FINANCIAL FRICTIONS AND INVESTMENT DYNAMICS IN MULTI-PLANT FIRMS

By

**Matthias Kehrig
University of Texas at Austin**

**Nicolas Vincent
HEC Montréal**

CES 13-56

November, 2013

The research program of the Center for Economic Studies (CES) produces a wide range of economic analyses to improve the statistical programs of the U.S. Census Bureau. Many of these analyses take the form of CES research papers. The papers have not undergone the review accorded Census Bureau publications and no endorsement should be inferred. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Republication in whole or part must be cleared with the authors.

To obtain information about the series, see www.census.gov/ces or contact Fariha Kamal, Editor, Discussion Papers, U.S. Census Bureau, Center for Economic Studies 2K132B, 4600 Silver Hill Road, Washington, DC 20233, CES.Papers.List@census.gov.

Financial Frictions and Investment Dynamics in Multi-Plant Firms^{*,**}

Matthias Kehrig[†] and Nicolas Vincent[‡]

June 2013

Abstract

Using confidential Census data on U.S. manufacturing plants, we document that most of the dispersion in investment rates across plants occurs within firms instead of across firms. Between-firm dispersion is almost acyclical, but within-firm dispersion is strongly procyclical. To investigate the role of firms in the allocation of capital in the economy, we build a multi-plant model of the firm with frictions at both levels of aggregation. We show that external financing constraints at the level of the firm can have important implications for plant-level investment dynamics. Finally, we present empirical evidence supporting the predictions of the model.

KEYWORDS: Investment, Plants vs. Firms, Q -Theory, Internal vs. External Capital Markets, Diversification Discount.

JEL CODES: E2, G3

*We are grateful to participants at the NBER Summer Institute, the Ifo Conference on Macroeconomics and Survey Data, the Midwest Macro Meetings and the Annual Meeting of German Economists Abroad, Gadi Barlevy, Marco Bassetto, Jeff Campbell, Russ Cooper, François Gourio, John Haltiwanger and Mark Wright for helpful discussions.

**Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. The latest version of this paper can be downloaded at <http://papers.ssrn.com/abstract=2285511>

[†]University of Texas at Austin, Department of Economics, 2225 Speedway Stop C3100, Austin, TX 78712, matthias.kehrig@austin.utexas.edu.

[‡]Institute of Applied Economics, HEC Montréal, 3000, chemin de la Côte-Sainte-Catherine, Montréal (Québec) H3T 2A7, nicolas.vincent@hec.ca.

1 Introduction

A considerable body of research has tried to understand the forces and frictions that shape capital investment decisions at the micro level. Yet, despite all this work, little consensus has emerged regarding the frictions that are economically important for investment dynamics. For example, some have argued that “technological” frictions such as a fixed cost of investing are crucial to replicate the lumpy aspect of investment documented empirically (see among others Caballero, Engel and Haltiwanger (1995); Cooper and Haltiwanger (2006); Gourio and Kashyap (2007)). Others have instead focused on the central role played by financing frictions arguing that they offer a natural explanation for the documented role of cash flow in investment regressions (see Fazzari, Hubbard and Petersen (1988); Gilchrist and Himmelberg (1995)). Arguably, a better understanding of the nature of frictions is of utmost importance as various frictions may have very different implications for macroeconomic aggregates. For example, some have argued that the introduction of fixed costs of investing in macroeconomic models has little implications for the behaviour of macro variables. On the other hand, if financial frictions are of first-order for micro-level investment dynamics, this may have important implications about the role of the financial accelerator for business cycles.

One striking feature of the empirical literature on investment is that it has been for the most part unconcerned with the organisational structure of firms: models are built around a single production unit, while empirical studies use either plant- or firm-level data depending on data availability. There are two main issues with this approach. First, because the level of aggregation used is application-specific, it makes it very difficult to compare the relative importance of the various frictions at play. For example, technological frictions such as investment irreversibilities or factory downtime costs matter for investment at the plant rather than the firm level as they are tied to the physical investment process at the plant as the location of production. Financial constraints, however, are more likely to arise in the firm-level context. Second, the level of aggregation may alter our assessment of the macroeconomic relevance of a given friction. For example, while the investment activity of plants may be lumpy, this picture changes a fair amount if one considers investment at the *firm level* rather than the *plant level*.¹ If this result is a product of a strategy by the firm to stagger capital expenditures across its plants, lumpiness may be of little importance for macroeconomic dynamics.

In this paper, we investigate how taking into account the multi-plant nature of firms informs us about the relative importance of various frictions for micro-level investment dynamics. First, using the data from the Annual Survey of Manufacturers between 1972 and 2010, we document that most dispersion dynamics in both investment and productivity occur *within* firms more so than they occur between firms. Most of the action within the firm is coming from plants with an investment spike, i.e. those undergoing large investment projects, in line with what others have found for the universe of plants (Doms and Dunne (1998), and Gourio and Kashyap (2007)).

Having shown that the allocation of capital across plants within the firm is a crucial dimension,

¹Eberly, Rebelo and Vincent (2012) show that investment is fairly smooth on the firm level.

we next build a multi-plant model of the firm. Plants face technical frictions such as fixed costs of investing or convex adjustment cost. Firms operate several plants and are subject to borrowing constraints. We simulate the model to show that firm-level external financing constraints can have a large effect on plant-level investment dynamics. For example, in the wake of a rise in the cost of borrowing, the firm tries to stagger large investment projects to minimise the need for external funds leading to a fall in the correlation of investment across plants but an increase in the autocorrelation of firm-level investment. Finally, in the last section, we use the Chicago Fed National Index of Financial Condition as a proxy for the availability of external funds and show that many of the predictions of the models are supported in the ASM data.

We see our project as a first step into modelling how the organisational structure of a firm impacts micro-level adjustment. In general, this firm-level dimension offers a new dimension to identify investment frictions, because it allows the researcher to study the joint investment dynamics of all plants within a firm. Some theoretical research has been done on the efficiency of internal versus external capital markets: [Gertner, Scharfstein and Stein \(1994\)](#); [Stein \(1997\)](#); [Malenko \(2012\)](#). With the exception of [Lamont \(1997\)](#); [Schoar \(2002\)](#); [Giroud \(forthcoming\)](#), empirical research on within-firm dynamics in general, is scarce. We attempt to fill this gap and provide also a theoretical explanation of the dynamics we observe in the data. Then, our paper is linked to previous research on the nature of adjustment cost ([Abel and Eberly \(1996\)](#); [Cooper and Haltiwanger \(2006\)](#); [Caballero, Engel and Haltiwanger \(1995\)](#); [Cooper, Haltiwanger and Power \(1999\)](#)) and their macroeconomic consequences ([Thomas \(2002\)](#); [Khan and Thomas \(2008\)](#)).

Our paper is organised as follows. In Section 2, we describe the data and show evidence on the importance of the within-firm dimension for investment dispersion. Section 3 describes our multi-plant model of the firm and analyses its predictions when an external financing constraint is introduced. In Section 4, we investigate whether the model predictions are borne out in the micro-level data. Section 5 concludes.

2 Empirical Motivation

2.1 Data

We use Census data on manufacturing establishment (plants) 1972-2010 from the Annual Survey of Manufactures and the Census of Manufactures. While these datasets limit our analysis to the manufacturing sector, it gives a much better representation of the firm universe than comparable datasets such as COMPUSTAT that capture only publicly traded firms which make up only 5% of the full sample in our dataset.² The goal of our analysis on this comparatively rich data is to compute and study the heterogeneity in investment rates across plants in the manufacturing sector. Investment heterogeneity indicates that some agent in the economy discriminates across

²[Davis et al. \(2006\)](#) have documented that publicly traded firms exhibit employment dynamics that are very different from privately held firms. Within the Census data, we can identify firms that are publicly traded from the COMPUSTAT-SSEL bridge. We will use this information to examine how publicly traded firms operate differently than privately held ones.

plants when allocating investment resources. We are interested how much heterogeneity there is, whether or not it fluctuates over the business cycle and which agent is relevant in making investment decisions. As for the latter, we want to distinguish across plants and firms as different agents, where the former is the smallest production unit within a firm and the latter has to compete in markets. This analysis should highlight how two fundamentally different institutions – markets versus firms – decide on an investment allocation decision.

The Census data allow us to measure annual investment, capital, output and cash flow (raw profits) for about 50k manufacturing establishments annually. For details about the data measurement and the imputation (mostly: capital) is described in [Kehrig \(2013\)](#). One shortcoming of our current analysis is that we do not have a measure of Tobin’s Q in the Census data. We plan to use corporate valuation techniques on the Census data and construct measures of Q in the future.

2.2 The Empirics of Investment Dispersion

To get an idea about the cross-sectional heterogeneity of investment, we compute the weighted cross-sectional variance of investment rates across plants in the economy:

$$\sigma_t = \sum_n \omega_{nt} \left[(i/k)_{nt} - \overline{(i/k)}_t \right]^2$$

where σ_t is the cross sectional *variance* at time t , n indicates the plant, i_{nt} and k_{nt} the investment and capital level of plant n , $\overline{(i/k)}_t$ the weighted average of investment rates and $\omega_{nt} = k_{nt}/K$ the share of the plant’s capital stock which we use as the weight of plant n . We focus on weighted dispersion because outliers will have a large impact on the measured dispersion.³

Such an exercise has been carried first by [Bachmann and Bayer \(2011\)](#) who document a procyclical dispersion on the *firm* level in German data. We focus on the dispersion on the *plant* level in U.S. manufacturing instead. The results are presented in [Figure 1](#).

Just visually inspecting the time series for dispersion and the aggregate investment rate, one can see a clear positive correlation: The contemporaneous correlation coefficient is 0.72 (see [Figure 4](#)), but investment dispersion continues to be positively correlated with a one-year lead and two-year lags. Using aggregate investment in manufacturing seems like the natural choice, but the positive correlation remains when using other measures of the cycle such as industrial production or GDP.

A natural question is to find out if this results is driven by the changes in the extensive margin, i.e. the measure of plants investing at all. The extensive margin can vary due to birth and death of plants (assuming that new-born young plants invest more to grow and dying plants are shrunk and invest less) and lumpy investment. The birth/death margin turns out to not matter, so we omit any discussion here. The lumpiness is a more serious concern because previous research has documented that investment is in fact lumpy (see [Doms and Dunne \(1998\)](#); [Cooper and Haltiwanger \(2006\)](#); [Gourio and Kashyap \(2007\)](#)). To test if lumpy investment drives the observed investment

³Alternatively, one may also use investment shares as weights.

FIGURE 1: INVESTMENT DISPERSION

heterogeneity, we limit the analysis to the plants that have strictly positive investment.⁴ These results are displayed in Figure 2.

Although investment dispersion changes a bit on average, the magnitude of this change is rather small and the overall cyclicity patterns are preserved as well. So we conclude that investment dispersion is not driven by the extensive margin. Although we focus on weighted investment dispersion here, this robustness check still gives a significantly positive correlation of the unweighted dispersion (albeit it is weaker).

Since the investment dispersion and its cyclicity are not a mere result of lumpy investment, we proceed and look at the agents that decide on investment. The Census data are collected at the plant level. Plants, in turn, are part of a firm that has the organisational control over the plants. These two levels of aggregation are also relevant in the sense that firms have to compete in markets while plants are active in the internal market of a firm. Also, investment is notoriously determined by frictions such as adjustment cost (see for example [Cooper and Haltiwanger \(2006\)](#); [Gourio and Kashyap \(2007\)](#)) and credit constraints. While the latter typically affect the firm, the former affect the plant. Table 1 shows that the average manufacturing firm in the U.S. operates on average 36 plants.⁵ This shows that there are possibly many agents within the firm that compete for investment funds. The observations labeled “LBD” come from the Longitudinal Business Database reports employment and firm affiliation of all active establishments (though no investment, output or capital data). This comparison indicates whether or not the ASM sample is very much distorted towards large multi-unit firms which does not seem to be the case.

⁴Following the literature, we define positive investment as plants with an investment of less than 1%.

⁵Our preferred specification is the average weighted by capital, so we get a sense of the relevant magnitude. Measured in an unweighted fashion, the average firm operates two establishments which obviously reflects the fact that most capital is concentrated in firms that operate many plants.

FIGURE 2: INVESTMENT DISPERSION

TABLE 1: NUMBER OF PLANTS PER FIRM IN U.S. MANUFACTURING

Dataset	Weights	Mean	Bottom Quartile	Median	Top quartile
ASM	unweighted	1.9	1	1	1
ASM	K -weighted	35.9	5	20	50
ASM	L -weighted	30.5	2	12	41
ASM	Y -weighted	34.0	4	18	47
LBD	unweighted	1.2	1	1	1
LBD	L -weighted	31.0	1	7	39

Note: ASM sample are the $\approx 50k$ annual plant-level observations in the ASM 1972-2010. The LBD sample contains all active manufacturing establishments 1976-2009. K -, L -, and Y -weighted refer to means and quartiles that are weighted by capital, employment and production, respectively. For data construction see [Kehrig \(2013\)](#).

Lastly, we dissect investment dispersion into movements between and within firms. This can easily be done by rewriting the dispersion across plants as

$$\begin{aligned}
\sigma_t &= \sum_n \omega_{nt} \left[(i/k)_{nt} - \overline{(i/k)}_t \right]^2 \\
&= \underbrace{\sum_j \omega_{jt} \left[(i/k)_{jt} - \overline{(i/k)}_t \right]^2}_{\sigma_t^B \text{ between}} + \underbrace{\sum_j \omega_{jt} \sum_n^{N_j} \tilde{\omega}_{njt} \left[(i/k)_{njt} - (i/k)_{jt} \right]^2}_{\sigma_t^W \text{ average within}}
\end{aligned} \tag{1}$$

where j denotes the firm which owns plant n , N_{jt} are the number of plants in firm j and $\tilde{\omega}_{njt} = \omega_{nt}/\omega_{jt}$ denotes the weight of plant n within its firm. We are interested in the time series of σ_t^B , the between-firm dispersion, and σ_t^W , the within-firm dispersion. They are displayed in Figure 3.

FIGURE 3: INVESTMENT DISPERSION BETWEEN AND WITHIN FIRMS

Clearly, there are key differences in these two time series: the within dispersion is quantitatively higher (it accounts for about 70% of the long-run overall dispersion) and appears to be the key driver of overall dispersion at business cycle frequency. Also, it is more correlated with the investment cycle as shown in Figure 4, but one can easily see that the ups and down in investment dispersion coincide with the ups and downs of with-firm investment dispersion.

TABLE 2: SUMMARY STATISTICS: INVESTMENT DISPERSION

Statistic	Total σ_t	Between σ_t^B	Within σ_t^W
Average	0.0083	0.0026	0.0058
Volatility	$1.62 \cdot 10^{-3}$	$0.71 \cdot 10^{-3}$	$1.39 \cdot 10^{-3}$
$Corr(I/K, \dots)$	0.73	0.25	0.72
$Corr(\sigma_t, \dots)$	1	0.52	0.91

FIGURE 4: CYCLICALITY OF INVESTMENT DISPERSION

Correlation of aggregate manufacturing investment (I/K) with overall investment dispersion (σ_t , left panel), between-firm (σ_t^B) and within-firm (σ_t^W) investment dispersion (right panel).

Lastly, we want to see what is driving the investment dynamics within a firm that cause this large and procyclical investment dispersion. We therefore decompose the within-firm dispersion further into a part that is driven by large investors (investment rate larger than 20%), small investors (<20%) and the margin at 0:

$$\sigma_t^W = \sum_{jt} \omega_{jt} \sigma_{jt}^W = \sum_j \omega_{jt} \left\{ \sum_n^{N_{jt}^{big}} \omega_{njt} \left[(i/k)_{njt}^{big} - \overline{(i/k)}_{jt} \right]^2 + \sum_n^{N_{jt}^{sm}} \omega_{njt} \left[(i/k)_{njt}^{sm} - \overline{(i/k)}_{jt} \right]^2 + \sum_n^{N_{jt}^0} \omega_{njt} \left[(i/k)_{njt}^0 - \overline{(i/k)}_{jt} \right]^2 \right\}$$

The results in Figure 5 show that plant that undergo a large investment project are driving almost the entire with-firm dispersion.

FIGURE 5: INVESTMENT DISPERSION

This results leads us to focus on these plants that undergo large investment projects. Why do firms select them? Are they more productive than the other plants within that same firm? We explore these questions in the next section.

2.3 Robustness

2.3.1 Entry and Exit of Plants

A high investment dispersion could come about because firms open new plants and investment in them heavily in booms. Entry of new plant is procyclical, so a procyclical investment dispersion could result from a procyclical entry and heavy start-up investment. Countercyclical exit has a similar effect. To control for that issue, we drop all establishments that are active for two years or less and will not exit within the next two years. The between-firm within-firm decomposition on the remaining “mid-age” plants looks similar to the one of the overall sample. The results are displayed in Figure 6.

FIGURE 6: INVESTMENT DISPERSION OF MID-AGE FIRMS

2.3.2 Are Firms Random Collection of Plants?

That most of the investment and productivity dispersion originates within firms proved to be a robust to explanations such as lumpy investment and entry and exit of production units. But before one jumps to conclusion about the (ir-)relevance of firms one should consider the possibility that firms are just random collection of plants. This means that plant-level investment could be purely driven by plant-specific characteristics. In that was true, firm affiliation would not matter because any firm manager would decide on the same investment. But if firm characteristics such as available credit matters or within-firm spill-overs or complementarities for plant investment, then firm affiliation does matter.

In order to test the hypothesis that firm affiliation is irrelevant, we construct 100 randomised samples where in each sample we replace the investment rate of a plant with a randomly drawn investment rate from another plant in the same year and the same 3-digit NAICS industry. We

call this perturbation Random Sample I. We keep the year and industry fixed to control for the fact that investment patterns and the nature of firm affiliation differ across industries. Of course, this is rather a coarse grid and we carry out a second, finer randomisation where we replace plants within 4-digit NAICS industries and the four Census Regions (Northeast, Midwest, West, South) to additionally control for regional investment patterns; we call this finer perturbation Random Sample II. Across these two random samples, we perform the same between- and within-firm decomposition of investment rates. Naturally, every random sample will yield a slightly different decomposition, so we use the standard deviation across the different random draws as error bands. The results are displayed in Table 3 and show that the value of the between- and within-firm variances in the actual data are outside the 95% error bands of the random draws. So we conclude that firms are not random collection of plants and that considering firm-level frictions matter for plant-level investment outcomes.

2.4 Summary statistics about investment and productivity

So far we have focused on cross-sectional moments in the investment data. Our goal is to explain the facts about within- versus between-firm investment and productivity dispersion and how frictions at both the plant and the firm level could deliver the served data facts. In order to do so, we will need a theory of a multi-unit firm and how it allocates investment resources across plants. Since this is computationally not an easy task – the state space increases exponentially in the number of plants, we will approach this problem by writing a model of a firm that operates two plants only. Given our data in Table 1 this is clearly an understatement of the average complexity, but we see this effort as a first step in writing a model of a more diversified firm. To guide the modelling, we present empirical facts about 2-plant firms.

Table 5 presents time series facts at the plant and the firm level about investment, average output per capital and cash flow (per capital). As is well-known, investment is lumpy at the plant level, so the autocorrelation is negative. At the firm level, investment is a bit smoother, but the difference is small – obviously because we consider the minimum-plant multi-unit firm and the observations are not weighted by capital or employment. In addition to the full sample, we present statistics broken down by the characteristic of whether or not a firm is publicly traded or not. This will play an important role in our analysis of firm-level (financial) frictions. not surprisingly, publicly traded firms invest more on average, are less volatile and more smooth.

Next, we focus on the joint statistics of 2-plant firms which are displayed in Table 6. We look at the (joint) likelihood of investment and the likelihood of investment spikes. Given that most of the investment heterogeneity is driven by units that invest more than 15%, the dynamics of plants with investment spikes look most important. According to our definition, about one in seven observations undergoes an investment spike. The share of investing firms is 75% which looks a bit large given previous work that established the importance of lumpy investment. This high number comes about because we consider only the ASM which is that subsample of the manufacturing sector that represents most of economic activity and size; so it's natural that most of these plants

TABLE 3: BETWEEN-WITHIN DECOMPOSITION OF K -WEIGHTED DISPERSION

Statistic	Actual Firms	Random Firms I	Random Firms II
σ	0.0084	0.0084	0.0084
σ^B	0.0026	0.0036 (0.0003)	0.0031 (0.0002)
σ^W	0.0058	0.0048 (0.0003)	0.0053 (0.0002)

Note: Table displays the average (over time) investment dispersion overall (σ_t) and its decomposition into between-firm (σ_t^B) and within-firm dispersion (σ_t^W). Random Firms I refers to a perturbation of plants and their investment rate within a given year and 3-digit NAICS industry, Random Firms II refers to a random perturbation of plants and their investment rate within given years, 4-digit NAICS industries and the four Census regions. The standard errors below the between and within firm component denote the time series average of the standard deviation across 100 bootstraps.

TABLE 4: BETWEEN-WITHIN DECOMPOSITION OF K -WEIGHTED DISPERSION OF APK!!!

Statistic	Actual Firms	Random Firms I	Random Firms II
σ	3.68	3.68	3.68
σ^B	1.446	1.582 (0.512)	
σ^W	2.24	2.10 (0.512)	

Note: Table displays the average (over time) investment dispersion overall (σ_t) and its decomposition into between-firm (σ_t^B) and within-firm dispersion (σ_t^W). Random Firms I refers to a perturbation of plants and their investment rate within a given year and 3-digit NAICS industry, Random Firms II refers to a random perturbation of plants and their investment rate within given years, 4-digit NAICS industries and the four Census regions. The standard errors below the between and within firm component denote the time series average of the standard deviation across 100 bootstraps.

TABLE 5: MOMENTS 2-PLANT FIRMS

Variable	i/k				y/k				cf/k			
	Mean	StD	CV	AR	Mean	StD	CV	AR	Mean	StD	CV	AR
<i>Plant level</i>												
All	0.094	0.085	0.90	-0.07	3.67	1.34	0.37	0.17	1.26	0.60	0.48	0.09
Private	0.094	0.085	0.90	-0.06	3.75	1.35	0.36	0.17	1.29	0.61	0.47	0.09
Public	0.097	0.086	0.88	-0.08	2.80	1.06	0.38	0.12	0.96	0.49	0.51	0.11
<i>Firm level</i>												
All	0.083	0.06	0.72	-0.06	3.06	0.87	0.28	0.18	1.05	0.40	0.38	0.11
Private	0.083	0.06	0.72	-0.06	3.12	0.88	0.28	0.18	1.07	0.41	0.38	0.11
Public	0.090	0.06	0.67	-0.02	2.41	0.74	0.31	0.15	0.81	0.34	0.42	0.09

Note: The statistics displayed above are ... **Describe!**

also tend to invest.

How rare is it that both plants in a firm undergo an investment spike at the same time? Very rare, this probability is less than 6% although it is considerably higher for publicly traded firms (above 9%). The probability that at least one plant in the firm spikes is much higher (23.8%, but even 28.4% for publicly traded firms). Because we are interested in the serial correlation of investment spikes within a firm, we also focus on the ample of firms that have at least one spiking plant and look at the time series properties of firms that have an investment spike. We compute the conditional probabilities (on having at least on spiking plant today) that firms last year had one, two or no spiking plant. Interestingly, slightly more than a third of firms that had one spiking plant today, had exactly one spiking plant last year, and 1 in 20 firms that has at least one spiking plant this year, had two spikes last year. This share is much higher for publicly traded firms (that are presumably not as credit constrained).

3 A Model of the Multi-Plant Firm

In this section, we describe, solve, simulate and analyse a simple model of a firm comprised of more than one plant. We study how various plant- and firm-level frictions interact with the optimal allocation of capital by the firm across its plants. At one extreme, the firm is a collection of disconnected plants: decisions are made on a plant-by-plant basis, without any interactions between them. We show that in the presence of frictions, the firm alters the size and timing of investment plans.

TABLE 6: SUMMARY STATISTICS FOR 2-PLANT FIRMS – UNWEIGHTED

Statistic	All	Privately held	Publicly traded
<i>Overall probabilities</i>			
$Pr((i/k) > 0.01)$	75.4%	75.2%	82.5%
$Pr((i/k) > 0.15)$	13.7%	13.6%	18.4%
<i>Within-firm probabilities</i>			
$Pr((i/k)_n < 0.01) \forall n = A, B$	10.3%	10.3%	9.2%
$Pr((i/k)_A \geq 0.01, (i/k)_B < 0.01)$	28.8%	29.0%	25.9%
$Pr((i/k)_n \geq 0.01) \forall n = A, B$	66.0%	65.7%	69.4%
$Pr((i/k)_n \geq 0.01)$ for one n	89.7%	89.7%	90.9%
$Pr((i/k)_n \geq 0.15) \forall n = A, B$	3.6%	3.5%	5.0%
$Pr((i/k)_n \geq 0.15)$ for one n	23.8%	23.5%	28.4%
$Pr((i/k)_{nt-1} > 0.15) \forall n = A, B$	5.8%	5.5%	9.2%
$Pr((i/k)_{nt-1} > 0.15)$ for one n	37.9%	37.3%	44.3%

3.1 The Problem of the Firm

We focus on the basic problem of a firm that operates two plants, A and B . This problem can be written in recursive form as:

$$V(z_A, z_B, k_A, k_B) = \max_{i_A, i_B} \{ \Pi_A + \Pi_B - i_A - i_B - \Theta(\mathbf{z}, \mathbf{I}, \mathbf{K}) + \beta EV(z'_A, z'_B, k'_A, k'_B) \}$$

where k_A and k_B are the beginning-of-period capital stocks of plants A and B respectively. z_n is the level of productivity of plant n , which could include a firm-specific component that is common to both plants. Π_A and Π_B are cash flows at the plant level, net of any fixed or variable costs (they will be discussed more in details later). Firm-level costs are summarised through the function Θ which may depend on the vectors of productivity ($\mathbf{z}$), investment ($\mathbf{I}$) and capital stock ($\mathbf{K}$).

Note that we are making the implicit assumption that plant-level profits are separable. Indeed, we will initially leave aside the possibility of interactions across plants, such as complementarities. The per-period cash flow function of plant n is given by:

$$\Pi_n = z_n k_n^\alpha - \phi - \theta(i_n, k_n)$$

where ϕ is an operating fixed cost and $\theta(i_n, k_n)$ is a cost function that may include variable and/or fixed investment adjustment costs.

The law of motion for the capital stock of plant n is standard:

$$k'_n = (1 - \delta) k_n + i_n$$

3.2 Plant- and Firm-Level Frictions

So far, we have only expressed investment frictions as general functions, $\Theta(\mathbf{z}, \mathbf{I}, \mathbf{K})$ and $\theta(i_n, k_n)$. Next, we describe more specifically the types of frictions we will study in our simulations.

As documented in a number of studies (REFS), investment dynamics at the plant level are characterised by lumpiness: multiple periods of inactivity (no investment) are followed by investment spikes. The traditional modelling feature used to reproduce this stylized fact is to introduce a fixed cost of investing: the firm must pay a certain cost, κ , if investment is greater than zero. Such costs can arise because investment activity has a disruptive effect on production activities in the short run, for example. This friction will be the central part of our plant-level cost function, $\theta(i_n, k_n)$. In addition, we include a traditional quadratic cost (REFS) adjustment cost, though mostly in order to make the plant-level profit function well behaved (VERIFY THIS: this friction will turn out to be of little interest in our context). To summarise, frictions at the plant level will be expressed as:

$$\theta(i_n, k_n) = \Psi \mathbf{I}\left(\frac{i_n}{k_n} > \vartheta\right) + \gamma \left(\frac{i_n}{k_n}\right)^2$$

where $\mathbf{I}$ is an indicator function equal to 1 if the plant investment rate is above ϑ (which is itself close to zero). Ψ and γ are parameters.

At the firm level, our focus is on the impact of financing frictions on investment dynamics. Our basic specification is simple: the firm needs to borrow if its capital expenditures are greater than the cash flows generated by the plants (net of all other costs). Any financing need leads to a firm-level cost that is exponentially increasing in the size of the financing needs, plus a fixed financing cost (for example, the managerial resources involved in dealing with the bank). One can think of this specification as a reduced form for a setup where the interest rate paid on the “loan” is increasing in its size due to default risk. More specifically, we use the following cost function:

$$\begin{aligned} \Theta(\mathbf{z}, \mathbf{I}, \mathbf{K}) &= \left[\eta \left(\frac{B}{K}\right)^2 + \zeta \right] K \cdot \mathbf{I}\{B > 0\} + \Phi \\ B &= i_A + i_B - (\Pi - \Phi) \\ K &= k_A + k_B \\ \Pi &= \Pi_A + \Pi_B \end{aligned}$$

where B represents the amount of external financing needed; $\mathbf{I}\{B > 0\}$ is an indicator function equal to 1 if B is positive; ζ is a fixed external financing cost; and Φ is the firm-level fixed operating cost. The firm is not financially constrained if both η and ζ are equal to zero.

3.3 Solving and simulating the model

The model is solved using a value function iteration procedure which is described in detail in Appendix XXX. We then simulate a panel of two-plant firms for T periods. In this section, we

present how the various frictions affect investment dynamics. To do so, we focus on a number of simple statistics that we compare between the base case and alternative scenarios.

3.4 The effect of firm-level financing frictions

3.4.1 Comparative statics

The decision of the firm described above reverts to two separate plant-level optimisation problems if $\eta = \zeta = 0$, i.e. if firm-level frictions are absent. We will start from such a situation and study how the introduction of financing frictions alter the properties of investment at both the plant and firm level.

The first column in Table 7 shows a selected number of moments for our baseline case ($\eta = \zeta = 0$). While the objective of this section is to analyse some comparative statics and not to model the average firm, we have tried to make sure that the simulated moments from our baseline are loosely in line with the empirical moments for 2-plant publicly traded firms. In particular, we used the coefficient of variation and serial correlation of output (as a fraction of the capital stock) to pin down the shock process at the firm and plant levels. We also adjusted the plant-level investment fixed and variable cost parameters (Ψ , ϑ and γ) in order to match the serial correlation of investment and some moments related to the prevalence of spikes.

Once again, it is important to notice that in the baseline scenario, the plants operate separately without any interaction. For example, the investment decision of plant A is in no way a function of the productivity shock or the investment timing of plant B. Hence, the fact that joint spikes are observed in 8% of the periods is simply by chance, not for optimality reasons from the firm's perspective.

Next, we look at the impact of the financial friction parameters η and ζ on plant and firm investment dynamics. The second column in Table 7 shows what happens when we turn on the fixed cost of borrowing, setting it to 1% of the capital stock. Not surprisingly, the firm is now much less likely to borrow: while the probability that capital expenditures were larger than funds available used to be around 35% in the baseline scenario (where the firm could borrow for free), it is now happening only 4% of the time. This seems to happen at least partly because the firm is trying to stagger investment activity across its two plants. The first hint of this is coming from the fact that the correlation between $(i/k)_{At}$ and $(i/k)_{Bt}$ falls from 0.28 to 0.14. In addition, the probability of observing an investment spike ($i/k > 0.15$) in both plants at the same time is halved (0.08 to 0.04) even though the unconditional likelihood of a spike at the plant level is slightly higher. The average size of investment spikes, however, is lower.

These findings are very similar if we instead introduce a quadratic borrowing cost. Focusing on the case where both types of frictions are activated (last column of Table 7), the impact on the optimal allocation of the firm is even more striking. Correlation between investment at the two plants actually turns sharply negative (from 0.28 to -0.33) yet the serial correlation of investment aggregated at the firm level increases significantly, from 0 in the baseline to 0.18 with financial

frictions. In addition, the probability that a spike in one plant is followed by a spike in the other plant rises (0.09 to 0.18), and the probability of a joint spike is now zero.

Next, we introduce a stochastic process for the degree of credit frictions. This will allow us to more closely compare the model predictions to the empirical results.

TABLE 7: IMPACT OF FINANCIAL FRICTIONS ON SELECTED MOMENTS FROM THE MODEL

	None	Fixed only	Convex only	Both
Fixed cost η	$\eta = 0$	$\eta = 0$	$\eta = 5$	$\eta = 5$
Convex cost ζ	$\zeta = 0$	$\zeta = 0.01K$	$\zeta = 0$	$\zeta = 0.01K$
$corr\left(\frac{i_{At}}{k_{At}}, \frac{i_{At-1}}{k_{At-1}}\right)$	-0.05	-0.07	-0.11	-0.10
$corr\left(\frac{i_t}{k_t}, \frac{i_{t-1}}{k_{t-1}}\right)$	0.00	0.04	0.17	0.18
$corr\left(\frac{i_{At}}{k_{At}}, \frac{i_{Bt}}{k_{Bt}}\right)$	0.28	0.14	-0.29	-0.33
$\Pr(\text{Borrow})$	0.43	0.04	0.43	0.05
$\Pr\left(\frac{i_{At}}{k_{At}} > 0.15\right)$	0.20	0.23	0.21	0.24
$\Pr\left(\frac{i_{At}}{k_{At}} > 0.15 \ \& \ \frac{i_{Bt}}{k_{Bt}} > 0.15\right)$	0.08	0.04	0.00	0.00
$\Pr\left(\frac{i_{At}}{k_{At}} > 0.15 \ \& \ \frac{i_{Bt-1}}{k_{Bt-1}} > 0.15\right)$	0.09	0.16	0.16	0.18

3.4.2 Introducing time-varying financial frictions

In this section we move away from comparative statics by incorporating a two-state process for the financial friction: the firm will alternate between states with low or high degrees of financial frictions. This exercise allows us to run on the simulated data regressions that are similar to the empirical specifications discussed in the next section where we use a time-series index of financial conditions to determine how investment dynamics are affected by changes in financial frictions.

For this exercise, we continue to set the fixed cost of borrowing, ζ , to 1% of the capital stock. This parameter will be time-invariant and can be interpreted for example as the management costs related to preparing a loan application and interacting with the financial intermediary. The other borrowing friction parameter, η , is time-varying: it can take the values 0 (low borrowing cost) or 5 (high borrowing cost) depending on the degree of credit tightness, and the probability of switching between the two states is equal to 0.1. All the other parameters of the model are unchanged.

We simulate the model and run regressions on the simulated data. The move to regression analysis is important if we want to link the predictions of the model with the data: while our model allows us to perfectly isolate the role of borrowing constraints, in the real world investment dynamics may be affected by multiple sources of heterogeneity unrelated to financial frictions. The use of controls is therefore crucial.

It should be noted that we are in no way trying to calibrate the size and relative importance of the firm-level financial frictions at this point. Therefore, what we are interested in determining

whether the model predictions are in line with what we observe in the data from a qualitative, not quantitative, standpoint.

Table 8 shows results for plant-level investment regressions. Each regression uses the investment-to-capital ratio for plant A at time t as the dependent variable (the two plants are perfectly symmetric in the model), $(i/k)_{At}$. In order to determine the impact of firm-level financial frictions on investment dynamics, we define a dummy variable, ς_t , equal to one if credit is tight in period t (i.e. $\eta = 5$) and zero otherwise.

The first two regressions of Table 8 focus on the role of output/cash flow variables in explaining movements in investment. Not surprisingly, plant-level investment is on average lower in periods of high borrowing costs ($\varsigma_t = 1$). In the top panel we can see that a 10% increase in the output-to-capital ratio of the plant raises its i/k ratio by about 1.3%. Interestingly, plant A's investment is also affected by plant B shocks, to a lesser degree. There are two possible reasons for this result. First, simultaneous increases in output at both plants A and B are potentially indicative of a firm-wide shock. Since firm shocks are more persistent than plant-level shocks, the optimal decision is to invest more.

Second, the internal finance channel is also at play: in a context where funds are scarce and borrowing costly (recall that $\zeta = 0.01K$ in all states), a good shock in plant B generates precious cash flow that can be used to finance investment in plant A. This is also evident in the results for the second regression of the same table. There, plant B's output is replaced by cash flow at the level of the firm, net of adjustment and fixed operating costs (notice that we do not use logs as cash flow is sometimes negative). Not only is plant A investment higher when firm cash flows are higher (conditional on the plant-specific shock), but we can see that this dependence on the firm's financial resources is particularly strong when credit is tight: the coefficient on cf_t/k_t more than doubles when $\varsigma_t = 1$. In other words, the existence of internal capital markets is particularly relevant when external financial constraints are more binding.

The last regression of Table 8 looks at the correlation in investment activity across plants within the firm. In periods where credit is cheap, $(i/k)_{At}$ is a positive function of investment activity in plant B, though the relationship is somewhat weak: a 1 percentage point increase in $(i/k)_{Bt}$ raises $(i/k)_{At}$ by less than 0.1 percentage point. The relationship, however, changes dramatically when credit is tight: in periods where $\eta = 5$, the same 1 p.p. change in investment at plant B leads instead to a *fall* of almost 0.5 percentage point in $(i/k)_{At}$ as investment activity in one plant crowds out investment elsewhere. With scarce funds, the firm selects the most profitable projects, postponing others as their marginal benefit is outweighed by the marginal cost of external funds.

The firm-level regression in Table 8 highlights another impact of financial frictions in multi-unit firms: the serial correlation of firm-level investment increases. As shown earlier in Table 7, our baseline calibration implies that the autocorrelation of i_t/k_t is slightly negative, which should not be too surprising as these firms are very small. However, a financial constraint shock makes firm-level investment significantly smoother. One potential explanation for this result is related to our earlier findings: when credit is tight, it makes the firm less likely to invest in both plants at the same

TABLE 8: PLANT-LEVEL INVESTMENT REGRESSIONS

Dependent variable		$(i/k)_{At}$
Constant		0.419***
Credit tightness	ς_t	-0.075***
Output plant A	$\log(y/k)_{At}$	0.129***
	$\varsigma_t \cdot \log(y/k)_{At}$	-0.014***
Output plant B	$\log(y/k)_{Bt}$	0.063***
	$\varsigma_t \cdot \log(y/k)_{Bt}$	-0.011**
R^2		0.33
Constant		1.647***
Credit tightness	ς_t	-0.074***
Output plant A	$(y/k)_{At}$	0.584***
Firm cash flow	$(cf/k)_t$	0.246***
	$\varsigma_t \cdot (cf/k)_t$	0.356***
R^2		0.34
Constant		1.882***
Credit tightness	ς_t	0.017***
i/k plant B	$(i/k)_{Bt}$	0.085***
	$\varsigma_t \cdot (i/k)_{Bt}$	-0.545***
Cash flow firm	$(cf/k)_t$	-0.007
	$\varsigma_t \cdot (cf/k)_t$	0.882***
R^2		0.40

Note: Dependent variable is i/k of plant A. $\varsigma_t = 1$ in periods of high borrowing cost, 0 otherwise. Controls such as capital stock or cash flow are included but not always reported. *, ** and *** indicate significance at the 10, 5 and 1 percent level respectively.

time in order to avoid costly borrowing. To understand why this may lead to higher autocorrelation of investment for the firm, consider the example of a positive firm-level shock. Since both z_A and z_B are now higher, both plants would now like to invest to reach their new optimal level of capital. But given limited cash flows, this implies that the firm would need to obtain costly funds on capital markets. Instead, it will sometimes find it optimal to stagger investment: plant A invests today, while plant B waits until tomorrow. By construction, this makes firm-level investment smoother. We show additional evidence for this kind of behaviour below.

TABLE 9: FIRM-LEVEL INVESTMENT REGRESSIONS

$$(i/k)_t = \beta_0 + \beta_1 \varsigma_t + \beta_2 \log(i/k)_{t-1} + \beta_3 \varsigma_t \cdot \log(i/k)_{t-1} + \varepsilon_t$$

Variable	Estimate	
Constant	β_0	0.419***
Credit tightness	β_1	-0.038***
Lagged firm investment	β_2	-0.034***
... interacted with credit tightness	β_3	0.125***
R^2	0.44	

Note: Dependent variable is i/k of the firm. $\varsigma_t = 1$ in periods of high borrowing cost, 0 otherwise. Controls such as capital stock and firm cash flow are included but not reported. *, ** and *** indicate significance at the 10, 5 and 1 percent level respectively.

In Table 10 we instead focus on investment spikes, having showed earlier that they were the main contributor to aggregate investment empirically. All regressions are linear probability models. In the first panel, the dependent variable is a dummy equal to 1 if the investment-to-capital ratio in plant A is greater than 15%, our threshold for a spike. Both output levels for plants A and B have a positive impact on the probability of a spike, in line with what we found earlier. Our focus, however, is on the spike indicator for plant B: very clearly, the spike activity in one plant does matter for the probability of a spike in the other. In periods of tight credit ($\varsigma_t = 1$), the probability of observing an investment spike in plant A is 0.5 percentage point lower if plant B is already spiking.

Another way to confirm this finding is to see whether occurrences of double spikes (i.e. spikes in both plants) is more or less likely when external funds are more costly. Conditional on firm cash flow and capital stock (coefficients not reported), the second panel of the same table shows that the probability of observing both plants spiking in the same period basically falls to zero when credit is tight.

Finally, the last regression in Table 10 revisits the question of staggered investment across plants. We regress the spike indicator for plant A on its own lag, the spike indicator for plant B at time t as well as its lagged value, both by itself and multiplied by the credit dummy, ς_t . Our focus is on this last variable: if the firm is more likely to stagger investment activity when borrowing is costly, the coefficient on the interaction term should be positive. This is what we obtain, with a

value of 0.036, though the effect seems to be relatively small.

All these results seem to indicate that the firm-level financial frictions, whether in the form of fixed or quadratic costs, alter significantly plant-level investment dynamics. In summary, the firm optimises the timing of investment spikes, making sure that joint spikes are avoided in order to minimise the need to borrow in a given period. Instead, following a positive firm-level shock affecting both plants, one plant spikes immediately while the other waits one period to invest. This, in turn, makes investment aggregated at the firm level more serially correlated in periods of tight credit. In the next section, we investigate whether these predictions are borne out empirically using a proxy for financial conditions.

TABLE 10: SPIKE ($i/k > 0.15$) REGRESSIONS

Dependent variable		$\mathbf{I}\{(i/k)_{At} > 0.15\}$
Constant		6.261***
Credit tightness	ς_t	-0.039***
Output plant A	$\log(y/k)_{At}$	0.504***
Output plant B	$\log(y/k)_{Bt}$	0.330***
Spike plant B	$\mathbf{I}\{(i/k)_{Bt} > 0.15\}$	-0.254***
	$\varsigma_t \cdot \mathbf{I}\{(i/k)_{Bt} > 0.15\}$	-0.261***
R^2		0.44
Dependent variable		$\mathbf{I}\{(i/k)_{At}, (i/k)_{Bt} > 0.15\}$
Constant		0.069***
Credit tightness	ς_t	-0.075***
R^2		0.13
Dependent variable		$\mathbf{I}\{(i/k)_{At} > 0.15\}$
Constant		6.651***
Credit tightness	ς_t	-0.108***
Spike plant B	$\mathbf{I}\{(i/k)_{Bt} > 0.15\}$	-0.478***
Lag spike plant A	$\mathbf{I}\{(i/k)_{At-1} > 0.15\}$	-0.048***
Lag spike plant B	$\mathbf{I}\{(i/k)_{Bt} > 0.15\}$	0.061**
	$\varsigma_t \cdot \mathbf{I}\{(i/k)_{Bt} > 0.15\}$	0.036**
R^2		0.48

Note: Dependent variable is spike dummy. $\varsigma_t = 1$ in periods of high borrowing cost, 0 otherwise. *, ** and *** indicate significance at the 10, 5 and 1 percent level respectively.

4 Credit Constraints and Investment: Empirical Evidence

4.1 Credit Constraints

In order to determine whether the predictions of the model with time-varying financial frictions are in line with what we observe in the data, we need some measure of financial conditions that covers a long enough time period. The National Financial Conditions Index (NFCI) from the Federal Reserve Bank of Chicago seems well suited to our purposes. The NFCI is a weighted average of a large number variables of financial activity, relative to their means. The index is therefore centered around zero by construction. We will be using the Adjusted NFCI (ANFCI), a version that isolates the component of financial conditions that is orthogonal to current economic conditions, allowing us to focus solely on the impact of fluctuations in credit tightness without worrying about endogeneity. Figure 7 shows the evolution of both the raw and adjusted NFCI.

FIGURE 7: CREDIT TIGHTNESS IN THE U.S. ECONOMY

Note: Annualised time series of the National Financial Conditions Index in the raw version (dashed line) and the version that is adjusted for endogenous responses of financial indicators to non-financial shocks (solid line). Shaded areas are NBER recessions.

As one can see, the adjusted indicator lines nicely with NBER recessions in the 1970's, 1980's and in 2008/09 while it does not increase before or during the 1991 and the 2001 recession – that financial indicators such as credit spreads did not catch those recessions is a well-known fact and we follow the literature to interpret these as non-financial recessions. In the subsequent analysis, we will use the ANFCI time series as an exogenous variable that constrains a firm's ability to borrow funds for investment purposes.

4.2 Investment Dynamics of Plants, of Firms and Within Firms

We first focus on the time series properties of investment rates at the plant level, the firm level and the joint investment dynamics of plants within firms. Table 7 is our benchmark.

4.2.1 Autocorrelation of plant-level investment

We start by examine the autocorrelation of investment at the plant and the firm level. Table 7 shows how plant-level investment becomes less and firm-level investment becomes more autocorrelated in the model. The latter is a response of firms where costly external credit induces the firm to smooth borrowing and thus investment. Predictions about the autocorrelation of plant-level investment are not as sharp and that is reflected in the data: Table 11 displays the estimates of regressing plant-level investment on the credit constraints indicator (denoted by ς_t). The estimates from the panel regression (our preferred specification) are not significant. A simple pooled OLS regression indicates that in tight credit times investment becomes more autocorrelated.

TABLE 11: CREDIT CONSTRAINTS AND PLANT- AND FIRM-LEVEL INVESTMENT

left panel (plant): $(i/k)_{nt} = \beta_0 + \beta_1(i/k)_{nt-1} + \beta_2\varsigma_t \cdot (i/k)_{nt-1} + \beta_3\varsigma_t + \varepsilon_{nt}$					
right panel (firm): $(i/k)_{jt} = \beta_0 + \beta_1(i/k)_{jt-1} + \beta_2\varsigma_t \cdot (i/k)_{jt-1} + \beta_3\varsigma_t + \varepsilon_{jt}$					
Coefficient	OLS	Panel	Coefficient	OLS	Panel
β_1	0.1590*** (0.0471)	-0.1571*** (0.0638)	β_1	0.1283*** (0.0565)	0.0356 (0.0614)
β_2	0.0145** (0.0069)	0.0109 (0.0095)	β_2	0.0264*** (0.0074)	0.0288*** (0.0075)
β_3	0.0004 (0.0008)	0.0017 (0.0011)	β_3	0.0003 (0.0007)	0.0003 (0.0007)
Controls	Yes	Yes	Controls	Yes	Yes
N	49k	17k	N	24k	9k

The predictions about firm-level investment are much sharper. Table 11 reveals that investment at the firm level significantly becomes more autocorrelated: A standard deviation to credit tightness almost doubles the autocorrelation. The effect at the firm level is much stronger than any change at the plant level where the autocorrelation increases only by 10%. Interestingly, tight credit itself does not significantly lower investment levels as the estimate of β_3 is not significantly negative as one may expect – and neither it is with the plant regression in Table 11.

4.2.2 Investment Correlation Within Firms

A particularly sharp implication of credit constraints is that investment across the two plants within the firm falls; Table 7 shows in fact that it may even become negative. This obviously reflects the fact that in times of tight credit the firm needs to scale back investment in general. So if one plants invests a lot, the the other one probably suffers when credit is tight.

Table 12 confirms this prediction of the model. The estimates of β_2 are significantly negative across both panel and OLS regressions. The estimate imply that a doubling of investment in the other plant lowers the investment rate in the other plant by two percentage points.

TABLE 12: CREDIT CONSTRAINTS AND WITHIN-FIRM INVESTMENT

$$(i/k)_{At} = \beta_0 + \beta_1(i/k)_{Bt} + \beta_2\varsigma_t \cdot (i/k)_{Bt} + \varepsilon_{nt}$$

Coefficient	OLS	Panel
β_1	0.0288*** (0.0017)	0.0275*** (0.0016)
β_2	-0.0035** (0.0014)	-0.0031** (0.0013)
Controls	Yes	Yes
N	66k	66k

4.2.3 Probability of investment spikes and spike size

We look at the likelihood of investment spikes in episodes of tight credit. It’s not obvious whether there will be more or less investment spikes when credit is tight. On the one hand, spikes will happen less often because tight credit limits overall investment resources. Then, tight credit results in a “lumpiness” effect because investment looks more lumpy. On the other hand, tight credit makes the firm smooth its borrowing so that investment spikes become more frequent. If that latter outcome prevails, one would expect investment spikes to become smaller; we label this latter effect the “smoothing effect.” We test both of these possible predictions and display the results of this regression in Table 13. The results are overall weak and borderline significant. But if at all, one sees that the probability of investment spikes increases and the level of investment spikes decreases significantly.

TABLE 13: CREDIT CONSTRAINTS AND INVESTMENT SPIKES

$$\mathbf{I}\{(i/k)_{nt} > 0.15\} = \beta_0 + \beta_1 \varsigma_t + \beta_2 \varsigma_t \cdot \mathbf{I}\{\text{Public}\} + \varepsilon_{nt}$$

$$(i/k)_{nt} = \tilde{\beta}_0 + \tilde{\beta}_1 \varsigma_t + \tilde{\beta}_2 \varsigma_t \cdot \mathbf{I}\{\text{Public}\} + \varepsilon_{nt} \quad \forall n, \text{ s.t. } (i/k)_{nt} > 0.15$$

Coefficient	OLS	Panel	Coefficient	OLS	Panel
β_0	0.0411*** (0.0131)	0.0519*** (0.0092)	$\tilde{\beta}_0$		
β_1	0.0002 (0.0012)	0.0003 (0.0011)	$\tilde{\beta}_1$		negative (significant)
β_2	-0.0055 (0.0067)	-0.0004 (0.0066)	$\tilde{\beta}_2$	positive (significant) $\Rightarrow$ publicly traded firms unaffected	
Controls	Yes	Yes	Controls	Yes	Yes
N	66k	66k	N	66k	66k

4.2.4 Joint distribution of investment spikes within firms

While predictions about single investment spikes are not conclusive, the model has fairly strong predictions about the *joint* distribution of investment spikes within firms. When credit is tight, a firm cannot allow both of its plants to undergo an investment spike in the same period – at least not if it’s financially constrained. As a consequence, the likelihood to see both plants undergoing an investment spike should drop significantly when credit is dear. This is clearly borne out in the data as Table 7 shows. An it does show up in the data as well: When credit gets tight, the likelihood to undergo an investment spike when the other plant undergoes one is 2% lower than when credit is loose. In that latter scenario, we can expect that one plant spiking raises the likelihood of the

other one spiking by 18% – probably reflecting a positive firm-specific productivity shock.

Note that this logic of reduced simultaneous spiker plants only applies to financially constrained firms. If a firm wasn’t financially constrained, we would expect the coefficient β_2 to be zero. To test for that hypothesis, we include an interaction term of credit tightness and the other plant spiking with a dummy variable that indicates whether or not the firm is publicly traded or not. The idea is that publicly traded firm probably are not affected by ς_t . As we can see, this is borne out in the data, albeit it’s borderline significant.

TABLE 14: CREDIT CONSTRAINTS AND FIRM-LEVEL INVESTMENT

$$\mathbf{I}\{(i/k)_{At} > 0.15\} = \beta_0 + \beta_1 \mathbf{I}\{(i/k)_{Bt} > 0.15\} + \beta_2 \varsigma_t \mathbf{I}\{(i/k)_{Bt} > 0.15\} + \beta_3 \varsigma_t \mathbf{I}\{(i/k)_{Bt} > 0.15\} \mathbf{I}\{\text{Public}\} + \varepsilon_{nt}$$

Coefficient	OLS	Panel
β_0	0.0378*** (0.0130)	0.0463*** (0.0089)
β_1	0.1846*** (0.0056)	0.1812*** (0.0056)
β_2	-0.0211*** (0.0051)	-0.0205*** (0.0051)
β_3	0.0237* (0.0132)	0.0201 (0.0135)
Controls	Yes	Yes
N	66k	40k

4.2.5 Serial correlation of single investment spikes

So what can firms that are credit constrained do if both of its plants are so productive that ideally they should btw undergo investment projects? If credit is tight and thus external finance particularly costly, then it may see no other possibility than to focus its funds on investing in one plant and postponing investment in the other plant. We call this spacing out of investment spikes “adjacent investment spikes” and test for them by regressing a dummy variable that indicates such “adjacent spike”-firms on credit conditions. The unconditional probability (without any especially tight credit) is about 10% (see Table 6) and the regression results in Table 15 tell us that this probability drops by 0.7%.⁶

⁶In this analysis, we consider firm that undergo exactly one spike today and one or two spikes last year. Instead, we should have restricted the sample to only those firms that have *at least* 2 spikes within a two year window and then see if they are more likely to stretch the at least 2 spikes out or if they choose to do it simultaneously and also consider interaction terms with the publicly traded dummy. Since this sample is smaller than the one we consider here, we see our estimates as a lower bound.

TABLE 15: CREDIT CONSTRAINTS AND SERIAL CORRELATION OF INVESTMENT SPIKES

$$\mathbf{I}\{(i/k)_{At}, (i/k)_{At-1} > 0.15\} = \beta_0 + \beta_1 \varsigma_t + \beta_2 \varsigma_t \cdot (i/k)_{Bt} + \varepsilon_{nt}$$

Coefficient	OLS	Panel
β_0	0.2148*** (0.0354)	0.1836*** (0.0274)
β_1	0.0069** (0.0030)	0.0069** (0.0032)
β_2	-0.0141 (0.0130)	-0.0141 (0.0133)
Controls	Yes	Yes
N	17k	9k

5 Conclusion

Investment is lumpy on the plant level, but smooth on the firm level. This is true for both publicly traded and privately held firms. While privately held firms seem to suffer from inefficient external capital markets, publicly traded firms seem to suffer from inefficient internal capital markets.

References

- Abel, Andrew B., and Janice C. Eberly. 1996. "Optimal Investment with Costly Reversibility." *Review of Economic Studies*, 63(4): 581–593.
- Bachmann, Rüdiger, and Christian Bayer. 2011. "Investment Dispersion and the Business Cycle." *Working Paper*.
- Caballero, Ricardo J., Eduardo M. R. A. Engel, and John C. Haltiwanger. 1995. "Plant-Level Adjustment and Aggregate Dynamics." *Brookings Papers on Economic Activity*, 1995(2): 1–54.
- Clementi, Gianluca, Julia K. Thomas, Aubhik Khan, and Berardino Palazzo. 2011. "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity." *Working Paper*.
- Cooper, Russell, John C. Haltiwanger, and Laura Power. 1999. "Machine Replacement and the Business Cycle: Lumps and Bumps." *American Economic Review*, 89(4): 921–946.
- Cooper, Russell W., and John C. Haltiwanger. 2006. "On the Nature of Capital Adjustment Costs." *Review of Economic Studies*, 73(3): 611–633.
- Davis, Steven J., John C. Haltiwanger, Ron Jarmin, and Javier Miranda. 2006. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." *NBER Macroeconomics Annual*, 21: 107–179.
- Doms, Mark, and Timothy Dunne. 1998. "Capital Adjustment Patterns in Manufacturing Plants." *Review of Economic Dynamics*, 1(2): 409–429.
- Eberly, Janice, Sergio Rebelo, and Nicolas Vincent. 2012. "What Explains the Lagged Investment Effect?" *Journal of Monetary Economics*, 59(4): 370–380.
- Eisfeldt, Andrea L., and Adriano A. Rampini. 2006. "Capital Reallocation and Liquidity." *Journal of Monetary Economics*, 53(3): 369–399.
- Eisfeldt, Andrea L., and Adriano A. Rampini. 2008. "Managerial Incentives, Capital Reallocation, and the Business Cycle." *Journal of Financial Economics*, 87(1): 177–199.
- Eisfeldt, Andrea L., and Tyler Muir. 2013. "Aggregate Issuance and Savings Waves." *Working Paper*.
- Fazzari, Steven M., R. Glenn Hubbard, and Bruce C. Petersen. 1988. "Financing Constraints and Corporate Investment." *Brookings Papers on Economic Activity*, 1988(1): 141–206.
- Gertner, Robert, David Scharfstein, and Jeremy Stein. 1994. "Internal versus External Capital Markets." *Quarterly Journal of Economics*, 109(4): 1211–1230.
- Gilchrist, Simon, and Charles P. Himmelberg. 1995. "The Role of Cash Flow in Reduced-Form Investment Equations." *Journal of Monetary Economics*, 36(3): 541–572.
- Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review*, 102(3): 1692–1720.

- Giroud, Xavier.** forthcoming. “Proximity and Investment: Evidence from Plant-Level Data.” *Quarterly Journal of Economics*.
- Giroud, Xavier, and Holger M. Müller.** 2012. “Capital and Labor Reallocation Inside Firms.” *Working Paper*.
- Gomes, João F.** 2001. “Financing Investment.” *American Economic Review*, 91(5): 1263–1285.
- Gomes, João F., Amir Yaron, and Lu Zhang.** 2006. “Asset Pricing Implications of Firms’ Financing Constraints.” *Review of Financial Studies*, 19(4): 1321–1356.
- Gourio, François, and Anil K. Kashyap.** 2007. “Investment Spikes: New Facts and a General Equilibrium Exploration.” *Journal of Monetary Economics*, 54: 1–22.
- Hennessy, Christopher A., and Toni M. Whited.** 2005. “Debt Dynamics.” *Journal of Finance*, 60(3): 1129–1165.
- Kehrig, Matthias.** 2013. “The Cyclicalities of Productivity Dispersion.” *Working Paper*.
- 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.
- Lamont, Owen A.** 1997. “Cash Flow and Investment: Evidence from Internal Capital Markets.” *Journal of Finance*, 52(1): 83–109.
- Malenko, Andrey.** 2012. “Optimal Design of Internal Capital Markets.” *Working Paper*.
- Papageorgiou, Theodore.** 2011. “Large Firms and Internal Labor Markets.” *Working Paper*.
- Schoar, Antoinette.** 2002. “Effects of Corporate Diversification on Productivity.” *Journal of Finance*, 57(6): 2379–2403.
- Stein, Jeremy C.** 1997. “Internal Capital Markets and the Competition for Corporate Resources.” *Journal of Finance*, 52(1): 111–133.
- Thomas, Julia K.** 2002. “Is Lumpy Investment Relevant for the Business Cycle?” *Journal of Political Economy*, 110(3): 508–534.

A.1 Productivity Dispersion

Productivity dispersion: Contribution of between-firm and within-firm dispersion

Legend:

- Total dispersion (Blue line)
- Between-firm dispersion (Red line)
- Within-firm dispersion (Green line)

The graph shows that total productivity dispersion is the sum of between-firm and within-firm dispersion. Both components show significant fluctuations over the 37-year period, with total dispersion generally higher than the individual components.

The results of the decomposition are displayed in Figure 8. Similar to investment of dispersion productivity dispersion within the firm is the major component of overall productivity dispersion. Both between- and within-firm productivity dispersion are cyclical. Within-firm productivity dispersion, however, plays a more significant role. About 60% of overall productivity dispersion originate within the firm. Moreover the cyclical changes in overall productivity dispersion are clearly driven by within-farm dispersion. While the cyclicity of within-firm dispersion is strong both for productivity and investment, the correlation with the cycle is different: investments dispersion is procyclical, but productivity dispersion is countercyclical. This tension between the benefits of investing (the marginal product of capital) and the cost of investment (the investment rate) has been noted in firm-level data (see for example Eisefeldt and Rampini (2006)). Here we document a much stronger pattern within the firm. Eisefeldt and Rampini (2006, 2008) have proposed countercyclical adjustment costs to resolve this puzzle. Our new findings point to frictions within the firm that play at least as important a role as frictions at the firm level.

A.2 Time variation in the productivity of spikers

In the above section on the investment dispersion within firms, we established that most of the dispersion comes from plants that undergo large investment projects. Are these plants chosen for large investment projects because they are more productive? The raw data – displayed in Table 16 – seem to suggest so.

TABLE 16: SUMMARY STATISTICS OF PLANTS WITH LARGE VS. SMALL INVESTMENT

	Average	Boom	Recession
i/k ($i/k > 0.15$):	0.41	0.42	0.39
i/k ($i/k \leq 0.15$):	0.039	0.041	0.037
$\log(y/k) : (x > 0.2)$	0.35	0.43	0.26
$\log(y/k) : (x < 0.2)$	0.03	0.09	-0.03