

The Gertler-Gilchrist Evidence on Small and Large Firm Sales

VV Chari, LJ Christiano and P Kehoe

July, 2013

In this note, we examine the findings of Gertler and Gilchrist, ('Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms,' *The Quarterly Journal of Economics*, vol. 109, no. 2, May 1994, pp. 309-340.) GG construct a measure of the sales of large and small establishments and argue that small establishment sales fall by more than large establishment sales in response to a contractionary monetary policy shock. This is interpreted as providing support for the notion that a monetary contraction affects the economy in part by making credit relatively tight for small firms. In this note, we investigate a closely related question. We ask what happens to the sales of large versus small firms during a business cycle contraction. The question is a different one, to the extent that shocks other than monetary policy also play an important role in triggering recessions. This makes whatever answer we find harder to interpret. At the same time, business cycle recessions are easier to identify in the data than are monetary policy shocks, and so the question we ask allows us to sidestep difficult identification questions.

Our analysis is broken into two parts. First, we confine ourselves to the data set constructed by Gertler and Gilchrist, which cover the period, 1959Q1-1991Q4. After reproducing their findings for the response of large and small firm sales to Romer-Romer monetary contractions, we find that small firm sales also drop relative to large firm sales during recessions. At the same time, there is substantial variation in sales across different business cycle episodes. The differences between small and large firms are less significant than they are after Romer-Romer episodes. We suspect that a formal test of the null hypothesis that large and small firm sales respond to recessions in the same way would not be rejected. In the second part, we construct a longer time series on large and small firm sales. This allows us to consider three additional recessions, two in the 1950s as well as the two 2000 recessions. We find that the additional business cycle episodes in our data sample overturn the results reported in the first part of our analysis. In particular, the response of large and small firm sales in a recession is roughly the same.

1 Some Key Dates

The analysis of Gertler-Gilchrist places special emphasis on several particular dates. We discuss these here. Figure 1 reports log US real GDP, after HP filtering. Grey areas indicate NBER business cycle peaks and troughs. Note how the NBER peaks typically occur several quarters after an actual turning point in the data.

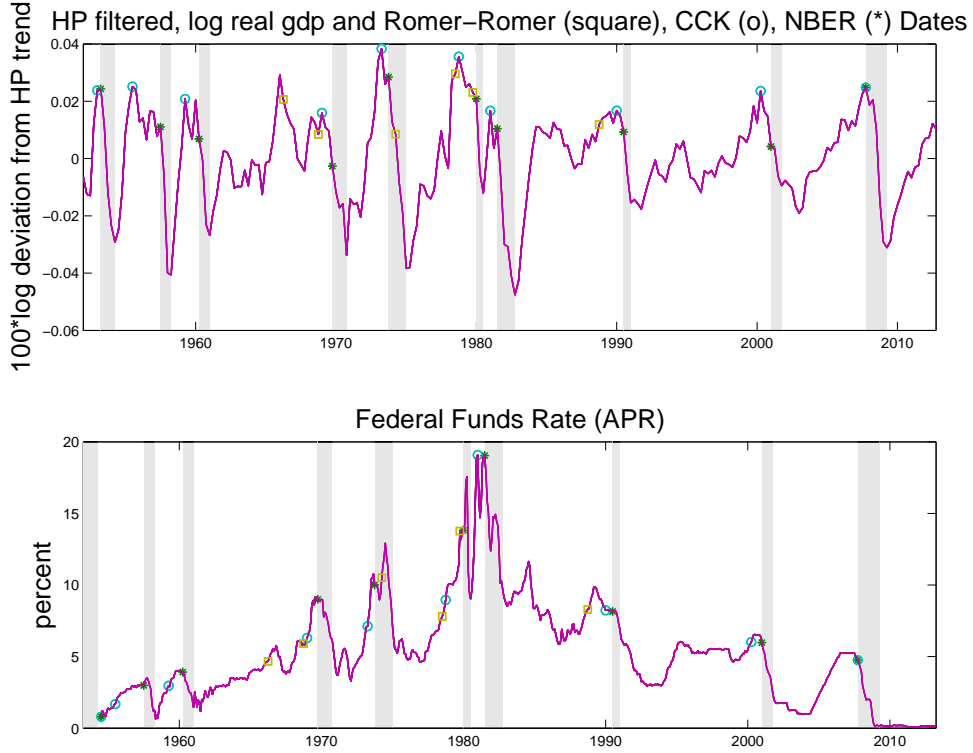


Figure 1

This motivated us to construct an alternative set of business cycle peaks. These are indicated by the circles, and we refer to these as ‘CCK business cycle dates’. They are local peaks in detrended output, after which output fell by at least several percent. The dates are 1953Q1, 1955Q3, 1959Q2, 1969Q1, 1973Q2, 1978Q4, 1981Q1, 1990Q1, 2000Q2, 2007Q4. The date, 1981Q1, satisfies our criterion that the economy fell by more than a couple of percent afterward. However, the expansion before 1981Q1 is very brief, and we thought of defining the 1981Q1 recession as part of a recession that actually began in 1978Q4. In effect, we follow the NBER in this episode in taking the position that there were two recessions at the end of the 1970s and the start of the 1980s. These occur in 1966Q2, 1968Q4, 1974Q2, 1978Q3, 1979Q3, 1988Q3. With the exception of the 1966 date, these dates were taken from <http://elsa.berkeley.edu/~dromer/papers/RomerandRomerDates.pdf>. The 1966 date was included in the Gertler-Gilchrist study and we include them here for that reason. Romer and Romer discuss the 1966 date in footnote 13 of their 1990 Brookings paper. (The 1974 and 1979 dates occur well after the economy has started to weaken and may perhaps therefore not be considered the primary cause for the business contractions in those two periods.) The second graph in Figure 1 shows the federal funds rate and the Romer and Romer dates.

2 Reproducing the Gertler-Gilchrist Results

We obtained the actual sales data, for large and small firms, that Gertler and Gilchrist used in their Figure II, page 321. These data are displayed in Figure 2. The actual data were provided by the authors in growth rates. The top panel displays the corresponding levels, obtained by cumulating the first differences and setting the initial condition to zero. Note how large firm sales displays strong growth until the early 1980s, whereupon it stabilizes. Small firm sales exhibit little growth throughout the period. The bottom panel of Figure 2 displays the level data, after linear detrending. The squares in the figure indicate the Romer and Romer dates and the circles indicate the CCK business cycle peaks.

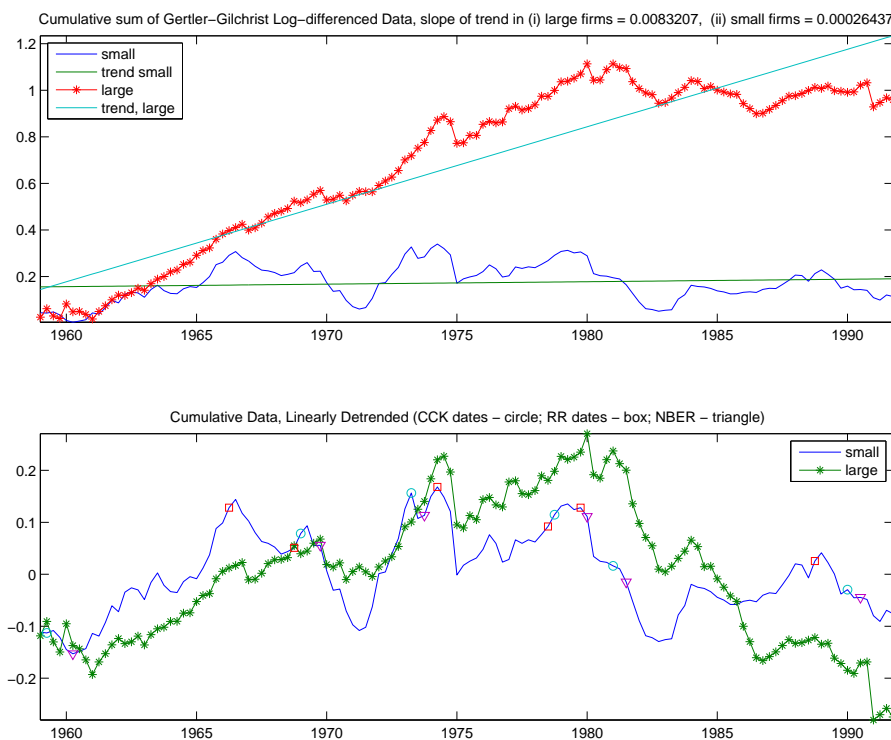


Figure 2

Consider the Romer-Romer dates (i.e., the squares) first. After the one in 1966 it is clear that small firm sales drop by more than large firm sales. The same is true for the 1968 date, although initially there is a transient rise in small firm sales. After 1974 small firm sales fall more than large firm sales, but the difference looks relatively small. The 1978 date does not fit into the general pattern. After this date the sales of small and large firms both increase until just before the next Romer-Romer date. After the 1979 date, the small firms seem to go down faster and more than the large firms.

Figure 3 displays a worm chart like Gertler-Gilchrist's Figure II, which displays the detrended data 8 quarters before and 12 quarters after a Romer-Romer date. All data are adjusted by a constant so that in period 0, the Romer-Romer date, the data are always normalized at zero. The starred line in the figure is the average of the individual curves. Note that the average cumulative drop in the level of sales by small firms after a contraction is about 5 percentage points greater than what it is for large firms.

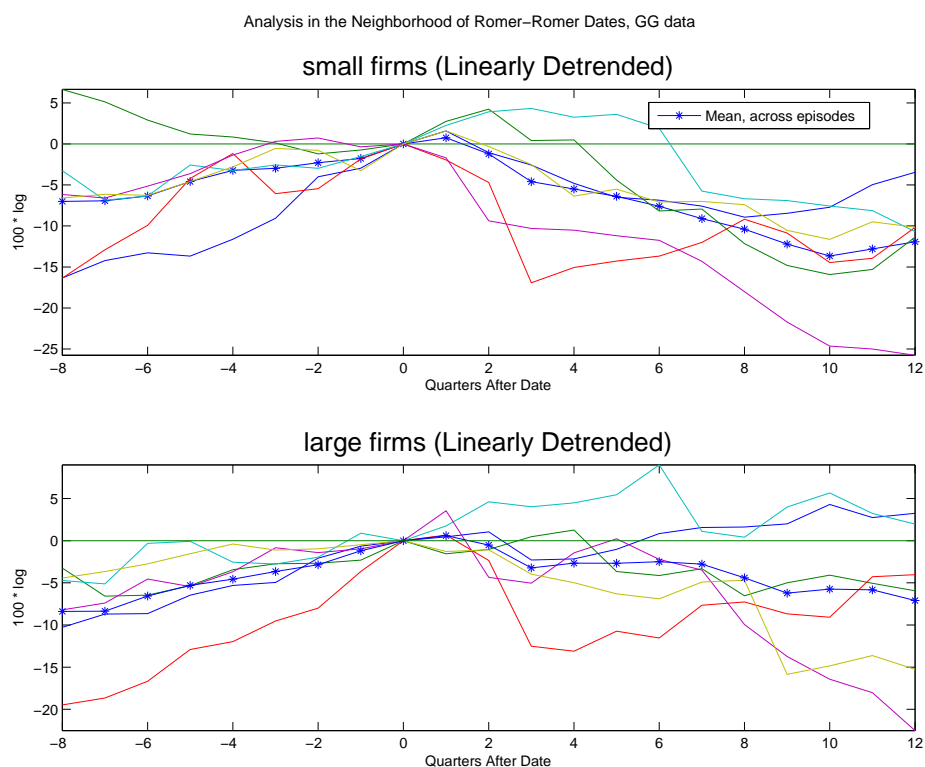


Figure 3

The results in our Figure 3 resemble the Gertler-Gilchrist, Figure II results closely. However, there is one difference. The data related to the 1979 episode appear different. Figure 2 shows that in that episode, sales of large firms are about 22 percent below peak and sales of small firms are a little over 25 percent below peak. (The 1979 episode corresponds to the lowest numbers at the end of the worm chart in Figure 3.) Gertler and Gilchrist, Figure II, report that at the end of the 1979 episode large firm sales are about 17 percent below peak and small firm sales are around 25 percent below peak. We are able to roughly reproduce the 1979 Gertler-Gilchrist worm episode if we assume the 1979 peak occurred in the third quarter.

We investigated the robustness of the above finding to the detrending strategy. In part, this is motivated by the evidence in Figure 2 that the trend in the levels of large and small

firms is not well captured by a simple linear trend. Repeating the calculations using the HP filter to detrend the cumulated levels, we obtain the following results:

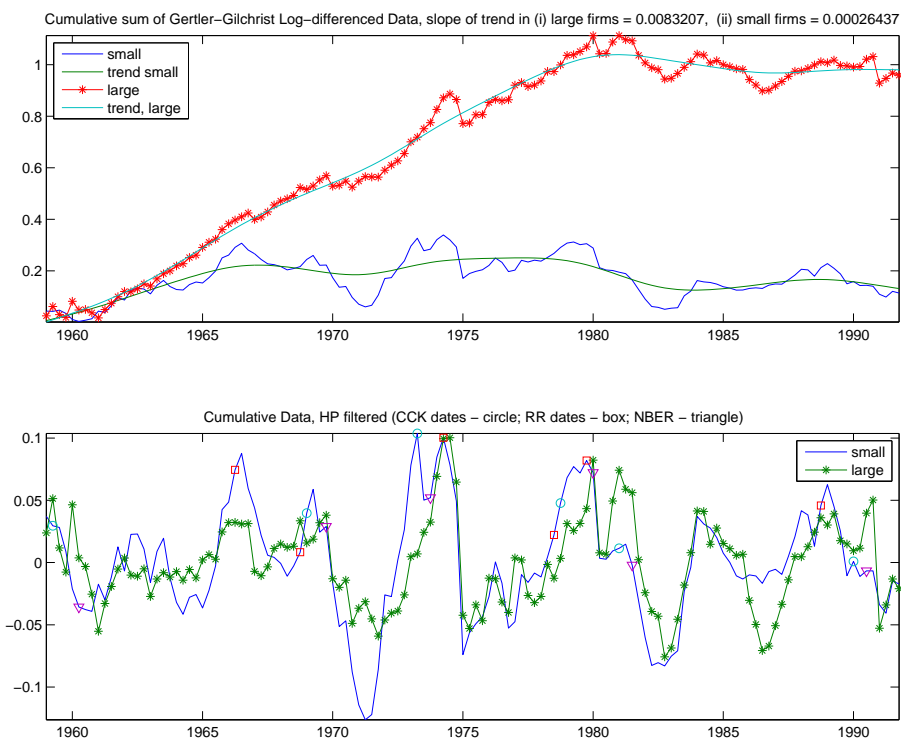


Figure 4

The behavior in Figure 4 of detrended large firm sales is somewhat sensitive to the detrending procedure, because of the apparent break in the trend of that series around 1980. As before, small firm sales drop a lot after the 1966 date. There is a small surge in small firm sales after the 1968 date. After the 1974 date, the fall in the two sales figures is very large - in the neighborhood of 17 percent - though the drop in small firm sales is slightly larger than the drop in large firm sales. The 1978 date remains completely inconsistent with the GG hypothesis, with large and small firm sales still surging afterward until the 1979 date. In the period after the 1988 Romer-Romer date, the drop in large and small firm sales seem to be roughly the same order of magnitude. Interestingly, after 1985 large firm sales are more volatile than small firm sales in a way that seems inconsistent at least with the spirit of the GG hypothesis. Overall, the results seem to favor the GG hypothesis less strongly with the change in detrending procedure.

Consistently with our informal analysis of Figure 4, when we compute the worm chart using the data in Figure 4, we find that the GG results are slightly weaker quantitatively

than the ones reported in Figure 3, Figure 5 below shows that the average cumulative drop in sales by small firms is about 7 percent by the third year after a Romer-Romer date, and the corresponding result of large firms is 5 percent. However, the spread between the mean cumulative drop in sales is larger in the quarters before. In addition, we see that across all Romer-Romer episodes small firms are always below zero after 8 quarters out, whereas large firms' sales are more often in the positive region. Overall, the impression from the following figure is that the cumulative drop in small firm sales is larger than it is for large firm sales, in the wake of a Romer-Romer date. There is a small quantitative difference with HP-filtering, but not a qualitative difference.

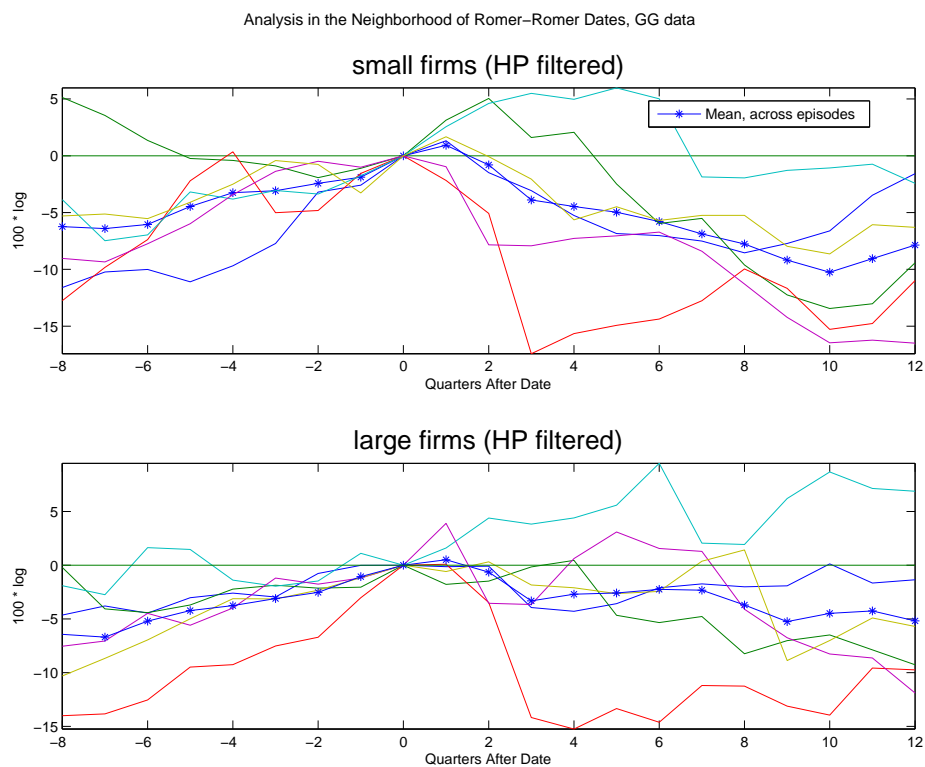


Figure 5

Next, we computed the worm chart calculations for CCK business cycle peaks, and reverting to linear detrending. The results are displayed in Figure 6:

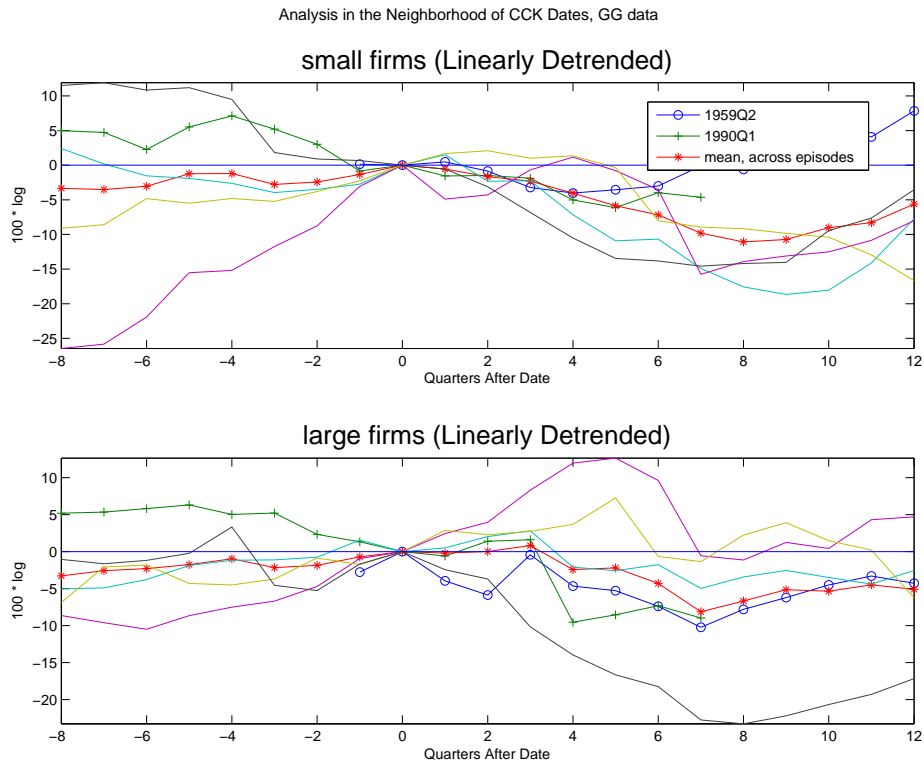


Figure 6

The first and last business cycle peaks are too near the ends of the data set, and so these episodes are truncated. The cross-episode mean averages just over the data that are available. It seems fair to conclude that these results are somewhat weaker than the Gertler-Gilchrist results, because there is a greater spread among the responses across episodes. Still, the averages that are reported indicate that small firm sales fall more than do large firm sales, after a CCK business cycle peak. In particular, in quarters 6 and 7, small firm sales are down by -4.5, -5.5, and -7.7 percent, respectively. At the same time, large firm sales are down by -2.33, -3.43, and -5.76 percent, respectively. Thus, the point estimates are consistent with the basic Gertler-Gilchrist results, though the differences seem less likely to be statistically significant, because of the high degree of variation across episodes.

We also examined the behavior of worms when the business cycle peaks correspond to NBER business cycle peaks and we continue to detrend linearly. The results are displayed in Figure 6a:

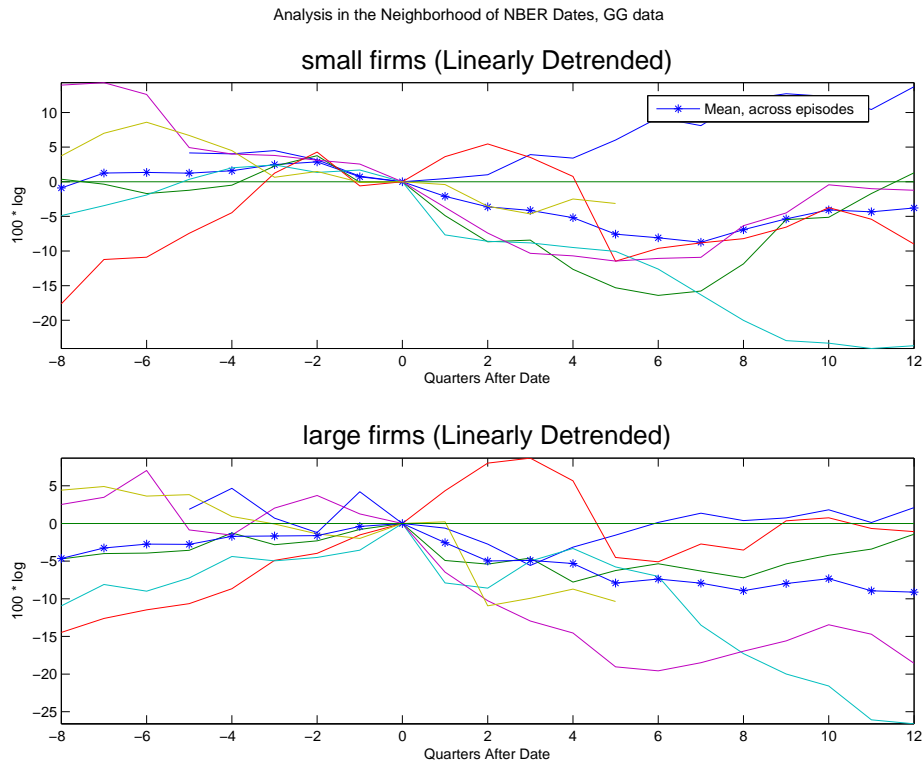


Figure 6a

Now, the results are essentially reversed. The large firms fall on average more during recessions. The difference in the averages is quite substantial and reaches 5 percentage points by three years after a business cycle peak.

3 Constructing Small and Large Firm Sales over a Larger Period

We expanded the Gertler-Gilchrist data set so that we can include more business cycle episodes in the analysis. We extended the data back to 1952Q1 and up to 2012Q4 by recovering data from various issues of the QFR. We constructed the data using a simple version of the procedure applied in Gertler-Gilchrist (see section 5 below for details.) The data are total dollar sales by manufacturing establishments in each of several size categories. A firm is in a given size category in a particular quarter if its assets are inside a specific range of values during that quarter. There are 7 - 9 size categories, with changes in the categories occurring four times in our data set. These changes reflect that the size categories are defined in nominal terms, so

that firms are drifting up through the size categories over time. The changes primarily involve consolidating the smallest size categories and expanding the largest.

Let x_1, x_2, \dots, x_{N_t} denote N_t asset categories. The asset cutoffs, x_i , and the number of asset categories, N_t , vary by year. The QFR reports the sales by establishments with assets less than x_1 , between x_1 and x_2 , and so on, finally with assets greater than x_{N_t} . For the interpolation procedure described below we require the top level of assets, say x_{N_t+1} . This information is not provided by the QFR and so we proxy x_{N_t+1} by the highest level of assets among all firms in COMPUSTAT in quarter t . The results are displayed in Figure 6a:

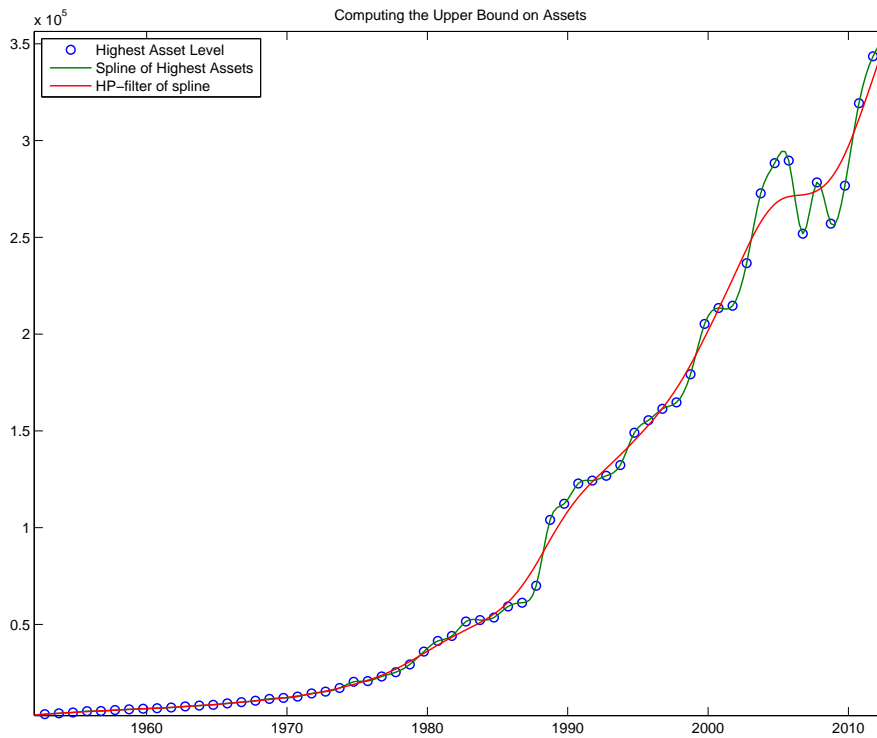


Figure 6a

In Figure 6a, the circles depict the actual asset levels and the line through the circles is a spline drawn through them. Notice the drop and volatility in assets of the top firm during the recent recession. The smooth line in the picture is the HP filter drawn through the circles. We use the points on the HP filter line as our proxy for the firm with the highest level of assets.

Following Gertler-Gilchrist, we define the sales of small firms as a weighted average of cumulative sales growth in the two asset size categories which straddle the 30th percentile of sales during period t . The weights are chosen so that the weighted average of cumulative sales in the two asset size classes average 30 percent of sales in period t . Specifically, we proceed as

follows. Let $S_{i,t}$ denote total sales of firms in size category i , in period t , after deflating by the GDP deflator. Let \bar{S}_{it} denote the sum of sales (scaled by all sales in period t) for asset size categories less than or equal to i , in period t :

$$\bar{S}_{it} = \frac{\sum_{j=1}^i S_{jt}}{\sum_{j=1}^{N_t} S_{jt}},$$

where N_t denotes all the size categories in period t and $i = 1, \dots, N_t$. We then compute

$$i_t = \min_{i \in \{1, N_t\}} \{\bar{S}_{it} > 0.3\}, \quad (1)$$

as well as the weight, ω_t , such that

$$\omega_t \bar{S}_{i_{t-1},t} + (1 - \omega_t) \bar{S}_{i_t,t} = 0.3.$$

Finally, small firm sales growth, s_t^S , is defined as

$$s_t^S = \omega_{t-1} \frac{\sum_{j=1}^{i_{t-1}-1} S_{j,t}}{\sum_{j=1}^{i_{t-1}-1} S_{j,t-1}} + (1 - \omega_{t-1}) \frac{\sum_{j=1}^{i_{t-1}} S_{j,t}}{\sum_{j=1}^{i_{t-1}} S_{j,t-1}}.$$

Large firm sales growth, s_t^L , is the complement of small firm sales growth. Specifically,

$$s_t^L = \omega_{t-1} \frac{\sum_{j=i_{t-1}}^{N_t} S_{j,t}}{\sum_{j=i_{t-1}}^{N_t} S_{j,t-1}} + (1 - \omega_{t-1}) \frac{\sum_{j=i_{t-1}+1}^{N_t} S_{j,t}}{\sum_{j=i_{t-1}+1}^{N_t} S_{j,t-1}}.$$

Note that the GDP deflator has no impact on the computation of the weights, ω_t , only on sales themselves. Also, note that the size categories over which sales growth is computed are the same for periods t and $t - 1$. Splicing was done when the size categories in $t - 1$ are different than they are in t .

The strategy we used to compute small and large firm sales is as follows. We compute ω_t and i_t as before. Let

$$A_{t-1}^c = \omega_{t-1} A_{i_{t-1}-1,t-1} + (1 - \omega_{t-1}) A_{i_{t-1},t-1},$$

where $A_{i_{t-1},t-1}$ denotes the level of nominal assets of the bin, i_{t-1} at $t - 1$. Denote the nominal asset cutoff at t to be

$$B_t^c = (1 + g_{t-1,t}^{NA}) A_{t-1}^c,$$

where $g_{t-1,t}^{NA}$ is defined as follows:

$$1 + g_{t-1,t}^{NA} = \frac{\text{Total nominal assets at } t}{\text{Total nominal assets at } t - 1}$$

Let

$A_{j_t} \sim$ asset bin cutoff in asset bin just above B_t^c

Let

$$\alpha_t = \frac{B_t^c - A_{j_t}}{A_{j_{t-1}} - A_{j_t}}$$

and

$$S_t^{\text{small}} = \alpha_t S_{j_{t-1},t} + (1 - \alpha_t) S_{j_t,t}$$

Then,

$$1 + g_{t,t+1}^{\text{small}} = \frac{S_t^{\text{small}}}{.3S_{t-1}^{\text{agg}}}$$

and

$$1 + g_{t,t+1}^{\text{large}} = \frac{S_t^{\text{agg}} - S_t^{\text{small}}}{.7S_{t-1}^{\text{agg}}}$$

To gain insight into these calculations, consider Figure 7a, which displays \bar{S}_{it} for $i = 1, \dots, N_t$ for each t within our data sets. The different graphs must be compared with caution, since the asset size categories differ across datasets. However, note how the distribution of production shifts increasingly towards ‘large’ firms.

Cumulative Distribution of Sales Across Asset size Categories

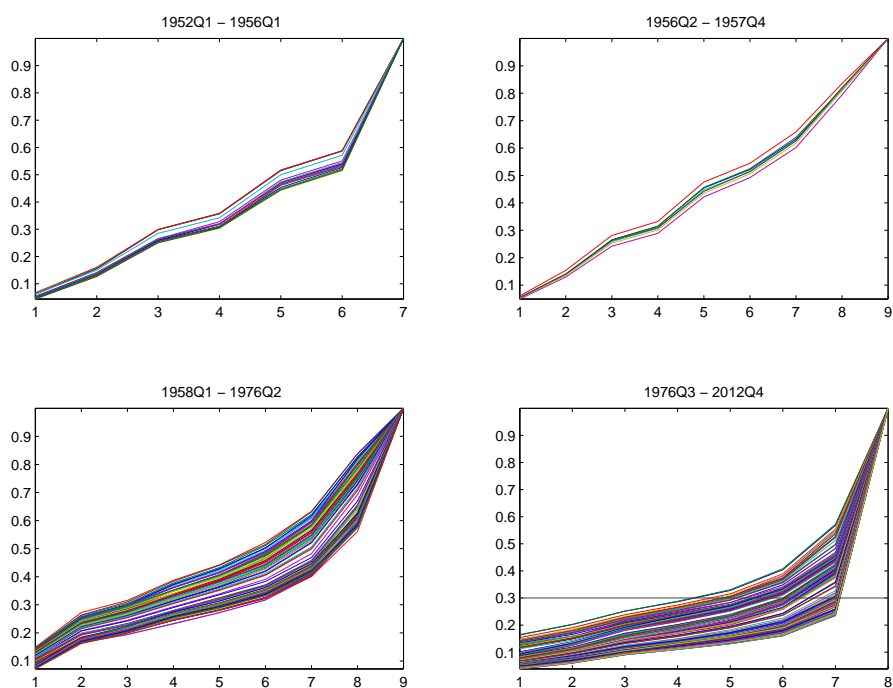


Figure 7a

A consequence of this shift is that in 2000Q4 and then in all quarters after 2003 we find $i_t = N_t$, the top asset size category. This can be seen in the following chart, which plots the fraction of sales of establishments up to the second largest establishment size, for the later part of our sample¹.

¹Let the size categories in the QFR be x_1, \dots, x_7 . As discussed in the text, we added a highest category, x_8 , using data on the firm with the greatest assets in COMPUSTAT. The object graphed in Figure 7aa is the sum of sales in categories 1 to 7, divided total sales.

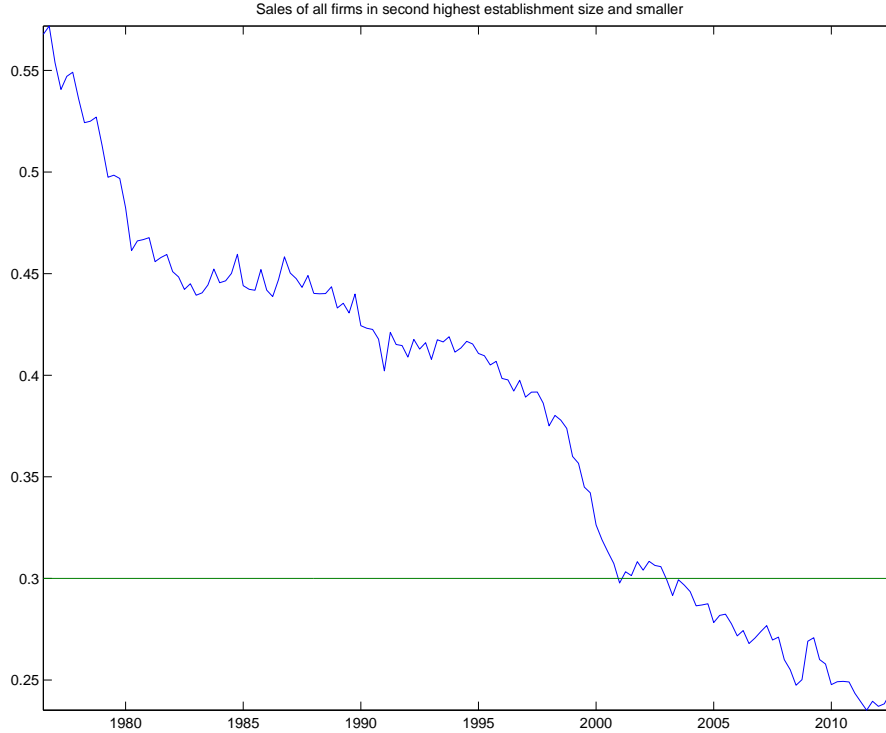


Figure 7aa

From Figure 7aa, it is evident that if we defined small firms with a 20 percent cutoff, then we would never encounter a case in which the highest size category contains the small firm cutoff. When the cutoff is contained inside the largest size category, this creates no problem for our computation of s_t^S . To compute s_t^L in this case, we set $\omega_t = 1$.

The QFR data we use were not seasonally adjusted. Consequently, the small and large firm sales growth rates have a pronounced seasonal. We seasonally adjusted the growth rates by replacing the original data with the error in the regression of the data on four seasonal dummies. We considered two strategies. In the first, we imposed that the dummies are constant throughout the sample. In the second, we allowed the dummies to take on different values before 1974 and after 1974. The latter strategy better accommodates a changing seasonal over time. In our work, we use data based on the second strategy.

Figure 7 below compares our computed sales growth data with Gertler and Gilchrist's. The two series are reasonably similar.

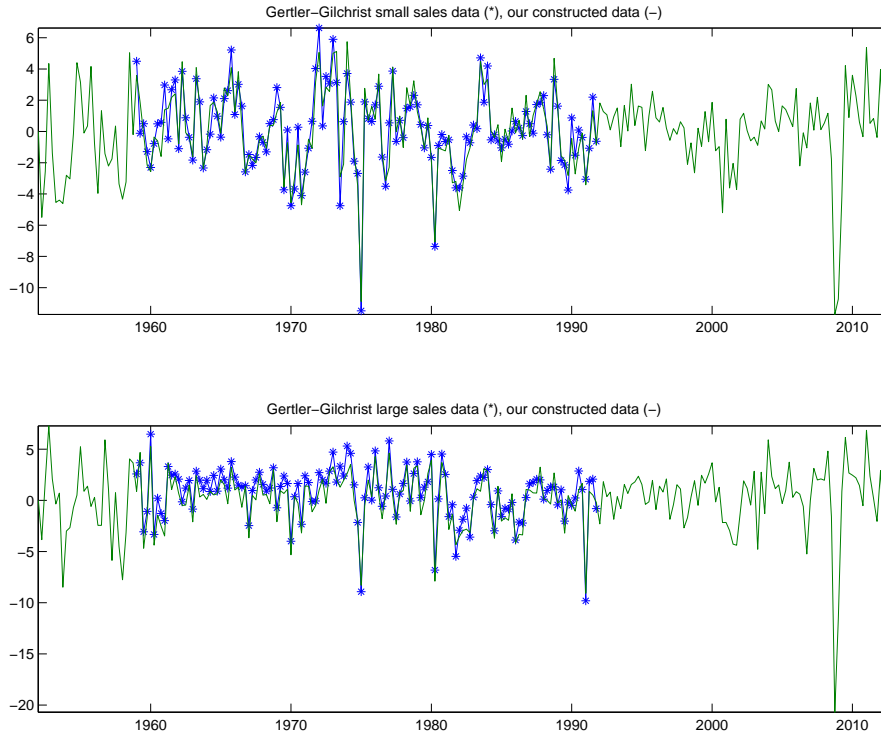


Figure 7

However, our data do reflect two adjustments that correct for suspicious spikes that do not appear in the Gertler-Gilchrist data. The adjustments were made by linear interpolation. The effects of the adjustments can be seen in Figure 7b:

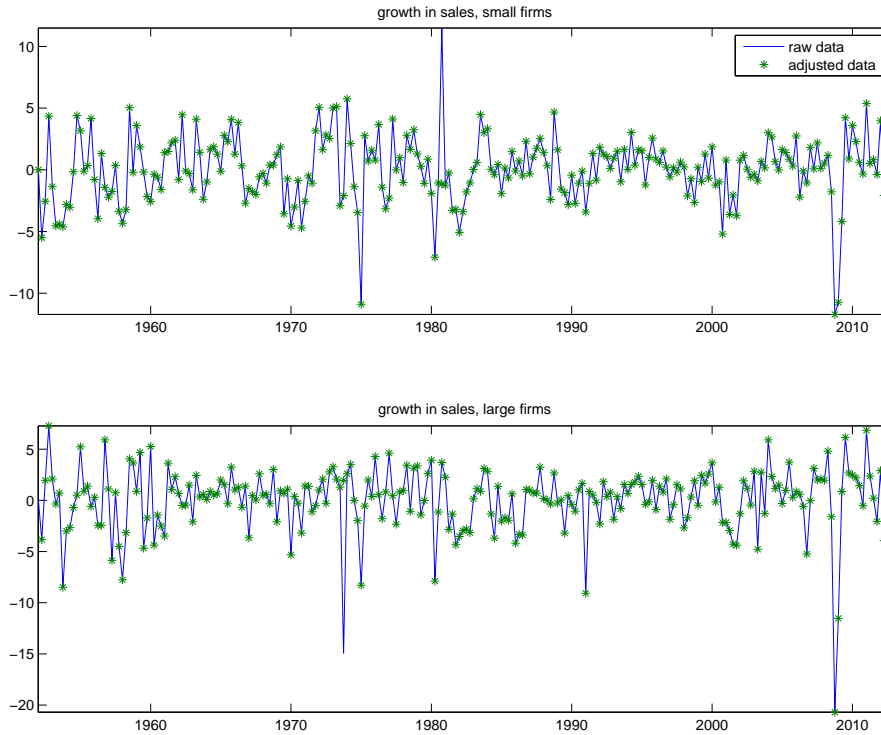


Figure 7b

Apart from the differences just mentioned, there are smaller differences between our constructed series and Gertler-Gilchrist's. Note that in the decade of the 1960s, our constructed large establishment growth rates are on average below Gertler and Gilchrist's.

Another way to compare the GG and our sales data is to first cumulate and then HP filter them. This comparison places less emphasis on the high frequency components of the data. The results are displayed in Figure 7c.

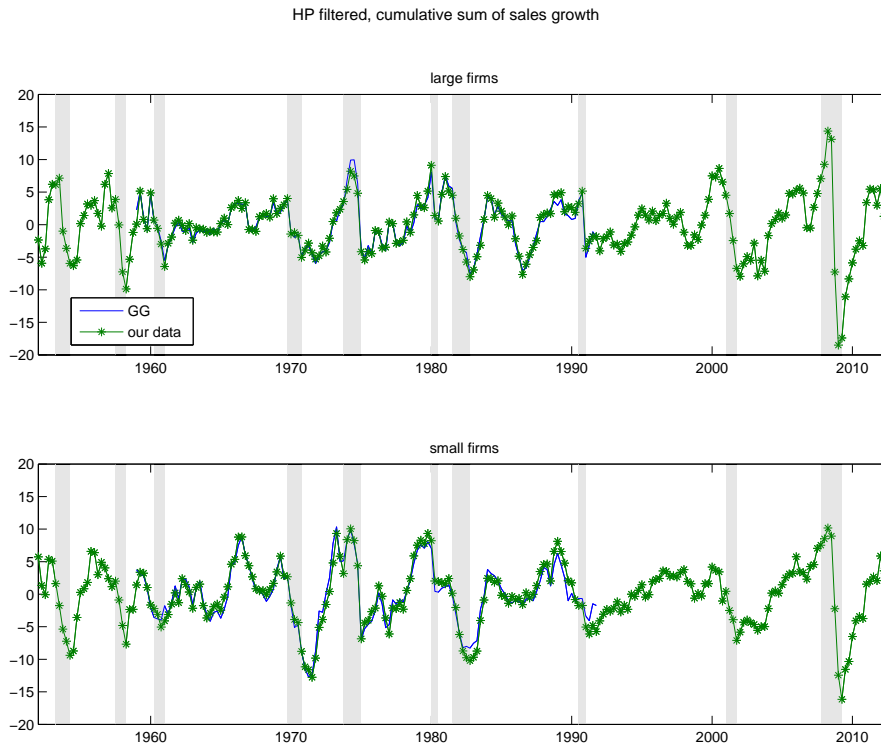


Figure 7c

Note that when viewed from this perspective the two series look very similar. Note too, how much larger the drop in large firm sales are in the most recent recession. From peak to trough of the sales data, they drop 35 percent. Small firm sales drop 25 percent from their peak to their trough during the recession. We put both sales series on the same graph, together with GDP in Figure 7d:

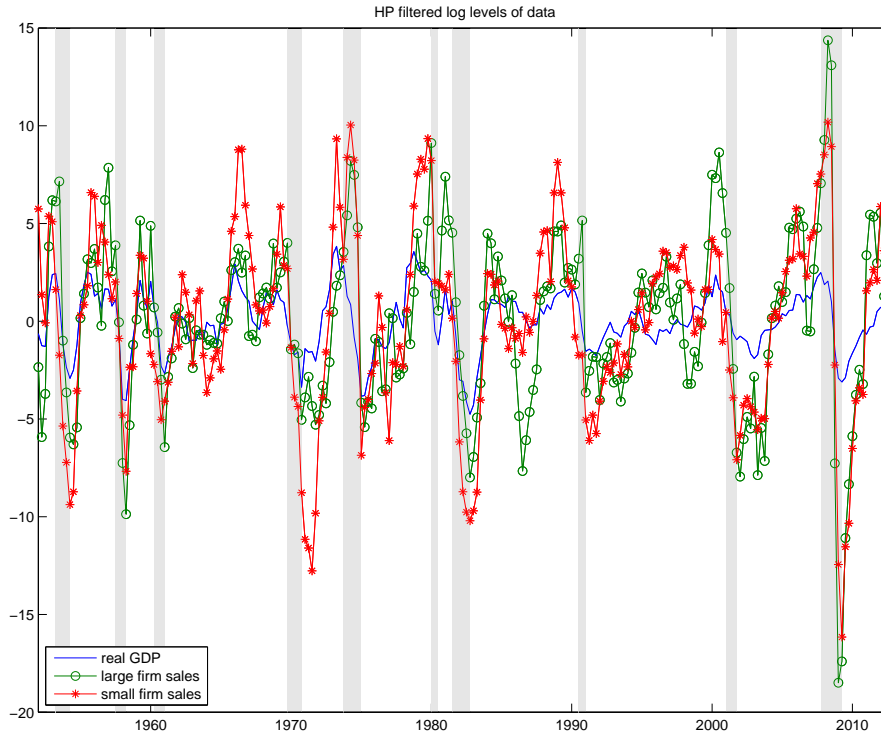


Figure 7d

A few differences are notable. First, both sales data are substantially more volatile than GDP. Second, in the neighborhood of the last two recessions, large firm sales appear to be more volatile, as we saw in Figure 7c. We investigate the relative volatility of the two sales data (after cumulating and HP filtering the data) in Figure 7e.

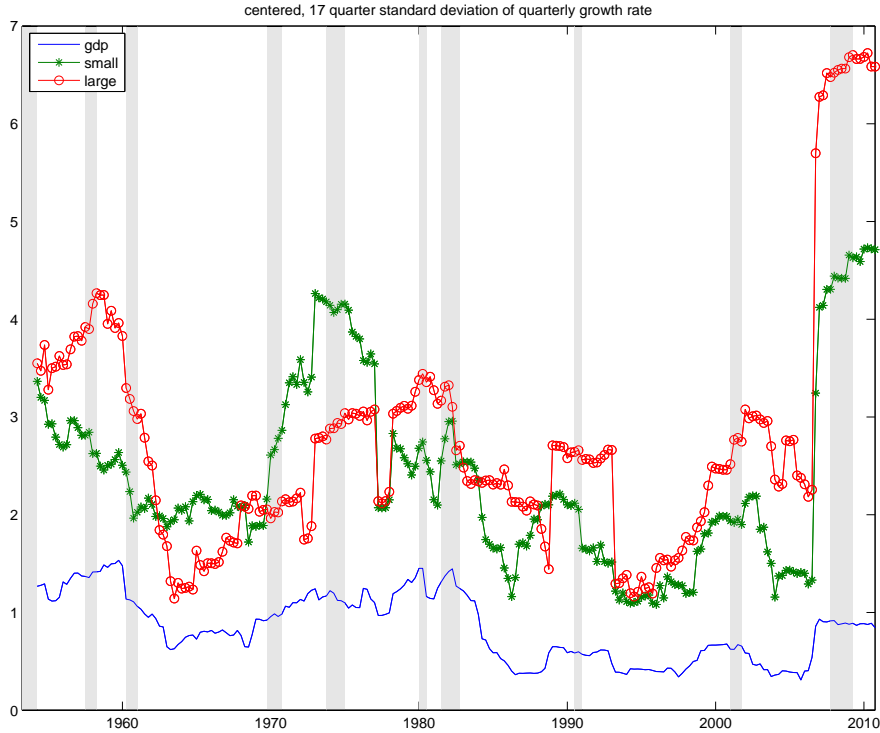


Figure 7e

There are several things to note in Figure 7e. First, both sales series are consistently more volatile than GDP. Second, the ‘Great Moderation’ observed for GDP - the reduction in volatility that started in the early 1980s - is also observed for sales. Third, during the 1959-1990 Gertler-Gilchrist sample, small firm sales are generally less volatile than large firm sales. This changed shortly before 1990, after which large firm sales are always more volatile than small firm sales. (Something worthy of additional investigation is the temporary jump in large establishment sales in the early 1990s.)

Finally, we compare the worms for each Romer-Romer date implied by our data and Gertler-Gilchrist’s data. These results are displayed in Figure 7f:

Comparing CCK and GG on Individual Worms, Romer-Romer dates and Linear Detrending

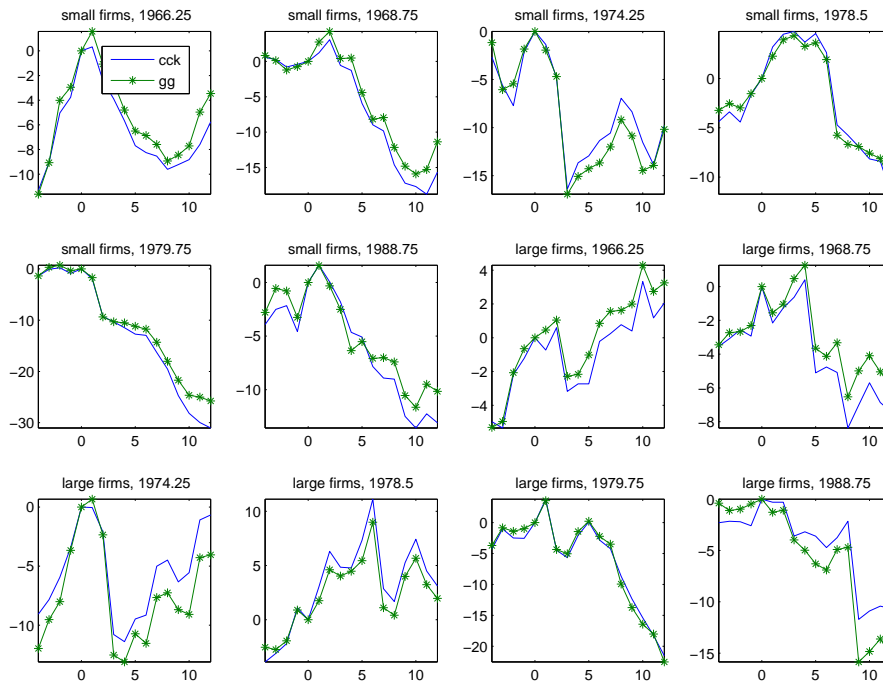


Figure 7f

These results are close.

4 Results Over the Longer Period

We begin our analysis by verifying that the Gertler-Gilchrist findings are roughly reproduced when we do their analysis on our data set. Thus, Figure 8 displays the cumulative sum of our sales data, together with their linear trend:worms for large and small firms using our data.

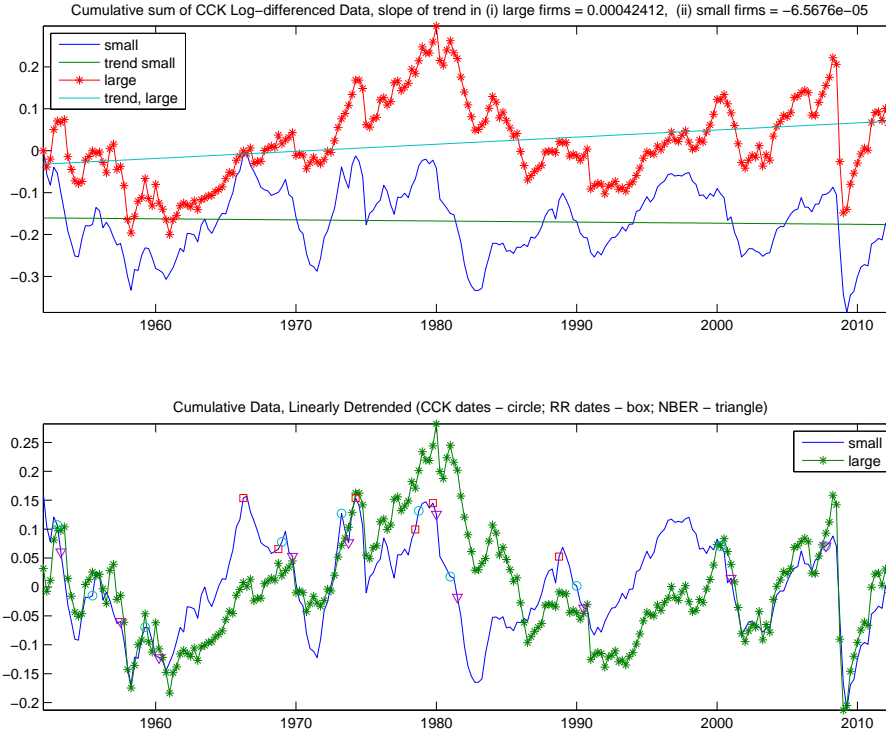


Figure 8

This figure is to be compared with Figure 2, which displays the Gertler-Gilchrist data. The trend for large firm sales is different. The bottom figure suggests that the deviations of the data from a linear trend may be similar and of course this is reinforced by the evidence in the previous section. Figure 9 displays the worm charts associated with the Romer-Romer dates using our data:

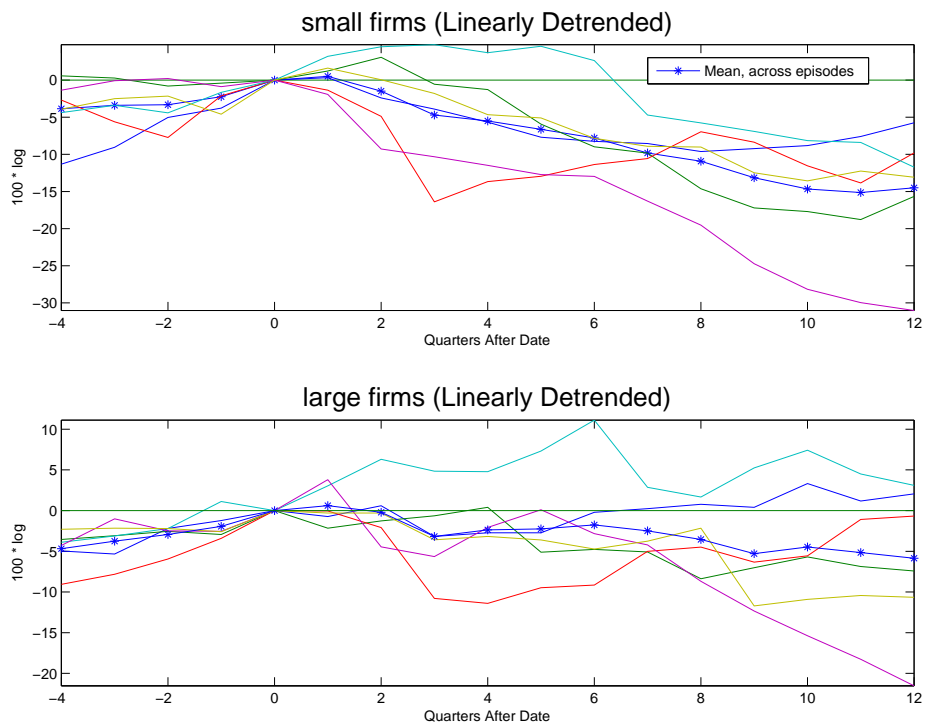


Figure 9

These results are even stronger than Gertler and Gilchrist’s results. The mean (across all episodes) drop in small firm sales is roughly 15 percent three years after a Romer-Romer date, while it is only about 5 percent for large firms. The spread is about 10 percentage points, in contrast with the spread of about 5 percent reported by Gertler and Gilchrist.

We now construct the analog of Figure 9 for CCK business cycle peaks. We obtain the following results:

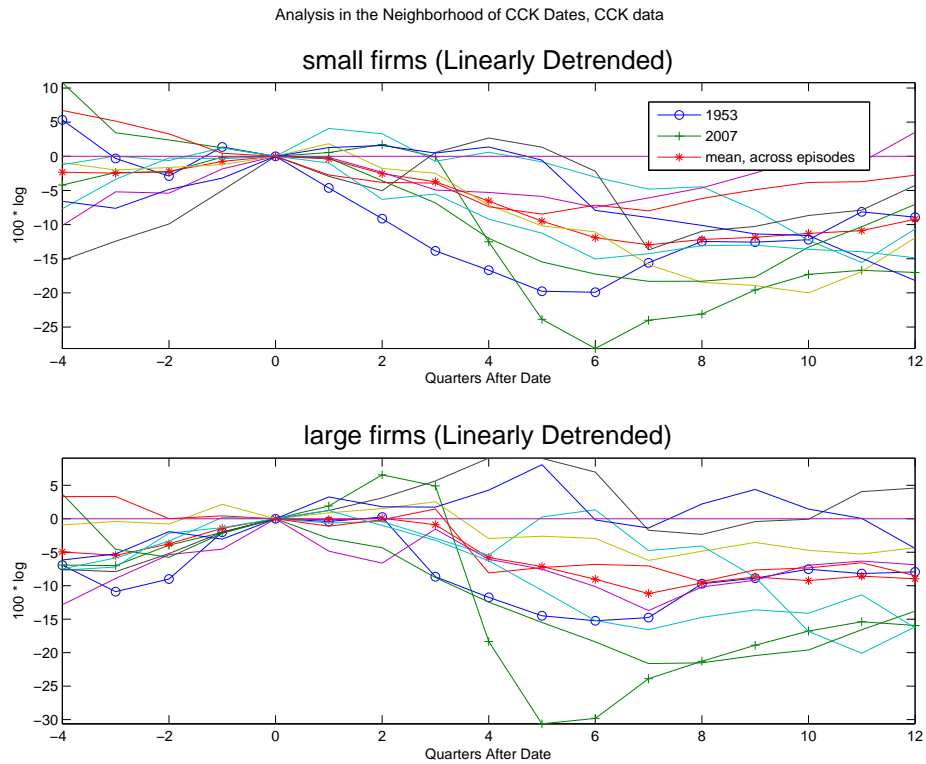


Figure 10

Now there is virtually no difference between the average responses of large and small firm sales.

To investigate robustness to detrending, we also computed the worms by detrending using the HP filter. The raw data and the HP trend and cycles for large and small firm sales are displayed in Figure 11:

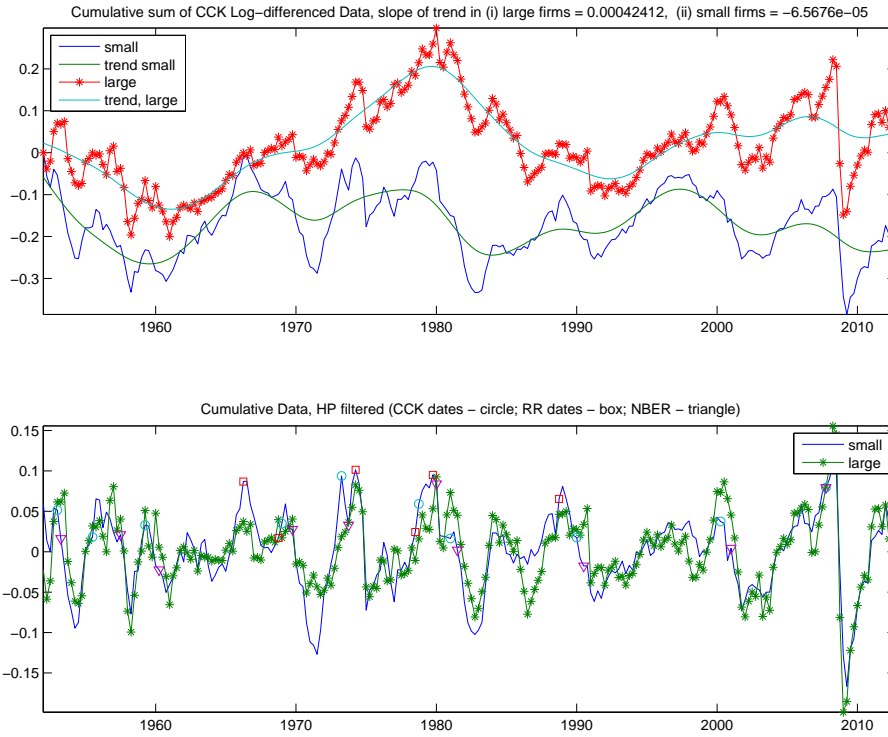


Figure 11

The worms based on HP detrending and CCK business cycle peaks are as follows:

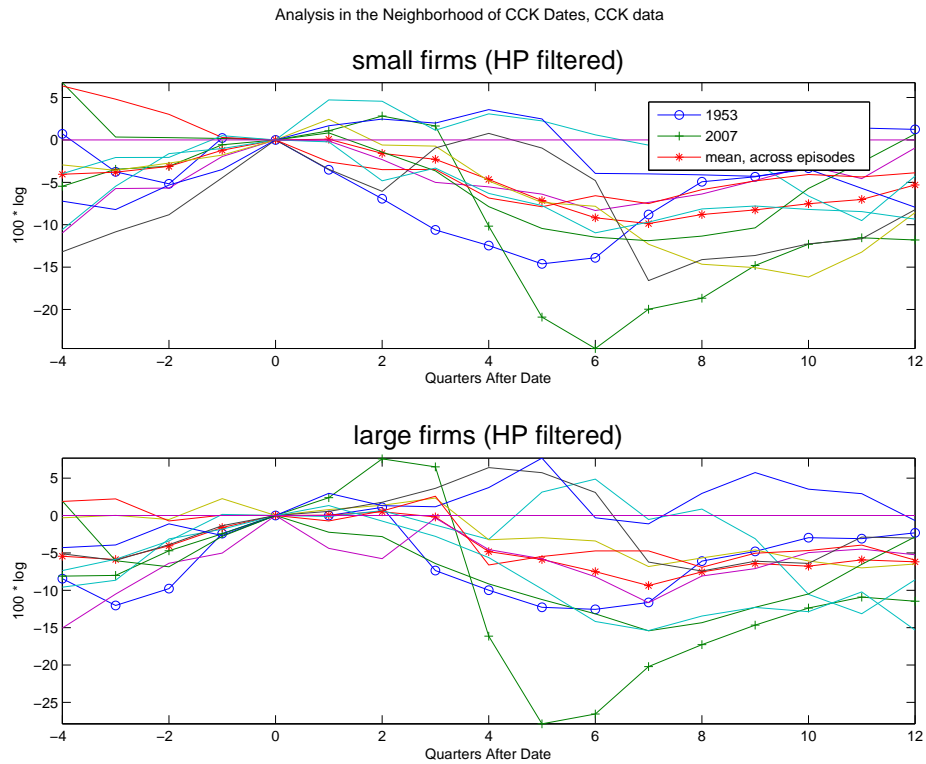


Figure 12

The results are similar to what we obtained with linear detrending. The drop, three years out from a business cycle peak, in large and small firm sales is roughly 5 percent.

We also examined the behavior of woms when the business cycle peaks correspond to NBER business cycle peaks. These are displayed in Figure 13:

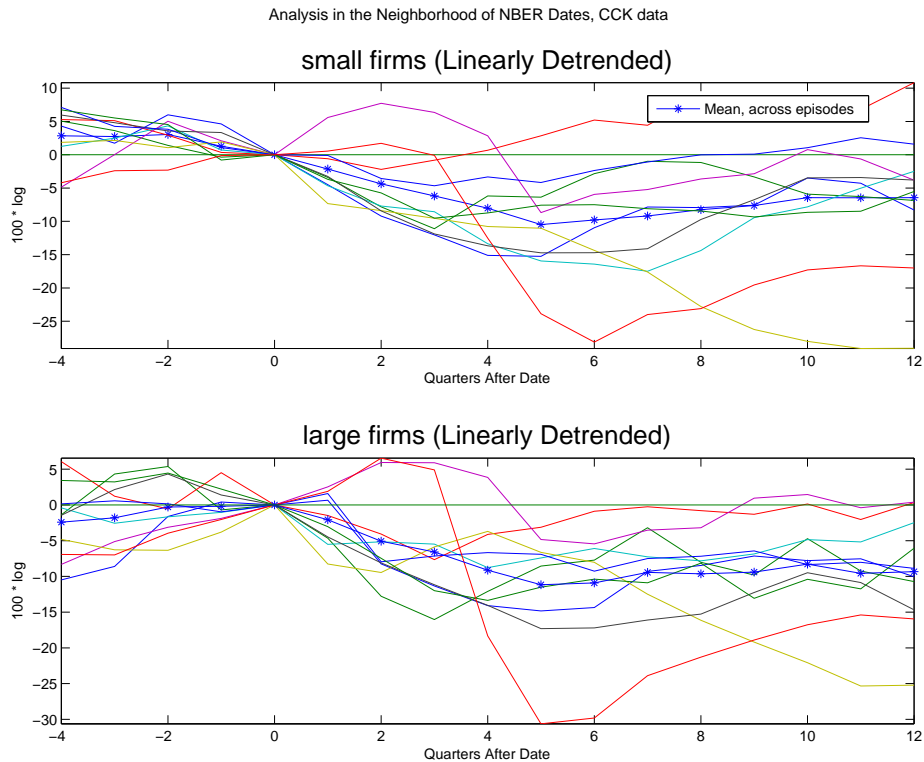


Figure 13

These results are similar to the ones based on the Gertler-Gilchrist data reported in Figure 6a. In particular, the large firms fall more, on average, after an NBER peak.

5 Data Sources and Collection

Our data on small and large firms were compiled from the Quarterly Financial Reports of Manufacturers. After 1988, these are available in electronic form from the ???. Prior to 1988, the data are not available in electronic form. These data were obtained from various issues, from 1950-1988, of the Quarterly Financial Reports. The electronic data can be used as presented, while the data before 1988 involve some complications.

Each quarterly report contains data for five quarters: the current and four preceding quarters. The data are subject to two kinds of revisions. We refer to these as numerical and methodological revisions. The numerical revisions correct errors such as reporting and data entry errors made by the data collecting authority. The methodological revisions involve changes in survey and sampling methodology and these are indicated with footnotes in the Quarterly Financial Reports.

Our procedure for computing growth rates for small and large establishments is described in section 3. That procedure involves comparing some variable (assets, sales, inventories, or employment) in two quarters, t and $t + 1$. The general principle that we follow is that we incorporate numerical revisions by using the most recent information on a variable, while ensuring that the same methodology is used for both quarters. The most recent information about quarters t and $t + 1$ appears in the quarter $t + 4$ and $t + 5$ issues of the Quarterly Financial Reports. In comparing t and $t + 1$, we take the quarter t observation from the most recent observation, the one reported in the QFR for quarter $t + 4$. With one exception, we take the quarter $t + 1$ observation from the $t + 5$ QFR. The exception is if there has been a methodological revision in the quarter $t + 5$ QFR. In this case, we take both the t and $t + 1$ observations from the $t + 4$ QFR. We do this in order to ensure that the period t and $t + 1$ observations are comparable in the sense of being reported using the same methodology. The problem is that the quarter $t + 5$ QFR does not report the revised quarter t observation. Note that the date t value of the variable used in comparing periods t and $t + 1$ may be different from that in comparing periods $t - 1$ and t .