Firm Cyclicality and Financial Frictions*

Alex Clymo† Filip Rozsypal‡

January 31, 2023

Abstract

Using administrative micro data we document how firms' sensitivities to business cycles differ by size and age. Among the youngest firms, small firms are more cyclical than large, but the reverse is true among older firms. The differences in cyclicality are large: "young and small firms" are more cyclical than large firms, who respond one-and-half to one to the aggregate business cycle. In contrast, "old and small" firms are closer to acyclical. High leverage firms are more cyclical than low leverage firms which—when combined with the age-profiles and cyclicalities of financial variables—suggests that financial frictions are likely to explain the excess cyclicality of "young and small" firms, but not of large firms. Augmenting a dynamic heterogeneous-firm model with heterogeneous returns-to-scale and entrant wealth allows it to replicate these findings, and implies that financial policies targeted at young firms become less effective in stimulating aggregate output while the opposite is true for direct labor subsidies.

Keywords: firm age, firm size, cyclicality, financial frictions

JEL codes: D22, E32, G32, L25

*The viewpoints and conclusions stated in this paper are the responsibility of the individual contributors, and do not necessarily reflect the views of Danmarks Nationalbank or Statistics Denmark. We thank Gian Luca Clementi, Michaela Elfsbacka Schmöller, Simon Gilchrist, Dino Palazzo, and Johannes Poeschl. as well as the participants of Firm Heterogeneity 2021 workshop and 2nd Research ESCB Cluster conference for their helpful comments and feedback. Karl Bertelsen Robak provided outstanding research assistance.

†Department of Economics, University of Essex. Wivenhoe Park, Colchester CO4 3SQ, United Kingdom. Email: `a.clymo@essex.ac.uk`

‡Research Unit, Danmarks Nationalbank, Langelinie Allé 47, 2100 København Ø, Denmark and Center for Macroeconomics (CFM). Email: `firo@nationalbanken.dk`

1 Introduction

There are systematic and significant differences how firms of different ages and sizes react to the business cycle. Why are certain firms more sensitive than others? Do size and age just act as a proxy for financial frictions which amplify responses to shocks, as suggested by Gertler and Gilchrist (1994)? In this paper we submit evidence that financial frictions do increase sensitivity to shocks, but mostly for young firms which constitute a relatively small fraction of aggregate GDP and the cyclicality of large firms comes from elsewhere.

Using firm-level administrative and balance sheet data from the universe of Danish firms, we make two main contributions. Firstly, we empirically document the distribution of firm cyclicality, as well as other characteristics, jointly over firm size and age. We find two sets of patterns; one for young firms which is consistent with financial frictions, and one for older firms for which we propose, and test empirically, an alternative mechanism stemming from firm-level differences in returns to scale. Secondly, we use a quantitative heterogeneous-firm model to show that our facts present a challenge to the basic financial frictions model, and alter model-based estimates of the quantitative effects of recession-fighting policies which target different groups of firms.

In our first set of empirical results we measure firm cyclicality by joint age-size bin by regressing firm-level growth rates on aggregate GDP growth. We measure cyclicality using the comovement of firm-level employment and sales with aggregate output, extending the methodology of Crouzet and Mehrotra (2020), to study the role of size and age. We find that "young and small" firms are the most cyclical in the economy, followed by large firms (regardless of age), while "old and small" firms are the least cyclical. Put differently, among young firms cyclicality decreases with size, while it increases with size for older firms. This finding highlights the importance of studying firm age and size together over the business cycle: Crouzet and Mehrotra (2020) find a limited role of size on cyclicality in their US dataset, but, with our richer dataset, we are able to show strong effects of size on cyclicality, which just happen to have opposite signs for young and old firms and hence offset each other in the aggregate.

Our second set of empirical results investigate the role of finance in driving cyclicality by firm age and size. We find suggestive evidence for financial frictions for younger firms, but much less so for older firms. Since we have financial data at the firm level even for small and young (including unlisted) firms, we are able to characterise the lifecycle and cyclical properties of financial variables across the whole firm distribution. Starting with the firm lifecycle, we find that leverage (debt over assets) is higher at young firms, even conditioning on firm size. This is consistent with models of the financial accelerator where younger, smaller firms are more financially constrained. Moreover, young firms also have the highest growth rates of debt and leverage, suggesting that they are actively trying to increase their debt, and hence more likely to be affected if the access to debt is restricted in a recession. Older firms, on the other hand, all have shrinking leverage ratios, showing that they are reducing their reliance on debt (presumably either by accumulating retained earnings or switching to equity financing). Finally, after controlling for age, firms of all sizes have very similar leverage ratios on average. Financial constraints are thus unlikely to bind more for larger firms than small firms, conditional on age, suggesting that a different mechanism is needed to explain the the positive effect of size on cyclicality of older firms.

We then move on to directly studying the relationship between leverage and cyclicality, and find that high leverage firms are more cyclical than low leverage firms. Conditioning on firm size does not change this finding, while conditioning on firm age reduces the excess cyclicality of high-leverage firms by between 30% and 90%. Combined with our lifecycle findings, we take this as evidence that differences in cyclicality by age are more likely to be related to finance than differences in cyclicality by size.[1]

Our theoretical contribution is to demonstrate the importance of capturing the heterogeneity of cyclicality along the full joint size and age distribution. We build a quantitative heterogeneous firm model which balances financial and real mechanisms to match our rich set of empirical moments. We build on the seminal work of Khan and Thomas (2013), who, building on insights of Bernanke and Gertler (1989); Kiyotaki and Moore (1997); Bernanke et al. (1999) and Jermann and Quadrini (2012), set up a heterogeneous firm model with financial frictions, which they use for business cycle analysis. We extend the model to match the cross-sectional age and size distributions of firms, and start by demonstrating that a standard model with financial frictions can be made to match the

[1] Given the nature of our dataset, our empirical investigation of financial frictions does not have an ambition to achieve causal identification. However, our results are consistent with evidence with direct identification of financial friction via firm-bank matches, such as Chodorow-Reich (2014), that finds that firms borrowing from financially distressed banks contracted more during the Great Recession, and that this effect is largest at smaller firms.

cross-sectional moments but is unable to qualitatively replicate the cyclicality results by joint age-size bin. We thus propose two extensions to the model in order to allow it to match our empirical evidence.

The first extension is permanent differences in returns to scale across firms. Making some firms have more decreasing returns to scale naturally makes them smaller, and also endogenously makes them less responsive to shocks and hence less cyclical, as in the data. Put differently, small firms are fundamentally different from large firms (think a local shop versus a large multinational corporation) and these differences in firm scope help explain why small firms are less cyclical than large.[2] We empirically explore the relationship between firm size and returns to scale and find that small firms indeed tend to have more decreasing returns to scale, enough to quantitatively match the difference in cyclicality by firm size among older firms through a non-financial channel.

Our second extension is differences in the initial net worth that firms enter with, depending on their size group. Specifically, the estimation finds that large entrants start less financially constrained than small entrants. This reduces the cyclicality of larger young firms, as in our data, giving a financial explanation for the size gradient among young firms. This additionally matches our lifecycle finding that larger entrants grow less fast than smaller entrants. With these two twists, our new model is able to replicate the cyclicality of firms by joint age-size bin from our data.

Finally, we show that our results have important bite, as the insights from the model can be used to better design stabilisation policies, both in terms of how much different policies can move aggregate output, but also which instruments to use to target specific types of firms. Our empirical results suggest that large young firms are less financially constrained than in the standard model. This dramatically reduces the power of an "age based" policy, which targets young firms by offering debt relief. In contrast, a policy of simply offering a wage bill subsidy to all firms becomes more powerful, since large firms now respond more strongly to the policy. These results show that it is crucial to understand and match the responsiveness of the full age-size distribution of firms in order to perform robust policy evaluation through the lens of structural models.

[2] Gavazza et al. (2018) also use permanent differences in returns to scale to drive differences in firm size, in a model of recruiting intensity calibrated to the US economy. They find that this helps match why recruiting intensity (measured as the vacancy yield) is more cyclical at small firms than large firms, by creating financially unconstrained small firms. We show how, among financially unconstrained (older) firms, differences in returns to scale can explain why large firms are more cyclical than small firms.

Indeed, governments across the world engage in various types of policies that target SMEs, sometimes explicitly mentioning financial constraints.[3] These policies often help firms by funding them directly or subsidising commercial loans, and better targeting might have noticeable aggregate welfare effects because the volume of funding provided is substantial.[4] For example, in 2021 the European Investment Fund committed €26.2 billion of equity injections and debt instruments via local European banks.[5] Similarly, the US State Small Business Credit Initiative approved in 2021 channeled $10 billion in credit and investment programs for small businesses, on top of the over $4 billion already spent by 2010's Small Business Lending Fund.[6] To the extent that SMEs might be likely to be financially constrained in standard models, these policies have a clear rationale. Our empirical and theoretical results suggest that paying careful attention to the joint age-size distribution could improve these policies further.

Related literature: There is a broad literature studying the effect of age and size on firm decisions and outcomes. Different papers find different, sometimes conflicting, results, which is partly driven by different samples of firms available (by age and size) in a given dataset. Gertler and Gilchrist (1994) investigate the cyclicality of small versus large firms and find that small firms are more sensitive to periods of credit market tightening than large firms. Khan and Thomas (2013) show that small firms contracted more than large firms during the financial crisis, and Gavazza et al. (2018) show that the vacancy yield was more cyclical at small than large firms during this same period. On the other hand, Moscarini and Postel-Vinay (2012) find that larger firms (in terms of number of employees) are more cyclical, when aggregate conditions are measured using the (HP-filtered) level of the unemployment rate. Similarly, Mian and Sufi (2014) show that larger establishments contracted more in areas with larger declines in house prices.

More recently, it has been shown that firm age is a more important predictor of both the average level and cyclicality of firm growth than firm size (see Fort et al. (2013); Haltiwanger et al. (2013), for evidence from the US). Fort et al. (2013) discuss the conflicting

[3]For example, Denmark spent an additional €130 million on supporting SMEs during COVID outbreak "granting tax referrals and comparable measures *to ease liquidity constraints of SMEs*" (emphasis added).

[4]It is also worth noting that SMEs often receive exemptions from various regulations, both actual and administrative which are harder to evaluate in monetary terms.

[5]See the European Investment Fund website for details.

[6]See Small Business Programs by the US Treasury

results by firm size, and add age to this analysis. They find that young firms are more cyclical than old firms, and that this difference is much more important than the differential between small and large firms. While they do not have direct financial data at the firm level, they use state-level house price data to argue that financial frictions may drive this result.[7] Compared to the latter two papers, we do not use any local identification for financial shocks, but instead directly measure how firm-level financial variables vary over age-size bins.

Due to data limitations, much of the knowledge about cyclicality and firm finance is based on large publicly traded firms. Sharpe (1994) uses Compustat data to document that high leverage firms are more cyclical than low leverage firms. Giroud and Mueller (2017) combine Compustat data with establishment-level employment data to show that the decline in house prices during the Great Recession, as investigated by Mian and Sufi (2014), was transmitted to declines in employment through high leverage firms. Conversely, Ottonello and Winberry (2020) use Compustat data and find that firms with low default risk, including those with low debt burdens, are the most responsive to monetary shocks. Relative to their paper, our sample includes non-listed firms, and thus younger firms who may behave differently to financial frictions than older, listed firms. Jeenas (2019) investigates the role of liquidity and leverage in driving heterogeneous investment dynamics, and finds that leverage ceases to be important once liquidity is controlled for. However, publicly traded firms are only a small subset[8] and as such are not representative of the whole firm population. Cloyne et al. (2019) use data for the US and UK to show that younger, non-dividend paying firms exhibit the largest and most significant changes in investment following monetary policy shocks. Beyond focusing on public firms, they also measure age as time since incorporation due to data availability, rather than foundation. In contrast, we are able to measure age since foundation. We also focus on overall cyclicality rather than the response to identified monetary policy shocks.

[7] Another strand of literature examines the cyclicality of firm financing, both in terms of empirics and also model building. For example, see Jermann and Quadrini (2012) (investigate the cyclicality of debt and equity issuance), Covas and Haan (2011) (the cyclicality of financing is different across firms of different sizes, with the procyclicality of equity issuance decreasing monotonically with firm size), Crouzet (2017) (the choice of bank and bond financing), Begenau and Salomao (2018) (firm size and debt/equity cyclicality), Jensen et al. (2017) (size and cyclicality of financing and probability of default), Nikolov et al. (2018) (size and source of financial constraints) or Poeschl (2023) (size and cyclicality of debt maturity).

[8] In the US there are around 4000 publicly traded firms (Gupta et al., 2021) in the population of over 5 million firms.

The two papers closest to our work are Crouzet and Mehrotra (2020) and Dinlersoz et al. (2018). Both study the US, and go beyond Compustat to achieve wider firm coverage, so their results are not based only on public firms. Crouzet and Mehrotra (2020) find that only the largest firms (99th percentile and above measured by assets) are less cyclical than the rest, which goes in the opposite direction to our results where cyclicality increases in size. However, there is a large difference in the definition of size groups, as they focus on the top 10% of firms (measured by assets) while we investigate firms across the whole size distribution (measured by employees). We extend their empirical specification by including firm age, and more importantly, the interactions of size and age. We also study the cyclicality of employment, thereby bridging to the literature focusing on employment fluctuations and financial frictions (Chodorow-Reich (2014); Duygan-Bump et al. (2015)). Dinlersoz et al. (2018) merge balance sheet data from Compustat and Orbis into the US Longitudinal Business Database (LBD). Similarly to our analysis, they are able to analyse both private and public firms and can measure firms' employment and age since foundation. They argue that small private firms are plausibly financially constrained both before and after the financial crisis, while larger private firms may have only become constrained during the crisis, and large public firms appear to never be financially constrained. Kochen (2022) also uses Orbis data to study firm finance over age and finds that young firms in developing countries are more constrained than those in developed countries.[9]

The rest of this paper is organised as follows. In Section 2 we discuss the data, construction of our key variables and the estimation specification we use. In Section 3 we present our empirical results, and in Section 4 we present our quantitative model findings. Finally, in Section 5 we conclude.

2 Data

Our dataset covers firms in Denmark between 1991 and 2019 at an annual frequency. From 2002 onwards, the coverage is nearly universal, with the exception of financial corporations. In order to analyse firm outcomes and financial balance sheet data together,

[9] Alternative mechanisms that could work on top of the mechanism that we describe in this paper could be based on earnings-based borrowing constraints (Drechsel, 2023), or entrepreneur-banker relationship building (Casiraghi et al., 2021).

we merge two datasets ("data registers") provided by Statistics Denmark (DST): the FIRE dataset ("Regnskabsstatistikken"), which broadly contains data on accounting variables, is merged with the FIRM dataset ("Firmastatistik"), containing data regarding economic, employment and accounting information at company level. To extend the coverage to 1990's we use the precursor of FIRM and FIRE in form of FIGT (Gammel Firmastatistik), covering 1992-1999. Additionally, for the period 2003-2019 we also use firm bank loans data (both short and long) from URTEVIRK register ("Udlånsrenter").

We use two different measures of employment. First, to classify firms into size bins, we use the firm headcount. For the purposes of computing cyclicality, we switch to an hours-based measure of employment which gives the total hours at the level of firms measured in multiples of full time equivalent workers. Sales ("Omsætning") are based either on balance sheet information or on VAT declarations. Our measure of debt is available from 2001 and contains both short and long term liabilities. Specifically, beyond short and long term debt ("Anden langfristet gæld" and "Andenkortfristet gæld"), we also include provisions ("Hensættelser") - unknown obligations such as deferred tax or pension obligations, and Long and short term debt to suppliers ("Langfristet/Kortfristet gæld til leverandører"). For assets, we consider a combination of intangible ("Immaterielle anlægsaktiver"), tangible ("Materielle anlægsaktiver") and financial assets ("finansielle anlægsaktiver").

The quality of this data is generally believed to be very high, as Statistics Denmark is a government agency, and most of the variables we use are originally collected by Denmark's tax authority, SKAT.[10] Additionally, DST also runs independent checks on the datasets. Individual firms are identified by a unique number that is generated at the time of registration. The merging of the datasets is done using this identifier, and thus provides exact matches. More information on data itself and the cleaning process is provided in Appendix A.1.

Our cleaned dataset is an unbalanced panel capturing employer firms in Denmark. Our baseline sample is firm-year observations containing both valid accounting (e.g. sales, employment) and balance sheet (e.g. debt, assets) data.[11] Due to the availability of balance sheet data our baseline sample starts in 2001 and so runs from 2001 to 2019 (effec-

[10] Sales, assets, liabilities, investment and information about employment based on payroll.

[11] Specifically, if we are computing a firm-level growth rate variable between $t-1$ and t we only require there to be valid debt data in period t to include the growth rate in our dataset.

tively 2002 to 2019 since our main empirical specifications use growth rates and require the existence of lagged data). Moreover, while balance sheet data is available for firms of all sizes, small firms are sampled less than large firms, making our sample stratified. Nonetheless, we still observe firms of all sizes and ages, both publicly listed and privately owned, with positive probability in our dataset.[12] This makes our dataset uniquely suited to studying the role of financial frictions across the whole distribution of firms, especially at younger and smaller firms that are not featured in datasets like COMPUSTAT. We also report results for extended samples that do not require the existence of balance sheet data and that start in 1991 in Appendix B.4. Our baseline dataset contains roughly 2 million firm-year observations, with approximately 100,000 firms per year (in the early 2000's the number of firms is around 80,000 but it grows to 120,000 at the end of the sample).

Given the start and end date of the underlying registers one might reasonably worry whether our results might be overweighting the role of finance due to the financial crisis. While it would be theoretically possible to extend our sample by using alternative datasets that cover different time periods, we believe that there was enough other variation in Danish business cycle that other shocks are also well represented. According to the OECD[13], our sample covers the following business cycle turning points: troughs in 1993M7, 1995M12, 1998M5, 2003M7, 2009M7, and 2014M4, and peaks in 1994M12, 1997M5, 2000M8, 2006M7, 2011M4, and 2019M6. Denmark thus experienced at least three recessions in our sample: the early 90's, early 2000's and the global financial crisis. This means that while certainly dominant, the financial crisis is not the only recession in our dataset driving variation in aggregate GDP.

2.1 Key variables

Employment and sales can be informative about the situations firms find themselves. In a world with stochastic demand, a firm finds itself with fluctuating revenue. In the presence of labour market frictions that firm might prefer to insulate its workforce. Firms

[12]The stratified sampling reduces the frequency of observations for small firms, but this should just reduce the power of our empirical results, rather than altering the point estimates we find, for this group of firms. This is supported by the fact that our main cyclicality results for employment and sales do not change when keeping the same time period but also including firm-year observations even if they are missing balance-sheet data. See Appendix B.4.

[13]See OECD turning points.

that are financially constrained are perhaps limited in the ability to do so. Following this motivating example, we focus on behavior of employment, sales and how they interact with some measures of financial frictions, allowing for heterogenous effect by size and age.

On the production side, we use data on sales, employment (both headcount, which we use for definition of size groups and the number of full time equivalent workers, which we use in the regressions), profits and investment. On the financial side, we use data on total assets, total liabilities, and the stock of debt of all maturities. We use the ratio of debt to assets as a measure of leverage. We additionally use data on a firm's sector of operation so the results are not driven by differences between sectors.

We define firm size by its lagged employment (headcount).[14] The firm size measure thus changes as the firm grows or shrinks as it ages and is hit by shocks. We sort firms into bins based on four quantile thresholds (0-30th, 30-60th, 60-90th and 90+) of size across the population of firms active that year. In Figure 12 in Appendix A.2 we plot the employment thresholds defining these size bins, and how they evolve over time. The thresholds are relatively stable over time, with a minor expansion of the largest firms at the end of the sample.[15] As in other countries, the firm size distribution is heavily skewed: the top size bin (containing only the 10% largest firms) represents over 70% of aggregate employment.

At the same time, we also sort firms into four age groups: 0-3, 4-8, 9-19, and 20+ years old. The number of firms within each age group changes over time, and as Figure 12 shows, cyclical fluctuations in entry create swings in cohort size that propagate over the age distribution. Firm age is measured from the moment the firm is registered.[16] This notion of age is thus the true age since foundation of the firm, which distinguishes us from other datasets which can only measure age since, for example, the firm was publicly listed on stock markets. As with our size measure, in our empirical work we do not work with age directly, but put firms into our four age bins.

For our cyclical measure, we use the standard growth rate of aggregate GDP, which we denote as $y_t \equiv \frac{GDP_t - GDP_{t-1}}{GDP_{t-1}}$, collected from the DST National accounts. Firm-level

[14] For robustness, we also follow Crouzet and Mehrotra (2020), and measure firm size by the value of its assets (we report the alternative results in Appendix B.3).

[15] For assets, the pattern is similar, see the figures in Appendix B.3.

[16] Given that it takes very little time to start a new firm in Denmark, we believe there is not a large need to formally register the firm long before the firm becomes economically active.

Figure 1: Share of firms across age and size bins

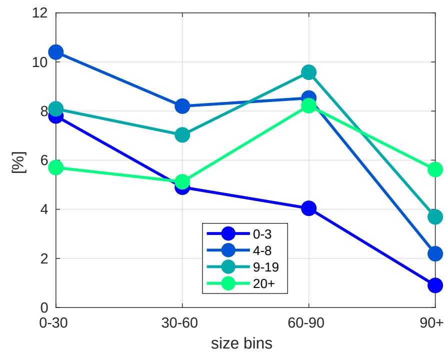

Fraction of observations in each joint age-size bin. Lines correspond to age bins and x-axis to size bins.

outcomes are defined the normalised growth rates suggested by Haltiwanger et al. (2013). Nominal variables are transformed into real by using aggregate CPI.

Each of the size-age bins contains different number of firms. Figure 1 plots the shares of all firm-year observations. Not surprisingly, the bin for the "youngest and largest" firms contains the smallest number of firms, roughly 0.8% of all observation. At the same time, this still means that meaningful number of firms do start very large. Most bins contain at least 4% of the sample, meaning that coverage is good across our age-size bins.

2.2 Estimation framework

To study the intricate interplay between firm size and age we allow for interactions between size and age bins. Therefore, the effect of being old, for example, is allowed to be different for small and large firms. Using the definition of groups from the previous section, we run a regression with a set of dummies controlling for the interaction of size and age. Formally, we run two types of regressions to get at the differences in levels and differences in cyclicality:

$$x_{i,t} = \sum_j \sum_k \alpha_{j,k} \mathbb{1}_{i \in I_t^j} \mathbb{1}_{i \in A(k)} + \sum_l \gamma_l \mathbb{1}_{i \in S(l)}, \tag{1}$$

$$\hat{g}_{x_{i,t}} = \sum_j \sum_k (\alpha_{j,k} + \beta_{j,k} y_t) \mathbb{1}_{i \in I_t^j} \mathbb{1}_{i \in A(k)} + \sum_l (\gamma_l + \delta_l y_t) \mathbb{1}_{i \in S(l)}, \tag{2}$$

where $\hat{g}_{x_{i,t}}$ denotes the firm-level normalised growth-rate of the variable of interest, such as turnover or employment at firm i. The indices j, k, and l index firm size bins, firm age bins, and firm sectors respectively.[17] $\mathbb{1}_{i \in I_t^j}$ is an indicator variable for firm i being in size group j at time t (and similarly $\mathbb{1}_{i \in A(k)}$ for age and $\mathbb{1}_{i \in S(l)}$ sector). Our baseline results are obtained by OLS estimation.[18] Primarily, we present the results graphically, combining the coefficients and plotting against the size, grouping the age bins by a line of age-specific colour.

The regression equation (1) is used to gain insight about the basic age-size distribution of variables of interest. We do so by grouping the coefficients α by age group and and plotting the values over firm size bin. The results should not strictly be interpreted as life-cycle profiles as the dataset is not a balanced panel due to firm exit.[19] We include sectoral controls so the results are to be interpreted as the within sector levels differences driven by age and size.

Equation (2) is used to study firm cyclicality. We regress firm level growth rates, $\hat{g}_{x_{i,t}}$, on aggregate growth, y_t, using dummy variables to separately estimate the cyclicality of different groups of firms. For each size bin, α_{jk} captures the marginal effect on the average growth rate of firms of being in that size bin. For these regressions, we are more interested in the β_{jk} parameters, which capture how the firm-level growth rates, $\hat{g}_{x_{i,t}}$, are differently related to the aggregate growth rate, y_t. The interpretation of β_{jk} is that a 1pp increase in aggregate growth is on average associated with a "β_{jk}"pp increase in firm-level growth for firms in size group j and age group k, on top of any additional effects captured by sector specific cyclicality. Thus, the β_{jk} capture the cyclicalities of each firm age-size group. Similarly, δ_l coefficients control the cyclicalities of the different sectors, to strip out the potentially differing average cyclicality of different industries. Thus, the effects when comparing coefficients from different, for example, age groups should be interpreted as within-industry effects. This specification is an extension of Crouzet and Mehrotra (2020)'s regression (their equation (1)) to include firm age categories interacted with the size bins.

[17] We use Danish 36 sector industrial classification DB07, based on NACE rev.2

[18] We also estimate the effects a corresponding specification with firm fixed effects, which hence only includes firms which move across size or age bin at least once over their life. The results are very similar to our baseline specification, as shown in Appendix B.5.

[19] For more discussion of the estimation of the effect of age using Danish firm micro data, see Andersen and Rozsypal (2021).

Table 1: Averages of Variables of Interest by Age and Size

	Age groups				Size groups			
	0-3	4-8	9-19	20+	0-30	30-60	60-90	90+
Employment	9.8	13.6	21.0	40.6	1.7	4.4	12.4	130.0
Sales	20993	32895	54765	121159	4857	10885	28497	363759
Assets	18172	32652	57960	141251	11022	19999	25445	375965
Debt	11590	19247	32340	74278	5553	11175	13917	208246
Equity	6581	13405	25620	66972	5469	8824	11527	167715
Bank loans	1128	2539	3970	8317	672	1059	2251	23699
Equity< 0	47.8	46.8	46.7	46.8	47.0	47.0	46.9	46.7
Bank loans> 0	50.3	60.1	63.9	68.2	48.0	59.8	69.1	80.4
D/A	0.85	0.79	0.70	0.62	0.76	0.75	0.73	0.68
C/A	0.17	0.16	0.15	0.11	0.18	0.15	0.13	0.09

Sales, assets, balance sheet debt, bank loans and equity in thousands of DKK (during the sample period 1000 DKK equaled 134 EUR due to the fixed exchange rate, or 150-200 USD). Reported numbers are the average values within in bin. Debt/assets (D/A = leverage) and Cash/Assets winsorized at 99.5th percentile. Only firms that do not exit in the current period. The fraction of firms with leverage equal to zero, negative equity and positive bank loans reported as percentages.

3 How does firm age and size determine firm outcomes?

In this section we present our empirical results, investigating firm averages and cyclicality across the joint firm age-size distribution, and the role of finance in driving these patterns. Our main focus is on the interaction between age and size, and we present results in a series of plots. These give the point estimates for the relevant coefficient from equations (1) and (2) for the $4 \times 4 = 16$ joint age-size bins.

3.1 Levels and growth rates of real variables

To give an overview of the distribution of various variables in the data, we provide a summary table with basic moments of the variables of interest over firm size or age in Table 1. The size distribution is wide, with the smallest (largest) bins having average employment of 1.7 (130) employees respectively, and similar differences for sales and financial variables. Firms also differ considerably by age: the youngest firm group has average employment of 9.8 while for the oldest it is 40.6.

Moving on to the joint age-size distribution, in Figure 2 we plot the average level and growth rate of sales and employment. Specifically, these are the coefficients α_{lk} from (1) and hence control for sectoral effects. Panels (a) and (b) show that there is little difference in average sales or employment across age bins within any size group. For employment, this is natural, as the size bins are defined by employment, but for sales this is not by construction.

Figure 2: Average levels and growth rates by size and age

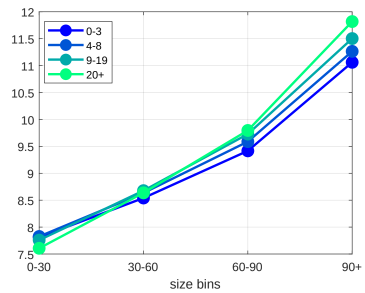

(a) (log) Sales

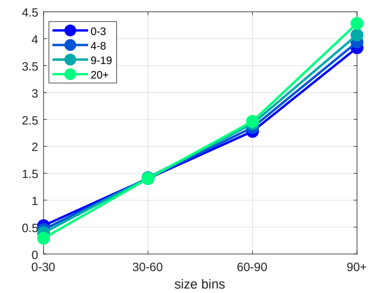

(b) (log) Employment

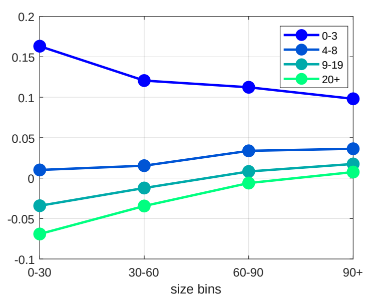

(c) Growth rate of sales

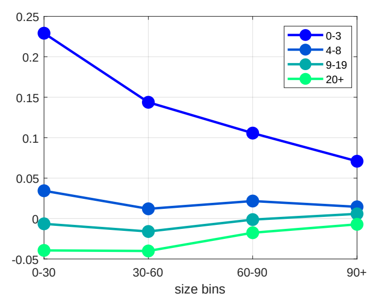

(d) Growth rate of employment

Average level and growth rate are computed as coefficient α_{lk} from regression (1) with log-level and growth rate as left-hand side variable respectively. Lines correspond to age bins and x-axis to size bins.

Panels (c) and (d) show the growth rates of the same variables by joint age-size group. In contrast to the levels, here we see a clear age effect, with the growth rate for all variables being the highest for the entrants. This suggests that on average firms start below their optimal size, and compared to entrants, older firms on average grow slower or shrink on average. "Old small" firms are thus fundamentally different from "young small" firms, as the former have reached their optimal size and are not growing anymore, while the later are growing by over 10% per year on average.

Moreover, we document a failure of Gibrat's law — the idea that growth rates should be independent of size — which is age dependent: Young firms (age 0-3) in the smallest age bin also grow the fastest, both for the input (employment) and also output (sales),

Figure 3: Cyclicality by Size and Age

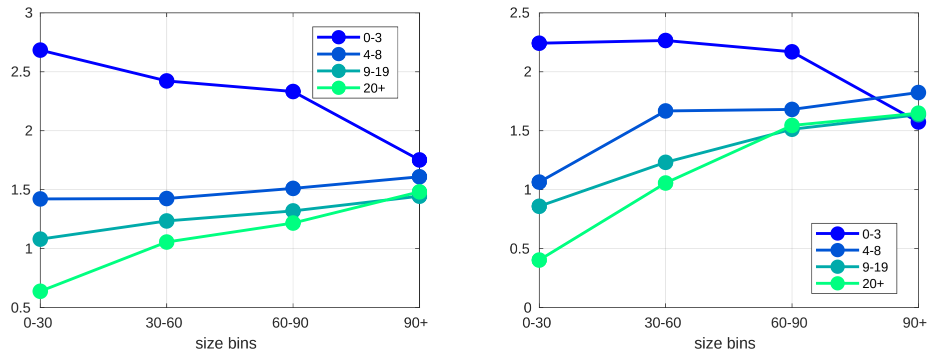

(a) Sales (b) Employment

Cyclicality is the coefficient β_{lk} corresponding to given size and age bin from regression (2). Lines correspond to age bins and x-axis to size bins.

suggesting that the dispersion of the starting size is actually larger than the dispersion in long run optimal size. For older firms this is reversed: larger firms grow faster than smaller firms for any age group beyond the entrants. In our view, these results highlight the importance of taking both age and size into account when understanding firm dynamics, and we show that this insight remains true when considering business cycle cyclicality in the following sections.

3.2 Cyclicality of real variables

In this section we investigate the cyclical sensitivity of firms by size and age. We do so without any reference to financial frictions, or other underlying causes of the differing levels of cyclicality. Thus, the results in this section are meant to be interpreted as theory free, and provide us with our basic stylised facts about firm cyclicality in the Danish economy. Following Crouzet and Mehrotra (2020) "cyclical sensitivity" refers to the extent that a worsening in aggregate conditions is systematically associated with declines in outcomes at firms of various groups.

We present our results from regression specification given by equation (2) for firm-level sales and employment. We have a separate cyclicality coefficient, $\beta_{j,k}$, for every age-size pair, which we plot in Figure 3. In the appendix, we give the regression results in Table 4. The 95% confidence intervals are narrow enough to statistically distinguish the difference in cyclicality between firms of age 0-3 and 20+ for the firms belonging to the smallest (0-30%) size bin.[20]

Panels (a) and (b) of Figure 3 show cyclicality by age-size bin for sales and employment respectively. The results show two general patterns which hold regardless of whether we measure cyclicality using sales or employment. First, younger firms are more cyclical across almost all of the size distribution. Second, the effect of size is different for entrants (the age 0-3 group) and for firms of all other sizes. For the entrants, the larger firms are on average less cyclical. In contrast, among all the older firms, larger firms tend to be more cyclical and gradient is the largest for the firms in the oldest age bin.

The corollary is that the difference in cyclicality between young vs. old firms is much larger among the smallest firms than among the largest firms. Indeed we detect a distinct non-linearity in the age-size relationship. In particular, for smaller firms, those in the 0-30th size percentile, the difference in cyclicality between young-small firms and old-small firms is dramatic. However, once we get the the largest (90%+) size bin, cyclicality is very similar for young and olds firms. In other words, all large firms are alike, but small firms can be very different, and small young firms and small old firms are not alike.

To place our results in context with the literature, it is interesting to compare our full joint age-size cyclicality results to simpler specifications which only investigate the role of size or age on cyclicality independently. We run these specifications, and present the results in Appendix B.2. When studied alone, we find that younger firms are more cyclical than old firms, consistent with the general view of the literature. But when we study size alone, we now find only a weakly positive relationship between size and cyclicality, which is less than half as strong as the positive relationship we find for the oldest firm group in Figure 3. Our more nuanced joint age-size regression naturally explains this result: Size alone does not predict cyclicality particularly strongly since the relationship between size and cyclicality has opposite signs for young and old firms, which roughly cancel out on average. This finding could help in reconciling the conflicting results about cyclicality by firm size discussed in the introduction.

[20] See Figure 13 in Appendix B.1.

3.3 Levels, growth rates, and cyclicality of financial variables

So far, we have presented results about the cyclicality of firms by joint age-size, without reference to any underlying theory which could explain these results. This places the results in a similar approach as the cyclicality results in Fort et al. (2013). In the following sections, we aim to go further, and use our firm-level financial data to provide evidence about the role of financial frictions in driving cyclicality by age and size.

In this section we analyse the level, growth, and cyclicality of firm debt and leverage by joint age-size bin, using the same approach (equations (1) and (2)) that we used for real variables. The results are presented in Figure 4. Starting with the levels of leverage (panel d), we find that younger firms are on average more leveraged than older firms, even conditioning on size. These firms are the likely candidates for being the most constrained, both because they already have the most (in relative terms) debt and the shortest track-record with lenders. The figure also reveals that, once age is controlled for, there is very little difference in leverage across firm groups: the size gradients are essentially flat, especially when compared to the vertical distance between the lines which represent the age gradient.[21] Since older firms are larger on average, this exercise shows the importance of studying age and size together when looking at finance: A simple correlation would find that smaller firms have higher leverage, but conditioning on joint age-size bins reveals that there is no within-age size gradient, and it is in fact young firms who are uniformly more leveraged.

Turning to the levels of debt (panel a), we see the reverse; older firms having more debt across all size groups. However, given what we already know about leverage (Debt/Assets) and size, this must simply represent the fact that larger firm have more assets. We can hypothesize that larger firms are probably less financially constrained than young-small firms. This conjecture is supported by the growth rate of leverage (panel e), which shows that young-small firms are the only firms with growing leverage. For all other firm groups, including all larger firms, they are shrinking their leverage on average. For large entrants (age 0-3 and size 60%+) the growth rate of both leverage is in line with the other large firms, which suggests entrants of this size are not being treated by banks any worse

[21] To the extent that there is a size gradient it has inconsistent signs across groups. For most firms the larger size seems to decrease the leverage (the exemption being the the smallest entrants and top 10% of firms in terms of size).

Figure 4: Average Levels, Growth Rates and Cyclicality of Debt and Leverage

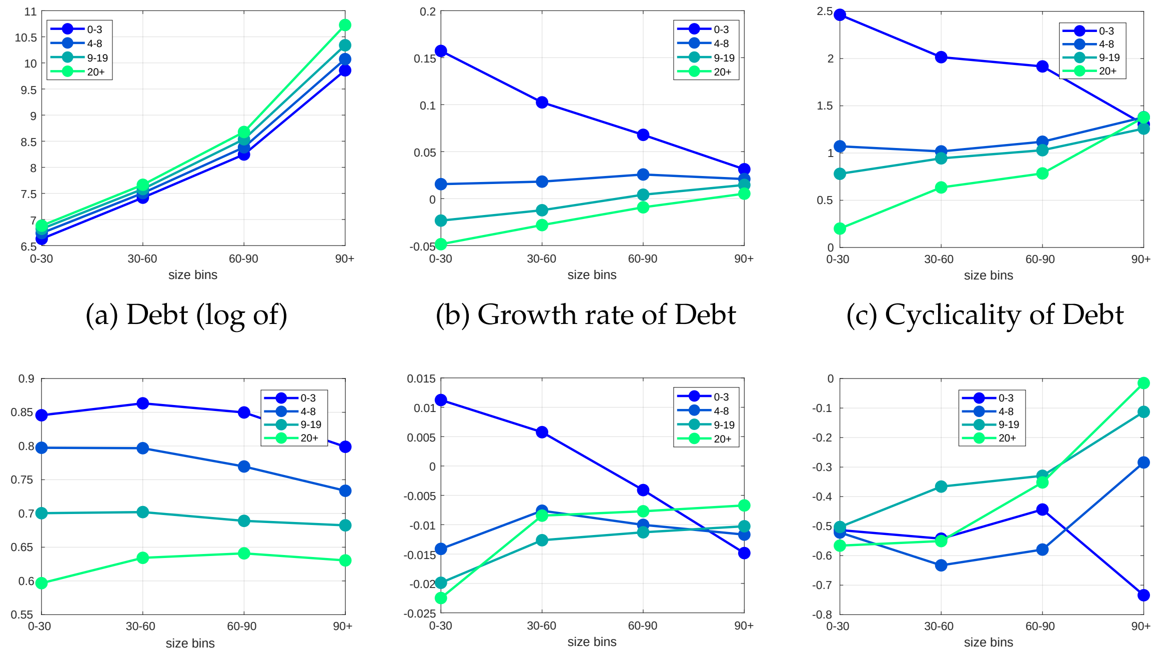

(a) Debt (log of)
(b) Growth rate of Debt
(c) Cyclicality of Debt
(d) Leverage (D/A)
(e) Growth rate of leverage
(f) Cyclicality of leverage

Average level and growth rate are computed as coefficient α_{lk} from regression (1) with log-level and growth rate as left-hand side variable respectively. Cyclicality is the coefficient β_{lk} corresponding to given size and age bin from regression (2). Lines correspond to age bins and x-axis to size bins. Leverage (Debt over assets (DA)) is winsorized at 99.5%.

than larger or older firms, and are likely not any more financially constrained. This distinction between the financial position of small versus large entrants will be important in our theoretical work.

Finally, turning to the cyclicality results, we see that debt (panel c) follows the same pattern as sales and other variables from Figure 3. However, leverage (panel f) seems to be countercyclical for almost all age-size groups, possibly because asset values decline more than debt in recessions. Given these valuation-based measurement issues in the cyclicality of leverage itself, we turn instead to studying whether the level of leverage affects the cyclicality of real variables in the next section.

3.4 How does finance affect firm cyclicality?

In this section we directly investigate the relationship between the level of a firm's leverage and the cyclicality of its real outcomes (sales and employment). We start with basic correlations before moving on how this interacts with the joint age-size distribution.

Basic regression: leverage only Consider the following regression specification:

$$\hat{g}_{x_{i,t}} = \sum_m (\omega_m + \psi_m y_t)\mathbb{1}_{i \in DA(m)} + \sum_l (\gamma_l + \delta_l y_t)\mathbb{1}_{i \in S(l)} \tag{3}$$

Here we use five bins—the leverage quintile $m = 1, ..., 5$ (lagged to avoid as much of simultaneity issues as possible)—to measure leverage $\mathbb{1}_{i \in DA(m)}$. This regression is a variation on (2), and controls for sector as before. In (3) we regress firm outcomes on the leverage bin and it's interaction with aggregate growth, with the coefficient ψ_m essentially giving a raw correlation between firm leverage level and firm cyclicality. We use employment growth and sales growth as the left hand side variables, and report results graphically in Figure 5(a). To facilitate comparisons across specifications, we subtract ψ_3 and so present cyclicality relative to the middle leverage bin.[22]

Figure 5(a) shows that firms with higher leverage are more cyclical, both measured in terms of sales growth and employment growth. This provides a direct link from finance to cyclicality, strongly suggesting that high leverage amplifies the business cycle by making more leveraged firms shrink more during recessions. The magnitude is large, and going from the bottom to top leverage quintile raises cyclicality by around 0.75 for employment and 0.2 for sales. This is equivalent to a non-trivial of the difference in cyclicality across certain firm age-size bins shown in Figure 3.

Interestingly, the figures also shows that leverage increases the cyclicality of employment more than the cyclicality of sales. This is consistent with our earlier motivating thought experiment: if cyclical fluctuations in sales are the result of fluctuations in demand and supply, this is at least to certain degree exogenous to what firms do which explains why leverage affects it relatively less. Employment, on the other hand, is a choice firm have to make, and firms might want to preserve their stock of workers if possible during recessions. If the firms with high leverage do not have the fiscal capacity to do so this would make the cyclicality of employment higher for firms with higher leverage.

Additive specification Having established that higher leverage is associated with higher cyclicality of real variables, a natural question is to what extent high leverage can explain the differences in cyclicality across age and size that we documented in Figure 3 and that motivate this paper. To do so, we first run a "horse race" type regression between

[22]Coefficients and standard errors for results from this section can be found in Table 3 in Appendix B.

Figure 5: The Effect of Leverage on Cyclicality of Sales and Employment

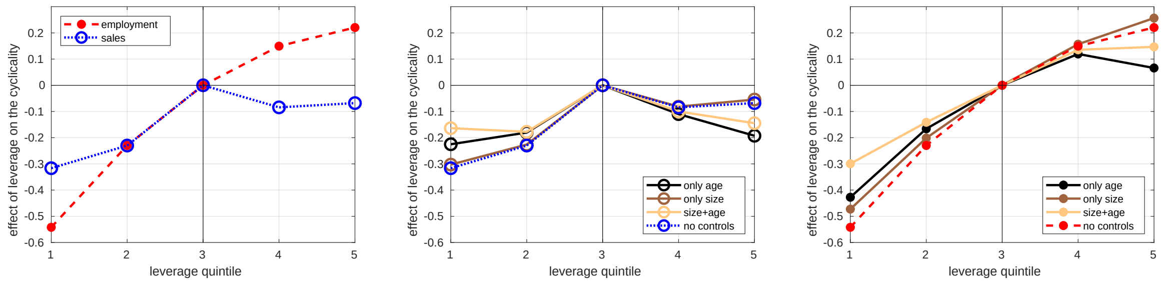

(a) No age-size controls (b) Sales: + Controls (c) Employment: + Controls

In panel (a), cyclicality is the coefficient ψ_m corresponding to given leverage bin from regression (3), with sales or employment growth as the left hand side variable. In panel (b) we repeat the sales cyclicality plot, with the additional lines giving ψ_m in regressions also controlling for age and size bins following (4). Panel (c) does the same for employment cyclicality. x-axis corresponds to leverage bins.

leverage and our original age-size bins:

$$\hat{g}_{x_{i,t}} = \sum_m (\omega_m + \psi_m y_t) \mathbb{1}_{i \in DA(m)} + \sum_j \sum_k (\alpha_{j,k} + \beta_{j,k} y_t) \mathbb{1}_{i \in I_t^j} \mathbb{1}_{i \in A(k)} + \sum_l (\gamma_l + \delta_l y_t) \mathbb{1}_{i \in S(l)} \quad (4)$$

In (4) we additionally include the original age-size interactions from our previous regression, so that the leverage coefficients ψ_m now give the effect of leverage on cyclicality which is not already explained by the age and size coefficients. In Figure 5(b) we plot the result for sales and in Figure 5(c) for employment. We also include the original regression (3) without controls, and versions of (4) where we control only for age and size bins, but not joint age-size bins.

The results suggest that finance does have some role to play in explaining the differences in cyclicality across firm age groups, but not size groups. Comparing the dashed "no controls" line and solid yellow "size + age" line reveals that the effect of leverage gets smaller (the gradient gets flatter) when age-size bins are also controlled for. Hence, some of the differences in cyclicality across age-size bin we previously documented are likely to be due to differences in leverage. However, further inspection of panels (b) and (c) reveals that it is only adding the age bins (solid black) which causes the leverage gradient to flatten, and not adding size bins (solid brown), which has almost no effect on the gradient. Adding the age bins alone reduces the gradient, measured as the difference between ψ_5 and ψ_1 by nearly 90% (0.02 vs. 0.2) for sales and 33% (0.25 vs. 0.75) for age, relative to the regression without any age-size controls. Our interpretation is that leverage, and financial frictions more generally, can explain why young firms are more cyclical than old firms, but not why large firms are more cyclical than small firms.

Triple interaction specification Our final empirical exercise is to explore the relationship between finance, joint age-size bin, and cyclicality in a triple interaction specification:

$$\hat{g}_{x_{i,t}} = \sum_m \sum_j \sum_k (\alpha_{j,k,m} + \beta_{j,k,m} y_t) \mathbb{1}_{i \in I_t^j} \mathbb{1}_{i \in A(k)} \mathbb{1}_{i \in DA(m)} + \sum_l (\gamma_l + \delta_l y_t) \mathbb{1}_{i \in S(l)} \tag{5}$$

In (5) the cyclicality of firms in finance bin m, size bin j and age bin k now measured by the coefficient $\beta_{j,k,m}$. Hence the effect of the finance bin m on cyclicality is allowed to differ for each age-size bin, which breaks down the average effect of finance on cyclicality in specifications (3) and (4) into the role of each age-size group. Since we are now estimating many more coefficients, we now restrict our size bins to terciles as some bins would otherwise have too few observations.

We present the results graphically in Figure 6 for employment cyclicality (i.e. where the left hand side variable is firm-level employment growth). In panels (a) to (c) we plot the coefficients $\beta_{j,k,m}$ for the low, middle, and highest leverage quintile ($m = 1, 3, 5$) respectively. In panels (d) to (f) we present the same results but now measuring financial position not by the level of leverage, but by the previous year growth rate of bank loans at the firm. The idea here is that firms that have recently increased their bank borrowing might have exhausted their borrowing capacity and hence be unable to borrow further, and so this serves as another potential measure of financial frictions.

We start by discussing panel (b), which plots cyclicality across joint age-size bins for the middle leverage quintile ($\beta_{j,k,3}$). We see that even within this narrow leverage rage, our basic fan-shaped picture from Figure 3 is clearly visible. This suggests that there is still some part of the basic differences in cyclicality across age-size bins which is *not* due to finance.

Firstly, in panel (e) we again plot cyclicality across joint age-size bins, but this time for the middle lagged bank debt growth quintile. Here we see that the positive relationship between size and cyclicality persists, but the excess cyclicality of young-small firms has now disappeared (i.e. the age 0-3 line no longer starts elevated and slopes down). This measure of financial frictions therefore suggests that all of the excess cyclicality of young-small firms is related to finance.

Secondly, going from left to right across the panels shows the effect of increasing lever-

Figure 6: Employment cyclicality: Heterogeneous Effect of Leverage by Size and Age

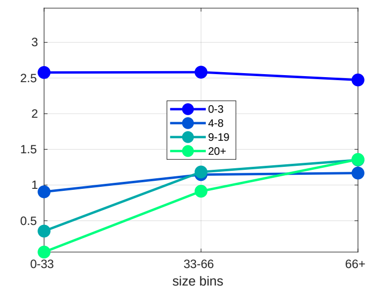

(a) Lowest leverage quintile

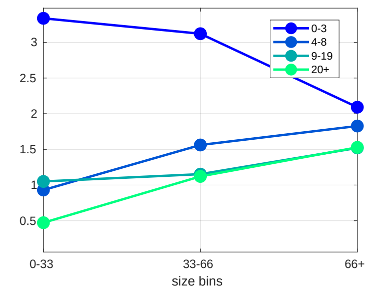

(b) Middle leverage quintile

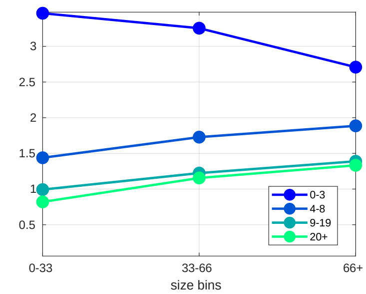

(c) Highest leverage quintile

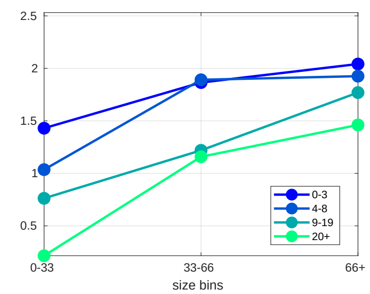

(d) Lowest %Δ loan quintile

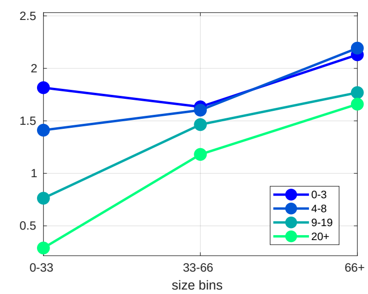

(e) Middle %Δ loan quintile

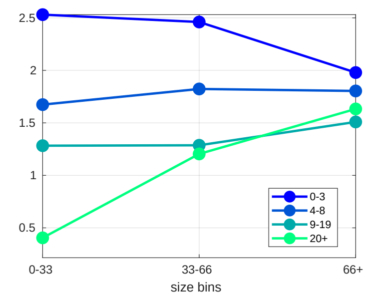

(f) Highest %Δ loan quintile

Cyclicality is the coefficient $\beta_{j,k,m}$ corresponding to given finance bin (m) size bin (j) and age bin (k) from regression (5), with employment growth as the left hand side variable. Panels (a) to (c) use leverage to measure finance, plotting $\beta_{j,k,m}$ for $j = 1, 3, 5$ respectively. Panels (d) to (f) do the same instead using past bank debt growth to measure finance. Lines correspond to age bins and x-axis to size bins.

age (moving from (a) to (c)) or increasing bank debt growth (panels (d) to (f)) on cyclicality. Comparing panels (a) and (c), for example, allows us to see how higher leverage changes cyclicality for each age-size bin individually. What stands out is that rising leverage seems to have the clearest effect on young-small firms, and less of an effect on older firms: as we move from (a) to (c), the age 0-3 line moves up and slopes down more, making the fan shaped pattern more pronounced. This is even clearer when looking at bank debt growth, as the original fan-shaped pattern only appears as we move to panel (f) and the age 0-3 line shifts up.

The comparisons in Figure 6 suggest that finance has the strongest effect on young-small firms. To show this more clearly, we summarise these same coefficients in a different way in Figure 7. Here we plot the *the financial cyclicality gradient*, which we define as the difference in cyclicality between firms moving from the highest to lowest finance bin: for any j, k, this is $\beta_{j,k,5} - \beta_{j,k,1}$. Panels (b) and (d) summarise the previously shown results for employment cyclicality with finance measured with leverage and bank debt growth respectively. Panels (a) and (c) present the same for sales cyclicality. Given the number of coefficients the results are somewhat noisy, but do suggest a story through the broad patterns.

Firstly, the effect of finance on cyclicality is typically largest for the young-small (age 0-3 size 0-33%) group across the most of the four panels. Hence, even if financial position did vary across two firms, this has the largest effect on cyclicality for young-small firms. This supports the idea that finance is a driver of cyclicality for young-small firms, more so than for other firm groups. Secondly, among large firms of all ages the financial cyclicality gradient is typically small, with coefficients for the $66+\%$ size group typically hovering close to zero. This suggests that our original finding in Figure 3 that large firms are more cyclical than small firms (among older firms) is not likely to be driven by financial frictions, and another explanation is needed. Finally, in some specifications, for example the cyclicality of employment across leverage bins in panel c, the financial cyclicality gradient is large even for old-small firms, and not just for young-small firms. This suggests that finance might have a role in driving cyclicality across the size distribution too, but this only occurs for the smallest firms which is why it was not picked up in specification (4). Nonetheless, since old-small firms are the least cyclical in the economy, to the extent that finance makes them more cyclical they were not that cyclical to begin with. This leaves finance as a strong suggested driver of high cyclicality only for the young-small firms in the economy.

3.5 Summary of empirical evidence

We have presented results on the cyclicality of firms by age and size, their interaction, and the role of finance. The central motivating finding is that cyclicality differs in non-monotone ways across the age-size distribution: "young and small" firms are the most cyclical, large firms are the second most cyclical, and "old and small" firms are the least cyclical. We then turn to understanding what drives this pattern.

The patterns that we observe for financial variables are in line with financial frictions affecting *young and small* firms more than the rest, and explaining why they are the most cyclical. In particular, young firms have higher leverage than old firms, while leverage does not vary by firm size once age is controlled for. Among young firms, the smallest

Figure 7: Financial cyclicality gradient

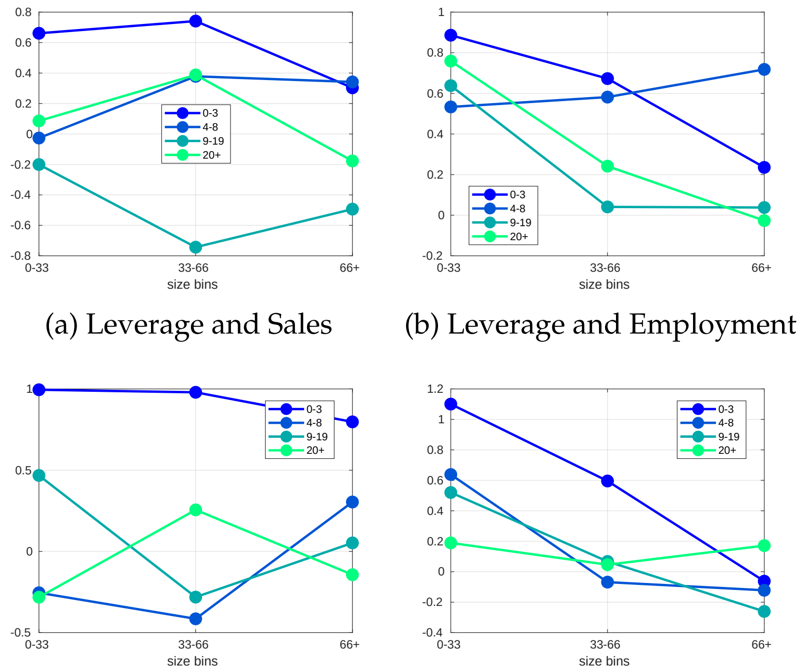

(a) Leverage and Sales (b) Leverage and Employment

(c) G.r. of loans and Sales (d) G.r. of loans and Employment

For any age-size bin, financial cyclicality gradient is the difference $\beta_{j,k,5} - \beta_{j,k,1}$ for size bin (j) and age bin (k) from regression (5). Panels in top row use leverage to measure finance, and panels in bottom row use past bank debt growth to measure finance. Panels in left column measure cyclicality of sales, and panels in right hand column measure cyclicality of employment. Lines correspond to age bins and x-axis to size bins.

firms have the highest leverage growth, while all large firms have declining leverage.

We find that firms with higher leverage tend to be more cyclical, supporting the idea that finance does drive cyclicality at some firms. Interestingly, this effect is stronger for employment than for sales which we interpret as a sign of demand shock smoothing by firms. Additionally controlling for age reduces the marginal effect of finance on cyclicality, while controlling for size does not, suggesting that financial frictions are more important for the age, not size, dimension of cyclicality. Finally, a triple interaction specification finds that higher leverage (or past growth of bank loans) has little effect on cyclicality of large firms, and only increases cyclicality at small firms (of all ages). This leads us to conclude that finance is likely to explain the higher cyclicality of young-small firms, but not the relatively high cyclicality of larger firms.

4 Quantitative model

In this section we build our quantitative model, which builds on classic heterogeneous-firm financial frictions models, such as Khan and Thomas (2013).[23]

4.1 Description of the model

The model features a continuum of heterogeneous firms, with both ex-ante and ex-post heterogeneity. There is a representative consumer, who owns firms and supplies labor. The model also features a final goods aggregating firm. The key extensions that we use to match our new stylized facts are differences in both financial frictions and returns to scale across firms. The model is set in continuous time $t \in [0, \infty)$ with an infinite horizon. We focus on the case without aggregate uncertainty, and conduct business-cycle experiments using unanticipated one-time shocks. The model is presented in steady state, for expositional simplicity, and we therefore drop the time subscript, t, in most of what follows.

4.1.1 Final goods producer

A continuum of heterogeneous firms $i \in [0, G]$ are our firms of interest, and we will refer to them simply as "firms" where it does not cause confusion. These firms produce a firm-specific intermediate good, q_i, using capital and labor. Their goods are all sold to a representative final goods producer who combines them to produce a composite final good, Q, which is used for consumption and investment. The final good is the numeraire, with price normalised to one in all periods. The final-goods production function is $Q = \left(\int_0^G q_i^\theta \, \mathrm{d}i\right)^{\frac{1}{\theta}}$ where Q is units of production of the final good. Define $\epsilon = 1/(1-\theta)$ as the elasticity of substitution between varieties and restrict to gross substitutability ($\epsilon \geq 1 \Leftrightarrow 0 \leq \theta \leq 1$). This ensures that intermediate goods firms have decreasing returns to scale in revenue, even if they have constant returns to scale in production.

The final-goods firm is a price taker in both the final and intermediates markets. Their profit is given by $\pi = \left(\int_0^G q_i^\theta \, \mathrm{d}i\right)^{\frac{1}{\theta}} - \int_0^G p_i q_i \, \mathrm{d}i$, where p_i is the price of each intermediate

[23] For examples of other work building on this framework, see Jo and Senga (2019), Ottonello and Winberry (2020), and the references therein.

firm's good. The final good producer's first order condition for intermediates gives

$$q_i = p_i^{-\epsilon} Q \tag{6}$$

This is the demand curve for the intermediate goods firms.

4.1.2 Intermediate goods firms (a.k.a. "Firms")

There is a mass G of firms which arises via firm entry and exit. Firms have both ex-ante and ex-post heterogeneity, are owned by the representative household, and discount the future at the interest rate, r.

At birth, firms draw a permanent "size type" $s = \{1, 2, ..., S\}$, which determines features which we wish to relate to firm size. Specifically, their returns to scale, η_s, depends on this size type, as well as a permanent component of their physical productivity, which we label z_s^S. To capture features of the firm lifecycle unrelated to financial frictions, we additionally introduce a "lifecycle shock", which we denote $g = \{1, 2\}$. All firms are born "young", with $g = 1$. At an exogenous rate α_G they transition to $g = 2$ and become "old". This shock controls a lifecycle component of their productivity, z_g^G, as well as firm exit rates, ζ_g, which we discuss further below. We normalize $z_2^G = 1$, so that z_1^G gives the productivity disadvantage of young firms. Finally, we allow for idiosyncratic shocks to firm productivity, z_j^J with $j = \{1, ..., J\}$ denoting discrete productivity levels with transition rates from j to j' of $\pi_{j,j'}^J$. All firms share the common production function

$$q = z \min\left\{k, \frac{l}{\alpha}\right\}^{\eta_s} \tag{7}$$

where $z \equiv z_s^S z_g^G z_j^J$ denotes overall productivity, which combines the size, lifecycle, and idiosyncratic components. Firms have Leontief production functions in capital and labor, with labor share determined by α.[24] If all firms had $\eta_s = 1$ then all firms would have constant returns to scale in production, and $\eta_s < 0$ denotes decreasing returns to scale. The demand curve is (6), and a firm's revenue is therefore $pq = z^\theta \min\left\{k, \frac{l}{\alpha}\right\}^{\eta_s \theta} Q^{1-\theta}$. Value added is equal to revenue: $y = pq$.

[24] The use of a Leontief production function is helpful in matching the wide size distribution in the data, when combined with financial frictions which directly affect the purchase of capital only, and not labor. With a Cobb Douglas production function, a financially constrained firm heavily substitutes from capital to labor while young. By ruling this out, the Leontief production function forces firms to maintain a fixed capital-labor ratio, so that financial frictions directly affect both capital and labor equally. This helps keep firms of size type s in the percentile group (0-30% and so on) that they are designed to match.

At the firm level, all factors of production can be adjusted freely without cost. We are in continuous time, and there is no time to build for capital. It is convenient to first optimise labor for a given level of capital. The Leontief production function gives the solution simply as $l(k) = \alpha k$, and $\pi(k, s, g, j) = (z_s^S z_g^G z_j^J)^\theta k^{\eta_s \theta} Q^{1-\theta} - \alpha w k$.

A firm's capital stock evolves through a standard accumulation equation. Given investment i per unit of time and depreciation rate δ we have: $\dot{k} = i - \delta k$. One unit of investment costs p_K units of final good. Old capital and investment are perfect substitutes for firms, so capital also trades at the price p_K.

Firms can borrow using a risk-free short-term bond b with interest rate r. They face a borrowing constraint which limits the amount they can borrow according to the amount they can post as collateral: $b \leq \lambda p_K k$, where recall that k is a firm's physical capital. The parameter λ controls the tightness of the borrowing limit, with smaller λ making the constraint tighter. In the business cycle experiments, we allow λ to evolve as an aggregate financial shock. A firm's net worth, n is defined as its assets less its liabilities: $n = p_K k - b$. Combining this with the borrowing limit gives $k \leq \frac{n}{p_K(1-\lambda)}$. Define a firm's leverage, ϕ, as $\phi \equiv p_K k/n$. Combining this with the borrowing limit gives the constraint as a constraint on leverage instead: $\phi \leq \bar{\phi}$, where $\bar{\phi} \equiv \frac{1}{1-\lambda}$ is the exogenous leverage limit.

A firm's net worth evolves according to

$$\dot{n} = \left(\frac{\pi(k, s, g)}{k} - (\delta + r) p_K \right) k + rn - d \tag{8}$$

where the first term is the net return on leveraged investment, and d denotes the dividend payout flow. We assume that firms cannot raise equity at all after the moment of birth, and so impose $d \geq 0$. We simplify the dividend payout policy, and impose that firms payout dividends only when net worth exceeds an exogenous level $\bar{n}$, and payout such that net worth remains at $\bar{n}$. Firms therefore pay no dividends while they are young, but then start paying out dividends when they are older and have achieved sufficient scale.

Firm exit is exogenous, and occurs at rate ζ_g. This is allowed to depend on the current lifecycle state, g, in order to match the data that young firms exit at a higher rate than old firms. When firms exit, they pay out their remaining net worth, n, as a final dividend.

The firm's problem can be stated recursively using a Hamilton Jacobi Bellman (HJB) equation. Optimized firm value, $v(n,s,g,j)$, can be expressed as

$$\begin{aligned} rv(n,s,g,j) = \max_{0 \le p_K k \le \bar{\phi} n} & d(n) + v_n(n,s,g,j)\left(\left(\frac{\pi(k,s,g,j)}{k} - (\delta + r)p_K\right)k + rn - d(n)\right) \\ & + \zeta_g\left(n - v(n,s,g,j)\right) + \mathbf{1}_{g=1}\alpha_G\left(v(n,s,2,j) - v(n,s,1,j)\right) \\ & + \sum_{j'} \pi^J_{j,j'}\left(v(n,s,g,j') - v(n,s,g,j)\right) + \alpha_\star\left(v^\star + n - v(n,s,g,j)\right) \end{aligned} \tag{9}$$

Here, $d(n)$ is the exogenous dividend payout policy for the current level of net worth. The v_n term is the drift in net worth, which depends on the capital choice and dividend payout. The terms on the second line concern the lifecycle: The ζ_g term captures firm exit, and the final term captures the transition from lifecycle state $g = 1$ to $g = 2$. In the final line, the first term captures transitions across idiosyncratic productivity states. The $\alpha_\star$ term captures the transition to an additional "superstar status", which is added to the model for calibration purposes. This shock is very rare and occurs on average later in life, selecting a few firms to have very high productivity, $z_\star$. See Appendix C for more details.

The firm investment policy in this setting can be expressed as an unconstrained optimal capital stock, which firms will achieve only if they are financially unconstrained. The first order condition with respect to capital is $v_n(n,s,g,j)\left(\pi_k(k,s,g,j) - (\delta + r)p_K\right) = \mu_k$, where $\mu_k \ge 0$ is the multiplier on the borrowing constraint. If a firm hits its borrowing constraint then we know that $k = \bar{\phi} n / p_K$. If a firm is rich enough to be unconstrained, then $\mu_k = 0$ and the capital FOC gives us $\frac{\pi_k(k,s,g,j)}{p_K} = \delta + r$. This gives the unconstrained investment policy if unconstrained, $k^{unc}(s,g)$, which has an analytic solution. The overall investment policy can then be simply expressed as $k(n,s,g) = \min\{\bar{\phi} n / p_K, k^{unc}(s,g)\}$.

Finally, denote by μ_0 the flow rate at which new firms enter, which is assumed to be constant. After entry, new firms draw their permanent size type, with γ^S_s denoting the probability of drawing type s. New entrants are endowed with some initial amount of net worth, n, from an initial equity injection by their owners. We suppose that firms start life with net worth equal to the fraction n^e_s of the net worth required to become financially unconstrained.[25] This is allowed to differ by size type, which will be important

[25] Let $n^{unc}(s)$ denote the amount of net worth required to become financially unconstrained for a firm of size type s. This is easy to calculate using the leverage constraint, as $k^{unc}(s,2,j) = \bar{\phi} n^{unc}(s,j)/p_K \rightarrow n^{unc}(s) = E_j k^{unc}(s,2,j) p_K / \bar{\phi}$, where we define $n^{unc}(s)$ as the level of net worth to be able to afford the unconstrained level of capital, given the borrowing constraint, when they reach lifecycle maturity ($g = 2$) and averaged across idiosyncratic shock levels. Entrants therefore start life with net worth equal to $n^e_s \times n^{unc}(s)$.

for matching the data on cyclicality by joint age-size bin.

4.1.3 Closing the model

Given the solution to the firm problem, we can simulate the endogenous firm distribution in steady state or in transitions. We can then calculate aggregates such as output and employment, and moments of the firm size and age distribution. We close the model by specifying how the prices that firms face (real wage, interest rate, and capital prices) are determined in a simple general equilibrium setting. Aggregate GDP is the sum over firms $Y \equiv \int_0^G y_i di$, and goods market clearing gives $Y = C + I$.

We assume that the representative household has instantaneous utility function over consumption, c, and labor supply, L^s, of $U(c, L^s) = c - (L^s/\chi)^{1+1/\sigma}/(1+1/\sigma)$ and discount rate ρ. This gives the equilibrium interest rate as a fixed constant $r = \rho$. The household's labor supply condition gives labor supply as a simple function of the wage: $L^s = \chi w^\sigma$. Finally, we suppose that investment goods can be produced one-for-one from the final good, giving a fixed equilibrium capital price of $p_K = 1$.

4.2 Result 1: Performance of a "steady state" calibration

In this section we describe what we call the "steady state" calibration, and show that it cannot match our new facts on cyclicality by joint firm age-size bin. This calibration generates a simple heterogeneous-firm model with financial frictions. The key idea is that this calibration targets only "steady state" moments of the firm distribution, and we will later contrast it with a "cyclical" calibration which additionally targets our new facts on cyclicality by age and size.

We loosely follow the calibration strategy of Khan and Thomas (2013), and so turn off three novel features of our model, which we will use later in our "cyclical" calibration. Firstly, we suppose that all firms have the same (constant) returns to scale, and so set $\eta_s = 1$ for all size types. Secondly, we attribute all employment growth of young firms to financial frictions, setting $z_1^G = z_2^G = 1$ so that the lifecycle component of productivity is constant. Finally, we do not explore how different firms may enter with different degrees of financial frictions, and suppose that $n_s^e = n^e$, so that all firms enter with the same fraction of unconstrained net worth.

4.2.1 "Steady state" calibration details

We start by describing our relatively standard parameters. We take one unit of time to be one year. We set the interest rate r to a 2% annual real interest rate, in line with the lower real interest rates seen in recent years. The capital depreciation rate δ is set to a 10% annual rate. We set θ to 0.9 to give a 10% markup in a frictionless model, as is standard in the New Keynesian literature. Firms have decreasing returns to scale in revenue, and so have well defined optimal sizes despite having $\eta_s = 1$. We choose the labor to capital ratio α to control the equilibrium quantity of employment, which is set to match the average firm size (total employment over total number of firms) in Denmark. The labor supply disutility χ is chosen to match a labor share of income of 60%. The labor supply elasticity η is set to 0.3, which implies that wages fall by 30% for a given change in employment. The entry rate μ_0 is chosen to normalize the mass of firms in steady state to one.

We set the leverage constraint to $\bar{\phi} = 3$. This implies that firms remain financially constrained only until around age 3 on average and therefore represents a relatively loose borrowing limit.[26] We set the level of net worth at which firms start paying out dividends to a large number.[27] We specify the idiosyncratic shock process as an AR(1) process discretised with $J = 2$ nodes using the Rowenhurst method. To remain close to Khan and Thomas (2013), we fix the annual autocorrelation at their value of 0.659, normalise the mean to one, and choose the standard deviation to match the standard deviation of idiosyncratic investment rates of 33.7% from Cooper and Haltiwanger (2006).[28]

Since we are interested in investigating cyclicality across the age and size distributions, a major goal of our calibration is to match these distributions well in steady state.[29] For the size distribution, we use our size types, s, to flexibly match the data. Specifically, we

[26] In our data, the highest average Debt/Asset ratios for any firm-age groups are around 0.8, for firms aged 0-3. This correspond to a leverage ratio of $1/(1-0.8) = 5$, so our choice of a maximum leverage of 3 is conservative in that we allow firms to take on slightly less data than in the data.

[27] Above the "minimum saving policy" (see Khan and Thomas (2013)), exactly whether or not firms pay out dividends has no effect on firms choices of employment and so on in steady state. We therefore choose that firms pay out for some $\bar{n}$ such that even the most productive firms can fund their unconstrained optimal capital with no debt ($\phi = 1$).

[28] The dataset of Cooper and Haltiwanger (2006) is a balanced panel of large manufacturing plants. To remain comparable with this data, we calibrate the standard deviation of investment rates only for a similar subset of firms in our model. See Appendix C for details.

[29] For the data used to calibrate distribution of number of firms and employment by firm age-size the model, we do not drop firms for whom we are missing data on debt. This ensures that we capture the number and size of firms in each age-size bin correctly, regardless of whether they have missing data on debt.

use $S = 4$ size bins to target the 0-30%, 30-60% 60-90%, and 90%+ percentile size bins in the data. We suppose the probability of being born in group $s = 1, 2, 3$ is 30% each, and $s = 4$ is 10%. By choosing the productivity levels appropriately each size type $s = 1, 2, 3, 4$ is therefore assigned to form the predominant mass of firms in each of the 0-30%, 30-60% 60-90%, and 90%+ size bins respectively.[30] We calculate these percentile-based size bins in our model exactly as in the data. We choose z_4^S to normalize aggregate GDP to 1, and choose the relative values of z_1^S, z_2^S, and z_3^S to match the average employment inside the 0-30%, 30-60%, and 60-90% size bins respectively. The fraction of firms in each percentile-based size bin is simply given by their definition.[31]. Permanent productivity heterogeneity across size types s allows us to match the very wide firm size distribution in the data, where 30% of firms have on average 1.95 employees, while the largest 10% have on average 146.76 employees.

Moving on to the firm age distribution, we target both the distribution of the number of firms by age (i.e. the exit rates) and the distribution of total employment by firm age. To target the exit rate we use data from Andersen and Rozsypal (2021), who calculate exit rates by firm age for the Danish economy. Using their data, we calculate that firms aged 0 have an exit rate 2.16 times higher than firms aged 16+, and firms aged 6 have an exit rate 1.33 times higher. We target these ratios, as well as an overall average exit rate of 8% per year, using the exit rates ζ_y and ζ_o and the speed at which firms transition from young to old, α_G. We target the distribution of employment by firm age in two ways: we match the average size of firms aged 0, and aged 20+ years old. We follow Khan and Thomas (2013) and use the initial net worth of entrants, here our parameter n_0, to target the average size of firms at age 0. Financial frictions are therefore used to explain firm growth in the early years of the lifecycle: age 0 firms have 9.35 employees on average in the data, while firms aged 4-8 have 16.44, because entrants start with only 35% of the net worth needed to reach their optimal size and grow as they overcome this friction. We use the rare "superstar firm" shock to target the high employment share of very old firms.

[30] Note that since firm grow over their lifetime due to financial frictions, not all firms in, for example, the 30-60% percentile bin in the model will be from the $s = 2$ type. However, type $s = 2$ firms form the vast majority of firms in that bin, which allows us to choose z_2^S to target average features of firms in that percentile bin. In the calibration, each size bin is composed between 80% and 100% of firms from the assigned type.

[31] In the data, the fraction of firms in, for example, the 0-30% bin is not exactly 30%, but is instead 37%. This is due to rounding, as we define our size bins based on the number of employees at the firm, which is an integer number, meaning that a discrete mass of firms may sit at the boundary.

Table 2: Firm distributions in the model and data

	Fraction of firms				Average employment			
Size	0-30	30-60	60-90	90+	0-30	30-60	60-90	90+
Model (s.s. cali)	0.30	0.30	0.30	0.10	2.01	5.85	16.61	138.02
Model (b.c. cali)	0.30	0.30	0.30	0.10	1.98	5.93	16.58	137.82
Data	0.36	0.26	0.28	0.10	1.95	5.65	15.90	146.15

(a) Size distribution

	Fraction of firms					Average employment				
Age	0	1-3	4-8	9-19	20+	0	1-3	4-8	9-19	20+
Model (s.s. cali)	0.08	0.19	0.21	0.27	0.25	9.41	13.63	18.54	20.25	33.41
Model (b.c. cali)	0.08	0.19	0.21	0.27	0.25	9.40	12.57	17.82	21.91	32.97
Data	0.05	0.18	0.25	0.25	0.27	9.35	11.90	16.44	21.82	32.95

(b) Age distribution

Firm age and size distributions in the model and data. “Model (s.s. cali)” refers to the model calibrated using the “steady state” calibration, and “Model (b.c. cali)” refers to the “cyclical” calibration. Size bins refer to percentile groups and age bins to age in years since birth. Average employment refers to total employment in the bin divided by the number of firms in the bin. In the data, the number of firms in, e.g., the 0-30% percentile bin is not exactly 30% of firms due to the fact that many small firms have exactly the same number of employees in the data, and hence lie on the boundaries of the sets.

Their productivity level, $z_\star$, is used to target the average employment of firms aged 20+ years old. Intuitively, the superstar shock is very rare, and therefore only occurs for firms on average when they are very old, selecting a few firms whose higher productivity leads them to have a size of nearly 600 employees later in life.

A complete list of parameter values is given in Table 6 in the Appendix C, as well as further details of firm policies in steady state.[32] Each parameter is adjusted to hit one moment, and we are able to hit all moments exactly, stopping when the error between all model and data moments falls below 5%. The model fits the firm age and size distributions extremely well, including in age bins which we did not target, as shown in Table 2.

[32]Specifically, in Figure 19 we plot the fraction of firms financially constrained by age and size bin.

4.2.2 Cyclical performance of the "steady state" calibration

We now show that the "steady state" calibration cannot match our new facts on cyclicality by joint age-size bin. As the cyclicality of different firm groups may depend on the shock hitting the economy, we begin by exploring the response of the economy to three different shocks. Given that we are primarily studying the Financial Crisis, we choose two financial shocks, and one real shock. In all cases, the size of the shock is chosen to generate a 1% fall in GDP which mostly dies out within three years, with the focus instead being on how the response to the shock differs across different firm groups.[33]

The results of this exercise are plotted in Figure 8, where we display the cyclicality of our joint firm age-size groups using the exactly the same regression approach that we previously applied to the data. We use employment as our firm-level outcome measure, which is regressed on real GDP growth using the specification (2). In Figure 8(d) we plot the data,[34] and in Figures 8(a) to (c) we plot results from the model. We plot the *relative* regression coefficients, defined as the regression coefficient for each age-size bin divided by the absolute value of the regression coefficient for the oldest-largest (age 20+ size 90%+) bin. Thus each value gives the bin's cyclicality relative to the oldest-largest age-size bin.[35]

We consider three aggregate shocks in the model. Panel (a) gives the results of an exogenous tightening of the borrowing constraint, represented as a reduction in the parameter $\bar{\phi}$. This shock was previously used in Khan and Thomas (2013), for example, to represent a financial crisis. Panel (b) gives the results an increase in the interest rate, r, charged to all firms. This represents either an increase in discount rates (Hall (2017)), or an increase in spreads charged to firms due to (for example) problems in the banking sector, modelled as an (unchanged) risk-free rate plus an increased spread charged to firms

[33] Specifically, suppose we shock a parameter x, by allowing it to vary with time, t. Then at time 0 the parameter unanticipatedly jumps to its new value, x_0, and then recovers back to its original steady-state value, x_{ss}, according to the deterministic process $\dot{x}_t = -\rho_x(x_t - x_{ss})$. We set $\rho_x = 0.9$, and compute the perfect foresight transitions of the economy to this shock using the so-called "MIT shock" approach. The shock has mostly died out within three years, and we simulate the transition for 20 years, confirming that raising this number, or choosing a finer time grid, has no effect on the results. Sample paths for the shocks and aggregates are plotted in Figure 20 in Appendix C.

[34] This simply repeats the previously-shown plot Figure 3(b).

[35] We use relative coefficients because there is a difference in the average level of firm cyclicality across model and data. For the non-relative values of the regression coefficients see Figure 3 for the data, and Appendix C.2 for the model experiments.

Figure 8: Cyclical response of age-size groups to shocks in the "steady state" calibration

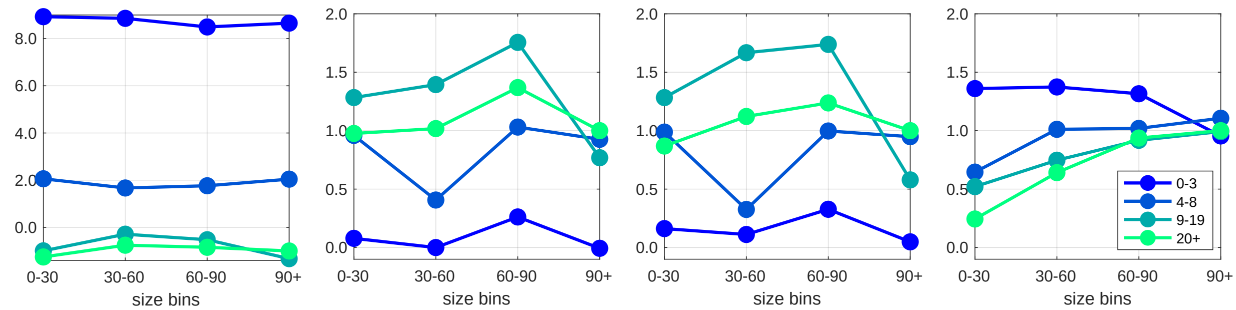

(a) Collateral ($\bar{\phi}$) (b) Discounts (r) (c) TFP (z) (d) Data

Panels give relative regression coefficients from regressions of firm-level growth rates of employment on aggregate GDP growth, computed from model simulated data for our recession experiments. Values are the regression coefficient for that age-size bin divided by the (absolute value of) the regression coefficient for the oldest-largest (age 20+ size 90%+) bin.The regressions are on firm-level data aggregated to the yearly level and treated in the same way as the data. Size bins are percentiles, and each line refers to a different firm age group. The final panel gives the results from real-world data, and the remaining panels from model data.

for borrowing.[36] Finally, Panel (c) gives the results of an aggregate TFP shock, which we model as a proportionally equal reduction in TFP at all firms. We include this shock in order to consider non-financial shocks, and how they compare to the two financial shocks.

Inspecting Figure 8 we see that this calibration of the model does not replicate the data in response to any of the three shocks considered. The collateral constraint shock (Panel (a)) comes closest, since in response to this shock young firms are more responsive than old firms, for a given size group. This is also true in the data: For example, in panel (d) we see that age 0-3 firms (blue line) are more cyclical than age 20+ firms (green line) for all size groups. In the model, young firms are more responsive to the collateral constraint shock than old firms because young firms are more likely to be financially constrained than old firms. For young firms, who have limited net worth, a reduction in borrowing forces them to reduce their investment in capital, and hence also their employment and output. For older firms, who have accumulated sufficient net worth to become financially unconstrained, a tightening of the borrowing constraint has no effect on their real variables, as they are already away from their borrowing constraints. However, the collateral constraint shock has an important failing relative to the data, which is that older firms are not responsive to the shock at all, whereas in panel (d) old and large firms (e.g. age 20+

[36]See Del Negro et al. (2017) and Gertler and Karadi (2011) for examples of representative agent models where a financial recession is modelled as an increase in financial spreads charged to all firms, either exogenously or endogenously.

and in the 90%+ size bin) are very cyclical in the data.[37]

This gradient within older firms—that larger older firms are very cyclical, while smaller older firms are not—is something that none of the three shocks can match. The discount rate (Panel (b)) and TFP (Panel (c)) shocks both generate that older firms are more cyclical than young. This is because these shocks affect the marginal incentives of firms, encouraging unconstrained firms to shrink in response to higher costs or lower productivity. This affects old firms more, since young firms are up against their borrowing constraints and hence unresponsive to such marginal incentives. But in response to these shocks all old firms (large or small) are equally cyclical, in contrast to the data. In fact, there is no monotonic size gradient within any age bin in all three of the model panels, in contrast to the data in Panel (d) where large firms tend to be more cyclical for the three oldest age groups, but less cyclical for the youngest age group. This implies that no combination of the three shocks can fully replicate the data in Panel (d), which represents a failing of the basic calibration according to our new data.[38] It is for this reason that we turn to our extended "cyclical" calibration, in order to understand the features needed to match the data.

4.3 Result 2: A "cyclical" calibration

In this section we describe what we call the "cyclical" calibration. This extends the model to include heterogeneous returns to scale and entrant net worth by firm size, to allow it to match our new facts on the cyclicality of firms by joint age-size bin. Most calibration targets remain the same as in the "steady state" calibration, and we describe the new features in the section below.

[37] Additionally, young firms are much too cyclical in response to the collateral shock: note that the coefficients on the youngest firms in Panel (a) are around 6, while in the data they are around 1.2. This is because a pure collateral constraint shock only affects young firms, while leaving older (and hence larger) firms unaffected, making the change in employment at young firms much larger than the change in aggregate GDP.

[38] While other types of shocks are also possible, we believe that these three shocks span quite well the types effects that shocks have on the firm age-size distribution in this model. In particular, we considered both shocks to the quantity (collateral constraint) and price (spread) of borrowing. We then considered a generic TFP shock, whose effects are similar (in terms of the firm age-size distribution) as any shocks to demand, productivity, or factor prices which affect firms in the same proportional manner. Since all firms in the basic model face the same factor prices, any shocks which transmit via TFP, demand, or factor prices will therefore have the same basic effects on the firm age-size distribution as the TFP shock.

4.3.1 "Cyclical" calibration details

Our "cyclical" calibration approach consists of choosing parameters of the model to match *both* steady state moments, *and* the cyclicality of firm age-size groups in the data from our regression results. Since the cyclicality of age-size groups depends on which shocks hit the economy (as we showed in the last section) we therefore jointly estimate the parameters of the model and the shocks which drive the recession episode.

Along with our previous targets, our calibration exercise targets 1) the relative cyclicality of firms by age and size (Section 3.2), 2) the size of the recession, and 3) the average growth rate of young-small firms. We choose two shock processes and two novel features of the model to match these firm level outcomes in this exercise. We jointly choose these nine new parameters to minimise the distance to all of the nine new moments discussed in this section and discuss the moments most closely related to each parameter in the text below.[39] When targeting the cyclicality of any firm age-size bin, we target the cyclicality *relative* to the largest, oldest firm group (90%+ size, age 20+) by targeting the ratio of their regression coefficients.

Firstly, we modify the calibration of the steady state of the model. In our "steady state" calibration, we used a common parameter n^e to adjust entry net worth to match the average size of aged 0 firms. In the "cyclical" calibration, we will use entry net worth to target cyclical moments, and so take a different approach. We now incorporate within-firm productivity growth with age, and instead calibrate $z_1^G < 1$ to match the average size of aged 0 firms.

We next describe the estimation of the shocks. In this section, we choose to focus only on financial shocks, and assume that the recession is driven jointly by the collateral and discount rate shock. We estimate the initial values of these shocks — denoted $\bar{\phi}_0$ and r_0 respectively — and allow them to fade back to steady state at the same rate of 0.9, as described in Footnote 33.[40] As small-young firms will be the most financially constrained

[39]Since calculating these moments requires simulating the recession experiment, we split the estimation into a two-layer procedure which exactly hits the steady state moments in an inner loop, and then uses a simulated minimum distance routine in the outer-loop, to minimize the sum of squared errors of the cyclical moments. More details are in Appendix C.1.

[40]As can be seen from Figure 8 panels (b) and (c), the discount rate shock and TFP shock generate very similar cyclicality across the age-size distribution. In this sense, the cyclicality data does not precisely identify the discount rate shock from the TFP shock, and we focus on the discount shock to provide a simple interpretation of the recession as a purely financial shock. However, other data do suggest that there was no great cyclicality of TFP during this period. In particular, we investigate the cyclicality of labor

and hence affected by the financial shock, we choose the size of the collateral constraint shock in order to match the relative cyclicality of small-young (0-30%, 0-3) firms, which is $2.24/1.64 = 1.36$ in the data in Figure 8(d). We also aim for a 5% total GDP fall an impact, comparable to the fall in GDP in Denmark in the first year of the Great Recession, and choose the size of the discount rate shock to match this fall, for a given size of collateral constraint shock.

We now turn to estimating cyclical dynamics by joint firm age-size bin. We target two key features of the data. Firstly, *among older firms* in the data, large firms are more cyclical than small firms. We argued that this is unlikely to be driven by financial frictions, and so we need another model feature to match this fact, which we match instead by calibrating the differing degrees of returns to scale across size types s. Crucially, differing degrees of returns to scale also imply different responsiveness to shocks, in a way that very natural meshes with our empirical findings. Small firms are likely to be small not because they are unproductive in a TFP sense, but because they are fundamentally very different businesses to larger firms (think a local shop versus Carlsberg Group), with a smaller business scope which could imply lower returns to scale. In terms of economic theory, firms with more decreasing returns to scale are also endogenously less responsive to shocks, which then gives a natural explanation for why larger firm are more cyclical in our data.[41] To reduce the number of free parameters, η_4 is chosen to normalize the average returns to scale to one.[42] We choose η_1, η_2, and η_3 to match the relative cyclicality of size bins 0-30%, 30-60%, and 60-90% among the oldest firm group (aged 20+).

Secondly, *among younger firms* in the data, large firms are less cyclical than small firms. We use this data to discipline the degree of financial frictions faced at birth by firm size, by allowing for differing net worth at birth for firms of different size types s. Specifically, we set n_2^e, n_3^e, and n_4^e to match the relative cyclicality of size bins 30-60%, 60-90%, and 90%+

productivity by firm, and find that the labor productivity of large firms is essentially acyclical, despite the high cyclicality of their value added and employment. We take this as evidence that a TFP shock is unlikely to be driving the employment of this firm group.

[41] To see this, consider a simple static model where a competitive firm produces output using capital only, with returns to scale α: $y = zk^{\alpha}$. They rent capital at price r. Their profit maximization problem $\max_k zk^{\alpha} - rk$ implies optimal capital $k = (\alpha z/r)^{\frac{1}{1-\alpha}}$. The elasticity of their capital choice to a change in productivity is $\frac{\partial \log k}{\partial \log z} = \frac{1}{1-\alpha}$. Thus the more decreasing returns to scale (lower α), the less responsive is the firm to changes in TFP (lower $\partial \log k / \partial \log z$). The same is true for a change in the factor price, r.

[42] By this we mean that, at the ergodic distribution, exogenously increasing all inputs proportionally at all firms by a factor λ raises aggregate output by a factor λ. Intuitively, this means that an appropriately weighted average of η_s across firms is equal to one.

among the youngest firm group (aged 0-3). By making n_2^e greater than n_1^e, for example, this makes type 2 entrants less financially constrained than type 1, and hence reduces cyclicality for larger firms within the youngest age group. Recall that the cyclicality of the aged 0-3, size 0-30% group was already "used" to calibrated the size of the financial shock. We therefore introduce one final moment to calibrate n_1^e. We choose this parameter to match the average employment growth rate of young (0-3) firms in the 0-30% size bin, since lower initial net worth implies that they start further below their optimal size and therefore will grow faster in their youth.[43]

A complete list of parameters is given in Table 6 in Appendix C. The model continues to give a very close match to the firm age and size distributions, as shown in Table 2. The values of the outer-loop moments in the data and the estimated model are given in Table 5, and the average error (computed as the square root of the mean squared error) is equal to 4.5%.

Of particular note is that we provide new independent empirical evidence in support of our finding that returns to scale are heterogeneous across firm size groups. In a separate empirical exercise, we split firms in the data by size and estimate revenue production functions using the Olley-Pakes (OP) and Levinsohn-Petrin (LP) methods (with Ackerberg-Caves-Frazer (ACF) correction). The results of this exercise support our assumption in the model: larger firms in the data do indeed tend to have on average larger returns to scale compared to smaller firms. The differences are quantitatively meaningful, with firms with 0-7 employees having returns to scale in revenue of around 0.85 while larger firms have returns to scale around 1.05 in the data. Coincidentally, this difference is very similar to the gap in revenue returns to scale that we independently reverse engineer in our model to match the cyclicality of firm size groups: For small firms we estimate revenue returns to scale of $\theta\eta_1 = 0.9 \times 0.7952 = 0.7157$ and for large firms $\theta\eta_4 = 0.9 \times 1.0407 = 0.9366$, giving a difference of around 0.25 in both data and model. We take this as supportive evidence that returns to scale differences could be important for understanding differences in cyclicality across firms size groups. The details are pro-

[43] This data is plotted in Figure 2(e). However, since that data excludes firms for whom we are missing data on debt, we target a slightly different version of this figure which includes all firms regardless of whether they have available debt data. This alternative sample is given, along with the model results, in Figure 10. Note that we can alternatively think of n_1^e being used to target the relative cyclicality of aged 0-3, size 0-30% firms, symmetrically with how n_2^e to n_4^e are chosen, since all parameters are jointly used to minimize the distance to all moments.

vided in Appendix B.6, along with discussion of other empirical work in support of this finding.

4.3.2 Cyclical performance of the "cyclical" calibration

We plot the results of our "cyclical" calibration in Figure 9, again using relative regression coefficients. Panel (d) gives the data, and panel (c) plots the regressions from the cyclical calibration. The model is able to match the key features of the data, which we targeted in the calibration. In particular, 1) young firms are more cyclical than old firms, and 2) among older firms, large firms are more cyclical than small firms, while the opposite is true among the youngest firms. Sample paths for the shocks and aggregates are plotted in Figure 21 in Appendix C.

Figure 9: Cyclical response of age-size groups in the "cyclical" calibration

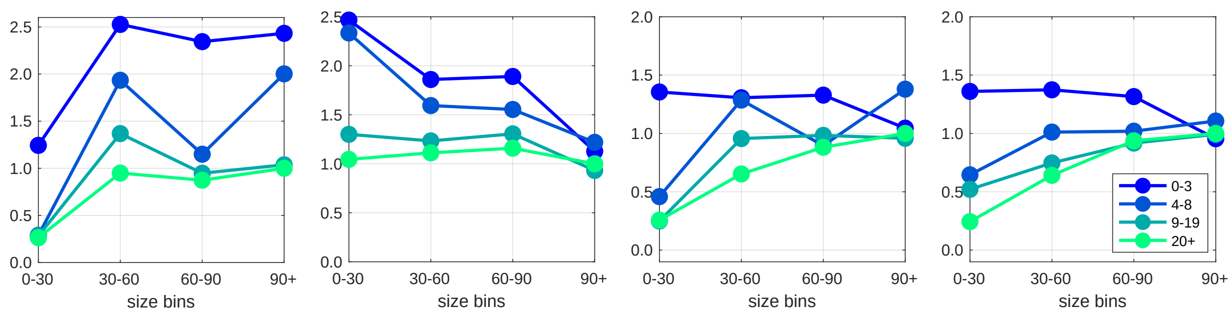

(a) Het. η_s (b) Het. n_s^e (c) Full model (d) Data

Panels give relative regression coefficients from regressions of firm-level growth rates of employment on aggregate GDP growth, computed from model simulated data for our recession experiments. Values are the regression coefficient for that age-size bin divided by the (absolute value of) the regression coefficient for the oldest-largest (age 20+ size 90%+) bin. The regressions are on firm-level data aggregated to the yearly level and treated in the same way as the data. Size bins are percentiles, and each line refers to a different firm age group. The final panel gives the results from real-world data, and the remaining panels from model data.

The "cyclical" calibration is able to match this data well due to the novel features we added to the model, and the mechanisms by which they work were described in the previous section. To show that these features contribute to cyclicality as discussed, in panels (a) and (b) we plot the cyclicality results for two recalibrated models which incorporate only one feature each. Panel (a) shows that adding heterogeneity in returns to scale makes large firms more cyclical. However, alone it also means that large-old firms are much too cyclical, as the excess cyclicality of the young does not decrease with size, as it does in the data. Panel (b) shows that adding heterogeneity in initial net worth makes the excess cyclicality of the young decrease with size. However, among old firms it does not make large firms more cyclical than small, as they are in the data. Panel (c) shows that putting these together in the full calibration yields the required match to the data.

One important feature of the data was not targeted, and instead serves as an untargeted test of the model. This is the fact that, among young firms, average growth rates are higher for smaller than for larger firms. In Figure 10(c) we plot the data, and in panel (b) we plot the results from our model, where only the value for age 0-3 size 0-30% firms was targeted. We see that the "cyclical" calibration captures the main pattern in the data, in particular the average growth rates of the smallest and largest firms in the age 0-3 group. In contrast, panel (a) shows that the "steady state" calibration fails to match this data because it generates equal growth rates when young for all firm size groups. In our new calibration, smaller firms are born with relatively less net worth, and will be more financially constrained early in life than larger firms.

Figure 10: Average growth rate of employment by age-size bin

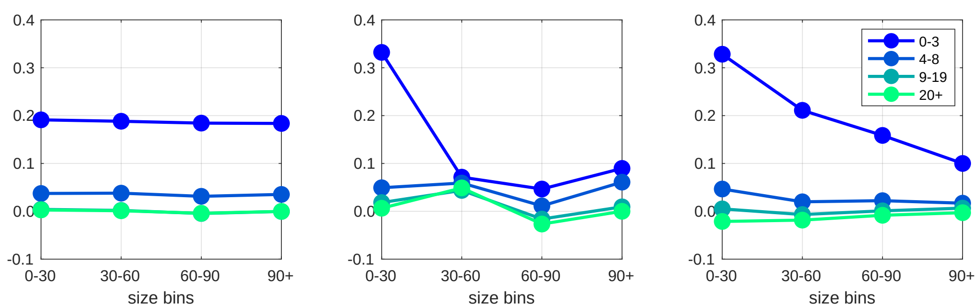

(a) "Steady state" cali (b) "Cyclical" cali (c) Data

Panels give regression coefficients from regressions of firm-level growth rates of employment on firm age-size dummies, computed from model simulated data. The regressions are on firm-level data aggregated to the yearly level and treated in the same way as the data. Size bins are percentiles, and each line refers to a different firm age group. The first panel gives results from the "steady state" calibration of the model, the second from the "cyclical" calibration, and the final from real-world data.

4.4 Result 3: Policy implications

Our final exercise is to investigate the implications of our results and model for business cycle policy. These are informed both by our empirical findings on the cyclicality of different firm groups, and the twists they implied to our calibrated model. In particular, we compare policy exercises in both the original "steady state" calibration, and in our new "cyclical" calibration model, which introduced the two new features of i) heterogeneous entrant net worth by size, and ii) heterogeneous returns to scale.

We consider two simple policies, meant to capture two different styles of possible policy intervention during a recession. The first policy we call an "incentive" type policy, which consists of a temporary subsidy to the firm's wage bill. In particular, we introduce a subsidy so that the government pays a fraction τ of all firms' wage bills, with $\tau = 0$ in steady state. We consider a temporary increase of the subsidy to 1% of the wage bill, which fades at rate 0.9 as did our business cycle shocks. We call this policy an incentive type policy because it changes the effective marginal cost of production for firms, and hence incentivises them to expand production. The second policy we call a "balance sheet" policy, which consists of giving debt relief to firms. In this policy, at time 0 the policymaker pays off a fraction x of all firms' debts, reducing their debt from b to $(1-x)b$, and hence increasing their net worth. This policy does not affect the marginal cost of finance for firms, but it does increase the net worth of firms, and thus the access to debt for financially constrained firms. We consider a one-off 20% debt forgiveness at time 0.

We start with the labor subsidy policy, shown in Figure 11(a). The left panel gives the response of aggregate output to the policy, showing that the policy is more effective in our new model than in the original calibration (3.8% output rise on impact vs. 2.9%). The centre and right panels give the regression coefficients measuring responsiveness by age-size groups. These reveal why the policy is more powerful in the new calibration, as the age-size responses are markedly different between the two models. In particular, in the new calibration, the policy now has a clear firm size dimension, with large firms being more responsive to the policy even conditioning on age. This creates a composition effect which boosts the aggregate response, as the firms who respond the most also happen to be large and hence more important for aggregate output. As with earlier results, this follows directly from the fact that large firms having less decreasing returns to scale makes them more sensitive to changes in their marginal costs.

We now turn to the debt forgiveness policy, shown in Figure 11(b). The response of aggregate output in the left panel now shows that the policy is instead less effective in our new model (0.2% output rise on impact vs. 3%). In both calibrations, the policy has a persistent effect on output which lasts many years, despite the policy being enacted only at time 0. Firm's financial positions are slow moving, and so by helping firms at time

Figure 11: Effect of two policies in standard vs calibrated model

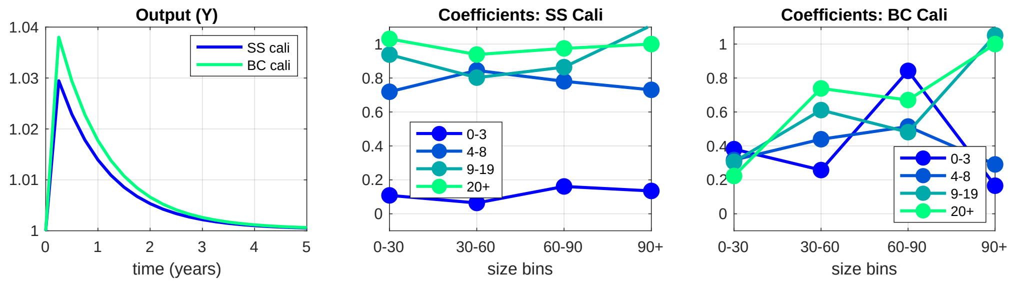

(a) Labor subsidy

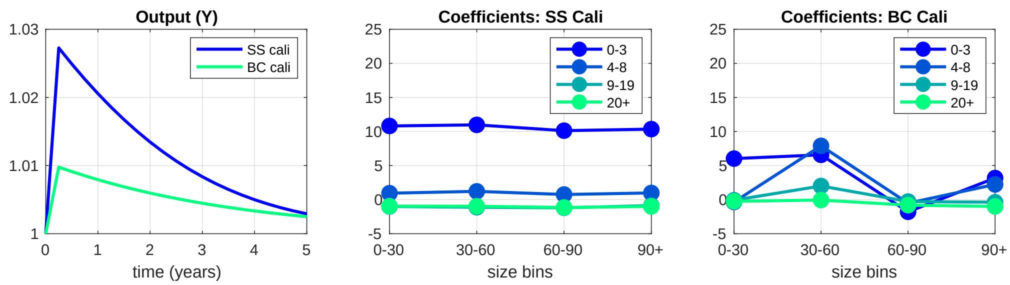

(b) Debt relief

The left panel gives the response of aggregate GDP to the policy experiment, in both the "steady state" and "cyclical" calibrations. The center and right panel give relative regression coefficients (relative to the oldest largest age-size bin) from regressions of firm-level growth rates of employment on aggregate GDP growth, computed from model simulated data for each calibration respectively.

0 they remain less financially constrained for the rest of their lifecycle. The regression results in the centre and right panels reveal why the policy is weaker in the new model. In both models, the policy has the largest effect on young firms, as these firms are more likely to be financially constrained and hence benefit from debt relief. However, in our new calibration, the data suggested that *large* young firms are less financially constrained. Hence the right panel finds that responsiveness to the policy is declining with size on average. This creates a composition effect which dampens the aggregate response, because the firms responding to the policy are now smaller on average than in the steady state calibration.

In conclusion, our empirical results inform changes to a standard heterogeneous firm model which have important policy implications. Some policies (incentive based) are more effective, and others are less (debt forgiveness), with the changes due to the fact that firms of different ages and sizes now respond differently to these policies.

5 Conclusion

In this paper we documented novel facts about the cyclicality firms by age and size, with particular attention paid to the interaction between the two, and the role of finance. Using high quality registry data from the universe of Danish firms, we first document that employment and turnover are more sensitive to the business cycle at younger firms than older firms, but that the relationship between size and cyclicality is more complicated. Among older firms, large firms are more cyclical than small, while among young firms, small firms are more cyclical than large.

These results are possible because our dataset contains explicit information about when firms are formed, allowing us to construct a high quality measure of firms actual age from legal inception. This distinguishes our dataset from other sources where it is either not possible to measure age, or only to do so from the age that firms go public. This allows us to look at the cyclicality of very young firms, which is where we find the strongest excess cyclicality. We additionally have data for firms of all sizes, allowing us to investigate cyclicality for even the smallest of firms. We use this data to additionally provide a detailed investigation of firm outcomes and growth rates across different size and age groups.

Given that our dataset contains detailed financial variables, we then investigate the role of finance in driving the excess cyclicality of different firm groups. We find that young firms have higher leverage than old firms, and hence are more likely to be financially constrained. They additionally are typically trying to expand their leverage, while leverage is typically shrinking at older firms. On the other hand, leverage ratios are remarkably similar across firm size groups, after controlling for firm age. Studying cyclicality by leverage, high leverage firms are more cyclical than low leverage firms, a finding which is dampened by controlling for age but not size. Taking these results together, we argue that the excess cyclicality of young firms is plausibly linked to financial frictions, while the same is less likely to be true for larger old firms.

We then use these insights to build a quantitative heterogeneous firm model, and investigate the extensions to a standard calibration needed to replicate our new facts. We find that standard calibrations struggle to match cyclicalities across age, size, and joint age-size bins at the same time. Part of the problem is that in standard models age and size are too closely linked, as young firms tend to be financially constrained and, hence, smaller. Two extensions bring the model closer to the data. Firstly, we introduce heterogeneous returns to scale, so that large firms have less decreasing returns to scale. This can parsimoniously explain why they are larger and more cyclical, and is consistent with our findings from a production function estimation exercise. Secondly, we allow larger firms to be born richer, and hence less financially constrained. This explains why, among large firms, cyclicality does not depend on firm age, as in the data. Together, these extensions bring the model's implications for cyclicality by joint age-size bin in line with the data. We finally use our model to investigate the effect of recession-fighting policies, and how they transmit through the firm age and size distributions. A key implication of these exercises is that properly matching the responsiveness of firms by age and size can have large effects on the policy implications of our models.

References

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2021). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies*, 89(1):45–86.

Andersen, S. G. and Rozsypal, F. (2021). What can old firms tell us about the effect of age on firm size. working paper.

Begenau, J. and Salomao, J. (2018). Firm Financing over the Business Cycle. *The Review of Financial Studies*, 32(4):1235–1274.

Bernanke, B. and Gertler, M. (1989). Agency Costs, Net Worth, and Business Fluctuations. *American Economic Review*, 79(1):14–31.

Bernanke, B., Gertler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, Part C, chapter 21, pages 1341–1393. Elsevier, 1 edition.

Casiraghi, M., McGregor, T., and Palazzo, B. (2021). Young Firms and Monetary Policy Transmission. IMF Working Papers 2021/063, International Monetary Fund.

Chodorow-Reich, G. (2014). The employment effects of credit market disruptions: Firm-level evidence from the 2008-09 financial crisis. *Quarterly Journal of Economics*, 129(1):1–59. Lead article.

Cloyne, J., Ferreira, C., Froemel, M., and Surico, P. (2019). Monetary policy, corporate finance and investment. Working Papers 1911, National Bureau of Economic Research.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the Nature of Capital Adjustment Costs. *The Review of Economic Studies*, 73(3):611–633.

Covas, F. and Haan, W. J. D. (2011). The Cyclical Behavior of Debt and Equity Finance. *American Economic Review*, 101(2):877–899.

Crouzet, N. (2017). Aggregate Implications of Corporate Debt Choices. *The Review of Economic Studies*, 85(3):1635–1682.

Crouzet, N. and Mehrotra, N. R. (2020). Small and large firms over the business cycle. *American Economic Review*, 110(11):3549–3601.

Del Negro, M., Eggertsson, G., Ferrero, A., and Kiyotaki, N. (2017). The great escape? a quantitative evaluation of the fed's liquidity facilities. *American Economic Review*, 107(3):824–57.

Dinlersoz, E., Kalemli-Ozcan, S., Hyatt, H., and Penciakova, V. (2018). Leverage over the Life Cycle and Implications for Firm Growth and Shock Responsiveness. NBER Working Papers 25226, National Bureau of Economic Research, Inc.

Drechsel, T. (2023). Earnings-Based Borrowing Constraints and Macroeconomic Fluctuations. *AEJ: Macroeconomics, forthcoming*, (pdr141).

Duygan-Bump, B., Levkov, A., and Montoriol-Garriga, J. (2015). Financing constraints and unemployment: Evidence from the great recession. *Journal of Monetary Economics*, 75(C):89–105.

Díez, F. J., Fan, J., and Villegas-Sánchez, C. (2021). Global declining competition? *Journal of International Economics*, 132:103492.

Fort, T. C., Haltiwanger, J., Jarmin, R. S., and Miranda, J. (2013). How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size. *IMF Economic Review*, 61(3):520–559.

Gavazza, A., Mongey, S., and Violante, G. L. (2018). Aggregate recruiting intensity. *American Economic Review*, 108(8):2088–2127.

Gertler, M. and Gilchrist, S. (1994). Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms*. *The Quarterly Journal of Economics*, 109(2):309–340.

Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34.

Giroud, X. and Mueller, H. M. (2017). Firm leverage, consumer demand, and employment losses during the great recession. *The Quarterly Journal of Economics*, 132(1):271–316.

Gupta, V., Koller, T., , and Stumpner, P. (2021). Reports of corporates'demise have beengreatly exaggerated. *McKinsey & Company Strategy & Corporate Finance Practice*.

Hall, R. E. (2017). High discounts and high unemployment. *American Economic Review*, 107(2):305–30.

Haltiwanger, J., Jarmin, R. S., and Miranda, J. (2013). Who Creates Jobs? Small versus Large versus Young. *The Review of Economics and Statistics*, 95(2):347–361.

Jeenas, P. (2019). Firm balance sheet liquidity, monetary policy shocks, and investment dynamics. Technical report, WP.

Jensen, T. L., Lando, D., and Medhat, M. (2017). Cyclicality and Firm Size in Private Firm Defaults. *International Journal of Central Banking*, 13(4):97–145.

Jermann, U. and Quadrini, V. (2012). Macroeconomic effects of financial shocks. *American Economic Review*, 102(1):238–71.

Jo, I. H. and Senga, T. (2019). Aggregate consequences of credit subsidy policies: Firm dynamics and misallocation. *Review of Economic Dynamics*, 32:68–93.

Khan, A. and Thomas, J. (2013). Credit shocks and aggregate fluctuations in an economy with production heterogeneity. *Journal of Political Economy*, 121(6):1055–1107.

Kiyotaki, N. and Moore, J. (1997). Credit Cycles. *Journal of Political Economy*, 105(2):211–248.

Kochen, F. (2022). Finance over the life cycle of firms. jmp.

Mian, A. and Sufi, A. (2014). What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223.

Moscarini, G. and Postel-Vinay, F. (2012). The contribution of large and small employers to job creation in times of high and low unemployment. *American Economic Review*, 102(6):2509–39.

Nikolov, B., Schmid, L., and Steri, R. (2018). The Sources of Financing Constraints. Swiss Finance Institute Research Paper Series 18-74, Swiss Finance Institute.

Ottonello, P. and Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, 88(6):2473–2502.

Poeschl, J. (2023). Corporate debt maturity and investment over the business cycle. *European Economic Review*, 152:104348.

Rovigatti, G. and Mollisi, V. (2016). PRODEST: Stata module for production function estimation based on the control function approach. Statistical Software Components, Boston College Department of Economics.

Sedláček, P. and Sterk, V. (2017). The growth potential of startups over the business cycle. *American Economic Review*, 107(10):3182–3210.

Sharpe, S. A. (1994). Financial Market Imperfections, Firm Leverage, and the Cyclicality of Employment. *American Economic Review*, 84(4):1060–1074.

APPENDICES

FOR ONLINE PUBLICATION ONLY

A Data appendix

A.1 Additional information on dataset building process

We use the following datasets provided by Statistics Denmark (DST). All data is at the yearly frequency.

- FIGT (“Gammel Firmastatistik”) + FIGF (“Gammel firmastatistik regnskabsdata”, 1992-1999), FIRM (“Firmastatistik”, 1999-2019): general firm-level data for sales and employment.
- URTEVIRK (2003-2019): bank loans to firms. This dataset includes information about all bank non-mortgage loans to companies aggregated from account level to firm level. It includes both secured and unsecured loans.
- FIRE (“Regnskabsstatistikken”, 2001-2019): Firm-level accounting data such as assets or debt of non-financial corporations. Information in FIRE comes *either a survey done by DST* with a rotating sample of approximately 9000 firms, *or directly tax authority SKAT*. The sampling of firms in the survey depends on their size: firms with more than 50 workers are always included, 20-49 are included for 5 years every 10 years, firms with 10-19 workers are included every 2 years every 10 years

Overall, we have universal coverage of Danish firms regarding employment and sales as well as financial variables for the period starting from 2001 until 2019. Due to the stratified sampling of balance sheet data from FIRE, for smaller firms we do not have information for every firm in every year, but we do have data that has positive coverage even up to the smallest firms. For the 90’s we are missing debt and asset information. Furthermore, the number of firms is lower with some sectors missing. However, the exact pattern why some firms are present and some not is not clear. Finally from 2003 we also have information about bank loans for the universe of firms.

Subject to some minimal threshold on economic activity,[44] all firms are legally obliged to report data to SKAT or DST, which are then collected in these databases. We drop all observations that we deem as inactive by our definition, that is firms that provide no information about employment, sales, value added, or profits.

We also drop all firms that never in their life employ more than one worker.[45] Finally, we also drop firms listed as non-profits as well as entities controlled by government at any level. In our baseline exercises, we include only firms that do not exit in the current or the next year. We thus do not separately investigate the role of firm entry or exit in driving cyclicality.

Sometimes, information about a particular variable for a given firm is missing in the aforementioned registers. This is more likely for for financial rather than real variables, for smaller firms and for firms in the process of exiting. The year of exit also causes problems for variables that measure stock at a given point in time, rather than annual average. For these reasons, we only consider observations for firms that are not exiting in a given year. Additionally, we require lagged information about assets and debt to be present for the regressions. This way we make sure that the estimated effects of including or excluding leverage controls are not the result of changing the set of firms in the sample.

Firm-level growth outcomes are defined by the normalised growth rates suggested by Haltiwanger et al. (2013): for any firm-level variable $x_{i,t}$, we measure growth from $t-1$ to t as

$$\hat{g}_{x_{i,t}} \equiv \frac{x_{i,t} - x_{i,t-1}}{\frac{1}{2}(x_{i,t} + x_{i,t-1})},$$

where i indexes firms and t years. As discussed by Haltiwanger et al. (2013), this growth rate, which uses the average of the current and past value as the denominator, rather than just the past value, is more robust and typically has better properties in firm-level data.

A.2 Evolution of size thresholds and age distribution

While the size thresholds are relatively stable over time, the number of firms in different age groups changes over time, both in absolute numbers but also in group size ranking.

[44] In most situation, firms that report employment that corresponds to less than 0.5 full time worker are considered inactive by DST, but still present in our data.

[45] We do this to eliminate sole proprietorship firms and also firms that exist due to tax optimisation purposes.

Pro-cyclical firm entry generates stronger and weaker cohorts which propagates over time to higher age groups. While interesting on its own (Sedláček and Sterk, 2017), we ignore the potentially link between average firm quality and the state of the business cycle at the time of entry.

Figure 12: Size Thresholds and Number of Firms in Different Age Bins

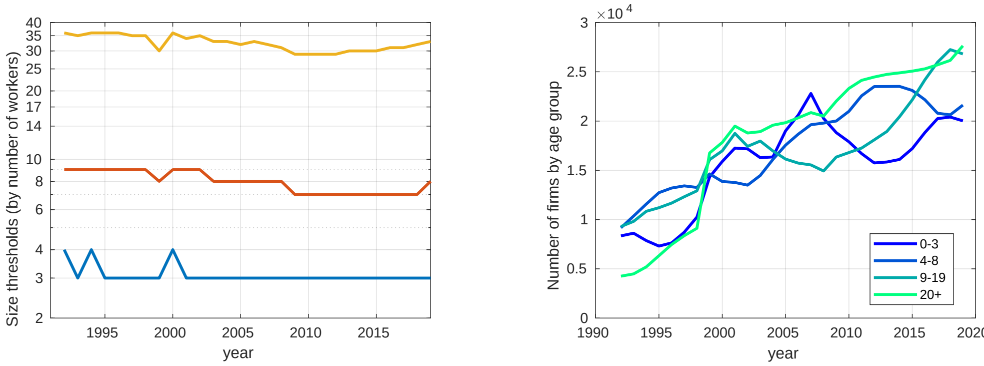

(a) Size thresholds (log scale)

(b) Number of firms in each age group

Panel (a) gives the size thresholds for our baseline empirical results. Firms at and below the blue line are in the 0-30% group, those between the red and blue lines are in the 30-60% group, and so on. Panel (b) gives the number of firms in each age bin, with age measured in years.

B Empirical analysis appendix

B.1 Confidence intervals around cyclicality results

In Figure 13 we repeat our main cyclicality plot, Figure 3, now including 95% confidence intervals around all coefficients. The confidence intervals are narrow enough to statistically distinguish the youngest versus the oldest firms among the smallest size bin group. Among the largest firm group, the cyclicality of firms of all ages are very similar in point estimate, and are not significantly different.

B.2 Cyclicality measured by size and age separately

In Figure 14 we compare our main cyclicality by joint age-size bin results, panels (a) and (d), with two sets of simpler results. In panels (b) and (e) we run regressions of cyclicality on size bins only, with no age bins or their interaction included. In panels (c) and (f)

Figure 13: Cyclicality of Sales and Employment with Confidence Intervals

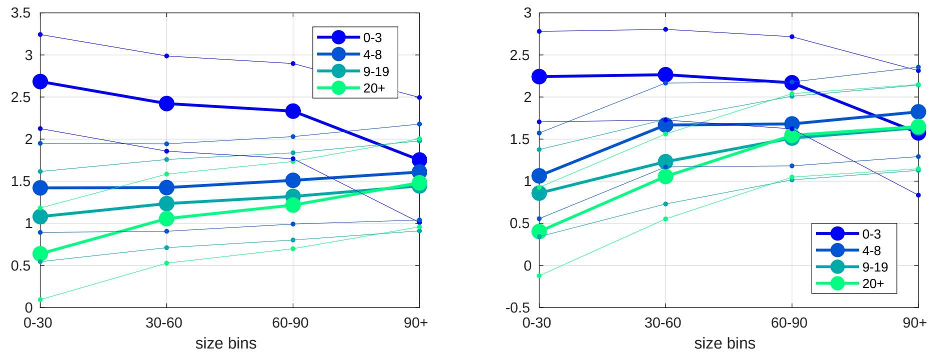

(a) Sales - baseline

(b) Employment - baseline

Cyclicality is the coefficient β_{lk} corresponding to given size and age bin from regression (2). Lines correspond to age bins and x-axis to size bins. Solid lines represent the mean estimated effect and the thin lines of the corresponding color represent $\pm$ 2 standard error confidence bars.

we run regressions of cyclicality on age bins only, with no size bins or their interaction included. As discussed in the main text, the gradient of cyclicality with respect to size alone (panels b and e) is relatively flat, because it averages the upwards and downwards slopes for different age groups (panels a and d). The numbers behind these results are also shown in Table 4.

B.3 Results with assets as the size sorting variable

In this section we show that defining size bins using a firm's assets rather than their employment leads to similar results to our baseline results from the text. In Figure 15 we repeat Figure 2 and Figure 3 but now defining size bins using assets rather than employment. The main patterns are still visible: Employment growth is higher at younger, especially young-small, firms, and cyclicality is increasing in size for older firms and decreasing in size for younger firms.

B.4 Results from different samples

In Figure 16 we consider two alternative samples. In panels (a) and (b) we repeat our cyclicality results for our baseline sample, which runs from 2001 and includes only firm-year observations with valid balance sheet data. The balance sheet data requirement restricts the sample because firstly balance sheet data is stratified and available with lower

Figure 14: Cyclicality by size and age separately

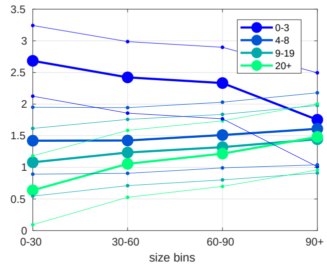

(a) Sales: baseline

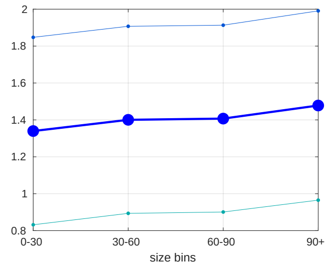

(b) Sales: by size

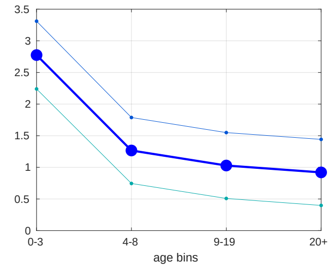

(c) Sales:by age

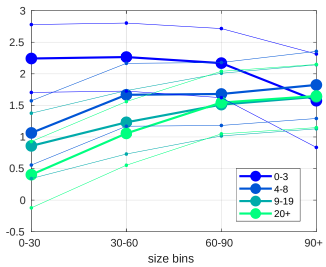

(d) Employment: baseline

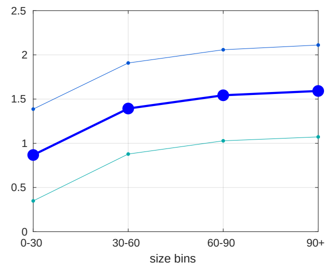

(e) Employment: by size

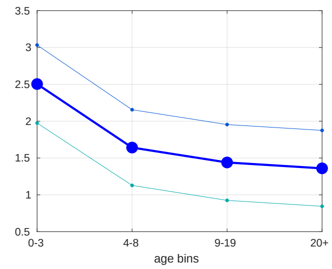

(f) Employment: by age

In panels (a) and (b), cyclicality is the coefficient β_{lk} corresponding to given size and age bin from regression (2). Lines correspond to age bins and x-axis to size bins. Solid lines represent the mean estimated effect and the thin lines of the corresponding color represent $\pm$ 2 standard error confidence bars. Panels (b) and (e) give results for a specification with size bins only but no age bins, and (c) and (f) for a specification with age bins only.

probability for small firms, and secondly, only available from 2001 onwards. We thus investigate the role of dropping this requirement on our main results for non-financial variables. In panels (c) and (d) we keep the baseline sample period, 2001 onwards, and now include even firm-year observations which have valid sales and employment data but not balance sheet data. We see that our main findings are preserved qualitatively, and with only minor quantitative changes. In panels (e) and (f) we do the same, and now also extend the sample back to 1991. Here the magnitudes change somewhat, and the results become a little more split between sales and employment. The finding that young firms are more cyclical is robust to this extended time period for both sales and employment. The finding that cyclicality is declining in size among young firms (age 0-3) is now more pronounced for sales, but still present for employment from size 30-60% onwards. The finding that cyclicality is increasing in size for older firms (age 20+) is more pronounced for employment, and sales have only a minor positive slope from size 30-60% onwards.

Figure 15: Average Levels, Growth Rates and Cyclicality with Size Defined by Assets

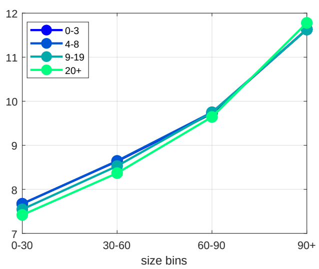

(a) log Sales

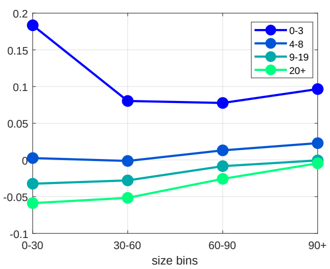

(b) Sales growth

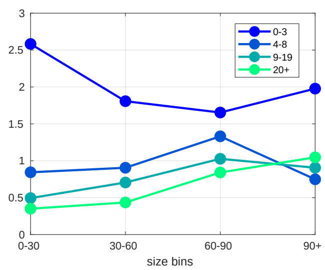

(c) Cyclicality: Sales

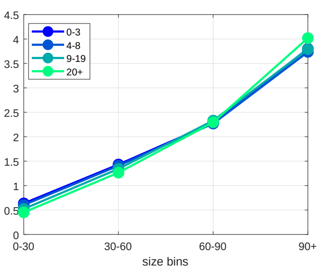

(d) log Employment

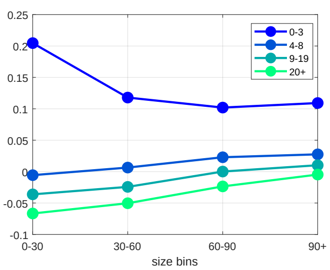

(e) Employment growth

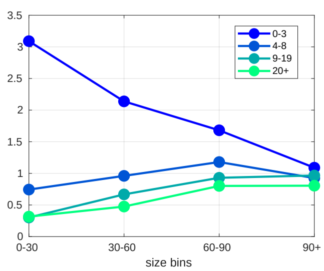

(f) Cyclicality: Employment

These figures repeat Figure 2 and Figure 3 but now defining size bins using assets rather than employment. Average level and growth rate are computed as coefficient α_{lk} from regression (1) with log-level and growth rate as left-hand side variable respectively. Cyclicality is the coefficient β_{lk} corresponding to given size and age bin from regression (2). Lines correspond to age bins and x-axis to size bins.

Figure 16: Cyclicality by Size and Age in different samples

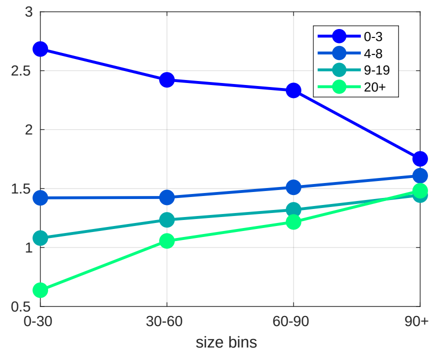

(a) Sales, baseline (2001-2019)

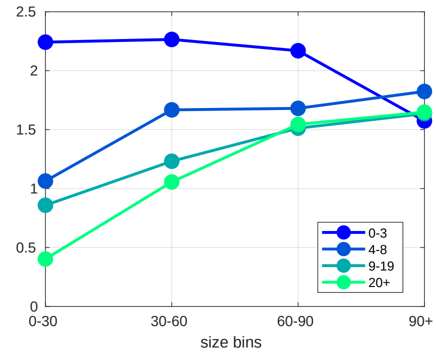

(b) Employment, baseline (2001-2019)

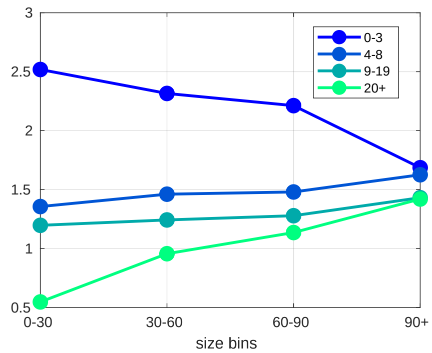

(c) Sales, alt. (2001-2019)

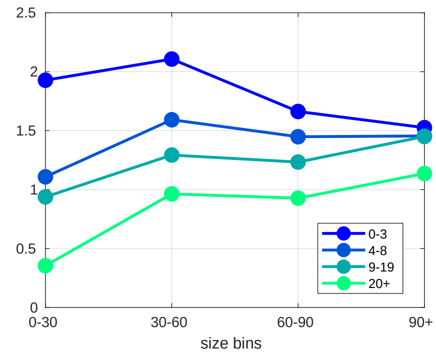

(d) Employment, alt. (2001-2019)

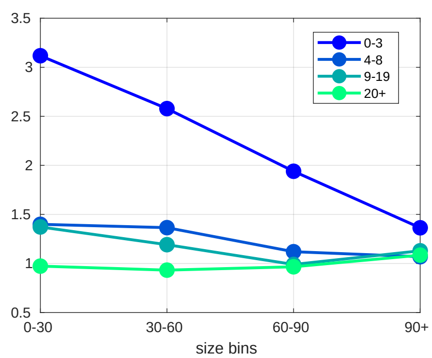

(e) Sales, alt. (1991-2019)

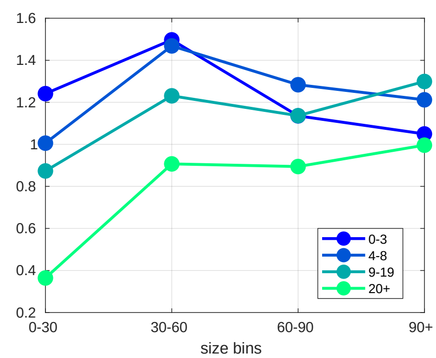

(f) Employment, alt. (1991-2019)

Cyclicality is the coefficient β_{lk} corresponding to given size and age bin from regression (2). Lines correspond to age bins and x-axis to size bins. Panels (a) and (b) give our baseline sample, which runs from 2001 (effectively 2002 as we require growth rates and lags) and includes only firm-year observations which have valid balance sheet data. Panels (c) and (d) keep the same time period but now include firm-year observations even if they are missing balance sheet data. Panels (e) and (f) do the same, and also extend the sample back to 1991.

B.5 Regression results

In this section we plot the regression results tables behind certain key figures. Table 3 gives the cyclicality regressions by leverage bin, with and without age-size controls, from regressions (3) and (4) shown in Figure 5. Table 4 gives the cyclicality regressions by joint age-size bin from regression (2) (as shown in Figure 3), as well as results by size and age separately (as shown in Figure 14). It is also contains a version of our main results controlling for firm fixed effects, which hence only includes firms which move across size or age bin at least once over their life, where the results are very similar to our baseline specification.

Table 3: Effects of Leverage on Cyclicality

	(1) sales	(2) sales	(3) sales	(4) sales	(5) employment	(6) employment	(7) employment	(8) employment
$1 \times$ y	-0.32***	-0.23***	-0.30***	-0.16**	-0.54***	-0.43***	-0.47***	-0.30***
	(-4.63)	(-3.30)	(-4.37)	(-2.37)	(-8.93)	(-7.12)	(-7.71)	(-4.93)
$2 \times$ y	-0.23***	-0.18***	-0.23***	-0.18***	-0.23***	-0.17***	-0.20***	-0.14**
	(-3.49)	(-2.75)	(-3.42)	(-2.70)	(-4.01)	(-2.93)	(-3.52)	(-2.50)
$4 \times$ y	-0.08	-0.11	-0.08	-0.10	0.15**	0.12**	0.16***	0.14**
	(-1.23)	(-1.61)	(-1.19)	(-1.47)	(2.46)	(1.99)	(2.58)	(2.26)
$5 \times$ y	-0.07	-0.19**	-0.05	-0.14*	0.22***	0.07	0.26***	0.15**
	(-0.88)	(-2.50)	(-0.70)	(-1.87)	(3.08)	(0.92)	(3.57)	(2.05)
Observations	649856	649856	649856	649856	649925	649925	649925	649925
Adjusted R^2	0.017	0.036	0.017	0.039	0.014	0.061	0.018	0.068
Size and Age controls	-	Age	Size	Age X Size	-	Age	Size	Age X Size

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

This table presents the coefficients ψ from different versions of regression (4), depending on the degree of controlling for size and age. The degree of leverage is captured by Debt/Assets quintile. Firms with median leverage (between 40th and 60th percentile) are treated as the base group. Standard errors are clustered at the firm level. Constant, sector-related and average growth rate related coefficients are omitted.

Table 4: Cyclicality of Sales and Employment by Firm Size and Age

	(1) sales	(2) employment	(3) sales	(4) employment	(5) sales	(6) employment	(7) sales	(8) employment	(9) sales	(10) employment
y	1.41***	1.46***	1.08***	1.42***	1.25***	1.38***	0.92***	1.10***	1.41***	1.44***
	(5.35)	(5.46)	(4.11)	(5.11)	(4.82)	(5.21)	(3.55)	(4.01)	(5.46)	(5.21)
0-30 × y	-0.85***	-1.14***	-0.71***	-1.02***	-0.40***	-0.83***			-0.17**	-0.67***
	(-6.50)	(-9.28)	(-5.14)	(-8.11)	(-5.19)	(-11.48)			(-2.28)	(-8.92)
30-60 × y	-0.42***	-0.55***	-0.36***	-0.52***	-0.28***	-0.35***			-0.09	-0.19***
	(-3.86)	(-6.35)	(-3.16)	(-5.86)	(-4.02)	(-5.51)			(-1.32)	(-3.01)
60-90 × y	-0.27***	-0.13*	-0.23**	-0.06	-0.18***	-0.17***			-0.08	-0.08
	(-2.89)	(-1.85)	(-2.42)	(-0.85)	(-2.82)	(-3.04)			(-1.32)	(-1.52)
0-3 × y	0.26	0.07	0.03	-0.11	1.41***	1.24***	1.71***	1.16***		
	(0.93)	(0.26)	(0.10)	(-0.34)	(16.54)	(13.88)	(19.64)	(12.45)		
4-8 × y	0.14	0.24*	0.06	0.08	0.41***	0.43***	0.34***	0.32***		
	(0.91)	(1.91)	(0.37)	(0.65)	(6.97)	(8.32)	(5.95)	(6.26)		
9-19 × y	-0.03	0.02	-0.14	-0.12	0.15***	0.07	0.10*	0.04		
	(-0.25)	(0.19)	(-1.16)	(-1.33)	(2.60)	(1.44)	(1.76)	(0.83)		
0-30 × 0-3	1.70***	1.66***	1.64***	1.41***						
× y	(5.15)	(4.85)	(4.63)	(3.70)						
0-30 × 4-8	0.65***	0.48***	0.84***	0.81***						
× y	(3.17)	(2.62)	(3.81)	(4.17)						
0-30 × 9-19	0.48**	0.26	0.76***	0.62***						
× y	(2.57)	(1.55)	(3.83)	(3.46)						
30-60 × 0-3	1.05***	1.20***	0.94***	0.79**						
× y	(3.27)	(3.63)	(2.75)	(2.15)						
30-60 × 4-8	0.23	0.35**	0.32	0.51***						
× y	(1.24)	(2.27)	(1.62)	(3.16)						
30-60 × 9-19	0.21	0.16	0.42**	0.45***						
× y	(1.26)	(1.24)	(2.42)	(3.29)						
60-90 × 0-3	0.79**	0.51	0.76**	0.15						
× y	(2.50)	(1.54)	(2.26)	(0.42)						
60-90 × 4-8	0.16	-0.03	0.15	-0.07						
× y	(0.90)	(-0.25)	(0.83)	(-0.49)						
60-90 × 9-19	0.13	-0.02	0.24	0.07						
× y	(0.93)	(-0.18)	(1.59)	(0.59)						
Observations	651292	669005	651292	669005	651292	669005	663743	678614	651292	669005
Adjusted R^2	0.041	0.132	0.026	0.070	0.039	0.123	0.049	0.133	0.016	0.039
Size and Age controls	Size x Age	Size x Age	Size x Age	Size x Age	Size + Age	Size + Age	only Age	only Age	only Size	only Size
Regression method	OLS	OLS	FE	FE	OLS	OLS	OLS	OLS	OLS	OLS

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Columns (1) and (2) present the cyclicality coefficients by joint age-size bin from regression (2). Standard errors are clustered at the firm level. Coefficients related to sectors omitted. The coefficients are relative to the base groups defined as the oldest age (20+) and the largest size (size percentile 90+). For robustness, we also estimate a version of regression (2) with firm fixed effects (and hence omitting the sectoral controls) and report the results in columns (3) and (4). The results are both qualitatively and quantitatively very similar. The remaining columns present restricted versions of this regression omitting age-size interactions (columns 5 and 6), and looking at age (columns 7 and 8) or size (columns 9 and 10 only.

B.6 Returns to scale and firm size in the data

To justify our assumption of differences in returns to scale across firm size groups, in this section we turn to production function estimation on our dataset. We estimate returns to scale as the sum of the factor elasticities across capital and labor estimated using the Olley-Pakes (OP) and Levinsohn-Petrin (LP) methods, both with the Ackerberg-Caves-Frazer (ACF) correction. Specifically, we split firms into three employment size bins: below 7, between 7 and 25 and above 25 workers. We only consider firms that stay within one size bin within the whole sample, to focus on firms near their steady state size and stay closer to the model concept of differences in returns to scale being a permanent characteristic of firms.[46] Then we estimate the production function independently for each size bin.[47]

Since our data contain only revenue and not physical out, we note that we are estimating revenue production functions, and hence returns to scale in revenue (not output). However, this is exactly the concept we are interested in our model, as cyclicality theoretically depends on revenue returns to scale (i.e. the combination of physical returns to scale and the CES demand curve elasticity) and not physical returns to scale, so this does not pose a problem.

In Figure 17 we plot the resulting sum of the labour and capital coefficients when firms from all sectors are pooled together in the estimation. We include sector-time fixed effects to absorb sectoral and business cycle differences across firms, and plot results for various choices of polynomial order in the control function. Figure 18 we plot results when the production function is estimated individually sector by sector. We control for year fixed effects in each regression, and set the polynomial order to 2. We use the Danish industrial classification which corresponds to 2 digit NACE specification. We plot sectors from “manufacturing”, “textiles” to “construction”, “trade” and “transportation”.

In all cases, independent of the estimation framework (OP or LP), the smallest firms have noticably lower returns to scale than the larger firm size groups. The results suggest that the smallest firms have on average decreasing returns to scale whereas the largest firms have increasing returns to scale. We typically find returns to scale of around 0.8 for the smallest firms, and above 1 and approaching 1.05 for the larger firm size groups.

[46]This requirement reduces the number of observations relative to the empirical exercise in the main body of the paper, and for that reason we also redefine the size bins.

[47]Specifically, we use PRODEST estimation routine by Rovigatti and Mollisi (2016).

Figure 17: Aggregate Returns to Scale

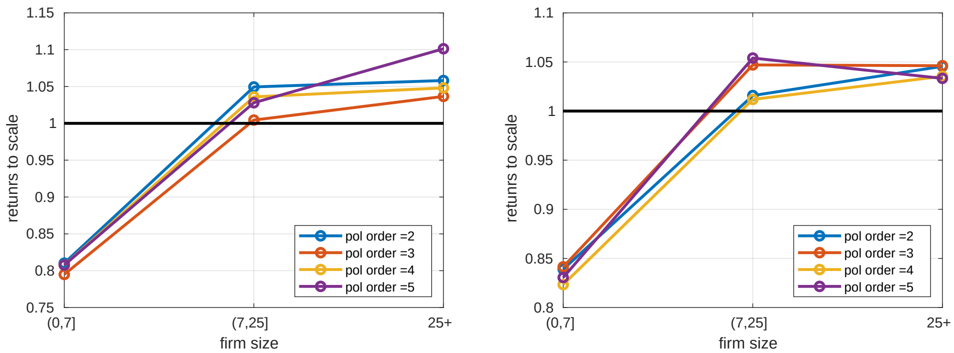

(a) Olley-Pakes method (b) Levinsohn-Petrin

The figures plot the returns to scale from our production function estimation exercise for each size bin. Returns to scale is the sum of the coefficients across the two inputs, capital and labour. Different lines show results for different polynomial order choices in the control function, between order 2 and 5. ACF correction is applied in all cases.

While not an exact match to the differences in returns to scale we reverse engineer in our estimation, the difference from small to large firms of around $1.05 - 0.8 = 0.25$ is very similar to the difference between the returns to scale of the largest and smallest firm size groups in our model. We take this as supportive evidence that returns to scale are an important driver of cyclicality differences across firm size groups.

When estimating value added production functions, the following data is used. For the OP method, the "proxy" variable is investment, whereas for the LP method, the proxy variable is intermediate inputs, which are calculated as value added less turnover. For the "free" variable, labour, full-time equivalent employees are used. The "state" variable, capital, is constructed by the perpetual inventory method using investments and the bookkeeping value of capital. Specifically, an initial capital level is determined as the highest of either reported capital or investments divided by an assumed depreciation rate of 10 percent. For subsequent years, capital is determined by the highest of either reported capital or the depreciated capital determined in the previous year plus investments. If there are any gaps in the series of any variable used in the production function estimation, the capital series is rebased with a new firm identifier. Sectoral deflation is applied to value added, intermediates, investments, and capital using sector level producer price index (PPI) data obtained from Statistics Denmark. Sectoral PPI data is published for varying levels of classification specificity. If several levels of PPI data are published for the sector of a given firm, the level with the highest frequency is used. If more levels

Figure 18: Sector-Specific Returns to Scale

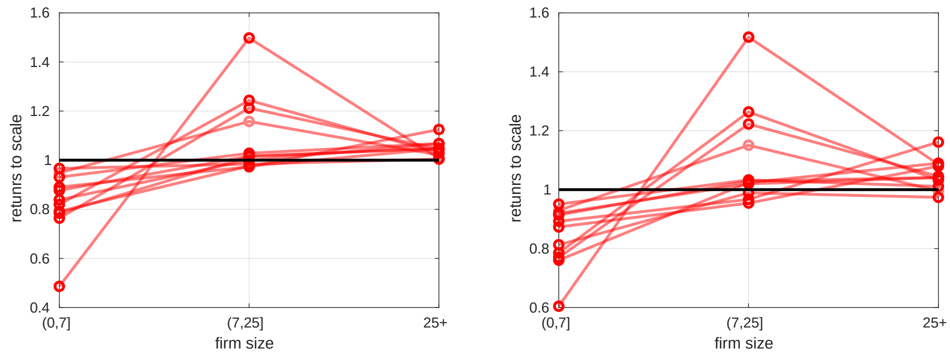

(a) Olley-Pakes (b) Levinsohn-Petrin

The figures plot the returns to scale from our production function estimation exercise for each size bin, now estimated sector-by-sector. Returns to scale is the sum of the coefficients across the two inputs, capital and labour. Different lines show results for different sectors. ACF correction is applied in all cases.

have the same frequency, the most specific level is used. In total, we have PPI data corresponding to 106,170 observations, i.e., 5.5 percent of the total sample, where two-digit sector classification data is used for 86 percent of observations.

Evidence from other papers Another way to interpret our heterogeneous returns to scale assumption is as a heterogeneous *demand elasticity* assumption. In particular, with our CES demand curve, firms face overall returns to revenue $\eta_s\theta$, which is the combination of returns to scale in production, η_s, and the demand elasticity parameter, θ. Therefore, to give smaller firms more decreasing returns to scale, we could equivalently held η_s equal across firms, and used heterogeneous demand elasticities, θ_s, across size groups. This would require setting $\theta_1 < \theta_2 < \theta_3 < \theta_4$, meaning that demand is more inelastic for small firms than large ($\epsilon = 1/(1-\theta)$). In our model, this would mean that small firms charge higher *markups* than large firms, based on the usual result that inelastic demand leads to higher markups: recall the standard static result that the optimal markup is equal to $\mu = \epsilon/(1-\epsilon) = 1/\theta$. This interpretation allows us to use data on average markups by firm size to interpret our assumption of more decreasing returns to scale at small firms. If markups are higher at small firms than large firms, this would provide additional support. Indeed, this appears to be the case in the data. Díez et al. (2021) compute markups for both private and public firms using the global Orbis dataset, for a set of firms accounting for 70% of global GDP. They find that there is a U-shaped relationship between markups and firm size, and that markups are decreasing with firm size for most of the size distribution: *"Contrary to common wisdom, we find that, unconditionally, smaller firms have higher markups even within narrowly defined industries—only when we focus on very large firms we do find a positive relation... markups first decrease with firm size and only when a (fairly large) size threshold is reached, markups start increasing with firm size"* (p2).

C Quantitative model appendix

C.1 Numerical solution details

We solve the model using continuous time numerical methods which draw on Achdou et al. (2021). We use their finite difference methods, and discretize the state variable n with a grid of 1000 nodes. Since the n grid is wide due to the permanent cross sectional heterogeneity between small and large firms, we place these nodes in a non-uniform way to allow more nodes at the low net worth levels experienced by small firms. Ergodic distributions and the aggregate simulations are calculated using the grid based simulation procedure that forms part of the Achdou et al. (2021) method.

To be comparable with the data when running regressions on model-simulated data, we construct time-aggregated yearly data for our regressions. This is done in such a way as to be comparible to our Danish data source. We first solve the transition path of the economy to our aggregate shocks exactly, use the grid-based simulation approach of the Achdou et al. (2021) method, iterating over guesses of aggregate price paths until the economy converges to the true transition path. This ensures an accurate solution to our transition experiments, which does not rely on simulated data from a finite number of firms.

We then construct a panel of 100,000 firms, accounting for entry and exit, which we simulate in response to the aggregate shock. The policy functions of these firms are the policies solved for exactly during the grid-based transition experiment. We aggregate the data up to yearly frequency to make firm-year observations, and regress this data on the growth rate of aggregate output, as done in our data work, using the same regression specification. Since we do not have a notion of industries in our model, we omit the sector dummies from our specification in the model-based regressions. We generate 20 years of data from the model to use for our regressions, which contains the single recession event driven by our MIT shock. Specifically, we allow for 5 years of data pre shock, and then 15 years of data from the moment the shock hits and through the economic recovery.

Table 5: Simulated Minimum Distance Details for "Cyclical" Calibration

Moment	Data	Model	Error	Associated parameter
Average employment growth age 0-3, size 0-30%	0.33	0.33	1.17%	n_1^e
Relative cyclicality age 0-3, size 0-30%	1.36	1.36	-0.35%	$\bar{\phi}_0$
Relative cyclicality age 0-3, size 30-60%	1.37	1.31	-4.96%	n_2^e
Relative cyclicality age 0-3, size 60-90%	1.32	1.33	0.95%	n_3^e
Relative cyclicality age 0-3, size 90%+	0.95	1.05	9.47%	n_4^e
Relative cyclicality age 20+, size 0-30%	0.24	0.26	5.17%	η_1
Relative cyclicality age 20+, size 30-60%	0.64	0.65	1.87%	η_2
Relative cyclicality age 20+, size 60-90%	0.94	0.88	-5.92%	η_3
5% peak GDP fall during recession	-0.05	-0.05	-0.37%	r_0
Average error (sqrt. of mean squared error)	–	–	4.50%	–

Targeted moments in the outer-loop simulated minimum distance estimation for the "cyclical" calibration. Associated parameter is illustrative only, as all parameters are jointly chosen to minimise the mean squared error of all moments.

Estimation of the idiosyncratic shock process: Our procedure broadly follows that of Khan and Thomas (2013). Firstly, the autocorrelation of idiosyncratic shocks is known to be hard to estimate, and Khan and Thomas (2013) choose an annual autocorrelation of 0.659, which we do too. Since we are in continuous time, we first specify that firms draw a new value of their idiosyncratic shock on average once a year, and that when they do it is drawn from a discretised AR(1) process with autocorrelation $\rho^I = 0.659$, mean $\mu^I = 1$, and unconditional standard deviation σ^I. This leaves the standard deviation to calibrate. The Cooper and Haltiwanger (2006) data is for a panel of large manufacturing firms, and since it is a long balanced panel these firms are also older on average. To be consistent with their sample, we thus calculate the standard deviation of investment rates in our model for larger firms, i.e. the largest size group $s = 4$. Since the firms are older in the data, they are likely to be financially unconstrained in our model and we therefore simplify our estimation by simulating a panel of firms who are financially unconstrained. We time-aggregate the simulated capital data to form yearly capital stocks and investment flows in line with how they are constructed in the data, and compute the standard deviation of investment rates on this data. We adjust σ^I to match the standard deviation of investment rates. Since we focus on financially unconstrained firms, the calibration of σ^I is independent of all other parameters in the estimation, apart from the returns to scale of large firms, η_4. To economise on parameters in the estimation, we pre-estimate σ^I by estimating it with a value of $\eta_4 = 1.05$ which is close to our estimated value in the cyclical calibration of the model. We use the same value of σ^I in both the steady state and cyclical calibrations.

Superstar firm details: As part of our calibration procedure, we allow a small number of firms to become "superstar" firms to match the importance of a few very large older firms in the data. We assign all firms a very small probability of becoming a superstar, which happens at rate $\alpha_\star$, so that only 0.5% of firms are superstars in steady state. When a firm becomes a superstar, it switches to a special superstar state with productivity $z_\star$ and returns to scale $\eta_\star = \eta_S$. Given the enormous change in optimal size that happens at this point, we allow firms to raise equity at the moment they become superstars, and allow them continuous access to equity from then on. They therefore become "Modigliani Miller" firms and their financial structure becomes undefined and they follow the efficient investment and production policies. We assume superstars hold a constant leverage rate, calibrated to that of the largest firms in the economy, and that they exit at a low rate of 1% per year, in line with the low exit rates of very large, old firms in the data. Since superstars face no financial frictions, their value can be expressed as $v^\star + n$ for some constant $v^\star$.

Estimation of the "steady state" calibration: The estimation of the "steady state" calibration is relatively simple, because it involves only steady state moments, and no cyclical moments. Parameters are either pre-set to a known value, or chosen to exactly hit one moment using an associated parameter. We use an iterative updating scheme, and stop once all moments are hit with 1% tolerance or less. There are 25 parameters of the model, which are given in the "Steady state" column of Table 6, with each associated moment given in the Source column.

Estimation of the "cyclical" calibration: The estimation of the "cyclical" calibration is more complicated, because we additionally target cyclical moments. For any parameter guess we must 1) solve the steady state of the model, 2) simulate the business cyclical experiment using an MIT shock, and 3) simulate a panel of firms to perform our regressions. We speed up the estimation using a two layer procedure.

In the "outer loop" we choose all parameters which are estimated on cyclical moments. These parameters are jointly chosen to minimize the distance from the cyclical moments using a numerical minimization routine (we use a pattern search algorithm). The nine parameters chosen in the outer loop are $(\eta_1, \eta_2, \eta_3, n_1^e, n_2^e, n_3^e, n_4^e, \bar{\phi}_0, r_0)$. Here $\bar{\phi}_0$ and r_0 refer to the value of the collateral constraint and discount rate shock at time 0. These are chosen to hit the following nine moments: 1) 5% aggregate output fall, 2) average growth rate of "age 0-3, size 0-30%" firm bin, 3) relative cyclicality of age 0-3 firms in all four size bins, and 4) relative cyclicality of age 20+ firms in the 0-30%, 30-60%, and 60-90% size bins.

In the "inner loop" we choose all parameters which are estimated on steady state moments. For any guess of the outer loop parameters, the inner loop chooses the inner loop parameters to exactly hit the inner loop moments (to a 5% tolerance). All parameters in the "Cyclical" column of Table 6 are chosen in the inner loop (apart from the nine outer loop parameters) with associated moment given in the Source column. Note that η_4 is chosen to impose aggregate constant returns to scale, which is done in the inner loop. Similarly, z_1^G is chosen to hit the average employment of aged 0 firms, which is also done in the inner loop.

The values of the moments in the data and the estimated model are given in Table 5. The estimation successfully matches all moments with errors of less than 10%, mostly less than 5%, and the average error (square root of the mean squared error) is equal to 4.5%.

Table 6: Model Parameters and Calibration

	Interpretation	"Steady state"	"Cyclical"	Source
	Parameters used in both calibrations:			
r	Discount rate	0.0202	0.0202	2% yearly real interest rate
δ	Depreciation rate	0.1054	0.1054	10% annual rate
θ	Substitution across varieties	0.9	0.9	10% markup in frictionless model
α	Labor-capital ratio in prod fun	9.1331	8.4815	Aggregate L
μ_0	Firm entry rate	0.0834	0.0834	Normal total mass of firms to one
$\bar{\phi}$	S.s. collateral limit	3	3	Maximum leverage
$\bar{n}$	Net worth where start paying dividends	59.9283	84.3044	Normalisation
χ	Labor disutility shifter	0.0114	0.0114	Labor share of income
σ	Labor supply elasticity	0.3	0.3	Real wage flexibility
α_s	Rate transition to superstar firm	5.1e-05	5.1e-05	0.5% of firms are superstar
$z_\star$	Superstar productivity	0.6393	0.4768	Employment share of firms age 20+
ζ_y	Exit rate when young ($g = 1$)	0.1415	0.1415	Exit rate age 0
ζ_o	Exit rate when old ($g = 2$)	0.0647	0.0647	Average exit rate 8% per year
α_G	Transition rate young to old	0.1964	0.1964	Exit rate age 6
σ^I	Std. idiosyncratic shocks	0.0234	0.0234	Std. investment rates
ρ^I	Autocorr. idiosyncratic shocks	0.6590	0.6590	Khan and Thomas (2013)
z_1^S	Productivity for type $s = 1$	0.3288	0.3137	Av. emp. size 0-30%
z_2^S	Productivity for type $s = 2$	0.3681	0.3454	Av. emp. size 30-60%
z_3^S	Productivity for type $s = 3$	0.4103	0.4000	Av. emp. size 60-90%
z_4^S	Productivity for type $s = 4$	0.5035	0.4183	Normalise $Y = 1$
γ_1^S	Fraction born type $s = 1$	0.3	0.3	Firms for 0-30% size bin
γ_2^S	Fraction born type $s = 2$	0.3	0.3	Firms for 30-60% size bin
γ_3^S	Fraction born type $s = 3$	0.3	0.3	Firms for 60-90% size bin
γ_4^S	Fraction born type $s = 4$	0.1	0.1	Firms for 90%+ size bin
	Parameters used "Steady state" calibration:			
η	Returns to scale (all firms)	1	–	All firms CRS
n^e	Net worth fraction of entrants	0.3543	–	Average employment of age 0 firms
z_1^G	Relative productivity of young	1	–	Not used
	Parameters used in "Cyclical" calibration:			
η_1	Returns to scale ($s = 1$)	–	0.7952	SMM (see Table 5)
η_2	Returns to scale ($s = 2$)	–	1.0407	SMM (see Table 5)
η_3	Returns to scale ($s = 3$)	–	0.9887	SMM (see Table 5)
η_4	Returns to scale ($s = 4$)	–	1.0407	Impose agg. economy has CRS
n_1^e	Net worth fraction of entrants ($s = 1$)	–	0.1937	SMM (see Table 5)
n_2^e	Net worth fraction of entrants ($s = 2$)	–	0.4664	SMM (see Table 5)
n_3^e	Net worth fraction of entrants ($s = 3$)	–	0.9188	SMM (see Table 5)
n_4^e	Net worth fraction of entrants ($s = 4$)	–	0.8297	SMM (see Table 5)
z_1^G	Relative productivity of young	–	0.9289	Average employment of age 0 firms
$\bar{\phi}_0$	Size of collateral constraint shock	–	-0.0926%	SMM (see Table 5)
r_0	Size of discount rate shock	–	0.1562%	SMM (see Table 5)

Parameters and calibration targets for the quantitative model. "Steady state" refers to the calibration of the model to steady-state moments only, from Section 4.2.1. "Cyclical" refers to the calibration to both steady state and business cycle moments, from Section 4.3.1.

C.2 Additional model tables and figures

Figure 19: Fraction of Firms Constrained by Age-Size Bin in the Model

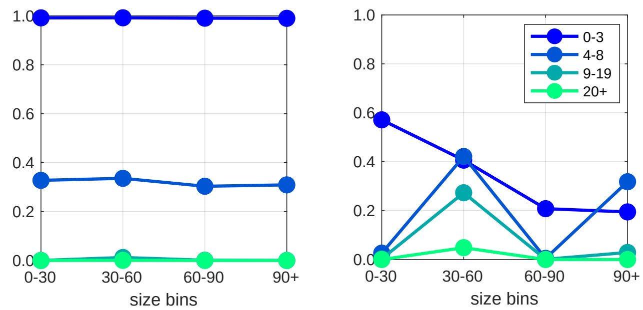

(a) "Steady state" cali (b) "Cyclical" cali

Panels give regression coefficients from regressions of a dummy of whether a firm is financially constrained ($\phi = \bar{\phi}$) on firm age-size dummies, computed from model simulated data. The regressions are on firm-level data aggregated to the yearly level and treated in the same way as the data. Size bins are percentiles, and each line refers to a different firm age group.

Figure 20: Effect of Various Shocks in the “Steady State” Calibration

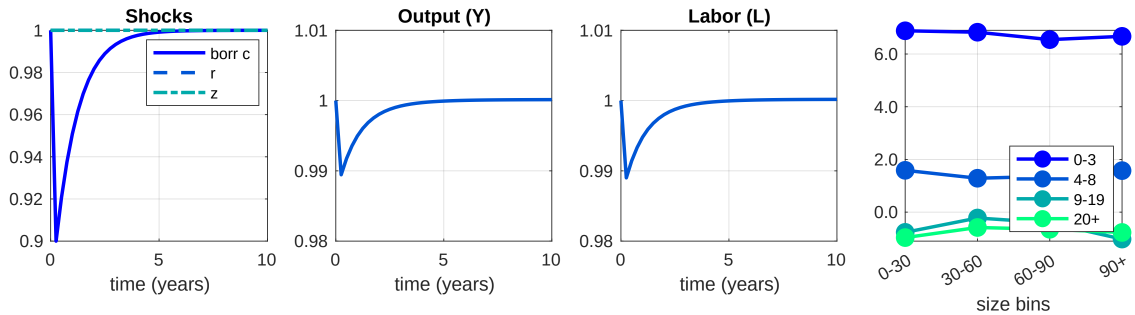

(a) Collateral constraint shock

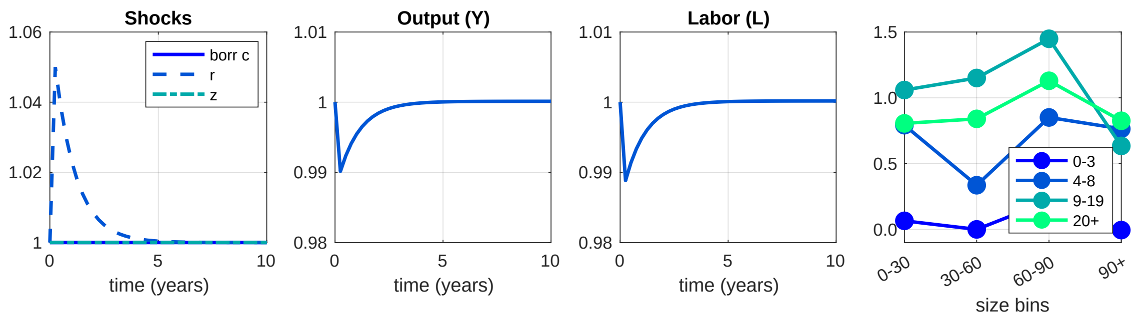

(b) Spread / discount rate shock

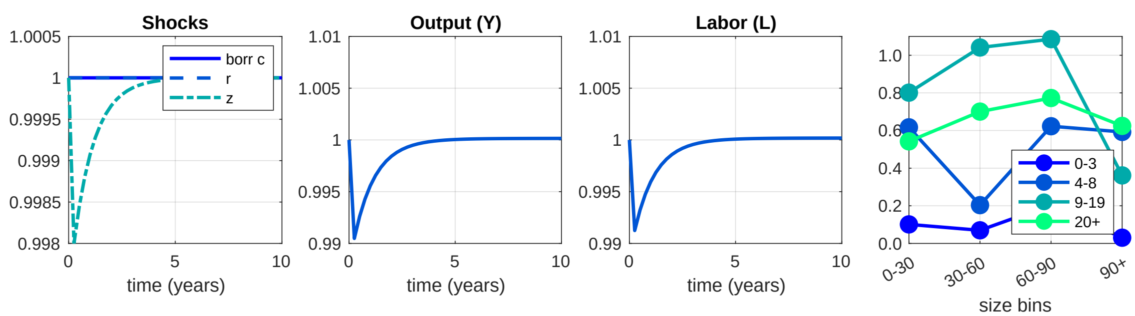

(c) TFP shock

This figure gives simulated aggregate paths and regression coefficients for various recession experiments. In each panel, the left plot gives the shock paths, the center two panels give the paths for aggregate output and labor, and the right panel gives the cyclicality of firm age-size groups computed using our regression approach. In the left panel, “borr c” refers to the path for $\bar{\phi}$, “r” to the path for r, and “z” to the path of the aggregate TFP shock.

Figure 21: Effect of Calibrated Shock Combination in Various Models

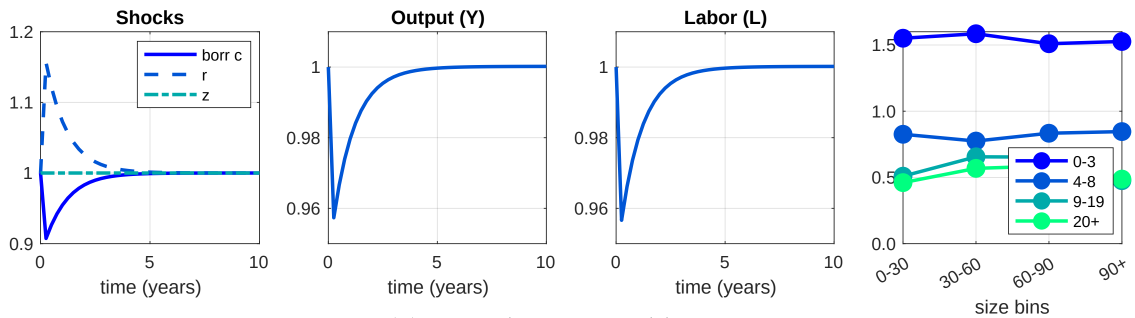

(a) “Steady state” calibration

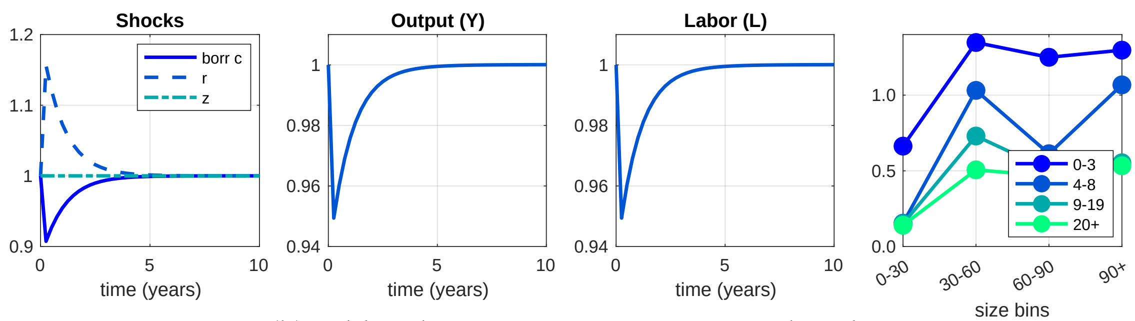

(b) Adding heterogeneous returns to scale only

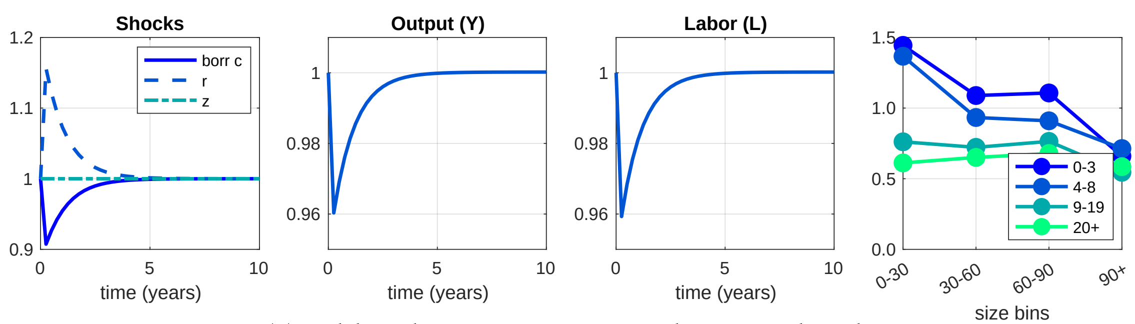

(c) Adding heterogeneous initial net worth only

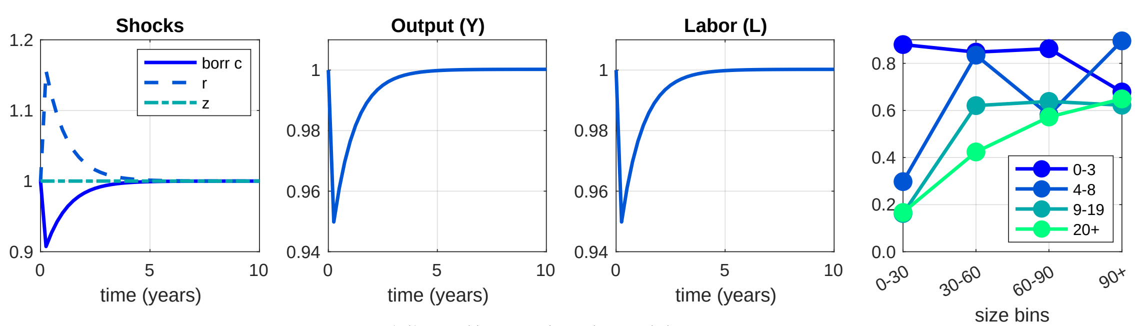

(d) Full “cyclical” calibration

This figure gives simulated aggregate paths and regression coefficients for various recession experiments. In each panel, the left plot gives the shock paths, the center two panels give the paths for aggregate output and labor, and the right panel gives the cyclicality of firm age-size groups computed using our regression approach. In the left panel, “borr c” refers to the path for $\bar{\phi}$, “r” to the path for r, and “z” to the path of the aggregate TFP shock.