by more than $\hat{C}_{t+1} - X_{t+1}$ because X_t is predetermined during the period of the shock while X_{t+1} is not. However, stronger habit formation also increases the incentive to smooth the consumption bundle \hat{C}_t , which may undo some of its effect on the real interest rate. In fact, Figure 3 shows that, without adjustment costs, this consumption smoothing undoes nearly all of the effect of habit formation on the real interest rate. The increased desire to smooth consumption implies that consumption is less responsive to the shock and therefore investment is more responsive. Since there are no adjustment costs, only a small change in the path of the real interest rate is required to induce firms to increase their investment.

It is only when habit formation is combined with adjustment costs, which impede the household's ability to smooth consumption over time, that the real interest rate falls in response to a TFP shock. In this case, a large decline in the real interest rate path is required to induce firms to increase their investment and therefore partially accommodate the household's desire for smooth consumption.²²

²¹ Due to Greenwood, Hercowitz, and Huffman (1988) preferences, the Euler equation is technically in terms of the consumption bundle $\hat{C}_t = C_t - \chi(N_t^{1+\eta}/(1+\eta))$ rather than consumption C_t itself. Quantitatively, the dynamics of this consumption bundle are similar to consumption itself.

²² This discussion makes clear that the response of the real interest rate to a productivity shock is informative about the degree of capital adjustment frictions and, therefore, the interest-sensitivity of investment. However, matching this impulse response does not imply realistic unconditional dynamics of the real interest rate; its volatility in the model is 0.15 percent (compared to 1.98 percent in the data) and its correlation with TFP is -0.66 (compared to -0.20 in the data). Adding more shocks would increase the volatility of the real interest rate and, therefore, require an even lower interest-sensitivity of investment to rationalize the observed volatility of investment in the data.

C. Relationship between the Real Interest Rate and the SDF

The dynamics of the risk-free rate that I target in my calibration are tightly related to the stochastic discount factor $\Lambda(z'; \mathbf{s})$ which firms use to value the benefits of investment. Note that the expected present value of the firm can be decomposed as

$$(14) \quad E[\Lambda(z'; \mathbf{s})v(\varepsilon', k', \xi'; \mathbf{s}') | \varepsilon, k, \mathbf{s}] = \frac{1}{1 + r(\mathbf{s})} E[v(\varepsilon', k', \xi'; \mathbf{s}') | \varepsilon, k, \mathbf{s}] \\ + \text{cov}(\Lambda(z'; \mathbf{s}), v(\varepsilon', k', \xi'; \mathbf{s}') | \varepsilon, k, \mathbf{s}),$$

where $r(\mathbf{s}) = (1/E[\Lambda(z'; \mathbf{s}) | \mathbf{s}]) - 1$ is the risk-free rate. The first term in the decomposition (14) captures the risk-free discounting of the value function, i.e., the implications of the stochastic discount factor (SDF) for intertemporal comparisons. The second term captures the covariance between the SDF and the value function, i.e., the implications of the SDF for risk. Hence, my calibration strategy directly targets the intertemporal component of the SDF and places no direct discipline on the risk component.²³

Capturing the dynamics of the risk component is outside the scope of this paper for two reasons. First, in order to capture movements in the expected risk premium, which are informative about the risk component, the macro asset-pricing literature often appeals to stochastic volatility or a time-varying market price of risk, both of which are not in the model. Second, in the data, the covariance between the excess return to equity and aggregate TFP is small and does not significantly fluctuate over time. To the extent that returns to equity reflect the returns to capital, these facts suggest that the scope for time-variation in risk component in response to TFP shocks is small.²⁴

D. Model Validation

Before presenting the main results of the paper, I show that the model performs well along dimensions that were not targeted in the calibration.

Micro Investment Behavior.—The stationary distribution of realized investment rates across firms is broadly comparable to the empirical distribution reported in Zwick and Mahon (2017). Figure 4 plots the histogram of investment rates in the model's stationary distribution. Due to the fixed cost, there is a large mass of observations with zero investment together with a large mass of observations with large positive investment spikes. Overall, the distribution is highly non-normal; it features both positive skewness (1.55, compared to 3.60 in the data) and excess kurtosis (6.77, compared to 17.6 in the data). Investment rates are slightly negatively autocorrelated in the model (−0.021) because large investment spikes are often followed by zero investment.

²³Note that the empirical analysis in Section IB ignores changes in the inflation risk premium, which are likely to be small on a quarter-to-quarter basis.

²⁴My model also does not match the average level of the equity risk premium in the data. Matching the level of the risk premium would increase the average level of the covariance term in (14), but not directly affect how the term responds to business cycle shocks.




FIGURE 4. DISTRIBUTION OF ANNUALIZED INVESTMENT RATES IN STEADY STATE

Notes: Histogram of investment rates in the model's steady state. Investment rates are time-aggregated to the annual level in order to compare to the data. The distribution features positive skewness (1.55 in the model compared to 3.60 in the data) and excess kurtosis (6.77 in the model compared to 17.60 in the data).

TABLE 6—UNCONDITIONAL BUSINESS CYCLE STATISTICS

Volatility			Cyclicalities		
Statistic	Data	Model	Statistic	Data	Model
$\sigma(Y_t)$ (percent)	1.61	1.77			
$\sigma(C_t)/\sigma(Y_t)$	0.50	0.71	$\rho(C_t, Y_t)$	0.79	0.99
$\sigma(I_t)/\sigma(Y_t)$	2.87	2.73	$\rho(I_t, Y_t)$	0.77	0.97
$\sigma(N_t)/\sigma(Y_t)$	1.15	0.71	$\rho(N_t, Y_t)$	0.86	1.00

Notes: All series have been logged and HP-filtered with smoothing parameter 1,600. *Data* refers to the aggregate US data, 1947:I–2016:IV, described in online Appendix Section II.A. *Model* refers to calibrated model.

Unconditional Business Cycle Statistics.—Table 6 shows that the model matches standard business cycle statistics roughly as well as the benchmark RBC model, despite the fact that the model has much stronger habit formation and adjustment costs. Investment is more volatile than output and consumption is less volatile than output in both the model and the data. The volatility of hours is lower in the model than in the data, which is a well-known defect of the benchmark RBC model as well. Finally, all macroeconomic aggregates are highly correlated with each other due to the fact that there is a single aggregate shock.

IV. Procyclical Responsiveness to Shocks

I now use the calibrated model to quantitatively analyze the implications of fixed costs for the dynamics of aggregate investment. Section IVA shows that fixed costs imply that aggregate investment is more responsive to shocks in expansions than in

recessions. Section IVB shows that quantitatively matching the dynamics of the real interest rate is key to generating this procyclical responsiveness in general equilibrium.

A. Role of Fixed Costs

I begin by describing how fixed adjustment costs generate the procyclical responsiveness of aggregate investment to an TFP shock. In order to isolate the role of firms' behavior in driving aggregate dynamics, I perform this analysis in "partial equilibrium," i.e., by aggregating firms' decision rules with prices held fixed at their steady state values. This analysis provides a natural benchmark against which to compare the general equilibrium results in Section IVB.

Procyclical Impulse Responses.—Figure 5 compares the response of aggregate investment to a TFP shock starting from two different points in the business cycle. The first starting point is an example recession generated by a history of -2.5 standard deviation shocks over the previous year, which generates 25 percent decline in annual investment in the full general equilibrium model and is comparable to the observed decline during the 2008 recession. The second point is an expansion generated by a history of $+2.5$ standard deviation positive shocks over the previous year, which generates a similarly sized increase in aggregate investment and is comparable to, e.g., the mid-1980s boom. I compute the response of aggregate investment to a $+1$ standard deviation shock starting from these two points. Since the model is nonlinear, I replicate this procedure over many simulations of the model and then take the average of these responses.

The left panel of Figure 5 shows the same shock generates more than double the amount of investment starting from the expansion than starting from the recession. The cumulative differences over time are smaller than the impact difference due to intertemporal substitution; starting from the expansion, some firms pull forward investment they would have done in the future.

Role of Fixed Costs.—The state dependence in these impulse responses is induced by more firms making an extensive margin investment starting from the expansion. In order to understand this result, first note that in steady state the average firm holds less capital than its target stock $k^a(\varepsilon, k; s)$ because of capital depreciation and convex adjustment costs (see footnote 26). Now consider a history of negative shocks which generates a recession. Since the negative shocks decrease the marginal product of capital, they decrease the target capital stock and therefore bring the average firm closer to its target. In this case, the probability of a firm paying its fixed cost, which is proportional to the adjustment threshold $\hat{\xi}(\varepsilon, k)$, falls.²⁵ Furthermore, additional shocks will have a relatively small effect on the adjustment probability $\hat{\xi}(\varepsilon, k)$.

On the other hand, a history of positive shocks will move the average firm even further below its target, i.e., induce $k^n(\varepsilon, k; s) \ll k^a(\varepsilon, k; s)$, and increase the adjustment probability $\hat{\xi}(\varepsilon, k; s)$. In this region of the state space, it turns out that further changes in $k^a(\varepsilon, k; s)$ have larger effects on the adjustment probability $\hat{\xi}(\varepsilon, k; s)$.

²⁵ Since firms' draws of the fixed cost ξ are i.i.d., for each value of productivity and capital (ε, k) a fraction $\hat{\xi}(\varepsilon, k; s)$ of firms will adjust while the remaining fraction will not.




FIGURE 5. PROCYCLICAL IMPULSE RESPONSES OF AGGREGATE INVESTMENT

Notes: Left panel plots the impulse responses to a one standard deviation positive shock to aggregate TFP computing following Koop, Pesaran, and Potter (1996). *Expansion* refers to a history of four 2.5 standard deviation shocks and *Recession* refers to history of four -2.5 standard deviation negative shocks. *Partial equilibrium* refers to aggregating firms' decisions holding prices fixed at their steady state values and *General equilibrium* refers to the full general equilibrium model. Since the model is nonlinear, I compute the impulse responses by (i) drawing a random series of aggregate shocks, (ii) adding the history of shocks to generate an expansion and recession, (iii) computing the difference between the simulation in which there is the additional shock and the original simulation, and (iv) repeating this procedure 200 times and taking the average of all the differences produced in step (iii). Right panel plots how the adjustment probability for firms (conditional on the average realization of idiosyncratic productivity ε) responds to a positive aggregate shock starting from steady state. The blue line (measured against the right axis) plots the steady state distribution of firms over capital k . The red lines (measured against the left axis) plot the probability paying the fixed cost and adjusting capital. The solid red line is in steady state and the dashed line is following a one standard deviation positive TFP shock (with prices held fixed at their steady state values).

Hence, the fact that the adjustment probability $\hat{\xi}(\varepsilon, k; \mathbf{s})$ is increasing in the distance from target $|k^a(\varepsilon, k; \mathbf{s}) - k^n(\varepsilon, k; \mathbf{s})|$ is the key source of procyclical responses to shocks (see Caballero and Engel 2007).

The right panel of Figure 5 plots how the adjustment probabilities of firms respond to a positive productivity shock starting from steady state (conditional on the average realization of idiosyncratic productivity ε). Consistent with this discussion, the adjustment probability is increasing in $|k^a(\varepsilon, k; \mathbf{s}) - k^n(\varepsilon, k; \mathbf{s})|$.²⁶ The positive shock increases the target capital stock for all firms and therefore shifts the adjustment probability function up and to the right.

B. Role of Prices in General Equilibrium

The left panel of Figure 5 shows that the procyclical responses described above survive in general equilibrium. The shock generates 23 percent more investment upon impact starting from the expansion than starting from the expansion than starting from the recession. The procyclicality of the impact response is lower in general equilibrium for two reasons. First, the real wage w_t is procyclical, which decreases the marginal revenue product of capital in response to the shock. Second, as Figure 6 shows, the real interest rate r_t falls by more in the recession than in the

²⁶The adjustment probability function is positive throughout the distribution of firms due to the existence of convex adjustment costs φ . Convex adjustment costs imply that, for high enough levels of capital k , the target capital stock is a decreasing function of current capital since capital decreases the marginal adjustment cost.




FIGURE 6. COUNTERCYCLICAL RESPONSIVENESS OF THE REAL INTEREST RATE

Notes: Impulse response of the real interest rate to a one standard deviation positive shock to aggregate TFP computed following Koop, Pesaran, and Potter (1996). *Expansion* refers to a history of four +2.5 standard deviation shocks and *Recession* refers to history of four −2.5 standard deviation negative shocks. Since the model is nonlinear, I compute the impulse responses by (i) drawing a random series of aggregate shocks, (ii) adding the history of shocks to generate an expansion and recession, (iii) computing the difference between the simulation in which there is the additional shock and the original simulation, and (iv) repeating this procedure 200 times and taking the average of all the differences produced in step (iii).

expansion, reflecting the fact that marginal utility growth falls by more starting from the recession. However, the procyclicality of the cumulative response is actually higher in general equilibrium than in partial equilibrium. This amplification is due to the fact that persistently low real interest rates following the shock weaken the intertemporal substitution motive for firms to pull forward their investment into the period of the shock.

In order to quantify the amount of time-variation in the impulse response function over a long simulation of the model, I follow Caballero and Engel (1999) and Bachmann, Caballero, and Engel (2013) and compute the “impact responsiveness index” RI_t ,

$$(15) \quad RI_t = 100 \times \log \left(\frac{I(z_t + \sigma_z, X_t, \mu_t) - I(z_t, X_t, \mu_t)}{I(\sigma_z, X^*, \mu^*) - I(0, X^*, \mu^*)} \right),$$

where $I(z, X, \mu)$ is aggregate investment given the aggregate state $\mathbf{s} = (z, X, \mu)$.²⁷ This responsiveness index measures the impact effect of a TFP shock at a given point in time, relative to the effect starting from steady state. I also compute the “cumulative responsiveness index” \widehat{RI}_t ,

$$(16) \quad \widehat{RI}_t = 100 \times \log \left(\frac{\hat{I}(z_t + \sigma_z, X_t, \mu_t) - \hat{I}(z_t, X_t, \mu_t)}{\hat{I}(\sigma_z, X^*, \mu^*) - \hat{I}(0, X^*, \mu^*)} \right),$$

²⁷ Bachmann, Caballero, and Engel (2013) use a more general measure that accounts for asymmetries in the response to a positive and negative shock. These asymmetries are small in my model, so I ignore them here for the sake of simplicity.

TABLE 7—FLUCTUATIONS IN RESPONSIVENESS INDEX OVER TIME

	Impact RI_t				Cumulative \widehat{RI}_t	
	95-5 ratio (percent)	90-10 ratio (percent)	$\rho(RI_t, Y_t)$	$\rho(RI_t, adj_t)$	95-5 ratio (percent)	90-10 ratio (percent)
<i>Benchmark calibration (PE interest elasticity $d \log I_t / dr_t = -6.74$)</i>						
Partial equilibrium	58.5	44.9	0.72	0.97	26.1	19.6
General equilibrium	26.3	20.3	0.99	0.65	16.2	12.7
GE dampening	55.0	54.8			37.9	35.2
<i>Khan and Thomas (2008a) Calibration (PE interest elasticity $d \log I_t / dr_t = -1.140.58$)</i>						
Partial equilibrium	1,032	419	0.29	0.62	168	106
General equilibrium	10.0	7.9	0.99	0.72	2.5	2.0
GE dampening	99.0	98.1			97.6	98.1

Notes: Impact responsiveness index RI_t defined in (15) and cumulative responsiveness index \widehat{RI}_t defined in (16) of the main text. adj_t computes the fraction of firms which pay their fixed cost. *Partial equilibrium* refers to aggregating firms' decision rules holding prices fixed at their steady state values. *General equilibrium* refers to full general equilibrium model. *Benchmark calibration* refers to calibrated model. *Khan and Thomas (2008a) calibration* refers to eliminating the convex adjustment costs ($\varphi \rightarrow 0$), reducing the fixed costs ($\bar{\xi} = 0.0083/4$), changing the idiosyncratic shock process ($\rho_\varepsilon = 0.859^{1/4}$ and $\sigma_\varepsilon = 0.022/4$), increasing the returns to scale ($\theta = 0.256$), eliminating the tax code ($\tau = 0$), eliminating habit formation ($\lambda = 0$), and using separable preferences between consumption and labor supply ($\log C_t - \chi(N_t^{1+\eta}/(1+\eta))$). *PE interest elasticity* is the effect of a one-time unexpected change in the real interest rate starting from steady state.

where $\hat{I}(z_t, X_t, \mu_t)$ is the amount of investment generated over the ten years following the shock. This long-run effect is closely related to the cumulative change in the capital stock and, therefore, output and consumption.

Table 7 shows that the model generates a significant amount of procyclical responsiveness over a long simulation. The impact responsiveness index RI_t is 26 percent higher than the fifth percentile; this variation is positively correlated with aggregate output and the amount of firms making an extensive margin investment, consistent with the mechanism described in Section IVA. Similarly, the ninety-fifth percentile of the cumulative responsiveness index \widehat{RI}_t is 16 percent higher than the fifth percentile.

Relationship to Proposition 1.—These quantitative results indicate that the model is far from the limiting case of Proposition 1, in which fixed costs are irrelevant for aggregate dynamics. In that case, small but extremely procyclical movements in the real interest rate are able to bring the dynamics of investment in line with the desires of the representative household because investment is extremely sensitive to changes in interest rates. In my calibrated model, the semi-elasticity of aggregate investment with respect to the real interest rate is only -6.74 , so general equilibrium has a smaller influence of aggregate investment dynamics.²⁸

Table 7 roughly replicates Khan and Thomas's (2008) calibration at the quarterly frequency and compares it to the predictions of my model.²⁹ The partial equilibrium

²⁸I compute the interest elasticity as the effect of an unexpected increase in the real interest rate starting from the nonstochastic steady state. I assume that the increase in the real interest rate is for one period and is not driven by any aggregate shock in the model (i.e., it is an exogenous increase in the interest rate).

²⁹I roughly replicate the Khan and Thomas (2008) calibration by eliminating the convex adjustment costs ($\varphi \rightarrow 0$), reducing the fixed costs ($\bar{\xi} = 0.0083/4$), changing the idiosyncratic shock process ($\rho_\varepsilon = 0.859^{1/4}$ and $\sigma_\varepsilon = 0.022/4$), increasing the returns to scale ($\theta = 0.256$), eliminating the tax code ($\tau = 0$), eliminating

version of this calibration generates enormous variation in both responsiveness indices; the ninety-fifth percentile of the impact responsiveness index RI_t is ten times as large as the fifth percentile, while the ninety-fifth percentile of the cumulative responsiveness index \widehat{RI}_t is four times as large as the fifth percentile. The partial equilibrium variation is so high because, without convex adjustment costs, firms are extremely sensitive to changes in the incentive to invest (as discussed in Section IA).³⁰ However, general equilibrium price movements eliminates roughly 99 percent of this variation in the responsiveness indices RI_t and \widehat{RI}_t . Consistent with the logic of Proposition 1, general equilibrium is powerful in this calibration because investment is extremely interest sensitive; the elasticity of aggregate investment with respect to the real interest rate is approximately $-1,140$. Furthermore, the remaining degree of variation in the responsiveness index RI_t is also present in the version of the model without fixed costs, which aggregates to a representative firm. Hence, in the Khan and Thomas (2008) model fixed costs are essentially irrelevant for aggregate dynamics relative to the representative firm model.³¹

Koby and Wolf (2020) provide a thorough analysis of the role of interest-sensitivity in driving the irrelevance results of the previous literature and argue that quantitative models should target the interest-sensitivity of investment. My calibration disciplines the interest-sensitivity of investment by separately targeting the volatility of investment and the dynamics of the interest rate in the data. The presence of convex adjustment costs φ is key to dampening the interest-sensitivity because it leads to an upward-sloping marginal cost curve for investment. Koby and Wolf's (2020) analysis applies to a much broader class of models with a richer set of adjustment frictions (e.g., financial frictions) and in which the relative price of investment is not tied as directly to the real interest rate as it is in my model. They discipline the price-sensitivity of investment by targeting the response of investment to the Bonus Depreciation Allowance estimated in Zwick and Mahon (2017). While my strategy is less direct than Koby and Wolf (2020), it does not require the additional structure necessary to accommodate Zwick and Mahon's (2017) difference-in-differences empirical specification. In addition, the precise dynamics of the real interest rate are quantitatively relevant in shaping the aggregate implications of micro-level frictions in models that feature a realistic price-sensitivity of investment. For example, as discussed above, the fact that the real interest rate is persistently low following a TFP shock amplifies the variation in the cumulative responsiveness index.

Role of Model Ingredients.—Table 8 decomposes the role of three key model ingredients in driving the procyclical responsiveness to shocks in my model. First, decreasing the size of the fixed costs decreases the variation in both responsiveness

habit formation ($\lambda = 0$), and using separable preferences between consumption and labor supply ($\log C_t - \chi(N_t^{1+\eta}/(1+\eta))$).

³⁰The fact that investment is so responsive to shocks in partial equilibrium also implies that the local approximation may become inaccurate. Therefore, the precise numbers for the partial equilibrium Khan and Thomas (2008) calibration should be interpreted with caution. I set the convex costs $\varphi = 0.005$ to generate some curvature and help stabilize the partial equilibrium dynamics.

³¹In fact, Khan and Thomas (2008) show that the entire distribution of aggregate investment rates is nearly identical in the two models.

TABLE 8—ROLE OF KEY MODEL INGREDIENTS

	PE $d\log I_t/dr_t$	Impact RI_t (percent)		Cumulative \widehat{RI}_t (percent)	
		95-5 ratio	90-10 ratio	95-5 ratio	90-10 ratio
Full model	-6.74	26.3	20.3	16.2	12.7
Smaller fixed costs	-7.70	21.8	16.9	6.2	5.0
Smaller convex costs	-8.38	18.7	14.5	10.5	8.3
Smaller habit	-6.74	25.9	19.8	15.8	12.2

Notes: Impact responsiveness index RI_t defined in (15) and cumulative responsiveness index \widehat{RI}_t defined in (16) of the main text. *Smaller fixed costs* keeps all parameters the same as the full model except decreases the upper bound on the distribution of fixed costs $\bar{\xi}$ by 50 percent. *Smaller convex costs* keeps all parameters the same as the full model except decreases the convex adjustment cost φ by 50 percent. *Smaller habit* keeps all parameters the same as in the full model except decreases the habit formation parameter λ by 50 percent. PE $d\log I_t/dr_t$ is the effect of a one-time unexpected change in the real interest rate in steady state.

indices, which is natural given that the fixed costs are a key source of the state dependence (as described in Section IVA).

Second, decreasing the size of the convex adjustment cost also decreases the variation in the responsiveness indices. This comparative static balances two opposing forces. On the one hand, lower convex costs make the extensive margin more responsive to shocks and, therefore, increases the degree of state dependence in partial equilibrium (see the partial equilibrium Khan and Thomas 2008 calibration in Table 7). On the other hand, lower convex costs also increase the interest-sensitivity of investment and, therefore, decrease the degree of state dependence in general equilibrium. This latter force quantitatively dominates the former force in the calibrated model.

Third, decreasing the strength of habit formation slightly reduces the variation in the responsiveness indices, but the effect is minor. This comparative static also balances two forces. First, weakening habit formation reduces the countercyclical-ity of the real interest rate, which in turn decreases the degree of state dependence. Second, weakening habit formation also decreases the procyclical responsiveness of the real interest rate to shocks shown in Figure 6. To understand this result, first note that the fact that investment is more responsive to shocks in expansions implies that consumption growth is more responsive in expansions as well. In turn, the fact that consumption growth is more responsive to shocks in expansions implies that marginal utility growth, and therefore the real interest rate, is also more responsive in expansions, which dampens the procyclical responsiveness of investment. Quantitatively, the strength of this mechanism is increasing in the strength of habit formation. While it is presumably possible to reverse engineer a functional form for habit formation that would eliminate these nonlinearities (see, for example, Campbell and Cochrane 1999), I do not pursue that approach here given that the ultimate impact on aggregate investment dynamics is small.³²

All that said, it important to emphasize that habit formation plays a crucial role in generating realistic interest rate dynamics and micro-level investment behavior. As Figure 3 shows, eliminating habit formation in my calibration would generate counterfactually procyclical interest rate dynamics. Matching the dynamics of the real

³²Consistent with this discussion, online Appendix Section VI shows that eliminating habit formation altogether further decreases the variation in the responsiveness index, but the effect is again minor.

interest rate with adjustment costs alone would require significantly larger adjustment costs and, therefore, generate counterfactual implications for investment. Hence, habit formation gives the model enough degrees of freedom to jointly match interest rate dynamics and investment behavior.³³

Online Appendix Section VI contains three additional robustness checks on the results in this section. First, it shows that the results also hold with separable preferences over consumption and labor supply. Second, it shows that the results hold when the tax code is eliminated. Third, it shows that the results are robust to changing the returns to scale. Gourio and Kashyap (2007) argue that lowering the returns to scale has a strong effect on the degree of state dependence; these results suggest that, conditional on a given interest-sensitivity of investment, the returns to scale are relatively unimportant in generating state dependence.

V. Implications for Stimulus Policy

In this section, I briefly study two implications of the model for investment stimulus policies. First, as with productivity shocks, the effectiveness of investment stimulus is state dependent and falls in recessions.³⁴ Second, the importance of the extensive margin at the micro level implies that targeting firms by their size can increase cost effectiveness up to 30 percent compared to existing size-independent policies. The goal of this second exercise is to illustrate how taking the micro-level behavior of investment seriously can affect policy design; the exercise abstracts from other considerations, such as financial frictions, which may affect the correlation of firms' size and their probability of adjustment.

I model investment stimulus as an exogenous shock to the tax-adjusted price of capital, $q(s)$, derived in Proposition 2. In this section only, I assume that the relative price is $q(s) = 1 - \tau(PV(s) + \omega)$, where ω is the investment stimulus shock. Online Appendix Section VII shows that the two most common investment stimulus policies in the United States, the investment tax credit and the bonus depreciation allowance, map into different values for the shock ω . I assume that the shock ω follows an AR(1) process:

$$\omega' = \rho_\omega \omega + \varepsilon'_\omega,$$

where $\varepsilon_\omega \sim N(0, \sigma_\omega^2)$. I choose the standard deviation of the shock $\sigma_\omega = 0.035$ to roughly match the size of a 50 percent Bonus Depreciation Allowance and the quarterly autocorrelation $\rho_\omega = 0.91$ to match a half-life of two years.

Figure 7 plots the impulse responses of aggregate investment and consumption to a one standard deviation positive stimulus shock starting from steady state. The shock immediately decreases the relative price of investment $q(s)$, which then increases investment. Since output is fixed upon impact, this higher investment must be met with lower consumption. Note that the stimulus shock is isomorphic to an

³³ Koby and Wolf's (2020) analysis avoids this issue by directly targeting the price-sensitivity of investment using variation in taxes rather than variation in the real interest rate. Indeed, they present a model without habit formation which nonetheless generates state-dependent impulse responses.

³⁴ State dependence in the effect of policy does not necessarily follow from the results in Section IV because policy shocks have different general equilibrium implications than productivity shocks.




FIGURE 7. AVERAGE IMPULSE RESPONSE TO INVESTMENT STIMULUS SHOCK

Note: Impulse response of aggregate investment and aggregate consumption to a one standard deviation positive investment stimulus shock ω .

investment-specific technological shock, which also induces a negative comovement between consumption and investment. Over time, persistently higher investment increases the capital stock, output, and therefore consumption.

State Dependent Effect of Investment Stimulus.—In order to quantify the degree of state dependence in response to the stimulus shock, I construct two policy responsiveness indices in the spirit of Section IV. The first measures the impact effect of the shock:

$$RI_t^\omega = 100 \times \log \left(\frac{I(z_t, \omega_t + \sigma_\omega, X_t, \mu_t) - I(z_t, \omega_t, X_t, \mu_t)}{I(0, \sigma_\omega, X^*, \mu^*) - I(0, 0, X^*, \mu^*)} \right),$$

where $I(z_t, \omega_t, X_t, \mu_t)$ is aggregate investment given the expanded aggregate state vector $\mathbf{s}_t = (z_t, \omega_t, X_t, \mu_t)$. The second index measures the cumulative effect of the shock:

$$\widehat{RI}_t^\omega = 100 \times \log \left(\frac{\hat{I}(z_t, \omega_t + \sigma_\omega, X_t, \mu_t) - \hat{I}(z_t, \omega_t, X_t, \mu_t)}{\hat{I}(0, \sigma_\omega, X^*, \mu^*) - \hat{I}(0, 0, X^*, \mu^*)} \right),$$

where $\hat{I}(z_t, \omega_t, X_t, \mu_t)$ is the cumulative amount of investment starting from aggregate state \mathbf{s}_t and reverting back to study state.

Table 9 shows that both the impact and cumulative effect vary significantly over time; the ninety-fifth percentile of the impact effect is 31 percent higher than the fifth percentile, and the ninety-fifth percentile of the cumulative effect is nearly 17 percent higher than the fifth percentile. Both indices are positively correlated with output, implying that the effectiveness of policy falls in recessions. A linear forecasting model, such as a VAR, would abstract from this state dependence and therefore be biased up in recessions.

TABLE 9—RESPONSIVENESS INDEX FOR INVESTMENT STIMULUS SHOCK

	95-5 ratio (percent)	90-10 ratio (percent)	75-25 ratio (percent)	$\rho(RI_t, \log Y_t)$
Impact RI_t^ω	31.1	24.6	10.5	0.89
Cumulative \widehat{RI}_t^ω	16.7	12.9	6.2	0.89

Notes: Time-variation in the responsiveness index for the stimulus policy shock. Impact responsiveness index RI_t^ω defined in (15) and cumulative responsiveness index \widehat{RI}_t^ω defined in (16) of the main text.

Increasing Cost Effectiveness with Micro-Targeting.—A general issue with investment stimulus policies is that much of their cost is due to subsidizing investment that would have been done even without the policy. Therefore, an important challenge to policymakers is identifying and disregarding this inframarginal investment. A key insight of my model is that, because investment primarily occurs along the extensive margin, most of the wasted subsidy to inframarginal investment is accounted for by subsidizing firms that would have made an extensive margin investment even without the policy. This insight simplifies the problem to identifying these inframarginal firms.

Figure 8 shows that a simple size-dependent implementation of this idea is quantitatively powerful. I now assume that the change in a given firm's relative price of investment $q(\mathbf{s})$ is

$$\omega \times \alpha_1 n(\varepsilon, k; \mathbf{s})^{\alpha_2},$$

where $n(\varepsilon, k; \mathbf{s})$ is the firm's employment and α_2 captures the weight of the policy on large firms. For each value of α_2 , I solve for the value of α_1 that makes the aggregate cost of the policy equal to the aggregate cost under the size-independent policy, i.e., with $\alpha_1 = 1$ and $\alpha_2 = 0$. Figure 8 shows that increasing the weight on large firms α_2 increases the amount of investment generated by the policy up to 30 percent. This occurs because, as in the data, small firms grow faster than the average firm (due to mean reversion in idiosyncratic productivity ε). In order to grow, these firms are more likely to invest, making them more likely to be inframarginal to the policy.³⁵

VI. Conclusion

In this paper, I have argued that accounting for the importance of the extensive margin in micro-level investment decisions matters for our understanding of aggregate investment dynamics because it implies that aggregate investment is more responsive to shocks in expansions than in recessions. Matching the dynamics of

³⁵ Of course, the quantitative effect of this size-dependent policy relies on this particular model of firm growth. Clementi and Palazzo (2016) have used a similar model to study firms' life cycles and argue that it provides a good fit to the data. However, other models may have different implications for the correlation between size and responsiveness to investment stimulus. For example, a model with financial frictions may imply that small firms are more likely to be financially constrained and therefore are more responsive to the policy. The goal of my exercise here is simply to illustrate the magnitude of the cost savings associated with micro-targeting firms along the extensive margin rather than strongly advocate for this particular size-dependent implementation.




FIGURE 8. COST EFFECTIVENESS OF SIZE-DEPENDENT POLICY

Notes: Aggregate investment generated by a size-dependent investment stimulus shock $\omega \times \alpha_1 n(\varepsilon, k; s)^{\alpha_2}$. For each value of α_2 , I solve for the value of α_1 that makes the aggregate cost of the subsidy equal to the aggregate cost with $\alpha_1 = 1, \alpha_2 = 0$.

the real interest rate is key to generating this result; in an otherwise standard RBC model, counterfactual movements in the real interest rate eliminate most of this procyclical responsiveness. More generally, these results show that data on interest rates place strong discipline on the role of general equilibrium in determining the aggregate implications of firm-level investment behavior.

REFERENCES

- Algan, Yann, Olivier Allais, and Wouter J. Den Haan. 2008. "Solving Heterogeneous-Agent Models with Parameterized Cross-Sectional Distributions." *Journal of Economic Dynamics and Control* 32 (3): 875–908.
- Bachmann, Rüdiger, Ricardo J. Caballero, and Eduardo M. R. A. Engel. 2013. "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model." *American Economic Journal: Macroeconomics* 5 (4): 29–67.
- Bachmann, Rüdiger, and Lin Ma. 2016. "Lumpy Investment, Lumpy Inventories." *Journal of Money, Credit, and Banking* 48 (5): 821–55.
- Bayer, Christian, and Volker Tjaden. 2016. "Large Open Economies and Fixed Costs of Capital Adjustment." *Review of Economic Dynamics* 21: 125–46.
- Beaudry, Paul, and Alain Guay. 1996. "What Do Interest Rates Reveal about the Functioning of Real Business Cycle Models?" *Journal of Economic Dynamics and Control* 20 (9–10): 1661–82.
- Berger, David, and Joseph Vavra. 2015. "Consumption Dynamics during Recessions." *Econometrica* 83 (1): 101–54.
- Boldrin, Michele, Lawrence J. Christiano, and Jonas D. M. Fisher. 2001. "Habit Persistence, Asset Returns, and the Business Cycle." *American Economic Review* 91 (1): 149–66.
- Caballero, Ricardo J., and Eduardo M. R. A. Engel. 1999. "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach." *Econometrica* 67 (4): 783–826.
- Caballero, Ricardo J., and Eduardo M. R. A. Engel. 2007. "Price Stickiness in Ss Models: New Interpretations of Old Results." *Journal of Monetary Economics* 54: S100–121.
- Caballero, Ricardo J., Eduardo M. R. A. Engel, and John C. Haltiwanger. 1995. "Plant-Level Adjustment and Aggregate Investment Dynamics." *Brookings Papers on Economic Activity* (2): 1–39.

- Campbell, Jeffrey R.** 1998. "Entry, Exit, Embodied Technology, and Business Cycles." *Review of Economic Dynamics* 1 (2): 371–408.
- Campbell, John Y., and John H. Cochrane.** 1999. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy* 107 (2): 205–51.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber.** 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review* 101 (3): 471–75.
- 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.
- Clementi, Gian Luca, and Berardino Palazzo.** 2015. "On The Calibration of Competitive Industry Dynamics Models." Unpublished.
- Clementi, Gian Luca, and Berardino Palazzo.** 2016. "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations." *American Economic Journal: Macroeconomics* 8 (3): 1–41.
- Cooper, Russell W., and John C. Haltiwanger.** 2006. "On the Nature of Capital Adjustment Costs." *Review of Economic Studies* 73 (3): 611–33.
- Cooper, Russell, and Jonathan Willis.** 2014. "Discounting: Investment Sensitivity and Aggregate Implications." Unpublished.
- Doms, Mark, and Timothy Dunne.** 1998. "Capital Adjustment Patterns in Manufacturing Plants." *Review of Economic Dynamics* 1 (2): 409–29.
- Edge, Rochelle M., and Jeremy B. Rudd.** 2011. "General-Equilibrium Effects of Investment Tax Incentives." *Journal of Monetary Economics* 58 (6–8): 564–77.
- Fernald, John.** 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Federal Reserve Bank of San Francisco Working Paper 2012-19.
- Gourio, Francois, and Anil K. Kashyap.** 2007. "Investment Spikes: New Facts and a General Equilibrium Exploration." *Journal of Monetary Economics* 54: S1–22.
- Greenwood, Jeremy, Zvi Hercowitz, and Gregory W. Huffman.** 1988. "Investment, Capacity Utilization, and the Real Business Cycle." *American Economic Review* 78 (3): 402–17.
- House, Christopher L.** 2014. "Fixed Costs and Long-Lived Investments." *Journal of Monetary Economics* 68: 86–100.
- House, Christopher L., and Matthew D. Shapiro.** 2008. "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation." *American Economic Review* 98 (3): 737–68.
- Khan, Aubhik, and Julia K. Thomas.** 2003. "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?" *Journal of Monetary Economics* 50 (2): 331–60.
- Khan, Aubhik, and Julia K. Thomas.** 2008. "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics." *Econometrica* 76 (2): 395–436.
- Koby, Yann, and Christian Wolf.** 2020. "Aggregation in Heterogeneous-Firm Models: Theory and Measurement." Unpublished.
- Koop, Gary, M. Hashem Pesaran, and Simon M. Potter.** 1996. "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics* 74 (1): 119–47.
- Krusell, Per, and Anthony A. Smith, Jr.** 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106 (5): 867–96.
- Miao, Jianjun, and Pengfei Wang.** 2014. "A Q-Theory Model with Lumpy Investment." *Economic Theory* 57 (1): 133–59.
- Reiter, Michael.** 2009. "Solving Heterogeneous-Agent Models by Projection and Perturbation." *Journal of Economic Dynamics and Control* 33 (3): 649–65.
- Syverson, Chad.** 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2): 326–65.
- Thomas, Julia K.** 2002. "Is Lumpy Investment Relevant for the Business Cycle?" *Journal of Political Economy* 110 (3): 508–34.
- Winberry, Thomas.** 2018. "A Method for Solving and Estimating Heterogeneous Agent Macro Models." *Quantitative Economics* 9 (3): 1123–51.
- Winberry, Thomas.** 2021. "Replication Data for: Lumpy Investment, Business Cycles, and Stimulus Policy." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.3886/E119560V1>.
- Xu, Qiping, and Eric Zwick.** 2020. "Tax Policy and Abnormal Investment Behavior." Unpublished.
- Zwick, Eric, and James Mahon.** 2017. "Tax Policy and Heterogeneous Investment Behavior." *American Economic Review* 107 (1): 217–48.