

Staff Working Paper/Document de travail du personnel 2016-51

Firm-Specific Shocks and Aggregate Fluctuations



by Leonid Karasik, Danny Leung and Ben Tomlin

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Bank of Canada Staff Working Paper 2016-51

November 2016

Firm-Specific Shocks and Aggregate Fluctuations

by

Leonid Karasik,¹ Danny Leung² and Ben Tomlin³

¹leo.karasik@gmail.com

²Statistics Canada
danny.leung@canada.ca

³Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
btomlin@bankofcanada.ca

Acknowledgements

We thank Isabelle Amano, Greg Bauer, Carlos Rosell and Alexander Ueberfeldt for their comments. The results have been institutionally reviewed to ensure that no confidential information is revealed. The views expressed in this paper are those of the authors and no responsibility for them should be attributed to Statistics Canada or the Bank of Canada.

Abstract

In order to understand what drives aggregate fluctuations, many macroeconomic models point to aggregate shocks and discount the contribution of firm-specific shocks. Recent research from other developed countries, however, has found that aggregate fluctuations are in part driven by idiosyncratic shocks to large firms. Using data on Canadian firms, this paper examines the contribution of large firms to industry-level fluctuations in gross output, investment and employment in the manufacturing sector. The data suggest that shocks to large firms can explain as much as 46% and 37% of the fluctuations in gross output and investment, respectively, but do not contribute to fluctuations in employment.

Bank topics: Business fluctuations and cycles; Firm dynamics; Market structure and pricing

JEL codes: E22; E23; E24; E3; L6

Résumé

Pour comprendre ce qui entraîne les fluctuations agrégées, de nombreux modèles macroéconomiques se tournent vers les chocs agrégés sans tenir compte de la contribution des chocs propres à l'entreprise. De récentes recherches menées par d'autres pays développés ont, toutefois, permis de constater que les fluctuations agrégées sont en partie attribuables aux chocs encaissés par les grandes entreprises. À l'aide de données sur des entreprises canadiennes, cet article examine la contribution des grandes entreprises aux fluctuations sectorielles de la production brute, de l'investissement et de l'emploi dans le secteur de la fabrication. Les données laissent supposer que les chocs encaissés par les grandes entreprises peuvent expliquer jusqu'à 46 % et 37 % des fluctuations de la production brute et de l'investissement, respectivement, sans toutefois contribuer aux fluctuations de l'emploi.

Sujets : Cycles et fluctuations économiques ; Dynamique des entreprises ; Structure de marché et fixation des prix

Codes JEL : E22; E23; E24; E3; L6

Non-Technical Summary

Motivation and Question

It is often assumed that shocks specific to individual firms cancel out in the aggregate and therefore do not contribute to aggregate fluctuations. Recent research, however, suggests that this may not be the case if markets are dominated by a small number of large firms. Given this context, this paper explores the extent to which shocks to individual firms contribute to annual fluctuations in sales, investment and employment growth in the Canadian manufacturing sector.

Methodology

Using firm-level data on the Canadian manufacturing sector from 2000-2012, firm-specific shocks to sales are defined as the difference between an individual firm's growth rate and the growth rate of the average of other firms in the same industry. The weighted average of these firm-specific shocks to sales is called the granular shock. The contribution of the granular shock to aggregate fluctuation in industry sales is obtained from regressing industry sales on the granular shock and examining the explanatory power. A similar approach is used to construct granular shocks for investment and employment, and examine their explanatory power. Several robustness checks are done to explore different assumptions in the calculation of the firm-level shocks (which provide ranges for the estimated effects of granular shocks) and to explore the time series versus cross-sectional components of the results.

Key Results and Contributions

The paper finds that within this framework, idiosyncratic shocks to large firms can account for at least 23% to 46% of the annual variation in manufacturing sales over the 2000-2012 period, and at least 15% to 40% of investment growth volatility. The findings for employment are inconclusive (which may be the result of the unique nature of firm-level employment adjustment). These results provide evidence that shocks to a relatively small number of large firms can be responsible for a significant share of annual variation in some important macroeconomic variables. This suggests that examining shocks to large firms can shed light on the evolution of macroeconomic variables over time and help forecast future growth.

1. Introduction

Do shocks to individual firms play a role in driving aggregate fluctuations? The conventional thinking, and, indeed, the assumption in many macroeconomic models, is that shocks to individual firms wash out in the aggregate.¹ Recent research, however, has revealed that firm-specific shocks do contribute to business-cycle fluctuations in markets dominated by a small number of very large firms (i.e., “granular” markets). Consider, for instance, the situation in the United States. Gabaix (2011) documented that nearly 30% of the U.S. gross domestic product (GDP) is generated by just 100 very large firms. His estimations show that shocks to these firms—or granular shocks, using his parlance—explain nearly half of the annual fluctuations in the U.S. GDP from 1952 to 2008.² The United States is not unique in this regard. The French economy is also dominated by a small number of very large firms, as documented by di Giovanni, Levchenko and Mejean (2014, hereafter DLM). Far from cancelling out in the aggregate, firm-specific shocks account for 80% of the annual variation in the aggregate sales of French firms from 1992 to 2007.

The work undertaken by Gabaix and DLM is part of a strand of research that seeks to uncover the micro origins of business cycles. Parallel research on that front has also found that business cycles are typically driven by industry-level fluctuations (Horvath 1998; Conley and Dupor 2003). As discussed by Foerster, Sarte and Watson (2011, hereafter FSW), aggregate fluctuations can be driven by movements in relatively small industries. To understand the causes of aggregate volatility, it is therefore necessary to obtain insights into the causes of industry-level fluctuations. This paper examines the extent to which idiosyncratic firm shocks—measured as the deviation of the growth rate of an individual firm from the average of other firms in the same industry—contribute to annual industry-level fluctuations in gross output, investment and employment in Canada’s manufacturing sector.³ Using firm-level data from 2000 to 2012, this paper finds that firm-specific shocks can account for 23% to 46% of the annual variation in gross output (i.e., sales) and 13% to 37% of the annual variation in investment at the industry level.⁴ Due to the high degree of idiosyncrasy in firm-level employment adjustment in response to economic shocks, it cannot be determined whether firm-specific shocks can explain variation

¹The assumption that firm-specific shocks disappear in the aggregate dates back to at least Lucas (1977). Examples of firm-specific shocks include labour strikes at a particular plant, a delay in obtaining inputs from a supplier, the adoption of new production techniques or inventory management methods, and managerial turnover. The standard assumption in many macroeconomic models is that business cycles are driven by aggregate shocks that affect most, if not all, firms.

²See Table II in Gabaix (2011).

³In the Canadian context, previous research by Leung, Rispoli and Chan (2012) found that while large firms (i.e., firms with more than 500 employees) comprise less than 1% of the total number of firms, they generate nearly half of the business-sector GDP. There is therefore reason to believe that changes in macroeconomic variables in Canada are (at least in part) caused by firm-specific shocks.

⁴These ranges are the result of the different assumptions used to estimate firm-specific shocks. Please refer to Section 3 for greater detail.

in industry-level employment.⁵

These results highlight the important role that firm-specific shocks play in generating movements in some—though possibly not all—macroeconomic variables. This constitutes an important shift in understanding the origin of aggregate fluctuations in Canada. It suggests that examining the activities of a small number of large firms can give policy-makers and economic forecasters a great deal of information about the current and future state of the economy.

The next section of this article describes the data that are used in this study. Section 3 discusses the means by which firm-specific shocks are measured, while Section 4 presents the results. Section 5 offers concluding remarks, along with a discussion about future research on this topic.

2. Data

This paper uses annual data on the operations of Canadian firms from 2000 to 2012. The data were obtained from Statistics Canada’s T2-LEAP database. Employment data were taken from the Longitudinal Employment Analysis Program (LEAP) database and were merged with financial data from the T2 tax filings of firms with the Canada Revenue Agency. The T2 file contains firm income statements, balance sheets and investments in tangible assets. For more detail on the data see Lafrance and Gu (2014) and Lafrance (2013). Firm-level data on gross output (measured by sales) and investment (measured by the sales cost of acquisitions of structures and machinery and equipment) are obtained from the T2 file. While the T2-LEAP database actually extends back to 1984, not all variables are available for the entire time period. For example, data for investment are available only from 2001 onwards.

While the T2-LEAP database covers the entire economy, attention is restricted to the manufacturing sector. It is a sector for which the aggregate data on employment, gross output (with total sales as a proxy), and investment in tangible capital from the T2-LEAP database have been confirmed to follow broadly the annual fluctuations of the published aggregates for the sector. It is therefore well suited for an examination of the sources of aggregate volatility in Canada. The manufacturing sector is composed of 86 industries classified at the 4-digit level of the North American Industry Classification System (NAICS).

It is tempting to think that aggregate fluctuations in manufacturing are driven primarily by fluctuations in a few large industries, such as automobile manufacturing, or food and beverage manufacturing. However, the results in FSW suggest that this is not the case for the United States. To provide a sense of the data and obtain greater insight into the origins of fluctuations

⁵This is expanded upon in Subsection 4.3.

within Canada’s manufacturing sector, the remainder of this section replicates the aggregate growth decomposition outlined in FSW. Let $r_{j,t}$ denote the log growth rate of a variable of interest for industry j in year t , and R_t denote the aggregate log growth rate of that variable for the entire manufacturing sector. The aggregate growth rate is simply the weighted average growth rates of the 86 industries:

$$R_t = \sum_{j=1}^{86} \tau_{t-1}^j r_{j,t}, \quad (1)$$

with τ_{t-1}^j denoting the share of industry j in aggregate manufacturing sales in year $t - 1$. Let $\bar{r}_t = \frac{1}{86} \sum_{j=1}^{86} r_{j,t}$ denote the average industry growth rate. Adding and subtracting \bar{r}_t to (1) results in the aggregate growth rate being a function of the growth rate in the average industry along with the weighted deviation from the average:

$$R_t = \bar{r}_t + \frac{1}{86} \sum_{j=1}^{86} \tau_{t-1}^j (r_{j,t} - \bar{r}_t). \quad (2)$$

The aggregate growth rate is therefore the sum of the equal-share component (\bar{r}_t) and the proportional-share component ($\frac{1}{86} \sum_{j=1}^{86} \tau_{t-1}^j (r_{j,t} - \bar{r}_t)$). As discussed by FSW, if aggregate growth is driven by a few large industries, then the variation in R_t should be driven by the variation in the proportional-share component. On the other hand, if aggregate variation is driven by growth in all industries, then the variation in R_t should be driven by the variation in the equal-share component. Table 1 presents the standard deviation for R_t , the equal-share component and the proportional-share component.

Table 1: Share weight decomposition of the standard deviation of manufacturing sales, investment and employment growth

	Sales	Investment	Employment
	standard deviation		
Aggregate growth	0.049	0.120	0.028
Equal-share component	0.045	0.052	0.030
Proportional-share component	0.021	0.133	0.009

Source: Statistics Canada, author’s calculations based on T2-LEAP data, 2000-2012.

In the cases of sales and employment growth, the standard deviations of the equal-share components more closely resemble in magnitude the standard deviations of the aggregate growth rates. The opposite is the case for investment, likely due to its lumpy nature—this difference will be revisited when considering the effects of granular shocks to investment.

3. Quantifying firm-specific and granular shocks

The identification of firm-specific shocks in this paper closely follows the methodology outlined by Gabaix (2011), which is also adopted by DLM. Let $g_{i,t}$ denote the growth rate of the variable of interest for firm i in year t . The variable of interest can be the sales, investment or employment of a firm. This paper uses the log growth rate to measure the growth of sales and employment between periods $t - 1$ and t . The midpoint growth rate is used to calculate the growth rate of investment.⁶ The growth rate of a firm consists of two components: one common to all firms in the industry (i.e., a macro shock) and one specific to the firm. Hence, the firm-specific shock is the portion of the growth rate $g_{i,t}$ that is unaccounted for by a common, industry-wide shock.

There are many possible ways to quantify the macro shock. Gabaix defines it as the average growth rate of a small subset of very large firms, implicitly acknowledging that macro shocks may affect large firms in a fundamentally different manner than smaller firms.⁷ On the other hand, DLM use the average growth rate of all firms as their measure of the common macro shock.⁸ While macro shocks are likely to have a heterogeneous impact on firms of different sizes, it is difficult to determine the extent to which this is the case, and whether there is variation across industries. Consequently, a combination of the two methodologies is used to identify reasonable bounds for the impact of firm-specific shocks on industry dynamics.

Let $M_{j,t}$ denote a set of firms in industry j in year t . $M_{j,t}$ can denote either the entire set of firms in the industry (as in DLM's work) or a small number of the largest firms in the industry (along the lines of Gabaix's methodology). The common macro shock, $\bar{g}_{j,t}^M$, is then the average (arithmetic mean) growth rate of these $M_{j,t}$ firms. The firm-specific shock is the component of the firm's growth rate that is not explained by this macro shock. Mathematically, it is the difference between the growth rate of the firm and the average growth rate of the $M_{j,t}$ firms:

$$\epsilon_{i,t}^M = g_{i,t} - \bar{g}_{j,t}^M. \quad (3)$$

⁶Due to the lumpy nature of investment and the fact that many firms report zero investment in a given year, using the log growth rate for investment would result in many dropped observations. Therefore, firm-level investment growth is defined as $g_{i,t} = \frac{2(X_{i,t} - X_{i,t-1})}{X_{i,t} + X_{i,t-1}}$, where $X_{i,t}$ is the level of investment for firm i in year t . This growth rate formulation has become standard in the analysis of firm dynamics (see Haltiwanger, Jarmin and Miranda, 2013; and Tornqvist, Vartia and Vartia, 1985).

⁷This is a reasonable assumption for several reasons. For instance, exchange-rate movements have a greater impact upon exporting firms, which are almost exclusively large firms. Such movements influence the domestic operations of exporters due to capacity constraints (Blum, Claro and Horstmann, 2013; Soderbery, 2014). In addition, Holmes and Stevens (2014) argue that large firms and small firms from the same industry produce fundamentally different types of products and target different types of customers. Aggregate shocks would therefore influence them in different ways.

⁸With their French data, DLM show that allowing firm sensitivity to aggregate and sectoral shocks to differ by firm size has little impact on the main results (see Subsection 4.4 in their paper).

The influence of a firm-specific shock on an industry aggregate is proportional to the size of the firm relative to the industry as a whole. The simplest measure of the size of a firm is its market share in the previous period, which is denoted by $s_{i,t-1}^j$. The overall impact of firm-specific shocks on industry aggregates is the weighted average of the firm-specific deviations from the average growth rate:

$$G_{j,t} = \sum_{i \in N_{j,t}} (s_{i,t-1}^j \cdot \epsilon_{i,t}^M). \quad (4)$$

$G_{j,t}$ is commonly referred to as the granular shock. $N_{j,t}$ denotes the set of firms in industry j that is used to construct the granular shock. Gabaix and DLM both set $N_{j,t} = M_{j,t}$. However, the former chose only a small number of very large firms, whereas the latter used all firms. The present study allows for $N_{j,t} \subseteq M_{j,t}$, to separate out the effects of shocks to very large firms from the definition of the macro shock.

The granular shock defined by equation (4) summarizes the aggregate impact of firm-specific shocks on a group of firms in industry j . If aggregate fluctuations are only caused by aggregate shocks, whether economy-wide or industry-wide, then the firm-specific deviations from the industry average growth rate $\bar{g}_{j,t}^M$ would cancel out in the aggregate. As emphasized by Gabaix, such shocks do not cancel out in granular markets because of the presence of large firms—the magnitude of a shock to a large firm is not offset by shocks to smaller firms. The impact of $G_{j,t}$ on industry dynamics therefore depends on the extent to which economic activity is concentrated within a small number of large firms (i.e., their combined market share), along with the dispersion of growth rates within the industry.

As mentioned, the manufacturing sector consists of 86 industries at the 4-digit NAICS level. The average industry contained over 550 firms in a typical year over the course of the sample period. When examining a smaller subset of large firms, this article looks into the contribution of firm-specific shocks to the 10 largest firms in each industry. This means that the focus is on a total of 860 firms that account for less than 2% of the total number of manufacturing firms.⁹ In an average year, this set of firms accounted for nearly three-quarters of aggregate sales, over 70% of aggregate investment, and nearly half of aggregate employment in the entire manufacturing sector.

The size of the 10 largest firms exhibits significant variation across industries, as their combined market share varies between 12.0% and 99.9%, with an average of 64%. Hence, the 10 largest firms may be subject to different types of shocks than smaller firms in some industries, but not in others. Consequently, the appropriate definition of $M_{j,t}$ to construct $\bar{g}_{j,t}^M$ may vary across industries, too. Disentangling these cross-industry differences is beyond the scope of this

⁹Since the size of a firm is based on its market share in the previous period, the composition of the 860 firms varies from year to year. Nonetheless, the annual turnover is rather small and stands at around 4% per year.

article.

Three broad approaches to construct the granular shock are therefore considered, and they are summarized in Table 2. The first two measures, granular shock 1 (GS1) and granular shock 2 (GS2), roughly correspond to the approaches of Gabaix and DLM, respectively. The third method is a hybrid of the two—it uses information on all firms to construct the macro shock $\bar{g}_{j,t}^M$ to derive the firm-level shock $\epsilon_{i,t}^M$, yet uses only the 10 largest firms to construct the granular shock $G_{j,t}$. This measure enables researchers to follow DLM’s approach in measuring the macro shock, while still allowing for the study of shocks to the largest firms. A priori, it is difficult to determine which of the three approaches provides the most appropriate measure of the granular shock, given the inherent difficulty in identifying the macro shock across multiple industries. Consequently, all three methods are used.

Table 2: Summary of granular shock measures

Granular shock measure	$N_{j,t}$	$M_{j,t}$
Granular shock 1	Ten largest firms	Ten largest firms
Granular shock 2	All firms	All firms
Granular shock 3	Ten largest firms	All firms

Note: $M_{j,t}$ is the set of firms used to calculate the macro shock. $N_{j,t}$ is the set of firms used to calculate the granular shocks.

4. Quantifying the contribution of granular shocks to aggregate fluctuations

As mentioned in the introduction, many macroeconomic models are premised on the assumption that shocks to individual firms cancel out in the aggregate. Gabaix pointed out that if this is the case, then it follows that a regression of any economic aggregate on the granular shock should yield an R-squared value close to zero.¹⁰ Alternatively, if the granular shock provides some explanatory power, the R-squared value should be positive.

Let $X_{j,t}$ denote an industry-level aggregate of interest. The three variables of interest are industry-level sales (which closely track gross output), investment and employment. The explanatory power of the granular shock on industry sales and employment can be captured by running the following univariate regression:

$$\ln X_{j,t} - \ln X_{j,t-1} = \beta_0 + \beta_1 G_{j,t} + u_{j,t}. \quad (5)$$

β_0 and β_1 are parameters to be estimated, and $u_{j,t}$ is an error term. The R-squared value from equation (5) reveals the amount of annual variation in the industry aggregate of interest that can

¹⁰The R-squared value summarizes the extent to which an independent variable explains variation in the dependent variable.

be explained by variations in the granular shock.¹¹ Separate panel regressions are run for sales and employment using the three granular shock measures outlined in Table 2. The regressions are estimated at the 4-digit NAICS level for 86 industries over a 12-year period, resulting in 1,032 observations. To account for correlation among the error terms for an industry over time, the standard errors are clustered at the 4-digit NAICS level in all of the regressions.¹²

The approach used to examine granular shocks to investment is slightly different. There are two important issues to consider when analyzing investment. First, as revealed in Table 1, in contrast to sales and employment, the proportional-share component is the dominant source of aggregate investment volatility. This means that an approach of running unweighted industry-level regressions—as in equation (5)—may not be suitable for thinking about the effects of granular shocks on aggregate investment volatility in the manufacturing sector. Second, the fact that firm-level investment is lumpy by nature means that the idiosyncratic shock can be particularly difficult to measure for investment. For example, many idiosyncratic shocks will be measured as negative due to the fact that firms will have zero investment growth from year to year, whereas average growth is often positive. To address these issues, the investment analysis will incorporate industry weights into the regression framework (where the weights are the industry shares in total investment) and the dependent variable will be the midpoint growth rate (as in Haltiwanger, Jarmin and Miranda, 2013, among others), which is consistent with how firm-level investment growth is defined in the previous section. That is:

$$\frac{2(X_{j,t} - X_{j,t-1})}{X_{j,t} + X_{j,t-1}} = \beta_0 + \beta_1 G_{j,t} + u_{j,t}. \quad (6)$$

Since data for investment are available only from 2001 onwards, there are 946 observations in the regressions for investment.

4.1. Sales

Table 3 presents estimation results for industry sales growth as the variable of interest. The most parsimonious granular shock measure (in terms of the data needed to measure the shocks), GS1, explains 23% of the annual variation in industry sales growth, while the least parsimonious measure, GS2, explains 46%. At first glance, it is difficult to determine whether the explanatory power doubles because of the inclusion of shocks to smaller firms or because of the difference in the measurement of firm-specific shocks. The GS3 results in Table 3 demonstrate that it is

¹¹It is possible that the granular shock is skewed by outliers. The sensitivity of the results to the Winsorization of firm-level growth outliers ($g_{i,t}$) at both the 5% level and the 10% level has therefore been examined. The results were not materially different from those presented below and are available from the authors upon request.

¹²The baseline regressions do not include industry fixed effects, and therefore capture both cross-sectional variation and time-series variation. Please refer to Subsection 4.4 for a discussion on disentangling these two sources of variation.

primarily because of the latter. Removing all but the 10 largest firms in each industry from the construction of the granular shock reduces the R-squared value by only 0.01, from 0.46 to 0.45. In other words, shocks to the 10 largest firms in each industry explain 45% of the annual variation in industry sales growth, while shocks to the remaining 98% of firms explain only 1% of the variation.

Table 3: Impact of firm-level shocks on aggregate sales growth

	Dependent variable, aggregate sales growth		
	Granular shock 1	Granular shock 2	Granular shock 3
Granular shock			
Coefficient	0.686**	0.685**	0.683**
Standard error	0.091	0.074	0.076
Constant			
Coefficient	0.010*	0.039**	0.033**
Standard error	0.005	0.006	0.006
Number of observations	1,032	1,032	1,023
R-squared	0.23	0.46	0.45

* significantly different from reference category ($p < 0.05$)

** significantly different from reference category ($p < 0.01$)

Source: Statistics Canada, authors' calculations based on T2-LEAP data, 2000-2012.

These results suggest that shocks to the vast majority of firms do cancel out, in line with the assumption of most macro models. However, shocks to a (relatively) small number of large firms do not cancel out, and this set of shocks partially drives annual variation in macroeconomic variables.

Estimating equation (5) using OLS will yield a coefficient estimate for β_1 that describes the influence of the granular shock on industry fluctuation in the average industry. As shown in Table 1, although the equal-share component accounts for most of the overall variation in sales, the proportional-share component stills plays a role. Therefore, these estimates may understate the contribution of granular shocks to aggregate fluctuations in the manufacturing sector, since large firms are typically concentrated in large industries (in fact, the correlation between an industry's share in total manufacturing output and the normalized Herfindahl index is positive). To address this concern, weighted regressions are run (using industry-level weights). It is found that the explanatory power of the granular residuals changes very little, and therefore only the unweighted results are reported.

4.2. Investment

Table 4 presents estimation results for industry investment growth (equation 6). The structure and layout of this table are identical to those of Table 3. As was the case with sales,

firm-level shocks are indeed an important source of annual fluctuations in industry investment. Granular shocks can explain between 13% and 37% of annual industry-level variation, depending on the methodology used to construct the granular shock. As was also the case for sales, although the explanatory power of GS2 exceeds that of GS1 for investment, the vast majority of the explanatory power of the former is attributable to the 10 largest firms in each industry. As shown in the GS3 results, removing all but the 10 largest firms from the construction of the granular shock lowers the R-squared by only 1 percentage point (from 0.37 to 0.36).

Table 4: Impact of firm-level shocks on aggregate investment growth

	Dependent variable, aggregate investment growth		
	Granular shock 1	Granular shock 2	Granular shock 3
Granular shock			
Coefficient	0.570*	0.768**	0.764**
Standard error	0.216	0.149	0.149
Constant			
Coefficient	-0.035*	-0.086**	-0.078**
Standard error	0.014	0.012	0.011
Number of observations	946	946	946
R-squared	0.13	0.37	0.36

* significantly different from reference category ($p < 0.05$)

** significantly different from reference category ($p < 0.01$)

Source: Statistics Canada, authors' calculations based on T2-LEAP data, 2001-2012.

4.3. Employment

Table 5 presents estimation results for industry employment growth. In marked contrast to the impact on industry sales and investment growth, granular shocks account for little of the annual variation in industry employment. The R-squared value for the three specifications ranges from 0.00 to 0.02, and the coefficient estimates for the granular shock are not statistically significant.¹³

At first glance, this result may seem surprising. However, it is consistent with a large body of literature that indicates that job creation and destruction rates at the firm level vastly exceed what is observed in the aggregate. This means that a large portion of the firm-specific employment shocks do cancel out in the aggregate.¹⁴

¹³As with sales, weighted regressions were run and the explanatory power of the granular residuals changes very little. These results, therefore, are not reported.

¹⁴See Baldwin (1995) and Rollin (2012) for evidence for Canada, and Haltiwanger, Jarmin and Miranda (2013) for evidence for the United States.

The low explanatory power of the granular shock can also be attributed to the greater inherent difficulty in identifying firm-specific shocks to employment, relative to sales and investment. Most firms are likely to respond to positive (or negative) macro shocks by increasing (or reducing) sales and investment. This is not necessarily the case for employment, as firms of different sizes respond in different ways to macro shocks (see Moscarini and Postel-Vinay, 2012; and Criscuolo, Gal and Menon, 2014).¹⁵ This may skew the average growth rate $\bar{g}_{j,t}^M$, resulting in a mis-estimation of the firm-specific shock $\epsilon_{i,t}^M$.

Table 5: Impact of firm-level shocks on aggregate employment growth

	Dependent variable, aggregate employment growth		
	Granular shock 1	Granular shock 2	Granular shock 3
Granular shock			
Coefficient	0.107	0.029	0.098
Standard error	0.071	0.043	0.071
Constant			
Coefficient	-0.026**	-0.027**	-0.026**
Standard error	0.004	0.004	0.004
Number of observations	1,032	1,032	1,023
R-squared	0.02	0.00	0.02

** significantly different from reference category (p<0.01)

Source: Statistics Canada, authors' calculations based on T2-LEAP data, 2000-2012.

A third reason for the low explanatory power is that employment is not concentrated among the largest firms to the same extent as sales and investment. As mentioned in Section 3, the 10 largest firms in each industry account for over 70% of aggregate sales and investment in a typical year, yet account for only half of aggregate employment. Hence, the shocks that affect the sales and investment decisions of large firms are much larger than those that affect the sales and investment decisions of other firms, and they do not cancel out in the aggregate. However, this may not be the case for employment. On the contrary, there is strong evidence that a large share of job creation and job destruction is generated by the entry of new firms and the exit of existing firms (see Haltiwanger, Jarmin and Miranda, 2013; and Rollin, 2012).¹⁶ Their omission

¹⁵Moscarini and Postel-Vinay (2012) found that employment growth in large firms is more responsive to macro shocks than employment growth in small firms. They attribute this dynamic to the greater ability of large firms to recruit skilled workers throughout the business cycle. Smaller firms, by contrast, find it difficult to do so during an economic boom, when workers have greater opportunities, and are therefore more hesitant to lay off workers in a downturn. Although the bulk of the paper by Moscarini and Postel-Vinay deals with U.S. firms, they show that their findings are also consistent with the employment decisions of large and small firms in Canada. Similarly, Criscuolo, Gal and Menon (2014) note that new firms, which are typically small, tend to create jobs (conditional on remaining in business) even in times of recession. The positive employment growth of young firms, the negative employment growth of large established firms, and the stagnant growth of smaller established firms could result in a severe mismeasurement of the macro shock $\bar{g}_{j,t}^M$ when it comes to employment.

¹⁶Rollin (2012) found that new firms were responsible for 16% of gross job creation, while exiting firms were

is another source of potential measurement error.

4.4. Time-series versus cross-sectional explanatory power

The results in Subsections 4.1 and 4.2 reveal that granular shocks explain a significant portion of the annual variation in industry sales and investment. This explanatory power can arise from cross-sectional variation, time-series variation, or both. Since the dataset consists of observations spanning only 12 years but 86 industries, it is probable that most of the variation arises from the cross-sectional component of the data. To disentangle the two effects, the following set of equations is estimated for sales:

$$\ln X_{j,t} - \ln X_{j,t-1} = \alpha + \gamma_j + \mu_{j,t}, \quad (7)$$

$$\ln X_{j,t} - \ln X_{j,t-1} = \alpha + \gamma_j + \delta G_{j,t} + \mu_{j,t}. \quad (8)$$

γ_j denotes a time-invariant industry fixed effect. The R-squared value from regression (7) reveals how much of the variation in the growth rate of the variable of interest is time-invariant. With this channel shut off, the increase in the R-squared value that arises from the inclusion of the granular shock in regression (8) is attributable to the time-series variation in the granular shock. A comparable set of (weighted) least squares regressions is estimated for investment, using the midpoint growth rate as in equation (6).

This set of regressions is estimated for both sales and investment using all three measures of the granular shock. The added explanatory power from the inclusion of the granular shock is essentially equivalent to the R-squared values presented in Tables 3 and 4. For example, industry fixed effects explain 12% of the annual variation in industry sales growth. The inclusion of the three granular shock measures raises the explanatory power to 0.33, 0.58 and 0.57, respectively. Therefore, the additional explanatory power of the three granular shocks is 0.21, 0.46 and 0.45, respectively, and this is in line with the results presented in Table 3. Similar results are observed for investment. Hence, most of the explanatory power identified in Tables 3 and 4 comes from the time-series component of the data. This suggests that examining shocks to large firms can shed light on the evolution of macroeconomic variables over time and help forecast future growth.

5. Discussion

The empirical results in this article provide clear evidence that shocks to a relatively small number of large firms are responsible for a significant share of the annual variation in some

responsible for 17% of gross job destruction in Canada from 2001 to 2009. These figures far exceed the market share of new firms. Hence, while the manufacturing sector in Canada is granular as far as sales and investment are concerned, it is not as granular in terms of employment.

important macroeconomic variables. These findings represent an important development in understanding the origins of aggregate fluctuations in Canada.

While this article answers the questions set out in the introduction, it also raises new ones. The empirical strategy is based on a range of possible estimates of the granular shock. More refined estimates would require an industry-by-industry analysis to determine how the effects of macroeconomic shocks are distributed across firms of different size categories within an industry. This is beyond the scope of this paper, but is worthy of further exploration.

Furthermore, the estimates of the explanatory power of granular shocks may be biased downward because of an additional source of bias in measuring firm-specific shocks. It is well known that industries are intimately connected to one another by input-output linkages (see Acemoglu et al., 2012; Carvalho, 2014; and Foerster, Sarte and Watson, 2011). Consequently, firm-specific shocks in one industry may spill over into other industries, resulting in further fluctuations. Because of such spillovers, a granular shock in one industry may be picked up as a macro shock in another. If this is the case, the industry-level variation that comes from this spillover should be characterized as originating from firm-specific shocks, rather than from macro shocks. Taking this propagation mechanism into account can improve the reliability of the estimates. In unreported results, the authors of this paper found significant linkages in the sales and employment dynamics of the industries that compose the manufacturing sector. Examining the extent to which input-output linkages cause granular shocks to spill over from one industry into others would therefore be a fruitful avenue of research and could help address the measurement issues raised in this paper.

Finally, it is important to distinguish between the focus of this paper—the determination of annual variations in industry outcomes—and long-term or trend growth. Despite the dominance of large firms in generating year-to-year variations, a significant share of long-run growth is caused by the entry of new firms and the expansion of small, young firms. Therefore, continued study of the entry and growth of new firms is important for improving the way long-run aggregate fluctuations are understood.

References

- Acemoglu, D., V. Carvalho, A. Ozgaldar, and A. Tahbaz-Salehi. 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica* 80(5): 1977-2016.
- Baldwin, J.R. 1995. *The Dynamics of Industrial Competition: A North American Perspective*. Cambridge: Cambridge University Press.
- Blum, B.S., S. Claro, and I.J. Horstmann. 2013. "Occasional and Perennial Exporters." *Journal of International Economics* 90(1): 65-74.
- Carvalho, V. 2014. "From Micro to Macro via Production Networks." *Journal of Economic Perspectives* 28(4): 23-47.
- Conley, T.G., and B. Dupor. 2003. "A Spatial Analysis of Sectoral Complementarity." *Journal of Political Economy* 111(2): 311-352.
- Criscuolo, C., P.N. Gal, and C. Menon. 2014. *The Dynamics of Employment Growth: New Evidence from 18 Countries*. OECD Science, Technology and Industry Policy Papers, no. 14. Paris: OECD Publishing.
- di Giovanni, J., A. Levchenko, and I. Mejean. 2014. "Firms, Destinations, and Aggregate Fluctuations." *Econometrica* 82(4): 1303-1340.
- Foerster, A., P.G. Sarte, and M.W. Watson. 2011. "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production." *Journal of Political Economy* 119(1): 1-38.
- Gabaix, X. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica* 79(3): 733-772.
- Haltiwanger, J., R. Jarmin, and J. Miranda. 2013. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95(2): 347-361.
- Holmes, T., and J. Stevens. 2014. "An Alternative Theory of the Plant Size Distribution, with Geography and Intra- and International Trade." *Journal of Political Economy* 122(2): 369-421.
- Horvath, M. 1998. "Cyclicality and Sectoral Linkages: Aggregate Fluctuations from Independent Sectoral Shocks." *Review of Economic Dynamics* 1(4): 781-808.
- Lafrance, A. 2013. *Firm Size and the Risk/Return Trade-off*. Economic Analysis Research Paper Series, no. 87. Statistics Canada Catalogue no. 11F0027M. Ottawa: Statistics Canada.
- Lafrance, A., and W. Gu. 2014. *Productivity Growth in the Canadian Broadcasting and*

- Telecommunications Industry: Evidence from Micro Data. Economic Analysis Research Paper Series, no. 89. Statistics Canada Catalogue 11F0027M. Ottawa: Statistics Canada.
- Leung, D., L. Rispoli, and R. Chan. 2012. Small, Medium-sized, and Large Businesses in the Canadian Economy: Measuring Their Contribution to Gross Domestic Product from 2001 to 2008. Economic Analysis Research Paper Series, no. 82. Statistics Canada Catalogue no. 11F0027M. Ottawa: Statistics Canada.
- Lucas, R.E. 1977. "Understanding Business Cycles." *Carnegie-Rochester Conference Series on Public Policy* 5(1): 7-29.
- Moscarini, G., and F. 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.
- Rollin, A.-M. 2012. Firm Dynamics: Employment Dynamics Arising from Firm Growth and Contraction in Canada, 2001 to 2009. The Canadian Economy in Transition, no. 24. Statistics Canada Catalogue no. 11-622-M. Ottawa: Statistics Canada.
- Soderbery, A. 2014. "Market Size, Structure, and Access: Trade with Capacity Constraints." *European Economic Review* 70: 276-298.
- Tornqvist, L., P. Vartia and Y.O. Vartia. 1985. "How Should Relative Changes be Measured?" *The American Statistician* 39(1): 43-46.