Small and Large Firms over the Business Cycle*

Nicolas Crouzet[†] Neil R. Mehrotra[‡]

This version: December 18, 2017

Abstract

Drawing from new, confidential data on income statements and balance sheets of US manufacturing firms, we provide evidence on the relationship between size, cyclicality and financial frictions. First, while sales and investment of smaller firms tend to fluctuate more over the business cycle, the difference is too small to have an impact on aggregates — especially given the high and rising degree of skewness of the firm size distribution. Second, the size effect remains unchanged when directly conditioning on firm-level proxies for financial strength; moreover, while there is a size effect for sales and investment, there is none for measures of external financing. This evidence suggests that the relative behavior of small firms may not be informative about the role of financing frictions in amplifying business cycles.

Keywords: Firm size, business cycles, financial accelerator

JEL Classification: E23, E32, G30

*The views expressed in this paper are those of the authors and do not necessarily represent the view of the US Census Bureau. All findings have been reviewed to ensure that no confidential information is disclosed. We thank seminar and conference participants at Barcelona GSE Summer Forum, Colby College, the Federal Reserve Board, the Inter-American Development Bank, the Minneapolis Fed, Northwestern Kellogg, the Society for Economic Dynamics, the Capital Markets group at the NBER Summer Institute, Minnesota Macro, Wharton, the Phladelphia Fed, the University of North Carolina, Pittsburg/Carnegie Mellon, Boston College, Boston University, UCLA, the NBER Monetary Economic meetings, Imperial College, and Duke Finance for useful comments. We also thank Manuel Amador, Cristina Arellano, Scott Baker, Effi Benmelech, V.V. Chari, Larry Christiano, Anna Cieslak, Janice Eberly, Mark Gertler, Simon Gilchrist, Kyle Herkenhoff, Ben Iverson, Ellen McGrattan, John Haltiwanger, Fabrizio Perri, Mitchell Petersen, Adriano Rampini, Ricardo Reis, and Julia Thomas for helpful discussion and useful suggestions. We thank Apoorv Gupta, Alice Jun, and Adriana Troiano for excellent research assistance.

[†]Northwestern University, Kellogg School of Management, e-mail: n-crouzet@kellogg.northwestern.edu

[‡]Brown University, e-mail: neil_mehrotra@brown.edu

1 Introduction

An important line of research in macroeconomics and corporate finance has sought to document cross-sectional differences in the response of firms to aggregate shocks. Following the work of Gertler and Gilchrist (1994), this literature has paid close attention to firm size. This focus was motivated by the idea that, since size may proxy for financial constraints, a higher sensitivity of small firms would provide evidence in favor of the "financial accelerator" — the view that financial frictions can amplify downturns.[1] However, largely because of data limitations, there remains vigorous debate about both the basic facts and their financial interpretation. More generally, relatively little is known about systematic differences in business-cycle sensitivities across firms.

In this paper, we bring new evidence to bear on these issues. We address three central questions. First, are small firms more cyclically sensitive than large firms, and if so, to what extent? Second, does this excess sensitivity substantially amplify aggregate fluctuations? Third, is this excess sensitivity a manifestation of cross-sectional differences in access to finance?

Our evidence comes from the confidential microdata underlying the US Census Bureau's Quarterly Financial Report (QFR), a survey which collects information on sales and financial liabilities of manufacturing, retail and wholesale trade firms. We use these micro records (income statements and balance sheets) in order to assemble a representative, quarterly panel of US manufacturing firms from 1977 to 2014. The resulting dataset is made up of approximately 900000 observations on approximately 80000 different firms. With this dataset, we then quantify the excess sensitivity of firms at the bottom of the size distribution, relate it to the behavior of aggregate quantities in our sample, and assess whether excess sensitivity is evidence of a financial amplification mechanism.

To our knowledge, this paper is the first to use this firm-level data in its panel format. The firm-level microdata of the QFR carry several advantages relative to both the publicly released version of the QFR, and to alternative firm-level datasets. Relative to the universe totals released in the public version of the QFR, most importantly, the firm-level data allows for an analysis in which controls can be introduced for financial factors that may be correlated with firm size and account for any measured excess sensitivity of small firms. Furthermore, we avoid a number of aggregation biases; in particular, the bins of asset value which are used to tabulate the publicly released totals are defined in nominal terms, and infrequentely adjusted over time, thus creating attrition in lower size group.

Relative to alternative firm-level datasets, the QFR micro-data also carry several advantages.

[1]The view that financial frictions may be responsible for the excess sensitivity of small firms in recessions is buttressed by an extensive corporate finance literature in which private and bank-dependent firms are often treated as being more financially constrained. Farre-Mensa and Ljungqvist (2016) provide an overview of measures of financial constraints commonly used in the corporate finance literature. Size is often used alone or as part of an index as a proxy for financial constraints - see Rajan and Zingales (1995), Almeida, Campello and Weisbach (2004), Whited and Wu (2006), and Hadlock and Pierce (2010).

Compustat, the typical firm-level data source on US firms, is limited to publicly traded firms; its sample is not representative of the cross-section of US firms. By contrast, the QFR panel is constructed by Census to accurately reflect the cross-section of US manufacturing firms. Having a representative sample of US firms matters not only when assessing the implications of excess sensitivity for aggregates, but also when linking it to financial constraints: Compustat indeed omits private, bank-dependent firms, precisely those the most likely to be financially constrained.[2]

Using the QFR microdata, we find evidence in favor of the excess sensitivity of small firms. On average over the sample, we find that the difference between sales growth of the bottom 99% of firms and the top 1% of firms exhibits a strong contemporary correlation with GDP. Our baseline estimate is that a 1% drop in GDP is associated with a 2.5% drop in sales at the top 1% of firms and a 3.1% drop in sales in the bottom 99%. The size asymmetry also appears in firm level regressions that control for industry and disaggregate firms into finer size quantiles. Though particular episodes differ, over the five recessions in our sample, sales at small firms contract more than sales at large firms.

Interestingly, the size effect is concentrated at the very top of the distribution - the top 0.5% of the size distribution; we find no evidence of large differences in the sales elasticity to GDP up to the 99.5th percentile. This finding, in and of itself, suggests that financial factors may not account for the size effect given the wide range of firm size over which there is no measurable size effect. Firm size in our data ranges from less than $200K for the smallest firms to $750 million (real 2009 dollars) for firms in the 99th percentile. Ex ante, it is not clear that financial frictions would be similarly severe over this large a range of firm sizes.[3]

The excess sensitivity we uncover for sales growth also holds for inventory growth and investment rates. Smaller firms exhibit stronger cyclical swings in inventory growth and investment, including both total investment and tangible investment (property, plant, and equipment). As with sales growth, this differential is concentrated at the top 0.5% of the asset distribution relative to all other firms. Within the bottom 99.5% of the firm size distribution, we find no difference in cyclicality.

We also extend this evidence on excess sensitivity to the analysis of firm-level responses to monetary policy shocks. Consistent with the results of Gertler and Gilchrist (1994), we find that the excess sensitivity pattern holds around Romer and Romer (1989) dates that appear in our sample.[4] Because there are only five Romer dates, and only two after 1988, we use an alternative

[2] The QFR is used as an input into calculations of corporate profits in the National Income and Product Accounts. Other advantages include the fact that firms report data at the quarterly frequency, disaggregate debt by source (banks, bond markets, commercial paper, and other sources of debt), and are instructed by Census to consolidate statements domestically, in contrast with Compustat where financial statements reflect global operations. See Section 2 for further discussion.

[3] An advantage of our data set relative to the QFR public releases is the ability to investigate for size asymmetries over a wide range of firm size.

[4] Our results however differ in the quantitative implications, particularly so in the second part of the sample after

method to gauge the effects of monetary policy. Specifically, we project firm-level responses of sales and investment on the identified monetary policy shock series of Romer and Romer (2004) (extended up to Wieland and Yang (2016) up to 2007), using a method analogous to Jordà (2005). This approach leads to results that are qualitatively consistent,with small firms more responsive to the shock, but which lacks statistical significance for most dependent variables, with the notable exception of inventories. The excess sensitivity of small firms to aggregate fluctuations overall thus seems stronger than in response to identified monetary policy shocks.

We then show that the excess sensitivity of small firms, while statistically significant, is in general too small in magnitude to have an effect on the cyclical behavior of aggregates. Our data allows us to construct counterfactual paths for aggregate sales growth, inventory growth and investment under the alternative assumption that firm-level cyclical sensitivities are the same in the cross-section, and to plot this counterfactual against realized aggregate sales growth. The difference between the two time series is difficult to detect. This finding is due to the extreme skewness of the distribution of sales and investment in the cross-section of firms. For instance, the top 0.5% of firms accounts for approximately 75% of total sales and 85% of total investment in the latter parts of the sample. Moreover, this concentration has been rising over the last 30 years implying that the relative importance of small firms for the cyclicality of aggregates has, if anything, been declining. Size-driven differences in cyclical sensitivities are too small to counterbalance this skewness.[5] To the extent that alternative monetary or fiscal policies could address this differential cyclicality, our results suggest that those policies would have little effect on aggregate fluctuations.

Our findings verifying the greater cyclicality of small firms beg the question of whether these differences in cyclicality are driven by a financial accelerator mechanism. Gertler and Gilchrist (1994) argued that small firms serve as a proxy for financially constrained firms as these firms exhibit greater bank dependence, cannot issue public debt, and face a higher degree of idiosyncratic risk. We verify that it is indeed the case that small firms do differ from large firms along these dimensions.[6] However, we provide two findings that cast doubt on whether the size difference is driven by a financial accelerator mechanism.

First, we introduce direct controls for balance sheet ratios emphasized in the financial frictions literature that should affect the cost and availability of external financing. We sort firms into leverage, liquidity and bank dependence categories. We also introduce dummies for whether a firm has accessed public debt markets in the past and whether it recently issued dividends. We

1990. We discuss our replication of their analysis in Section 4.6 and Appendix D, and explain the differences in magnitudes we uncover.

[5]Our results with respect to skewness echo Gabaix (2011), but we nevertheless find that the average/typical firm behaves over the cycle in much the same way as the aggregates which are dominated by the behavior of the largest firms.

[6]These average differences in capital structure across size groups is, however, dwarfed by variations in capital structure within each size group.

show that none of these controls eliminates the size differential that we document; additionally, the quantitative magnitude of the size differential is almost unchanged (except when one controls for access to bond markets, in which case the size effect is magnified). We also run triple-interaction regressions, where firm size categories are interacted with measures of financial strength. Effectively, we double sort firms by size and a proxy of financial strength. We find the size effect remains present within both the "constrained" and "unconstrained" group, and that its magnitude is largely unchanged.

Second, in order to address the possibility that size is simply a better proxy for financing constraints than the other ones we consider, we look for additional testable predictions of the view that the size effect reflects financial frictions. Specifically, we study the effects of an aggregate, non-financial shock in a simple heterogeneous firm model of investment with financial frictions. We set up the model in such a way that size perfectly proxies for the severity of financial frictions. We start by noting that, if firms have no access to external financing, the model predicts that small and constrained firms will respond *less* than large and unconstrained firms to aggregate shocks. This result emerges because the shock tends to have a stronger effect on the unconstrained optimal size of unconstrained firms than it does on the net worth and investment capacity of constrained firms.

In order to match the size effect on investment, we therefore extend the model to allow for debt financing. We show that, so long as the borrowing constraint faced by firms is sufficiently procyclical, investment at small firms will respond more to aggregate shocks. However, a side-effect of introducing a strongly procyclical borrowing constraint is that debt financing at small firms also becomes more procyclical than at large firms.

We then contrast these predictions of the model to the data. We construct cumulative changes in sales, investment, and the stock of debt around the recessions in our sample using a simple event study framework. The framework allows us to condition on size, and to assess whether the responses which we document differ between small and large firms. While we document a statistically significant difference in the response of sales and investment across size groups, we find no such difference in the response of debt. Total debt, bank debt and, particularly, short-term debt all behave very similarly among small and large firms in the onset of recessions, contrary to the model's predictions. Overall, neither the regression evidence nor the behavior of debt provide strong support in favor of the view that the size effect is a reflection of financial constraints. It should be noted that the peak in debt financing at small firms occurs closer to recession dates than for large firms, but the differences is not significant at all horizons.

In addition to investigating whether the size effect is driven by access to financing in recessions, we also search for non-financial explanations. We find some limited evidence in favor of a diversification hypothesis - that is small firms are more cyclically sensitive because their customer base is not as well-diversified as large firms. We show that, within 3-digit manufacturing industries, the magnitude of the size effect is correlated with export exposure and downstream diversification.

In the first case, industries that have greater export exposure (as measured by exports as a share of total output) exhibit a larger size effect. To the extent that the largest firms are exporters and international business cycles are imperfectly correlated with the US business cycle, demand at the the largest firms in high-export industries are buffered relative to industries with less export exposure.

In the second case, downstream diversification is measured via a Herfindahl index that measures how broadly an industry's production is used across the economy. Under the assumption that the largest firms within 3-digit industries with a high Herfindahl are the ones selling across industries, then these firms may be better insulated via diversification across customers. This evidence is suggestive but consistent with a non-financial explanation for the size effect.

We conclude by directly examining the recession behavior of firms sorted by financial strength instead of size. We use the same five financial strength indicators as described earlier: leverage, liquidity, bank dependence, access to public debt markets, and dividend issuance. Leverage, liquidity and bank dependence groups all display a behavior qualitatively consistent with the financial accelerator narrative; for example, inventories of bank-dependent firms fall somewhat more during the early stages of recessions. However, in all cases, the difference is not statistically or economically significant. Firms with access to public debt markets display, if anything, a higher sensitivity to recessions. Only the behavior of dividend-issuing firms is significantly different from that of non-dividend issuing firms. Overall, this exercise suggests that these simple proxies for financial strength do not tend to be associated with a higher degree of responsiveness during recessions.

It is worth emphasizing some limits to the scope of our findings. Our data does not allow us to measure employment; thus, we cannot assess the possibility that labor hoarding may differ across small and large firms during recessions. In Section 4, we use firm counts in our data to estimate that the top 1% of firms in our sample have at least 2500 employees. Based on this cutoff, using Business Dynamics Statistics data, firms with less than 2500 employees in manufacturing account for a substantial share of employment (over 50%). Therefore, if the differential sales elasticity we find carries over to employment, small firms may be more relevant for employment fluctuations. Likewise, we cannot rule out large excess sensitivity among non-manufacturing firms, which account for a substantial fraction of business-cycle fluctuations in value added and employment.

The remainder of the paper is organized as follows. Section 2 details the construction of the QFR data set and provides summary statistics for small and large firms. Section 3 provides time series and regression evidence on the response of small and large firms over the business cycle, in recessions, and after Romer and Romer (1989) dates. Section 4 analyzes the aggregate implications of size asymmetries between small and large firms. Section 5 presents findings on whether the size differences we document are evidence of a financial accelerator. Section 6 concludes.

1.1 Related Literature

Our analysis most closely relates to a literature examining the business cycle fluctuations of small and large firms. This literature, beginning with Gertler and Gilchrist (1994), utilizes the public releases of the QFR data to examine the cyclicality of sales at small and large firms. Gertler and Gilchrist (1994) showed that small firms are more sensitive than large firms in response to monetary policy shocks, but, more recently, Chari, Christiano and Kehoe (2013) argue that this differential cyclicality does not hold across all recessions. Using the Gertler & Gilchrist methodology, Kudlyak and Sanchez (2017) show that large firms contract more than small firms in the Great Recession. We are able to replicate the findings of each of these papers using our data set and the Gertler & Gilchrist methodology for classifying large and small firms; we discuss in Section 3 the reason for differences in our results versus this literature.

Given that employment data by firm size is relatively more plentiful than sales or investment data, a larger literature has examined size asymmetries in employment and job flows over the business cycle and sought to quantify the effects of credit supply shocks in the Great Recession. Moscarini and Postel-Vinay (2012) examines differences in job creation between small and large firms over the business cycle while Fort et al. (2013) and Mehrotra and Sergeyev (2016) consider the behavior of job flows and employment by firm size and age. Fort et al. (2013) argue that employment at small-young firms are more sensitive than large-mature firms and appear particularly sensitive to changes in house prices. Using Compustat data, Sharpe (1994) finds that higher leverage firms shed sales and employment faster than lower leverage firms while also finding evidence of a separate size asymmetry. Recent work by Ottonello and Winberry (2017) shows that low leverage firms respond more strongly to monetary policy shocks than high leverage firms among Compustat firms.

A broad empirical literature has examined the role of disruptions to firm credit supply as a driver of particular recessions; much of this work uses firm size as a proxy for financial constraints. Bernanke, Lown and Friedman (1991), Bernanke and Blinder (1992) and Kashyap, Lamont and Stein (1994) all consider the role of a credit channel in explaining specific downturns. In the Great Recession, Chodorow-Reich (2014) finds the largest effects of the credit shock due to Lehman Brothers bankruptcy at small and medium sized firms. Mian and Sufi (2014) use establishment size as a proxy for financing frictions in examining the effect of falling house prices on credit supply. Using a heterogenous firm dynamics model, Khan and Thomas (2013) show that a credit shock generates a sharper fall in employment at financial constrained firms consistent with the behavior of employment small and large firms in the Great Recession. The focus of the bulk of this literature is on the effects of an identified credit supply shock; in this paper, we provide broader, though less precisely identified, evidence on the role of financial frictions across different business cycles. Recent work by Bergman, Iyer and Thakor (2015) investigates the presence of a financial accelerator in the farming sector using exogenous temperature shocks.

We also relate to a literature that studies the cyclicality of firm financing in aggregate and in the cross-section. Jermann and Quadrini (2012) investigates the cyclicality of overall corporate debt and equity, while Covas and Den Haan (2011) argues that the cyclicality of equity financing differs with firm size. Begenau and Salomao (2015) analyze the cyclicality of financing in Compustat data and consider implications in a quantitative firm dynamics model, while Crouzet (2017) studies the implications of substitution between bank and bond financing for aggregate investment. Likewise, Shourideh and Zetlin-Jones (2012) consider differences in the reliance on external financing of small and large firms and provide evidence on the financing of private firms in the UK. Gopinath et al. (2015) draws on balance sheet data for small firms in southern European countries to assess the role of integration and capital misallocation in the 2000s.[7] In contrast to these papers, our data set captures the cyclicality of financing at small, nonpublic firms in the US that are not present in Compustat.

2 Data

2.1 The Quarterly Financial Report

The Quarterly Financial Report (QFR) is a survey of firms conducted quarterly by the US Census Bureau. The survey covers several sectors of the US economy: mining, manufacturing, and wholesale and retail trade firms. Since 2009, the survey has been broadened to include a selected set of firms in service industries. Surveyed firms are required to report an income and balance sheet statement each quarter. Data collected by the QFR is used by Bureau of Economic Analysis as an input in estimates of corporate profits for the national income and product accounts, as well as in various other official statistical publications such as the Flow of Funds.[8]

The QFR data is a stratified random sample. This sample is created using corporate income tax records provided by the Internal Revenue Service (IRS) to the Census Bureau. Any manufacturing firm that files a corporate income tax return (Form 1120) with assets over \$250K may be included in the QFR sample. The random stratification is done by size, meaning that firms above certain size thresholds are included in the QFR sample with certainty, while smaller firms are sampled randomly. Since 1982, firms with more than \$250 million in book assets are sampled with certainty; the microdata therefore includes the universe of such firms. Firms with between \$250K and \$250 million in assets are instead sampled randomly, so that the microdata contains only a representative sample. Specifically, each quarter, a set of firms with between \$250K and \$250 million in book assets is randomly drawn and included in the sample for the following 8 quarters. At the same time, approximately 1/8th of the existing sample stops being surveyed. For the \$250K-\$250 million

[7] See Kalemli-Ozcan et al. (2015) for details on the European firm level balance sheet data used in the paper.

[8] The QFR has its origins in World War II as part of the Office of Price Administration. The survey was administered by the Federal Trade Commission until 1982 when it was transferred to the Census Bureau.

dollar group, the microdata is thus a rotating panel, akin to the Current Population Survey (CPS). The exact coverage of the sample relative to the population of firms varies across quarters, but is typically in the neighborhood of 5-8%. For instance, in 2014q1 (the last quarter of our sample), the QFR surveyed 8122 manufacturing firms, out of an estimated population of 136205. Of these surveyed firms, 3700 had less than \$10 million in assets, 2768 had between \$10 and \$250 million in assets, and 1654 had more than \$250 million in assets.

Firms which are part of the rotating random sample receive a simplified ("short") form requiring them to report their income statement and balance sheet for the quarter. Firms which are sampled with certainty receive a somewhat more detailed ("long") form, which requires them to provide more information on the composition of their debt and their financial assets.[9] Based on the underlying sample frame, the Census Bureau then assigns sampling weights to each firm in order to generate population estimates of quantities of interest.[10]

2.2 Data construction

The micro files of the QFR required substantial initial work in order to construct a usable panel data set. This is because, in comparison to other Census datasets like the Longitudinal Business Database, the QFR microdata almost never been used by researchers, and to our knowledge, not at all since the move to the NAICS classification, in 2000.[11] The Census Bureau provided raw data files from 1977q1 to 2014q1, but these data files were not linked across quarters. To compute investment rates and growth rates, firms had to be linked across quarters. In general, a survey identifier was available; however, changes in the encoding format of the survey identifiers on a number of quarters required us to match firms based on other identifiers. To do so, we relied on the employer identification number (EIN) of firms along with matches based on firm name and location of firm headquarters.

Between 1994 and 2000, the raw Census data files were missing sampling weights. We used public releases of the QFR that contain statistics of the number of firms by strata to reconstruct sampling weights over this period.[12] These weights were also adjusted so that aggregate assets for

[9] The QFR short and long forms are available at http://www.census.gov/econ/qfr/forms.html.

[10] To be more precise, the QFR uses post-stratification sampling weights which are adjusted to reflect potential changes in the composition of size and industry stratum of the firm after the stratum is formed. As a result, sampling weights may vary slightly within firm over the duration of the panel. A detailed exposition of the survey stratification and the methodology used for estimating universe totals is available at https://www.census.gov/econ/qfr/documents/QFR_Methodology.pdf .

[11] The only instance of the use of the QFR microdata of which we are aware is Bernanke, Gertler and Gilchrist (1996), who use the pre-2000 microdata to compare firm-level to aggregate growth in sales.

[12] Aggregates of the QFR are publicly available at https://www.census.gov/econ/qfr/historic.html. In a given quarter, the Census Bureau releases a set of tables by asset size class and industry; one of these tables provides the number of firms by industry and asset size class. For an example, see Table L in http://www2.census.gov/econ/qfr/pubs/qfr09q1.pdf.

manufacturing firms match assets as publicly reported by the Census Bureau. Between 1977 and 1994 and post 2000, we find that, using the Census Bureau's sampling weights, aggregate sales and assets match the publicly available releases.

In addition to linking the firm observations across quarters and imputing sampling weights, we also drop miscoded observations and keep only firms with strictly positive assets and balance sheet data that balances correctly.[13] Less than 0.1% of firm-quarter observations have balance sheets for which the sum of liabilities and equity does not match reported assets within less than 0.01% suggesting that the data suffers from limited misreporting. The cleaned data set we work with contains about 1.5 million firm-quarter observations between 1977q1 and 2014q1, of which about 900K are manufacturing firms.[14]

In this paper, we will focus on two samples. The summary statistics and the time series that do not require the computation of growth rates are built off the full sample of approximately 900K firm-quarter observations for manufacturing firms. We use a different sample for computing growth rates or investment rates: we then require firms to have reported data for the four quarters prior to the observation date, in order to be able to compute the year-on-year changes in quantities of interest. For the majority of small firms, which are tracked for 8 quarters, taking year over year growth rates eliminates approximately half of the observations. This second sample with firm-level growth rates for manufacturing firms contains approximately 460,000 observations.[15]

2.3 Advantages of the QFR

Before discussing the summary statistics of the data set, it is worth comparing the QFR to alternative firm-level data sets and discussing some of its advantages and drawbacks.

The primary firm-level financial data set is Compustat. Relative to Compustat, the main advantage of the QFR is that it constitutes a representative sample of the population of US manufacturing firms, given that the sampling frame is drawn from IRS administrative data and response is mandatory. In particular, it includes private, smaller, bank-dependent firms, which are not covered by Compustat but nevertheless constitute the typical firm in the population. Since these firms are those most likely to suffer from frictions arising from limited access to capital markets, the QFR is a particularly attractive source of information to answer the questions on which this paper focuses.

[13] A final issue was that the data did not have a codebook. Because the contents of variables in the micro-data files were not always named in an unambiguous manner, this meant that it was sometimes not possible to match with certainty variables to survey response items in the short and long form. In order to deal with this issue, we matched the exact dollar values of ambiguously named variables to public reports of corporations with similar consolidation rules as those required by the QFR.

[14] Currently, we have not analyzed the non-manufacturing part of the data set, since firms with less than $50 million in assets are not sampled, but we plan to do so in future work.

[15] The growth rate sample is more than half the full sample due to the presence of large, continually sampled (long-form) firms.

There are three other differences between Compustat and the QFR. First, the income and balance sheet data is reported at a quarterly frequency facilitating business cycle analysis. While a quarterly version of Compustat exists, most analyses (including those focusing on business-cycle facts) use the annual version of the data. The quarterly data in the QFR is updated by firms with high frequency: for example, in any quarter less than 2% of (unweighted) firm-level inventory observations are identical to the previous quarter. Second, the QFR asks firms to classify their liabilities into bank and non-bank liabilities, and for larger firms, to provide estimates of bonds and commercial paper outstanding.[16] This additional firm-level data on the composition of debt by source is not directly available in standard annual versions of the Compustat data set, and requires further merges with other datasets in order to be computed. Lastly, as an input into the national accounts, the QFR asks firms for a domestic consolidation of the financial statements. For firms with significant global operations, a substantial fraction of income may be earned outside the US and a significant fraction of assets may be located outside the US. In principle, the QFR more accurately measures activity within the US relative to Compustat.

For smaller European firms, the Amadeus data set provides income and balance sheet data. In comparison to our data set, the Amadeus data set has greater industry coverage, but has a shorter time span (since 2000) and provides data at an annual frequency (see Kalemli-Ozcan et al. (2015)). Alternative US data sets that provide data on small, private firms includes the Survey of Small Business Finances and Sageworks (see Asker, Farre-Mensa and Ljungqvist (2011)). However, neither data set provides the coverage, frequency, or time horizon which the QFR does.[17]

2.4 Summary Statistics

Table 1 provides summary statistics on key real and financial characteristics for small and large manufacturing firms. These statistics are constructed by grouping firms into quantiles of current book assets, computing moments within bins, and averaging across quarters from 1977q1 to 2014q1. In contrast to public releases of the QFR, which are published by fixed nominal size bins, our definition of size groups adjusts over time with inflation and growth. All nominal values are deflated by the BEA price index for manufacturing, normalized to 1 in 2009q1.[18]

Table 1, panel A clearly illustrates the high degree of skewness in both sales and assets. The top 0.5% of firms in the size distribution have assets of $6.7 billion and sales of $ 1.5 billion annually. By contrast, firms within the bottom 90% of the size distribution have just $2 million in assets and $1.2 million in sales. The resulting extreme degree of concentration of sales and assets

[16] The QFR also require larger firms to provide a highly detailed overview of their financial assets, including, among others, cash and demand deposits inside and outside the US and Federal and local government debt owned. We do not use this data in this paper.

[17] Asker, Farre-Mensa and Ljungqvist (2011) report that the Sageworks database contains financial statement data for about 95,000 firms from 2002 to 2007.

[18] The series is available at `http://bea.gov/industry/gdpbyind_data.htm`.

among the top 0.5% is discussed in more detail in Section 4. As discussed there, investment is also very skewed. However, Table 1 shows that investment rates are comparable across size classes, so that the skewness of investment primarily reflects the skewness of the asset size distribution rather than differences in investment intensity. Finally, note that sales growth is substantially faster at the largest manufacturing firms over this period; consequently, asset concentration in the manufacturing industry has increased markedly over the past 35 years.

Table 1, panel B provides key financial ratios by firm size categories. A standard measure of leverage - the debt to asset ratio — generally decreases across firm size categories. However, a standard measure of liquidity - the cash to asset ratio - is also highest among smaller firms. Overall, net leverage (debt less cash over assets) is fairly stable across size classes providing no evidence that smaller firms carry greater leverage. However, we do find that smaller firms have more short-term debt and bank debt (as a share of total debt), and rely more on trade credit than larger firms consistent with Gertler and Gilchrist (1994).

One clear difference between large and small firms — particularly among the largest 0.5% of firms — is the intangible asset share. Firms in the survey report separately property, plant, and equipment (tangible assets) from other long-term assets. A high share of intangible assets likely reflects the accumulation of goodwill due to past acquisitions, so that the sharp increase in intangible asset share across size classes underscores the importance of acquisitions for growth at the very largest firms.[19] Table 2 contains a more complete decomposition of firms' balance sheet by size groups. This decomposition shows that the higher share of intangible assets does not come at the expense of a lower share of tangible long-term assets among large firms, but rather a substantially lower fraction of short-term assets (receivables and inventory) relative to small firms. Thus, both on the liability and the asset side, large firms' maturity structure is longer than short firms'.

It is worth emphasizing that, despite differences across size classes in various real and financial characteristics, there remains tremendous heterogeneity within size classes. Table 3 provides an approximate interquartile range for sales growth, leverage, and liquidity.[20] For sales growth and leverage, the approximate interquartile range within size bins dwarfs the differences across size bins. The interquartile range narrows for larger size classes, but nevertheless remains substantial. It is worth noting that a substantial fraction of firms carry zero leverage; these zero leverage firms tend to be concentrated in the bottom 90% of the size distribution.

To summarize, on average over the sample, small firms tend to have similar net leverage as large firms, but rely more extensively on bank debt, short-term debt and trade credit. Moreover,

[19] Even for firms with low or zero intangible asset share, the market value of the firm may differ substantially from the book value of the firm. However, our data only contains book value of assets; for most firms in our sample, which are private, no measure of market value is readily available.

[20] Due to data disclosure restrictions, we provide averages above and below the median within size classes, rather than the exact 25th and 75th percentiles.

sales and assets display an extreme degree of concentration among the very largest firms, and increasingly so over time, given their faster average growth rate. Finally, within size classes, firms display substantial heterogeneity in capital structure and firm growth rates. We next turn to differences in cyclical behavior across these size groups.

3 The excess sensitivity of small firms

This section measures the extent to which small firms display "excess sensitivity" to aggregate fluctuations. By "excess sensitivity", we mean that a worsening in aggregate conditions is associated with systematically bigger declines in sales and investment among small firms than among large firms. We quantify excess sensitivity in two ways. First, we construct size-specific correlations of firm-level growth with business cycle fluctuations. Second, we measure size-specific impulse responses to an exogenous tightening of monetary policy.

3.1 The excess sensitivity of small firms to business cycles

3.1.1 Measurement

Appendix A describes the sample selection, the size groupings, and the measures of firm-level growth which we use throughout this section. Three features of this measurement framework are worth emphasizing. First, we measure the sensitivity of *firm-level* growth to aggregate conditions. We thus sort on size at the firm level, and fully control for industry effects (and, in later sections, for firm-level differences in capital structure). This is distinct from previous work on the QFR data, which was limited to measuring the growth of aggregates by nominal size bins, due to the formatting of the public releases of the QFR. The connection between firm-level and aggregate growth is discussed in greater detail in Section 4.

Second, we base our size groups on *quantiles* of the *lagged* empirical distribution of *book assets*. We use *quantiles* — for example, the bottom 99% versus the top 1% — because they are immune to long-run upward size drift due to inflation and real growth. This problem arises when using fixed nominal thresholds as in the public QFR releases. Classifying firms by their *lagged* position in the size distribution helps alleviate the cyclical effects of reclassification bias emphasized in Moscarini and Postel-Vinay (2012).[21] Finally, we use *book assets* because, among the possible measures of size in our data, it is the most stable at higher frequencies. In particular, unlike sales, it does not display substantial seasonal variation at the firm level.

Third, in our baseline estimates, we measure growth among the sample of surviving firms. In particular, we do not take into account the effect of differences in the cyclical sensitivities of the

[21] If firms tend to cross the threshold from small to large during expansions, measures of the relative growth rate of large firms using their ex-post size will be biased upward.

rate of exit of small and large firms. Our baseline results should thus be thought of as capturing the intensive margin differences in the revenue cyclicality of small and large firms. We discuss the impact of exit on our estimates in Section 3.2.

3.1.2 The behavior of sales

Figure 1 shows the time series for the average growth rate of sales of two size groups, the bottom 99% (denoted by $\hat{g}_t^{(\text{small})}$), and the top 1% (denoted by $\hat{g}_t^{(\text{large})}$). Each series is the year-on-year equal-weighted average growth rate of sales among firms belonging to each of the two size groups one year prior.[22]

The most striking feature of these two series is perhaps how closely they track each other (their sample correlation is 0.93). In particular, from 1987 to 1990, 1995 to 2000, and 2002 to 2007, it is difficult to distinguish growth rates across these groups visually. Nevertheless, there are periods of notable divergence. The two periods which stand out the most are 1982q3-1984q1 - the recovery from the Volcker recessions - and 2008q3-2009q4 - the early stages of the Great Recession. In the first instance, the growth rate of small firms far outpaced that of large firms; in the second instance, it was markedly lower. The recovery of the 1990-1991 recession also features a slightly faster growth rate of small firms. Thus, even though visually the common cyclical component in small and large firms' growth stands out most, one cannot rule out that sales growth contains a size-dependent cyclical component.

Figure 2 shows that the difference between small and large firms' average growth rate is positively correlated to GDP growth. This figure plots the time series $\Delta\hat{g}_t \equiv \hat{g}_t^{(small)} - \hat{g}_t^{(large)}$ against year-on-year changes in real GDP. The estimated slope coefficient of the bivariate simple OLS between the two series is 0.60, with a White standard error of 0.11. The economic interpretation of this coefficient is that, for every percentage point decline in GDP, sales decline, on average, by 0.6% more among small firms than they do among large firms.[23]

Table 4 reports estimates of the semi-elasticity of firm-level growth to GDP growth, and confirms the visual impressions conveyed by Figure 2. The model estimated is:

$$g_{i,t} = \sum_{j\in\mathcal{J}} (\alpha_j + \beta_j \Delta GDP_t)\, \mathbf{1}_{i\in\mathcal{I}_t^{(j)}} + \sum_{l\in\mathcal{L}} (\gamma_l + \delta_l \Delta GDP_t)\, \mathbf{1}_{i\in\mathcal{L}} + \epsilon_{i,t}. \tag{1}$$

[22]The specific definition of the time-series reported for the small firm group is given in equation (17) of appendix A, for the interquartile range $(k_1, k_2) = (0, 99)$. The large firm group corresponds to $(k_1, k_2) = (99, 100)$. Unless otherwise noted, all series are deflated by the BEA's chain type price index for manufacturing value added (`bea.gov/industry/gdpbyind.htm`) before computing growth rates.

[23]This correlation is robust to alternative measures of the business cycle: growth rate of overall industrial production or manufacturing IP or the change in the unemployment rate. This correlation also holds for subsamples before and after 1992 and excluding either the Volcker recovery or the Great Recession. However, the correlation becomes insignificant if both the Volcker recovery and the Great Recession are excluded.

Here, i identifies a firm and t identifies a quarter. The dependent variable, $g_{i,t}$, is the year-on-year growth log change in sales. The set $\mathcal{I}_t^{(j)}$ is a size group; for instance, firms below the 90th percentile of the distribution of book assets four quarters ago.[24] Additionally, $\Delta GDP_t = \log\left(\frac{GDP_t}{GDP_{t-4}}\right)$ is the year-on-year growth rate of GDP, and $\mathcal{L}$ is a set of industry dummies.[25] The two main differences between this regression and the simple visual evidence are that this specifications allows for four different size groups (the bottom 90%, 90-99%, 99% to 99.5% and the top 0.5%), instead of two, and that it controls for industry effects.

The first column of Table 4 reports estimates of the difference $\beta_j - \beta_{(0,90)}$, for the size groups $j \in \{(90, 99), (99, 99.5), (99.5, 100)\}$. For these three size groups, the difference is negative, consistent with the view that small firms are more sensitive to aggregate fluctuations. The size effect thus does not simply reflect cyclical differences across industries. The results of Table 4 also reveal that the cross-sectional differences in cyclical sensitivity are most notable among the top 0.5%, which represents approximately 500 firms in each quarter. In particular, relative to the baseline group (where book assets average $2m$\$), the cyclical sensitivity of sales growth among in the 90-99% (where book assets average $49m$\$) is not statistically different; for the 99-99.5% (where book assets average $626m$\$), the cyclicality is slightly smaller, but the significance is marginal.[26] It is really only at the very top that the difference emerges. We have experimented with more size classes; within the bottom 90% of the firm size distribution, we find no evidence of differences in cyclical sensitivity. It is also worth noting that the adjusted R-squared for this regression is quite low, indicating that, despite the obvious common component between small and large firms, there is considerable heterogeneity in sales growth at the firm level.

Figure 3 conveys a similar message but reports estimates of the absolute cyclical sensitivity of each size group. Specifically, it plots the average marginal effect of ΔGDP_t at the mean, for each size group (including the 0-90% group), as well as the unconditional cyclical sensitivity (the red line). The only group with a statistically different elasticity from the unconditional cyclical sensitivity is the top 0.5%. Moreover, note that the absolute magnitude of the elasticities to GDP growth are substantially larger than the cross-group difference. This fact will be important in Section 4 when we consider the aggregate implications for sales of the cross-group difference in elasticities.

A simple summary of the evidence on sales is the following: when GDP growth drops by 1%, the largest firms' sales drop by approximately 2.5%, while the smallest firms' sales drop by approximately 3.1%. This effect is statistically significant, and driven by differences between the top 0.5% and the rest of the firms. We next turn to the evidence on inventory and fixed investment.

[24] See appendix A for a formal definition of the size groups.

[25] The baseline regression results are reported by classifying firms into durable and non-durable industries. Results are unchanged when using NAICS 3-digit industries. Section 5 further discusses the size effect within NAICS 3-digit industries.

[26] Standard errors are clustered at the firm level.

3.1.3 The behavior of investment

The time-series for inventory growth and investment in fixed assets reported in Figure 4 also displays comovement across small and large firms, but to a lesser extent than sales (the respective sample correlations between the small and large time series are 0.64 and 0.52). For inventory, the episodes of notable divergence between small and large firms are two recoveries: the 1983-1985 recovery and the aftermath from the Great Recession. These two episodes convey a mixed message. In particular, in the aftermath of the Great Recession, inventories at large firms actually recovered more quickly.

For fixed investment, the most striking fact is that contractions in fixed investment seem to occur with a lag at larger firms. This is particularly visible during the Volcker recessions. Slowdowns in investment also persist longer; in the aftermath of the 2000-2001 recession, the turning point for investment among large firms occurred approximately 4 quarters later for large firms than for small firms.[27]

The regression evidence, reported in Table 4, provides a clearer picture than the long time series. The second and third columns report estimates of model (1) when the dependent variable is either inventory growth (second column) or the fixed investment rate (third column). Consistent with the behavior of sales, inventory growth of the top 0.5% of firms has a significantly smaller conditional elasticity to GDP growth.[28] The economic magnitude of the effect is large: for the bottom 90%, the average marginal effect of a 1% drop in GDP is a 1.9% drop in inventory, about double the effect for the top 0.5%.

The results for fixed investment are, if anything, starker. The difference between the 99-99.5% and the 99.5-100% groups and the bottom group are both statistically significant. In terms of economic magnitudes, a 1% drop in GDP is associated with a 0.9% drop in investment among the $(0, 99)$ group, relative to a baseline investment rate of approximately 26.0%. Among the $(99, 100)$ group, the investment drop is more muted: 0.15%, relative to a baseline investment rate of approximately 21%. The small estimated elasticity of investment to aggregate conditions among larger firms is likely driven by the fact large firms seem to cut investment with a lag, as mentioned before.

Nevertheless, the overall message is the same as for sales: inventory growth and investment rates among small firms are substantially more sensitive to business cycles than among large firms.[29]

[27] This lag structure also accounts for the fact that the contemporaneous correlation of GDP growth and investment is not significantly positive among the largest firms in the QFR sample, as documented in Table 4. Appendix E discusses this lag in more detail, and shows that it is also present in both the annual and quarterly Compustat data.

[28] As was also the case for sales, the estimated difference in elasticities between the bottom 90% and the top 0.5% lines up with the results of a simple OLS regression of the difference in inventory growth between the top 1% and the bottom 99%, which delivers a slope coefficient of approximately 0.7. Results are not reported but available upon request.

[29] Like sales, the presence of a size effect for inventories and investment holds within subsamples except for when both the Volcker recovery and Great Recession are dropped.

3.2 Firm age and exit

Fort et al. (2013) argue that the business cycle behavior of firm employment depends crucially on firm age (as opposed to simply firm size). Our data set does not have an indicator for firm age. To proxy for firm age in the QFR, we group firms (starting in 1982) into those that first appeared at least five years ago in the sample, and the rest. We re-estimate the size effect in the sample of firms at least five years of age.[30] This procedure has a clear drawback - firms older than five years that are only sampled once will be incorrectly classified as young. Subject to this caveat, we find that the size effect within the subsample of mature firms is approximately 80% in magnitude of what it is in the overall sample. Therefore, the size effect we document is not solely driven by firm age.

Our baseline results focus on the sample of surviving firms. This is primarily because the variables explaining non-response are not continuously available prior to 2000, so that we cannot confidently distinguish between true exits, mergers/acquisitions, and non-response prior to that date. We re-estimated the size effect in the sample of all firms-quarter observations *including* unanticipated non-responses, which account for approximately 3.5% of observations and using the Davis, Haltiwanger and Schuh (1996) bounded growth rates in order to include exiting firms. While the point estimates for the size effect is higher including exit, it is not statistically different from the estimate excluding imputed exit (but still using the bounded growth rates). This result is driven by the fact that in this data, the imputed exit rate among the bottom 99% group is not substantially more volatile at business cycle frequencies than among the top 1% group.

3.3 The excess sensitivity of small firms to monetary policy shocks

So far, we have presented evidence on the elasticity of firm sales to the US business cycle by firm size. One concern with these unconditional correlations is that they may mask important differences across firm size in the response to particular types of macroeconomic shocks. That is, some part of business cycle fluctuations may be driven by shocks that have a uniform effect across firm size, while other shocks exhibit stronger effects across firm size. In particular, Gertler and Gilchrist (1994) focus on the response of small and large firms after monetary policy shocks as identified in Romer and Romer (1989).

We start by examining the response of sales, inventories, and investment after the Romer and Romer dates in our sample. Figure 5 compares the cumulative change in sales of the top 1% and bottom 99% of firms, after a Romer and Romer (1989) date. Here, we define size using the top 1% vs. bottom 99% of firms in the one-year lagged asset distribution. We use the five Romer and Romer (1989) dates provided by the updated evidence in Kudlyak and Sanchez (2017): 1978q3, 1979q4, 1988q4, 1994q2 and 2008q3. As seen in Figure 5, we verify that sales, inventories, and

[30]There are a nontrivial number of observations for small firms which are sampled in distinct periods; that is, a firm is sampled for 8-12 quarters and appears several years later resampled again for 8-12 quarters.

investment contract faster after the Romer and Romer dates in our sample at small firms relative to large firms. These findings are qualitatively consistent with those of Gertler and Gilchrist (1994). Section 4.6 and Appendix D contains more details on the comparison between our results and the findings of Gertler and Gilchrist (1994).

To further gauge the effect of monetary policy shocks, we examine the response of sales by firm size groups to the monetary shock series constructed in Romer and Romer (2004) and updated by Wieland and Yang (2016). We construct the responses by firm size group using a projection method analogous to Jordà (2005). Our specification is:

$$\begin{aligned}\Delta y_{i,t,t+h} &= \sum_{j\in\mathcal{J}}\left(\alpha_j^{(h)}+\beta_j^{(h)}rr_{t-1,t}\right)\mathbf{1}_{i\in\mathcal{I}_t^{(j)}}+\sum_{l\in\mathcal{L}}\left(\gamma_l^{(h)}+\delta_l^{(h)}rr_{t-1,t}\right)\mathbf{1}_{i\in\mathcal{L}} \\ &+ \sum_{j\in\mathcal{J}}\sum_{q=1}^{4}\left(\mathbf{1}_{i\in\mathcal{I}_t^{(j)}}\times\mathbf{1}_{\mathbf{q}(t)=q}\right)s_{j,q}^{(h)}+\phi^{(h)}X_t+\epsilon_{i,t,h}\end{aligned} \tag{2}$$

y is the log of sales (or other variable of interest), i is firm, t is the quarterly date, h is horizon, $\mathcal{J}$ are size groups, $rr_{t-1,t}$ is the shock, $\mathcal{L}$ is industries, $\mathbf{q}(t)$ is the quarter (1 through 4) associated with date t, and X_t is a set of macroeconomic controls. We classify firms into two size groups, the (0,99) and the (99,100) groups. Our macro controls include unemployment, CPI, commodity prices, and the Fed funds rate, allowing for two lags. Our industry groups are the durable and non-durable sectors.[31] The primary coefficient of interest is $\beta_j^{(h)}$, which is the response of sales in size group j at horizon h to the monetary policy shock $rr_{t-1,t}$.

As discussed in Romer and Romer (2004), the monetary policy shock is measured using the deviation of the implemented Fed funds rate from internal forecasts prior to the meeting date. The updated time series time series is monthly, from 1969m1 to 2007m12. The sample stops thereafter because of the binding zero lower bound. We aggregate this time series to the quarterly frequency by taking the cumulative sum of the shock for each quarter, and using the end-of-quarter monthly value. We then use the quarterly time series from 1977q3 to 2007q4; our projection estimates thus exclude the response to monetary policy shocks that occured during or after the Great Recession. In response to a 1 percentage point innovation to the shock, similar projection methods using aggregate data indicate that the Federal Funds rate increases by 1.9 percentage points on impact, and mean-reverts back to zero within the first three quarters. The response of aggregate variables is strong and persistent: the trough in the response of industrial production is -1.1% (four quarters out) and the peak response of unemployment is a 0.35 percentage points (also four quarters out). The response of the CPI is slightly weaker, although it eventually declines by -0.5% two years out.[32]

[31] Results are qualitatively unchanged when using NAICS 3-digit sub-sectors instead.

[32] Results for Jorda projections using aggregate data are available from the authors upon request. Note that an alternative approach would be to use the series identified using high-frequency variation in Fed Funds futures around monetary policy announcement dates, as in Bernanke and Kuttner (2005), Gürkaynak, Sack and Swanson (2005) and Gertler and Karadi (2015). The time series for these shocks is only available from 1990m1 onwards, but does

Figure 6 shows the response of sales, inventories, and investment to the Romer and Romer shock series. The point estimates show that sales growth falls somewhat faster at small firms relative to large firms, consistent with our findings for the elasticity of firm sales growth with respect to the business cycle. However, the difference between sales growth at the top 1% and the bottom 99% is not statistically significant for most quarters. The evidence for a size effect is stronger for inventory growth, with small firm's inventory contracting while large firms inventory continues to expand after the shock. In this case, the difference between the small and large firms are statistically significant. Investment rates, like sales growth, are more sensitive at small firms, but the difference is again not statistically significant.

Overall, the effect of monetary policy shocks is qualitatively consistent with the findings in Gertler and Gilchrist (1994), but the differences across size groups are not statistically significant for sales or investment. To avoid attrition bias (since small firms are sampled for 8-12 quarters), we estimated the Jorda specification in firm-level data up to a horizon of only 8 quarters. To obtain a longer horizon, we also estimate the Jorda specification on average sales growth (inventory growth, investment rates) within firm-size classes; these projections amount to pooling firm-level data by size class before estimating the effect of monetary policy shocks. Our findings are essentially unchanged: there is lower sensitivity at the top, but it is not statistically significant.

This section has established that relative to large firms, small firms tend to exhibit a higher sensitivity of sales growth, inventory growth and investment to aggregate fluctuations. Quantitatively, a 1% point fall in GDP is associated with a 3.1% point drop in sales among the bottom 99% of firms, but only a 2.5% fall among the top 1%; the differences in elasticities are larger for inventory and fixed investment, and all are statistically significant. This excess sensitivity is also visible, although statistically insignificant, in response to identified monetary policy shocks. The remainder of the paper asks two questions: are these differences relevant for aggregate fluctuations, and are these differences driven by a financial accelerator mechanism?

cover the Great Recession period. The results from such an analysis, also available from the authors upon request, are qualitatively consistent with those obtained using the Romer-Romer shocks, in that point estimates indicate that small firms display excess sensitivity, but are not also statistically significant. However, one drawback of using these shocks is that, in a Jorda Projection framework, they lead to an expansionary response of aggregates, as pointed out by Ramey (2016). This is also true in our firm-level data, where innovations to the shock series are associated with overall increases in sales, inventories and, to a lesser extent, investment, abeit weaker at small than at large firms.

4 Aggregate implications

4.1 A simple decomposition of aggregate growth

Appendix B shows that the growth rate of any aggregate variable of interest between quarters $t-4$ and t, denoted by G_t, can be decomposed as:

$$\begin{aligned} G_t &= \hat{g}_t^{\text{(large)}} \\ &+ s_{t-4}\left(\hat{g}_t^{\text{(small)}} - \hat{g}_t^{\text{(large)}}\right) \\ &+ \hat{cov}_t. \end{aligned} \tag{3}$$

Here, $s_{t-4} = \frac{X_{t-4}^{(small)}}{X_{t-4}}$ is the initial fraction of x accounted for by small firms, and $\hat{g}_t^{(small)}$ and $\hat{g}_t^{(large)}$ are the cross-sectional average growth rates considered in the previous section.[33] The term $\hat{cov}_t$ is itself a weighted average of two terms:

$$\hat{cov}_t = \hat{cov}_t^{\text{(large)}} + s_{t-4}\left(\hat{cov}_t^{\text{(small)}} - \hat{cov}_t^{\text{(large)}}\right).$$

Each of the two terms $\hat{cov}_t^{\text{(small)}}$ and $\hat{cov}_t^{\text{(large)}}$ can be interpreted as cross-sectional covariances between firms' initial shares in their group, and their subsequent growth.[34] These terms capture the idea that if firms that are initially large in their group also grow faster, then total growth will tend to outpace firm-level growth in that group (and vice-versa if initially large firms grow more slowly). In principle, differences in the covariance terms between small and large firms may also by relevant for understand the contribution of small firms to fluctuations in aggregate growth. Note that this decomposition is only correct if the set of firms entering aggregate sales is held constant from t to $t-4$; thus, it should be thought of as a decomposition of growth of *surviving* firms and does not reflect any effect of entry or exit.

The decomposition (3) attributes aggregate growth in the variable of interest to three different sources: firm-level growth of large firms $\hat{g}_t^{\text{(large)}}$; differential firm-level growth between small and large firms $\Delta\hat{g}_t \equiv \hat{g}_t^{\text{(small)}} - \hat{g}_t^{\text{(large)}}$; and a term capturing the covariance between initial size and growth $\hat{cov}_t$. It clarifies the intuitive fact that, in order to matter, the growth differential $\Delta\hat{g}_t$ must be large relative to small firms' initial share s_{t-4}. Additionally, the decomposition indicates that business-cycle variation in $\hat{cov}_t$ could offset the effect of firm-level growth on aggregates. The

[33] This section analyzes a decomposition for the same log growth rates as discussed in the previous section, up to the approximation $log(1+x) \approx x$. Appendix C derives a similar decomposition for the commonly used growth rates $\tilde{g}_{i,t} = \frac{x_{i,t}-x_{i,t-4}}{\frac{1}{2}\left(x_{i,t-4}+x_{i,t}\right)}$, introduced by Davis, Haltiwanger and Schuh (1996). The appendix reproduces the same decomposition using these growth rates and show that all the results of this section are unchanged.

[34] Specifically, $\hat{cov}_t^{\text{(j)}} = \sum_{i\in\mathcal{I}_t^{\text{(j)}}} \left(w_{i,t-4} - \frac{1}{\#\mathcal{I}_t}\right)\left(g_{i,t} - \hat{g}_t^{(j)}\right)$, where j is small or large firms and where $w_{i,t-4}$ is the four-quarter lagged share of the total value of the variable of interest accounted for by firm i. This term is a cross-sectional covariance up to a normalizing factor. Appendix B contains more details on the decomposition and its interpretation.

relationship between firm-level and aggregate growth thus depends on the properties of s_{t-4} and $\hat{cov}_t$ in the data.

4.2 The covariance between initial size and growth

In order to clarify the contribution of the term $\hat{cov}_t$ to business-cycle variation in G_t, it is useful to note that the analogous decomposition to (3) also holds within each firm group, namely:

$$\begin{aligned} G_t^{(\text{small})} &= g_t^{(\text{small})} + \hat{cov}_t^{(\text{small})}, \\ G_t^{(\text{large})} &= g_t^{(\text{large})} + \hat{cov}_t^{(\text{large})}. \end{aligned} \tag{4}$$

Let Y_t be a business-cycle indicator; for instance, $Y_t \equiv \Delta GDP_t$. We can then write the correlation between $G_t^{(\text{small})}$ and Y_t as:

$$corr(G_t^{(\text{small})}, Y_t) = \frac{\sigma_{\hat{g}_t^{(\text{small})}}}{\sigma_{G_t^{(\text{small})}}} corr\left(\hat{g}_t^{(\text{small})}, Y_t\right) + \frac{\sigma_{\hat{cov}_t^{(\text{small})}}}{\sigma_{G_t^{(\text{small})}}} corr\left(\hat{cov}_t^{(\text{small})}, Y_t\right). \tag{5}$$

Here, σ_Z denote the standard deviation of variable Z. Equation (5) breaks down the correlation between G_t and Y_t into a component originating from firm-level growth, and a component originating from the covariance term. Of course, the same holds for large firms, and for firms overall.

Table 5 reports the values of the different elements of the right-hand side of (5), when the variable of interest is sales. It shows that the covariance terms - whether it be for small firms, large firms or all firms - have a limited (although non-zero) contribution to business-cycle variation in aggregate growth. Of course, these terms are non-zero on average; in fact, their sample means are 0.13, 0.29 and 0.23 for small, large and all firms, respectively. The large *average* difference in the covariance term between small and large firms has a substantial effect on trends. Namely, for small firms, cumulative average firm-level growth tracks fairly closely the path of aggregates; by contrast, for large firms, cumulative firm-level growth falls far short of the trend in aggregates, as documented in Figure 7, reflecting the rise in concentration.

But both the correlation to GDP growth of these covariance terms, and their standard deviation relative to aggregate sales growth G_t, are substantially smaller than for the cross-sectional average firm-level growth rates. For example, for large firms, the correlation between aggregate sales growth and GDP growth is 0.62 in the sample; this can be broken down into a contribution of $0.64 = 0.83 \times 0.77$, coming from the term $\frac{\sigma_{\hat{g}_t^{(\text{large})}}}{\sigma_{G_t^{(\text{large})}}} corr\left(\hat{g}_t^{(\text{large})}, Y_t\right)$, and $-0.02 = 0.45 \times (-0.05)$, coming from the term $\frac{\sigma_{\hat{cov}_t^{(\text{large})}}}{\sigma_{G_t^{(\text{large})}}} corr\left(\hat{cov}_t^{(\text{large})}, Y_t\right)$. This simple decomposition thus suggest that, up to first order, business-cycle variation in the covariance terms contribute little to aggregate growth; instead, average firm-level growth is the dominant factor.

4.3 The relative importance of small firms

Figure 8 reports the level (left column) and the share (right column) of total sales, inventory, fixed investment, and total assets of the bottom 99% of firms by size. The right column, in particular, corresponds to the time-series s_t defined above. As previously, size groups are defined relative to the one-year lagged distribution of assets. Two points about these time series are worth emphasizing.

First, the relative importance of the bottom 99% is, on average, small. Their average share of total sales, inventory, fixed investment, and total assets, are, respectively, 26.4%, 27.8%, 11.8% and 16.0% in this sample. The particularly low share for assets reflects the extreme degree of skewness of the firm size distribution; by contrast, the fact that the share of sales is higher is consistent with the fact that smaller firms are less capital-intensive. Nevertheless, this skewness presents a first hurdle for the excess sensitivity of small firms to substantially affect aggregates.

Second, movements in the average shares seem dominated by a long-term downward trend, not business-cycle variation. The share of sales of the bottom 99% falls from 35.6% in 1977q3 to 20.4% in 2014q1, while their share of assets falls from 25.6% to 9.0%; this decline is secular over the period with an acceleration around the 2000's. This is not to say that cyclical movements in small firms' shares are completely absent: for instance, the raw correlation $corr\left(s_{t-4}, \Delta GDP_t\right)$ is approximately 0.37 in the sample. While substantial cyclicality of the share could, in principle, offset its low average level and magnify the term $\Delta\hat{g}_t$, Figure 8 suggests that this unlikely to be the case in the data.

Overall, business cycle variation in aggregate growth is primarily driven by firm-level growth, not by the residual covariance term. Additionally, the share of small firms in total sales is low, so that differences in firm-level growth of small firms relative to large ones are unlikely to be reflected in aggregate fluctuations. We next quantify this statements more precisely, by constructing counterfactual paths for aggregate growth and analyzing their business cycle behavior.

4.4 Counterfactuals

We start by constructing the counterfactual time series:

$$G_t^{(1)} = G_t - s_{t-4}\left(\hat{g}_t^{(\text{small})} - \hat{g}_t^{(\text{large})}\right) \tag{6}$$

This time series nets out the contribution of firm-level growth differentials between the small and large firms — the second term of the decomposition (3). One could also net out differentials in the covariance terms; the counterfactual time series obtained would then simply be the aggregate growth rate among large firms:

$$\begin{aligned} G_t^{(2)} &= G_t - s_{t-4}\left(\hat{g}_t^{(\text{small})} - \hat{g}_t^{(\text{large})}\right) \\ &\quad - s_{t-4}\left(\hat{cov}_t^{(\text{small})} - \hat{cov}_t^{(\text{large})}\right) \\ &= G_t^{(large)}. \end{aligned} \tag{7}$$

We construct a third and final counterfactual, which assumes that small firms' growth has no cyclical component:

$$G_t^{(3)} \quad = G_t^{(large)} + s_{t-4}\left(G^{(small)} - G_t^{(large)}\right), \tag{8}$$

where $G^{(small)}$ is the sample average of $G_t^{(small)}$. This last counterfactual not only remove all *excess* cyclical sensitivity of small firms relative to large ones (as the first and second counterfactuals do); it removes the cyclicality of small firms altogether.

We are interested in the differences in cyclicality between these aggregate time series. As in the previous section, our simple metric for cyclicality are the estimates of the slope term in an OLS regression of G_t, $G_t^{(1)}$, $G_t^{(2)}$ and $G_t^{(3)}$ on the annual log-change in real GDP.

Table 6 reports the estimated slopes of the actual and counterfactual aggregate growth series for sales, inventory, fixed investment, and total assets. For sales (first line), the actual and counterfactual elasticities are close: the point estimates differ by approximately 13 basis points, and this difference is not statistically significant. The economic interpretation of this difference is that, all other things equal, if the elasticity of small firms' sales growth were equal to that of large firms, aggregate sales' elasticity to GDP growth would only be only 5% smaller. The second counterfactual series is even closer, indicating that cyclical variation in the difference between the covariance terms between small and large firms is, if anything, dampening aggregate fluctuations. The same conclusion holds for inventory; and it holds, in even stronger terms, for investment and for total assets. This is perhaps unsurprising given the high degree of concentration documented previously. Unsurprisingly — since it removes *all* cyclicality from small firms' aggregate growth — the third and last counterfactual results in lower estimates of elasticities to GDP growth. Nevertheless, the estimated elasticities remain substantial (all at least 70% of the actual elasticity), again highlighting the disproportionate contribution of large firms' growth to the aggregate.

Note that the results are consistent with a simple rule of thumb: the aggregate impact of small firms' excess sensitivity is equal to the product of the typical share of the small firms, multiplied by the excess sensitivity of small firms. For sales, for example, the results of the previous section indicate that the difference in elasticities to GDP growth between small and large firms (the excess sensitivity of small firms) is approximately 0.6. The results reported in Figure 8 indicate small firms' share is, on average, approximately 25%. The product of the two is: $0.6 \times 0.25 = 15$ bps, or approximately the difference between the estimated and counterfactual elasticities (13 bps). The fact that this rule of thumb delivers approximately the same result as the computation reported in Table 6 indicates that both cyclical movements in the covariance term and cyclical variation in small firms' share, have a limited impact on the cyclical fluctuations in aggregate growth. Figure 9 drives home this last point, by reporting the three time series G_t, $G_t^{(1)}$ and $G_t^{(2)}$ for sales. The three overlap and are visually indistinguishable.

4.5 Sales versus employment

While we have shown that the contribution of small firms to aggregate fluctuations in sales, inventories, and investment is quite small, we are unable to offer a similar calculation for employment given that firms do not report employment in this survey. However, we can estimate the employment threshold for large firms using data from the Census Bureau's Business Dynamics Statistics. There are roughly 1000 firms in the top 1% of our sample. In the BDS, the top 1000 firms in 2014 correspond to those firms with over 2500 employees. Likewise, given that firms are only sampled if their assets exceed $250K, we estimate that firms with less than 10 employees are not sampled. In 2014, firms with over 2500 employees account for 43% of manufacturing employment (only counting firms with at least 10 employees). Thus, the degree of skewness in employment is considerably less than that of sales, inventories, and investment. Moreover, the share of manufacturing employment at firms with 2500 employees has been falling over time (from about 55% in the early 1980s).

Thus, to the extent that small and large firms differ in their elasticity of employment growth to GDP, these differences are likely to be relevant for overall employment fluctuations in manufacturing. In Figure 10, we use BDS data to compute employment growth rates in manufacturing for all firms (with at least 10 emps) and for firms with more than 2500 employees. The two series move together but the degree of correlation is far weaker than shown in Figure 9. Moreover, the degree of skewness in employment in sectors outside manufacturing is somewhat lower, suggesting that differences in the business cycle sensitivity of small and large firms maybe considerably more important for employment.

4.6 Comparison to Gertler and Gilchrist (1994)

We conclude this section by comparing our results with those of Gertler and Gilchrist (1994), who use the public releases of the QFR between 1954 and 1990 to show that, around six dates identified by Romer and Romer (1989) as monetary policy tightenings, the sales and inventory of small firms tend to contract substantially more than those of large firms. The paper in particular indicates that, on average, small firms "account for between 55 and 60 percent of the drop in manufacturing sales" in the three year period subsequent to the monetary policy contraction.

In appendix D, we replicate their analysis on the portion of our data which overlaps with their original sample (1977q3 to 2014q1). There are two primary differences between their analysis and the one proposed in this section: the methodology used to compare small and large firms, and the period spanned by the analysis. The appendix shows while the methodological differences between analyses (specifically, our use of the cross-sectional distribution of book assets to define firm size, and our focus on equal-weighted growth rates) account for some of the difference in magnitude between our findings, most of it seems attributable to the period spanned by our analysis. Specifically, using their methodology on the three Romer dates for which our data overlap (1978q3, 1979q4 and 1988q4), in the 1977q3-1990q4 sample, we find that small firms account for 42% of the drop in manufacturing sales, a figure similar to, if slightly smaller than, that reported by Gertler and Gilchrist (1994).

However, in the full sample, using the Gertler and Gilchrist (1994) methodology on the three original Romer dates, plus the two additional dates identified by Kudlyak and Sanchez (2017) as further "Romer" dates (1994q2 and 2008q3), we find that the excess response of small firms is much more muted, and only accounts for 22% of the total drop in manufacturing sales. Moreover, this result is in line with the relative importance of small firms in total sales, and with the average magnitude of the excess sensitivity of small firms which we documented in the previous section. This divergence between earlier and later sample periods may reflect either changes in the conduct of monetary policy, changes in the underlying mechanism by which the monetary policy change affects small firms' sales, or changes in the relative importance of small firms (as discussed in this section). In any instance, the three earlier Romer dates in our sample, and in particular the Volcker recessions, seem to have more sharply affected small firms than subsequent episodes did, leading to the bulk of the discrepancies between our results and the results of Gertler and Gilchrist (1994).

This section has shown that the excess sensitivity of small firms has a limited impact on the behavior of aggregates in our sample. This is primarily because the relative importance of the bottom 99% is small and declining in the data, and secondarily because the difference between aggregate and firm-level growth - a residual capturing the covariance between initial size and growth - displays very little cyclical variation, so that firm-level and aggregate growth closely track each other.

5 The financial origins of excess sensitivity

As mentioned in the introduction, the early financial accelerator literature emphasized a variety of mechanisms whereby recessions, including ones not originating in the financial sector, could be worsened due to the presence of financial frictions. In this section, we investigate whether the size asymmetry we have documented should be interpreted as evidence of such financial amplification. We start by including various proxies for balance sheet strength in our size regressions; we find that the size effect remains significant and, in most cases, is quantitatively unchanged. However, it is possible that size is simply a better proxy for financial constraints. Therefore, we consider a benchmark model of firm investment with financing constraints where small firms, by construction, are financially constrained. In the simplest version of the model, an aggregate shock actually implies excess sensitivity of *large* firms relative to *small* firms. With a more general specification of the borrowing constraint, it is possible to generate excess sensitivity of sales and investment for small relative to large firms. However, in this case, small firms should exhibit more cyclical financing than large firms in recessions. As we show, our data does not bear this out.

5.1 The size effect and other proxies for financial constraints

We start by examining how estimates of the size effect vary when controlling for observable financial characteristics at the firm level. We start by estimating the following "horse-race" regressions:

$$\begin{aligned} g_{i,t} &= \sum_{j \in \mathcal{J}} (\alpha_j + \beta_j \Delta GDP_t) \, \mathbf{1}_{i \in \mathcal{I}_t^{(j)}} + \sum_{l \in \mathcal{L}} (\gamma_l + \delta_l \Delta GDP_t) \, \mathbf{1}_{i \in \mathcal{L}} \\ &+ \sum_{k \in \mathcal{K}} (\zeta_k + \eta_k \Delta GDP_t) \, \mathbf{1}_{\left\{ i \in \mathcal{F}_t^{(k)} \right\}} + \epsilon_{i,t}. \end{aligned} \tag{9}$$

In these regressions, the size controls are identical to the baseline estimation of section 3: size groups, indexed by j, are defined using lagged firm size, and results for 90-99th percentile, 99th to 99.5th percentile, and top 0.5% are reported relative to the baseline 0-90% group. As before, we also include indicators for durable and non-durable manufacturing.[35] In contrast to the baseline regression, $k \in \mathcal{K}$ now indexes groups of our measures of financial strength. We consider five different measures of financial strength: bank-dependence, leverage, liquidity, access to public debt markets, and dividend issuance.

Column (1) in Table 7 controls for the degree of bank-dependence in the size regression. Our measure of bank dependence is the share of bank debt in total debt. This variable has a bimodal distribution, with some firms nearly fully reliant on bank debt and some firms (including zero leverage firms) have no reliance on bank debt. We sort firm into low bank dependence firms (with a bank share of less than 10%), intermediate bank dependence firms (between 10% and 90%), and high bank dependence firms (over 90%).

Column (2) in Table 7 controls for leverage. We split the sample into four bins: firms with zero debt, firms with a debt to asset ratio of less than 15%, firms with a debt to asset ratio of between 15% and 50%, and firms with debt to asset ratio over 50%. Firms with leverage less than 15% approximately account for the bottom quarter of the leverage distribution, while firms above 50% account for approximately the top quarter.[36]

Column (3) controls for liquidity. We consider three liquidity classes: cash to asset ratio of less than 1%; cash to asset ratio between 1% and 20%; cash to asset ratio above 20%. As with leverage, we choose fixed thresholds that approximate the bottom and top quartiles.[37]

Column (4) controls for access to public debt markets. Specifically, we classify a firm-quarter observation as having access to public debt markets if the same firm has ever reported some positive liability in either commercial paper or long-term bonds. Because it relies only on responses from

[35] Our results hold when controlling for NAICS 3-digit industries.

[36] We use fixed thresholds given the absence of a time trend in leverage.

[37] The cash to asset ratio for the median firm in the QFR dataset rises starting around 2005. The top quartile of the cash to asset distribution, however, is fairly stable over time, rising only slightly toward the end of the sample.

the long-form survey, this variable is most informative for the largest firms (it is equal to zero for firms receiving the short-firm survey). As documented by Faulkender and Petersen (2005), even among publicly traded firms, only a minority have access to public debt market, so that there is meaningful variation in this measure among large firms.

Finally, column (5) controls for dividend issuance. A firm-quarter observation is classified as a dividend issuer if it issued dividends in the year prior to the quarter of observation. About half of firm-quarter observations in the regression sample are dividend issuers.

For bank-dependence, leverage, liquidity, and dividend issuance, the coefficients on GDP interacted with size class — particularly the top 0.5% — remain significant, and in magnitude, similar to the baseline regression. Thus, none of these controls changes the estimates of the size effect. The exception is market access, but the change in the size coefficient is inconsistent with the financial accelerator view. One would expect firms with market access to have a lower degree of sensitivity to the business cycle, and therefore the size effect to fall in magnitude once one controls for market access. Instead, we find that it rises, suggesting that firms with access to public debt markets are, if anything, more cyclically sensitive than other large firms. This result appears again in section 5.4; it may be due to firms with more cyclical investment opportunities choosing to tap bond markets at the beginning of recoveries.

In any case, the main message of Table 7 is that the excess sensitivity of small firms survives, and is in fact almost unchanged (or even amplified) after controlling for the five simple proxies for financial constraints.

We next turn to triple-interaction regressions. These regressions are meant to answer the following question: is the size effect weaker among groups of financially stronger firms? In order to measure financial strength, we use the same five ratios as in the previous horse-race regressions. We estimate a regression of the same form as (9), but where observations are effectively double sorted by their position in the firm size distribution and bins of a measure of financial strength. As in previous regressions, we also include industry fixed effects, and interactions of industry effects and GDP growth.

Results are reported in Table (8). In this table, all estimates of the size effect are expressed relative to the bottom $[0, 90]$ group.[38] The first column is the baseline regression without triple interaction - the same regression as in Table (4). The coefficient -0.60, for instance, indicating that the sales elasticity to GDP of firms in the $[99, 99.5]$ group is 0.6 points lower than that of firms in the $[0, 90]$ group.

The second and third columns report similar elasticities when size and bank dependence categories are interacted. The estimates are organized by bank dependence groups; in order to keep the table readable, we have kept only two groups for bank dependence. Firm-year observations

[38] This is with the exception of regressions conditioning on bond market access where results are reported relative to the $[0, 99]$ group as there are too few observations with bond market access in the $[0, 90]$ group.

in the low bank-dependence group had a ratio of bank debt to total debt below 0.9 in the prior year, whereas firms in the high bank-dependence group had a ratio of bank debt to total debt over 0.9.[39] The reported coefficients denote relative elasticities within each bank dependence group. The estimates suggest that among banks with low to moderate bank dependence, the size estimate has the same sign, and a similar magnitude as in the unconditional regressions. Among highly bank-dependent, the size effect is slightly *smaller*, although the high minus low difference (reported in the right column) is not statistically significant. Had the size effect been a reflection of financial constraints, one might have expected it to be much weaker among firms with access to other sources of financing than bank debt; instead, it is somewhat stronger.

The following columns repeat this exercise for other proxies for financial constraints.[40] While results differ across measures of financial constraints, it is worth noting that, with the exception of the last indicator — firms' dividend issuance behavior — measures of the size effect are never statistically different across groups of financial strength proxies. Directionally, the estimates of the relative size effect for leverage and dividend issuance groups are consistent with the view that the size effect is weaker among financially stronger firms; on the other hand, estimates using liquidity and bond market access are not. Overall, the lack of significance in the cross-group differences in the size effect paired with its significance within group bolster the view that the size effect may not be financially driven.

5.2 The behavior of debt

Our finding that alternative proxies of balance sheet strength and financial constraints do not eliminate the size effect may be driven by the fact that size is a superior proxy for financial constraints than leverage, liquidity, dividend issuance, access to public debt markets or the bank share. In order to obtain additional testable predictions of the idea that the size effect proxies for financial constraints, we build a simple model where firm size is, by construction, a perfect indicator of financial constraints. A key prediction of the model is that excess sensitivity of investment, if it is driven by financial constraints, should also transalte into excess sensitivity of debt issuances among small firms. We then construct the real and financial responses to an aggregate shock of firms of different sizes, and compare them to the data.

A baseline model The model is set in discrete time. Firms maximize the present discounted value of future payouts to equityholders, and use the constant discount rate $\frac{1}{1+r}$. The problem of

[39]In order to avoid creating non-overlapping groups, which would complicate disclosure of results, we are limited to using a grouping by financial strength indicators that is a coarser version of the grouping of Table (7).

[40]For leverage, we split the sample above and below 0.5. For liquidity, we use a 0.01 cash to asset ratio as the threshold between low and high liquidity. These choices correspond approximately to the top quartile of the distribution of leverage and the bottom quartile of the distribution of the cash to asset ratio. Again, for disclosure reasons, these groupings are coarser versions of the sorts used in Table (7).

a surviving firm, indexed by i, in period t, is:

$$\mathbf{V}_t(k_{i,t}) = \max_{k_{i,t+1}} \quad \eta n_{i,t} + (1-\eta)\left(n_{i,t} - k_{i,t+1} + \frac{1}{1+r}\mathbf{V}_{t+1}(k_{i,t+1})\right)$$

$$\text{s.t.} \quad \begin{aligned} n_{i,t} &= z_t k_{i,t}^{\zeta} + (1-\delta)k_{i,t} \\ [\lambda_{i,t}] \quad 0 &\leq n_{i,t} - k_{i,t+1} \end{aligned}$$

Here, $k_{i,t}$ are the firm's assets in place. The firm's operating profits are given by $\pi_{i,t} = z_t k_{i,t}^{\zeta}$, with $0 < \zeta < 1$ denoting the curvature of the profit function with respect to assets, and z_t is an aggregate shock, which may capture aggregate changes in productivity, demand, or the cost of inputs.[41] Finally, $n_{i,t}$ is the firm's net worth, which is equal to the sum of its operating profits and the depreciated value of its capital stock.

There are two financial frictions in this environment. The first is that payouts to equityholders must be positive, that is, $n_{i,t} \geq k_{i,t}$. The frictionless model is one where, by contrast, payouts to equityholders can take any sign, without affecting their marginal benefit (or cost): $n_{i,t} \gtreqless k_{i,t}$. The second is that firms are not allowed to borrow. Firms are therefore completely internally financed. Note that another way to express the financial constraint is that $\pi_{i,t} \geq i_{i,t} = k_{i,t+1} - (1-\delta)k_{i,t}$, so that operating profits must fully cover investment in each period. The shadow value of internal funds is $\nu_{i,t} = 1 + \lambda_{i,t}$; a firm is constrained, if and only if, $\nu_{i,t} > 1$. The stark assumption of pure internal financing is a useful benchmark, which we relax below.

Finally, with probability η, a surviving firm exogenously exits at the beginning of the period. In this case, equityholders receive the firm's net worth as a payout. In order to focus the analysis on intensive margin responses, we assume that replacement of each exiting firm occurs at a exogenously determined level of assets, k_e.

In its stationary equilibrium ($z_t = z$ for all t), the frictionless model has the simple solution:

$$k_{i,t+1} = k^* \equiv \left(\frac{\zeta z}{r+\delta}\right)^{\frac{1}{1-\zeta}}, \quad \forall i, t. \tag{10}$$

At this value for $k_{i,t+1}$, the expected discounted marginal product of capital is equal to 1. In the frictionless model, all surviving firms have the same size. By contrast, in its stationary equilibrium, the solution to the model with frictions is:

$$k_{i,t+1} = \begin{cases} n_{i,t} & \text{if } n_{i,t} < k^* \\ k^* & \text{if } n_{i,t} \geq k^* \end{cases}. \tag{11}$$

So long as $n_e = z k_e^{\zeta} + (1-\delta)k_e < k^*$, the stationary equilibrium also features a cross-section of firms of different sizes: firms are born small relative to their desired size k^*, must save to reach it, and may fail to reach their optimal size due to the exogenous exit shock.

[41]The curvature in the profit function may originate either in decreasing returns in production, or in monopoly power. Depending on which specific microfoundation for the profit function is chosen, z_t will be given by a specific combination of aggregate productivity, the real wage rate, and aggregate demand for the industry's product.

The effects of an aggregate shock We consider the perfect foresight response of the model to a shock to z_t. Specifically, we assume that at time $t = -1$, $z_t = z$, and that the model is in its stationary equilibrium. Moreover, at time 0, firms learn that the future path of z_t, for $t \geq 0$, will be $z_t = z \exp\left(-\rho^t \epsilon\right)$, where $\epsilon > 0$ is a shock to productivity, and ρ is the persistence of the shock. This exercise is meant to approximate the response of the economy to a mean-reverting decline in productivity.

What are the cross-sectional implications of this shock? The top panel of figure 11 shows the perfect foresight response of output to a temporary decline in z_t, starting from the steady-state described by (11).[42] In the model with frictions, the most responsive firms are the largest ones — there is excess sensitivity, but it comes with the opposite sign as in the data.

Why are large firms more sensitive? The aggregate shock has two effects: it lowers all firms' net worth $n_{i,t} = z_t k_{i,t}^{\zeta} + (1-\delta)k_{i,t}$; but it also reduces the optimal unconstrained size of firms,

$$k_{t+1}^* = \left(\frac{\zeta z_{t+1}}{r+\delta}\right)^{\frac{1}{1-\zeta}}.$$

When the shock hits the economy, initially unconstrained firms (those with $n_{i,0} \geq k^*$) find themselves with financial slack: even though their net worth falls, it still remains above the new unconstrained threshold, $\underline{n}_1 = k_1^*$. As a result, these firms respond by paying out excess cash, and shrinking to $k_{i,1} = k_1^*$. By contrast, most constrained firms start from a point where $n_{i,0} < \underline{n}_1 = k_1^*$. That is, these firms are below their optimal size *even after* the aggregate shock. These firms' responses then only reflect changes in net worth. Because net worth is a linear function of the aggregate shock, whereas the optimal size is a convex function of the aggregate shock, the optimal size response tends to be larger than the net worth response.[43] Financial frictions, in this case, work like an adjustment cost, and moderates the response of quantities.

Adding pro-cyclical external financing The previous example shows that restricted access to external finance alone is not sufficient to generate a size effect. We next add debt financing to the model, and allow the borrowing constraint to be a function of both the firm's net worth and,

[42] The calibration of the model is described in Appendix F; in particular, the choice of the exogenous exit rate and the entry size imply that in steady-state, 1% of firms are unconstrained. The path of the shock is $z_t = z\exp(-\rho^t\epsilon)z$; in all figures, we use $\rho = 0.8$ and $\epsilon = 0.01$.

[43] The appendix shows that a necessary condition for the response of net worth to be smaller than the response of the optimal investment target is that $\frac{\rho}{1-\zeta} \geq \frac{r+d}{\zeta}$. This condition is met in our calibration, under reasonable values of the curvature parameter ζ, it will be met so long as the aggregate shock is not too transitory. It is clear that a purely transitory shock ($\rho = 0$) would only have a net worth effect, and hence only cause constrained firms to respond.

crucially, of the aggregate shock. The firm's objective is now:

$$
\begin{aligned}
\mathbf{V}_t(k_{i,t}, b_{i,t}) \quad &= \quad \max_{k_{i,t+1}, b_{i,t+1}} \quad \eta n_{i,t} + (1-\eta)\left(n_{i,t} - k_{i,t+1} + b_{i,t+1} + \frac{1}{1+r}\mathbf{V}_{t+1}(k_{i,t+1}, b_{i,t+1})\right) \\
& \qquad n_{i,t} \quad = \quad z_t k_{i,t}^{\zeta} + (1-\delta)k_{i,t} - (1+r)b_{i,t} \\
\text{s.t.} \quad & \qquad b_{i,t+1} \quad \leq \quad \mathbf{b}(n_{i,t}; z_t) \\
& \qquad n_{i,t} + b_{i,t+1} \quad \geq \quad k_{i,t+1}
\end{aligned}
$$

where $\mathbf{b}(\cdot,\cdot)$ — the borrowing constraint — is a function of both the firm's net worth and the aggregate shock z_t. As before, firms cannot raise equity (i.e. issue negative dividends).[44]

The solution to the firm's problem is similar to the case with no borrowing; the details are specified in Appendix (F). Firms with high levels of net worth invest at the optimal level k^*_{t+1}, while firms with insufficient net worth are either partially or fully constrained. Partially or fully constrained firms do not issue any dividends. Fully constrained firms utilize all their borrowing capacity; that is, $k_{i,t+1} = n_{i,t} + \mathbf{b}(n_{i,t}, x_{t+1})$. Partially constrained firms invest at the currently optimal level, but pay zero dividends. There need not be partially constrained firms in equilibrium; the situation only occurs when fundamentals are such that firms may be constrained tomorrow, for example if z_t is rising sharply over time.[45]

As before, we construct the response to a one-time unanticipated and mean-reverting decline in z_t, and compare the responses of small and large firms. The bottom panel of 11 displays the sales, investment, dividend issuance and debt financing response of small and large firms. These responses are constructed under the assumption that the borrowing constraint is "sufficiently" elastic with respect to the aggregate shock in order to generate excess sensitivity of investment among small firms.[46] Under this assumption, small firms will cut back on investment faster, and subsequently experience larger declines in sales, than large firms. It is straightforward to understand why a highly procyclical borrowing constraint is necessary to generate excess sensitivity of small firms in the model. Constrained firms' investment is given by their total financing capacity:

$$k_{i,t+1} = n_{i,t} + \mathbf{b}(n_{i,t}, z_t),$$

while unconstrained firms' investment is simply the optimal path $k^*_{t+1} = \left(\frac{\zeta z_{t+1}}{r+\delta}\right)^{\frac{1}{1-\zeta}}$. The latter is a convex function of the aggregate shock; intuitively, so long as the borrowing function is chosen so that the total borrowing capacity $n_{i,t} + \mathbf{b}(n_{i,t}, z_t)$ is a "more" convex function of the aggregate shock, the investment response of small/constrained firms will be larger.

[44] Additionally, we restrict attention to solutions which satisfy the following transversality condition: $\lim_{t\to+\infty} (1+r)^{-t}\mathbf{V}_t(k_{i,t+1}, b_{i,t+1}) \leq 0$.

[45] The appendix provides detailed conditions under which the partially constrained regime exists. It is worth noting that it never exists in steady-state.

[46] The appendix derives a simple sufficient condition on the elasticity of the borrowing constraint with respect to the aggregate shock that ensures the model generates excess sensitivity for investment.

However, a byproduct of the assumption that the borrowing constraint displays substantial procyclicality is that debt financing flows among small firms should also respond strongly to the aggregate shock. The bottom panel of figure 11 reports the cumulative change in debt among small and large firms. The contraction in debt among small firms is deeper and more protracted than among large firms. This is the financial flipside of the excess sensitivity of investment which the model generates. The model thus suggests that if small firms display excess sensitivity in investment because of financial constraints, then, we should also expect to find excess sensitivity in debt flows.

The behavior of debt during recessions In order to compare model and data, we compute the cumulative change in variables of interest in a 15-quarter window around the beginning of a recession. Let $g_{i,t}$ denote one of the outcome variables of interest (year-on-year sales growth, inventory growth); we estimate the model:

$$g_{i,t} = \alpha + \beta \mathbf{1}_{\left\{i \in \mathcal{I}_t^{(0,99)}\right\}} + \sum_{k=-4}^{10} \left(\alpha_k + \beta_k \mathbf{1}_{\left\{i \in \mathcal{I}_t^{(0,99)}\right\}} \right) \mathbf{1}_{\{t+k \in \mathcal{H}\}} + \epsilon_{i,t} \tag{12}$$

where $i \in \mathcal{I}_t^{(0,99)}$ is the set of small firms, defined as the bottom 99% of the lagged distribution of book assets, and $\mathcal{H}$ is one of four recession start dates: $\mathcal{H} = \{1981q3, 1990q3, 2001q1, 2007q4\}$. We then construct cumulative responses by size: $\{c_{\mathcal{L},k}\}_{k=-4}^{10}$ and $\{c_{\mathcal{S},k}\}_{k=-4}^{10}$ for large and small firms respectively:

$$c_{\mathcal{L},k} = \sum_{j=-4}^{k} (\alpha + \alpha_j) - \sum_{j=-4}^{0} (\alpha + \alpha_j),$$
$$c_{\mathcal{S},k} = c_{\mathcal{L},k} + \sum_{j=-4}^{k} (\beta + \beta_j) - \sum_{j=-4}^{0} (\beta + \beta_j),$$

as well as the associated standard errors. Note, in particular, that in order to avoid overlapping event windows, we only consider the second of the two recession start dates of the early 1980s.

Figure 12 reports the cumulative path of sales, inventory and fixed capital and the associated +/- 2 standard error bands. The behavior of sales is qualitatively consistent with the baseline regression: the cumulative drop in sales following the onset of the recession is substantially larger for the bottom 99% of firms, and the difference is statistically significant. The behavior of inventory investment and fixed investment is also qualitatively consistent with the baseline regressions; however, the differences are not statistically different across size groups, except for the cumulative decline in large firms' inventory at long lags. Perhaps the most striking qualitative feature of investment behavior is that the decline of investment among large firms seem to lag that of small firms by three to four quarters.[47] This lag is not visible in the sales response.[48]

Figure 13 repeats this exercise for cumulative changes in total debt, bank debt and short-term debt. Here, short-term debt is measured as debt with maturity one year or less normalized by assets lagged four quarter, and bank debt is short and long-term bank loans normalized by assets lagged four quarters.

In contrast with the predictions of the model, we find little difference in the behavior of debt financing at small and large firms. The left panel of Figure 13 shows that it is difficult to observe sharp differences in the behavior of debt overall. Given that the behavior of overall debt may mask significant movements in important components of debt, we also display the response of bank debt and short-term debt. The cumulative decline in bank and short-term debt is initially more pronounced among large firms, though not statistically different; eventually, the reduction in debt actually becomes bigger among large firms. The response of short-term debt among small is particularly, and strikingly, difficult to separate from that of large firms.

The behavior of debt in response to monetary policy shocks Arguably, monetary policy shocks may more directly impact that the cost of external borrowing and induce procyclical fluctuations in borrowing cost.[49] In turn, these episodes may provide a better example of the exogenous movements in borrowing capacity of firms which we analyze in the model. We thus next extend our analysis of Section 3, where we examined the effect of monetary policy shocks on firm borrowing by firm size categories using Jorda projection methods, to three additional dependent variables: the ratio of total debt to assets, the ratio of bank debt to assets and the ratio of short-term debt to assets. In effect, we are tracing out the impulse response of firm borrowing to an identified monetary policy shock.

Figure 14 shows the cumulative change in each of these debt ratios after a exogenous tightening in monetary policy. In the case of total debt and bank debt, the point estimates show that financing to small firms falls somewhat more than financing to large firms at most horizons, but the difference

[47] Aggregate fixed capital formation, in the QFR data, lags real GDP growth by three to four quarters as well: the contemporaneous correlation with year-on-year real GDP growth is 0.19, while the three-quarter lagged correlation is 0.59. This is consistent with the recession behavior documented in Figure 12, since, as discussed below, large firms account for between 80-90% of total fixed capital formation in the QFR data.

[48] Also in contrast to the sales response, the lack of statistical significance suggests that the excess sensitivity documented in the baseline regressions is driven by recoveries, rather than recessions. This is partly visible in Figure 12: the relative response of small firms' inventory at 10 and more quarters out is statistically different at that stage, when recoveries are already under way. In undisclosed results, we verify that restricting the sample to the onset of recessions indeed leads to insignificant estimates of excess sensitivity for inventory and fixed capital investment.

[49] The financial accelerator mechanism works through balance sheet effects where a fall in the price of capital goods reduces firm net worth and raises borrowing costs. Various credit channels of monetary policy (the bank lending channel, bank credit channel, and net worth channel) each emphasize how changes in monetary policy transmit to firms by raising the cost of borrowing.

between small and large firms is not significant. In the case of short-term debt, the response of large firms exceeds that of small firms at all horizons though, again, the difference is insignificant. In comparison to the evidence at recession dates, the excess sensitivity of debt is even harder to discern, bolstering our conclusion that the size effect does not reflect the effect of financial frictions. As in Section 3, we can also estimate impulse response over longer time horizons by taking average debt growth by firm size classes and then applying the Jorda method. We can also use an alternative series of monetary policy shocks from Gertler and Karadi (2015). In both cases, the point estimates are either inconsistent with the view that small firms are subject to tighter borrowing constraints after monetary policy shocks, or the differences between small and large firms' responses are statistically insignificant.

Overall, while it is clearly stylized, the model captures the basic mechanisms whereby financing constraints can amplify the business cycle. A sufficiently procyclical supply of external financing is needed to cause financially constrained firms to reduce investment relative to unconstrained firms. This should also manifest itself as procyclical debt financing flows, more so at small and constrained firms. However, the observed behavior of debt financing, particularly the behavior of short-term debt, does not appear to be more procyclical at small than at large firms.

5.3 Alternative interpretations of the size effect

We next investigate alternative explanations for the size effect, by exploiting variation in the size effect across 3-digit manufacturing industries. The size effect survives within 3-digit industries, but displays substantial heterogeneity and is attenuated or reversed in some of the smaller subindustries. First, we find no correlation between the size effect at the three digit level and a measure of external financial dependence based on Kaplan and Zingales (1997).[50] The absence of any correlation strengthens our view that the size effect is not drive by financial frictions.

We explore two alternative, nonfinancial hypotheses for the size effect: international exposure and downstream diversification. Using BEA input-output tables, we construct a measure of export intensity - total exports divided by gross output.[51] We find a positive correlation between the size effect and export intensity; industries with a high export share exhibit a stronger differences in the cyclicality between large and small firms. This correlation is consistent with a diversification hypothesis: large firms (which are more likely to be firms that export) are less exposed to the US business cycle and, to the extent that business cycles across countries are imperfectly correlated, these firms are relatively insulated from fluctuations in US GDP. The left panel of Figure 15 plots the correlation between the size effect and export exposure.

[50] Specifically, we use the values for NAICS industries in Appendix Table A2 in Duygan-Bump, Levkov and Montoriol-Garriga (2015).

[51] We compute this measure every five years from 2000 to 2015 and take averages over this period.

We can test the diversification hypothesis in another way. We examine the correlation of a measure of downstream diversification with the size effect at the 3-digit level. Our measure of downstream diversification is a Herfindahl index using the share of industry X's gross output used by industry Y. A high value of the Herfindahl indicates low diversification - industries that supply relatively little to other industries as inputs. We find a modest negative correlation between the Herfindahl index and the size effect; those industries that are more diversified exhibit a greater difference between large and small firms. The right panel of 15 plots the correlation between our diversification measure and the size effect across 3-digit NAICS.

The correlation is further strengthened if NAICS 336 (motor vehicles), which is an outlier in concentration, is dropped. This correlation is consistent with the hypothesis that large firms in industries with a well-diversified customer base are able to diversify across industries magnifying the difference between large and small firms. The implicit assumption is that the number of customers expands with firm size. While the evidence is clearly limited, our findings for 3-digit industries suggests that diversification could explain part of the size effect and represents a mechanism distinct from financial frictions.

5.4 Non-size evidence of a financial accelerator

The regression results we have discussed so far suggest that the size effect may not be evidence in favor of the financial accelerator mechanism. We conclude this section by documenting whether firms respond heterogeneously to recessions when conditioning directly on balance sheet characteristics, instead of size. Specifically, we provide event study plots comparing the evolution firm sales, inventories, and tangible investment around recessions, separating firms in groups of leverage, liquidity, bank-dependence, access to bond markets, and dividend issuance.

Figure 16 depicts the evolution of firms sales, inventories, and fixed capital comparing zero leverage firms (which account for roughly 20% of firm-quarter observations), and firms with positive leverage; we classify firms based on their four-quarter lagged debt to asset ratio. As the plots show, the evolution of sales and investment at the two groups of firms is largely indistinguishable during recessions. The same holds true for liquidity: when sorting firms into low liquidity (firms with a cash to asset ratio of less than 0.2) and high liquidity (firms with a cash to asset ratio of greater than 0.2), we also find largely indistinguishable cumulative responses of sales, inventories, and investment.

The last row of Figure 16 sorts firms into bank-dependent and non-bank-dependent. The former are defined as firms with more than 90% of debt in the form of bank loans four quarters past. While bank dependent firms do qualitatively experience a sharper contraction in their sales and investment than non-bank dependent firms, the differences are, again, not statistically significant. Results based on leverage sorts would appear to be inconsistent with a financial accelerator mechanism.

Under the financial accelerator mechanism, higher leverage firms should experience increases in the cost of external financing during recessions, leading to a faster decline in factor inputs and production relative to firms that do not rely on external financing. By contrast, the evidence provided above suggests that there is no sharp difference in the behavior of higher-leverage firms during recessions.

Figure 17 provides the event study plots for firms sorted on public debt market access (top row) and dividend issuance (bottom row). Firms with a history of accessing public debt markets contract their sales and inventories *faster* than firms with no history of market access. The financial accelerator mechanism would predict the opposite, as firms with access to bond markets should better be able to smooth sales and inventories over the business cycle. Moreover, the point estimates suggest that investment falls faster at firms without market access, but that the difference is not statistically significant. By contrast, firms sorted on dividend issuance do display statistically significant differences for inventory and investments in recessions: firms that issued dividends during the prior year also reduce inventories and investment more gradually than firms that didn't.

Overall, these findings provide only weak evidence for the presence of a financial accelerator. Standard measures of balance sheet strength do not predict excess sensitivity for financially weaker firms in recessions; market access seems to be associated with a magnified sensitivity to recessions. Dividend issuance appears somewhat more promising, with non-dividend issuing firms cutting inputs faster in recessions than dividend issuing firms. Further research is needed to determine to what extend dividend issuance is a good proxy for financial constraints, as opposed to future investment opportunities.

6 Conclusion

This paper has brought new evidence to bear on the question of whether, and why, cross-sectional differences in exposure to business cycles might be related to firm size. This evidence, though limited to the manufacturing sector, has the advantage of covering a representative sample of the population of US firms, at the quarterly frequency, over a period spanning the last 5 recessions. Moreover, this evidence allows one to directly link real decisions of firms to their financial strength, which the literature on firm dynamics and business cycles has argued is a key determinant of heterogeneous responses to aggregate conditions.

We find strong evidence that smaller firms tend to be more sensitive to aggregate conditions than large firms, consistent with previous literature. Our point estimate suggests that a 1% drop in GDP is associated with a 2.5% contraction in sales for firms in the top 1% of the size distribution, but a 3.1% contraction for firms in the bottom 99%.

Our evidence however casts doubt on the commonly accepted interpretations of this finding. First, we show that the effect is at least as much about expansions as it is about recessions, and furthermore, that it is mostly accounted for by the top 0.5%, with the rest of firms in the distribution having statistically indistinguishable sensitivities. Second, the degree of concentration of sales and investment is dramatic; by the latter parts of the sample, for instance, the top 0.5% of firms account for about 75% of sales. As a result, the excess sensitivity of smaller firms is insufficient to substantially affect the volatility of aggregates; we estimate that, absent excess sensitivity, the elasticity of aggregate sales to GDP growth in our sample would only be about 0.15 points smaller, from a baseline of 2.30.

Finally, we provide evidence that this excess sensitivity cannot easily be accounted for by financial factors: the behavior of debt (in particular short-term debt) during recessions does not significantly differ among small firms; and moreover, controlling for proxies for financial constraints does not eliminate our estimated size effect. More generally, firm groups conditioning directly on these proxies does not exhibit a substantial difference in cyclical sensitivity either.

These results suggest two potential directions to further test the hypothesis that the excess sensitivity of small firms is financial in nature. First, while it is notoriously difficult to measure financial constraints, the broader question of whether small firms are more financially constrained could be explored in more detail using this data; differential exposure in the timing of either tax or banking reforms is a potential avenue of research. Second, the results on sorts of firms by groups of financial strength (leverage, liquidity, dividend issuance) reported here could be interpreted from the standpoint of a more detail structural model with heterogeneous firms and financial frictions than the one provided in this paper. We leave these issues — and the broader question of the extent to which the financial accelerator contributes to amplifying business cycles — to future research.

References

Almeida, Heitor, Murillo Campello, and Michael S. Weisbach. 2004. "The Cash Flow Sensitivity of Cash." *Journal of Finance*, 59(4): 1777–1804.

Asker, John, Joan Farre-Mensa, and Alexander Ljungqvist. 2011. "Comparing the Investment Behavior of Public and Private Firms." National Bureau of Economic Research.

Begenau, Juliane, and Juliana Salomao. 2015. "Firm Financing over the Business Cycle." Mimeo. University of Minnesota.

Bergman, Nittai K., Rajkamal Iyer, and Richard T. Thakor. 2015. "Financial Accelerator at Work: Evidence from Corn Fields." National Bureau of Economic Research.

Bernanke, Ben S., and Alan S. Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission." *American Economic Review*, 82(4): 901–921.

Bernanke, Ben S., and Kenneth N. Kuttner. 2005. "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance*, 60(3): 1221–1257.

Bernanke, Ben S., Cara S. Lown, and Benjamin M. Friedman. 1991. "The Credit Crunch." *Brookings Papers on Economic Activity*, 21(2): 205–247.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist. 1996. "The Flight to Quality and the Financial Accelerator." *Review of Economics and Statistics*, 78(1): 1–15.

Chari, V.V., Lawrence Christiano, and Patrick Kehoe. 2013. "The Behavior of Large and Small Firms Over the Business Cycle." Mimeo. University of Minnesota.

Chodorow-Reich, Gabriel. 2014. "The Employment Effects of Credit Market Disruptions: Firm-Level Evidence from the 2008-2009 Financial Crisis." *Quarterly Journal of Economics*, 129(1): 1–59.

Covas, Francisco, and Wouter J. Den Haan. 2011. "The Cyclical Behavior of Debt and Equity Finance." *American Economic Review*, 101(2): 877–899.

Crouzet, Nicolas. 2017. "Aggregate Implications of Corporate Debt Choices." *Forthcoming, Review of Economic Studies.*

Davis, Steven J., John Haltiwanger, and Scott Schuh. 1996. "Small Business and Job Creation: Dissecting the myth and Reassessing the Facts." *Small Business Economics*, 8(4): 297–315.

Duygan-Bump, Burcu, Alexey Levkov, and Judit Montoriol-Garriga. 2015. "Financing Constraints and Unemployment: Evidence from the Great Recession." *Journal of Monetary Economics*, 75: 89–105.

Farre-Mensa, Joan, and Alexander Ljungqvist. 2016. "Do Measures of Financial Constraints Measure Financial Constraints?" *Review of Financial Studies*, 29(2): 271–308.

Faulkender, Michael, and Mitchell A. Petersen. 2005. "Does the Source of Capital Affect Capital Structure?" *Review of Financial Studies*, 19(1): 45–79.

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

Gabaix, Xavier. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica*, 79(3): 733–772.

Gertler, Mark, and Peter Karadi. 2015. "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics*, 7(1): 44–76.

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

Gopinath, Gita, Sebnem Kalemli-Ozcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez. 2015. "Capital Allocation and Productivity in South Europe." National Bureau of Economic Research.

Gürkaynak, Refet S., Brian Sack, and Eric Swanson. 2005. "The Sensitivity of Long-term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review*, 95(1): 425–436.

Hadlock, Charles J., and Joshua R. Pierce. 2010. "New Evidence on Measuring Financial Constraints: Moving Beyond the KZ Index." *Review of Financial Studies*, 23(5): 1909–1940.

Jermann, Urban, and Vincenzo Quadrini. 2012. "Macroeconomic Effects of Financial Shocks." *American Economic Review*, 102(1): 238–271.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses Local Projections." *American Economic Review*, 95(1): 161–182.

Kalemli-Ozcan, Sebnem, Bent Sorensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yesiltas. 2015. "How to Construct Nationally Representative Firm Level Data from the ORBIS Global Database." National Bureau of Economic Research.

Kaplan, Steven N., and Luigi Zingales. 1997. "Do Investment-Cash Flow Sensitivities Provide Useful Measures of Financing Constraints?" *Quarterly Journal of Economics*, 112(1): 169–215.

Kashyap, Anil K., Owen A. Lamont, and Jeremy C. Stein. 1994. "Credit Conditions and the Cyclical Behavior of Inventories." *Quarterly Journal of Economics*, 109(3): 565–592.

Khan, Aubhik, and Julia K. Thomas. 2013. "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity." *Journal of Political Economy*, 121(6): 1055–1107.

Kudlyak, Marianna, and Juan M. Sanchez. 2017. "Revisiting Gertler-Gilchrist Evidence on the Behavior of Small and Large Firms." *Journal of Economic Dynamics and Control*, 77: 48–69.

Mehrotra, Neil R., and Dmitriy Sergeyev. 2016. "Financial Shocks and Job Flows." Mimeo. Brown University.

Mian, Atif, and Amir Sufi. 2014. "What Explains the 2007-2009 Drop in Employment?" *Econometrica*, 82(6): 2197–2223.

Moscarini, Giuseppe, and Fabien Postel-Vinay. 2012. "The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment." *American Economic Review*, 102(6): 2509–2539.

Ottonello, Pablo, and Thomas Winberry. 2017. "Financial Heterogeneity and the Investment Channel of Monetary Policy." Mimeo. University of Chicago.

Rajan, Raghuram G., and Luigi Zingales. 1995. "What Do We Know about Capital Structure? Some Evidence from International Data." *Journal of Finance*, 50(5): 1421–1460.

Ramey, Valerie A. 2016. "Macroeconomic Shocks and Their Propagation." *Handbook of Macroeconomics*, 2: 71–162.

Romer, Christina D., and David H. Romer. 1989. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." *NBER Macroeconomics Annual*, 4: 121–170.

Romer, Christina D, and David H Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 1055–1084.

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

Shourideh, Ali, and Ariel Zetlin-Jones. 2012. "External Financing and the Role of Financial Frictions over the Business Cycle: Measurement and Theory." Mimeo. University of Pennsylvania.

Whited, Toni M., and Guojun Wu. 2006. "Financial Constraints Risk." *Review of Financial Studies*, 19(2): 531–559.

Wieland, Johannes F, and Mu-Jeung Yang. 2016. "Financial Dampening." National Bureau of Economic Research.

A Measurement framework

For clarity, the following paragraphs provide the details of the way in which we construct the size classification and growth measures used in section 3.

Sample selection Let i index firms and t index quarters. Let $x \in X$ index variables of interest; in the analysis, we use $X = \{\text{sales}, \text{inventory}, \text{NPPE stock}, \text{assets}\}$. Let:

$$\mathcal{I}_t(x) \quad \equiv \{ \quad i \quad \text{s.t.} \quad x_{i,t-4} > 0 \quad \text{and} \quad x_{i,t} > 0 \quad \} \tag{13}$$

We restrict attention to firms with strictly positive values of the variables of interest so as to compute log growth rates (see below). In order to be able to construct a consistent sample across variables of interest, we only consider firms $i \in \mathcal{I}_t$, where:

$$\mathcal{I}_t \equiv \bigcap_{x \in X} \mathcal{I}_t(x).$$

Size classification Let $a_{i,t}$ denote book assets. For every quarter t, we compute a set of percentiles,

$$\mathcal{P}_t = \left\{ \bar{a}_t^{(k)} \right\}_{k \in K},$$

where $K \subset [0, 100]$, $\bar{a}_t^{(0)} = 0$ and $\bar{a}_t^{(k)} = +\infty$. These percentiles are computed using the distribution of book assets of *all* firms, not only those firms $i \in \mathcal{I}_t$. Moreover, these percentiles are obtained using the Census-provided cross-sectional sampling weights $z_{i,t}$. We then define:

$$\mathcal{I}_t^{(k_1,k_2)} = \left\{ \quad i \in \mathcal{I}_t \quad \text{s.t.} \quad a_{i,t-4} \in \left[\bar{a}_{t-4}^{(k_1)}, \bar{a}_{t-4}^{(k_2)} \right[\quad \right\}. \tag{14}$$

In the case of the simple sample split between bottom 99% and top 1%, the small and large firms groups are defined as:

$$\begin{aligned} \mathcal{I}_t^{(\text{small})} &= \mathcal{I}_t^{(0,99)}, \\ \mathcal{I}_t^{(\text{large})} &= \mathcal{I}_t^{(99,100)} = \mathcal{I}_t \setminus \mathcal{I}_t^{(0,99)}. \end{aligned} \tag{15}$$

Growth rates For any $i \in \mathcal{I}_t$, we define growth rates as:

$$g_{i,t}(x) = \begin{cases} \log\left(\frac{x_{i,t}}{x_{i,t-4}}\right) & \text{if} \quad x \in \{\text{sales}, \text{inventory}, \text{NPPE stock}, \text{assets}\} \\ \frac{nppe_{i,t} - nppe_{i,t-4} + dep_{i,t-4,t}}{nppe_{i,t-4}} & \text{if} \quad x = \text{fixed investment}. \end{cases} \tag{16}$$

We focus on log growth-rates because they are easier to use in the decomposition of aggregate growth into firm-level growth rate discussed in section 4. Annual differences (instead of quarterly differences) are the main specification both because they are consistent with the size classification (which is based on one-year lags, so as to adequately capture initial size), and because they neutralize the issue of seasonal variation in the variables of interest. Cross-sectional averages of growth rates are then defined as:

$$\begin{aligned} \hat{g}_t^{(k_1,k_2)}(x) &\equiv \frac{1}{\mathcal{Z}_{t-4}^{(k_1,k_2)}} \sum_{i\in\mathcal{I}_t^{(k_1,k_2)}} z_{i,t-4}\, g_{i,t}(x) \\ Z_{t-4}^{(k_1,k_2)} &\equiv \sum_{i\in\mathcal{I}_t^{(k_1,k_2)}} z_{i,t-4}. \end{aligned} \tag{17}$$

and $z_{i,t-4}$ are the Census-provided cross-sectional sampling weights. Throughout, we analyze cross-sectional average time-series after de-meaning them (since the focus in not on long-term trends, but rather on the cyclicality of growth); we do not use any further detrending or filtering.

Robustness Our results for sales, inventory, the stock of net property, plant and equipment are robust to using half-growth rates of the form $2\frac{x_{i,t}-x_{i,t-4}}{x_{i,t}+x_{i,t-4}}$. Qualitatively and quantitatively, results do not change substantially whether one uses the one-year lagged or current weights in computing average growth rates of the form 17. Since the sample is titled toward larger firms, carrying the analysis using unweighted data ($z_{i,t} = 1, \forall (i,t)$) leads to qualitatively identical results, but somewhat smaller magnitudes.

B Decompositions of aggregate growth

Assume that all observations are equally weighted, that is:

$$z_{i,t} = 1 \quad \forall (i,t).$$

Let $\mathcal{I}_t^{(\text{small})} \subset \mathcal{I}_t$ denote the set of indexes of small firms, and $\mathcal{I}_t^{(\text{large})} = \mathcal{I}_t \setminus \mathcal{I}_t^{(\text{small})}$ be the set of large firms.[52] For some variable of interest $x \in \{\text{sales}, \text{inventory}, \text{NPPE stock}, \text{assets}\}$, and for some quarter t, define:

$$\begin{aligned} X_t &= \sum_{i\in\mathcal{I}_t} x_{i,t}, & X_{t-4} &= \sum_{i\in\mathcal{I}_t} x_{i,t-4}, & G_t &= \frac{X_t}{X_{t-4}}, \\ X_t^{(\text{small})} &= \sum_{i\in\mathcal{I}_t^{(\text{small})}} x_{i,t}, & X_{t-4}^{(\text{small})} &= \sum_{i\in\mathcal{I}_t^{(\text{small})}} x_{i,t-4}, & G_{t-4}^{(\text{small})} &= \frac{X_t^{(\text{small})}}{X_{t-4}^{(\text{small})}}, \\ X_t^{(\text{large})} &= \sum_{i\in\mathcal{I}_t^{(\text{large})}} x_{i,t}, & X_{t-4}^{(\text{large})} &= \sum_{i\in\mathcal{I}_t^{(\text{large})}} x_{i,t-4}, & G_{t-4}^{(\text{large})} &= \frac{X_t^{(\text{large})}}{X_{t-4}^{(\text{large})}}. \end{aligned} \tag{18}$$

These are simply totals for all firms and by group, along with their growth rates. Let:

$$s_{t-4} = \frac{X_{t-4}^{(\text{small})}}{X_{t-4}}$$

[52] See appendix A for a formal definition of the size classification. Here, we refer to an arbitrary size classification, so long as it constitutes a partition of $\mathcal{I}_t$; in the counterfactuals that are reported next, we will focus on partition between the bottom 99% and top 1% by lagged book assets.

be the initial fraction of the aggregate value of x accounted for by small firms. Define the following firm-level growth rates and shares by:

$$\begin{aligned} g_{i,t} &= \frac{x_{i,t}}{x_{i,t-4}} \\ w_{i,t-4} &= \begin{cases} \frac{x_{i,t-4}}{X_{t-4}^{(\text{small})}} & \text{if } i \in \mathcal{I}_t^{(small)} \\ \frac{x_{i,t-4}}{X_{t-4}^{(\text{large})}} & \text{if } i \in \mathcal{I}_t^{(large)} \end{cases} \end{aligned} \tag{19}$$

First, note that the total growth of x for small firms (the growth rate $G_{t-4}^{(\text{small})}$ defined above) can be decomposed as:

$$G_t^{(small)} = \hat{g}_t^{(\text{small})} + c\hat{o}v_t^{(\text{small})}, \tag{20}$$

where:

$$\begin{aligned} \hat{g}_t^{(small)} &= \frac{1}{\#\mathcal{I}_t^{(\text{small})}} \sum_{i \in \mathcal{I}_t} g_{i,t} \\ c\hat{o}v_t^{(\text{small})} &= \sum_{i \in \mathcal{I}_t^{(\text{small})}} \left(w_{i,t-4} - \frac{1}{\#\mathcal{I}_t} \right) \left(g_{i,t} - \hat{g}_t^{(small)} \right). \end{aligned} \tag{21}$$

The first term in this decomposition, $\hat{g}_t^{(\text{small})}$, is the cross-sectional average growth rate of the variable x. (Up to a constant and up to the approximation $log(x) \approx x - 1$ for x close to 1, this is the same variable as reported, for instance, in figure 1 for sales.) The second term can be interpreted as an (un-normalized) covariance, since $\frac{1}{\#\mathcal{I}_t^{(\text{small})}} = \frac{1}{\#\mathcal{I}_t^{(\text{small})}} \sum_{i \in \mathcal{I}_t^{(\text{small})}} w_{i,t-4}$. It captures the dependence between initial size (as proxied by the initial share of total size, $w_{i,t-4}$) and subsequent growth (as measured by $g_{i,t}$). Note that this decomposition is exact in any subset of $\mathcal{I}_t$; it holds for large firms as well, for example. Second, note that since $X_t = X_t^{(small)} + X_t^{(large)}$ and $X_{t-4} = X_{t-4}^{(small)} + X_{t-4}^{(large)}$, the following simple shift-share decomposition holds:

$$\begin{aligned} G_t &= s_{t-4} G_t^{(\text{small})} + (1 - s_{t-4}) G_t^{(\text{large})} \\ &= G_t^{(\text{large})} + s_{t-4} \left(G_t^{(\text{small})} - G_t^{(\text{large})} \right). \end{aligned} \tag{22}$$

Combining the two equations, we obtain the decomposition:

$$\begin{aligned} G_t &= \hat{g}_t^{(\text{large})} \\ &+ s_{t-4} \left(\hat{g}_t^{(\text{small})} - \hat{g}_t^{(\text{large})} \right) \\ &+ c\hat{o}v_t, \end{aligned} \tag{23}$$

where the covariance term $c\hat{o}v_t$ is given by:

$$c\hat{o}v_t = c\hat{o}v_t^{(\text{large})} + s_{t-4} \left(c\hat{o}v_t^{(\text{small})} - c\hat{o}v_t^{(\text{large})} \right).$$

C Decomposition of aggregate growth using DHS growth rates

This section replicates the decomposition results of section 4 using an alternative set of measures of growth at the firm level: the bounded growth rates introduced by Davis, Haltiwanger and Schuh (1996) (henceforth DHS). For any variable x, these growth rates are given by:

$$\tilde{g}_{i,t} = \frac{x_{i,t} - x_{i,t-4}}{\frac{1}{2}\left(x_{i,t} + x_{i,t-4}\right)} \in [-2, 2] .$$

These growth rates are a second-order accurate approximation to the standard growth rate $\frac{x_{i,t}}{x_{i,t-4}} - 1$ in a neighborhood of 1; furthermore, they are bounded, and moments of the distribution of these growth rates are therefore not too sensitive to outliers.

Using the same steps as outlined in appendix B, it is straightforward to verify that the following decomposition holds exactly:

$$\tilde{G}_t = \hat{\tilde{g}}_t^{(large)} + \tilde{s}_{t-4}(\hat{\tilde{g}}_t^{(large)} - \hat{\tilde{g}}_t^{(small)}) + c\hat{\tilde{o}}v_t^{(large)} + \tilde{s}_{t-4}(c\hat{\tilde{o}}v_t^{(large)} - c\hat{\tilde{o}}v_t^{(small)}),$$

where:

$$\begin{aligned}
\tilde{G}_t &= \frac{X_t - X_{t-4}}{\frac{1}{2}(X_t + X_{t-4})} \\
\tilde{s}_{t-4} &= \frac{X_t^{(small)} + X_t^{(large)}}{X_t + X_{t-4}} \\
\hat{\tilde{g}}_t^{(small)} &= \frac{1}{\#\mathcal{I}_t^{(\text{small})}} \sum_{i \in \mathcal{I}_t} \tilde{g}_{i,t} \\
c\hat{\tilde{o}}v_t^{(\text{small})} &= \sum_{i \in \mathcal{I}_t^{(\text{small})}} \left(\tilde{w}_{i,t-4} - \frac{1}{\#\mathcal{I}_t}\right)\left(\tilde{g}_{i,t} - \hat{\tilde{g}}_t^{(small)}\right),
\end{aligned}$$

and $\hat{\tilde{g}}_t^{(large)}$, $c\hat{\tilde{o}}v_t^{(\text{large})}$ are similarly defined. In this decomposition, the weights appearing in the covariance terms are given by:

$$\tilde{w}_{i,t} = \frac{x_{i,t} + x_{i,t-4}}{\sum_{i \in \mathcal{I}_t} x_{i,t} + x_{i,t-4}}.$$

Thus, they capture not the initial size of the firm relative to other firms initially in the same size group, but its average size over the period between $t-4$ and t, relative to the average size of firms initially in the same size group.

When we apply this decomposition to the same sample as in section 4, the two key results of the analysis using log growth rates still hold. First, the covariance terms in the decomposition account for a very small fraction of the overall correlation between aggregate growth and GDP growth; the lion's share of that correlation, instead, comes from the cross-sectional average components, $\hat{\tilde{g}}_t^{(small)}$ and $\hat{\tilde{g}}_t^{(large)}$. Table 9 makes this point; its contents are almost identical to those of Table 9 in the main text. Second, estimated elasticities of counterfactual time series for aggregate growth attempting to remove either the "excess sensitivity" or the cycliality of small firms overall are very close to the actual elasticities of time series for aggregate growth. Table 10 reports these results;

again, they are almost identical to the results from the same exercise conducted using log growth rates, and reported in Table 6 in the main text. The reason for the similarity between these results is simple: these two growth rates are very highly correlated at the firm level, in the sample of continuing firms used throughout in the main text. The results of section 4 thus do not critically depend on the use of log growth-rates for the construction of the decomposition of aggregate growth.

D Details on the comparison to Gertler and Gilchrist (1994)

This appendix compares our results to those of Gertler and Gilchrist (1994). That paper studies the behavior of small and large firms around dates identified by Romer and Romer (1989) as exogenous contractions in monetary policy. In this appendix, we replicate their analysis on the QFR micro-data for the period 1997q3-2014q1. There are two important differences between our analysis and theirs: the methodology, and the sample period analyzed. We start by discussing these differences, and then provide a reconciliation of their results to ours.

D.1 The methodology of Gertler and Gilchrist (1994)

The analysis of Gertler and Gilchrist (1994) centers around computing the cumulative change in revenue of an "aggregate" small and "aggregate" large firm. Revenues of the "aggregate" small firm are defined as the total sales of the group of firms which, starting from the smallest (by assets), accounting for a cumulative 30% of total sales at any point in time. Conversely, the revenues of the "aggregate" large firms are the total sales of firms which, starting from the largest (by assets), account for a cumulative 70% of revenue. This methodology differs from our analysis in two main ways: first, it focuses on aggregate, not firm-level growth; second, it results in a different definition of the relative importance of small and large firms. For completeness, what follows is a formal description of the construction of these series.

Let x denote nominal assets, let $\left\{x^{(1)},...,x^{(n)}\right\}$ denote the QFR's nominal asset bins' cutoffs, and let y denote nominal sales. For each quarter t, define $\underline{x}_t$ by:

$$\underline{x}_t = \max\left\{ x \in \left\{x^{(1)},...,x^{(n)}\right\} \Big/ \frac{\sum_{x_{i,t}\leq x} y_{i,t}}{Y_t} \leq 0.3 \right\}$$

Furthermore, let $\underline{x}_t^+$ be the cutoff immediately above $\underline{x}_t$ in the list $\left\{x^{(1)},...,x^{(n)}\right\}$. Compute the weight w_t such that:

$$w_t \frac{\sum_{x_{i,t}\leq \underline{x}_t} y_{i,t}}{Y_t} + (1-w_t)\frac{\sum_{x_{i,t}\leq \underline{x}_t^+} y_{i,t}}{Y_t} = 0.3$$

The growth rate of small firms' sales between time $t-1$ and t is then defined as:

$$G_t^{(small,GG)} = w_t \frac{\sum_{\{i/x_{i,t}\leq \underline{x}_t\}} y_{i,t}}{\sum_{\{i/x_{i,t-1}\leq \underline{x}_t\}} y_{i,t-1}} + (1-w_t)\frac{\sum_{\left\{i/x_{i,t}\leq \underline{x}_t^+\right\}} y_{i,t}}{\sum_{\left\{i/x_{i,t-1}\leq \underline{x}_t^+\right\}} y_{i,t-1}}.$$

The growth rate of large firms is defined analogously, using the cumulative sum of sales over the remaining bins of asset size. In our implementations of the GG methodology, we use four-quarter lagged growth rates, in order to remove seasonality in our data. Moreover, consistent with GG, we de-mean the small and large growth series before computing cumulative growth rates.

D.2 The choice of dates

The analysis of Gertler and Gilchrist (1994) also differs from ours in that it focuses on specific dates around which monetary policy changes stance. The outcome measured is then the average cumulative change in the revenue of the "aggregate" small and large firm defined above, across these dates. There are six such "Romer" dates in their analysis; only three directly overlap with our sample: 1978q3, 1979q4 and 1988q4. The recent analysis of Kudlyak and Sanchez (2017) has proposed adding two other dates to this list: 1994q2 and the credit crunch of 2008q3. In our comparison, we will therefore repeat their analysis on the 3 dates which directly overlap, and then on the set of 5 "Romer" dates used by Kudlyak and Sanchez (2017).

D.3 Replication and comparison

Figure 18 replicates the Gertler-Gilchrist analysis on the overlapping portion of our sample: 1977q3 to 1990q4. The lines reported in each panel are averages over the three Romer-Romer dates of 1978q3, 1979q4 and 1988q4. The top left panel plots the path of sales when small and large firms are defined as we do in the main text: using percentiles of the lagged distribution of assets, and reporting equal-weighted (as opposed to value-weighted) growth rates. The cumulative change in sales is between -13.8% for small firms and -6.3% for large firms under this methodology. The second panel repeats this exercise, but moving from equal- to value-weighted growth rates. Results are very similar, consistent with the evidence, in section 4, that the covariance term which connects equal- and value-weighted growth rates does not have a strong cyclical component. The black line in this graph is the cumulative change in total sales in the sample over these dates. The cumulative aggregate sales decline for sales overall is -8.8%, versus -6.2% for large firms. Thus, small firms substantially "amplify" the response of aggregate sales (by 42%, or $\frac{8.8-6.2}{6.2}$). Finally, the third panel exactly replicates the methodology of Gertler and Gilchrist (1994), in particular using the aggregated micro-data in the same format as original published by the QFR. It finds substantially the same differential response as the left and middle panels. Thus, for this sample period, the results are quantitatively and qualitatively consistent across methodologies, and lead to the conclusion that small firms "amplify" the response of aggregate sales by about 40%.

Figure 19 next replicates the Gertler-Gilchrist analysis on the 1977q3-2014q3 sample. The lines reported in each panel are averages over the fixe Romer-Romer dates of 1978q3, 1979q4, 1988q4, 1994q2 and 2008q3, as updated by Kudlyak and Sanchez (2017). The top left panel, which reports the cumulative change in sales using the same methodology as we do in the main text, leads to a sales decline of -9.2% for small firms and -5.6% for large firms. The second panel, using value-weighted growth rates, shows an aggregate decline of -5.9% for large firms, -9.7% for small firms, and -7.2% overall. The black line in this graph is the cumulative change in total sales in the sample over these dates. The last panel, using the methodology of Gertler and Gilchrist (1994), finds approximately the same results. The three methodologies therefore again deliver the same results. However, small firms (under all three methods) are now reponsible for a smaller amount of amplification: 22% $\left(=\frac{7.2-5.9}{5.9}\right)$ instead of 42%.

This discussion suggests that most of the difference between our conclusions and the conclusions of Gertler and Gilchrist (1994) primarily stems not from methodological distinctions, but from differences in the periods which we study. Their focus on specific dates differs from our approach of measuring an average difference in business-cycle (or monetary shock) sensitivity. In particular, the Volcker recessions lead to a particularly sharp excess response of small firms, which likely dominates the original findings of Gertler and Gilchrist (1994). Most importantly, the tendency for small firms to respond more to monetary policy tightenings may have declined over the second half of the sample, as also argued by Kudlyak and Sanchez (2017). This difference in the response to monetary shock across periods may reflect either changes in the conduct of monetary policy, or changes in the transmission of these shocks to small firms.

E The cyclicality of investment rates

In the QFR data, two cyclical properties of firm-level investment stand out. First, the contemporaneous correlation of firm-level investment with GDP growth, after controlling for industry effects, is slightly negative among the top 0.5% of firms, as reported in Table 4. Second, during recessions, the decline in investment among the top 1% of firms lags that of the bottom 99% of firms by 2-4 quarters, as indicated by the right panel of Figure 11. This appendix argues that the lag structure in investment among the largest firms can also be documented in two analogous data sources: the manufacturing segments of the annual and quarterly versions of Compustat.[53]

E.1 Data construction and summary statistics

Annual data Our source for the annual version of Compustat is the monthly update of the Fundamentals Annual file.[54] In order to obtain up-to-date industry identifiers, we merge this file

[53] Replication code for this exercise is available from the authors upon request.

[54] We use the latest version of the `funda` file, available on WRDS at: `/wrds/comp/sasdata/nam/funda.sas7bdat`. We use only firm-year observations with strictly positive assets (variable `at`) and which satisfy the four standard screens INDL for industry format, STD for data format, D for population source and C for consolidation. The company file we use is the latest version available at: `/wrds/comp/sasdata/nam/company/company.sas7bdat`.

with the Company file; whenever the 3-digit NAICS historical code is missing, we fill it with the next most recent available observation, using the Company file NAICS as the last (year 2017) NAICS observation.

In order to facilitate comparison with the QFR results, we focus on the following measure of investment:

$$ik_{i,t} = \frac{k_{i,t} - k_{i,t-1} + dep_{i,t}}{k_{i,t-1}}.$$

Here, $k_{i,t}$ is the stock of net property, plant and equipment reported on the balance sheet of firm i in year t, and $dep_{i,t}$ is depreciation reported in the firm's year t income statement.[55] Both $k_{i,t}$ and $dep_{i,t}$ are deflated using the BEA price index for manufacturing, as in the main text; the results also hold when using the BEA's 3-digit NAICS annual price indices to deflate nominal values. We keep firm-year observations in sample if (a) t is between 1977 and 2014; (b) the firm-year observation is incorporated in the US (variable `fic` from the company file equal to "USA"); (c) the 3-digit NAICS code is between 311 and 339 in sample; (b) $k_{i,t}$ is non-missing and weakly larger than 1m$; (c) $dep_{i,t}$ is non-missing and weakly positive.

Each year, we create four size groups, corresponding to the four quartiles of the sample distribution of book assets. The average size of firms in each group over the 1977-2014 sample is reported in Table 11, after deflating book assets by the manufacturing price index. As in the main text, firms are then grouped according to their one-year lagged position in the firm size distribution. Relative to the overall sample, the regression sample is the subset of firm-year observations such thatthe firm is also present in sample one year prior; (b) total depreciation $dep_{i,t}$, in nominal terms, is weakly smaller than the one-year lagged stock of net property, plant and equipment. This latter criterion helps filter very large positive observation of $ik_{i,t}$. The resulting annual sample has 72363 firm-year observations.

Quarterly data We follow a similar procedure to construct the quarterly sample.[56] The fundamentals quarterly file does not contain NAICS 3-digit identifiers. Whenever possible, we use the 3D-NAICS identifier at the annual frequency, as described above; otherwise, we use the identifier from the company files. As in the QFR data, we construct year-on-year investment rates at the quarterly frequency for each firm: $ik^q_{i,t} = \frac{k^q_{i,t} - k^q_{i,t-4} + dep_{i,t-4,t}}{k^q_{i,t-4}}$. Here, t now denotes a quarter; $k^q_{i,t}$ denotes the net stock of property, plant and equipment (variable `ppentq`) deflated by the price index for manufacturing; we interpolate the annual time series in order to obtain quarterly data. The variable $dep_{i,t-4,t}$ denotes *total* depreciation over the preceding year, which we compute by taking the sum of reported depreciation in the four quarters up to and including quarter t. As in the annual data, we only keep observations for which $dep_{i,t-4,t} \geq k^q_{i,t-4}$ in nominal terms. Finally,

[55] We use fiscal year, variable `fyear`, to date our observations; replacing by the calendar year which most overlaps the firm's fiscal year does not change our results.

[56] We use the latest version of the `fundq` file, available on WRDS at: `/wrds/comp/sasdata/nam/fundq.sas7bdat`.

we keep only observations with fiscal years between 1984 and 2014, since little data is available at the quarterly frequency prior to 1984. The resulting quarterly sample has 186784 firm-quarter observations.

Summary statistics Table 11 reports summary statistics for the average size and the average investment rate in the three different samples. QFR firms in the size-groups 1-2 (corresponding to the bottom 99% of the QFR distribution of book assets) are substantially smaller, on average, than firms in the bottom two size groups of the Compustat samples (the bottom 50% of the Compustat distribution of book assets). However, firms in group 4 (the top 0.5% of firms in QFR, and the top 25% of firms in Compustat) have comparable sizes (approximately 7bn$ on average). Measured investment rate among smaller firms (groups 1-3) are somewhat lower in the QFR than they are in Compustat; however, for the top size group, they have the same average magnitude. This suggests that the top quartile of Compustat firms represents relatively well the top 0.5% of firms in the QFR, those with a differential investment behavior.

E.2 The cyclical properties of investment

We first document unconditional estimates of the cyclicality of investment across size groups in Compustat data sources, and compare them to the QFR estimates. We use the same framework as in the main text, described in equation (1), in order to quantify this cyclicality; in particular, we use year-on-year GDP growth as our proxy for the state of the business cycle, and we control for durable/non-durable industry effects and their interaction with the year-on-year GDP growth. (The results are unchanged when controlling for 3D-NAICS effects in the same way). Table 12 reports the results, along with the estimates of the coefficients in the QFR data, which are identical to those reported in Table 4.

In both the quarterly and the annual Compustat, the baseline coefficient has the same magnitude and the opposite sign as the coefficient for the largest size group, group 4. In both cases, one cannot reject that the sum of the two coefficients is equal to 0.[57] The baseline industry group corresponding to the coefficient reported in the first line of Table 12 are firms in the durable sector; however, estimates of the average marginal effect of GDP growth on investment (not reported) convey the same message. In annual data, the point estimate for the average marginal effect is 0.066, with a 95% confidence interval of $[-0.118; 0.245]$; in quarterly data, those numbers are -0.057 and $[-0.297, 0.182]$. Thus, in Compustat data as well as in QFR data, investment at the largest firms does not display a significantly positive correlation with contemporaneous GDP growth.

We next turn to the question of whether investment declines among large firms also display a lag in Compustat data. We estimate the same simple event study response for investment as the

[57]The t-statistic for the tests are -0.24 in annual data and 1.41 in annual data, respectively.

one described in section 5.2 of the main text, using the Compustat quarterly sample. In order to focus on the lag among the largest firms in the data, we trace out the cumulative investment rates of the top size group — groups 3 and 4 from Table 12 — and the bottom size group — groups 1 and 2 from Table 12. Figure 20 reports the results. As in the QFR data, investment lags the start of the recession: the peak of the cumulative investment rate occurs three quarters after the start of the recession in both size groups. Moreover, there is a sharper slowdown in the invsetment rate among the bottom size group (the cumulative investment rate is between quarters 0, when the recession starts, and 3 is smaller in the bottom size groups than in the top size groups). The difference in lags between the top and the bottom size groups is less visible than in the QFR data. The fact that the typical size of firms in the bottom size groups is substantially larger in the QFR than in Compustat may explain this discrepancy.

Overall, these findings indicate that Compustat data shares the two salient features of the QFR investment rates — the fact that the very largest firms do not display a positive contemporaneous correlation with GDP growth, and the fact that investment declines seems to lag the beginning of recessions.

F Model details

F.1 The model with no external finance

Sufficient conditions for excess sensitivity First note that, in the stationary equilibrium of the model, the (gross) growth rate of the capital stock of a constrained firm is given by:

$$\begin{aligned} g_{i,cons} &= \tfrac{k_{i,t+1}}{k_{i,t}} \\ &= \tfrac{n_{i,t}}{k_{i,t}} \\ &= \tfrac{zk_{i,t}^{\zeta}+(1-\delta)k_{i,t}}{k_{i,t}} \\ &= 1-\delta+zk_{i,t}^{1-\zeta} \\ &\geq 1-\delta+\tfrac{1}{\zeta}(r+\delta) \equiv g_{cons} \end{aligned}$$

where the last line comes from the fact that $k_{i,t} \leq k^*$. Note that $g_{cons} > 1$. By contrast, in steady-state, the (gross) growth rate of unconstrained firms is $g_{uncons} = 1$.

Now consider a firm which is constrained at $t = -1$ and stays constrained at $t = 0$, when the shock occurs. Following similar steps, the gross growth rate of its capital stock will be given by:

$$\begin{aligned} g_{cons}^{(0)} &= \tfrac{k_{i,1}}{k_{i,0}} \\ &= \tfrac{n_{i,1}}{k_{i,0}} \\ &= \tfrac{z\exp(-\epsilon)k_{i,0}^{\zeta}+(1-\delta)k_{i,0}}{k_{i,0}} \\ &= 1-\delta+z\exp(-\epsilon)k_{i,0}^{1-\zeta} \\ &\geq 1-\delta+\tfrac{1}{\zeta}(r+\delta)\exp(-\epsilon) \\ &\approx g_{cons}-\tfrac{1}{\zeta}(r+\delta)\epsilon \end{aligned}$$

Thus, the drop in growth relative to g_{cons} is approximately:

$$\Delta g_{cons.} = -\frac{1}{\zeta}(r+\delta)\epsilon.$$

By contrast, for unconstrained firms, it is straightforward to see that the drop in growth relative to g_{unc} is:

$$\Delta g_{unc.} = -\frac{\rho}{1-\zeta}\epsilon.$$

Thus, for sales growth among large firms to fall more, relative to trend, that growth among small firms, it must be the case that:

$$\frac{\rho}{1-\zeta} \geq \frac{1}{\zeta}(r+\delta),$$

which holds in the calibration we study. Note here that in both the data and the model, growth among small and large firms is measured relative to its long-run average. The "de-trending" used in this derivation is approximate, in that it substitutes the long-run average growth rate of small firms for its lower bound, $g_{cons.}$, instead of the actual cross-sectional average growth rate of small firms in steady-state. However, the impulse responses reported are constructed using the actual long-run average growth rate of small firms in the stationary steady-state; this does not change the conclusion that small firms do not display excess sensitivity in this model.

Calibration of the model We construct a quarterly calibration of the model; in particular, we set $\zeta = 0.8$, $\delta = \frac{0.20}{4}$ and $r = \frac{0.02}{4}$. Additionally, we set:

$$z = \left(\frac{\zeta}{\delta+r}\right)^{-1},$$

This normalization implies that the steady-state size of unconstrained firms satisfies $\log(k^*) = 0$.

Given a value for the entry size k_e such that $k_e < \bar{k}$, there exists a unique integer $N \geq 2$ such that:

$$\mathbf{n}^{N-1}(k_e) < k^* \quad , \quad \mathbf{n}^{N}(k_e) \geq k^*,$$

where $\mathbf{n}(k) \equiv x^{1-\zeta}k^{\zeta} + (1-\delta)k$, and $\mathbf{n}^j(.)$ is the j-th iterate of $\mathbf{n}$. The stationary distribution is then a discrete distribution $\{\mu_j\}_{j=0}^N$, with $\sum_{j=0}^N \mu_j = 1$, supported on $N+1$ points $\{k_j\}_{j=0}^N$, where:

$$k_j = \begin{cases} \mathbf{n}^j(k_e) & \text{if } 0 \leq j \leq N-1 \\ k^* & \text{if } j = N \end{cases} \tag{24}$$

Given the exit rate η, and a mass of entering firms M, the distribution is given by:

$$\mu_j = \begin{cases} (1-\eta)^j M & \text{if } 0 \leq j \leq N-1 \\ \frac{(1-\eta)^N}{\eta} M & \text{if } j = N \end{cases} \tag{25}$$

We normalize $M = \frac{1}{\eta}$, so that the total mass of firms is 1 in steady-state. We then pick the entry size k_e to be $k_e = (0.001)k^*$, similar to the $p50/p99$ ratio of book assets in the QFR. Given that $\log(k^*) = 0$, this requires $\log(k_e) = \log(0.001)$. Given this choice of k_e, $N(k_e)$ is determined; given the calibration above, we have $N = 113$. We then pick η so that, in steady-state, 1% of firms are unconstrained: $\frac{(1-\eta)^N}{\eta} = 0.01$. This choice allows us to think of the size-conditional impulse response reported in the main text as also reflecting the behavior of constrained and unconstrained firms. Given all other parameters, matching this target requires $\eta = 0.040$. This exit rate is somewhat higher than what is observed among the firms of the balanced QFR panel. With a lower curvature of the profit function, however, it is straightforward to obtain lower implied exit rates; moreover, the qualitative implications of the model are independent of the particular value chosen for η.

F.2 The model with debt financing

Characterization of optimal policies The following lemma, and the figure that accompanies it, gives the solution to the problem of the firm with financial constraints. For brevity, the proofs of the lemma and the others that follow are omitted, but they are available from the authors upon request.

Lemma 1 (Constrained solution)**.** *Assume that the borrowing constraint is* $\mathbf{C}^1$ *and satisfies:*

$$\frac{\partial \mathbf{b}}{\partial n_{i,t}}(n_{i,t}, z_{t+1}) \geq 0, \quad \mathbf{b}(0, z_{t+1}) = 0;$$

$$\frac{\partial \mathbf{b}}{\partial z_{t+1}}(n_{i,t}, z_{t+1}) \geq 0.$$

Let $\{\underline{n}_t\}_{t\geq 0}$ *be the unique solution to:*

$$\begin{aligned} \underline{n}_t &= \max\left(\mathbf{c}^{-1}\left(k^*_{t+1}; z_{t+1}\right), -\left(\frac{1}{\zeta}-1\right)(\delta + r_b)k^*_{t+1} + \frac{1}{1+r_b}\underline{n}_{t+1}\right), \\ \lim_{t\to+\infty} (1+r_b)^{-t}\underline{n}_t &\leq 0, \end{aligned} \tag{26}$$

where $\mathbf{c}\left(n,z\right) \equiv n+\mathbf{b}(n,z)$ *is the maximum investment capacity of a firm with net worth* n*, conditional on aggregate productivity being equal to* z*. The solution to the firm's problem takes one of three forms, corresponding to three regions for net worth:*

- ***If*** $n_{i,t} < \mathbf{c}^{-1}\left(k^*_{t+1}; z_{t+1}\right)$*, the firm is constrained:*

  $$k_{i,t+1} = \mathbf{c}\left(n_{i,t}, z_{t+1}\right), \quad d_{i,t} = 0, \quad \frac{1}{1+r_b} b_{i,t+1} = \mathbf{b}(n_{i,t}, z_{t+1}), \quad \mathbf{V}_t(k_{i,t}, b_{i,t}) < \mathbf{V}_t^{(unc)}\left(k_{i,t}, b_{i,t}\right).$$

  Investment is strictly smaller than the optimal unconstrained level: $k_{i,t+1} = \mathbf{c}\left(n_{i,t}, z_{t+1}\right) < k^*_{t+1}$*. The marginal value of net worth is strictly above* 1*.*

- ***If*** $n_{i,t} \in \left[\mathbf{c}^{-1}\left(k^*_{t+1}; z_{t+1}\right), -\left(\frac{1}{\zeta}-1\right)(\delta+r_b)k^*_{t+1} + \frac{1}{1+r_b}\underline{n}_{t+1}\right]$*, the firm is partially constrained; it invests at the currently optimal scale, but issues no dividends:*

  $$k_{i,t+1} = k^*_{t+1}, \quad d_{i,t} = 0, \quad \frac{1}{1+r_b} b_{i,t+1} = n_{i,t} - k^*_{t+1}, \quad \mathbf{V}_t(k_{i,t}, b_{i,t}) < \mathbf{V}_t^{(unc)}\left(k_{i,t}, b_{i,t}\right).$$

  The marginal value of net worth is strictly above 1*.*

- ***If*** $n_{i,t} > \underline{n}_t$*, the firm is fully unconstrained, can invest at the optimal scale today and at all future dates:*

  $$k_{i,t+1} = k^*_{t+1}, \quad d_{i,t} \geq 0, \quad \frac{1}{1+r_b} b_{i,t+1} \leq \mathbf{b}(n_{i,t}, x_t), \quad \mathbf{V}_t(k_{i,t}, b_{i,t}) = \mathbf{V}_t^{(unc)}\left(k_{i,t}, b_{i,t}\right).$$

  The marginal value of net worth is equal to 1*.*

The lemma says that there are three possible regions for firms' policies: either firms are constrained, in that they issue no dividends, borrow as much as possible, and invest below the optimal level today; or, they are partially constrained, in that they issue no dividends, but invest at the optimal level today and borrow less (strictly) than the maximum possible; or, they are fully unconstrained. Firms move up across these three regions as their net worth increases.

In the constrained region, investment today is entirely constrained by the firms' investment capacity,

$$k_{i,t+1} = \mathbf{c}(n_{i,t}, z_{t+1}) = n_{i,t} + \mathbf{b}(n_{i,t}, z_{t+1}) < k^*_{t+1}.$$

So the responsiveness of these firms' investment to shocks depend on their effect on *current* net worth, and potentially future productivity. By contrast, in the partically constrained and unconstrained region, investment today depends only on fundamentals tomorrow $k_{i,t} = k^*_{i,t+1}$.

The partially constrained region need not exist. Namely, for it to exist, it needs to be the case that:

$$\mathbf{c}^{-1}\left(k^*_{t+1}; z_{t+1}\right) < -\left(\frac{1}{\zeta}-1\right)(\delta+r_b)k^*_{t+1} + \frac{1}{1+r_b}\underline{n}_{t+1}.$$

The right-hand side of this equation is the level of net worth necessary today in order to be able to implement the unconstrained optimal plan *starting tomorrow*; the left-hand side is the level of net worth necessary to implement the unconstrained optimal level of investment *today*. So, the partially constrained region only exists if the fundamentals process is such that firms will need high(er) levels of net worth in the future in order to implement the unconstrained plan. Most likely, that will be when fundamentals are low today relative to what they will be in the future.

It is immediate to see that there are no partially constrained firms in the stationary steady-state of the model. Additionally, one can rule out the possibility by imposing some restrictions on the aggregate process $\{z_t\}_{t\geq 0}$ and on the borrowing constraint $\mathbf{c}$.

Lemma 2. *Let:*

$$\forall t \geq 0, \quad g_t \equiv -\left(\frac{1}{\zeta}-1\right)\frac{r_b+\delta}{1+r_b}\frac{k^*_{t+1}}{\mathbf{c}^{-1}\left(k^*_{t+1},z_{t+1}\right)}+\frac{1}{1+r_b}\frac{\mathbf{c}^{-1}\left(k^*_{t+2},z_{t+2}\right)}{\mathbf{c}^{-1}\left(k^*_{t+1},z_{t+1}\right)}. \tag{27}$$

Assume that $\{z_t\}_{t\geq 0}$ is increasing and bounded from above, and that $\{g_t\}_{t\geq 0}$ is strictly decreasing. Let:

$$T \equiv \min\{ \quad t\geq 0 \quad s.t. \quad g_t \leq 1 \quad \}.$$

Then the net worth threshold $\{\underline{n}_t\}_{t\geq 0}$ is given by:

$$\underline{n}_t = \begin{cases} -\left(\frac{1}{\zeta}-1\right)\frac{r_b+\delta}{1+r_b}k^*_{t+1}+\frac{1}{1+r_b}\underline{n}_{t+1} & \text{if} \quad t\leq T-1, \\ \mathbf{c}^{-1}\left(k^*_{t+1},z_{t+1}\right) & \text{if} \quad t\geq T. \end{cases} \tag{28}$$

In particular, if $g_0 \leq 1$, then the unconstrained threshold is always given by:

$$\underline{n}_t = \mathbf{c}^{-1}\left(k^*_{t+1},z_{t+1}\right);$$

as a result, firms are never partially constrained.

This lemma essentially places a restriction on the fundamentals of the model that ensures that the unconstrained threshold $\underline{n}_t$ does not grow "too fast" in the wake of the shock. The calibration below (and the particular functional form for $\mathbf{c}$ chosen) satisfy the restriction provided by lemma 5. This ensures that firms are always completely constrained, or completely unconstrained, which simplifies the analysis of the model.

F.3 Borrowing constraint and sufficient conditions for excess sensitivity

We assume that the borrowing constraint is given by:

$$\mathbf{b}(n_t,z_{t+1}) = \left(\frac{1}{\theta}\left(\frac{z_{t+1}}{z}\right)^{\alpha}-1\right)n_t \quad , \quad \alpha\geq 0, \quad \theta\leq 1.$$

This parametrization captures some of the limit cases we are interested in. As $\theta \to 0$, the frictionless model obtains; when $\theta = 1$ and $\alpha = 0$, firms cannot borrow and the baseline model (with the addition of saving) obtains. Finally, the parameter α controls the sensitivity of the borrowing threshold to the aggregate shock, z_t; when $\alpha = 0$, the borrowing constraint only depends on net worth, and not on the shock; when $\alpha \in]0,1]$, the borrowing constraint is a concave function of the aggregate shock; and when $\alpha \in]1,+\infty]$, it is a convex function of the aggregate shock. Having a specific functional form will also allow us to plot impulse responses of the model.

Note that, given the functional form chosen for the borrowing constraint, the parameter α is irrelevant to the determination of the steady state. In what follows, we use:

$$\theta = 0.8,$$

implying a debt-to-asset ratio of about 0.2 in the version of the model with borrowing constraints that do not vary with productivity. This figure is consistent with the average net debt-to-asset ratio which we documented in the QFR data. We leave other parameters unchanged relative to the baseline model without borrowing.

The parameter α controls the ability for the model to generate excess sensitivity of sales and investment. To see this, first note that, following the same steps as in the model without borrowing, an approximation to the growth rate of constrained firms in the stationary steady-state of the model is:

$$g_{cons} = \frac{1}{\theta}\left(1 - \delta + \frac{1}{\zeta}(r+\delta)\right).$$

The impact growth rate on impact, on the other hand, can be bounded from below by:

$$g_{cons}^{(0)} \quad \geq \tfrac{1}{\theta}\exp(-\alpha\epsilon)\left(1-\delta+\tfrac{1}{\zeta}(r+\delta)\exp(-\epsilon)\right)$$

Thus, the impact response of growth among constrained firms, relative to the long-run steady-state, is:

$$\Delta g_{cons} = -\frac{1}{\zeta}(r+\delta)\epsilon - \frac{1}{\theta}\alpha\epsilon\left(\frac{1}{\zeta}(r+\delta)+(1-\delta)\right) + o(\epsilon).$$

The impact response of unconstrained firms is the same as in the previous model. Thus, excess sensitivity of small/constrained firms firms will obtain so long as:

$$\frac{\rho}{1-\zeta} \leq \frac{1}{\zeta}(r+\delta) + \frac{1}{\theta}\alpha\left(\frac{1}{\zeta}(r+\delta)+(1-\delta)\right),$$

and in particular, for sufficiently high values of α. In the reported impulse responses, we use $\alpha = 5$, which ensures that this condition holds.

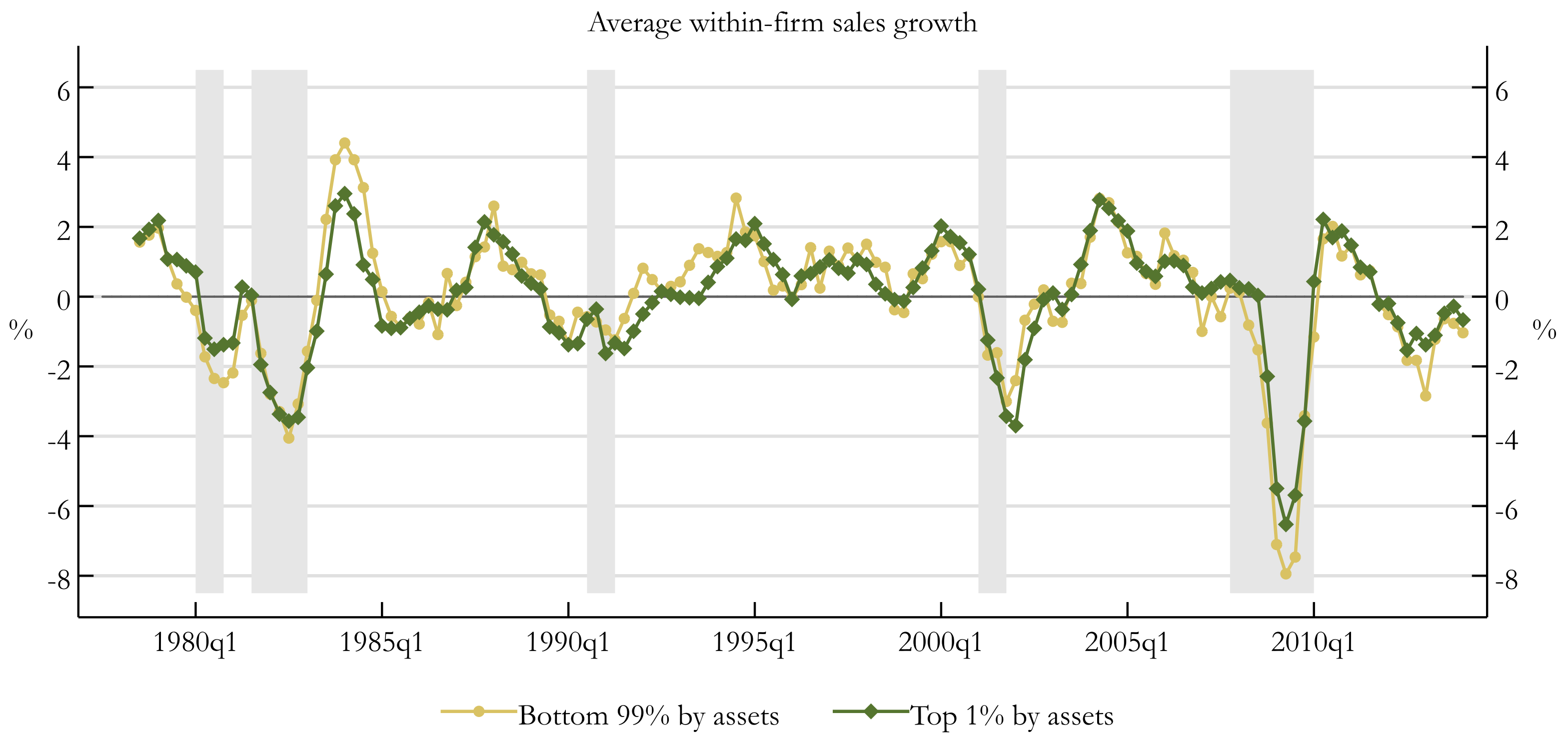

Figure 1: Average firm-level growth rate of sales of small (yellow, round markers) and large (green, diamond markers) firms. Small firms are those belonging to the bottom 99% of the one-year lagged distribution of book assets, while large firms are those belonging to the top 1% of the one-year lagged distribution of book assets. See appendix A for details on the construction of size groups and growth rates. Times series are demeaned.

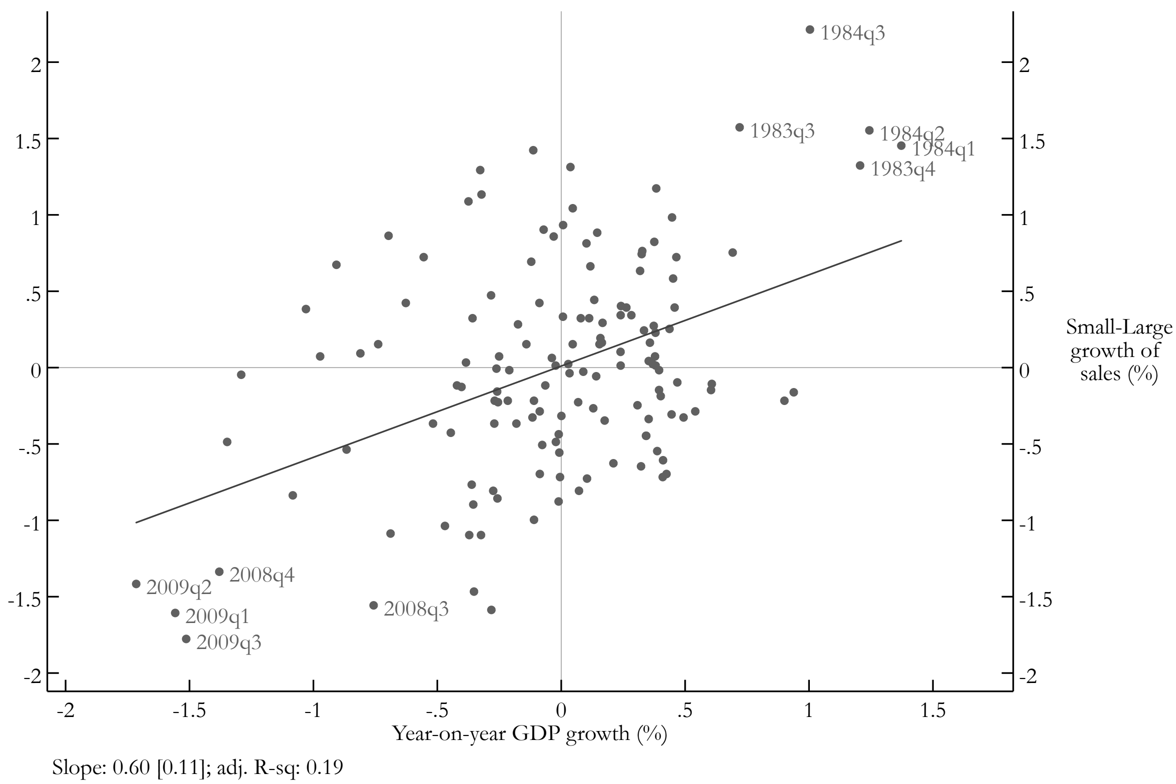

Figure 2: Difference between average growth rate of sales $\hat{g}_t^{(small)}(\text{sales}) - \hat{g}_t^{(large)}(\text{sales})$ (vertical axis) and year-on-year GDP growth (horizontal axis). Both series are demeaned. White standard errors in brackets.

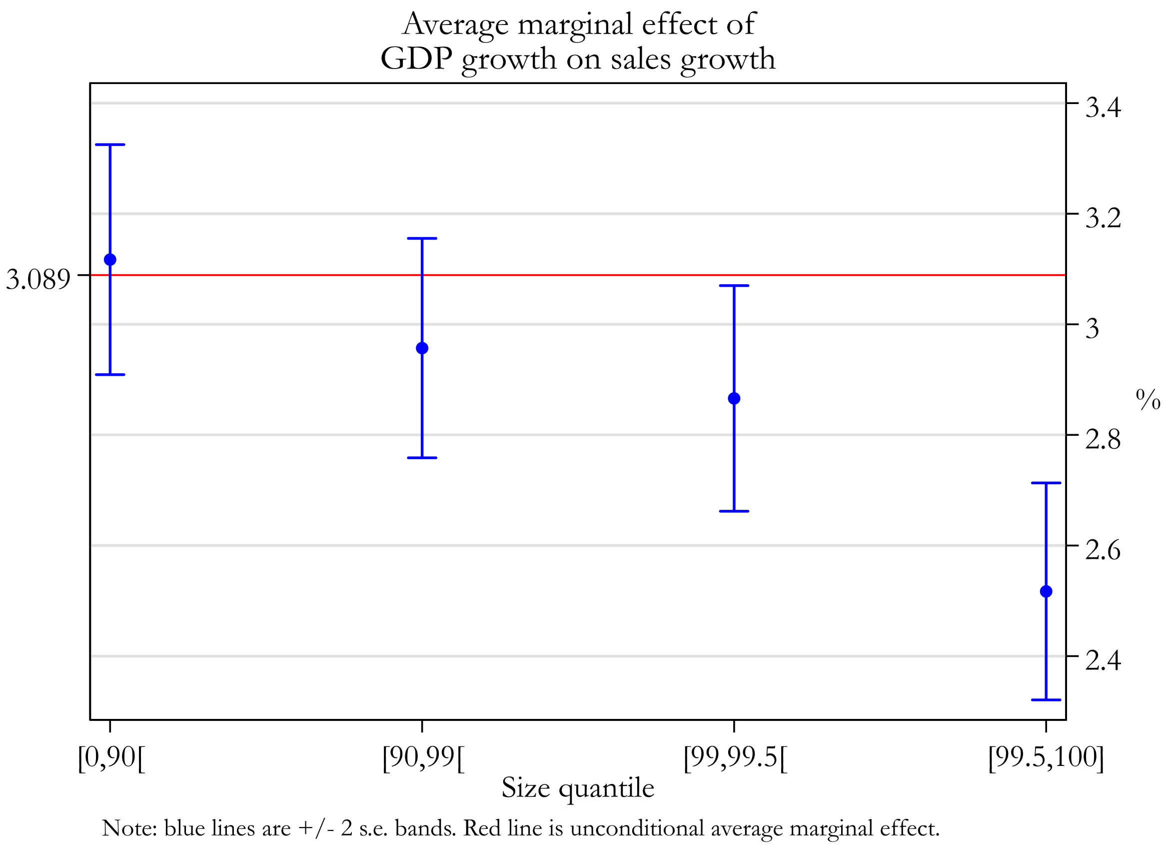

Figure 3: Marginal effects of GDP growth on sales growth, by size group (blue boxplots), and unconditionally (red line). The marginal effects are computed using estimates of model (1).

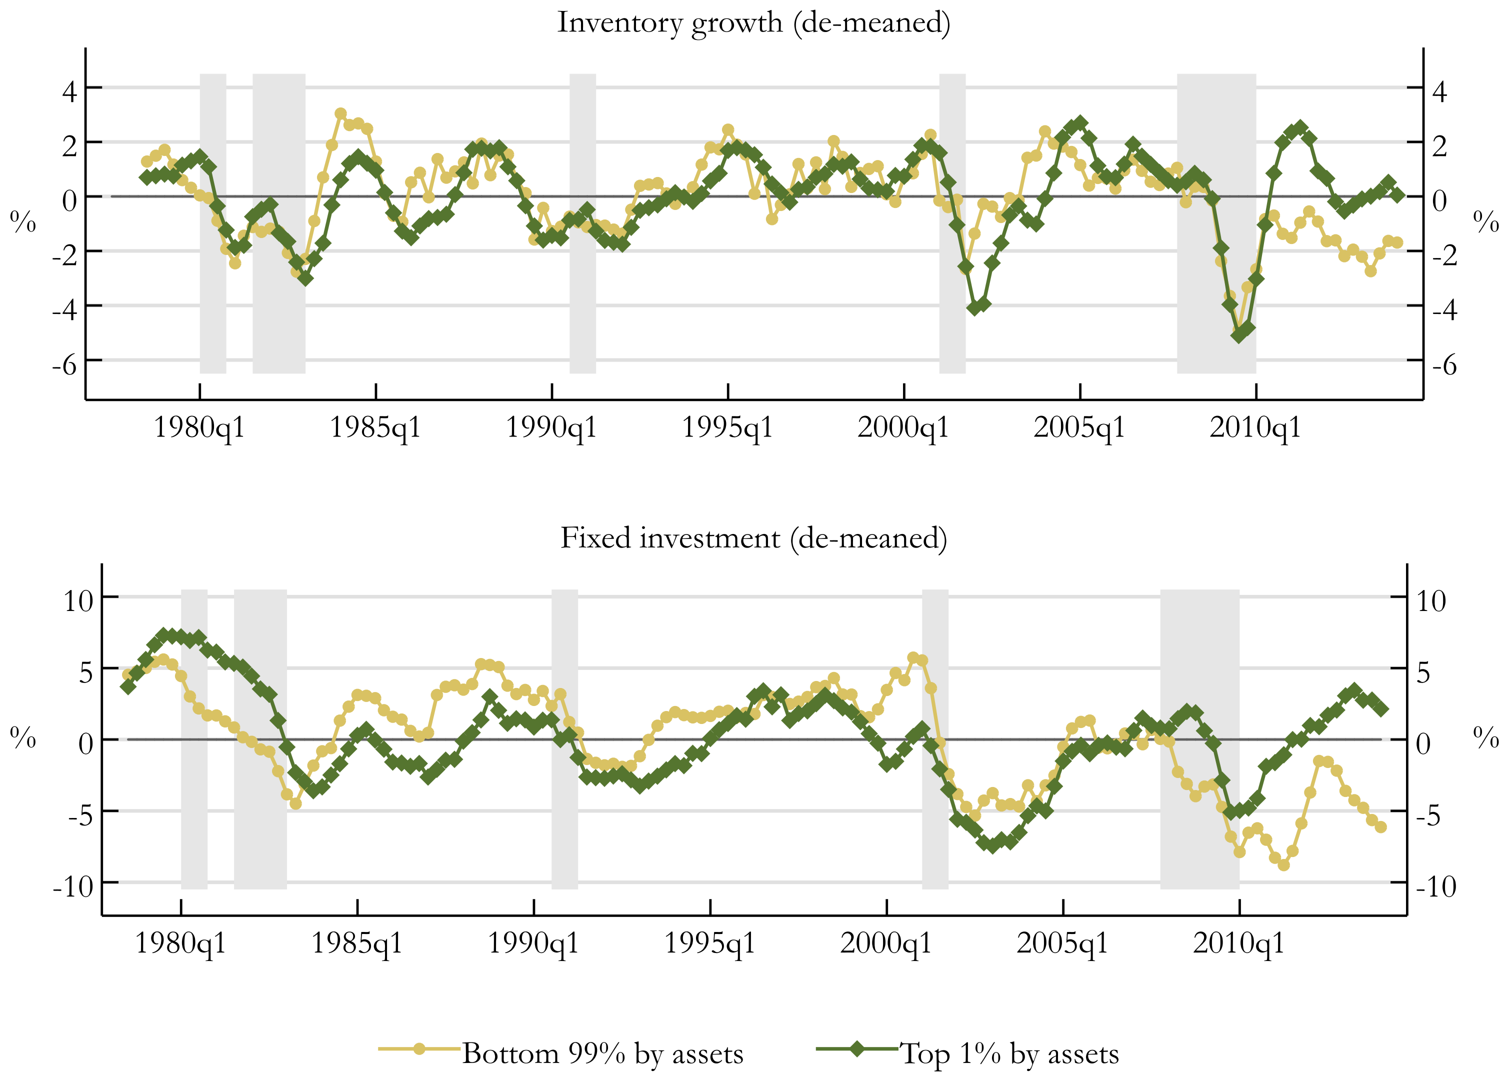

Figure 4: Average firm-level growth rate of small (yellow, round markers) and large (green, diamond markers) firms; top: inventory growth rate; bottom: fixed investment rate. Small firms are those belonging to the bottom 99% of the one-year lagged distribution of book assets, while large firms are those belonging to the top 1%. See Appendix A for details on the construction of size groups and growth rates. All series are demeaned.

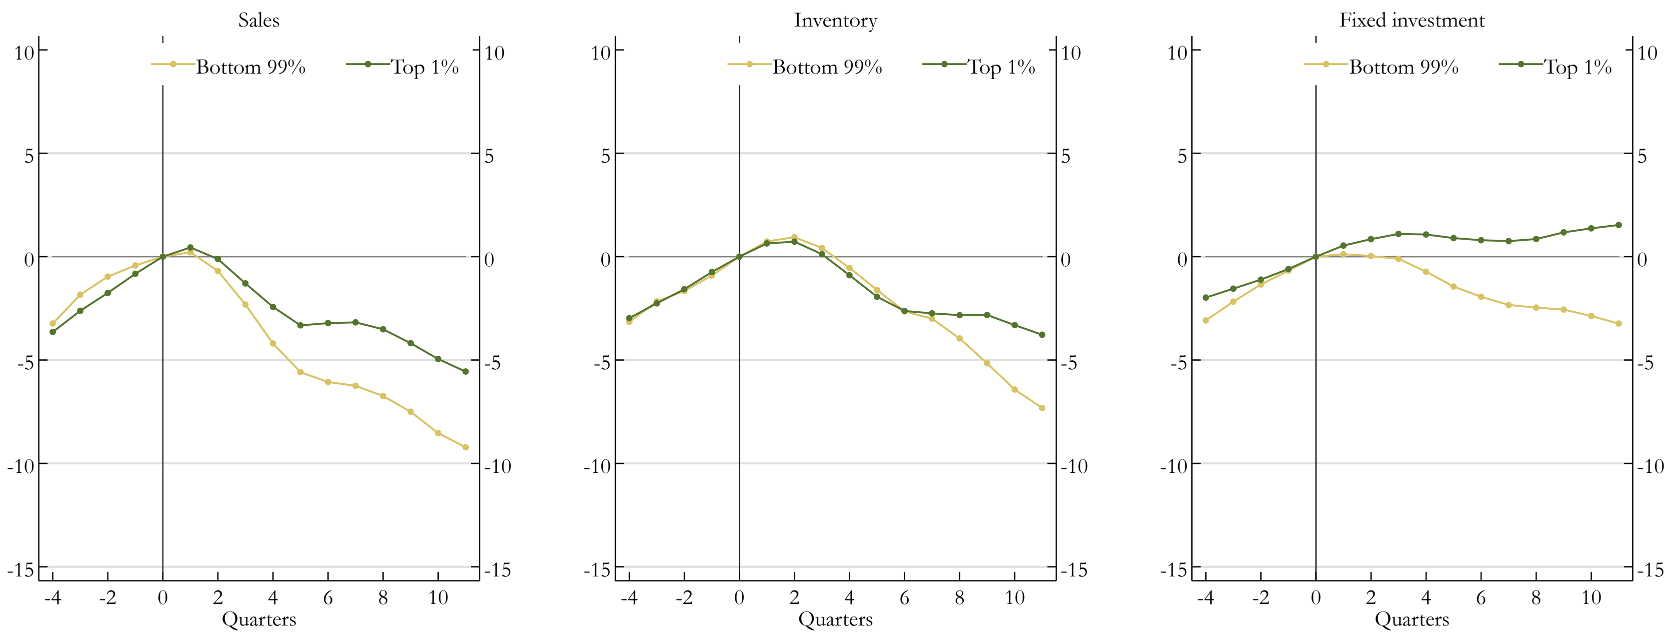

Figure 5: Cumulative change of sales, inventory and fixed assets around the five updated Romer dates of 1978q3, 1979q4, 1988q4, 1994q2 and 2008q3. All panels use the equal-weighted growth rates and the (0,99)/(99,100) size classification. Data is from 1977q3 to 2014q1. All time-series are de-meaned.

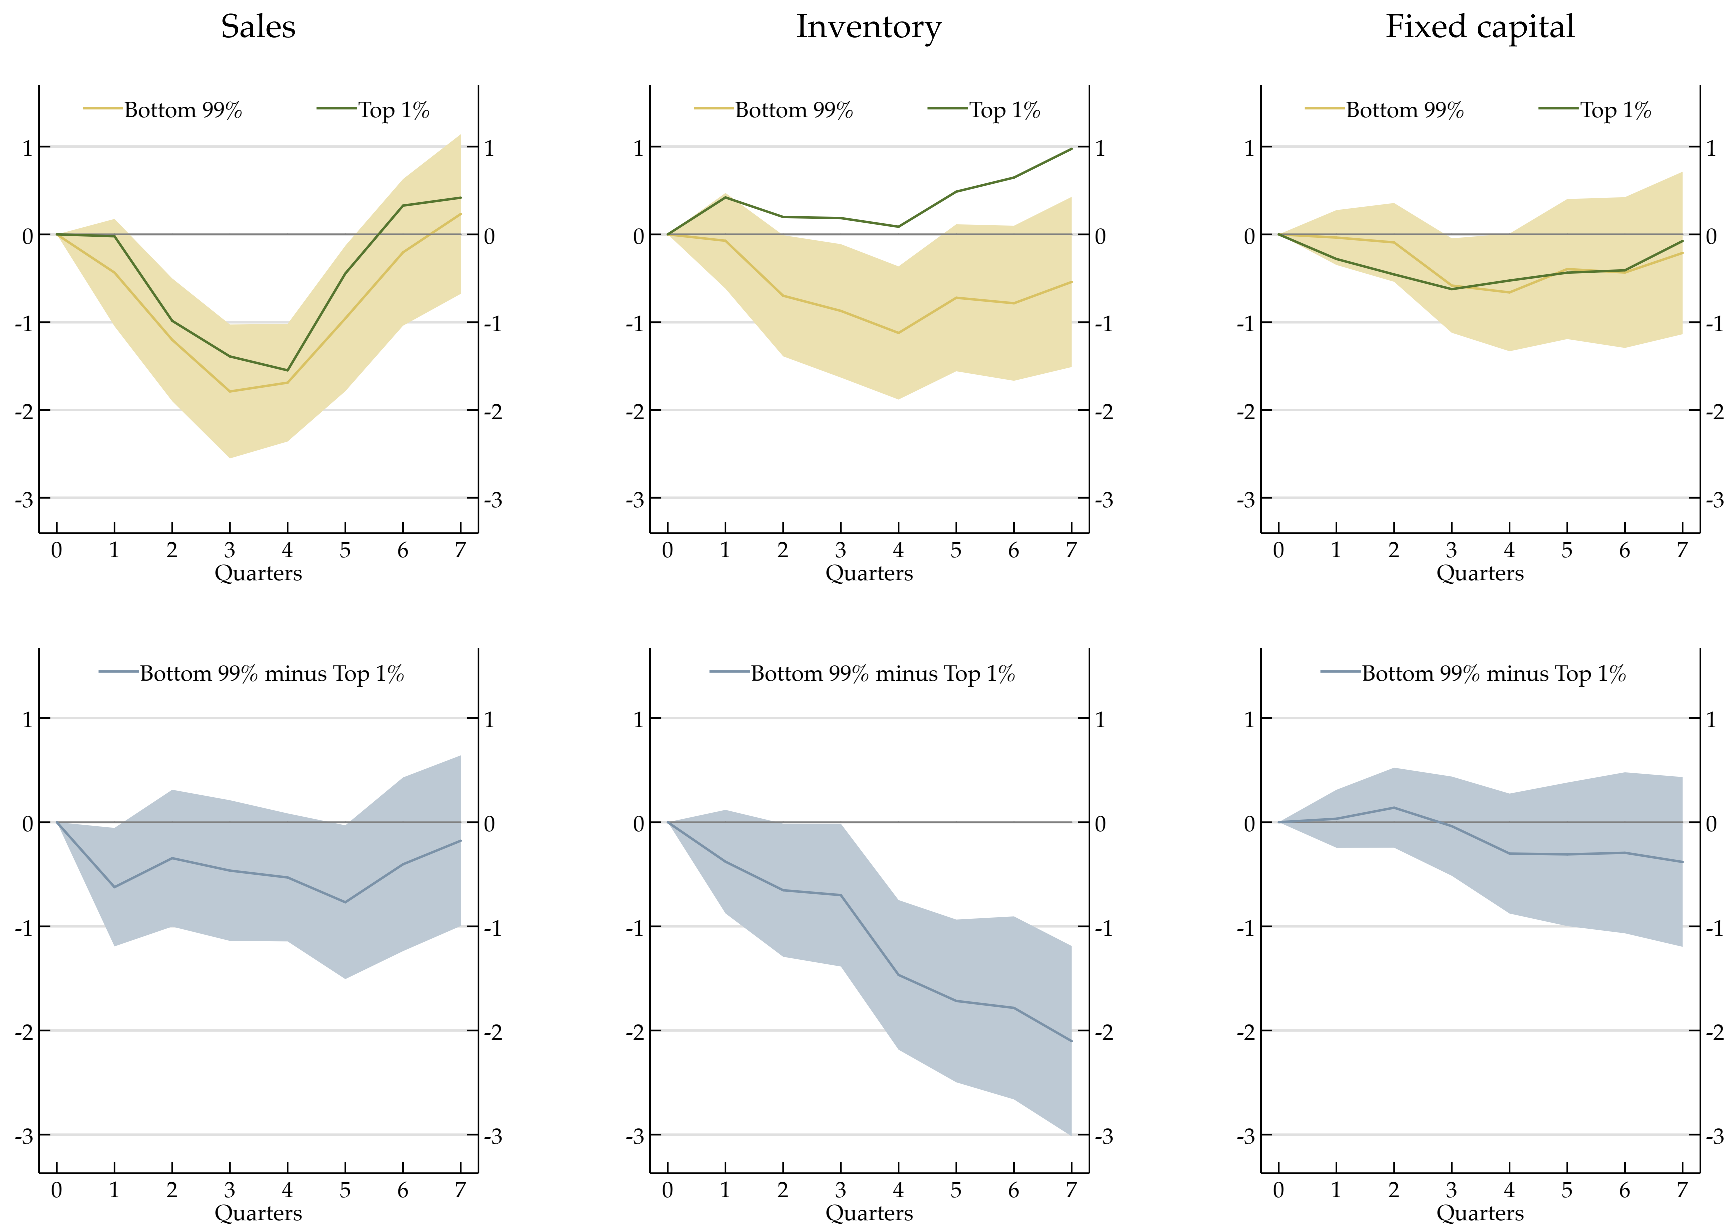

Figure 6: Firm-level response of sales, inventory and fixed capital to an innovation to the Romer and Romer (2004) shock. The estimated specification is model 2. The top row of graphs reports the average marginal effect at the mean of a one percentage point increase in $rr_{t-1,t}$, for the bottom 99% and top 1% size group. The yellow shaded area is the 95% confidence interval; standard errors are clustered at the firm-level and heteroskedasticity-robust. The bottom row of graphs reports the difference in the OLS coefficients $\beta^{(h)}_{(0,99)} - \hat{\beta}^{(h)}_{(99,100)}$, along with its 95% confidence interval. Data is from 1977q3 to 2007q4.

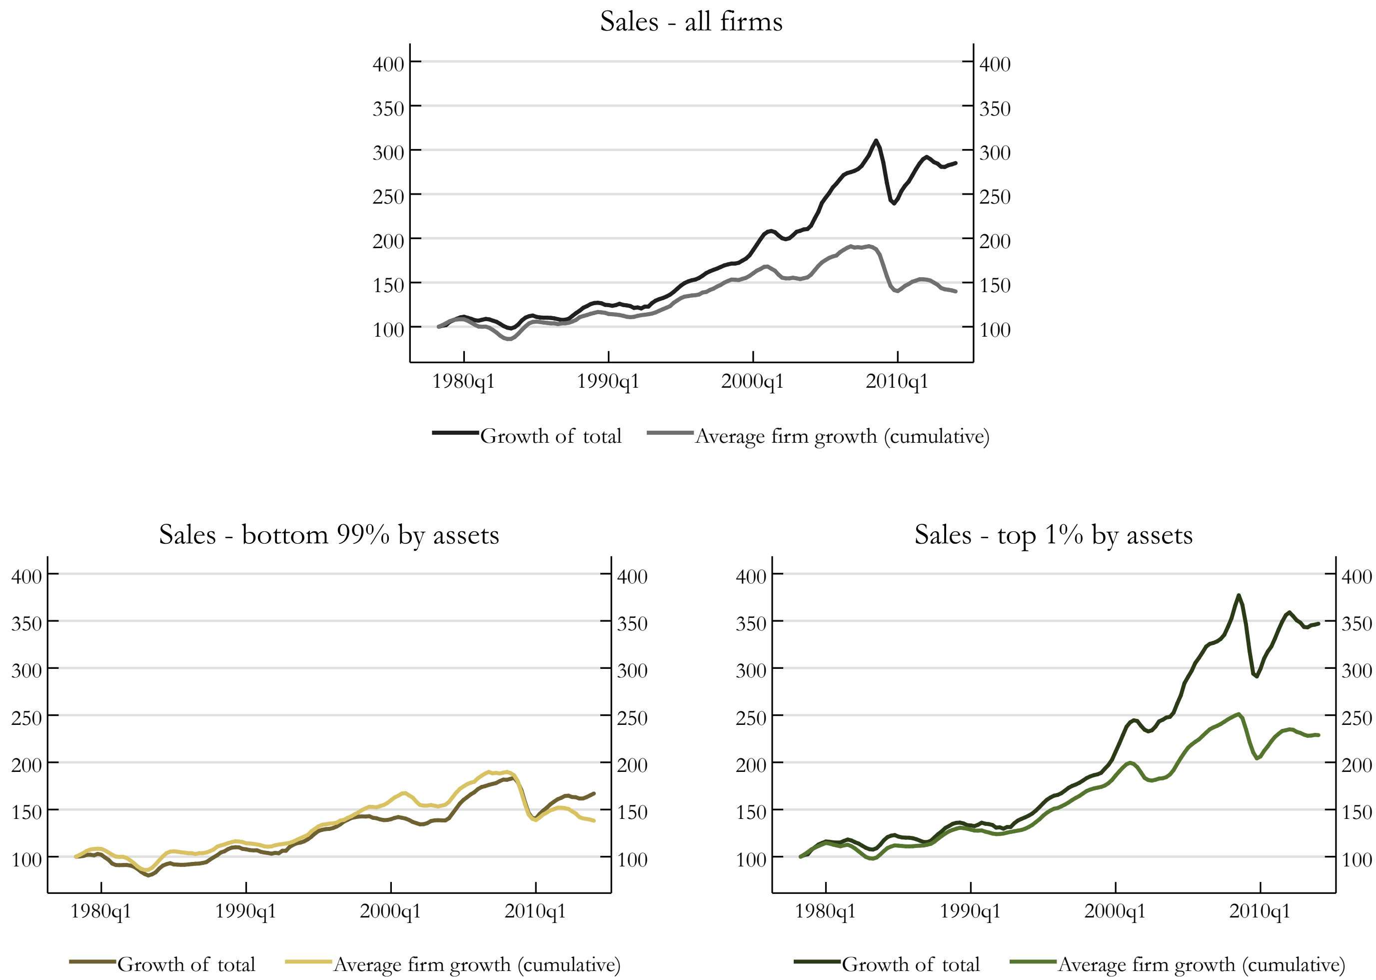

Figure 7: Aggregate sales and average within-firm cumualtive growth rate of sales. Each panel reports total annual sales (the cumulative value of $\hat{G}_t$), and the cumulative average growth rate of sales (the cumulative value of $\hat{g}_t$), for a different group of firms. All series are normalized to 100 in 1978q1.

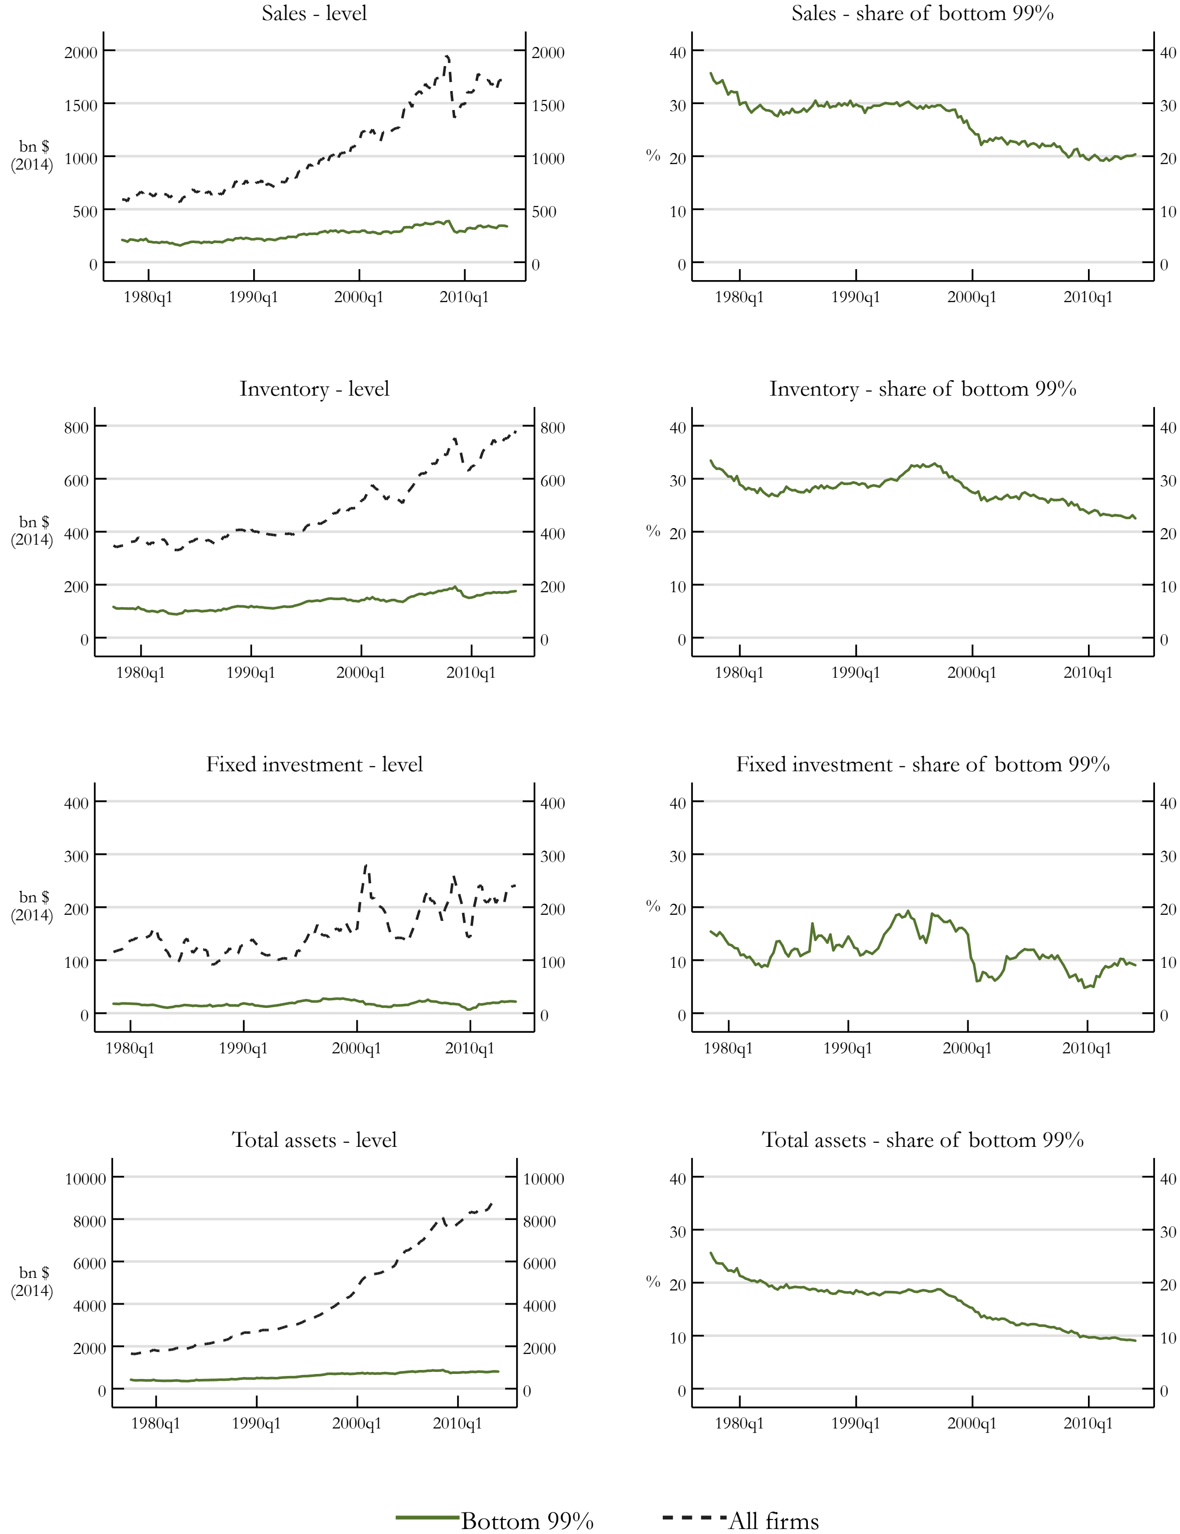

Figure 8: Concentration of sales, inventory, fixed investment, and total assets in the US manufacturing sector. The left column reports total nominal values for the bottom 99% and top 1% of firms by size. All series are deflated by the BEA price index for manufacturing, normalized to 1 in 2009q1; the series is available at http://bea.gov/industry/gdpbyind_data.htm. Series are unfiltered. The right column reports the share of the bottom 99% (the ratio of the corresponding graph in the left column).

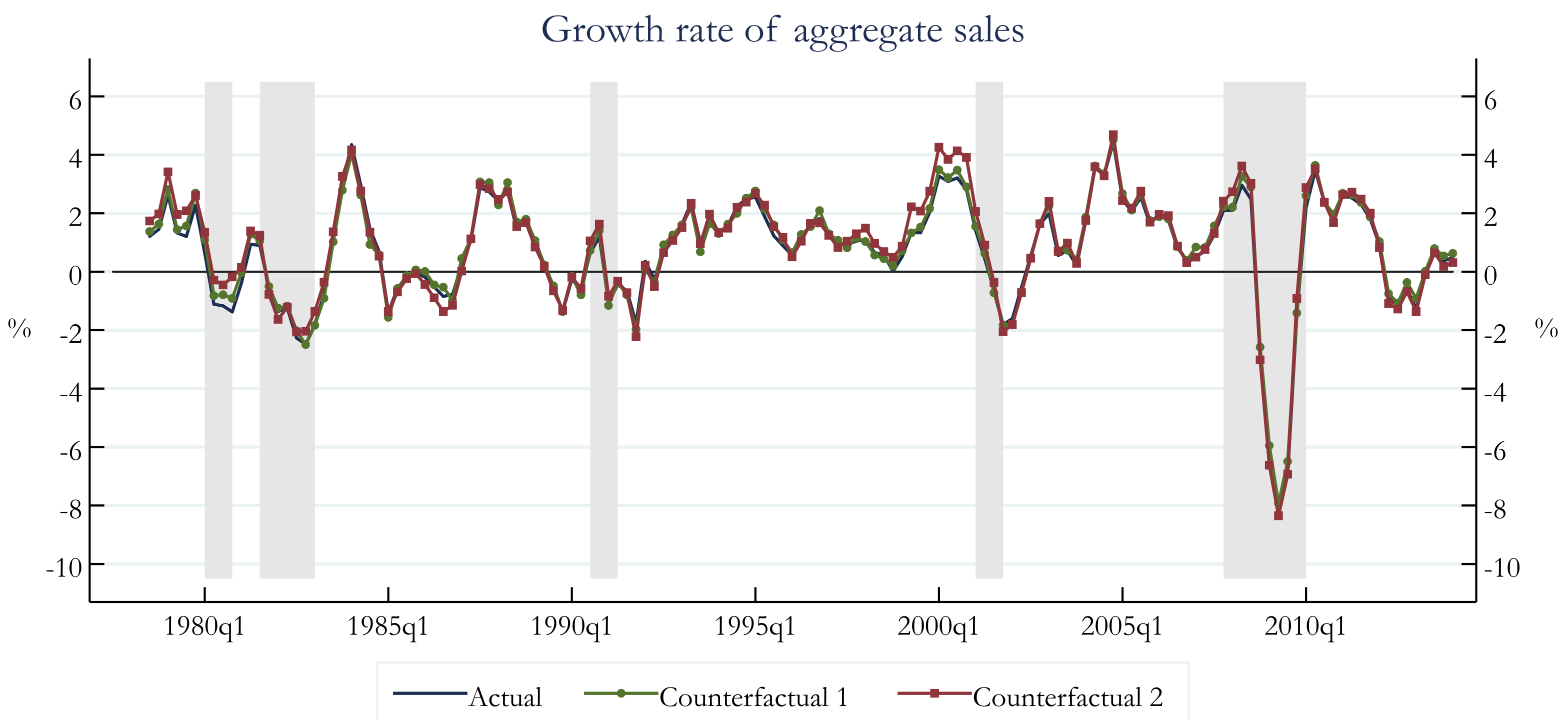

Figure 9: Aggregate growth rate of sales G_t (solid blue line), counterfactual growth rate 1 $G_t^{(1)}$, and counterfactual growth rate 2 $G_t^{(2)}$.

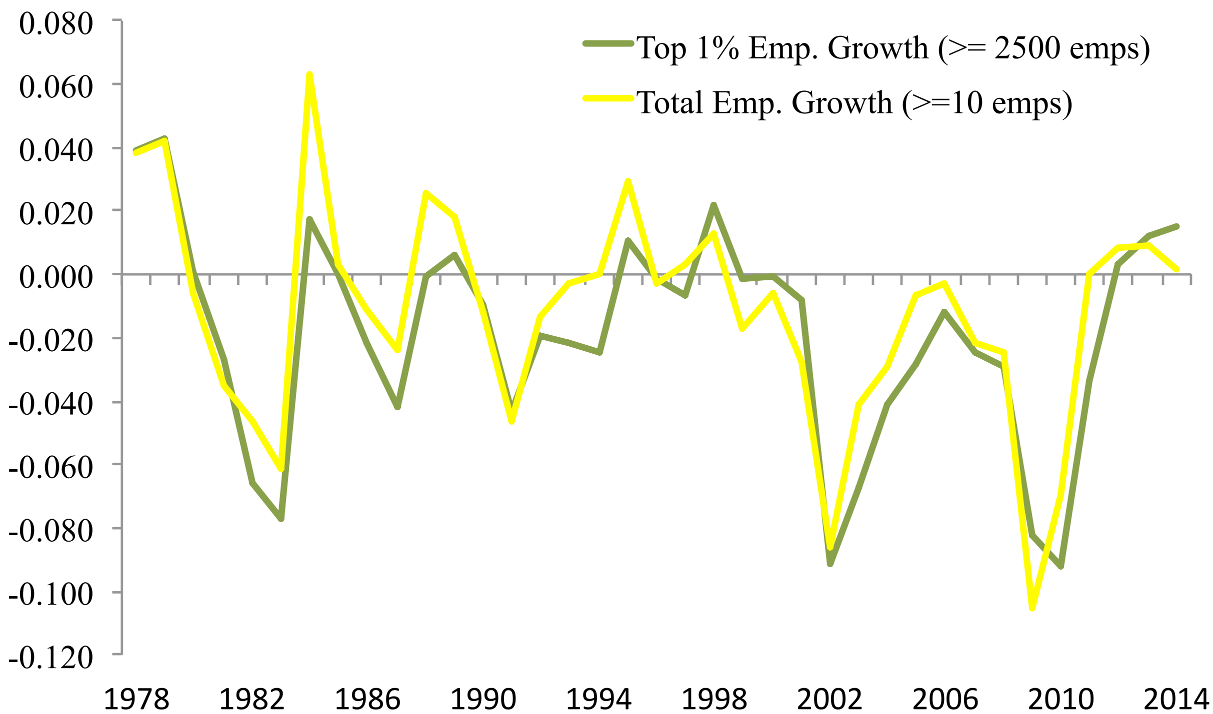

Figure 10: The green line displays annual employment growth for the estimated top 1% of manufacturing firm. The yellow line displays annual employment growth for all manufacturing firms with over 10 employees (our estimate of the portion of manufacturing employment captured in our data set).

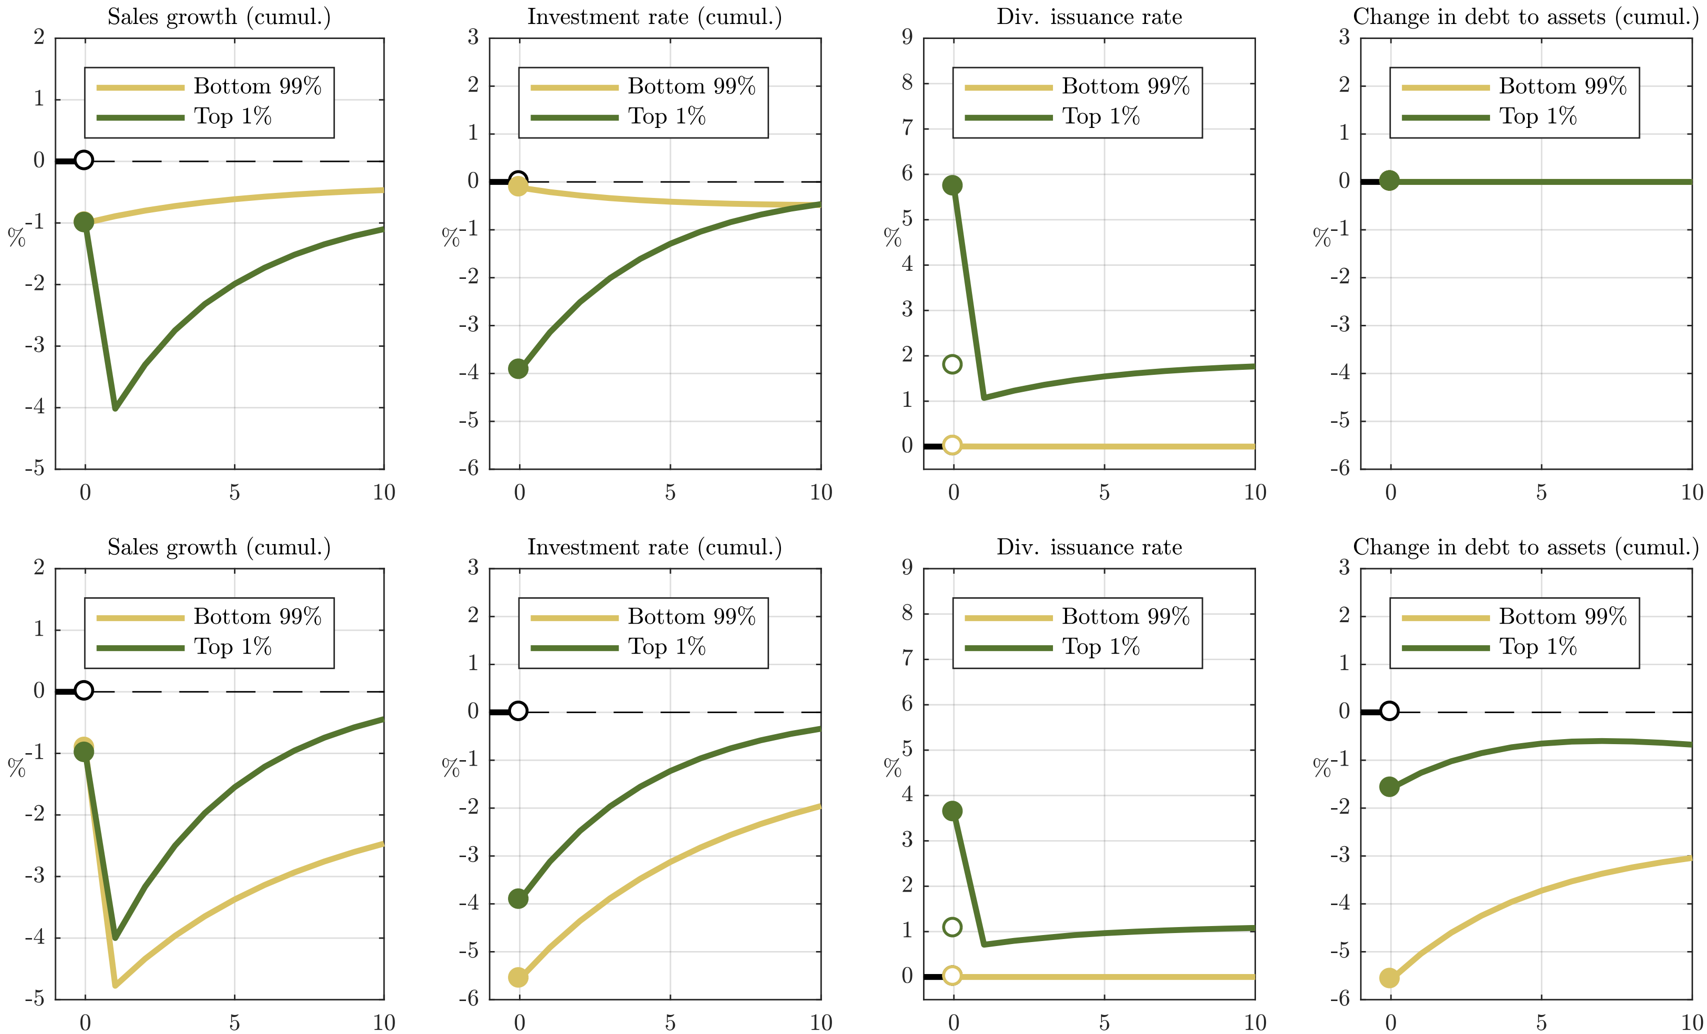

Figure 11: Impulse responses to an aggregate shock in the models of section 5.2. The green lines correspond to firms in the top 1% of the one-quarter lagged distribution of book assets, and the yellow lines correspond to firms in the bottom 99%; book assets in the model are defined as $k_{i,t}$. The top row reports impulse responses in the model with no external financing. The bottom row show the impulse responses in the model with borrowing.

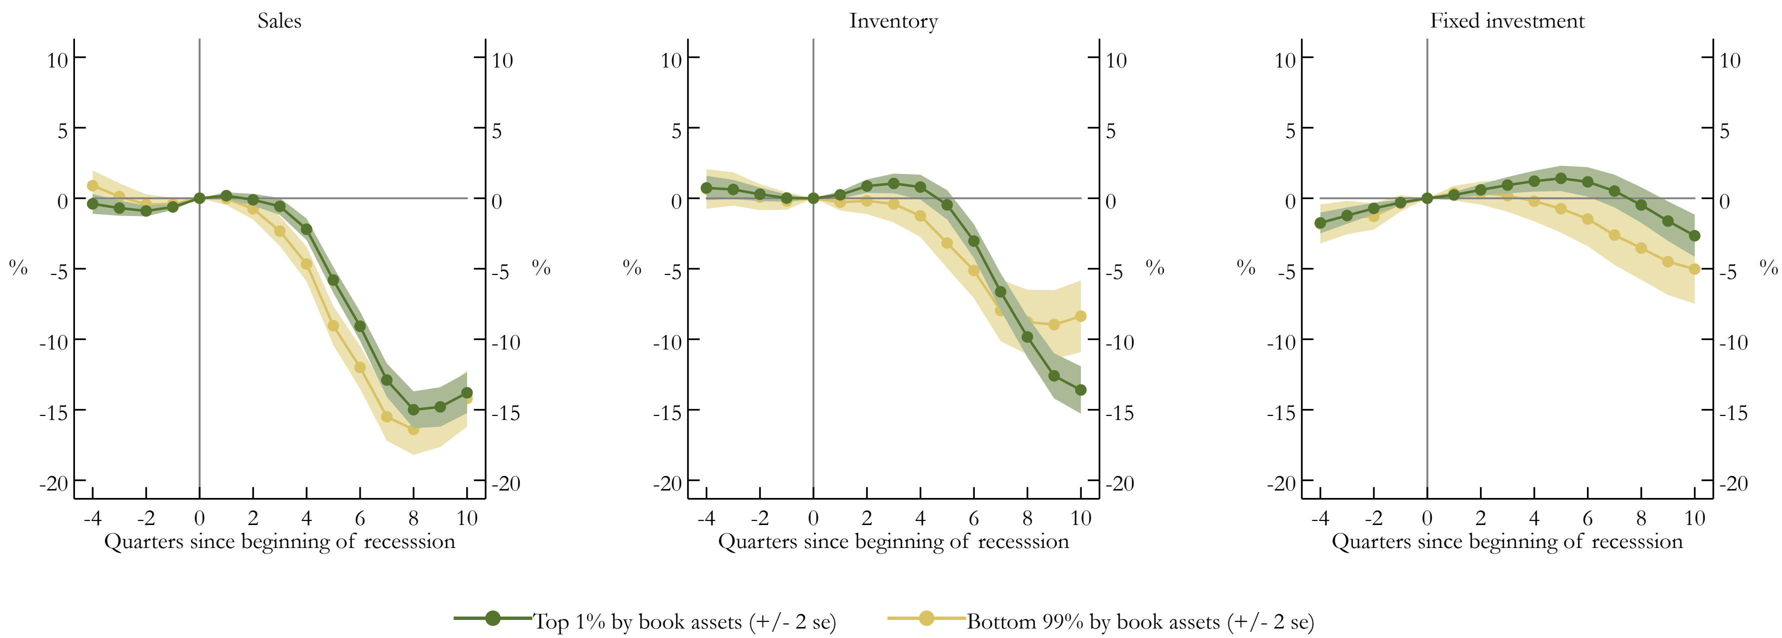

Figure 12: The behavior of sales, inventory and fixed capital after the start of a recession. Each graph reports the cumulative change in a variable of interest after the beginning of a recession; see section 5.2 for details on the estimation. Shaded areas are +/- 2 standard error bands. All growth rates are computed year-on-year and expressed at the quarterly frequency. Recession start dates are 1981q3, 1990q3, 2001q1, and 2007q4.

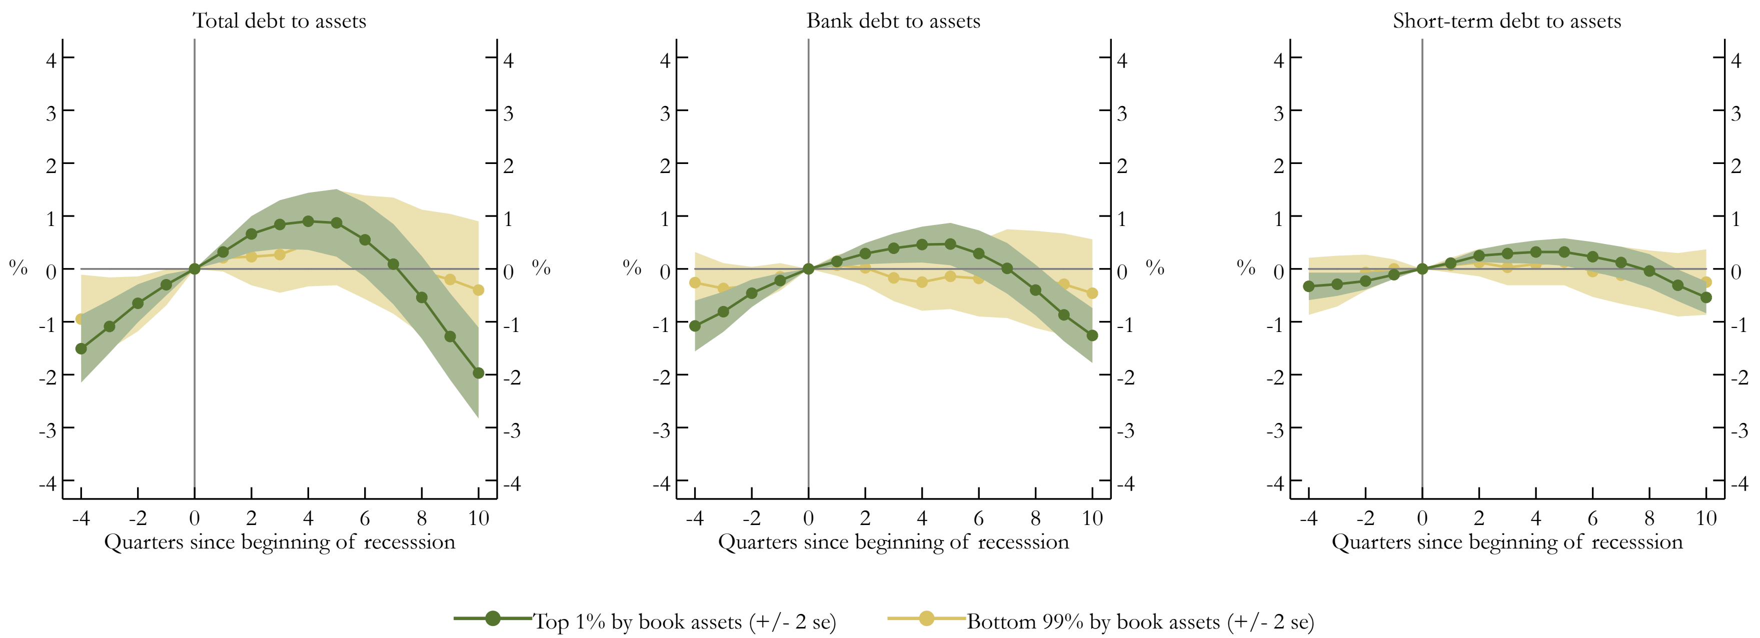

Figure 13: The behavior of debt overall, bank debt, and short-term debt after the start of a recession. Each panel reports changes relative to quarter 0 (the recession start date), computed using the cumulative sum of average growth rate of each size group. Growth rates at the firm-level are computed as $\frac{x_{i,t}-x_{i,t-4}}{assets_{i,t-4}}$, where $x \in \{\text{all debt}, \text{bank debt}, \text{short-term debt}\}$. Size groups are defined with a four-quarter lag.

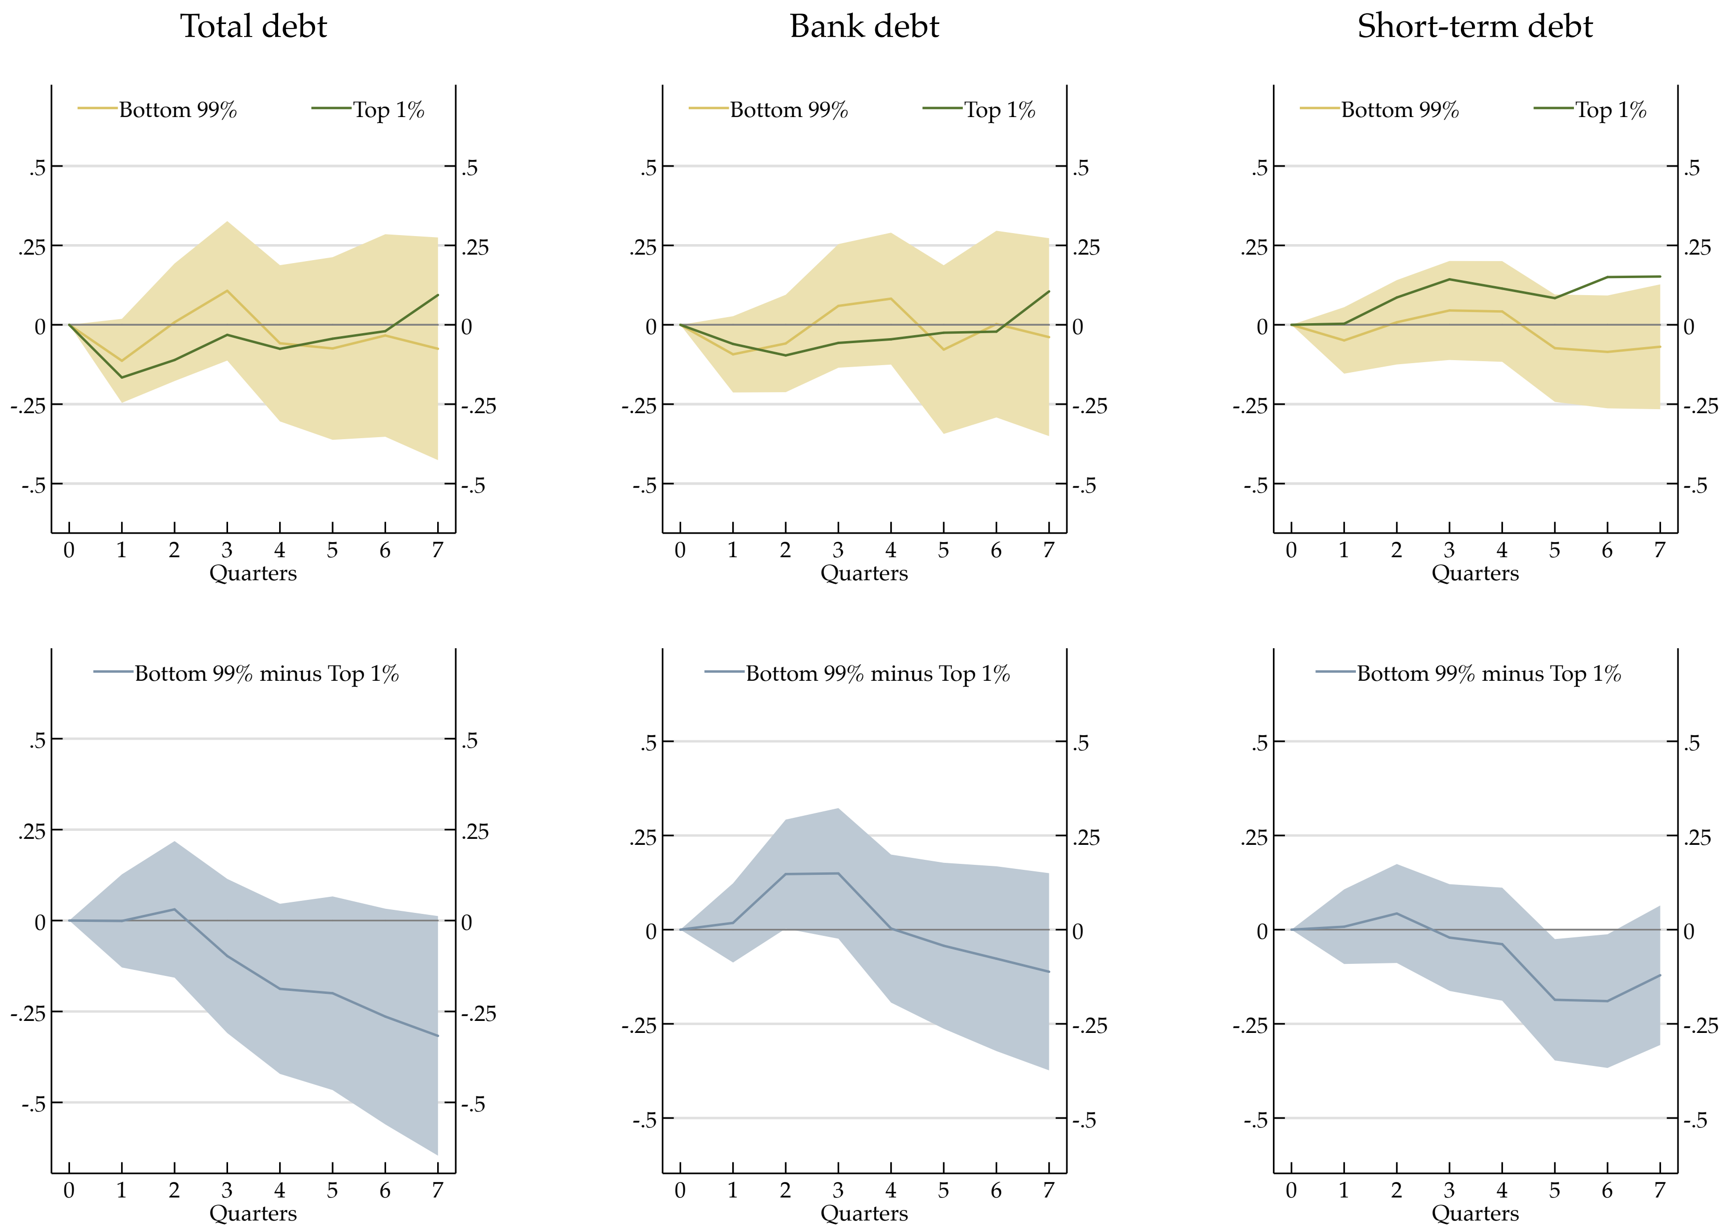

Figure 14: Firm-level response of the ratios of total debt, bank debt and short-term debt to assets to an innovation to the Romer and Romer (2004) shock. The estimated specification is model 2. The top row of graphs reports the average marginal effect at the mean of a one percentage point increase in $rr_{t-1,t}$, for the bottom 99% and top 1% size group. The yellow shaded area is the 95% confidence interval; standard errors are clustered at the firm-level and heteroskedasticity-robust. The bottom row of graphs reports the difference in the OLS coefficients $\beta^{(h)}_{(0,99)} - \hat{\beta}^{(h)}_{(99,100)}$, along with its 95% confidence interval. Data is from 1977q3 to 2007q4.

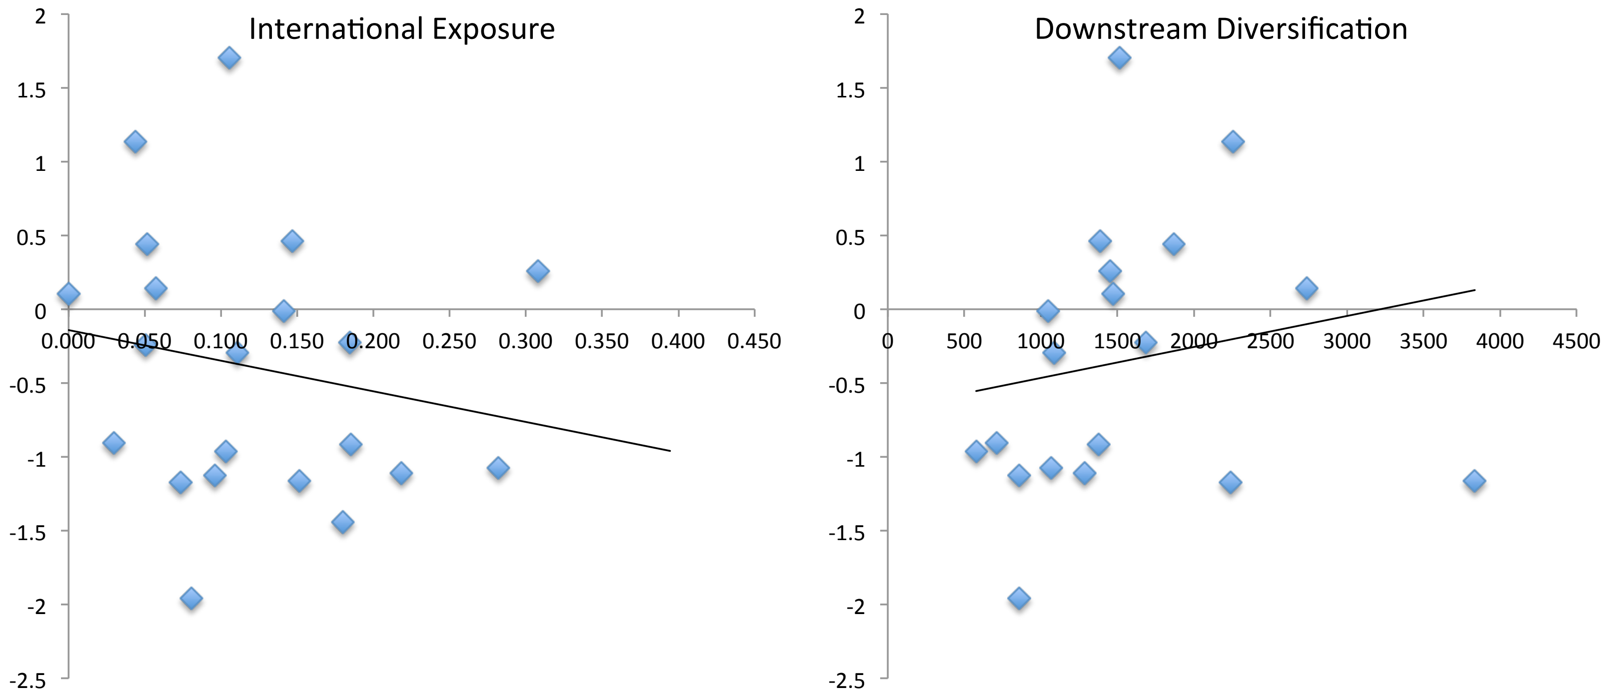

Figure 15: The left-hand panel shows a scatterplot of the size effect by 3-digit industry (y-axis) against the measure of export exposure by 3-digit industry. The size effect is the difference between the elasticity of the top 1% to US GDP growth and the elasticity of the bottom 99% within 3-digit industry. A more negative number indicates that larger firms within a given industry are less sensitive than smaller firms in that industry. Export exposure is measured as total exports divided by total gross output. The right-hand panel shows a scatterplot of the size effect against downstream diversification. Downstream diversification is measured as a Herfindahl index as the share industry X's gross output used by industry Y.

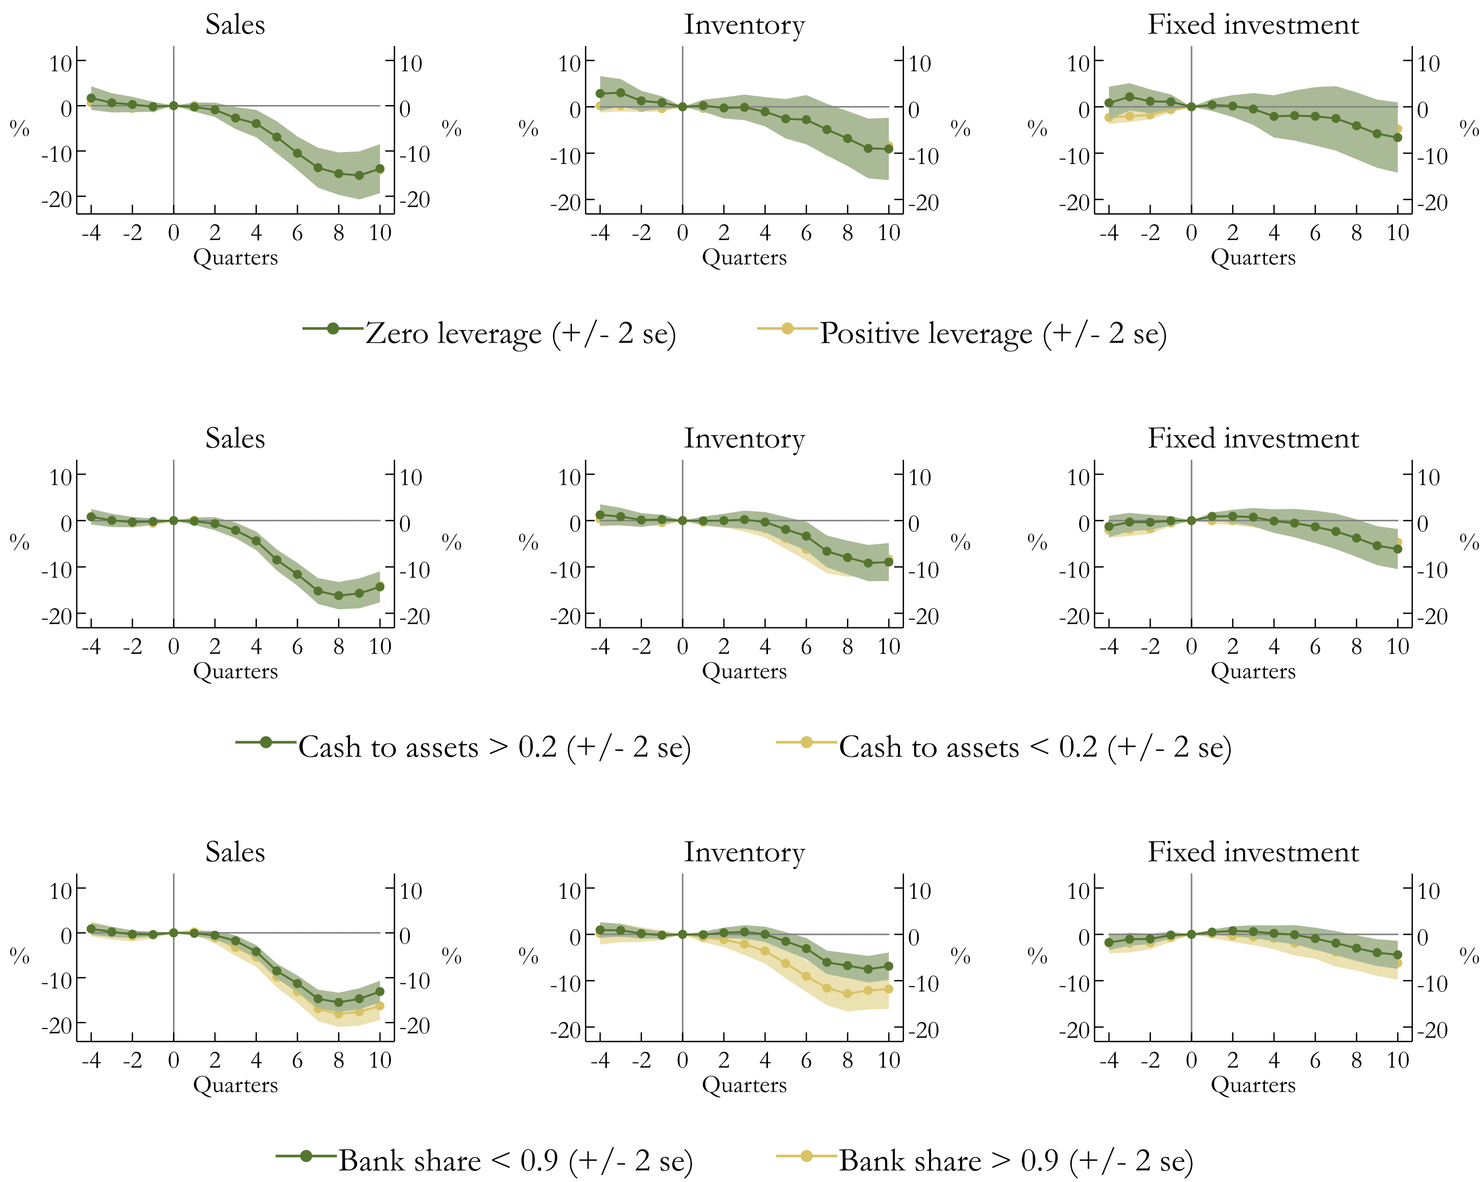

Figure 16: Sales, inventory and fixed capital after the start of a recession, across firms sorted by leverage, liquidity and bank dependence. Each graph reports the cumulative change in a variable of interest after the beginning of a recession; see section 5.2 for details on the estimation. Shaded areas are +/- 2 standard error bands. Variable definitions are given in appendix (A). Top row: firms sorted based on lagged leveraged; middle row: firms sort based on lagged cash-to-asset ratio; bottom row: firms sorted on bank dependence.

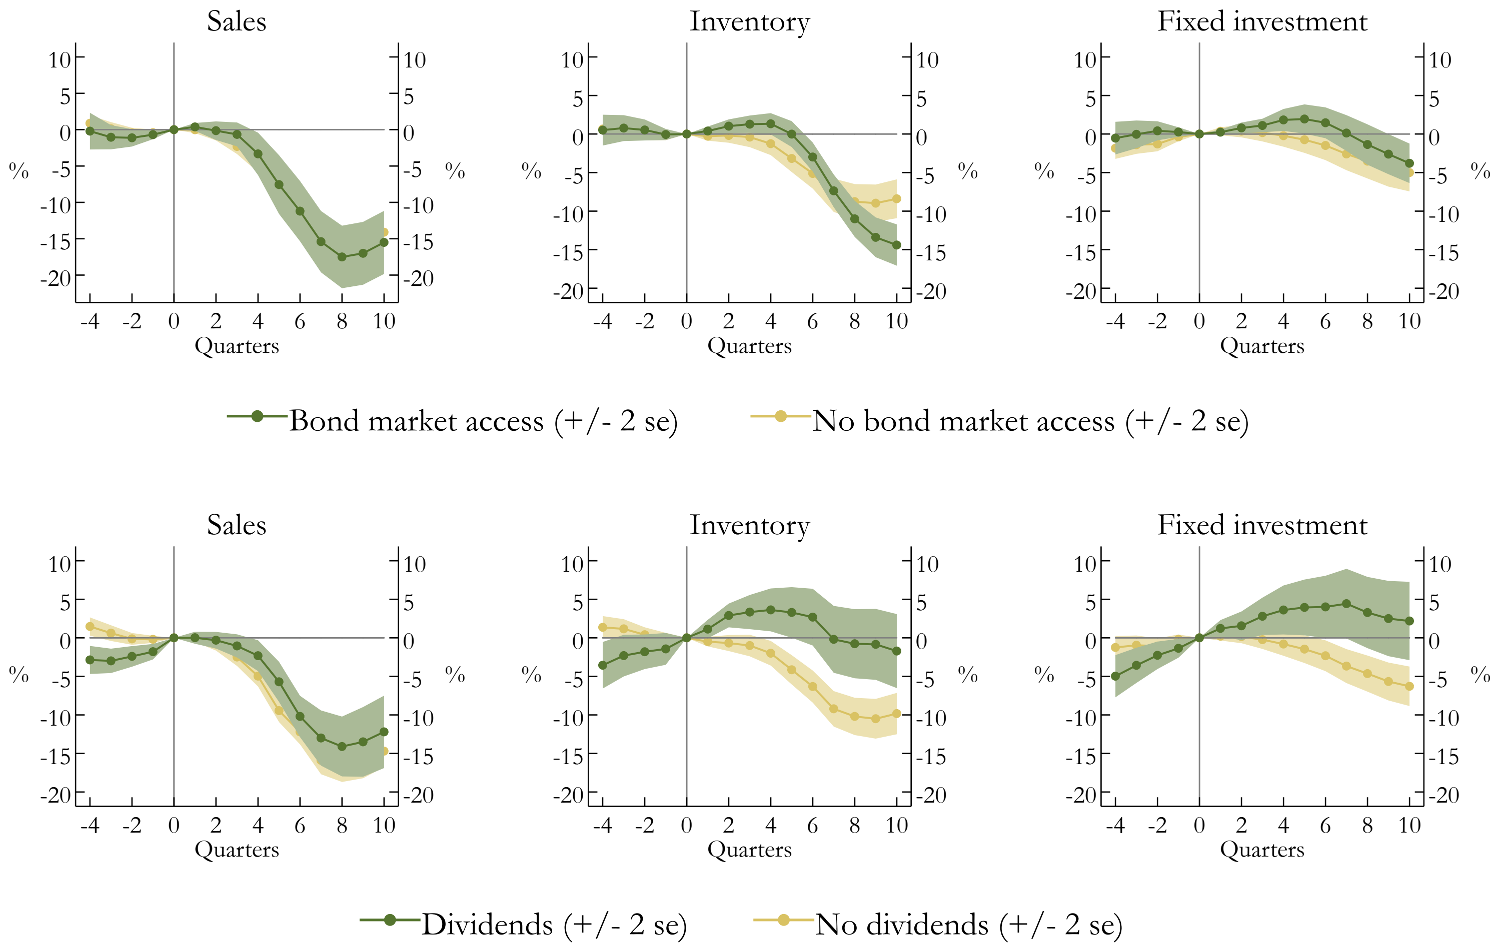

Figure 17: Sales, inventory and fixed capital after the start of a recession, across firms sorted by market access and dividend issuance. Each graph reports the cumulative change in a variable of interest after the beginning of a recession; see section (5.2) for details on the estimation. Shaded areas are +/- 2 standard error bands. Variable definitions are given in appendix (A). Top row: firms sorted based on lagged access to bond market; bottom row: firms sort based on lagged dividend issuance.

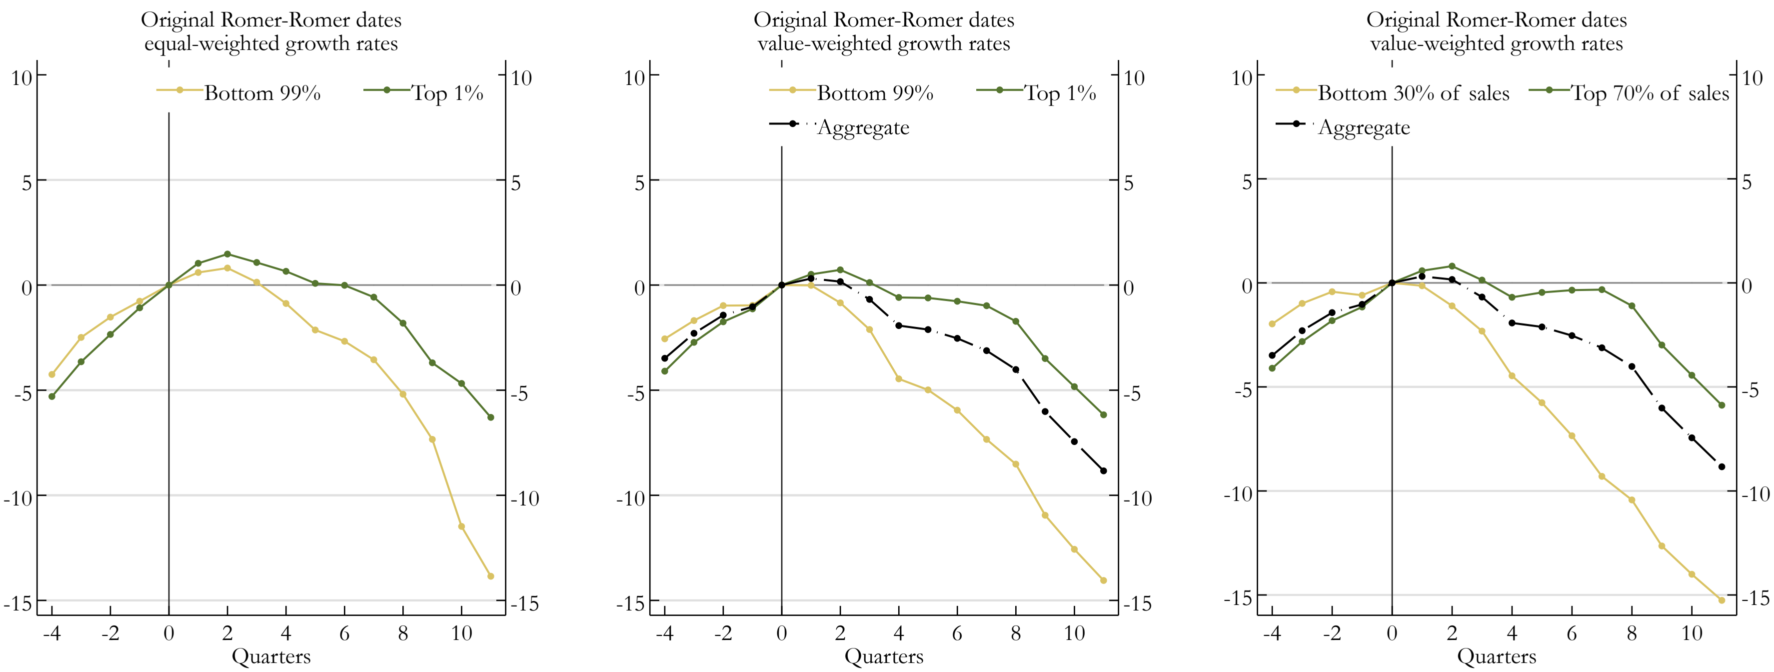

Figure 18: Cumulative change of sales around the three original Romer dates of 1978q3, 1979q4 and 1988q4, using different methodologies. The top left panel uses the equal-weighted growth rates and the (0,99)/(99,100) size classification from the main text. The middle panel uses value-weighted growth rates and the same size classification. The right panel uses the size classification and growth rate construction of Gertler and Gilchrist (1994), as described in appendix D. We use data from 1977q3 to 1990q4, the overlapping portion of our and Gertler and Gilchrist (1994)'s sample. All time-series are de-meaned.

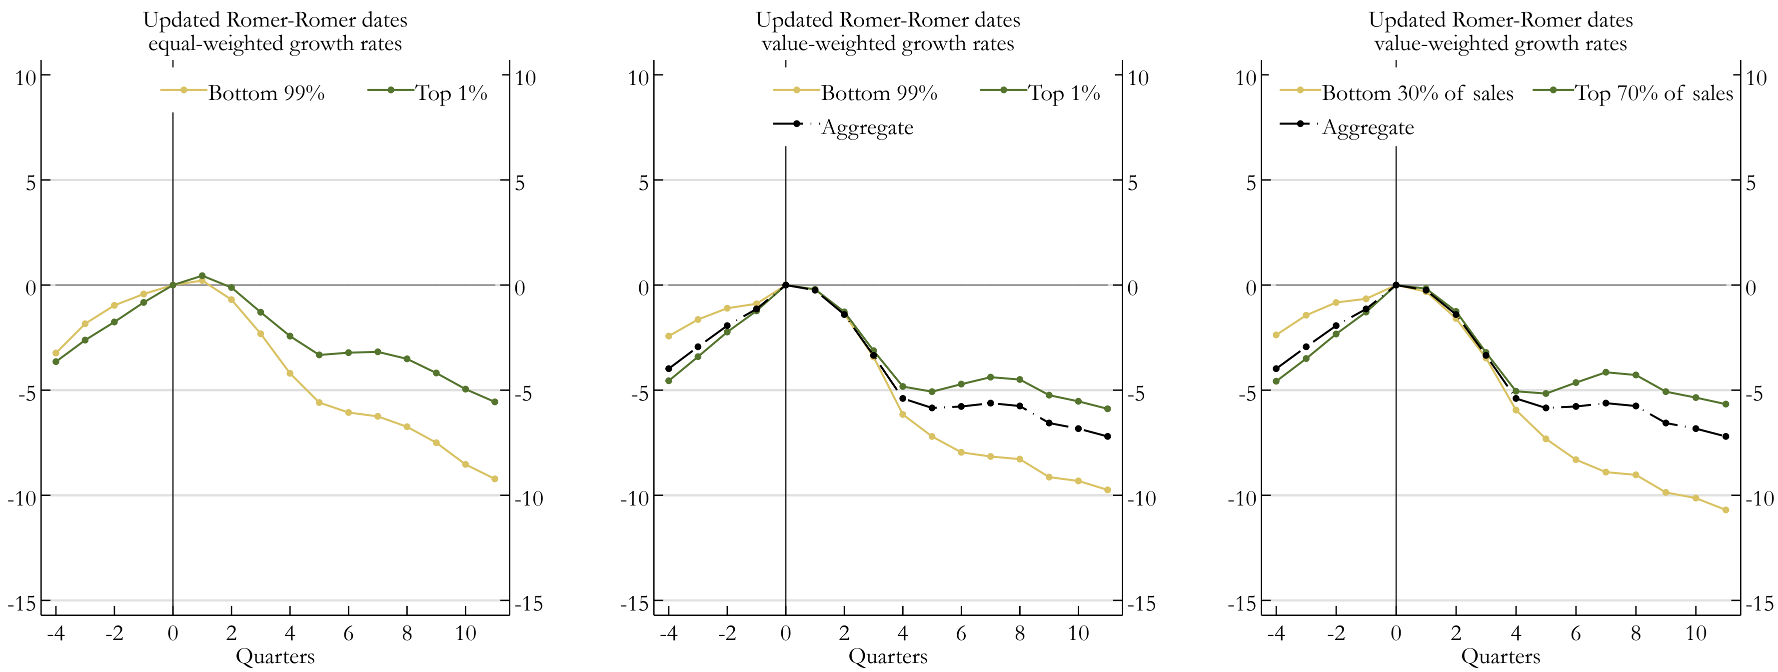

Figure 19: Cumulative change of sales around the five updated Romer dates of 1978q3, 1979q4, 1988q4, 1994q2 and 2008q3, using different methodologies. The top left panel uses the equal-weighted growth rates and the (0,99)/(99,100) size classification from the main text, and is identical to the top left panel of Figure 5. The middle panel uses value-weighted growth rates and the same size classification. The right panel uses the size classification and growth rate construction of Gertler and Gilchrist (1994), as described in appendix D. We use data from 1977q3 to 2014q1. All time-series are de-meaned.

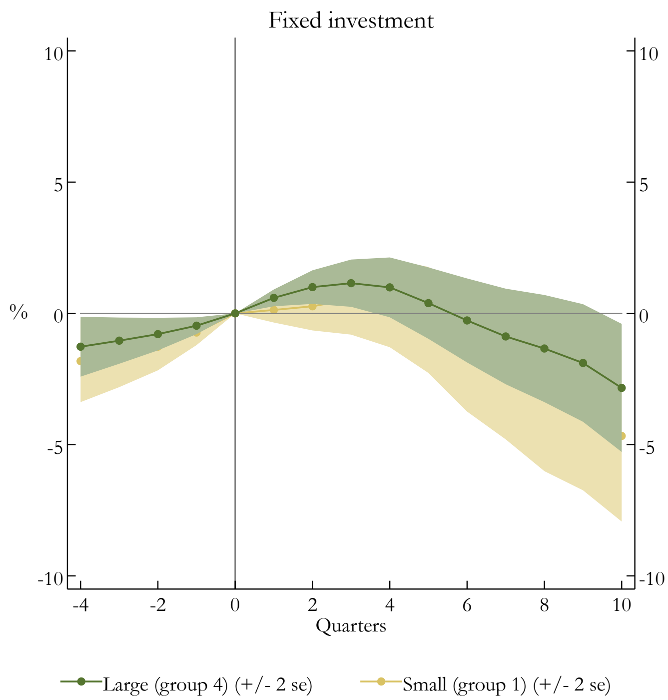

Figure 20: The behavior of fixed investment after the start of a recession in the quarterly Compustat sample. The graph reports the cumulative investment rate relative to the beginning of the recession; see section 5.2 for details on the estimation. Shaded areas are +/- 2 standard error bands. See appendix A for details on the definition of size groups. Recession start dates are 1981q3, 1990q3, 2001q1, and 2007q4.

Panel A: size and growth

Size group	*0-90th*	*90-99th*	*99-99.5th*	*>99.5th*
Assets ($ mil.)	$2.0	$48.8	$626.0	$6766.3
Sales ($ mil., quarterly)	$1.2	$18.8	$181.1	$1420.8
Sales growth (year-on-year)	0.19%	4.58%	4.34%	4.08%
Investment rate (year-on-year)	26.50%	24.91%	21.89%	20.36%

Panel B: financial characteristics

Size group	*0-90th*	*90-99th*	*99-99.5th*	*>99.5th*
Debt to asset ratio	0.35	0.29	0.30	0.28
Cash to asset ratio	0.15	0.10	0.07	0.06
Short-term debt (fraction of total debt)	0.33	0.33	0.20	0.18
Bank debt (fraction of total debt)	0.48	0.57	0.43	0.28
Trade credit (fraction of total liabilities)	0.32	0.27	0.17	0.13
Intangible assets (fraction of total assets)	0.05	0.11	0.26	0.36
Zero leverage (% of tot. firm-quarter obs.)	20%	13%	8%	3%
Negative book equity (% of tot. firm-quarter obs.)	5%	<1%	<1%	<1%
Bank dependent (% of tot. firm-quarter obs.)	26%	29%	20%	10%

Table 1: Real and financial firm characteristics, by size group. Assets and sales are averages from 1977q1 to 2014q1 within category expressed in real 2009 dollars; values are deflated using the price index for value added in manufacturing, available from the Bureau of Economic Analysis at http://bea.gov/industry/gdpbyind_data.htm. All other variables are ratios as described in the main text. See Appendix A for details on the construction of size groups.

Size group	*0-90th*	*90-99th*	*99-99.5th*	*>99.5th*
Assets				
Financial assets, incl. cash	0.149	0.099	0.074	0.055
Short-term assets				
Receivables	0.284	0.229	0.165	0.124
Inventory	0.218	0.241	0.172	0.130
Other	0.040	0.037	0.042	0.041
Long-term assets				
Net property, plant and equipment	0.269	0.288	0.289	0.287
Other, incl. intangibles	0.050	0.106	0.259	0.362
Liabilities				
Debt				
Due in 1 year or less				
Bank debt	0.083	0.083	0.032	0.016
Non-bank debt	0.035	0.019	0.019	0.028
Due in more than 1 year				
Bank debt	0.107	0.111	0.110	0.072
Non-bank debt	0.123	0.079	0.141	0.179
Trade payables	0.156	0.123	0.085	0.071
Other, incl. capital leases	0.099	0.121	0.187	0.233
Equity	0.393	0.463	0.426	0.416

Table 2: Average balance sheet, by size group. All numbers are expressed as fraction of total book assets. Fractions may not add up to 1 due to rounding. Financial assets are the sum of cash and deposits, treasury and federal agency securities, and all other financial assets. Other short-term assets include pre-paid expenses and income taxes receivable. Non-bank debt inculdes commercial paper, bonds, and other short- and long-term notes. Other liabilities include tax liabilities and capital leases. Definitions of the variables in terms of QFR items from survey forms 300, 201, and 200 are available upon the authors on request. See Appendix A for details on the construction of size groups.

Size group	*0-90th*	*90-99th*	*99-99.5th*	*>99.5th*
Sales growth, $<$ p25	-26.27%	-16.59%	-12.66%	-10.97%
Sales growth	0.19%	4.58%	4.34%	4.08%
Sales growth, $>$ p75	26.77%	25.83%	21.41%	19.19%
Leverage, $<$ p25	0.00	0.01	0.04	0.07
Leverage	0.35	0.29	0.30	0.28
Leverage, $>$ p75	0.47	0.39	0.39	0.36
Liquidity, $<$ p25	0.00	0.00	0.00	0.00
Liquidity	0.15	0.10	0.07	0.06
Liquidity, $>$ p75	0.20	0.13	0.10	0.07

Table 3: Approximate inter-quartile ranges for selected variables, by firm size group. All variables are averages from 1977q1 to 2014q1 within size group. Leverage is defined as the ratio of debt to assets, while liquidity is defined as the ratio of cash to assets. See Appendix A for details on the construction of the size groups. Exact percentiles are not reported in order to preserve data confidentiality.

	Sales	Inventory	Fixed investment
GDP growth	3.700***	2.065***	0.912***
	(0.000)	(0.000)	(0.000)
$[90, 99] \times$ GDP growth	-0.160	-0.107	-0.299*
	(0.260)	(0.538)	(0.057)
$[99, 99.5] \times$ GDP growth	-0.251*	-0.299*	-0.687***
	(0.080)	(0.097)	(0.000)
$[99.5, 100] \times$ GDP growth	-0.600***	-0.730***	-1.257***
	(0.000)	(0.000)	(0.000)
N	≈ 460000	≈ 460000	≈ 460000
nr. firms	≈ 60000	≈ 60000	≈ 60000
adj. R^2	0.025	0.006	0.003
industry controls	yes	yes	yes
s.e. clustering	firm-level	firm-level	firm-level

Table 4: Estimate of regression of the baseline model (1) for sales growth, inventory growth, and the fixed investment rate. See Appendix A for details on the construction of the dependent variable and size groups. The investment rate is computed as $\frac{nppe_{i,t} - nppe_{i,t-4} + dep_{i,t-4,t}}{nppe_{i,t-4}}$, where $dep_{i,t-4,t}$ is cumulative reported depreciation between $t-4$ and t. All values are deflated by the quarterly manufacturing price index. $*$, $**$ and $***$ indicate 10%, 5% and 1% significance levels, respectively with p-values reported in parentheses.

	Small firms	Large firms	All firms
$corr(G_t, Y_t)$	0.68	0.62	0.65
$\frac{\sigma_{\hat{g}_t}}{\sigma_{G_t}}$	1.02	0.83	0.89
$corr(\hat{g}_t, Y_t)$	0.84	0.77	0.80
$\frac{\sigma_{c\hat{o}v_t}}{\sigma_{G_t}}$	0.54	0.45	0.41
$corr(c\hat{o}v_t, Y_t)$	-0.32	-0.05	-0.15

Table 5: Decomposition of the correlations of aggregate sales growth among all firms, small firms, and large firms, to GDP growth. See section 4.2 for details on the decomposition.

	Actual β	Counterfactual 1 $\beta^{(1)}$	Counterfactual 2 $\beta^{(2)}$	Counterfactual 3 $\beta^{(3)}$
Sales	2.293	2.154	2.270	1.701
	(0.342)	(0.342)	(0.366)	(0.298)
Inventory	0.919	0.719	0.770	0.552
	(0.226)	(0.250)	(0.226)	(0.167)
Fixed investment	0.584	0.569	0.569	0.506
	(0.145)	(0.151)	(0.148)	(0.135)
Total assets	0.876	0.787	0.838	0.692
	(0.121)	(0.129)	(0.119)	(0.105)
Observations	143	143	143	143

Table 6: Cyclical sensitivities of aggregate sales, inventory, fixed investment, and total assets. Each line reports the estimated slope in regressions of the form $G_t = \alpha + \beta \log\left(\frac{GDP_t}{GDP_{t-4}}\right) + \epsilon_t$, where G_t is an aggregate growth rate of sales, inventory, fixed investment, or total assets. The first column uses the actual time series G_t; the second column, the counterfactual time series $G_t^{(1)}$; the third column, the counterfactual time series $G_t^{(2)}$; and the fourth column, the counterfactual time series $G_t^{(3)}$. Heteroskedasticity robust standard errors in parentheses.

	Baseline	(1)	(2)	(3)	(4)	(5)
$[90, 99] \times$ GDP growth	-0.160	-0.189	-0.195	-0.162	-0.193	-0.176
$[99, 99.5] \times$ GDP growth	-0.251^{*}	-0.257^{*}	-0.321^{**}	-0.282^{*}	-0.490^{***}	-0.247
$[99.5, 100] \times$ GDP growth	-0.600^{***}	-0.563^{***}	-0.675^{***}	-0.640^{***}	-1.097^{***}	-0.594^{***}
Bank share [0.10,0.90] $\times$ GDP growth		0.300				
Bank share $< 0.10 \times$ GDP growth		-0.315				
Leverage [0.15,0.50] $\times$ GDP growth			-0.126			
Leverage (0,0.15] $\times$ GDP growth			-0.474^{*}			
Leverage $= 0 \times$ GDP growth			-0.630^{**}			
Liquidity [0.01,0.20] $\times$ GDP growth				0.228		
Liquidity $> 0.20 \times$ GDP growth				-0.101		
Market access $\times$ GDP growth					0.826^{**}	
Dividend issuance $\times$ GDP growth						0.087
N	≈ 460000	≈ 460000	≈ 460000	≈ 460000	≈ 460000	≈ 460000
nr. firms	≈ 60000	≈ 60000	≈ 60000	≈ 60000	≈ 60000	≈ 60000
adj. R^2	0.025	0.025	0.025	0.025	0.025	0.025
industry controls	yes	yes	yes	yes	yes	yes
s.e. clustering	firm-level	firm-level	firm-level	firm-level	firm-level	firm-level

Table 7: Estimate of regression of the regression model (9) for sales growth on firm size and proxies for financial constraints. Each column is a separate regression. All coefficients are reported relative to a baseline group; for size, the baseline group is the $[0, 90]$ group. See Appendix A for details on the construction of the dependent variable and size groups, and text for description of each proxy for financial constraints. Standard errors clustered at the firm level. $*$, $**$ and $***$ indicate 10%, 5% and 1% significance levels, respectively with p-values reported in parentheses.

Panel A	**Baseline**	**Bank dependence**			**Leverage**			**Liquidity**		
		Low	High	Diff	Low	High	Diff	High	Low	Diff
$[90, 99] \times$ GDP growth	-0.160	-0.094	-0.286	-0.192	-0.195	0.010	0.205	-0.257	0.179	0.436
$[99, 99.5] \times$ GDP growth	-0.251^{*}	-0.270	-0.184	0.085	-0.182	-0.570	-0.387	-0.384^{**}	0.042	0.376
$[99.5, 100] \times$ GDP growth	-0.600^{***}	-0.616^{***}	-0.429^{**}	0.187	-0.530^{***}	-0.974^{***}	-0.444	-0.604^{***}	-0.482^{*}	0.122
N	≈ 460000		≈ 460000			≈ 460000			≈ 460000	
nr. firms	≈ 60000		≈ 60000			≈ 60000			≈ 60000	
adj. R^2	0.025		0.025			0.025			0.025	
industry controls	yes		yes			yes			yes	
s.e. clustering	firm-level		firm-level			firm-level			firm-level	

Panel B	**Baseline**	**Bond market access**			**Dividend issuance**		
		Yes	No	Diff	High	low	Diff
$[90, 99] \times$ GDP growth	-0.160				-0.178	-0.146	0.032
$[99, 99.5] \times$ GDP growth	-0.251^{*}	-2.907	-0.421^{***}	2.486	-0.150	-0.385	-0.235
$[99.5, 100] \times$ GDP growth	-0.600^{***}	-3.568^{*}	-0.832^{***}	2.736	-0.395^{**}	-0.758^{*}	-0.363^{**}
N	≈ 460000		≈ 460000			≈ 460000	
nr. firms	≈ 60000		≈ 60000			≈ 60000	
adj. R^2	0.025		0.025			0.025	
industry controls	yes		yes			yes	
s.e. clustering	firm-level		firm-level			firm-level	

Table 8: Triple-interaction regressions. The dependent variable is sales growth. The columns marked baseline, bank dependence, leverage, liquidity, bond market access, and dividend issuance each correspond to one regression. For each financial indicator, coefficients are reported by sub-groups corresponding to firms which are less likely to be financially constrained (left column) and firms which are more likely to be financially constrained (middle column). The coefficients shown are differences in elasticities to GDP growth relative to the $[0, 90]$ size group *within each sub-group*. That is, the coefficient -0.616 in the right column of the bank dependence regression indicates that *within* firms with low bank dependence, the top 0.5% of firms has an elasticity of sales growth -0.616 points smaller than the $[0, 90]$ group. The last column, denoted Diff, reports the difference across groups of the size effect, as well as its significance level. The Bond market access regressions only compare the $[0, 99]$ group to others in order to avoid violating disclosure limits as there are too few observations with a bond issuance history in the $[0, 90]$ size group. See Appendix A for details on the construction of the dependent variable and size groups, and text for description of each proxy for financial constraints. Standard errors clustered at the firm level. $*$, $**$ and $***$ indicate 10%, 5% and 1% significance levels, respectively with p-values reported in parentheses.

	Small firms	Large firms	All firms
$corr(\tilde{G}_t, Y_t)$	0.68	0.61	0.65
$\frac{\sigma_{\hat{g}_t}}{\sigma_{\tilde{G}_t}}$	0.97	0.81	0.86
$corr(\hat{\tilde{g}}_t, Y_t)$	0.84	0.77	0.80
$\frac{\sigma_{c\hat{o}v_t}}{\sigma_{\tilde{G}_t}}$	0.51	0.45	0.41
$corr(c\hat{\tilde{o}}v_t, Y_t)$	-0.26	-0.04	-0.10

Table 9: Decomposition of the correlations of aggregate sales growth among all firms, small firms, and large firms, to GDP growth, constructed using DHS growth rates. See section C for details on the decomposition.

	Actual β	Counterfactual 1 $\beta^{(1)}$	Counterfactual 2 $\beta^{(2)}$	Counterfactual 3 $\beta^{(3)}$
Sales	2.285	2.174	2.263	1.696
	(0.339)	(0.339)	(0.362)	(0.296)
Inventory	0.918	0.758	0.768	0.554
	(0.225)	(0.250)	(0.226)	(0.184)
Fixed investment	0.583	0.576	0.567	0.507
	(0.145)	(0.151)	(0.148)	(0.135)
Total assets	0.876	0.791	0.857	0.694
	(0.121)	(0.129)	(0.745)	(0.107)
Observations	143	143	143	143

Table 10: Cyclical sensitivities of aggregate sales, inventory, fixed investment, and total assets using DHS growth rates. Each line reports the estimated slope in regressions of the form $\tilde{G}_t = \alpha + \beta \log\left(\frac{GDP_t}{GDP_{t-4}}\right) + \epsilon_t$, where G_t is an aggregate growth rate of sales, inventory, fixed investment, or total assets. The first column uses the actual time series G_t; the second column, the counterfactual time series $\tilde{G}_t^{(1)}$; the third column, the counterfactual time series $\tilde{G}_t^{(2)}$; and the fourth column, the counterfactual time series $\tilde{G}_t^{(3)}$. All these series are constructed as in (4), but using the decomposition of aggregate growth based on DHS growth rates rather than log growth rates. Heteroskedasticity robust standard errors in parentheses.

	Size group			
	(1)	**(2)**	**(3)**	**(4)**
Assets (2009 m$)				
QFR	2.0	48.8	626.0	6766.3
Compustat (annual)	22.6	94.4	375.7	7348.9
Compustat (quarterly)	23.6	102.3	409.6	7989.8
Investment rate				
QFR	26.50%	24.91%	21.89%	20.36%
Compustat (annual)	30.93%	32.00%	27.94%	21.87%
Compustat (quarterly)	28.83%	31.89%	28.90%	22.69%

Table 11: Summary statistics for the QFR sample and the two Compustat samples. Each column corresponds to a different size group. For QFR data, size groups are defined as in the main text. For Compustat (annual and quarterly), size groups are quartiles of the distribution of book assets (variable `at` in the annual files and `atq` in the quarterly files). Assets are nominal book values deflated by the BEA price deflator for manufacturing value added, as in the main text. See appendix E for details on the construction of the annual and quarterly Compsutat samples and the computation of investment rates.

	QFR	Compustat (annual)	Compustat (quarterly)
GDP growth	0.912***	1.082***	0.537***
	(0.000)	(0.000)	(0.002)
Size group 2 × GDP growth	−0.299*	−0.235	−0.103
	(0.057)	(0.234)	(0.644)
Size group 3 × GDP growth	−0.687***	−0.329*	−0.250
	(0.000)	(0.082)	(0.248)
Size group 4 × GDP growth	−1.257***	−0.921***	−0.572***
	(0.000)	(0.000)	(0.004)
N	≈ 460000	72363	186784
nr. firms	≈ 60000	6550	5944
adj. R^2	0.003	0.022	0.017
industry controls	yes	yes	yes
s.e. clustering	firm-level	firm-level	firm-level

Table 12: Investment cyclicality by size in the QFR data (first column) and for the annual and quarterly Compustat samples (second and third columns. The baseline coefficient (first line) refers to firms in the durable sector. All values are deflated by the quarterly manufacturing price index. $*$, $**$ and $***$ indicate 10%, 5% and 1% significance levels, respectively with p-values reported in parentheses.