The Gertler-Gilchrist Evidence on Small and Large Firm Sales

VV Chari, LJ Christiano and P Kehoe

January 2, 2007

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 2000 recession. 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 the Gertler-Gilchrist places special emphasis on several particular dates. We discuss these here. The first figure below reports log US real GDP, after HP filtering. The observations market '*' come in couples, with the first one indicating an NBER business cycle peak and the second a trough. Note how these peaks typically occur several quarters after an actual turning point in the data.

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 are local peaks in detrended output, after which output fell by at least several percent. The dates are 1953Q2, 1955Q3, 1960Q1, 1966Q1, 1969Q1, 1973Q2, 1978Q4, 1981Q1, 1990Q1, 2000Q2. There were two dates, 1955Q3 and 1981Q1, that we were unsure of. In the first of these episodes, the economy fell by less than 2 percent in the subsequent year, and then exhibited a brief period of strength in 1956Q4, before then plunging by several percent. We could have perhaps picked 1956Q4 as the business cycle peak instead. The second 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.

The boxes indicate Romer and Romer dates. These occur in 1966, Quarter 2, December 1968, April 1974, August 1978, October 1979, and August 1988. 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.

Figure 1

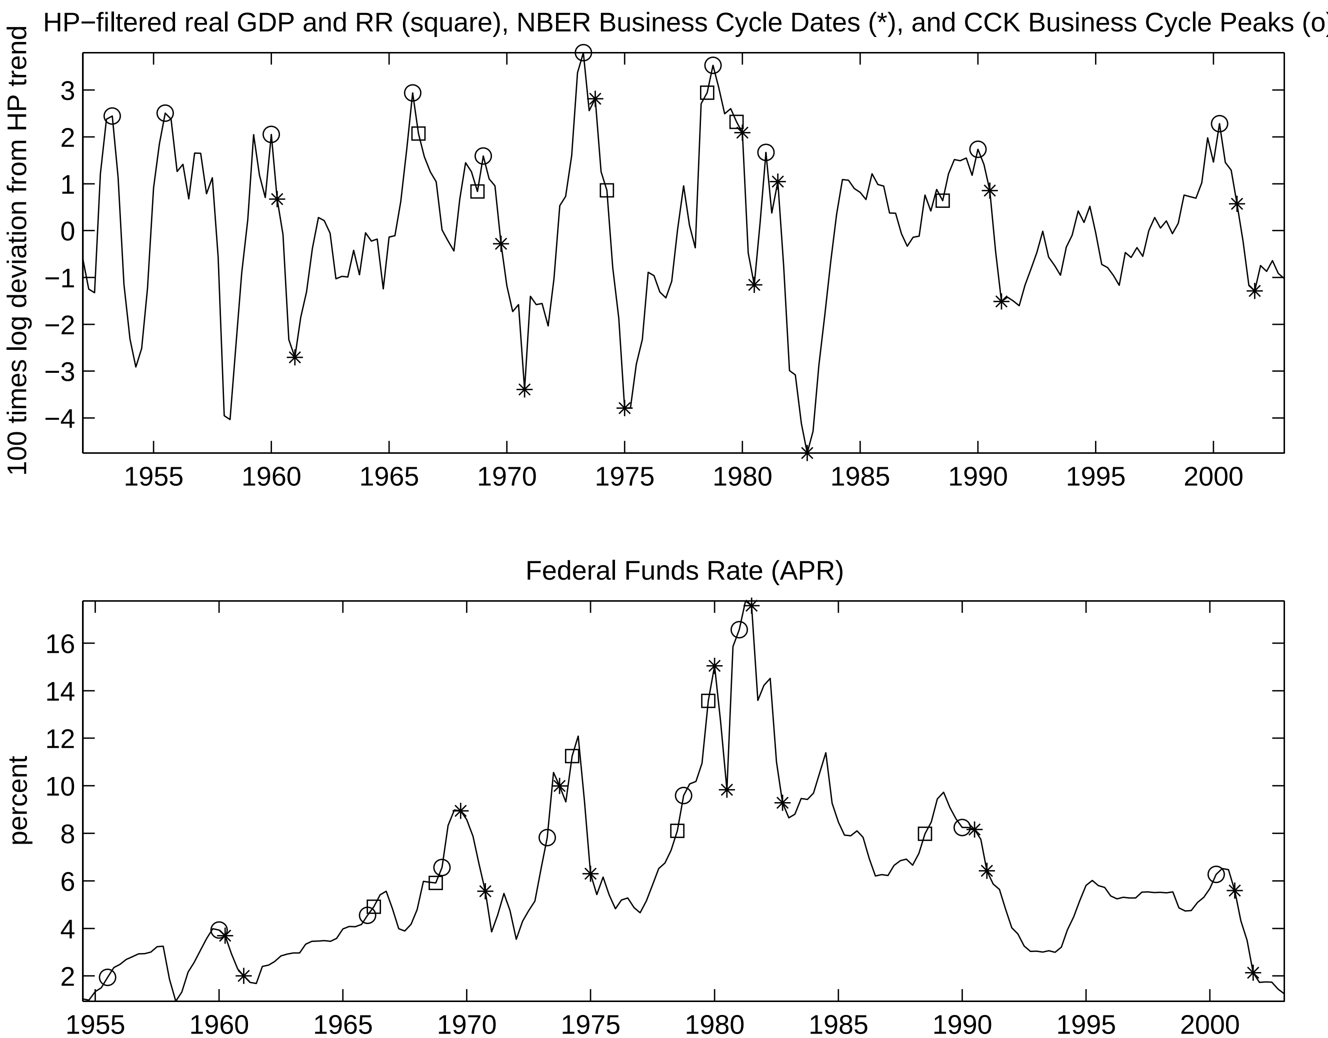

The second graph shows the federal funds rate. Note that the Romer and Romer dates occur while interest rates are rising. Two dates - April 1974 and October 1979 - look somewhat suspicious, as dates when the Fed resolved to initiate a battle against inflation. Each of these two dates occur late during a sharp rise in the interest rate. One might expect that a Fed decision to sharply raise the interest rate would occur near the beginning of an upswing in the interest rate, not near the end.

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 below. 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 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

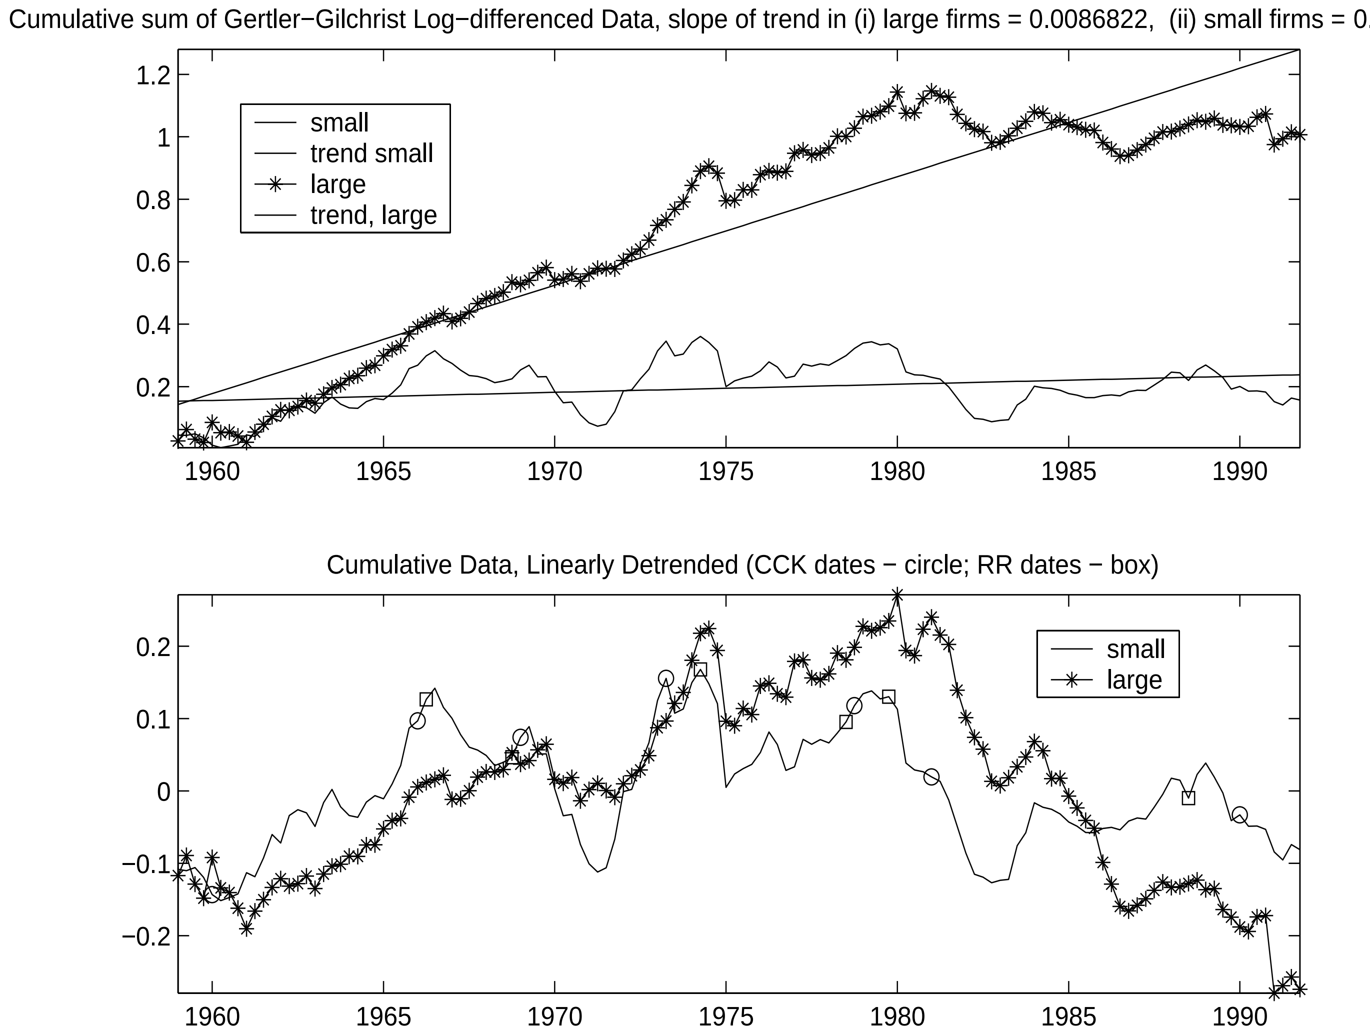

Consider the Romer-Romer dates (i.e., the squares) first. After the one in 1966 it is clear that small firm sales drop by more than large firm sales. The same is true for the 1968 date, although initially there is a transient rise in small firm sales. After 1974 small firm sales fall more than large firm sales, but the difference looks relatively small. After the 1978 date there is actually a surge in small firm sales.

But, the surge in large firm sales is a bit stronger and the small firms slow down and falter before the large firms do. After 1979, the small firms seem to go down faster and more than the large firms. However, one could perhaps debate this by pointing to the fact that large firms ultimately fall much further. Also, the part of the 1978 date that is consistent with GG is actually the same at the 1979 date. So, counting 1978 and 1979 separately seems like double counting. In the wake of the 1988 date, it is not clear which goes down more.

We produced 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

Analysis In Neighborhood of Romer and Romer Dates GG data

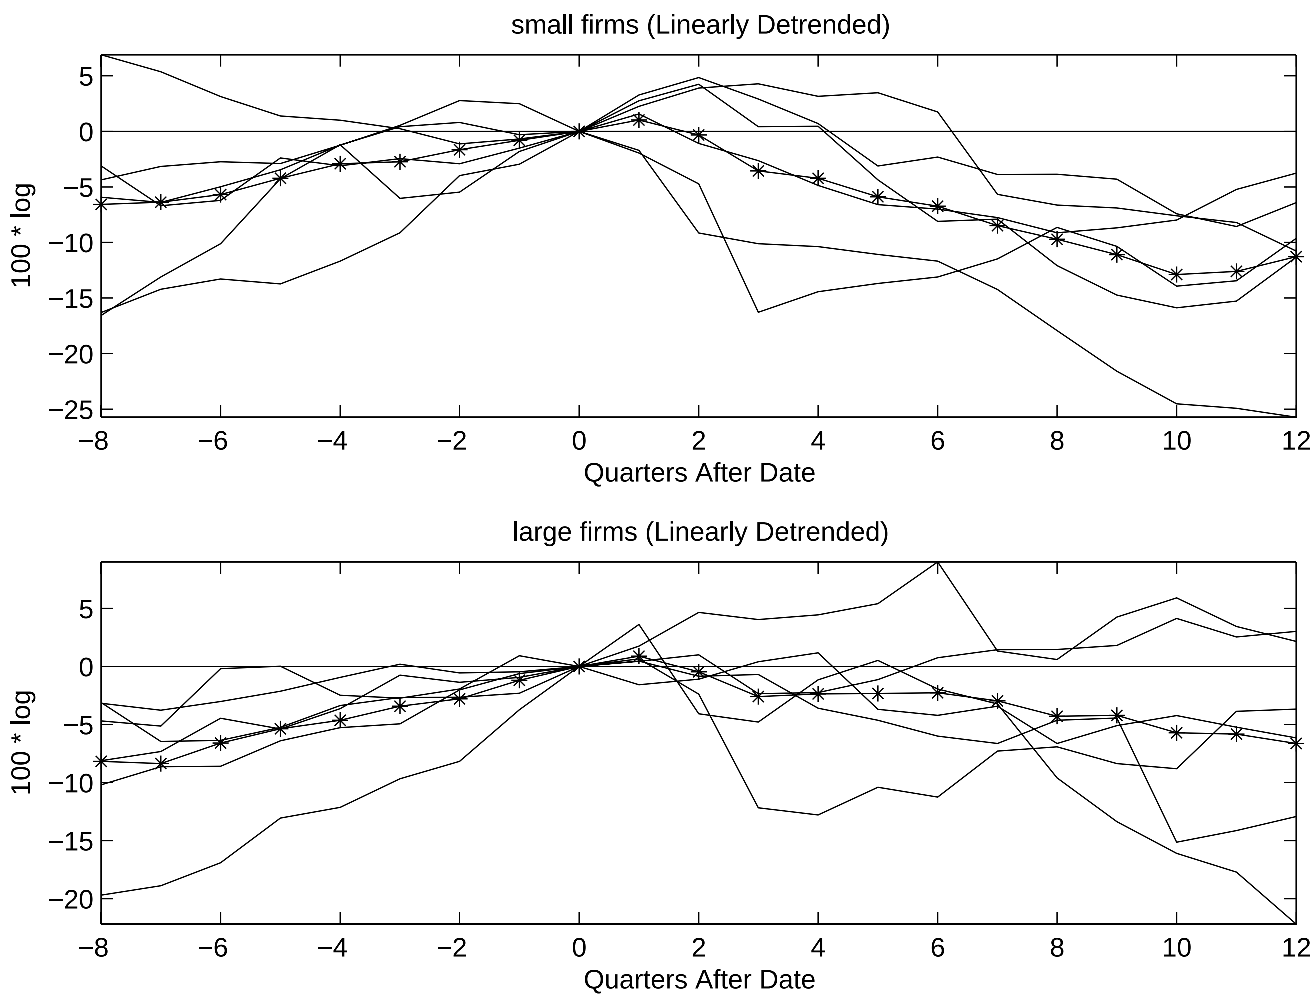

We investigated the robustness of the above finding to the detrending strategy. This is motivated by the evidence in Figure 2 that the trend in the levels of large and small firms is not well captured by a simple linear trend. Repeating the calculations in Figure 2, using the HP filter to detrend the cumulated levels, we obtain the following results:

Figure 4

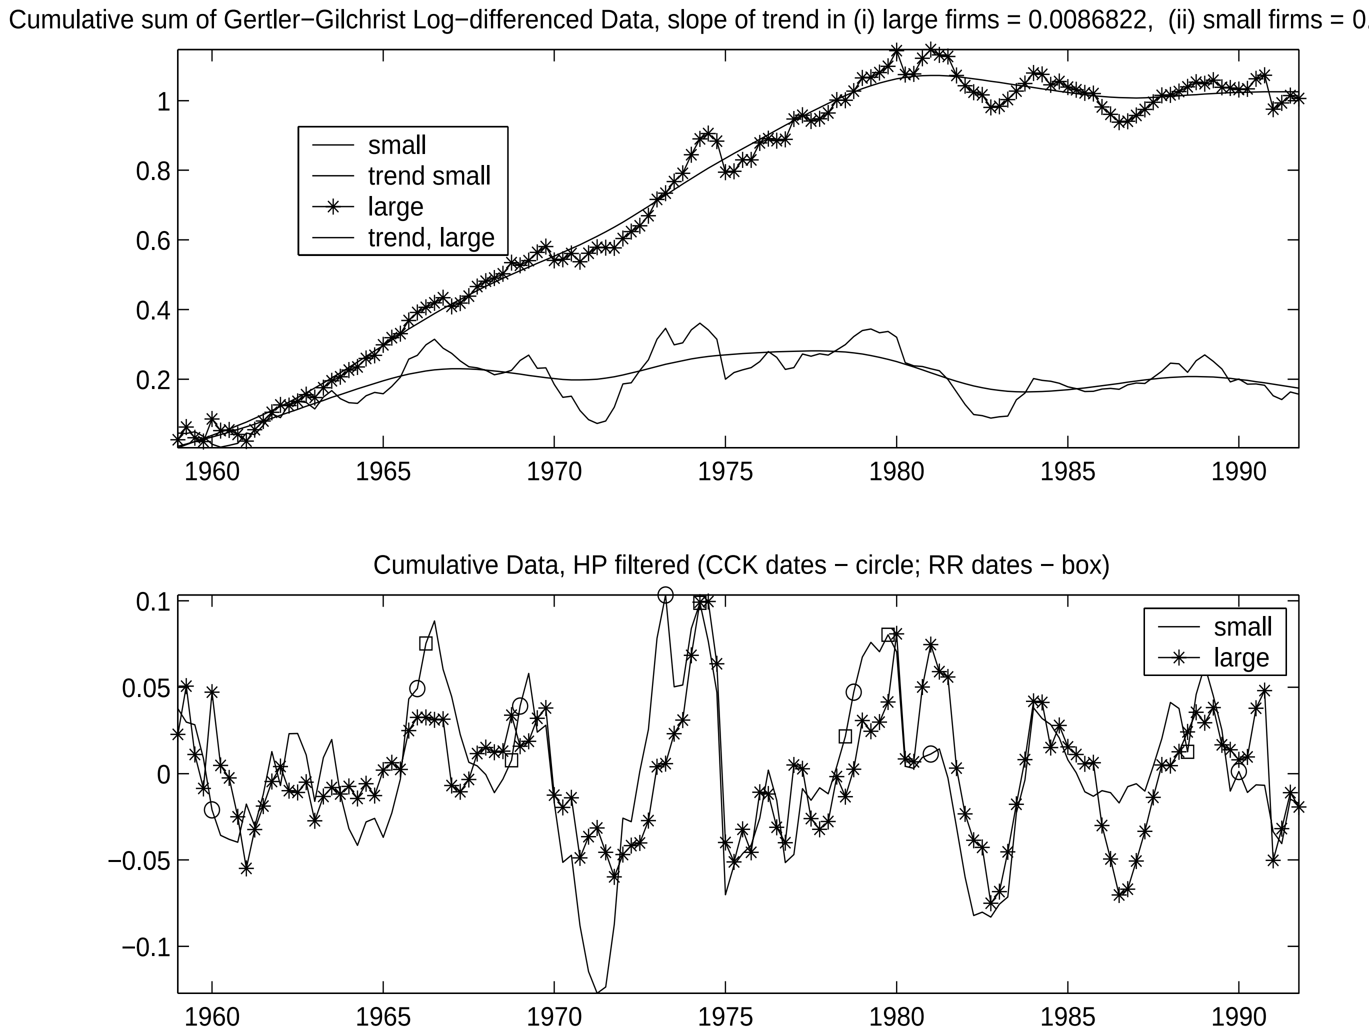

The behavior of detrended large firm sales is somewhat sensitive to the detrending procedure, because of the apparent break in the trend of that series around 1980. As before, small firm sales drop a lot after the 1966 date. There seems to be a bigger surge in small firm sales after the 1968 date, though overall perhaps it can be said that after that date small firm sales drop by a lot more than large firm sales.

After the 1974 date, the fall in the two sales figures seem to be of similar magnitude. The 1978 data is much less consistent with the GG hypothesis with the change in detrending, because the surge in sales by small firms thereafter is so great. In the immediate aftermath of the 1979 date, large and small firms seem to react in a similar way. This too, seems not consistent with the GG hypothesis. The fact that after 1985 large firm sales are more volatile than small firms sales also seems inconsistent with the spirit of the GG hypothesis. In the period after the 1988 Romer-Romer date, the overall movements in large and small firm sales are similar, except that the former are still more volatile than the latter. Overall, the results seem to favor the GG hypothesis less strongly with the change in detrending procedure.

Somewhat inconsistently with our informal analysis of Figure 4, when we compute the worm chart using the data in Figure 4, we obtained once again results that favor the GG conclusion (see Figure 5). 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.

Figure 5

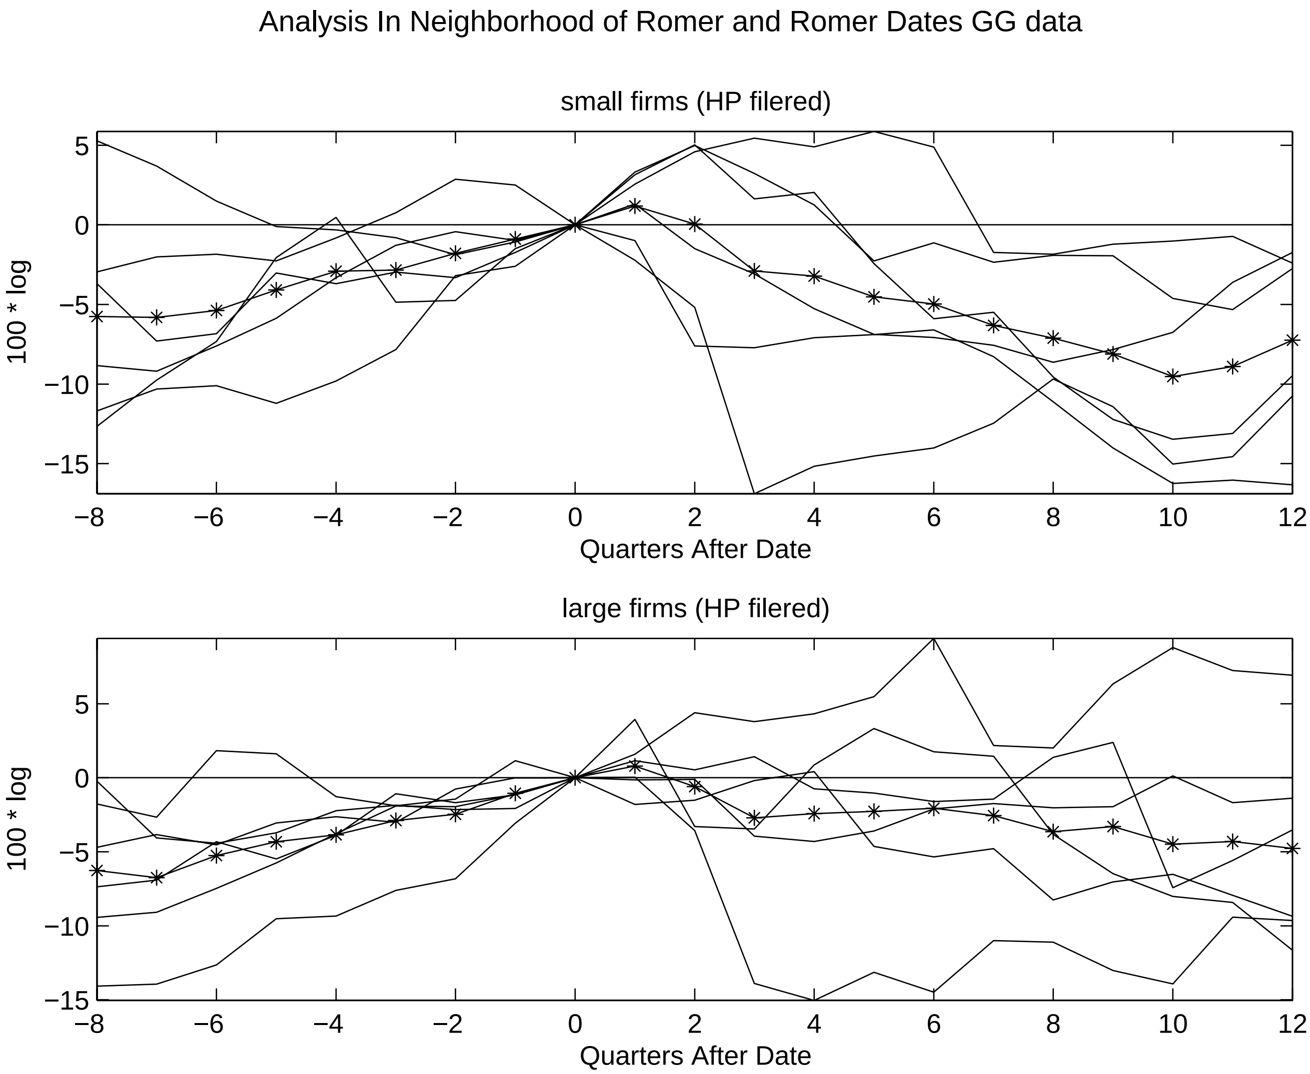

We repeated the worm chart calculations for CCK business cycle peaks and obtained the following results:

Figure 6

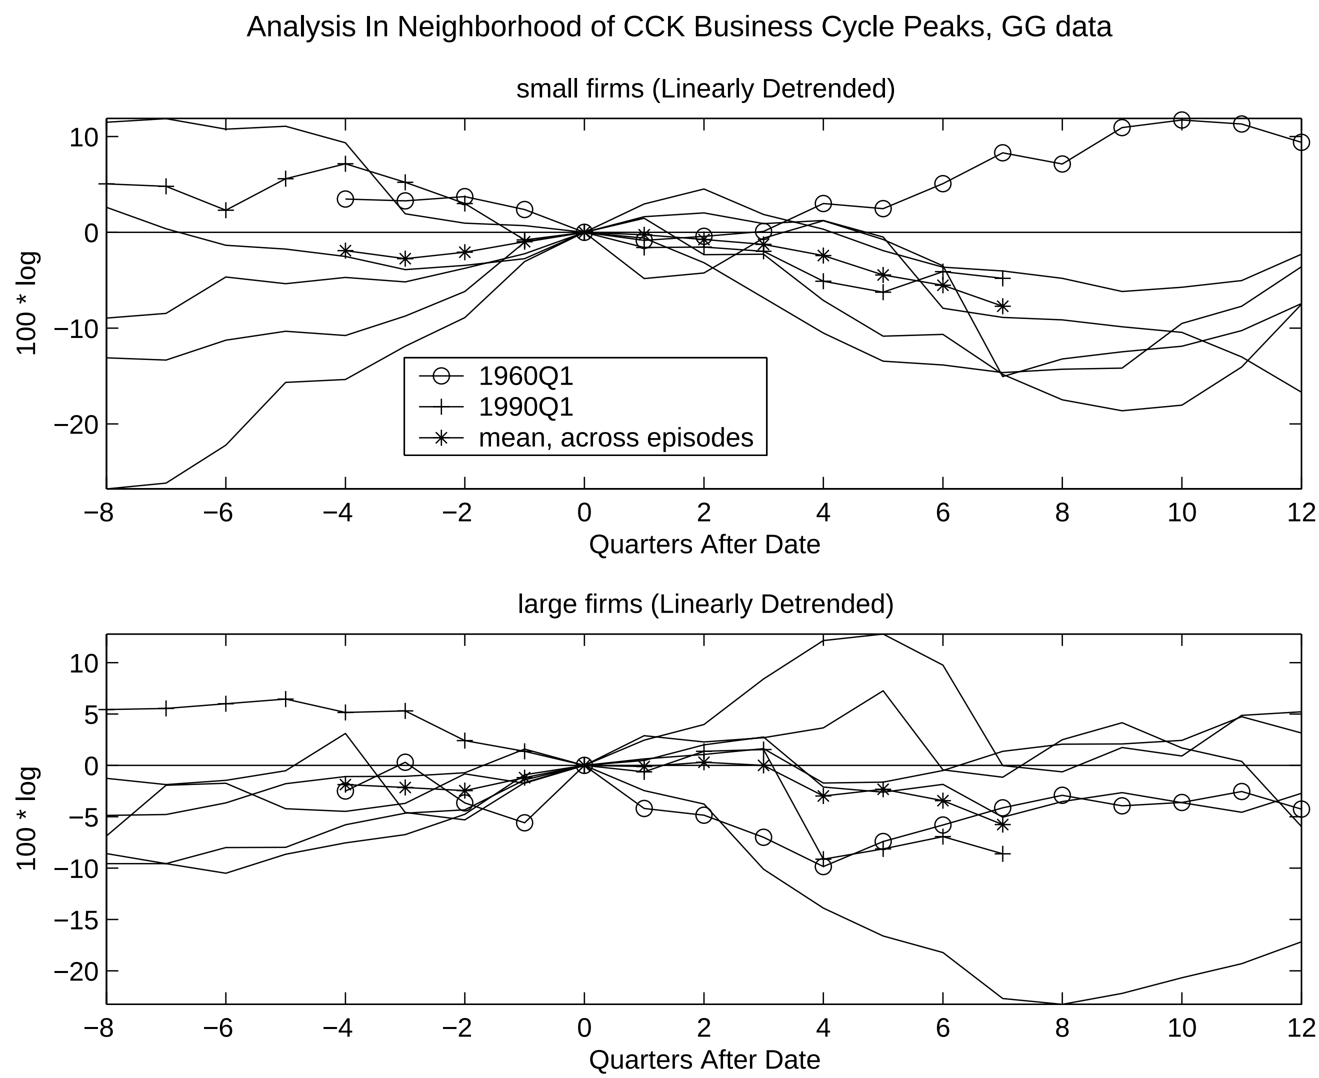

Note that the starred line does not extend to the end or to the beginning of the responses. This is because the first and last business cycle peaks are too near the ends of the data set. It seems fair to conclude that these results are somewhat weaker than the Gertler-Gilchrist results, because there 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.

3. Bringing More Business Cycle Episodes into the Analysis

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 2000Q3 by recovering data from various QFRs. We constructed the data using the simplest version of the procedure applied in Gertler-Gilchrist. 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 occuring 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.

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 $i = 1, ..., N_t$. We then compute

$$i_t = \min_{i \in \{1, N_t\}} \left\{ \bar{S}_{it} > 0.3 \right\}, \tag{3.1}$$

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

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

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

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

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.

An alternative strategy for computing ω_t in (3.1) would have used total assets instead of total sales to construct $\bar{S}_{it}$. In this case, firm size would have been based purely on assets.

To gain insight into these calculations, consider Figure 7a, which displays $\bar{S}_{it}$ for $i = 1, ..., N_t$ 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. Indeed, in the last date of our last dataset, we find $i_t = N_t$, the top asset size category. In this case, for the purpose of constructing s_t^L we set $\omega_t = 1$.

Figure 7a

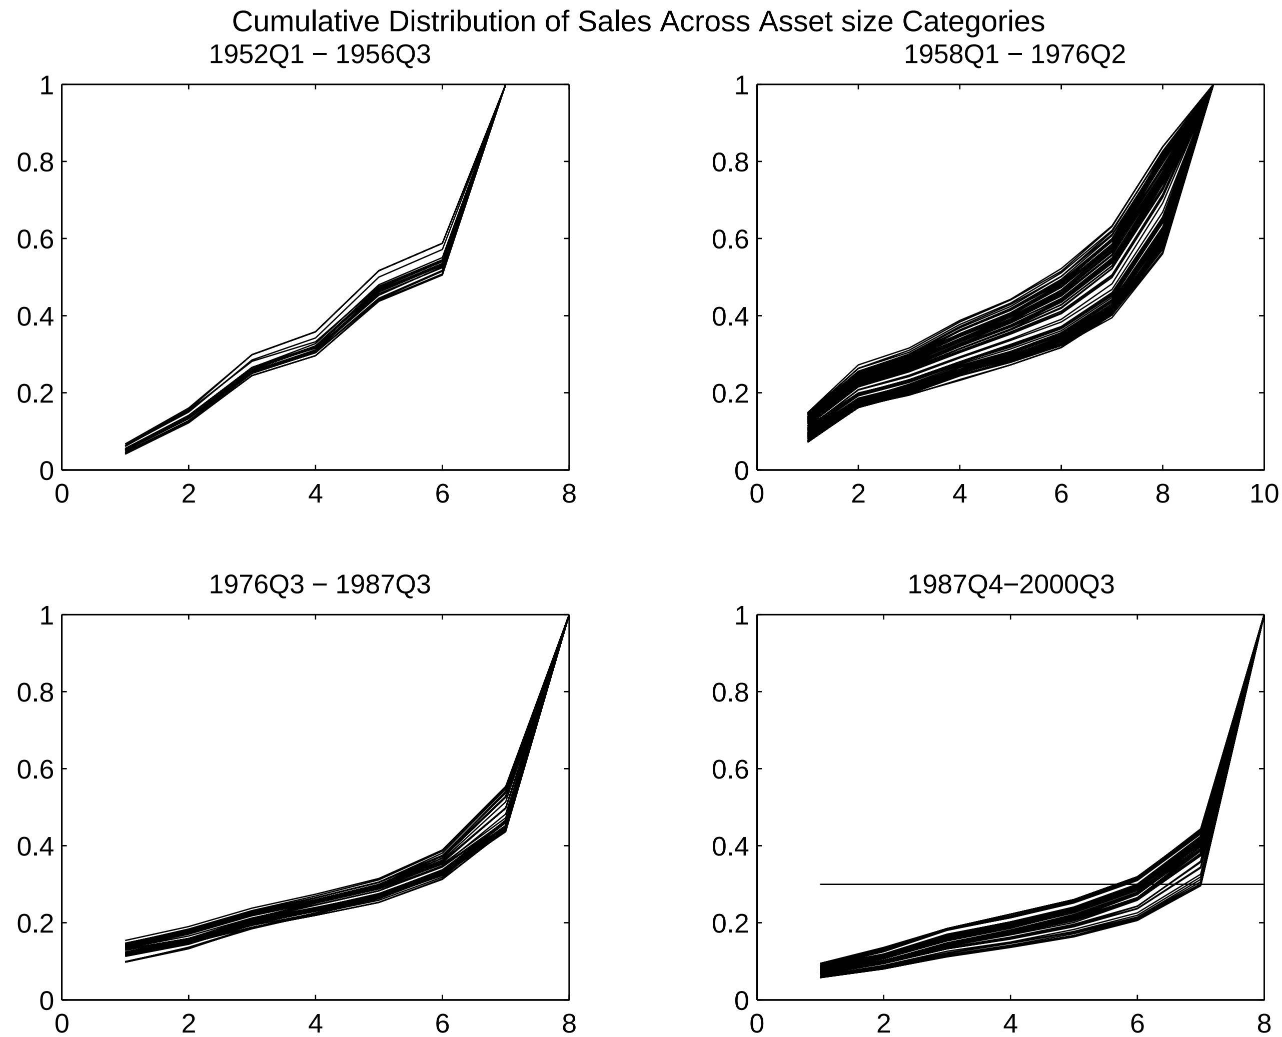

Figure 7 below compares our computed sales growth data with Gertler and

Gilchrist's. The two series are reasonably similar. The ratio of the standard deviation of our measure of small firm sales growth, to the standard deviation of Gertler and Gilchrist's measure is 1.04. The corresponding statistic for large firm sales growth is 0.93. The correlation between the two measure of small sales growth is 0.81 and the correlation of the two measures of large firm sales growth is 0.86. (These statistics were computed over the sample that the two data series share.)

Figure 7

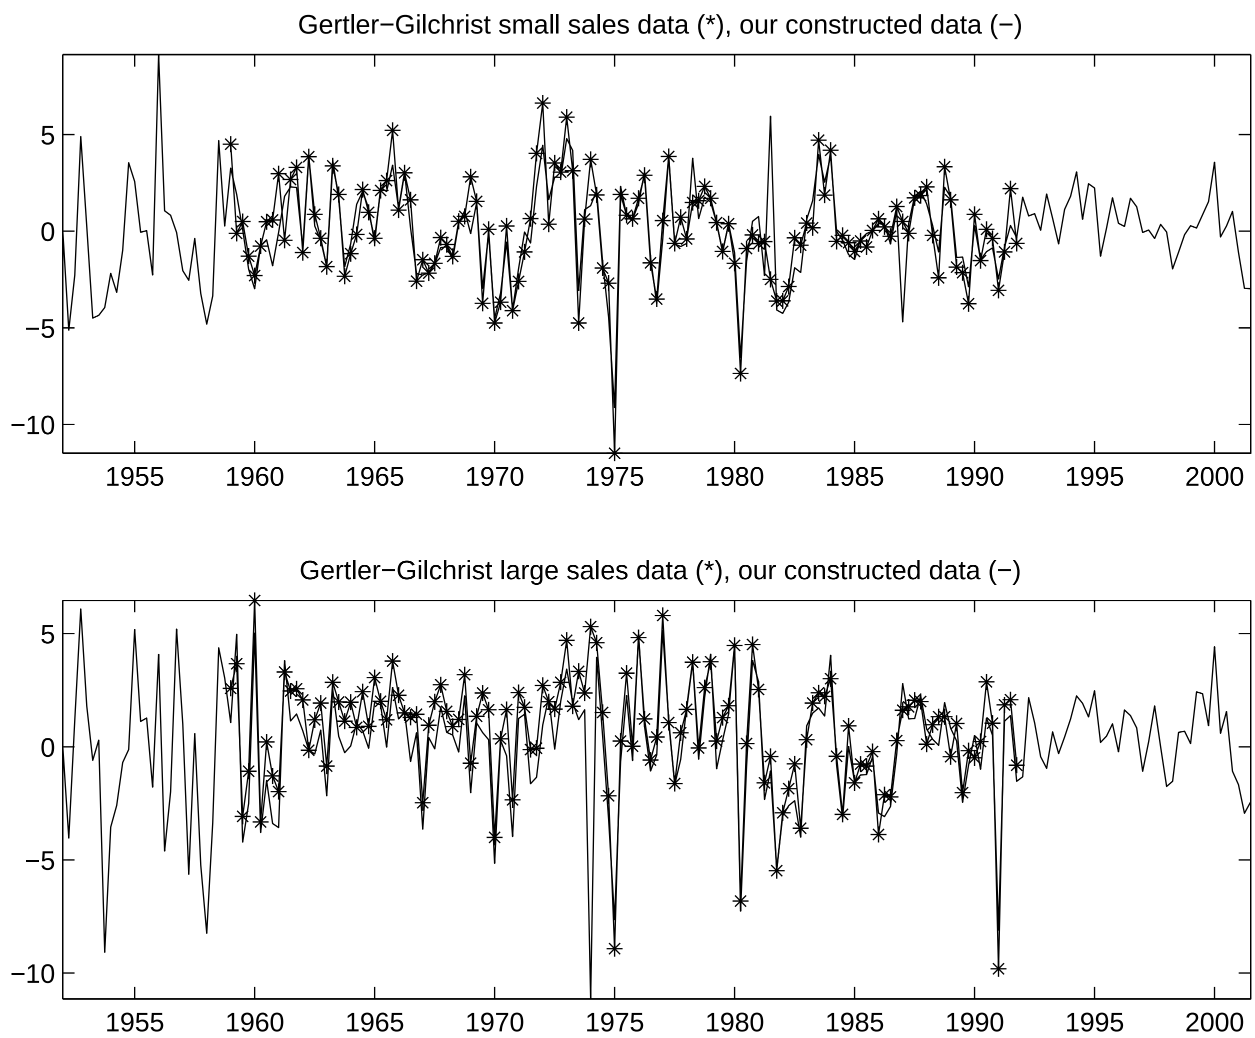

There are a few isolated instances in which there is a sharp difference between the two estimates of sales growth. In each case, the differences reflect a sharp, suspicious, jump in our constructed series. In the case of small firm sales growth, there are two such instances. The small firm sales growth jump in 1981Q3 reflects that there are 25 and 30 percent jumps in nominal sales in the second and third-smallest size categories in 1981Q3 over 1981Q2 (these growth rates are expressed at quarterly rates!) . The small firm sales growth drop in 1987Q1 reflects that there is a 24 percent drop in nominal sales of firms in the smallest size category (<5 million in assets) 1987Q1 over 1986Q4 (again, at a quarterly rate). We suppose that these changes reflect measurement error or data revision. Turning to the growth rate of large firm sales, note that there is one instance, 1974Q1, in which our constructed measure drops sharply. This reflects that there is a 21 percent drop in nominal sales by the largest size category of firms in 1974Q1. Sales then jump back up by 13 percent in 1974Q2, and so we suspect that this represents measurement error. When replaced these suspicious observations with linearly interpolated values, the spikes were attenuated, but not greatly.

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 small firm growth rate is on average below Gertler and Gilchrist's.

We now look at the HP-filtered data. The top panel displays the cumulated level of our measure of the sales of large and small firms after HP-filtering. HP filtered, log GDP is also displayed. Note that the volatility of the two series is very similar. Small firms seem to exhibit more volatility in the 1960s and 1970s, but in the rest of the sample there are times when large firms are more volatile. The bottom panel shows the same results for the Gertler-Gilchrist data. Note that here there is somewhat more that small firms exhibit greater volatility. The relative volatility of small firm sales is 2.8 versus 2.1 for large firms. In part, this reflects the different samples, however. When we recalculated the relative volatility of our measure of sales over the Gertler-Gilchrist time period, we obtained that small firms' relative volatility is 2.3 and large firms' volatility is 2.6.

Figure 7b

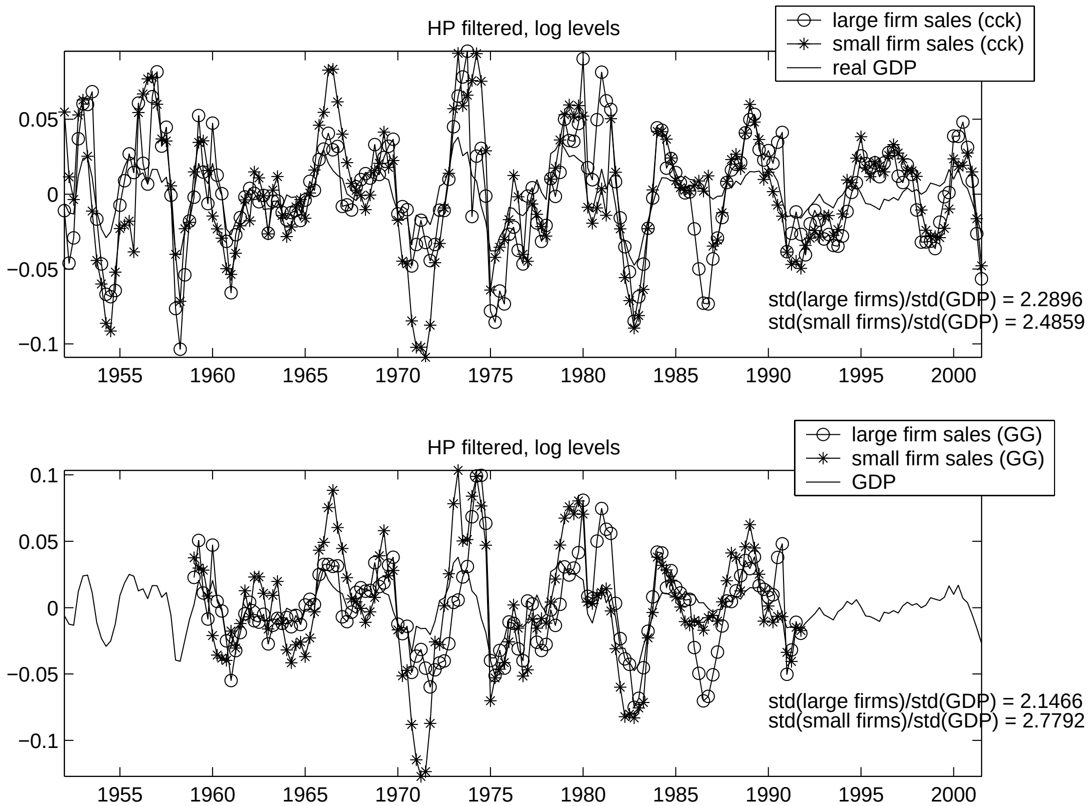

We did a direct comparison, based on HP filtered levels data, between the Gertler-

Gilchrist and our measures of sales. These are displayed in Figure 7c:

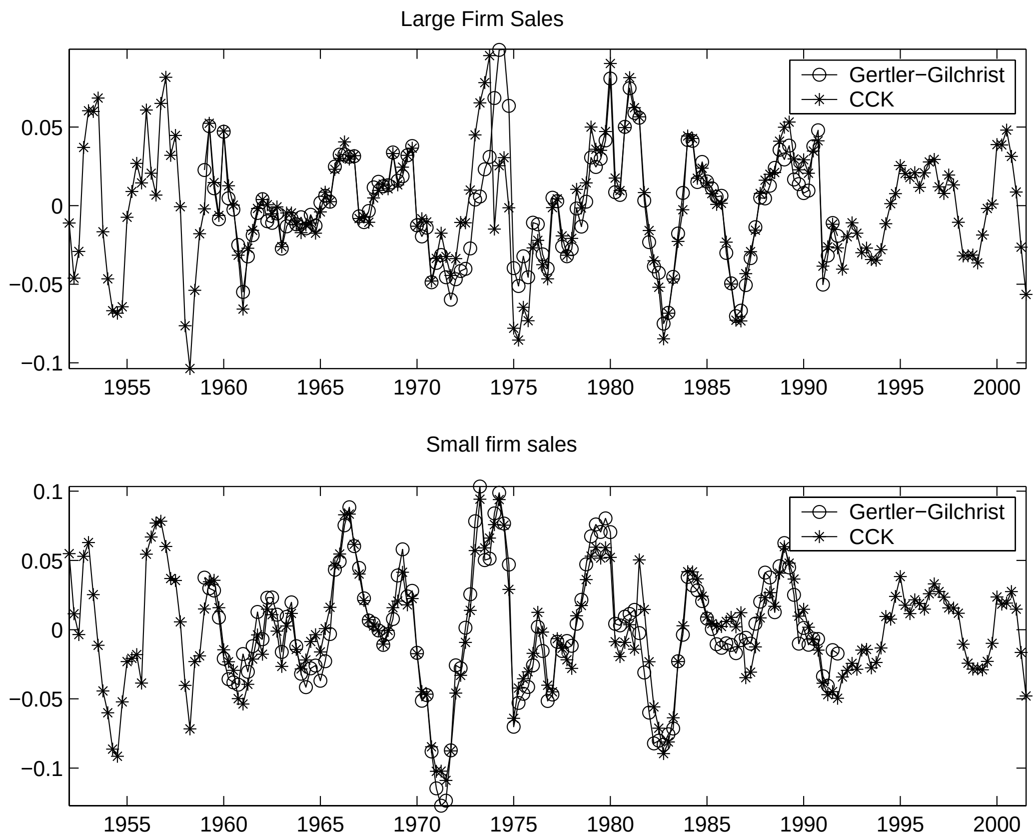

A few differences are notable. In the case of large firm sales in the early 1970s our measure and the Gertler-Gilchrist's measure are slightly out of phase.

Figure 8 displays the cumulative levels of the data:

Figure 8

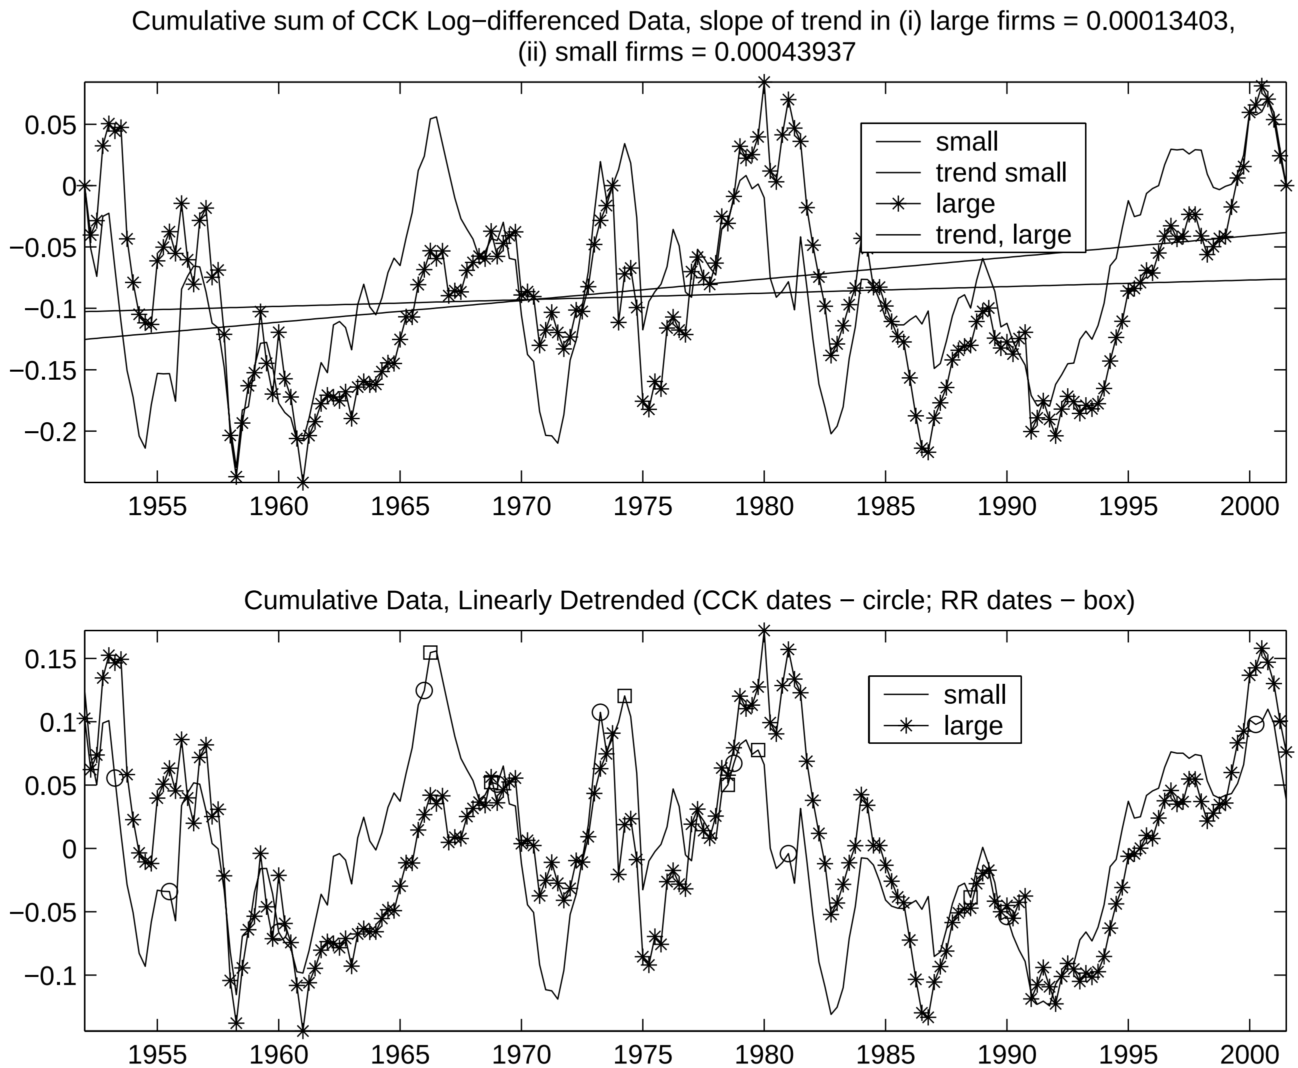

In contrast to what we found with Gertler and Gilchrist's data, there is virtually no trend in either of our level series. This represents a particularly sharp change for the sales of large establishments, since in Gertler and Gilchrist's data these display strong growth. That our level sales data don't display a trend is not surprising, in view of our relatively low estimate of the growth rate of sales in large establishments (see Figure 7).

Figure 9 displays the worm chart for Romer and Romer episodes, i.e., the object most comparable to Gertler and Gilchrist's analysis:

Figure 9

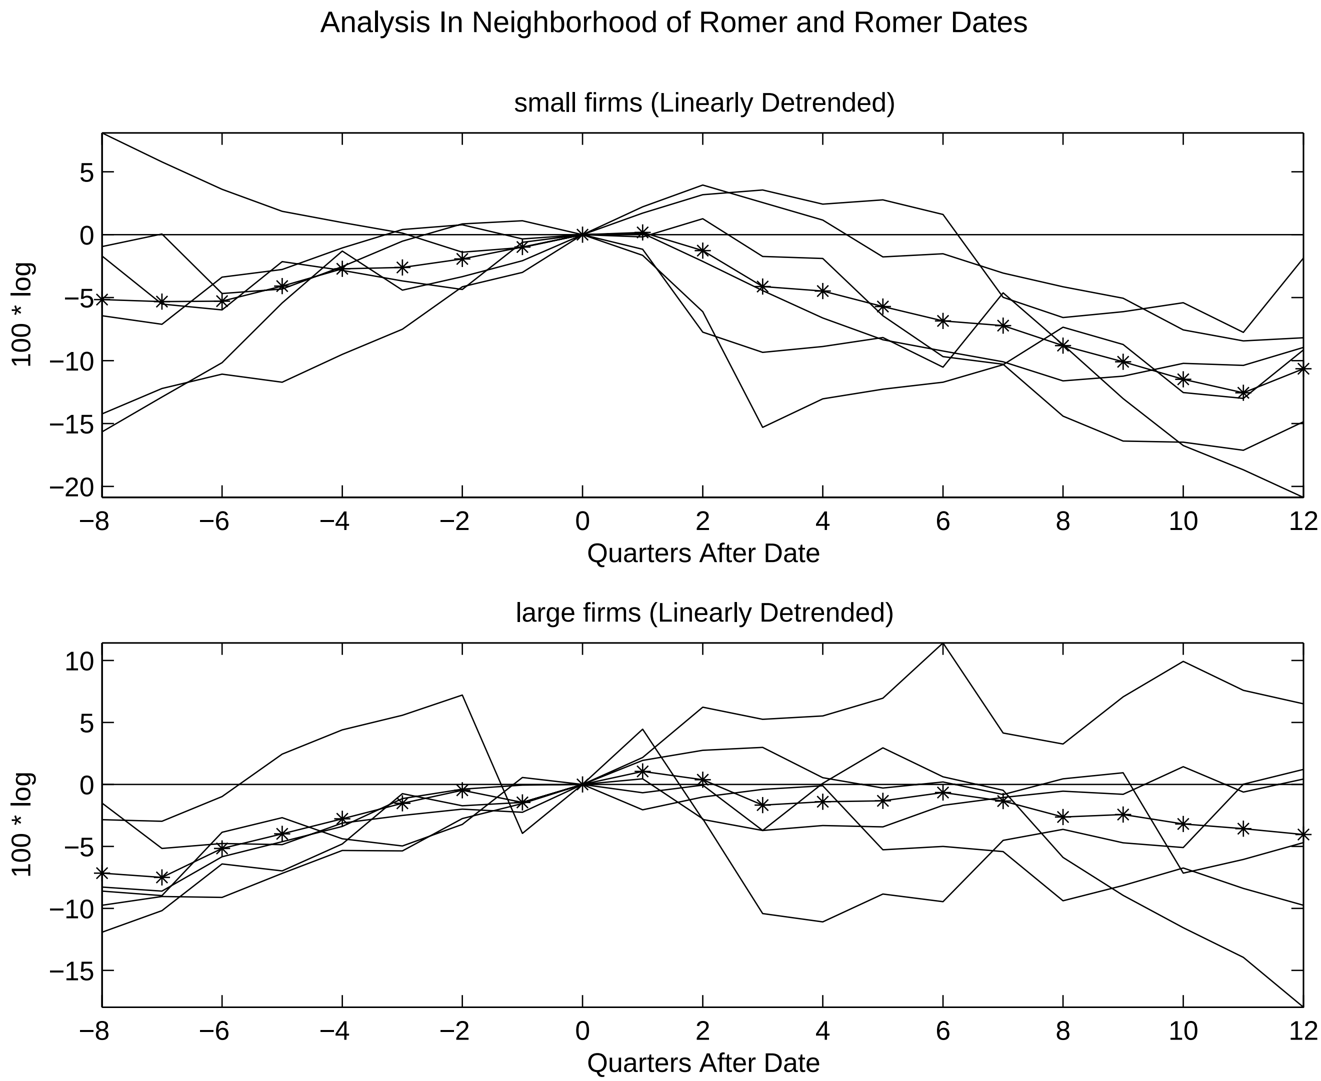

Note that the results are essentially unchanged relative to GG's analysis (see Figure 5). There is a clear tendency for small firms sales to decline by more than large firm sales. This finding suggests that the differences between our constructed data and Gertler and Gilchrist's are not large from the point of view of the questions that interest us.

Figure 10 considers the CCK business cycle dates with linear detrending:

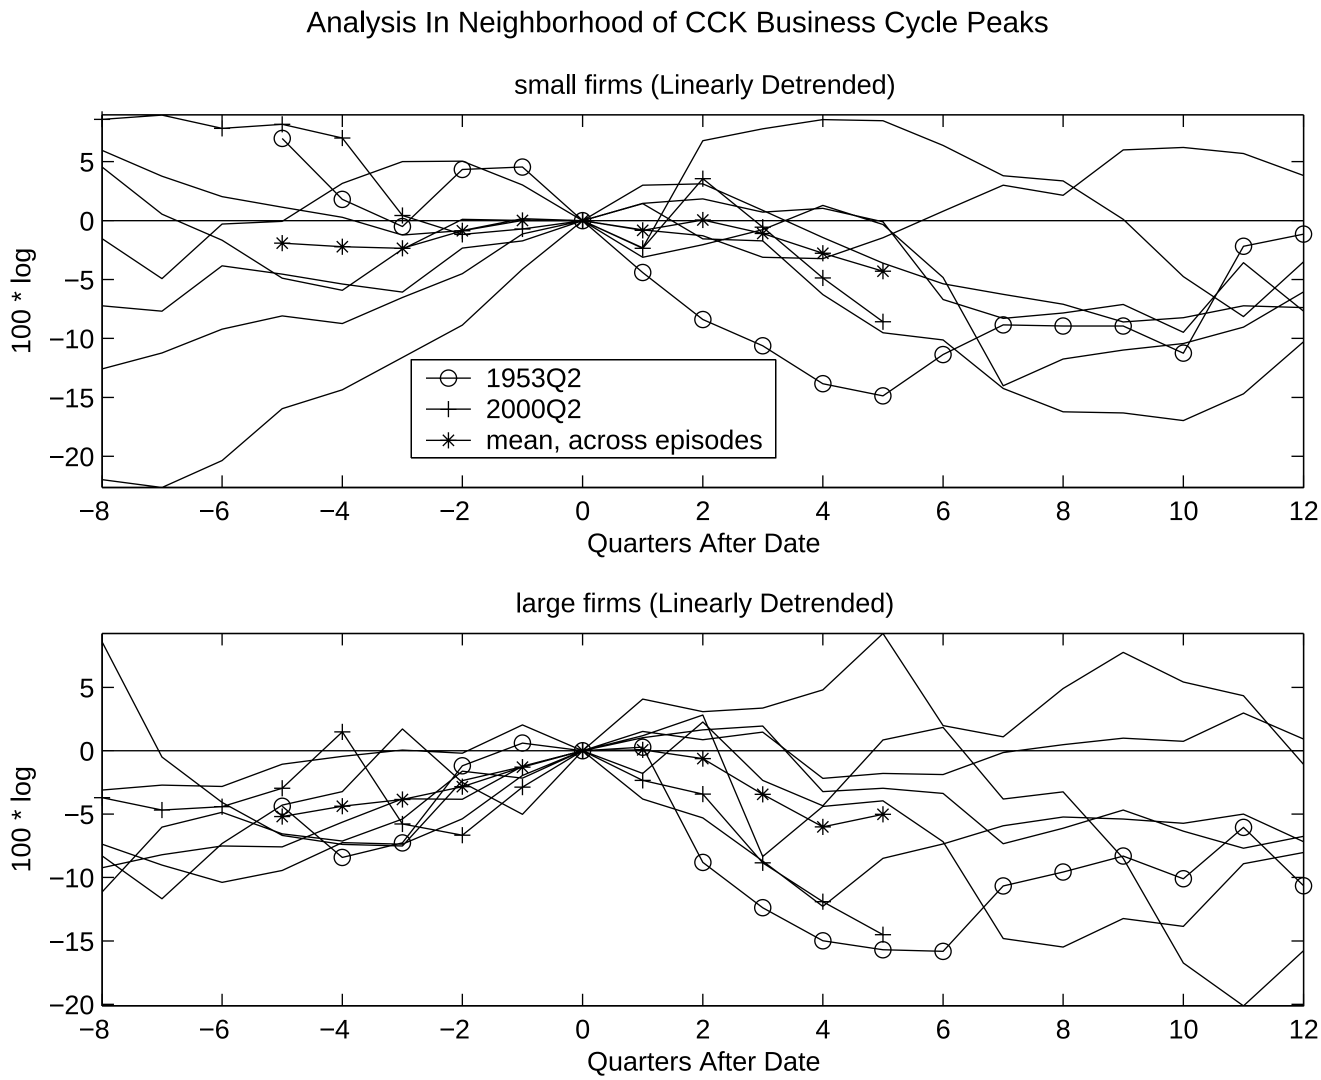

These results are considerably less supportive of the Gertler-Gilchrist hypothesis that the sales growth of small establishments falls more than the sales growth of large establishments in a recession (with HP filtering, the results are not noticeably different). Note that the distribution of responses for large establishments is shifted down relative to that of small establishments. Because the first and last CCK business cycle peaks are too close to the boundaries of the dataset, the mean responses could not be performed over the whole range (i.e., the starred line does not extend all the way to the beginning and end). However, over the few responses that we could compute, large and small firm sales differ by very little. When we drop the first (1953Q2) and last peaks (2000Q2) from the analysis then mean responses can be computed over the whole range of responses, and are displayed in Figure 11:

Figure 11

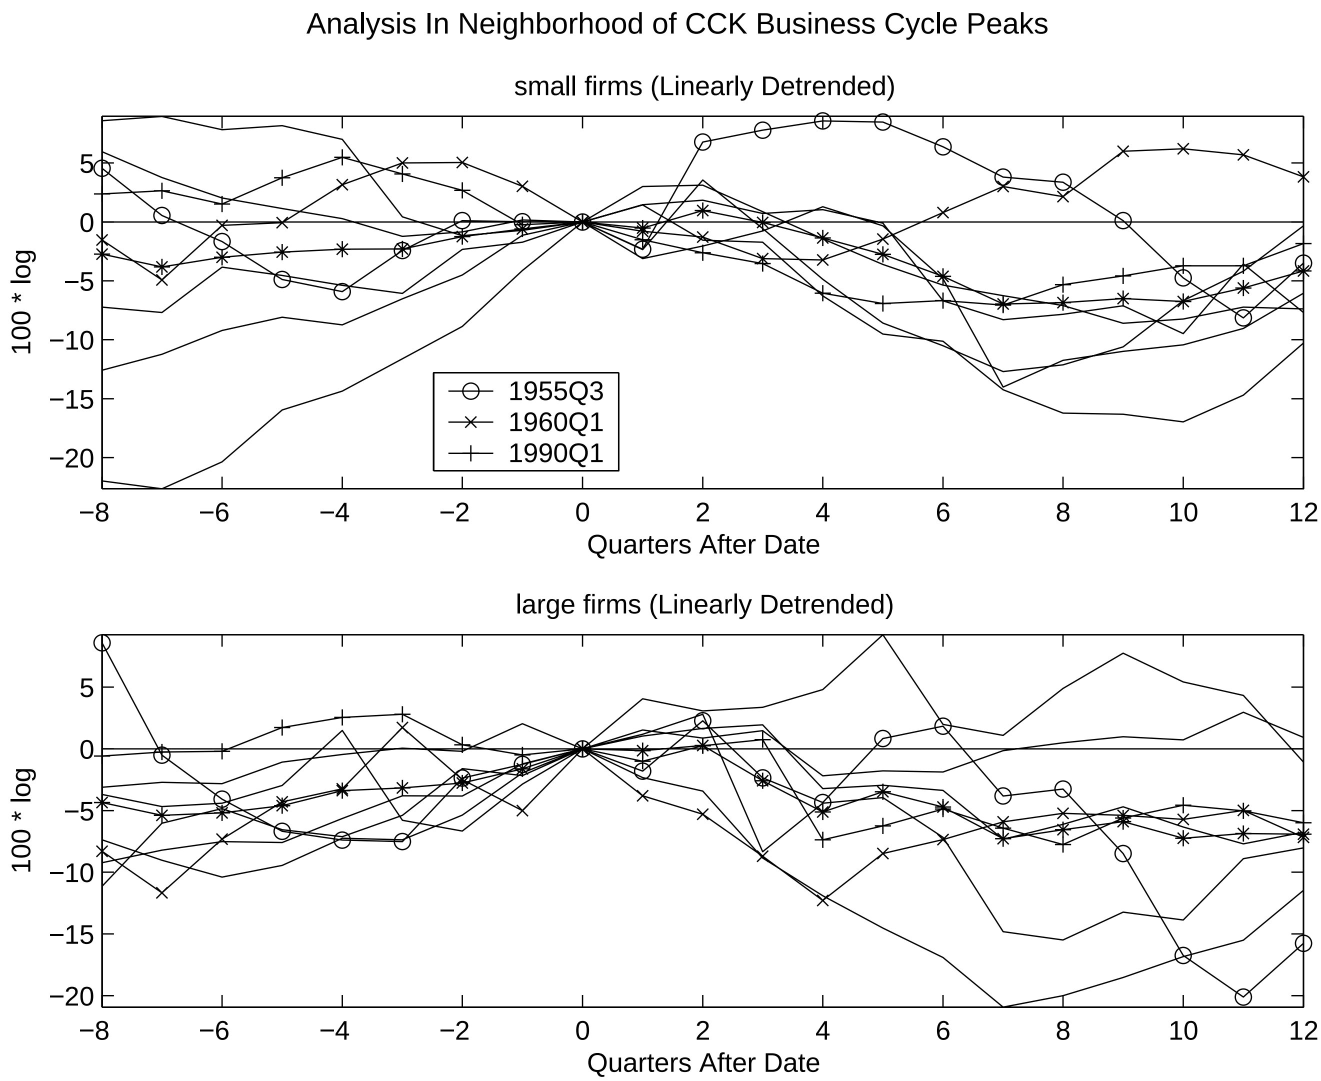

Two things are worth noting in Figure 11. First, as in Figure 10, the responses display more variation across episodes than the ones reported by Gertler and Gilchrist (see Figure 5). Second, the mean drop in the sales growth of large establishments is greater than the mean response of small establishments (see the starred lines).

These results may at first seem puzzling. When we computed worm charts for CCK business cycle peaks using the Gerlter and Gilchrist data, we obtained their basic result that the growth of sales in large establishments falls less during a recession than the growth of sales in small establishments (see Figure 6). What accounts for the difference between the results in Figures 6 and 11? In principle, it could be something about the construction of the sales data, or about the business cycle episodes studied. We argue that the latter is the right answer.

The three business cycle episodes included in Figure 11 which are not fully used in Figure 6 are the ones associated with peaks in 1955Q3, 1960Q1, and 1990Q1. These episodes are indicated in Figure 11. Consistent with what we saw in Figure 8, all these episodes are ones in which small firms' sales do not drop more than large firm sales. Indeed, in the first of these episodes, small firm sales actually surge as the economy enters the recession. Consistent with the view that these episodes have an important impact on the analysis, we recover the Gertler-Gilchrist findings when we drop them:

Figure 12

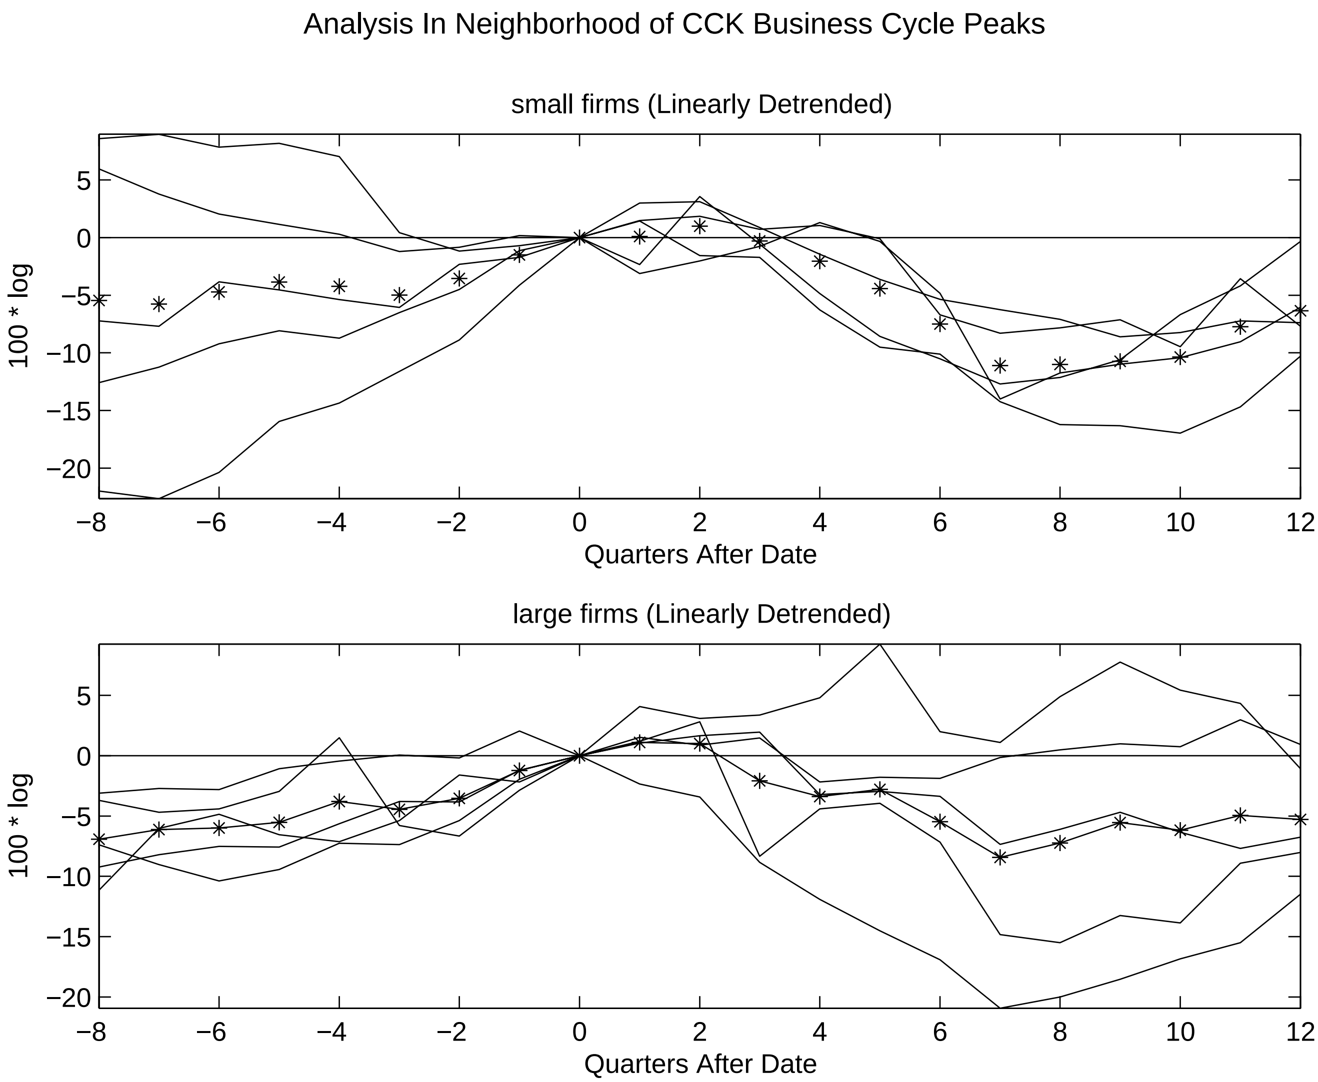

With the 1955Q3, 1960Q1, 1990Q1 episodes excluded from the analysis, the cross-business-cycle-episode mean of small firm sales falls by as much as 10 percent by 7 quarters after the business cycle peak while the mean of large firm sales falls by at most around 6 percent. Generally, the impression from Figure 12 is that small establishment sales falls by more than large establishment sales. Figures 11 and 12 indicate that when all business cycles are introduced in the analysis, there is not a significant difference in the behavior of large and small firm sales in recessions.

Figure 13 shows what happens after NBER business cycle peaks:

Figure 13

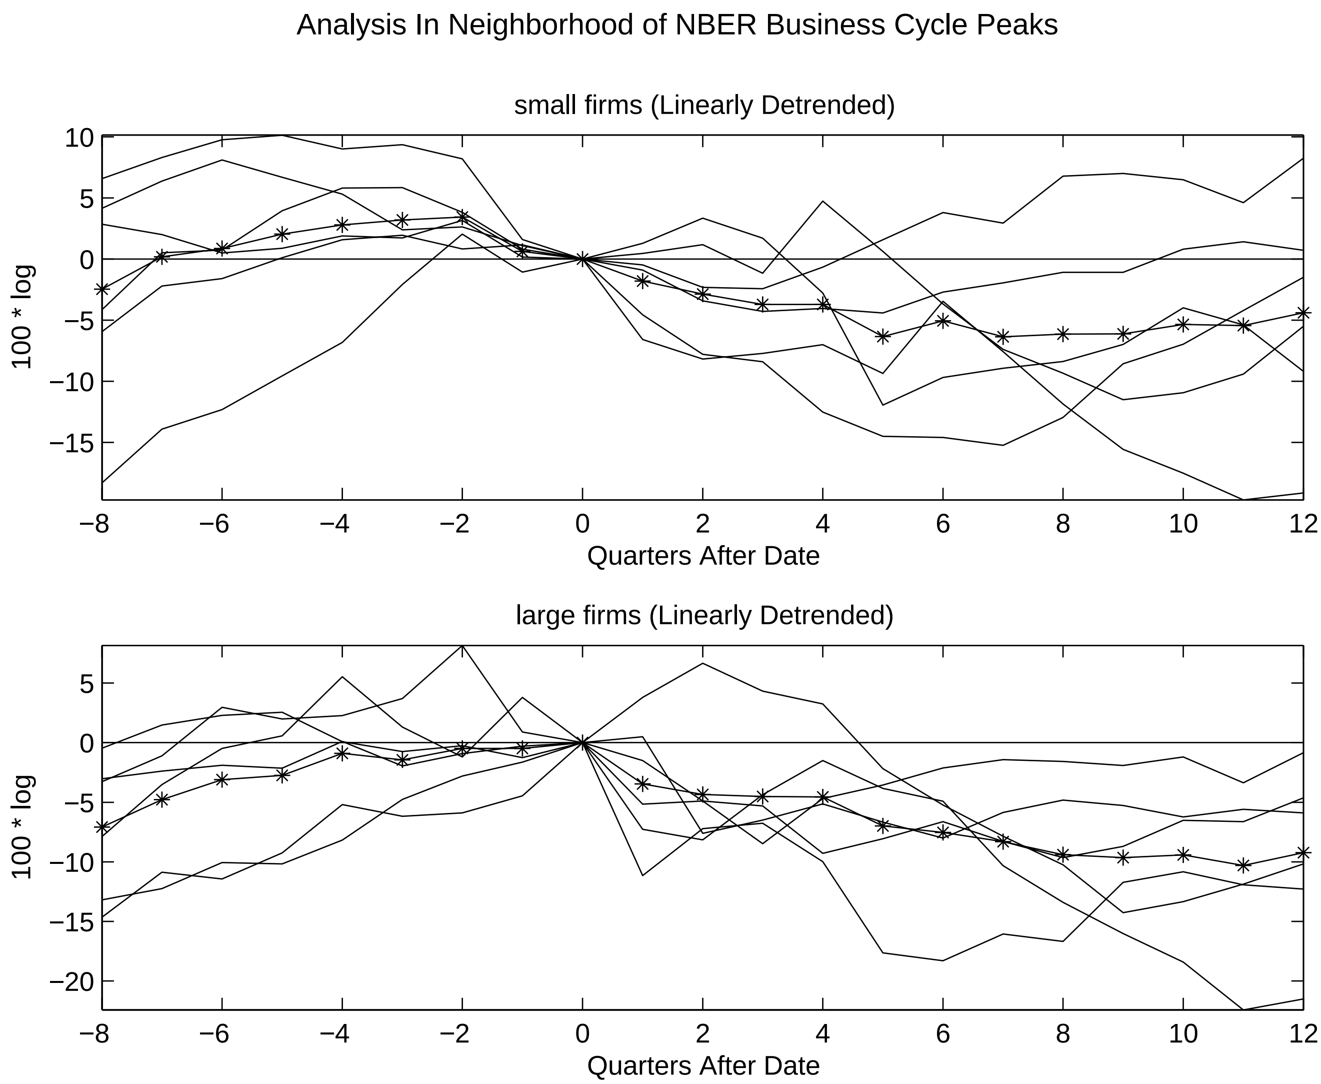

These results are also consistent with the proposition that there is no particular difference in the response of the sales of small establishments in recessions.