ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

https://www.econometricsociety.org/

Econometrica, Vol. 88, No. 6 (November, 2020), 2473–2502

FINANCIAL HETEROGENEITY AND THE INVESTMENT CHANNEL OF MONETARY POLICY

PABLO OTTONELLO

Department of Economics, University of Michigan and NBER

THOMAS WINBERRY

Booth School of Business, University of Chicago and NBER

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website http://www.econometricsociety.org or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

Econometrica, Vol. 88, No. 6 (November, 2020), 2473–2502

FINANCIAL HETEROGENEITY AND THE INVESTMENT CHANNEL OF MONETARY POLICY

PABLO OTTONELLO
Department of Economics, University of Michigan and NBER

THOMAS WINBERRY
Booth School of Business, University of Chicago and NBER

We study the role of financial frictions and firm heterogeneity in determining the investment channel of monetary policy. Empirically, we find that firms with low default risk—those with low debt burdens and high "distance to default"— are the most responsive to monetary shocks. We interpret these findings using a heterogeneous firm New Keynesian model with default risk. In our model, low-risk firms are more responsive to monetary shocks because they face a flatter marginal cost curve for financing investment. The aggregate effect of monetary policy may therefore depend on the distribution of default risk, which varies over time.

KEYWORDS: Monetary policy, firm heterogeneity, investment, financial frictions, New Keynesian.

1. INTRODUCTION

AGGREGATE INVESTMENT is one of the most responsive components of GDP to monetary shocks. Our goal in this paper is to understand the role of financial frictions in determining this investment channel of monetary policy. Given the rich heterogeneity in financial positions across firms, a key question is: which firms are the most responsive to changes in monetary policy? The answer to this question is theoretically ambiguous. On the one hand, financial frictions generate an upward-sloping marginal cost curve for investment, which dampens the response of investment to monetary policy for firms that are more severely affected by financial frictions. On the other hand, monetary policy may flatten out this marginal cost curve, for example, by increasing cash flows or improving collateral values, which amplifies the response of investment for affected firms. This latter view is the conventional wisdom in the literature, informed by applying the financial accelerator logic across firms.

We address the question of which firms respond the most to monetary policy using new cross-sectional evidence and a heterogeneous firm New Keynesian model. Our empirical work combines monetary shocks, measured using the high-frequency event-study approach, with quarterly Compustat data. We find that investment done by firms with low default risk is significantly and robustly more responsive to monetary policy than investment done by firms with high default risk. Motivated by this evidence, our model embeds

Pablo Ottonello: pottonel@umich.edu
Thomas Winberry: Thomas.Winberry@chicagobooth.edu

We thank four anonymous referees, Andy Abel, Adrien Auclert, Cristina Arellano, Neele Balke, Paco Buera, Simon Gilchrist, Nils Gornermann, Chris House, Erik Hurst, Alejandro Justiniano, Greg Kaplan, Rohan Kekre, Aubhik Khan, John Leahy, Yueran Ma, Alisdair McKay, Fabrizio Perri, Felipe Schwartzman, Linda Tesar, Julia Thomas, Joe Vavra, Ivan Werning, Toni Whited, Christian Wolf, Arlene Wong, and Mark Wright for helpful conversations. We thank Alberto Arredondo, Mike Mei, Richard Ryan, Yuyao Wang, Liangjie Wu, and especially Samuel Stern for excellent research assistance. This research was funded in part by the Michigan Institute for Teaching and Research in Economics. Ottonello gratefully acknowledges financial support from the University of Michigan and Winberry gratefully acknowledges financial support from Chicago Booth.

© 2020 The Econometric Society https://doi.org/10.3982/ECTA15949

a heterogeneous firm investment model with default risk into the benchmark New Keynesian environment and studies the effect of a monetary shock. In our calibrated model, firms with low default risk are more responsive to monetary policy, similar to our empirical estimates. We perform a simple calculation to show how these heterogeneous responses imply that the effect of monetary policy may become smaller when default risk in the economy is high. At the same time, we find that all firms affected by default risk in our model are *more* responsive to monetary policy than they would be in a version of our model without any default risk at all, consistent with Bernanke, Gertler, and Gilchrist (1999).

Our baseline empirical specification estimates how the semielasticity of firm investment with respect to a monetary policy shock depends on two measures of the firm's default risk: its leverage ratio and its "distance to default" (which estimates the probability of default from the values of equity and liabilities). We control for firm fixed effects to capture permanent differences across firms and control for sector-by-quarter fixed effects to capture differences in how sectors respond to aggregate shocks. Conditional on our set of controls, leverage is negatively correlated with distance to default and credit rating, and distance to default is positively correlated with credit rating. Therefore, we view low leverage and high distance to default as proxies for low default risk.

We find that having one standard deviation lower leverage implies that a firm is approximately one-fourth more responsive to monetary policy and that having one standard deviation higher distance to default implies that the firm is one-half more responsive. These differences across firms persist for up to 3 years after the shock and imply large differences in accumulated capital over time. Consistent with the idea that default risk drives these heterogeneous responses, borrowing costs and the use of external finance increase by less for high-risk firms than for low-risk firms following a monetary expansion.

In order to interpret these empirical results, we embed a model of heterogeneous firms facing default risk into the benchmark New Keynesian framework. These firms invest in capital using either internal funds or external borrowing; they can default on their debt, leading to an external finance premium. There is also a group of "retailer" firms with sticky prices, generating a New Keynesian Phillips curve linking nominal variables to real outcomes. We calibrate the model to match key features of firms' investment, borrowing, and lifecycle dynamics in the microdata. Our model generates realistic behavior along nontargeted dimensions of the microdata, and the peak responses of aggregate investment, output, and consumption to a monetary policy shock are broadly in line with the peak responses estimated in the data by Christiano, Eichenbaum, and Evans (2005).

We simulate a panel of firms from our calibrated model and find that firms with low measured default risk are more responsive to monetary policy, as in the data. These heterogeneous responses reflect how monetary policy directly changes the expected return on capital, which drives the response of low-risk firms, and indirectly changes cash flows and recovery values, which drive the response of the high-risk firms. Since low-risk firms are more responsive overall, our empirical results indicate that the direct effects of monetary policy dominate the indirect ones.

Finally, we quantify how changes in the distribution of default risk may alter the aggregate effect of monetary policy by fixing the firm-level response to monetary shocks while varying the initial distribution of firms. We find that a monetary shock generates an approximately 30% smaller change in the aggregate capital stock starting from a distribution with 50% less net worth than the steady state distribution. This calculation suggests a potentially important source of time-variation in monetary transmission: monetary policy is less powerful when default risk is higher.

Related Literature. Our paper contributes to five strands of literature. The first studies the role of financial frictions in the transmission of monetary policy to the aggregate economy. Bernanke, Gertler, and Gilchrist (1999) embed the financial accelerator in a representative firm New Keynesian model; we build on Bernanke, Gertler, and Gilchrist (1999)'s framework to include firm heterogeneity. Consistent with their results, we find that the response of aggregate investment to a monetary shock is larger in our model than in a model without financial frictions at all. However, among the 99.4% of firms affected by financial frictions in our model, those with low default risk are more responsive to a monetary shock than those with high default risk, creating the potential for distributional dependence.

Second, we contribute to the literature that studies how the effect of monetary policy varies across firms by showing that firms with low default risk are more responsive to monetary policy. Other studies argue that the firm-level response also depends on size (Gertler and Gilchrist (1994)), liquidity (Jeenas (2019)), or age (Cloyne et al. (2018)). Online Appendix C shows that our results are robust to controlling for these other firm characteristics. Our results do not necessarily contradict these other studies; instead, we simply study different features of the data.

Third, we contribute to the literature which incorporates microlevel heterogeneity into the New Keynesian model. To date, this literature has focused on how household-level heterogeneity affects the consumption channel of monetary policy; see, for example, McKay, Nakamura, and Steinsson (2016); Kaplan, Moll, and Violante (2018); Auclert (2019); or Wong (2019). We instead explore the role of firm-level heterogeneity in determining the investment channel of monetary policy.[1] In contrast to the consumption channel, we find that both direct and indirect effects of monetary policy play a quantitatively important role in driving the investment channel. The direct effect of changes in the real interest rate are smaller for households because they have a consumption-smoothing motive which firms lack.

Fourth, we contribute to a growing literature which argues that monetary policy is less effective in recessions by suggesting that changes in the distribution of default risk are another reason monetary policy may become less effective. Tenreyro and Thwaites (2016) estimated a nonlinear time-series model and find that monetary policy shocks have a smaller impact on real economic activity in recessions than in normal times. Vavra (2013) and McKay and Wieland (2019) provided models in which monetary policy is less powerful in recessions due to changes in the distribution of price adjustment or durable expenditures.

Finally, we contribute to the literature which studies the role of financial heterogeneity in determining the business cycle dynamics of aggregate investment by introducing sticky prices and studying the effect of monetary policy shocks. Our model of firm-level investment builds heavily on Khan, Senga, and Thomas (2016), who study the effect of financial shocks in a flexible price model. We extend the model to include capital quality shocks and a time-varying price of capital in order to generate variation in lenders' recovery value of capital, as in the financial accelerator literature.

[1] Reiter, Sveen, and Weinke (2013) showed that a model with firm heterogeneity and fixed capital adjustment costs generates a counterfactually large and short-lived response of investment to monetary policy because, conditional on adjusting, investment is extremely interest-sensitive in their model. We dampen the interest-sensitivity of investment using financial frictions and convex adjustment costs to aggregate capital. Koby and Wolf (2020) dampened the interest-sensitivity using convex adjustment costs at the firm level and found that the fixed costs generate state-dependent responses to monetary policy.

Road Map. Our paper is organized as follows. Section 2 provides empirical evidence that the firm-level response to monetary policy varies with default risk. Section 3 develops our heterogeneous firm New Keynesian model to interpret this evidence. Section 4 provides a theoretical characterization of the channels through which monetary policy drives investment in our model. Section 5 then calibrates the model and verifies that it is consistent with key features of the joint distribution of investment and leverage in the microdata. Section 6 uses the model to study the monetary transmission mechanism. Section 7 concludes. The online Appendix may be found at Ottonello and Winberry (2020) and the Supplementary Materials may be found in the replication file.

2. EMPIRICAL RESULTS

We document that firms with low default risk—proxied by low debt burdens and high distance to default—are significantly more responsive to changes in monetary policy than other firms in the economy.

2.1. *Data Description*

Our sample combines monetary policy shocks with quarterly Compustat data.

Monetary Policy Shocks. We measure monetary shocks using the high-frequency, event-study approach pioneered by Cook and Hahn (1989). Following Gurkaynak, Sack, and Swanson (2005) and Gorodnichenko and Weber (2016), we construct our shock $\varepsilon_t^{\mathrm{m}}$ as

$$\varepsilon_t^{\mathrm{m}} = \tau(t) \times (\mathtt{ffr}_{t+\Delta_+} - \mathtt{ffr}_{t-\Delta_-}), \tag{1}$$

where t is the time of the monetary announcement, $\mathtt{ffr}_t$ is the implied Fed Funds Rate from a current-month Federal Funds future contract at time t, Δ_+, and Δ_- control the size of the time window around the announcement, $\tau(t) \equiv \frac{\tau_m^n(t)}{\tau_m^n(t)-\tau_m^d(t)}$ is an adjustment for the timing of the announcement within the month, $\tau_m^d(t)$ denotes the day of the meeting in the month, and $\tau_m^n(t)$ the number of days in the month. We focus on a window of $\Delta_- =$ fifteen minutes before the announcement and $\Delta_+ = 45$ minutes after the announcement. Our shock series begins in January 1990, when the Fed Funds futures market opened, and ends in December 2007, in order to focus on conventional monetary policy. During this time, there were 164 shocks with a mean of approximately zero and a standard deviation of 9bp.

We time aggregate the high-frequency shocks to the quarterly frequency in order to merge them with our firm-level data. We construct a moving average of the raw shocks weighted by the number of days in the quarter after the shock occurs (see Supplemental Materials A for details). Our time aggregation strategy ensures that we weight shocks by the amount of time firms have had to react to them. Table I indicates that these "smoothed" shocks have similar features to the original high-frequency shocks. For robustness, we also use the alternative time aggregation of simply summing all the shocks that occur within the quarter, as in Wong (2019). Table I shows that the moments of these alternative shocks do not significantly differ from the moments of the smoothed shocks, and Supplemental Materials B shows that our main results are robust to using this alternative form of time aggregation.

TABLE I

SUMMARY STATISTICS OF MONETARY POLICY SHOCKS[a]

	High Frequency	Smoothed	Sum
Mean	−0.0185	−0.0429	−0.0421
Median	0	−0.0127	−0.00509
S.D.	0.0855	0.108	0.124
Min	−0.463	−0.480	−0.479
Max	0.152	0.233	0.261
Observations	164	71	72

[a]Summary statistics of monetary policy shocks for the period 1/1/1990 to 12/31/2007. “High frequency” shocks are estimated using the event study strategy in (1). “Smoothed” shocks are time aggregated to a quarterly frequency using the weighted average described in Supplemental Materials A. “Sum” refers to time aggregating by simply summing all shocks within a quarter.

Firm-Level Variables. We draw firm-level variables from quarterly Compustat, a panel of publicly listed U.S. firms. Compustat satisfies three key requirements for our study: it is quarterly, a high enough frequency to study monetary policy; it is a long panel, allowing us to use within-firm variation; and it contains rich balance-sheet information, allowing us to construct our key variables of interest. To our knowledge, Compustat is the only U.S. dataset that satisfies these three requirements. The main disadvantage of Compustat is that it excludes privately held firms.[2] In Section 5, we calibrate our economic model to match a broad sample of firms, not just those in Compustat.

Our main measure of investment is $\Delta \log k_{jt+1}$, where k_{jt+1} is the book value of the tangible capital stock of firm j at the end of period t. We use two measures of a firm's financial position to proxy for default risk. First, we measure leverage ℓ_{jt} as the firm's debt-to-asset ratio, where debt is the sum of short term and long term debt and assets is the book value of assets. Second, we measure the firm's distance to default dd_{jt} following Gilchrist and Zakrajšek (2012). Distance to default dd_{jt} has been shown by Schaefer and Strebulaev (2008) to account well for variation in corporate bond prices due to default risk and is widely used in the finance industry. In order to validate these proxies, we correlate them with credit rating cr_{jt}, measured as S&P's long-term issue rating of the firm. Supplemental Materials A provides details of our data construction.

Panel (a) of Table II presents simple summary statistics of the final sample used in our analysis. The mean distance to default implies that a 6 standard deviation shock over a given year will drive the average firm to default, in line with Gilchrist and Zakrajšek (2012). We winsorize our sample at the top and bottom 0.5% of observations of investment, leverage, and distance to default in order to ensure our results are not driven by outliers.

Panel (b) of Table II shows the correlation structure of leverage, distance to default, and credit rating. Higher leverage is positively correlated with a smaller distance to default and a lower credit rating, indicating that higher debt burdens are associated with higher default risk. Firms with higher distance to default also have higher credit ratings, validating our interpretation of distance to default. Panel (c) of Table II shows that these

[2]The main attractive alternatives, covering a much broader set of firm sizes than Compustat, are the datasets constructed in Crouzet and Mehrotra (2020) (using data from the Quarterly Financial Reports) and in Dinlersoz et al. (2018) (combining data from the U.S. Census Longitudinal Business Database, Orbis, and Compustat). However, the dataset in Crouzet and Mehrotra (2020) only follows small firms for 8 quarters, which limits the ability to use within-firm variation, while the dataset in Dinlersoz et al. (2018) contains data for small firms at an annual frequency.

TABLE II
SUMMARY STATISTICS OF FIRM-LEVEL VARIABLES[a]

(a) Marginal Distributions

	$\Delta \log k_{jt+1}$	ℓ_{jt}	dd_{jt}	$\mathbb{1}\{\mathrm{cr}_{jt} \geq A\}$
Mean	0.004	0.263	5.788	0.025
Median	−0.004	0.201	4.742	0.000
S.D.	0.091	0.348	5.082	0.156
95th Percentile	0.128	0.719	15.213	0.000
Observations	343,276	343,276	238,531	343,276

(b) Correlation Matrix (raw variables)[b]

	ℓ_{jt}	dd_{jt}	$\mathbb{1}\{\mathrm{cr}_{jt} \geq A\}$
ℓ_{jt}	1.00		
dd_{jt}	−0.39	1.00	
	(0.00)		
$\mathbb{1}\{\mathrm{cr}_{jt} \geq A\}$	−0.02	0.20	1.00
	(0.00)	(0.00)	

(c) Correlation matrix (residualized)[b]

	ℓ_{jt}	dd_{jt}	$\mathbb{1}\{\mathrm{cr}_{jt} \geq A\}$
ℓ_{jt}	1.00		
dd_{jt}	−0.28	1.00	
	(0.00)		
$\mathbb{1}\{\mathrm{cr}_{jt} \geq A\}$	−0.01	0.05	1.00
	(0.00)	(0.00)	

[a]Summary statistics of firm-level variables for the period 1983q3 to 2014q4. $\Delta \log k_{jt+1}$ is the change in the capital stock, ℓ_{jt} is the ratio of total debt to total assets, dd_{jt} is the firm's distance to default, and $\mathbb{1}\{\mathrm{cr}_{jt} \geq A\}$ is an indicator variable for whether the firm's credit rating is above an A. Panel (a) computes summary statistics of these variables in our sample before winsorization, Panel (b) computes their pairwise correlations, and Panel (c) computes pairwise correlations of the residuals from the regression $x_{jt} = \alpha_j + \alpha_{st} + \mathbf{\Gamma}' Z_{jt-1} + e_{jt}$, where $x_{jt} \in \{\ell_{jt}, \mathrm{dd}_{jt}, \mathbb{1}\{\mathrm{cr}_{jt} \geq A\}\}$, α_j is a firm fixed effect, α_{st} is a sector-by-quarter fixed effect, and Z_{jt-1} is a vector of firm-level controls containing sales growth, size, current assets as a share of total assets, and an indicator for fiscal quarter.

[b] p-values in parentheses.

correlations are all also true conditional on the controls in our baseline regression specification (2) below.

2.2. *Heterogeneous Responses to Monetary Policy*

Specification. We estimate variants of the baseline empirical specification

$$\Delta \log k_{jt+1} = \alpha_j + \alpha_{st} + \beta\big(x_{jt-1} - \mathbb{E}_j[x_{jt}]\big)\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}' Z_{jt-1} + e_{jt}, \tag{2}$$

where α_j is a firm j fixed effect, α_{st} is a sector s by quarter t fixed effect, $\varepsilon_t^{\mathrm{m}}$ is the quarterly monetary policy shock, $x_{jt} \in \{\ell_{jt}, \mathrm{dd}_{jt}\}$ is the firm's leverage ratio or distance to default, $\mathbb{E}_j[x_{jt}]$ is the average value of x_{jt} for a given firm over the sample, Z_{jt-1} is a vector of controls, and e_{jt} is a residual.[3] Our main coefficient of interest is β, which measures how the semielasticity of investment $\Delta \log k_{jt+1}$ with respect to monetary shocks $\varepsilon_t^{\mathrm{m}}$ depends on the within-firm variation in the financial position $x_{jt} - \mathbb{E}_j[x_{jt}]$ (we discuss the rationale for using within-firm variation in financial position below). We do not estimate the specification with credit ratings $x_{jt} = \mathrm{cr}_{jt}$ because the within-firm variation in credit ratings is limited.

[3]The sectors s we consider, based on SIC codes, are: agriculture, forestry, and fishing; mining; construction; manufacturing; transportation communications, electric, gas, and sanitary services; wholesale trade; retail trade; and services. We do not include finance, insurance, and real estate, and utilities.

Throughout, we cluster standard errors two ways to account for correlation within firms and within quarters.

We control for a number of factors that may simultaneously affect investment and financial position but which are outside the scope of our economic model in Section 3. The firm fixed effects α_j capture permanent differences in investment behavior across firms and the sector-by-quarter fixed effects α_{st} capture differences in how broad sectors are exposed to aggregate shocks. The vector Z_{jt-1} includes the level of the financial position x_{jt-1}, total assets, sales growth, current assets as a share of total assets, and a fiscal quarter dummy. The vector Z_{jt-1} also includes the interaction of financial position with the previous quarter's GDP growth in order to control for differences in cyclical sensitivities across firms.[4]

Online Appendix A.1 shows that using the interaction of within-firm variation in financial position with the monetary shock $(x_{jt-1} - \mathbb{E}_j[x_{jt}])\varepsilon_t^{\mathrm{m}}$ ensures that our results are not driven by permanent heterogeneity in responsiveness across firms. This choice is motivated by our economic model in Section 3, in which firms are ex ante homogenous. In contrast, firms in the data may be ex ante heterogeneous in how they respond to monetary policy according to their financial position x_{jt}. For example, firms in risky markets may be permanently more exposed to interest rate fluctuations and also permanently more likely to default. If we had instead interacted the *level* of financial position with the monetary shock $x_{jt}\varepsilon_t^{\mathrm{m}}$, then our results would be partly determined by such permanent differences in responsiveness. By demeaning financial position within firms, $(x_{jt-1} - \mathbb{E}_j[x_{jt}])\varepsilon_t^{\mathrm{m}}$, our estimates are instead driven by how a *given* firm responds to monetary policy when it has higher or lower default risk than usual.

Results. Table III reports the results from estimating the baseline specification (2). We perform two normalizations to make the estimated coefficient β easily interpretable. First, we standardize the firm's demeaned leverage $\ell_{jt} - \mathbb{E}_j[\ell_{jt}]$ and distance to default $\mathrm{dd}_{jt} - \mathbb{E}_j[\mathrm{dd}_{jt}]$ over the entire sample, so their units are standard deviations in our sample. Second, we normalize the sign of the monetary shock $\varepsilon_t^{\mathrm{m}}$ so that a positive value corresponds to a cut in interest rates.

The first three columns in Table III show that firms with lower leverage and higher distance to default are more responsive to monetary shocks $\varepsilon_t^{\mathrm{m}}$. Column (1) implies that a firm has approximately a 0.7 units lower semielasticity of investment to monetary policy when it is one standard deviation more indebted than it typically is in our sample. Adding firm-level controls Z_{jt-1} in Column (2) does not significantly change this point estimate; therefore, we focus on specifications with firm-level controls Z_{jt-1} for the remainder of the paper. Column (3) shows that a firm has approximately a 1.1 unit higher semielasticity when it is one standard deviation further from default than usual. Column (4) shows that leverage is rendered statistically insignificant conditional on distance to default, indicating that our results are primarily driven by distance to default, which we view as our most direct measure of default risk.

[4]If the monetary shock $\varepsilon_t^{\mathrm{m}}$ is truly exogenous, then this control would be unnecessary in large samples. However, we find that the largest shocks occur at the beginning of the two recessions in our small sample. Failing to incorporate this fact may bias our results if firms with different financial positions are differentially exposed to business cycle events. Online Appendix A.2 shows that controlling for differences in cyclical sensitivities strengthens the differential responses to monetary shocks, but that our results are qualitatively robust to excluding those controls as well.

TABLE III

HETEROGENEOUS RESPONSES OF INVESTMENT TO MONETARY POLICY[a]

	(1)	(2)	(3)	(4)	(5)
leverage × ffr shock	−0.69	−0.57		−0.26	−0.14
	(0.29)	(0.27)		(0.35)	(0.58)
dd × ffr shock			1.14	1.01	1.16
			(0.41)	(0.40)	(0.47)
ffr shock					2.14
					(0.61)
Observations	219,402	219,402	151,027	151,027	119,750
R^2	0.113	0.124	0.141	0.142	0.151
Firm controls	no	yes	yes	yes	yes
Time sector FE	yes	yes	yes	yes	no
Time clustering	yes	yes	yes	yes	yes

[a]Results from estimating $\Delta \log k_{jt+1} = \alpha_j + \alpha_{st} + \beta(x_{jt-1} - \mathbb{E}_j[x_{jt}])\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}' Z_{jt-1} + e_{jt}$, where α_j is a firm fixed effect, α_{st} is a sector-by-quarter fixed effect, $x_{jt} \in \{\ell_{jt}, \mathrm{dd}_{jt}\}$ is leverage or distance to default, $\mathbb{E}_j[x_{jt}]$ is the average of x_{jt} for firm j in the sample, $\varepsilon_t^{\mathrm{m}}$ is the monetary shock, and Z_{jt-1} is a vector of firm-level controls containing x_{jt-1}, sales growth, size, current assets as a share of total assets, an indicator for fiscal quarter, and the interaction of demeaned financial position with lagged GDP growth. Standard errors are two-way clustered by firms and quarter. We have normalized the sign of the monetary shock $\varepsilon_t^{\mathrm{m}}$ so that a positive shock corresponds to a decrease in interest rates. We have standardized $(\ell_{jt} - \mathbb{E}[\ell_{jt}])$ and $(\mathrm{dd}_{jt} - \mathbb{E}[\mathrm{dd}_{jt}])$ over the entire sample. Column (5) removes the sector-quarter fixed effect α_{st} and estimates $\Delta \log k_{jt+1} = \alpha_j + \alpha_{sq} + \gamma\varepsilon_t^{\mathrm{m}} + \beta(x_{jt-1} - \mathbb{E}_j[x_{jt}])\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}_1' Z_{jt-1} + \mathbf{\Gamma}_2' Y_{t-1} + e_{jt}$, where Y_t is a vector with four lags of GDP growth, the inflation rate, and the unemployment rate.

Column (5) removes the sector-by-quarter fixed effects in order to estimate the average effect of a monetary shock:

$$\Delta \log k_{jt+1} = \alpha_j + \alpha_{sq} + \gamma\varepsilon_t^{\mathrm{m}} + \beta\big(x_{jt-1} - \mathbb{E}_j[x_{jt}]\big)\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}_1' Z_{jt-1} + \mathbf{\Gamma}_2' Y_{t-1} + e_{jt}, \quad (3)$$

where α_{sq} is a sector s by quarter q seasonal fixed effect and Y_t is a vector with four lags of GDP growth, the inflation rate, and the unemployment rate. The average investment semi-elasticity is roughly 2.[5] Hence, our interaction coefficients in the previous columns imply an economically meaningful degree of heterogeneity.

Dynamics. In order to estimate the dynamics of these differential responses across firms, we estimate the Jorda (2005)-style local projection of specification (2):

$$\log k_{jt+h} - \log k_{jt} = \alpha_{jh} + \alpha_{sth} + \beta_h\big(x_{jt-1} - \mathbb{E}_j[x_{jt}]\big)\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}_h' Z_{jt-1} + e_{jth}, \quad (4)$$

where $h \geq 1$ indexes the forecast horizon. The coefficient β_h measures how the cumulative response of investment in quarter $t + h$ to a monetary policy shock in quarter t depends on the firm's demeaned financial position $x_{jt-1} - \mathbb{E}_j[x_{jt}]$ in quarter $t - 1$. We use the cumulative change in capital on the left-hand side in order to easily assess the implications of our estimates for the capital stock itself. Online Appendix B.1 estimates a dynamic version of the specification (3) without the sector-time fixed effects α_{sth} and shows that the average firm's response is persistent, peaking 2 to 4 quarters after the shock.

[5]Assuming an annual depreciation rate of $\delta = 0.1$, this estimated coefficient implies that a one percentage point cut in the interest rate increases annualized investment by 20%, which is at the upper end of estimated user-cost elasticities in the literature (see, e.g., Zwick and Mahon (2017)).

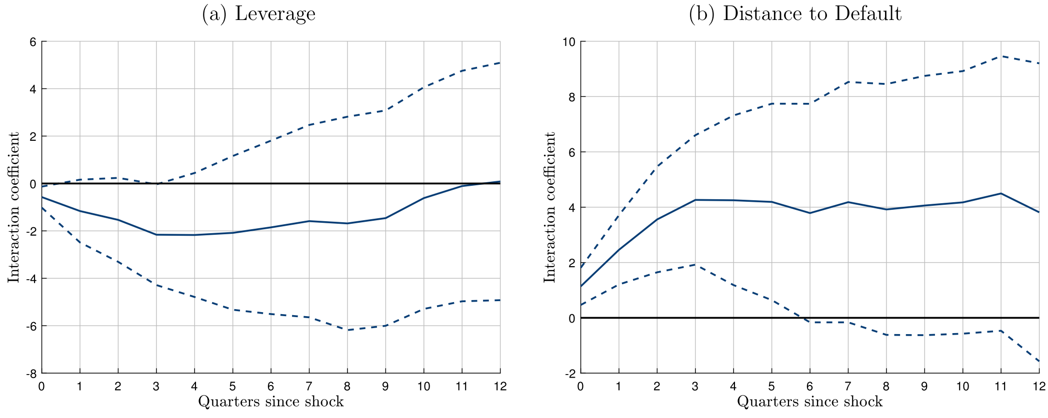

FIGURE 1.—Dynamics of differential response to monetary shocks. Notes: dynamics of the interaction coefficient between financial positions and monetary shocks over time. Reports the coefficient β_h over quarters h from $\log k_{jt+h} - \log k_{jt} = \alpha_{jh} + \alpha_{sth} + \beta_h(x_{jt-1} - \mathbb{E}_j[x_{jt}])\varepsilon_t^{\mathrm{m}} + \Gamma_h' Z_{jt-1} + e_{jth}$, where all variables are defined in the notes for Table III. Dashed lines report 90% error bands.

Figure 1 shows that firms with low leverage and high distance to default are consistently more responsive to the shock for up to 3 years later. Panel (a) shows that the peak of the differences by leverage occurs after 4 quarters and that the differences disappear after 12 quarters. Panel (b) shows that the differences by distance to default are larger and significantly more persistent than for leverage. However, these long-run differences are imprecisely estimated with large standard errors, so we focus on the impact effect of the shock for the rest of the paper.

Additional Empirical Results. Online Appendix B and the Supplemental Materials B contain three sets of additional empirical results. The first set of additional results contains a number of robustness checks of our main results. Online Appendix B.3 shows that the results hold when controlling for the information channel of monetary policy using Greenbook forecast revisions (following Miranda-Agrippino and Ricco (2018)) and that the results hold when we start the sample in 1994 rather than 1990. We also perform robustness checks regarding firm-level heterogeneity, including controlling for lagged investment, controlling for interactions of the monetary shock with other firm-level covariates (such as sales growth, future sales growth, size, or liquidity), and investigating other indices of financial constraints.

The second set of results includes some additional analysis of the data. First, as described above, Online Appendix B.1 estimates the dynamics of the average response to monetary policy. Second, Online Appendix B.2 shows that the heterogeneous responses are primarily driven by expansionary shocks. Third, Supplemental Materials B shows that our results hold if we measure leverage using only short term debt, only long term debt, only other liabilities, or using leverage net of liquid assets, though the estimates are less precise for individual categories of debt.

The third set of additional results, in Online Appendix C, relates our work to various strands of the existing literature. First, we show that small firms, measured using Gertler and Gilchrist (1994)'s methodology, are more responsive to monetary shocks in our sample; our results are robust to controlling for this effect. Second, we show that older firms are slightly less responsive to monetary shocks, consistent with recent work by Cloyne et al. (2018); again, our results are robust to controlling for this effect. Third, we reconcile our results with recent work by Jeenas (2019), who argues that low-leverage firms are *less* responsive to monetary policy over longer horizons. We argue that these results are largely driven by permanent heterogeneity in responsiveness, which is outside the scope of our analysis. We also show that our results are not driven by heterogeneity in liquidity across firms, which Jeenas (2019) emphasizes. Supplemental Materials C also shows that our results are not driven by differences in firm-level sales volatility.

3. MODEL

We now develop a heterogeneous firm New Keynesian model in order to interpret this cross-sectional evidence and study its aggregate implications. We describe the model in three blocks: an investment block, which captures heterogeneous responses to monetary policy; a New Keynesian block, which generates a Phillips curve; and a representative household, which closes the model.

3.1. *Investment Block*

The investment block contains a fixed mass of heterogeneous firms that invest in capital subject to financial frictions. It builds heavily on the flexible-price model developed in Khan, Senga, and Thomas (2016). Besides incorporating sticky prices, we extend Khan, Senga, and Thomas (2016)'s framework in three ways. First, we add idiosyncratic capital quality shocks, which help us match observed default rates in the data. Second, we incorporate aggregate adjustment costs in order to generate time-variation in the relative price of capital, as in the financial accelerator literature (e.g., Bernanke, Gertler, and Gilchrist (1999)). Third, we assume that new entrants have lower initial productivity than incumbents, which helps us match lifecycle dynamics.

Production Firms. Time is discrete and infinite. There is no aggregate uncertainty; in Sections 4 and 6 below, we study the transition path in response to an unexpected monetary shock. Each period, there is a fixed mass 1 of production firms. Each firm $j \in [0, 1]$ produces an undifferentiated good y_{jt} using the production function $y_{jt} = z_{jt}(\omega_{jt}k_{jt})^{\theta}l_{jt}^{\nu}$, where z_{jt} is an idiosyncratic total factor productivity shock, ω_{jt} is an idiosyncratic capital quality shock, k_{jt} is the firm's capital stock, l_{jt} is the firm's labor input, and $\theta + \nu < 1$. The idiosyncratic TFP shock follows a log-AR(1) process $\log z_{jt+1} = \rho \log z_{jt} + \varepsilon_{jt+1}$, where $\varepsilon_{jt+1} \sim N(0, \sigma^2)$.

The capital quality shock is i.i.d. across firms and time and follows a truncated log-normal process with support $[-4\sigma_\omega, 0]$, where σ_ω is the standard deviation of the underlying normal distribution. This process implies that with some probability p_ω, no capital quality shock is realized ($\log \omega_{jt} = 0$), but with probability $1 - p_\omega$, capital quality is drawn from the region of a normal distribution within $[-4\sigma_\omega, 0]$. The capital quality shock also affects the value of the firm's undepreciated capital at the end of the period, $(1-\delta)\omega_{jt}k_{jt}$. We view capital quality shocks as capturing unmodeled forces which reduce the value of the firm's capital, such as frictions in the resale market, breakdown of machinery, or obsolescence.[6]

The timing of events within each period is as follows:

[6]Mechanically, the capital quality shocks allow the model to generate positive default risk for a large cross-section of firms. In our model, the value of a firm is dominated by the value of its undepreciated capital stock; without risk to this stock, our model would have the counterfactual prediction that only firms with very low net worth would have positive probability of default.

(i) A mass $\overline{\mu}_t$ of new firms enter the economy. We assume that the mass of new entrants is equal to the mass of firms that exit the economy so that the total mass of production firms is fixed in each period t. Each of these new entrants draws idiosyncratic productivity z_{jt} from the time-invariant distribution $\mu^{\text{ent}}(z) \sim \log N(-m\frac{\sigma}{\sqrt{(1-\rho^2)}}, \frac{\sigma}{\sqrt{(1-\rho^2)}})$, where $m \geq 0$ is a parameter governing the average productivity of new entrants.[7] New entrants are endowed with k_0 units of capital from the household and have no debt. They then proceed as incumbent firms.

(ii) Idiosyncratic shocks to TFP and capital quality are realized.

(iii) With probability π_d each firm receives an i.i.d. exit shock and must exit the economy after producing.

(iv) Each firm decides whether or not to default. If a firm defaults, it immediately and permanently exits the economy. In the event of default, lenders recover a fraction of the firm's capital stock (described in more detail below) and equity holders receive a zero payoff. The fraction of a defaulting firms' capital not recovered by its lenders is transferred lump-sum to households. In order to continue operation, the firm must pay back the face value of its outstanding debt, b_{jt}, and pay a fixed operating cost ξ in units of the final good.

(v) Continuing firms produce by hiring labor l_{jt} from a competitive labor market at real wage w_t. Firms sell their output to retailers (described below) in a competitive market at relative price p_t expressed in terms of the final good (which is our numeraire, described below). At this point, firms that received the i.i.d. exit shock sell their undepreciated capital and exit the economy.

(vi) Continuing firms purchase new capital k_{jt+1} at relative price q_t. Firms have two sources of investment finance, each of which is subject to a friction. First, firms can issue new nominal debt with real face value $b_{jt+1} = \frac{B_{jt+1}}{P_t}$, where B_{jt+1} is the nominal face value and P_t is the nominal price of the final good. Lenders offer a price schedule $Q_t(z_{jt}, k_{jt+1}, b_{jt+1})$ for this debt (we derive this price schedule below). Second, firms can use internal finance by lowering dividend payments d_{jt} but cannot issue new equity, which bounds dividend payments $d_{jt} \geq 0$.[8]

We write the firm's optimization problem recursively. The individual state variables of a firm are its total factor productivity z and its net worth

$$n = \max_l p_t z(\omega k)^\theta l^\nu - w_t l + q_t(1-\delta)\omega k - b\frac{1}{\Pi_t} - \xi,$$

where $\Pi_t = \frac{P_t}{P_{t-1}}$ is realized gross inflation. Net worth n is the total amount of resources available to the firm other than additional borrowing. Conditional on continuing, the real

[7]The parameter m is motivated by the evidence in Foster, Haltiwanger, and Syverson (2016) that young firms have persistently low levels of measured productivity.

[8]The nonnegative dividend constraint captures two key facts about external equity documented in the corporate finance literature. First, firms face significant costs of issuing new equity, both direct flotation costs and indirect costs. Second, firms issue external equity very infrequently. This specific form of the nonnegativity constraint is widely used in the macroliterature because it allows for efficient computation of the model in general equilibrium.

equity value $v_t(z, n)$ solves the Bellman equation[9]

$$\begin{aligned} v_t(z,n) = \max_{k',b'} & \, n - q_t k' + \mathcal{Q}_t(z,k',b')b' \\ & + \mathbb{E}_t\big[\Lambda_{t+1}\big(\pi_d \chi^1\big(\hat{n}_{t+1}(z',\omega',k',b')\big)\hat{n}_{t+1}(z',\omega',k',b') \\ & \times (1-\pi_d)\chi^2_{t+1}\big(z',\hat{n}_{t+1}(z',\omega',k',b')\big)v_{t+1}\big(z',\hat{n}_{t+1}(z',\omega',k',b')\big)\big)\big] \\ \text{such that } & n - q_t k' + \mathcal{Q}_t(z,k',b')b' \geq 0, \end{aligned} \tag{5}$$

where $\hat{n}_{t+1}(z', \omega', k', b') \equiv \max_{l'} p_{t+1}z'(\omega'k')^{\theta}(l')^{\nu} - w_{t+1}l' + q_{t+1}(1-\delta)\omega'k' - b'\frac{1}{\Pi_{t+1}} - \xi$ is the net worth implied by k', b', and the realization of z' and ω'; $\chi^1(n)$ and $\chi^2_t(z, n)$ are indicator variables taking the value of zero if the firm defaults, conditional on exogenously exiting and not exiting; and Λ_{t+1} is the stochastic discount factor.

Proposition 1 characterizes the decision rules which solve this Bellman equation.

PROPOSITION 1: *Consider a firm at time t that is eligible to continue into the next period, has idiosyncratic productivity z, and has net worth n. The firm's optimal decision is characterized by one of the following three cases*:

(i) ***Default***: *there exists a threshold $\underline{n}_t(z)$ such that the firm defaults if $n < \underline{n}_t(z)$. These firms cannot satisfy the nonnegativity constraint on dividends.*

(ii) ***Unconstrained***: *there exists a threshold $\overline{n}_t(z)$ such that the firm is **financially unconstrained** if $n > \overline{n}_t(z)$. Unconstrained firms follow the "frictionless" capital accumulation policy $k'_t(z, n) = k^*_t(z)$ which solves $q_t = \mathbb{E}_t[\Lambda_{t+1} MRPK_{t+1}(z', k^*_t(z))|z]$, where $MRPK_{t+1}(z', k') = \mathbb{E}_{\omega'}[\frac{\partial}{\partial k'}(\max_{l'} p_{t+1}z'(\omega'k')^{\theta}(l')^{\nu} - w_{t+1}l' + q_{t+1}(1-\delta)\omega'k')]$ is the return on capital to the firm. Unconstrained firms are indifferent over any combination of b' and d such that they remain unconstrained for every period with probability one.*

(iii) ***Constrained***: *firms with $n \in [\underline{n}_t(z), \overline{n}_t(z)]$ are **financially constrained**. Constrained firms' optimal investment $k'_t(z, n)$ and borrowing $b'_t(z, n)$ decisions solve the Bellman equation (5). Constrained firms also pay zero dividends because the value of resources inside the firm, used to lower borrowing costs, is higher than the value of resources outside the firm.*

PROOF: See Supplemental Materials E. *Q.E.D.*

Lenders. There is a representative financial intermediary that lends resources from the representative household to firms at the firm-specific price schedule $\mathcal{Q}_t(z, k', b')$. If the firm defaults on the loan in the following period, the lender recovers a fraction α of the market value of the firm's capital stock $q_{t+1}\omega'k'$. The debt price schedule prices this

[9]Firms which receive the exogenous exit shock have simple decision rules. Those that do not default simply sell their undepreciated capital after production. Since these firms cannot borrow, they default whenever net worth $n < 0$.

default risk competitively:

$$\mathcal{Q}_t(z,k',b') = \mathbb{E}_t\bigg[\Lambda_{t+1}\frac{1}{\Pi_{t+1}}\Big(1-\big(1-\big(\pi_d\chi^1\big(\hat{n}_{t+1}(z',\omega',k',b')\big)\\ +(1-\pi_d)\chi^2_{t+1}\big(z',\hat{n}_{t+1}(z',\omega',k',b')\big)\big)\big)\\ \times\Big(1-\min\Big\{\frac{\alpha q_{t+1}(1-\delta)\omega'k'}{b'/\Pi_{t+1}},1\Big\}\Big)\Big)\bigg], \tag{6}$$

3.2. *New Keynesian Block*

The New Keynesian block of the model is designed to parsimoniously generate a New Keynesian Phillips curve relating nominal variables to the real economy.

Retailers and Final Good Producer. There is a fixed mass of retailers $i \in [0,1]$. Each retailer produces a differentiated variety $\widetilde{y}_{it}$ using the heterogeneous production firms' good as its only input: $\widetilde{y}_{it} = y_{it}$, where y_{it} is the amount of the undifferentiated good demanded by retailer i. Retailers set a relative price for their variety $\widetilde{p}_{it}$ but must pay a quadratic price adjustment cost $\frac{\varphi}{2}(\frac{\widetilde{p}_{it}}{\widetilde{p}_{it-1}}-1)^2 Y_t$, where Y_t is the final good. The retailers' demand curve is generated by the representative final good producer, which has production function $Y_t = (\int \widetilde{y}_{it}^{\frac{\gamma-1}{\gamma}}\,\mathrm{d}i)^{\frac{\gamma}{\gamma-1}}$, where γ is the elasticity of substitution over intermediate goods. This final good is the numeraire.

The retailers and final good producers aggregate into the familiar New Keynesian Phillips Curve:

$$\log \Pi_t = \frac{\gamma-1}{\varphi}\log\frac{p_t}{p^*} + \beta\mathbb{E}_t\log\Pi_{t+1}, \tag{7}$$

where $p^* = \frac{\gamma-1}{\gamma}$ is the steady state relative price of the heterogeneous production firm output.[10] The Phillips Curve links the New Keynesian block to the investment block through the relative price p_t. When aggregate demand for the final good Y_t increases, retailers must increase production of their differentiated goods. Because of the nominal rigidities, this force increases demand for the heterogeneous firms' goods y_{it}, which increases their relative price p_t and generates inflation through (7).

Capital Good Producer. There is a representative capital good producer who produces new aggregate capital using the technology $\Phi(\frac{I_t}{K_t})K_t$, where I_t are units of the final good used to produce capital, $K_t = \int k_{jt}\,\mathrm{d}j$ is the aggregate capital stock at the beginning of the period, $\Phi(\frac{I_t}{K_t}) = \frac{\hat{\delta}^{1/\phi}}{1-1/\phi}(\frac{I_t}{K_t})^{1-1/\phi} - \frac{\hat{\delta}}{\phi-1}$, and $\hat{\delta}$ is the steady-state investment rate.[11] Profit maximization pins down the relative price of capital as

$$q_t = \frac{1}{\Phi'\Big(\dfrac{I_t}{K_t}\Big)} = \left(\frac{I_t/K_t}{\hat{\delta}}\right)^{1/\phi}. \tag{8}$$

[10]We focus directly on the linearized formulation for computational simplicity.

[11]Because the capital quality shock follows a truncated log-normal process, the steady-state investment rate is $\hat{\delta} = (1-(1-\delta)\mathbb{E}[\omega])(1+\frac{k_0\overline{\mu}}{K^*})$, where K^* is the steady-state capital stock and $\overline{\mu}$ the steady-state level of new entrants. For more details, see Supplemental Materials D.

Monetary Authority. The monetary authority sets the nominal risk-free interest rate R_t^{nom} according to the Taylor rule $\log R_t^{\text{nom}} = \log \frac{1}{\beta} + \varphi_\pi \log \Pi_t + \varepsilon_t^{\text{m}}$, where $\varepsilon_t^m \sim N(0, \sigma_m^2)$, φ_π is the weight on inflation in the reaction function, and ε_t^{m} is the monetary policy shock.

3.3. *Representative Household and Equilibrium*

There is a representative household with preferences over consumption C_t and labor supply L_t represented by the expected utility function

$$\mathbb{E}_0 \sum_t^\infty \beta^t (\log C_t - \Psi L_t),$$

where β is the discount factor and Ψ controls the disutility of labor supply. The household owns all firms in the economy. We study perfect foresight transition paths with respect to aggregate states, so the stochastic discount factor and nominal interest rate are linked through the Euler equation for bonds, $\Lambda_{t+1} = \frac{1}{R_t^{\text{nom}}/\Pi_{t+1}}$.

An equilibrium is a set of value functions $v_t(z, n)$; decision rules $k_t'(z, n)$, $b_t'(z, n)$, $l_t(z, n)$; measure of firms $\mu_t(z, \omega, k, b)$; debt price schedule $Q_t(z, k', b')$; and prices w_t, q_t, p_t, Π_t, Λ_{t+1} such that (i) all firms optimize, (ii) lenders price default risk competitively, (iii) the household optimizes, (iv) the distribution of firms is consistent with decision rules, and (v) all markets clear. Supplemental Materials D precisely defines an equilibrium of our model.

4. CHANNELS OF MONETARY TRANSMISSION

Before performing the quantitative analysis, we theoretically characterize the channels through which monetary policy affects investment in our model. This exercise identifies the key sources of heterogeneous responses across firms, which motivates our calibration in Section 5.

Monetary Policy Experiment. We study the effect of an unexpected innovation to the Taylor rule ε_t^{m} followed by a perfect foresight transition back to steady state. This approach allows for clean analytical results because there is no distinction between ex ante expected real interest rates and ex post realized real interest rates. We focus on financially constrained firms as defined in Proposition 1 because they make up more than 99.4% of the firms in our calibration.

Impact on Decision Rules. The optimal choice of investment k' and borrowing b' satisfy the following two conditions:

$$q_t k' = n + \frac{1}{R_t(z, k', b')} b', \tag{9}$$

$$\left(q_t - \varepsilon_{Q,k'}(z, k', b') \frac{Q_t(z, k', b')b'}{k'}\right) \frac{R_t^{\text{sp}}(z, k', b')}{1 - \varepsilon_{R,b'}(z, k', b')}$$
$$= \frac{1}{R_t} \mathbb{E}_t\left[\text{MRPK}_{t+1}(z', k')\right]$$

$$+\frac{1}{R_t}\frac{\mathbb{C}ov_t\big(\mathrm{MRPK}_{t+1}(z',\omega'k'),1+\lambda_{t+1}(z',\hat{n}_{t+1}(z',\omega',k',b'))\big)}{\mathbb{E}_t\big[1+\lambda_{t+1}(z',\hat{n}_{t+1}(z',\omega',k',b'))\big]}$$
$$-\frac{1}{R_t}\mathbb{E}_{\omega'}\big[v^0_{t+1}(\omega',k',b')g_z\big(\underline{z}(\omega',k',b')|z\big)\hat{\underline{z}}_{t+1}(\omega',k',b')\big], \tag{10}$$

where R_t is the real risk-free rate between t and $t+1$, $R_t(z,k',b')=\frac{1}{Q_t(z,k',b')}$ is the firm's implied interest rate schedule, $\varepsilon_{Qk'}(z,k',b')$ is the elasticity of the bond price schedule with respect to investment k', $R^{\mathrm{sp}}_t(z,k',b')=R_t(z,k',b')/R^{\mathrm{nom}}_t$ is a measure of the borrowing spread, $\varepsilon_{Rb'}(z,k',b')$ is the elasticity of the interest rate schedule with respect to borrowing, $\lambda_t(z,n)$ is the Lagrange multiplier on the nonnegativity constraint on dividends, $\underline{z}_t(\omega,k,b)$ is the default threshold in terms of productivity (which inverts the net worth threshold defined in Proposition 1), $v^0_t(\omega,k,b)\equiv v_t(\underline{z}_t(\omega,k,b),\hat{n}_t(\underline{z}_t(\omega,k,b),\omega,k,b))$ is the value of the firm evaluated at the default threshold, $g_z(z'|z)$ is the density of z' conditional on z, and $\hat{\underline{z}}_{t+1}(\omega',k',b')\equiv\frac{\partial\underline{z}_{t+1}(\omega',k',b')}{\partial k'}+\frac{\partial\underline{z}_{t+1}(\omega',k',b')}{\partial b'}(q_t-\varepsilon_{Q,k'}\frac{Q_t(z,k',b')b'}{k'})\frac{R_t(z,k',b')}{1-\varepsilon_{R,b'}(z,k',b')]}$. Condition (9) is the nonnegativity constraint on dividends and condition (10) is the intertemporal Euler equation. The expectation and covariances in this expression are only taken over the states in which the firm does not default.

The marginal cost of capital is the product of two terms. The first term, $q_t-\varepsilon_{Q,k'}(z,k',b')\frac{Q_t(z,k',b')b'}{k'}$, is the relative price of capital goods q_t net of the interest savings due to higher capital, $\varepsilon_{Q,k'}(z,k',b')\frac{Q_t(z,k',b'}{b}'k'$. The interest savings result from the fact that, all else equal, higher capital decreases expected losses due to default to the lenders. The second term in the marginal cost of capital is related to borrowing costs, $\frac{R^{\mathrm{sp}}_t(z,k',b')}{1-\varepsilon_{R,b'}(z,k',b')}$. A higher interest rate spread or a higher slope of that spread results in higher borrowing costs.

The marginal benefit of capital is the sum of three terms. The first term, $\frac{1}{R_t}\times\mathbb{E}_t[\mathrm{MRPK}_{t+1}(z',k')]$, is the expected return on capital discounted by the real interest rate. The second term, $\frac{1}{R_t}\frac{\mathbb{C}ov_t(\mathrm{MRPK}_{t+1}(z',\omega'k'),1+\lambda_{t+1}(z',\omega',k',b'))}{\mathbb{E}_t[1+\lambda_{t+1}(z',\omega',k',b'))]}$, captures the covariance of the return on capital with the firm's shadow value of resources. The third term captures how additional investment affects the firm's default probability and, therefore, the value of the firm. In our calibration, this term is negligible because the value of the firm close to the default threshold, $v^0_{t+1}(\omega',k',b')$, is essentially zero.

Figure 2 plots the marginal benefit and marginal cost schedules as a function of capital accumulation k'. In order to illustrate the key economic mechanisms, we compare how these curves shift following an expansionary monetary policy shock for two extreme examples of firms. These firms share the same level of productivity but the first firm has high net worth and is currently risk-free (though it is still constrained in the sense of Proposition 1), while the second has low net worth and is risky constrained.

Risk-Free Firm. The left panel of Figure 2 plots the two schedules for the risk-free firm. The marginal cost curve is flat when capital accumulation k' can be financed without incurring default risk, but becomes upward sloping when the borrowing required to achieve k' creates default risk and, therefore, a credit spread. The marginal benefit curve is downward sloping due to diminishing returns to capital. In the initial equilibrium, the firm is risk-free because the two curves intersect in the flat region of the marginal cost curve.

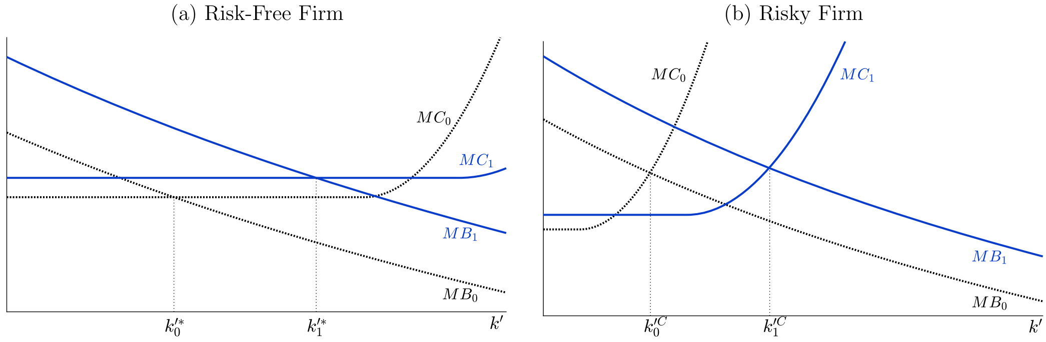

FIGURE 2.—Response to monetary policy for risk-free and risky firms. Notes: Marginal benefit and marginal cost curves as a function of capital investment k' for firms with same productivity. Left panel is for a firm with high initial net worth and right panel is for a firm with low initial net worth. Marginal cost curve is the left-hand side of (10) and marginal benefit the right-hand side of (10). Dashed black lines plot the curves before an expansionary monetary policy shock, and solid blue lines plot the curves after the shock.

The expansionary monetary shock shifts both the marginal benefit and marginal cost curves. The marginal benefit curve shifts out for two reasons. First, the shock decreases the real interest rate, which decreases the firm's discount rate R_t and, therefore, increases the discounted return on capital. Second, the shock also changes the relative price of output p_{t+1}, the real wage w_{t+1}, and the relative price of undepreciated capital q_{t+1} due to general equilibrium effects. In our calibration, these changes increase the return on capital $\text{MRPK}_{t+1}(z, k')$ and, therefore, further shift out the marginal benefit curve. Third, the shock also affects the covariance term and the change in default threshold, which further shifts out the marginal benefit curve.

The marginal cost curve shifts up because the increase in aggregate investment demand increases the relative price of capital q_t. In the new equilibrium, the firm has increased its capital and remains risk-free because the marginal benefit and marginal cost curves still intersect along the flat region of the marginal cost curve.

Risky Firm. The right panel of Figure 2 plots how the marginal benefit and marginal cost curves shift for the risky firm. Because this firm has low initial net worth n, it needs to borrow more than the risk-free firm to achieve the same level of investment. Hence, its marginal cost curve is upward-sloping over a larger region of net worth.

The key difference between the risky and the risk-free firm is how monetary policy shifts the marginal cost curve. As with the risk-free firm, the curve shifts up because the relative price of capital q_t increases, but there are now two additional effects. First, monetary policy increases net worth n, which decreases the amount the firm needs to borrow to finance any level of investment and, therefore, extends the flat region of the marginal cost curve. The increase in net worth can be decomposed according to

$$\frac{\partial \log n}{\partial \varepsilon_t^{\mathrm{m}}} = \frac{1}{1-\nu-\theta}\left(\frac{\partial \log p_t}{\partial \varepsilon_t^{\mathrm{m}}} - \nu \frac{\partial \log w_t}{\partial \varepsilon_t^{\mathrm{m}}}\right)\frac{\iota_t(z, \omega k)}{n} + \frac{\partial \log q_t}{\partial \varepsilon_t^{\mathrm{m}}}\frac{q_t(1-\delta)\omega k}{n} + \frac{\partial \log \Pi_t}{\partial \varepsilon_t^{\mathrm{m}}}\frac{b/\Pi_t}{n}, \tag{11}$$

where $\iota_t(z, \omega k) = \max_l p_t z(\omega k)^\theta l^\nu - w_t l$. This expression (11) contains three ways that monetary policy affects cash flows. First, monetary policy affects current revenues by changing the relative price of output p_t net of real labor costs νw_t. Second, monetary policy affects the value of firms' undepreciated capital stock by changing the relative price of capital q_t. Finally, monetary policy changes the real value of nominal debt through inflation Π_t.

The second key difference in how monetary policy affects the risky firm's marginal cost curve is that it flattens the upward-sloping region, reflecting reduced credit spreads. Recall that, in the event of default, lenders recover $\alpha q_{t+1}\omega_{jt+1}k_{jt+1}$ per unit of debt; since the shock increases the relative price of capital q_{t+1}, it also increases the recovery rate, which reduces credit spreads. This channel is reminiscent of the "financial accelerator" in Bernanke, Gertler, and Gilchrist (1999). Monetary policy also decreases the probability of default, but this effect is quantitatively small in our calibration.

Whether the risky firm is more or less responsive than the risk-free firm depends crucially on the size of these two shifts in the marginal cost curve. Theoretically, they may or may not be large enough to induce the risky firm to be more responsive to monetary policy than the risk-free firm. The goal of our calibration is to quantitatively discipline these shifts using our model.[12,13]

Relationship to Other Papers. The simple framework in Figure 2 provides a powerful tool to organize various results in the existing literature. Bernanke, Gertler, and Gilchrist (1999) developed a model in which firms' production functions are constant returns to scale, which results in a horizontal marginal benefit curve for investment. The level of investment is determined by the point at which this curve intersects the upward-sloping region of the marginal cost curve. Therefore, movements in the marginal cost curve have a stronger effect on how investment responds to monetary policy shocks, increasing the strength of the financial accelerator channel described above.

Jeenas (2019) developed a model in which firms face collateral constraints and a fixed cost of issuing debt but can accumulate liquid financial assets. This model implies that firms face two kinked marginal cost curves for financing investment: one corresponding to using liquid assets (which is flat until these assets are exhausted, and then becomes vertical) and another corresponding to new borrowing (which is flat—at a higher lever due an exogenous spread in borrowing costs—until firms reach the collateral constraint, at which point it becomes vertical). Optimal investment is determined by the mix of these two marginal costs curves that firms use when financing their investment. Many firms do not find it worthwhile to issue new debt, so their marginal source of investment finance

[12]Online Appendix B.2 shows that the heterogeneous responses to monetary policy we find in the data are primarily driven by expansionary shocks. While we do not emphasize that result due to its wide standard errors, it is potentially consistent with the analysis in Figure 2. Suppose that high-risk firms tend to position themselves at the point where their marginal cost curve just begins to be upward sloping. Then an expansionary shock will move these firms forward along the upward-sloping part—dampening their response relative to low-risk firms—while a contractionary shock will move them backward along the flat part—not dampening their response.

[13]We can use this analysis to conjecture how incorporating long-term debt would affect our results. If we were to increase the maturity of debt but hold all other parameters fixed, then we would of course decrease default probabilities (since firms will have to roll over less debt each period) and potentially flatten out the marginal cost curve. Therefore, we would also need to recalibrate the parameters in order to match the same average probability of default as in the current model. We expect that this recalibration would also imply a similar slope of the default probabilities with respect to borrowing and, therefore, a similar slope for the marginal cost curve. However, the marginal cost curve may become more responsive to monetary shocks if the resulting inflation significantly decreases the real value of long-term debt (as in Gomes, Jermann, and Schmid (2016)).

is liquid assets. Therefore, firms that have more liquid assets have a larger flat region of their liquid-asset-cost curve and are more responsive to monetary policy.

Cloyne et al. (2018) argued that young firms are more responsive to monetary shocks in the U.S. and the U.K. While we show in Online Appendix C.2 that age does not drive our empirical results, one can nevertheless interpret their findings through the lens of our model. One possible interpretation is that young firms face a steeper marginal cost curve than old firms, but young firms' marginal cost curves are also more sensitive to monetary policy. Another interpretation is that young firms' marginal benefit curves are themselves more responsive to monetary policy, for example if their product demand is more cyclically sensitive.

5. PARAMETERIZATION

We now calibrate the model and verify that its steady state behavior is consistent with key features of the microdata.

5.1. *Calibration*

We calibrate the model in two steps. First, we exogenously fix a subset of parameters. Second, we choose the remaining parameters in order to match moments in the data.

Fixed Parameters. Table IV lists the parameters that we fix. The model period is one quarter, so we set the discount factor $\beta = 0.99$. We set the coefficient on labor $\nu = 0.64$. We choose the coefficient on capital $\theta = 0.21$ to imply a total returns to scale of 85%. Capital depreciates at rate $\delta = 0.025$ quarterly. We choose the elasticity of substitution in final goods production $\gamma = 10$, implying a steady state markup of 11%. This choice implies that the steady state labor share is $\frac{\gamma-1}{\gamma}\nu \approx 58\%$, close to the U.S. labor share reported in Karabarbounis and Neiman (2013). We choose the coefficient on inflation in the Taylor rule $\varphi_\pi = 1.25$, in the middle of the range commonly considered in the literature. We set the price adjustment cost parameter $\varphi = 90$ to generate a Phillips Curve slope equal to 0.1, as in Kaplan, Moll, and Violante (2018). Finally, we set the curvature of the aggregate

TABLE IV

FIXED PARAMETERS[a]

Parameter	Description	Value
Household		
β	Discount factor	0.99
Firms		
ν	Labor coefficient	0.64
θ	Capital coefficient	0.21
δ	Depreciation	0.025
New Keynesian Block		
ϕ	Aggregate capital AC	4
γ	Demand elasticity	10
φ_π	Taylor rule coefficient	1.25
φ	Price adjustment cost	90

[a]Parameters exogenously fixed in the calibration.

TABLE V

FITTED PARAMETERS[a]

Parameter	Description	Value
Idiosyncratic shock processes		
ρ	Persistence of TFP (fixed)	0.90
σ	SD of innovations to TFP	0.03
σ_ω	SD of capital quality	0.04
Financial frictions		
ξ	Operating cost	0.04
α	Loan recovery rate	0.54
Firm lifecycle		
m	Mean shift of entrants' prod.	3.12
k_0	Initial capital	0.18
π_d	Exogenous exit rate	0.01

[a]Parameters chosen to match the moments in Table VI.

adjustment costs $\phi = 4$ following Bernanke, Gertler, and Gilchrist (1999). This level of adjustment costs roughly matches the peak response of investment relative to the peak response of output estimated in Christiano, Eichenbaum, and Evans (2005).

Fitted Parameters. We choose the parameters listed in Table V to match the empirical moments reported in Table VI. The first set of parameters governs the idiosyncratic shocks (ρ, σ, and σ_ω), the second set governs the frictions to external finance (ξ and α), and the third set governs the firm lifecycle (m, k_0, and π_d). None of the statistics that we

TABLE VI

CALIBRATION TARGETS AND MODEL FIT[a]

Moment	Description	Data	Model
Investment behavior (annual)			
$\sigma(\frac{i}{k})$	SD investment rate	0.33	0.37
Financial behavior (annual)			
$\mathbb{E}$[default rate]	Mean default rate	3.0%	3.0%
$\mathbb{E}[\frac{b}{k}]$	Mean gross leverage ratio	0.34	0.49
Frac$(b > 0)$	Firms w/ positive debt	0.81	0.70
Firm Growth (annual)			
N_1/N	Share of employment in age ≤ 1	0.03	0.02
N_{1-10}/N	Share of employment in age $\in (1, 10)$	0.21	0.36
N_{11+}/N	Share of employment in age ≥ 10	0.76	0.62
Firm Exit (annual)			
$\mathbb{E}$[exit rate]	Mean exit rate	8.7%	8.8%
M_1/M	Share of firms at age 1	0.10	0.08
M_2/M	Share of firms at age 2	0.08	0.06

[a]Empirical targets in the calibration. See main text for data sources.

target are drawn from Compustat; later on, when we compare our model to the empirical results from Section 2, we will account for the selection of firms into Compustat.[14]

We target the dispersion of plant-level investment rates in Census microdata reported by Cooper and Haltiwanger (2006), which places discipline on the degree of idiosyncratic risk faced by firms.[15] We target a number of statistics related to firms' use of external finance. Following Bernanke, Gertler, and Gilchrist (1999), we target a mean default rate of 3% as estimated in a survey of businesses by Dun and Bradstreet. We target an average firm-level gross leverage ratio of 0.34 from the microdata underlying the Quarterly Financial Reports, as reported in Crouzet and Mehrotra (2020). We also target the share of firms with positive debt from Crouzet and Mehrotra (2020) in order to maintain a realistic distinction between gross and net leverage.[16]

The final two sets of moments are informative about firm lifecycle dynamics. We target the share of employment in firms of age less than 1 year, between 1 and 10 years, and over 10 years, all of which are informative about how quickly young firms grow. We also target the average exit rate and the share of firms in the economy at age 1 and 2, which is informative about the exit rate of young firms. All of these statistics are computed from the Business Dynamics Statistics (BDS), the public-release sample of statistics aggregated from the Census' Longitudinal Business Database (LBD).

Table VI shows that our model matches the targeted moments reasonably well despite the fact that it is overidentified. The model roughly matches the dispersion of investment rates, which captures the degree of idiosyncratic risk faced by firms. The model also matches the average default rate, but slightly overpredicts the average gross leverage ratio and underpredicts the fraction of firms with positive debt. The model matches the share of employment in young firms, but somewhat overpredicts the share of employment in 1–10 year old firms.

The calibrated parameters in Table V are broadly comparable to existing estimates in the literature. Idiosyncratic TFP shocks are less persistent and more volatile than measured aggregate TFP shocks, consistent with direct measurements of plant- or firm-level productivity. The calibrated loan recovery rate is 0.54, as in Khan, Senga, and Thomas (2016). New entrants start with significantly lower productivity and capital than the average firm. The capital quality shock process implies that there is a $p_\omega = 0.59$ probability of receiving a zero shock $\log \omega_{jt} = 0$. Online Appendix D contains a formal discussion of identification using the local elasticities of moments with respect to parameters as well as the elasticities of estimated parameters with respect to moments (computed using the tools from Andrews, Gentzkow, and Shapiro (2017)). The financial frictions in our calibrated model affect nearly all firms in the stationary distribution. Using the classification from Proposition 1, 52.8% of firms are risky constrained, 47.5% of firms are risk-free constrained, and 0.6% of firms are unconstrained.

[14]At each step of the moment-matching process, we choose the disutility of labor supply Ψ to generate a steady state employment rate of 60%.

[15]We prefer to use the plant-level data from Cooper and Haltiwanger (2006), rather than firm-level data from other sources, because Cooper and Haltiwanger (2006) carefully construct measures of retirement and sales of capital to measure negative investment, which is important in our model because capital is liquid. Cooper and Haltiwanger (2006)'s sample is a balanced panel of plants that have survived at least 16 years; to mirror this sample selection in the model, we condition on firms that have survived for 20 years, and our calibration results are robust to different choices of this cutoff.

[16]We do not target credit spreads because observed credit spreads in the data are driven not only by the risk-neutral pricing of default risk, as in our model, but also by aggregate risk premia, which are outside our model. Consistent with this idea, the average annual credit spread in our model is 0.7%, compared to the 2.4% spread of BAA corporate bonds over the 10-year Treasury bond in the data.

5.2. *Financial Heterogeneity in the Model and the Data*

Online Appendix D analyzes firms' decision rules in steady state and identifies two key sources of financial heterogeneity across firms. The first source is lifecycle dynamics; firms are born below their optimal scale, that is, $k_0 < k^*(z)$, and need to grow their capital stock. These young firms initially borrow in order to accumulate capital, increasing their risk of default and therefore borrowing costs. The second source of financial heterogeneity is TFP shocks z; a positive shock increases the firm's optimal scale $k^*(z)$, which again induces debt-financed capital accumulation. We show that firms more affected by these financial frictions have a positive "marginal propensity to invest" out of net worth.

The lifecycle dynamics of firms in our model are in line with the key features of the data emphasized by the firm dynamics literature. Panel (a) in Figure 3 compares the distribution of firm growth rates in steady state to the establishment-level data from the Business Employment Dynamics (BED) data, reported in Davis et al. (2010). The model matches the empirical distribution of growth rates fairly well except for the large mass of growth rates within $(-0.05, 0.05)$ in the data. This discrepancy is driven by the fact that 15% of firms have growth rates of exactly zero; these observations likely correspond to small, nongrowing establishments, which are outside our model. Panel (b) in Figure 3 shows that the model produces a negative correlation between age and growth, as in the data (and comparable to the model in Clementi and Palazzo (2016)).

Online Appendix D and Supplemental Materials F further analyze the behavior of the model in steady state and compares it to the data. First, we further analyze the lifecycle dynamics of firms. Second, we show that the joint distribution of investment and leverage rates in our model is comparable to Census and Compustat data. Finally, we compare our model's sample of public and private firms to the data.

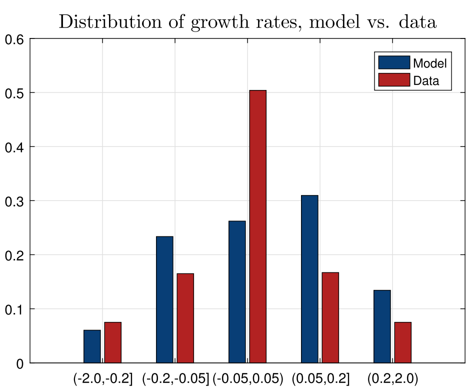

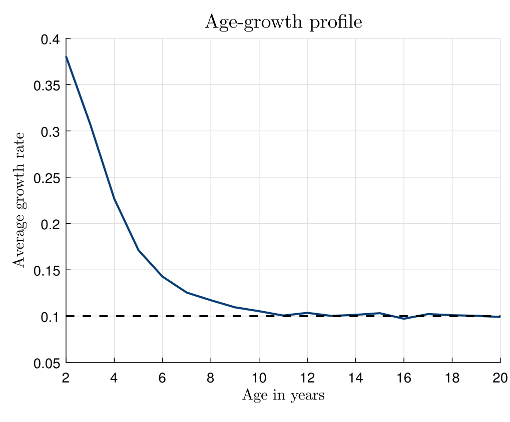

FIGURE 3.—Firm lifecycle dynamics, model vs. data. Notes: Panel (a) plots a histogram of the distribution of quarterly growth rates in the model versus the data. In the model, we measure the firm's growth rate as $\frac{l_{jt+1}-l_{jt}}{0.5(l_{jt+1}+l_{jt})}$ where l_{jt} is employment. "Data" is the empirical distribution of quarterly establishment growth rates in the Business Employment Dynamics (BED) data, reported in Davis et al. (2010). Panel (b) plots the average firm-level growth rate as a function of age in steady state. We add 0.1 to the model's growth profile to account for the fact that our model does not feature trend growth. We exclude the first year of growth since firms in our model are born significantly below optimal scale; the average growth rate in year 1 is nearly 1.

6. QUANTITATIVE MONETARY POLICY ANALYSIS

We now quantitatively analyze the effect of a monetary policy shock $\varepsilon_t^{\mathrm{m}}$. Section 6.1 begins the analysis by computing the aggregate impulse responses to an expansionary shock in our calibrated model. Section 6.2 then studies the heterogeneous effects of monetary policy across firms and shows that, consistent with the empirical results from Section 2, firms with high default risk are less responsive to monetary policy. Finally, Section 6.3 performs a simple calculation to show that the aggregate effect of monetary policy may depend on the distribution of default risk across firms.

The economy is initially in steady state and unexpectedly receives a $\varepsilon_0^{\mathrm{m}} = -0.0025$ innovation to the Taylor rule which reverts to 0 according to $\varepsilon_{t+1}^{\mathrm{m}} = \rho_m \varepsilon_t^{\mathrm{m}}$ with $\rho_m = 0.5$. We compute the perfect foresight transition path of the economy as it converges back to steady state.

6.1. *Aggregate Response to Monetary Policy*

Figure 4 plots the responses of key aggregate variables to this expansionary shock. The shock lowers the nominal interest rate and, because prices are sticky, also lowers the real interest rate. The lower real interest rate stimulates investment demand by shifting out the marginal benefit of investment, as discussed in Section 4. It also stimulates consumption demand from the household due to the standard intertemporal substitution. The higher aggregate demand for goods changes other prices in the economy, further shifting the marginal benefit and marginal costs curves for investment. Overall, investment increases by approximately 1.4%, consumption increases by 0.4%, and output increases by 0.5%. These magnitudes are broadly in line with the peak effects of monetary policy shocks estimated in Christiano, Eichenbaum, and Evans (2005); for a similarly-sized change in the nominal interest rate, they find that investment increases by approximately 1%, consumption increases by 0.2%, and output increases by 0.5%.[17]

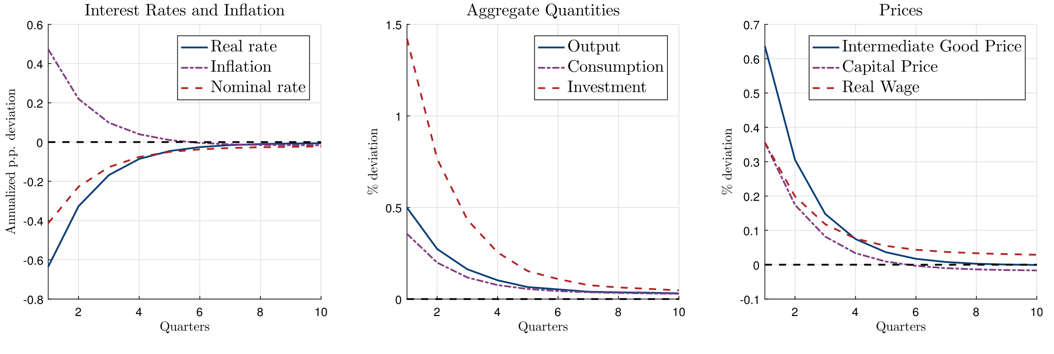

FIGURE 4.—Aggregate responses to expansionary monetary shock. Notes: Aggregate impulse responses to a $\varepsilon_0^{\mathrm{m}} = -0.0025$ innovation to the Taylor rule which decays at rate $\rho_m = 0.5$. Computed as the perfect foresight transition in response to a series of unexpected innovations starting from steady state.

[17]Our model does not generate the hump-shaped aggregate responses emphasized by Christiano, Eichenbaum, and Evans (2005). We could do so by incorporating adjustment costs to investment rather than capital. However, in order to be consistent with the hump-shaped responses of consumption and employment, we would also need to add habit formation and potentially labor adjustment costs. While interesting, this extension is outside the scope of this paper, whose goal is to focus on the role of financial heterogeneity in monetary transmission using an otherwise basic New Keynesian model.

6.2. *Heterogeneous Responses to Monetary Policy*

Model-Implied Regression Coefficients. In order to directly compare our model to the data, we simulate a panel of firms in response to a monetary shock and estimate our empirical specification (2) on the simulated data:[18] $\Delta \log k_{jt+1} = \alpha_j + \alpha_{st} + \beta(x_{jt-1} - \mathbb{E}_j[x_{jt}])\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}' Z_{jt-1} + e_{jt}$. We account for the sample selection into Compustat by conditioning on firms that have survived at least 7 years, which is around the median time to IPO reported in Wilmer Cutler Pickering Hale and Dorr LLP (2017). Online Appendix D shows that the behavior of the model's public firms versus private firms compares fairly well to the data along certain dimensions. We assume that the high-frequency shocks $\varepsilon_t^{\mathrm{m}}$ that we measure in the data are the innovations to the Taylor rule in the model. We estimate the regressions using data from 1 year before the shock to 10 quarters after the shock.

We estimate the empirical specification (2) using leverage ℓ_{jt} as the measure of financial position x_{jt} for two main reasons. First, it is not obvious how to map our model to measured distance to default because measured distance to default is based on the volatility of firms' equity values, which are partly driven by equity risk premia due to aggregate shocks. Second, there is a tight relationship between leverage and default risk in the model. This relationship occurs because there is a monotonic relationship between leverage and net worth (shown in Figure 24 in Supplemental Materials F), and firms only default when net worth falls below the default threshold $\underline{n}_t(z)$.

Columns (1) and (2) of Table VII shows that high-leverage firms are less responsive to monetary policy in the model, as in the data. In the data, a firm with one standard deviation more leverage than the average firm has an investment semielasticity that is approximately -0.57 percentage points lower than the average firm; in the model, that firm has an approximately -1.47 lower semielasticity (which is just outside the 95% confidence interval of the empirical estimate). The R^2 of the regression is lower in the data than in the model, indicating that the data contain more unexplained variation than the model.

TABLE VII

EMPIRICAL RESULTS, MODEL VS. DATA[a]

	Standardized		Not Standardized	
	Data (1)	Model (2)	Data (3)	Model (4)
demeaned leverage × ffr shock	−0.57	−1.30	−3.12	−6.84
	(0.29)		(1.47)	
Firm controls	yes	yes	yes	yes
Time FE	yes	yes	yes	yes
R^2	0.12	0.58	0.12	0.58

[a]Column (1) show the results from running the specification $\Delta \log k_{jt} = \alpha_j + \alpha_{st} + \beta(\ell_{jt-1} - \mathbb{E}_j[\ell_{jt}])\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}' Z_{jt-1} + e_{jt}$, where all variables are defined in the notes for Table III. Column (2) estimates this empirical specification on the simulated data. In the model, we use time fixed effects rather than sector-time fixed effects and we do not include the subset of control variables Z_{jt} which are outside our model, such as fiscal quarter. The sample period is 4 quarters before the monetary shock through 10 quarters after the shock. To mirror the sample selection into Compustat, we condition on firms that have survived at least 7 years. Columns (3) and (4) do not standardize $\ell_{jt} - \mathbb{E}_j[\ell_{jt}]$.

[18]In the model, we use time fixed effects rather than sector-time fixed effects because our model does not contain multiple sectors. In addition, we do not include the subset of control variables Z_{jt} which are outside our model, such as fiscal quarter.

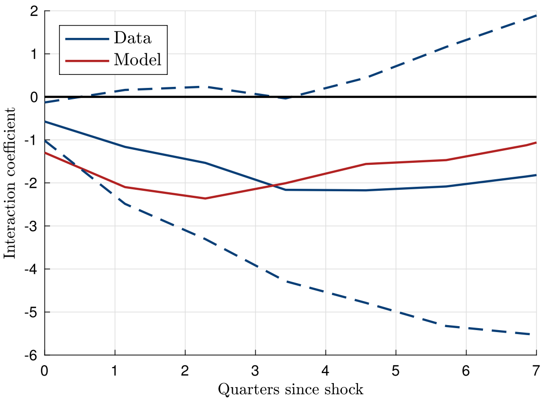

FIGURE 5.—Dynamics of differential responses, model vs. data. Notes: dynamics of the interaction coefficient between leverage and monetary shocks. Reports the coefficient β_h over quarters h from $\log k_{jt+h} - \log k_{jt} = \alpha_{jh} + \alpha_{sth} + \beta_h(\ell_{jt-1} - \mathbb{E}_j[\ell_{jt}])\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}'_h Z_{jt-1} + \mathbf{\Gamma}'_{2h}(\ell_{jt-1} - \mathbb{E}_j[\ell_{jt}])Y_{t-1} + e_{jt}$, where all table notes from Columns (1) and (2) of Table VII apply. Dashed lines report 90% error bands.

Columns (3) and (4) estimate the same regression except that they do not standardize the leverage variable (so that the magnitudes of the coefficients can be interpreted as a reduced-form elasticity of the responsiveness with respect to leverage). The coefficient increases by roughly an order of four in both the model and the data, consistent with the fact that the dispersion of leverage in our model is similar to the data.

Figure 5 shows that the dynamics of the differential responses of investment are persistent in the model, consistent with the data. In this figure, we estimate the local projection (4) on our model-simulated data using standardized leverage. Quantitatively, the model's differences mostly stay within the data's 90% confidence interval up to 8 quarters after the shock.

These heterogeneous responses indicate that high-leverage firms are positioned on the upward-sloping part of their marginal cost curve from Figure 2, and that the shifts in that curve are quantitatively dominated by the shift out in the marginal benefit curve. This mechanism has two implications for the data, both of which are confirmed in Figure 6. First, the top panel shows that firms with high default risk—either those with high leverage or low distance to default—see an increase in their borrowing costs relative to firms with low default risk. In the data, we measure borrowing costs as average interest payments relative to lagged liabilities; therefore, the response of interest payments after 8 to 10 quarters (which is approximately the average maturity of debt in Compustat) most closely corresponds to interest payments on new borrowing in our model.[19]

The second implication of our model's mechanism is that low-risk firms can afford to access more external finance following a monetary shock; Figure 6 shows that this implication also holds in the data. We define external financing flows as the sum of the change

[19]The borrowing costs of high-leverage firms also rise relative to the borrowing costs of low-leverage firms in our model, but the quantitative magnitude of this spread is smaller than in the data. However, the marginal cost curve in the model is determined by both measured spreads and the shadow value of the net worth ($\lambda_t(z, n)$ from Section 4). We do not compare the quantitative implications of the model to the data because our model is not calibrated to match credit spreads, the empirical response of credit spreads to monetary policy are contaminated by changes in risk premia, and credit spreads do not directly correspond to the shadow value $\lambda_t(z, n)$.

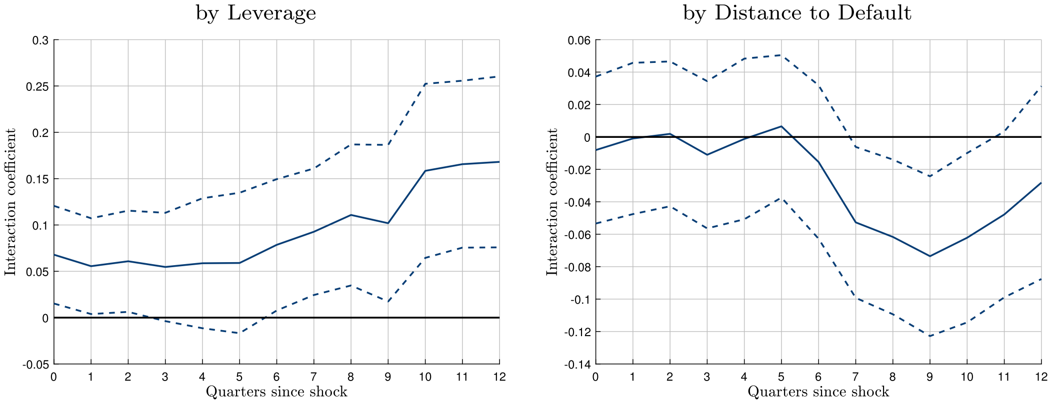

Response of External Financing Flows

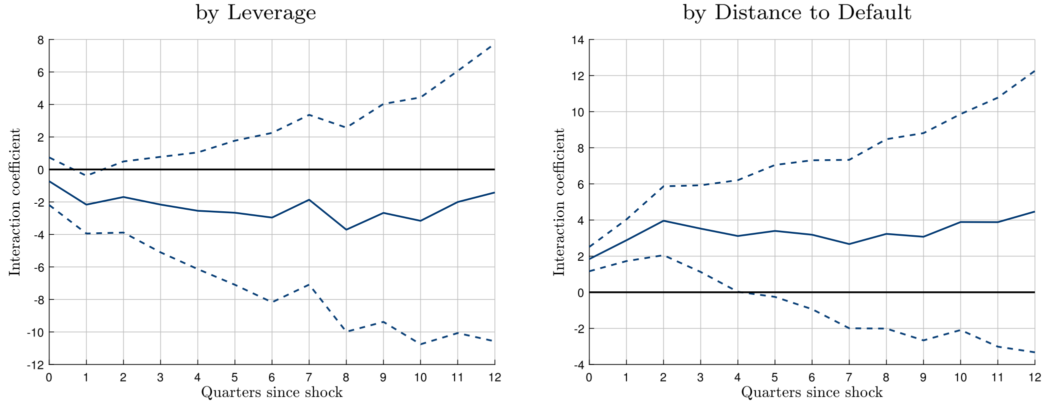

FIGURE 6.—Testable implications of model mechanism for the data. Notes: Reports the coefficient β_h from $y_{jt+h} = \alpha_{jh} + \alpha_{sth} + \beta_h(x_{jt-1} - \mathbb{E}_j[x_{jt}])\varepsilon_t^{\mathrm{m}} + \mathbf{\Gamma}'_h Z_{jt-1} + e_{jth}$, where y_{jt+h} is either the average interest rate in period $t+h$ or the external financing flows between periods $t+h$ and t, and all notes from Columns (1) and (2) of Table VII apply. Dashed lines report 90% error bands.

in total debt plus the change in total equity, scaled by assets. We then estimate the dynamic empirical specification (4) with this measure of financing flows on the left-hand side. We include equity because it is a source of external finance for firms in the data. Figure 6 shows that, in the data, firms with low leverage or high distance to default access more external finance following a monetary shock.

Decomposition of Channels Driving Heterogeneous Responses. In order to better understand the sources of these heterogeneous responses across firms, we now decompose the channels through which monetary policy affects firms' investment into three different channels. First, we compute the "direct effect" of monetary policy by feeding in the path of the real interest rate R_t and hold all other prices fixed at steady state. Second, we feed in the series of the relative price of capital q_t but keep all other prices fixed. Finally, we feed in all other prices in the model—the relative price of output p_t, the real wage w_t, and inflation Π_t—but keep the real interest rate R_t and relative price of capital q_t fixed. Figure 7 plots the semielasticity of investment to each of these series in the initial period of the shock, conditional on a particular level of idiosyncratic TFP.

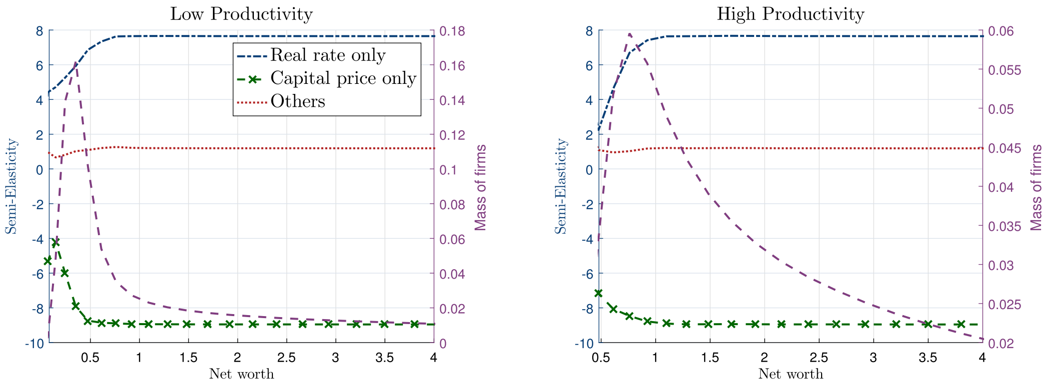

FIGURE 7.—Decomposition of semielasticity of capital to monetary policy shock. Notes: Semielasticity of capital in response to monetary shock and the initial distribution of firms. "Real rate only" refers to feeding in the path of the real interest rate R_t but keeping all other prices fixed at steady state. "Capital price only" refers to feeding in the path of the relative price of capital q_t but keeping all other prices fixed at steady state. "Others" refers to feeding in all other prices but keeping the real rate and capital price fixed at steady state. Dashed purple line is the initial distribution of firms. Left panel plots for low level of idiosyncratic productivity z. Right panel plots for high level of idiosyncratic productivity z.

The results in Figure 7 indicate that the heterogeneous responses in our model are driven by the fact that firms with high default risk face a steeper marginal cost curve for financing investment. The decrease in the real interest rate shifts out the marginal benefit curve in Figure 2; firms with higher default risk—which, in Figure 7, have low net worth—are less responsive to this change because they have a steeper marginal cost curve. The increase in the relative price of capital has two offsetting effects. On the one hand, it makes new capital more expensive and, therefore, shifts the marginal cost curve up for all firms; on the other hand, it increases the recovery rate of lenders in the event of default, which flattens out the marginal cost curve for firms with low net worth. On net, the former force outweighs the latter, but firms with low net worth are more strongly affected by the latter. Finally, the increase in all other prices also shifts out the marginal benefit curve by increasing the revenue product of capital; once again, this effect is offset for low net worth firms by their steeper marginal cost curve.[20]

The fact that both the direct and indirect effects play a quantitatively important role in driving the investment channel of monetary policy contrasts with Auclert (2019)'s and Kaplan, Moll, and Violante (2018)'s decomposition of the consumption channel. In the context of a household's consumption-savings problem, they find that the contribution of the direct effect of lower real interest rates is small relative to the indirect general equilibrium effects of higher labor income. In our model, direct interest rate effects are stronger because firms are more price-sensitive than households. In fact, without any financial frictions at all, the partial equilibrium elasticity of investment with respect to interest rates would be nearly infinite (see Koby and Wolf (2020) for a discussion of the role of interest-elasticities in heterogeneous firm macromodels). In contrast, households are less price sensitive because of consumption-smoothing motives, so these direct effects are less important.

[20]The changes in these other prices also increase net worth by increasing cash flow and increasing the relative value of undepreciated capital, which moves firms along the x-axis of Figure 7.

TABLE VIII

AGGREGATE RESPONSE DEPENDS ON INITIAL DISTRIBUTION[a]

(everything rel. to steady state)	Bad distribution	Medium distribution
Avg. capital response	0.67	0.84
Avg. net worth	0.48	0.75
Frac. risky constrained	1.37	1.17

[a] Dependence of aggregate response on initial distribution. We compute the change in aggregate capital for different initial distributions as described in the main text. “Bad distribution” corresponds to $\widehat{\omega} = 1$ and “Medium distribution” corresponds to $\widehat{\omega} = 0.5$.

6.3. *Aggregate Implications of Financial Heterogeneity*

In this subsection, we illustrate two ways in which financial heterogeneity matters for understanding the aggregate transmission mechanism. We first show that the aggregate effect of a given monetary shock may be smaller when the initial distribution of firms features higher default risk. Nevertheless, we show that the aggregate effect of monetary policy is *larger* in our model than in a comparable version of the model without any financial frictions (which collapses to a representative firm).

State Dependence of Aggregate Transmission. In order to illustrate the quantitative scope for state dependence, we fix the semielasticity of capital with respect to monetary policy as a function of firms' state variables and vary the initial distribution of firms. We vary the initial distribution of firms $\mu(z, n)$ by taking the weighted average of two reference distributions. The first reference distribution is the steady-state distribution $\mu^*(z, n)$. The second reference distribution $\widetilde{\mu}(z, n)$ assumes that the conditional distribution of net worth for every level of productivity is equal to the distribution of net worth conditional on a low realization of productivity in steady state. We then compute the initial distribution as a weighted average of these two reference distributions, $\mu(z, n) = \hat{\omega}\widetilde{\mu}(z, n) + (1 - \hat{\omega})\mu^*(z, n)$.

Table VIII shows that the average response of capital accumulation is 33% smaller starting from the low net-worth distribution $\widetilde{\mu}(z, n)$ than starting from the steady state distribution $\mu^*(z, n)$. In that distribution, average net worth is 52% lower and there are 37% more risky constrained firms in the low-net worth distribution than in the steady state distribution. Placing a weight $\hat{\omega} = 0.5$ on the steady state distribution increases the aggregate capital response, but it is still 16% below the response starting from steady state.

These results suggests a potentially powerful source of time-variation in the aggregate transmission mechanism: monetary policy is less powerful when net worth is low and default risk is high. Of course, a limitation of this analysis is that we have varied the initial distribution exogenously. The natural next step in this analysis is to incorporate various business cycle shocks into our model and study the shapes of the distributions that actually arise in equilibrium. We also emphasize that aggregate state dependence is an implication of the microlevel behavior of our model and has not been validated using aggregate time-series evidence.[21]

[21] Tenreyro and Thwaites (2016) provided time-series evidence that monetary shocks are less powerful in recessions, which is broadly consistent with the implication of our model to the extent that firm-level net worth falls in recessions. However, it is difficult to estimate the contribution of default risk alone in driving this result given that the changes in the distribution are slow-moving and highly correlated with other relevant factors in the time series.

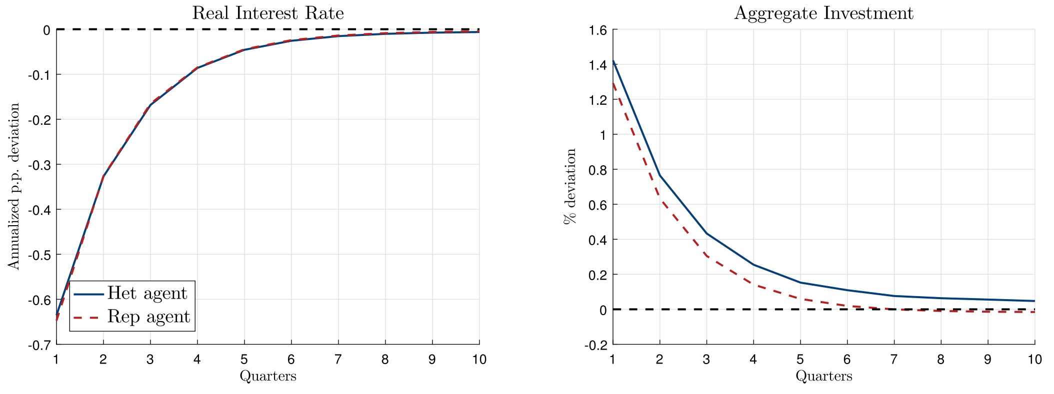

FIGURE 8.—Aggregate impulse responses in full model vs. rep firm model. Notes: "Het agent" refers to calibrated heterogeneous firm model from the main text. "Rep agent" refers to a version of the model in which the heterogeneous production sector is replaced by a representative firm with the same production function and no financial frictions.

Comparison to Frictionless Model. We now compare our full model to a model in which we eliminate financial frictions. We do so by removing the nonnegativity constraint on dividends; in this case, the investment block of the model collapses to a financially unconstrained representative firm (see Khan and Thomas (2008) Appendix B). Figure 8 shows that the impact effect of monetary policy on investment is larger in our full model than in the representative firm benchmark. Hence, despite the fact that risky constrained firms are less responsive than risk-free constrained firms, both types of constrained firms are more responsive than in a model without financial frictions because expansionary monetary policy increases firms' net worth.

7. CONCLUSION

In this paper, we have argued that financial frictions dampen the response of investment for firms with high default risk. Our argument had two main components. First, we showed in the microdata that firms with high leverage or low credit ratings invest significantly less than other firms following a monetary policy shock. Second, we built a heterogeneous firm New Keynesian model with default risk that is quantitatively consistent with these empirical results. In the model, monetary policy stimulates investment through a combination of direct and indirect effects. High-risk firms are less responsive to these changes because their marginal cost of investment finance is steeper than that of low-risk firms. The aggregate effect of monetary policy is primarily driven by these low-risk firms, which suggests a novel form of state dependence: monetary policy may be less powerful when default risk in the economy is higher.

Our results may be of independent interest to policymakers who are concerned about the distributional implications of monetary policy across firms. An often-discussed goal of monetary policy is to provide resources to viable but credit constrained firms. Many policymakers' conventional wisdom, built on the financial accelerator mechanism, suggests that constrained firms will significantly increase their capital investment in response to expansionary monetary policy. Our results imply that, instead, expansionary policy will stimulate the less risky firms in the economy to invest.

REFERENCES

ANDREWS, I., M. GENTZKOW, AND J. SHAPIRO (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of Economics*, 132 (4), 1553–1592. [2492]

AUCLERT, A. (2019): "Monetary Policy and the Redistribution Channel," *American Economic Review*, 109 (6), 2333–2367. [2475,2498]

BERNANKE, B. S., M. GERTLER, AND S. GILCHRIST (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework," Handbook of Macroeconomics, 1341–1393. [2474,2475,2482,2489,2491, 2492]

CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45. [2474,2491,2494]

CLEMENTI, G. L., AND B. PALAZZO (2016): "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations," *American Economic Journal: Macroeconomics*, 8 (3), 1–41. [2493]

CLOYNE, J., C. FERREIRA, M. FROEMEL, AND P. SURICO (2018): "Monetary Policy, Corporate Finance and Investment," National Bureau of Economic Research w25366. [2475,2481,2482,2490]

COOK, T., AND T. HAHN (1989): "The Effect of Changes in the Federal Funds Rate Target on Market Interest Rates in the 1970s," *Journal of Monetary Economics*, 24, 331–351. [2476]

COOPER, R. W., AND J. C. HALTIWANGER (2006): "On the Nature of Capital Adjustment Costs," *The Review of Economic Studies*, 73, 611–633. [2492]

CROUZET, N., AND N. MEHROTRA (2020): "Small and Large Firms Over the Business Cycle," Working paper, persistent link, https://EconPapers.repec.org/RePEc:cen:wpaper:18-09. [2477,2492]

DAVIS, S. J., R. J. FABERMAN, J. C. HALTIWANGER, AND I. RUCKER (2010): "Adjusted Estimates of Worker Flows and Job Openings in JOLTS," *Labor in the New Economy*, 71, 187. [2493]

DINLERSOZ, E., S. KALEMLI-OZCAN, H. HYATT, AND V. PENCIAKOVA (2018): "Leverage Over the Life Cycle and Implications for Firm Growth and Shock Responsiveness," National Bureau of Economic Research w25226. [2477]

FOSTER, L., J. HALTIWANGER, AND C. SYVERSON (2016): "The Slow Growth of New Plants: Learning About Demand?" *Economica*, 83, 91–129. [2483]

GERTLER, M., AND S. GILCHRIST (1994): "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *The Quarterly Journal of Economics*, 109 (2), 309–340. [2475,2481]

GILCHRIST, S., AND E. ZAKRAJŠEK (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102 (4), 1692–1720. [2477]

GOMES, J., U. JERMANN, AND L. SCHMID (2016): "Sticky Leverage," *American Economic Review*, 106 (12), 3800–3828. [2489]

GORODNICHENKO, Y., AND M. WEBER (2016): "Are Sticky Prices Costly? Evidence From the Stock Market," *The American Economic Review*, 106, 165–199. [2476]

GURKAYNAK, R. S., B. SACK, AND E. T. SWANSON (2005): "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," International Journal of Central Banking. [2476]

JEENAS, P. (2019): "Firm Balance Sheet Liquidity, Monetary Policy Shocks, and Investment Dynamics," Working paper. [2475,2482,2489]

JORDA, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *The American Economic Review*, 95, 161–182. [2480]

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108 (3), 697–743. [2475,2490,2498]

KARABARBOUNIS, L., AND B. NEIMAN (2013): "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 129, 61–103. [2490]

KHAN, A., AND J. THOMAS (2008): "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," *Econometrica*, 76, 395. [2500]

KHAN, A., T. SENGA, AND J. THOMAS (2016): "Default Risk and Aggregate Fluctuations in an Economy With Production Heterogeneity," Working paper, persistent link, https://EconPapers.repec.org/RePEc:red:sed017:889. [2475,2482,2492]

KOBY, Y., AND C. K. WOLF (2020): "Aggregation in Heterogeneous-Firm Models: Theory and Measurement," Working paper. [2475,2498]

MCKAY, A., AND J. F. WIELAND (2019): "Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy," National Bureau of Economic Research w26175. [2475]

MCKAY, A., E. NAKAMURA, AND J. STEINSSON (2016): "The Power of Forward Guidance Revisited," *American Economic Review*, 106 (10), 3133–3158. [2475]

MIRANDA-AGRIPPINO, S., AND G. RICCO (2018): "The Transmission of Monetary Policy Shocks," CEPR discussion paper no. DP13396, available at SSRN, https://ssrn.com/abstract=3302647. [2481]

OTTONELLO, P., T. WINBERRY (2020): "Supplement to 'Financial Heterogeneity and the Investment Channel of Monetary Policy'," *Econometrica Supplemental Material*, 88, https://doi.org/10.3982/ECTA15949. [2476]

REITER, M., T. SVEEN, AND L. WEINKE (2013): "Lumpy Investment and the Monetary Transmission Mechanism," *Journal of Monetary Economics*, 60 (7), 821–834. [2475]

SCHAEFER, S. M., AND I. A. STREBULAEV (2008): "Structural Models of Credit Risk Are Useful: Evidence From Hedge Ratios on Corporate Bonds," *Journal of Financial Economics*, 90 (1), 1–19. [2477]

TENREYRO, S., AND G. THWAITES (2016): "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 8 (4), 43–74. [2475,2499]

VAVRA, J. (2013): "Inflation Dynamics and Time-Varying Volatility: New Evidence and an SS Interpretation," *The Quarterly Journal of Economics*, 129 (1), 215–258. [2475]

WILMER CUTLER PICKERING HALE AND DORR LLP (2017): "2017 IPO Report," Technical report. Available at https://www.wilmerhale.com/en/insights/publications/2017-ipo-report. [2495]

WONG, A. (2019): "Refinancing and the Transmission of Monetary Policy to Consumption," working paper. [2475,2476]

ZWICK, E., AND J. MAHON (2017): "Tax Policy and Heterogeneous Investment Behavior," *American Economic Review*, 107 (1), 217–248. [2480]

Co-editor Giovanni L. Violante handled this manuscript.

Manuscript received 17 December, 2017; final version accepted 19 June, 2020; available online 14 July, 2020.