

How Important is the Intensive Margin of Labor Adjustment?*

Thijs van Rens[†]

CREI and Universitat Pompeu Fabra

October 2011

Abstract

Using new quarterly data for hours worked in OECD countries, Ohanian and Raffo (2011) argue that in many OECD countries, particularly in Europe, hours per worker are quantitatively important as an intensive margin of labor adjustment, possibly because labor market frictions are higher than in the US. I argue that this conclusion is not supported by the data. Using the same data on hours worked, I find evidence that labor market frictions are higher in Europe than in the US, like Ohanian and Raffo, but also that these frictions seem to affect the intensive margin at least as much as the extensive margin of labor adjustment.

*Discussion of “Aggregate Hours Worked in OECD Countries: New Measurement and Implications for Business Cycles” by Lee Ohanian and Andrea Raffo, prepared for the April 2011 Carnegie-Rochester Conference on Public Policy.

[†]I gratefully acknowledge the hospitality of the Department of Economics at Northwestern University and financial support from the Spanish Ministry of Science and Innovation (Juán de la Cierva fellowship, grant ECO2008-01665 and Consolider-Ingenio grant CSD2006-00016), the Ministry of Education (José Castillejo fellowship), the Generalitat de Catalunya (DURSI grant 2009SGR1157), and the Barcelona Graduate School of Economics.

1 Introduction

Ohanian and Raffo (2011) present new quarterly data series for total hours worked in 14 OECD countries over the past 50 years. Employment data were available for these countries, and data on total hours were available for the US, but since labor markets in at least some of these countries are quite different, there is no reason to believe that total hours and employment would comove in the same way as in the US. Therefore, these new data have the potential to substantially contribute to our understanding of business cycles across countries.

Using the new data, the paper documents a large number of stylized facts regarding the business cycle properties of total hours and its components employment and hours per worker. In addition to a standard set of business cycle statistics about the volatility of labor and its comovement with output and labor productivity, the results include business cycle diagnostics, using a similar framework as in Hall (1997), Cole and Ohanian (2004) and Chari, Kehoe, and McGratten (2007). This exercise focuses on whether recessions are primarily periods of low productivity wedges or high labor wedges. The paper does not try to distill its findings into one or two main conclusions, instead leaving it up to the reader to decide what are the patterns that emerge from all these facts.

In this discussion, I focus on what I consider to be the most important result in the paper: that the intensive margin (movements in hours per worker) is an important part of labor adjustment in many countries. This result is surprising given the conventional wisdom that the majority of labor adjustment happens along the extensive margin (movements in employment), which is based on the simple observation that, in the US data, fluctuations in employment explain at least two thirds of the business cycle volatility in total hours, see e.g. King and Rebelo (1999). Ohanian and Raffo argue that this result is not true in other countries and conclude that “employment is a poor proxy for labor input in many OECD countries, as changes in hours per workers are about as large as changes in employment” (p.3). I will argue that the new data on total hours worked do not support this conclusion.¹

There are good reasons to believe that the intensive margin of labor adjustment may be more important in other OECD countries than in the US. Labor market frictions tend to be higher in Europe, as the OECD Employment Protection Legislation (EPL) index referred to in the paper makes clear (Section 2.4). If these labor market frictions affect the extensive margin more than the intensive margin, for example because EPL takes the form of hiring and firing costs, then we would expect that in Europe a larger fraction

¹There are many other results in the paper, mostly about the robustness of findings in previous studies, which are interesting in their own right but which I do not discuss here. For example, Lee and Andrea show that the findings of Ohanian (2010) that there were remarkable differences between the US and other advanced economies in the Great Recession are largely unchanged when using total hours instead of employment. They also find that the decline in the procyclicality of labor productivity around the Great Moderation in the US, as documented in Gali and Gambetti (2009) and Gali and van Rens (2010), was present in other OECD countries as well.

of labor adjustment takes place along the intensive margin. However, whether this is true is of course ultimately an empirical question.

Many of the facts the paper documents are related to each other. For example, labor productivity is just output minus labor, so the information in the second moments of labor and productivity is the same information as the information in the second moments of labor and output. More importantly, comparing wedges from the diagnostic framework calculated from total hours versus employment does not use much of the information in the data and boils down to documenting the fraction of labor adjustment that occurs along the extensive margin, duplicating the same information in the data in different ways. Thus, much of the overwhelming richness of facts in the paper reduces to one and the same observation. On the other hand, there is additional information in these new data that the paper does not exploit. If we are interested in frictions on labor adjustment, presumably there is as much to learn from the persistence as from the volatility in the data, and we would want to explore the dynamics of labor adjustment.

I attempt to address both these issues in this discussion. I start by summarizing the results in the paper that relate to the importance of the intensive margin of labor adjustment in Section 2. In Section 3, I formalize the argument that many of the results from the business cycle diagnostics reflect (non-linear transformations of) the same information in the data. Then, in Section 4, I document this information in a non-theoretical way that is better comparable to the unconditional second moments and try to extract additional information about labor market frictions from the data by exploring the dynamics of labor adjustment. I find that, although the data are consistent with labor market frictions being higher in Europe than in the US, these frictions seem to affect the intensive margin at least as much as the extensive margin of labor adjustment.

2 Summary of the findings in Ohanian and Raffo (2011)

Three sets of facts in Ohanian and Raffo (2011) pertain to the importance of hours per worker as an intensive margin of labor adjustment across countries. Ohanian and Raffo conclude from these facts that the intensive margin of labor adjustment is more important in other countries, in particular in Europe, than in the US, which is consistent with labor market frictions along the extensive margin being higher in Europe than in the US. However, in all three cases, the evidence is mixed.

First, Ohanian and Raffo (Section 3) document that the volatility of total hours is much larger than volatility of employment. As a result, hours-based labor productivity (output per hour) is much more procyclical than employment-based labor productivity (output per worker). This finding is consistent with frictions along the extensive margin, inducing firms to use the intensive margin of labor adjustment. However, although the volatility of total hours varies substantially across countries, the ratio of the volatility

employment over total hours is similar for Europe and the US, inconsistent with the idea that labor market frictions are higher in Europe.²

Second, the cyclical pattern of labor wedges is quite different in Europe and the US, whereas productivity wedges behave similarly in both continents (Section 4.2). In Europe, labor wedges are less strongly correlated with output and fluctuations in labor wedges lag fluctuations in output by more than in the US. However, when we compare hours-based and employment-based labor wedges, we find that the ratio of the employment and hours-based wedges is remarkably similar across countries, in particular in Europe and the US, casting doubt on the interpretation that these patterns are driven by employment protection being much higher in Europe.

Third, in the US, the Great Recession was associated with a historically large labor wedge, whereas the productivity wedge barely moved. In other countries, in particular in Europe, the exact opposite was true (Section 4.3). These results, which are obtained from a peak-to-trough analysis for the Great Recession, are basically unaffected using labor wedges constructed from total hours instead of employment.

3 Business cycle diagnostics of labor adjustment

For each set of facts from the business cycle diagnostics exercise described above, there are two types of results: comparisons of the cyclicity of wedges between the US and Europe, and comparisons between wedges constructed using employment data and wedges constructed using data on total hours worked. While suggestive, a direct comparison of wedges needs to be interpreted with caution. There are many differences between the US and Europe. And since the wedges are calculated using data on output, capital and consumption as well as labor input, differences that are not related to the labor market affect the wedges. For example, the finding that labor wedges increase more in recessions in Europe than in the US may indicate that in Europe labor falls more in recessions, or that output falls less. The latter might occur for instance because product market frictions (price rigidities) are larger in Europe than in the US.

Comparisons of the cyclical behavior of productivity and labor wedges constructed using employment data to wedges constructed using data on total hours, in the US versus Europe, are much cleaner than direct comparisons of wedges between the two continents. Like a difference-in-difference estimator, this comparison ‘controls’ for any differences between Europe and the US that are not related to the labor market, in particular to

²This picture changes a bit if we look at hours per worker instead of employment. Hours per worker are substantially more volatile in Europe than in the US. The reason for this discrepancy is that standard deviations of employment and hours per worker do not add up to the standard deviation of total hours. This is not only because of the HP filter, as the paper notes in footnote 11, but also because the relative importance of the intensive and extensive margins is not constant over time, resulting in a correlation of less than one between employment and hours per worker. For this reason, in Section 4, I report peak-to-trough changes rather than standard deviations.

differences in the relative importance of the intensive and extensive margins of labor adjustment. However, the flip-side of this observation is that the information contained in the productivity wedges is the same information that is contained in the labor wedges, and that the same information can be seen directly from the data on employment and total hours, without using data on output or other variables, and without having to use a specific model structure or assume values for the structural parameters. To formalize this point, consider the expressions for the productivity and labor wedges as derived in the paper.

The productivity wedge aims to measure the quality of the production technology, and equals the Solow residual, see equation (5) in the paper,

$$Z_t = \frac{Y_t}{K_t^\theta L_t^{1-\theta}} \quad (1)$$

where Y_t is output, K_t is capital and L_t is a measure of labor input, either total hours worked L_t^H or employment L_t^N . The parameter θ is the capital share in production and needs to be calibrated. Now suppose we construct this productivity wedge using data on total hours Z_t^H and using employment data Z_t^N , take logs and evaluate the difference between both wedges.

$$\log Z_t^H - \log Z_t^N = -(1 - \theta) (\log L_t^H - \log L_t^N) \quad (2)$$

The difference in the productivity wedges is proportional to the difference in total hours worked versus employment, i.e. to hours per worker.

The labor wedge is meant as a measure of labor market frictions, and measures the ratio of the marginal rate of substitution between leisure and consumption over the marginal product of labor, see equation (6) in the paper,

$$X_t = \frac{B}{1 - \theta} \frac{C_t}{Y_t} L_t^{\frac{1+\varepsilon}{\varepsilon}} \quad (3)$$

where C_t is consumption, B is a parameter of the utility function determining the (marginal) utility from leisure relative to the utility from consumption and ε is the Frisch elasticity of labor supply. Taking logs and evaluating the difference between the labor wedge constructed using total hours worked X_t^H and using employment X_t^N , the difference between the two wedges is again proportional to the difference between total hours and employment.

$$\log X_t^H - \log X_t^N = \frac{1 + \varepsilon}{\varepsilon} (\log L_t^H - \log L_t^N) \quad (4)$$

The comparison of the employment and hours-based labor wedges contains exactly the same information as comparing the employment and hours-based productivity wedges,

or directly comparing the movements in employment and hours.³ The fact that productivity and labor wedges seem to give a different picture is due entirely to the fact that comparing productivity wedges dampens the difference in the adjustment in employment versus hours ($1 - \theta$ is smaller than one) whereas comparing labor wedges amplifies it ($\frac{1+\varepsilon}{\varepsilon}$ is larger than one).

4 The (non)importance of the intensive margin

Since they contain the same information as the raw data on employment and hours, there is little reason to report business cycle statistics for productivity and labor wedges. In this section, I summarize the evidence for the importance of hours per worker as an intensive margin of labor adjustment using only the raw data on employment and total hours worked. Instead of standard deviations, as in the paper, I report peak-to-trough changes. I show that this analysis gives a very similar picture as the standard deviations in Table 4, but peak-to-trough changes have the advantage that changes in employment and hours per worker add up to changes in total hours. After summarizing the evidence in the paper in this way, I explore the dynamics of labor adjustment by describing graphically the responses of total hours and its margins of adjustment in an average recession.

The first two columns in Table I summarize the results of the peak-to-trough analysis of movements in employment and total hours worked in the US and Europe. As in the paper, the numbers for Europe are the average of the three largest continental European countries: Germany, France and Italy. The table reports the drop in total hours and employment. The most intuitive way to summarize this information is as the fraction of the adjustment in total hours that is due to employment, $\Delta \log L_t^N / \Delta \log L_t^H$. This fraction is reported in the third column.

The fraction of adjustment that occurred along the extensive margin was larger in the US than in Europe for the average recession over the 1960-1984 period and in the Great Recession, consistent with the hypothesis that labor market frictions are higher in Europe. However, the picture is reversed for the average recession over the 1985-2007 period, when the extensive margin in Europe looked more flexible than in the US.

The changes in the US-Europe differences over time are driven by large changes in the labor adjustment mechanism in Europe. The importance of the extensive versus the intensive margin is relatively stable in the US. Part of the reason for this finding

³In order to convince myself that this is true, I took the peak-to-trough changes in employment and hours-based productivity wedges and used equation (2) to back out the implied difference in the peak-to-trough fall in employment versus hours, assuming $\theta = 0.36$ as in the paper. I also took the peak-to-trough changes in employment and hours-based labor wedges and used equation (4) to calculate the approximate implied difference in the fall in employment versus hours, assuming $\varepsilon = 1$ as in the paper. Then, I compared these numbers to the difference in the peak-to-trough changes in hours and employment. The implied differences from the productivity and labor wedges match up exactly with the data (up to rounding error).

may be sampling error. Because the paper does not report standard errors, it is hard to judge how seriously we should take the individual numbers. However, to get an idea of the amount of uncertainty, the table reports separate numbers for the three individual countries that go into the European average. The differences between European countries are substantially larger than the difference between the European average and the US.

How do the results from the peak-to-trough analysis line up with the unconditional second moments for employment and hours? If the fraction of labor adjustment that happens along the extensive margin is constant across fluctuations, we would expect the standard deviation of log employment relative to total hours to equal the ratio of the drop in employment relative to total hours from peak to trough in a recession.⁴ Columns 3 through 4 in the table report the standard deviations as in Table 4 in the paper, and their ratio. Broadly speaking, the numbers are similar as those from the peak-to-trough analysis. The contribution of the extensive margin to labor adjustment in Europe is smaller than in the US over the 1960-1984 period, but larger over the 1984-2007 period.

The fact that the contribution of the extensive margin to labor adjustment is similar in a peak-to-trough analysis and in the unconditional second moments suggests that the importance of the extensive and intensive margins is relatively constant and the dynamics of adjustment along each margin are similar. I now explore these dynamics explicitly. This is interesting because the shape of the responses, in particular their persistence, is informative about labor market frictions. If it is true that labor market frictions are higher in Europe than in the US, as Ohanian and Raffo argue, then we would expect that labor adjustment is more persistent in Europe than in the US. If it is further true that labor market frictions work mostly on the extensive margin, then we would expect that in Europe the intensive margin is more important, not only in terms of the overall adjustment, but particularly in the short run.

To evaluate these predictions, Figures I and II graph the adjustment of total hours and its components employment and hours per worker in an average recession in the US and the three largest continental European economies. These graphs are constructed by regressing the first difference of log hours (or log employment or log hours per worker) on 20 lags of a dummy variable indicating an NBER or ECRI peak date. The coefficient on the q^{th} lag of the dummy measures the change in hours q periods after the start of the recession, averaged over all recessions in the sample period. I then calculate the cumulative change since the start of the recession and plot it against q . Since I find no significant differences in these responses across subsamples, I do this for the entire sample period excluding the Great Recession, 1960-2007.

⁴If the fraction of adjustment along the extensive margin, $\Delta \log N / \Delta \log H$, is constant, then it must be that $\log N = \lambda \log H$, where $\lambda = \Delta \log N / \Delta \log H$ is the contribution of the extensive and $1 - \lambda$ the contribution of the intensive margin to the adjustment of total hours worked. Then, $sd(\log N) / sd(\log H) = \lambda$.

Figure I plots the response of total hours worked in an average recession with one standard error bands. The standard error bands measure the uncertainty about the average response. The fact that the standard errors are relatively large reflects the fact that there are large differences between recessions in the same country. The main difference between the US and the European countries is that the fall in hours is much less persistent in the US, with the recovery setting in after 5 quarters on average. There are large and significant differences between the European countries in the depth of the average recession.

Figure II decomposes the response of total hours into adjustments along the extensive and the intensive margin. In all four countries considered, employment closely tracks the cyclical behavior of total hours worked, while the intensive margin of labor adjustment plays a relatively minor role. In particular, there is no evidence that the intensive margin is more important in the short run versus the long run, or in Europe versus the US.

Figure III explores whether there might be an asymmetry in labor adjustment in recoveries versus recessions. In this figure, I plot the average response of hours worked and its components q quarters after a trough date. The intensive margin seems to be slightly more important in recoveries than in recessions, at least in the US and Germany. However, the extensive margin is still responsible for most of the adjustment in total hours.

The finding that the reduction in total hours worked is (much) more persistent in Europe than in the US is consistent with the hypothesis that labor market frictions are larger in Europe. The finding that the intensive margin is relatively unimportant both in Europe and in the US and regardless of the horizon suggests that adjusting the intensive margin is costly as well. It seems the data would indicate that frictions on both the extensive and the intensive margins of labor adjustment are higher in Europe.

5 Conclusions

Ohanian and Raffo (2011) make a valuable contribution to the literature on business cycle fluctuations on the labor market, by providing new quarterly data on total hours worked in 14 OECD countries. The construction of these new data is carefully done. There is a lot to be learned from these data, and it seems likely that they will be widely used.

In terms of substantive results, I focused on the importance of the intensive margin of labor adjustment in Europe versus the US. I tried to organize the findings in the paper concerning this issue, to relate these findings to each other, and to evaluate whether the data support the conclusion that higher labor market frictions make the intensive margin of labor adjustment more important in Europe. I find no evidence for this conclusion. The data are consistent with labor market frictions being higher in Europe than in the

US, but these frictions are at least as high along the intensive as along the extensive margin.

References

- Chari, V. V., P. J. Kehoe, and E. R. McGratten (2007). Business Cycle Accounting. *Econometrica* 75(3), 781–836.
- Cole, H. and L. E. Ohanian (2004). The US and UK Great Depressions Through the Lens of Neoclassical Growth Theory. *American Economic Review* 92(2), 28–32.
- Gali, J. and L. Gambetti (2009). On the Sources of the Great Moderation. *American Economic Journal: Macroeconomics* 1(1), 26–57.
- Gali, J. and T. van Rens (2010). The Vanishing Procyclicality of Labor Productivity. CREI working paper.
- Hall, R. E. (1997). Macroeconomic Fluctuations and the Allocation of Time. *Journal of Labor Economics* 15(1), S223–S250.
- King, R. G. and S. T. Rebelo (1999). Resuscitating Real Business Cycles. In I. J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics, volume*, pp. 927–1007. 1B.
- Ohanian, L. E. (2010). The Economic Crisis from a Neoclassical Perspective. *Journal of Economic Perspectives* 24(4), 45–66.
- Ohanian, L. E. and A. R. Raffo (2011). New Measures of Hours Worked in the OECD: Implications for Business Cycles, 1960-2010. *Journal of Monetary Economics, this issue*.

Table I

	<i>peak-to-trough</i>			<i>standard deviation</i>		
	Hours	Empl	ext.m.	Hours	Empl	ext.m.
1960:Q1-1984:Q4						
US	-3.6	-2.8	77%	0.85	0.72	85%
Europe	-4.7	-2.6	56%	0.76	0.50	66%
Germany	-7.6	-4.0	53%	0.72	0.57	79%
France	-3.2	-1.9	59%	0.84	0.49	59%
Italy	-3.2	-1.9	60%	0.71	0.43	61%
1985:Q1-2007:Q4						
US	-2.2	-1.6	74%	1.23	0.96	78%
Europe	-3.4	-3.1	92%	0.80	0.70	87%
Germany	-4.0	-3.3	82%	0.61	0.65	107%
France	-1.8	-1.3	74%	0.92	0.66	72%
Italy	-4.2	-4.6	109%	0.88	0.78	88%
Great Recession						
US	-7.0	-5.5	79%			
Europe	-3.5	-1.5	42%			
Germany	-2.7	0.6	-22%			
France	-1.4	-1.2	86%			
Italy	-6.4	-3.8	60%			

Peak-to-trough changes are log deviations. Changes for the 1985:Q1-2007:Q4 period do not include the Great Recession, even for countries (US, Italy), in which the recession started in 2007. Standard deviations are relative to the standard deviation of log real output and are calculated after HP filtering the data with a smoothing parameter of 1600.

Figure I
Labor adjustment in average NBER/ECRI recessions, 1960:Q1-2007:Q4
Response total hours with standard error bands (dash)

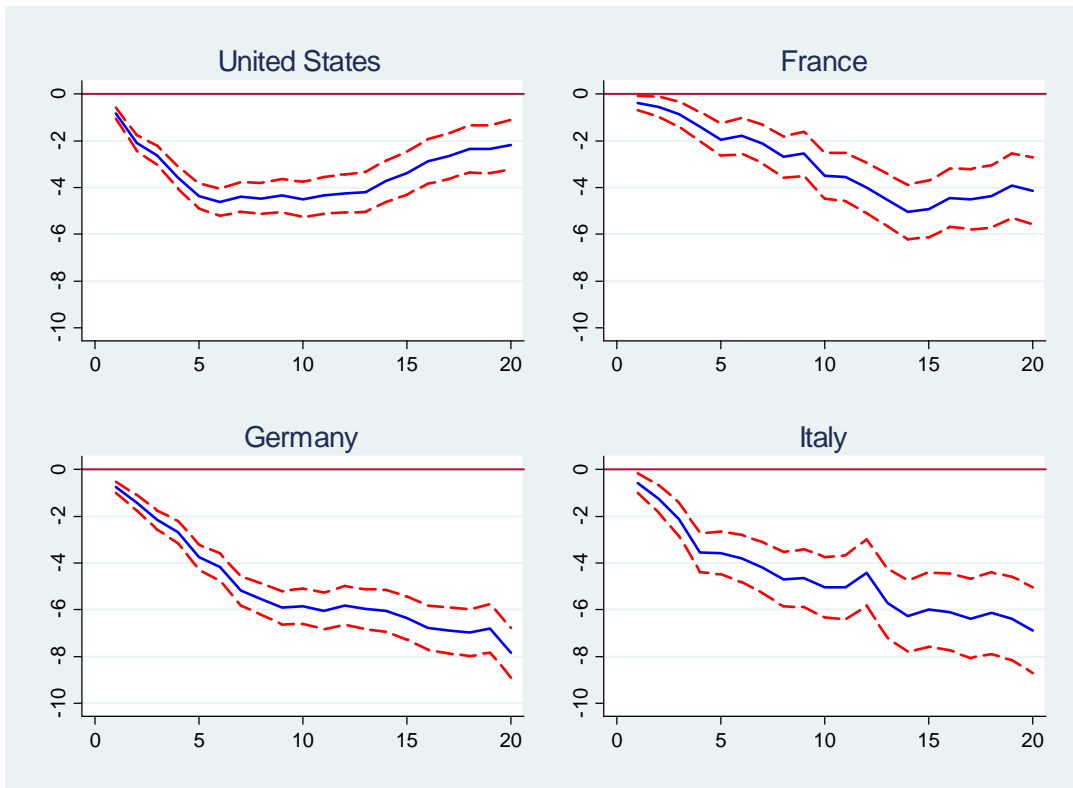


Figure II

Labor adjustment in average NBER/ECRI recessions, 1960:Q1-2007:Q4
Response total hours (solid), employment (dash) and hours per worker (dot)

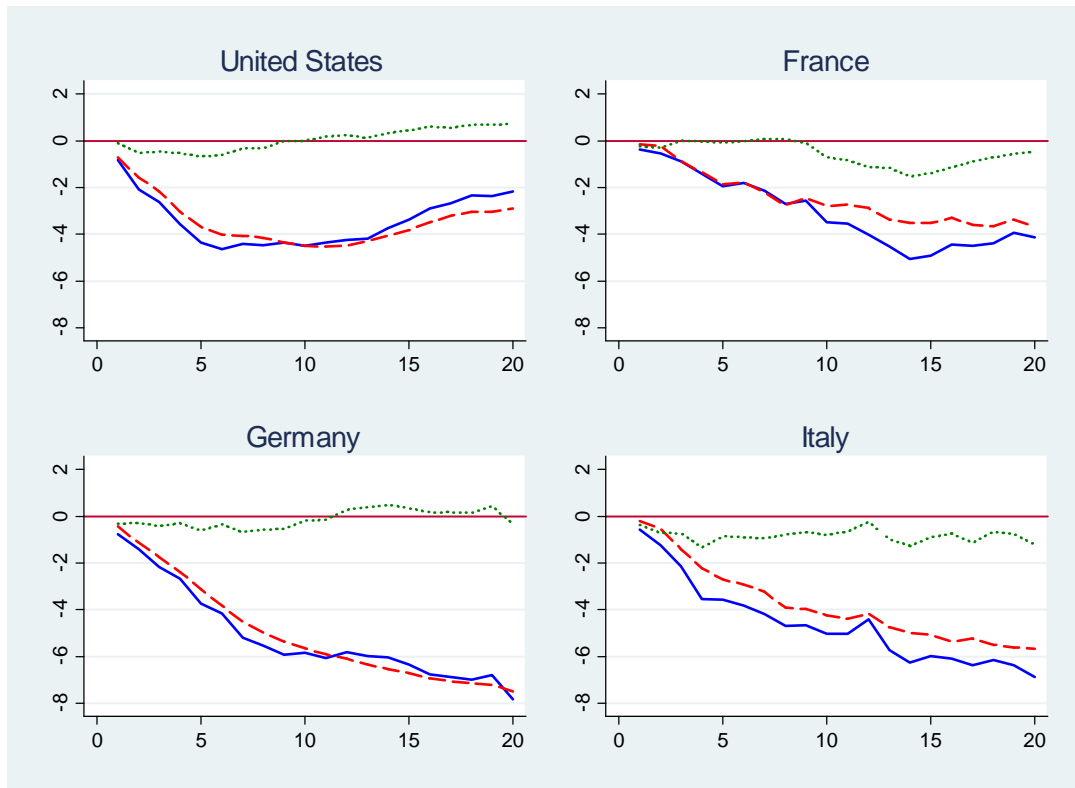


Figure III

Labor adjustment in average NBER/ECRI recovery, 1960:Q1-2007:Q4
Response total hours (solid), employment (dash) and hours per worker (dot)

