SFB 649 Discussion Paper 2006-042

Discussion of "The Source of Historical Economic Fluctuations: An Analysis using Long-Run Restrictions" by Neville Francis and Valerie A. Ramey

Harald Uhlig*

* School of Business Administration and Economics,
Humboldt-Universität zu Berlin;
Tilburg University; Bundesbank and CEPR

This research was supported by the Deutsche Forschungsgemeinschaft through the SFB 649 "Economic Risk".

http://sfb649.wiwi.hu-berlin.de
ISSN 1860-5664

SFB 649, Humboldt-Universität zu Berlin
Spandauer Straße 1, D-10178 Berlin

SFB 649 ECONOMIC RISK BERLIN

"The Source of Historical Economic Fluctuations: An Analysis using Long-Run Restrictions" by Neville Francis and Valerie A. Ramey

Discussion by Harald Uhlig
School of Business Administration and Economics
Humboldt Universität zu Berlin,
Tilburg University, Bundesbank and CEPR
uhlig@wiwi.hu-berlin.de

Date: May 8th, 2006, originally issued 2004

Footnote: I am grateful to Neville Francis and Valerie A. Ramey for sharing their data with me. This research was supported by the Deutsche Forschungsgemeinschaft through the SFB 649, "Economic Risk". Keywords: technological progress, business cycles, hours worked, labor productivity, vector autoregressions, identification, long-run restrictions, medium-run restrictions. JEL codes: E22, E23, E24, E32, C32

Abstract:

This paper discusses the paper "The Source of Historical Economic Fluctuations: An Analysis using Long-Run Restrictions" by Neville Francis and Valerie A. Ramey. It argues that these authors have made great progress both in the precise measurement of labor input as well as determining the effect of productivity shocks on labor, but a number of questions remain. As for measurement, the issue of schooling needs further work. As for calculating the long-run impact of labor productivity shocks, unreasonable results emerge for the response of the capital stock, if included in the VAR. Using medium-term identification delivers more reasonable results.

1. The issue

What is the consequence of technological progress? To the lay person, the answer often seems obvious: technological progress is both a virtue and a vice. It is a virtue, because it has made our lifes more comfortable. But it is also a vice because people are loosing jobs. Factories, where humans used to work together in the past to create the products to be sold, are now instead filled with machines and the occasional operator, pushing some buttons now and then, and with the other workers out of a job. The latter aspect is viewed as a negative aspect of technological progress and often dominates public debates.

The economist gives a dramatically different assessment. According to the data presented by Francis and Ramey, labor productivity in the private sector has increased by the factor 13.5 between 1889 and 2002. Francis and Ramey show, that hours worked in the private sector per capita really has not changed all that much, once one takes into account e.g. trends in schooling or government employment. Thus, using a standard Solow growth accounting exercise, i.e. postulating that capital and output are on the steady state growth path and labor is constant, and postulating a Cobb-Douglas production function with a capital share of one third, one finds that total factor productivity has increased by a factor of 5.7 between 1889 and 2002. I.e, even if the (quantity of the) capital stock per worker had not changed at all, a worker now would need only 10 minutes to produce what his great-great-grandfather would have taken an hour to do. And with the additional capital accumulated in the mean time, that number shrinks to less than five minutes. The welfare gains from this dramatic increase in productivity are obvious and self-evident. The virtuous aspects of technological progress dominate by far.

But what about the vice, that people might loose jobs? Take again the evidence produced by Francis and Ramey, that hours worked in the private sector have not changed much over the last century. Clearly then, over the long haul, technological progress has not led to machines replacing workers: rather, the desires of humans for consumption relative to enjoying leisure have risen with the technological progress (and therefore the opportunity costs of leisure) we have made. We need to work just as hard as more than 100 years ago to keep up. And surely, even if one does not buy into the evidence produced by Francis and Ramey and rather believes that hours have shown a secular decline, few economists would interpret this as evidence that technological progress is the culprit for the high unemployment rates observed in some of the modern welfare states. Instead, the secular technological progress is viewed as freeing up time for people to enjoy as leisure. This macroeconomic perspective may be complicated by redistributional issues - some workers may benefit more from "replacement" by technological progress than others, and some may even indeed suffer from it - but overall the conclusion remains: technological progress is a virtue.

So, why should we care what technological progress - or better, surprise changes in the rates of technological progress - do to the labor market in the short run? After all, this is the question which Francis and Ramey seek to answer, extending the research agenda of Blanchard and Quah (1989) and Galí (1999). That research is challenging the bold claim initially made by Kydland and Prescott (1982) and others, that random fluctuations in technological progress are the cause of business cycles. (As an aside, let it be pointed out that "technology" in the Arrow-Debreu sense simply refers to the possibility of turning inputs into output, so that the term "technology shock" simply refers to any changing production function. Instead, the recent literature as well as the Francis-Ramey paper has focussed on the common language interpretation of the word "technology", i.e. patents, engines, microchips. I shall follow along for the purpose of this discussion.). That research helped to explain the challenging observation that labor productivity is procyclical. Demand-driven theories need to be worked pretty hard to cough up this and other key business cycle facts.

Some of these business cycle facts are listed in table 1, using the Francis-Ramey data from 1889 to 2002. Note the positive correlations between output on the one hand and labor as well as labor productivity on the other, no matter how and when one measures it. However, also note the low correlations between labor and labor productivity: for 1950 to 2002, that correlation appears to be near zero. This correlation is low and lower than the numbers usually given in the literature, see e.g. Cooley (1995). This simple statistic already sheds a lot of light on the results by Francis and Ramey. Clearly, if there are three statistics - output, labor and productivity - which are not simultaneously highly correlated with each other, two rather than one source of randomness are needed to explain most of it. Since it is labor and productivity, which show low correlation, whatever explains the movements in productivity won't explain the movements in labor and vice versa (which is essentially also what Francis and Ramey find in their VAR estimates). So, this alone merits much deeper investigation: is it really true that labor productivity and hours worked show low correlations at business cycle frequencies? This is a bold claim, which is only implicit in this paper, and which could lead to a change in our thinking about business cycles, if it holds up to scrutiny.

In light of these tables, the key business cycle question is: what explains the high labor-output and productivity-output correlations, while generating low correlation between labor and labor productivity? If one buys into the findings by Galí and now by Francis-Ramey, business cycle theories driven by technology shocks cannot be it. Thus, perhaps rather than the welfare question as to whether technological progress is a good thing (it is, in practically all reasonable models), the issue at stake is: what explains business cycles.

There is a long literature criticizing the claim of the real business cycle paradigm already, though. Even if technology shocks were to lead to initial increases in labor, they may contribute little to the variance of output, see e.g. Christiano, Eichenbaum and Vigfusson (2004) or Altig et al (2002). Variance decompositions rather than impulse responses might be the most interesting object of investigation here: Francis and Ramey provide it in their interesting table 4. More importantly, it is all too easy to rule out explanations of business cycles - instead, we need good, convincing theories explaining them! I keep on being surprised how easily the model by Hansen (1985) - which nowadays should be considered as a strikingly simple model - explains many of the key business cycle facts. That model sets a standard, and every graduate student in economics should learn it well. Is there a similarly beautiful, alternative explanation in sight, which works even better? I do not think there is, and even if there were, it does not seem that the profession has decided to raise that one on its shield yet as the new key paradigm of business cycles. Perhaps, too much effort has gone into shooting down a model that works ok. More work should instead go into providing a model that works better.

The issue thus cannot be as to whether technology shocks explain business cycles or not. Instead, the issue is whether one can properly identify technology shocks, using long run restrictions, and what they imply about labor movement in the short run. Let me thus turn to the contribution of Francis and Ramey. It is two-fold. First, they carefully put together a long-run data set, which is an interesting object of investigation in its own right. Second, they use a variety of VAR specifications to analyze the question at hand, thereby addressing some of the issues raised in the recent debates.

2. Long-run data

The first contribution of Francis and Ramey is to put together a set of long time series, and to carefully account for the hours worked in the private sector, addressing the issue of changes in education and the role of the government. This is meant to achieve three goals. First, it should provide more data to study the question at hand: clearly, this is always a good idea. The alternative (which should be pursued!) is to investigate more countries rather than longer time series, but certainly, having longer time series cannot hurt.

The second goal is to provide a more balanced view of the changes in labor input, leading to some surprising results (and as for the third goal, see the next section). Francis and Ramey make the bold claim that hours per capita in the private sector have not changed very much over more than a century, see e.g. their figure A. The average for 1950 to 2002 is not even 7% below the average for 1889 to 1940. Their adjustment of the standard labor force accounting comes from three sources.

First, rather than just considering everyone above age 16, they implicitely account for child labor by subtracting only children in school from the population above age 4. But one could go further. Probably, one ought to e.g. also account for changes in the death rates in early childhood due to a variety of diseases. Figure A shows the decline in children death rates: unfortunately, I only had data for ages 1 to 4 rather than 5 to 16 at hand. Nonetheless, that figure shows a dramatic fall in death rates by a factor of 60: while 2% of all children between age 1 and 4 died in 1900, only 0.03% of them did so in 2000, see also the right-top panel of figure 3 in the Francis-Ramey paper. It is plausible that there also was a dramatic decline in death rates for children between the ages 5 and 16. Certainly, the demographics in 1889 was much more heavily tilted towards younger people and children than it was in 2000, so that one is dividing by a larger number in 1889 than in 2000. This is ok, if indeed all children above age 4 not in school were indeed working. But I tend to think that this is a somewhat extreme view.

That this could matter can be glanced from figures B through D. In figure B, the fraction of children between ages 5 and 16 compared to the total population is plotted: one can see the large swings between nearly 25% at the top and slightly above 15% at the bottom. The numbers are very high towards the beginning of the sample and during the "baby boom" years. The numbers become more dramatic, once one subtracts out the population above 65 and government employees in the denominator, see figure C. Now, the numbers swing between 21% and 30%. These swings would matter a lot, if all children between age 5 and 16 were accounted as part of the work force.

Thus, Francis and Ramey rightly subtract school enrolment. We do this too both in the numerator and denominator in figure D. Now, the numbers turn negative, though: from a peak near 3% to the lowest point at near -18%. One can already see this in figure 3 of the Francis-Ramey paper that this must be so: according to their numbers, nearly a quarter of the total population was enrolled in school during the last 20 years. Clearly then, this must include people who are not "too young to work". Francis and Ramey effectively make the extreme assumption, that everyone enrolled in school is not working. Given these large fractions of the population enrolled in school, one may have some serious doubts here. Finally, much of the apparent rise in hours worked in the recent 20 to 30 years is due to female labor force participation. To me, it is plausible, though, that a sizeable share of women have worked hard in the private sector previously as well, but outside official government statistics at the turn of the century. I can only hope that the original data used by Francis and Ramey has taken this properly into account, but I am skeptical.

In sum, one has to wonder whether the calculations by Francis and Ramey really are the final word on measuring the actual work force. Probably, they are not. Nonetheless, their findings are fascinating. According to conventional wisdom, hours worked per capita have declined over the last 100 years or so. Francis and Ramey have taken a first and very useful stab at that issue, and challenged that wisdom. This really calls out for an intensive analysis of the data to settle this issue.

3. VAR estimates

The third goal of providing a longer data set is to take more confidence in imposing long-run restrictions. But here, I believe, one should be careful. Classical econometric time series analysis makes one believe that there is a big difference between an exact unit root and a root smaller than one. But this difference concerns the hypothetical exercise of considering ever longer data sets, governed by some given stochastic law of motion. However, a different view is more sensible. Most of the time, the data is of given length, and one has to make inference about the various competing possibilities for the roots governing the process, see e.g. Sims-Uhlig (1991). Or - as is argued here - one is indeed provided with longer time series, but then one has a hard time believing that the stochastic properties do not change. Simple econometric models are plausibly understood to be parsimonious descriptions of key features of the data, rather than true, genuine data generating mechanisms which remain unchanged across several centuries. Given the difficulties described above in truly calculating the work force and thus in truly calculating labor productivity, the impression that more than a century of data really helps in separating the unit root parts from the non-unit-root parts may thus be just a misleading illusion. More data helps, yes. But with more data, one ought to consider a richer set of econometric possibilities. In particular, slow variations in the regression coefficients would invalidate the inference made here.

My interpretation of the evidence presented is therefore different. Evidence for a unit root in the data should be interpreted as evidence for some very persistent feature. Whether the effect gradually dies out after 100 years or whether it does not, is probably not of relevance. The "technology shock" identified in this paper is then simply that shock which leads to the most persistent changes in labor productivity among all the shocks one can consider. This is interesting.

I am concerned that it should make a difference as to whether hours are regarded as stationary or not. Since OLS provides consistent estimates of VAR coefficients, regardless of whether the variables are stationary or not, one can identify this most persistent shock certainly also in a VAR in levels, even if hours have a unit root. Conversely, if hours do not have a unit root, then a VAR which uses hours in first differences is misspecified. Christiano, Eichenbaum and Vigfusson (2004) have nicely demonstrated, how the level results encompass the results for first-differences, and that one therefore ought to trust the level results rather than the results for first-differences, regardless of what the unit root tests for labor show. I find their argument in favour of the level specification more plausible than the arguments given in the Francis-Ramey paper in favour of the first-difference specification. To put it succinctly: VARs should typically be estimated in levels, unless there are very good theoretical reasons, not to. And one should use Bayesian methods, which can deal with the uncertainty regarding the presence of unit roots in a very natural and practical manner, see Uhlig (1994). Certainly, detrending with a quartic trend strikes me as something which may potentially be rather misleading. Are we allowed to extrapolate that trend towards infinity? And what does it do to inference about cause and effects in VARs, if current data is "cleansed" from a trend, which in turn is estimated with the help of future data?

The level results look particularly damaging to the benchmark real business cycle perspective, though, see figure 5. Apparently, the permanent shocks to labor productivity - which Francis and Ramey call and identify as technology shocks - are only a minor cause of output fluctuations (see the bottom row in figure 5) and thus, it is no surprise that they also do not do much to labor and even lead to a short, initial decline, see the middle panel. Interestingly, the sample here matters a lot. Figure 6A shows, that labor shows practically no initial response to a permanent productivity shock for postwar data, while it shows a large negative response in prewar data. The first of these findings is one of the key points in Christiano, Eichenbaum and Vigfusson (2004): the level specification overturns the findings of Galí (1999) for postwar data. So, Francis and Ramey come to Galí's rescue by showing that the level specification makes things "worse" for the real business cycle school, once one takes that level specification to prewar data. This point has also been made by myself in Uhlig (2004), using a previous version of the Francis-Ramey data set.

There are three potential conclusions one can draw: at this point, a reader should feel free to choose any one of them. First, the econometric model does not remain stable over time - which seems to me to further invalidate the whole idea of using long-run identification. Second, policy and in particular labor market regulations have changed over time, leading to different behaviour, see also Galí, López-Salido and Vallés (2003). The particular change here presents a bit of a challenge. Armchair reasoning would suggest that labor markets were more flexible in prewar years than in postwar years. Furthermore, financial markets may have been less efficient back then. If so, can we think of models explaining the different responses, documented by Francis and Ramey? That strikes me as an interesting research agenda.

Third, perhaps the stochastic properties are fairly stable, but inference based on long-run identification simply is too fragile. In Uhlig (2004), I argue for a variety of reasons to rely on medium-run identification instead. The technical details are in that paper, but the idea is this. While long-run identification finds that shock (or shock direction), which explains as much as possible of the variance of the revision of the long-run (more precisely, the infinite-horizon) forecast in productivity, medium-run identification seeks that shock which explains as much as possible of that variance for some medium-run forecast revision, say, three years out. The results are in figure E. Now, the impulse responses for labor remain a lot more stable, which makes me want to trust these results more than the results from long-run identification. According to these results, labor does not fall much in the prewar years either in response to a technology shock. Actually, labor does not react much at all - which is effectively a restatement of the low labor-productivity correlations of table 1.

A careful reader might also note the larger numbers. According to the data I received from Francis and Ramey, labor productivity growth has an annual standard deviation of 3.1% for the entire sample. The MLE standard deviation of the one-step ahead prediction error for productivity is 2.7%, when using 4 lags and a constant in a bivariate VAR and logs of the data. The squares of the impulse-responses of productivity at date 0 should add up to this number, if one uses shocks one standard deviation in size, so I suspect that some rescaling somewhere is the cause of the difference between the results here and in Francis and Ramey. More importantly, it should be noted that this standard deviation is about 3 times as large in prewar data than in postwar data, pointing to changing stochastic properties across the sample. This again casts doubt on long-run identification. It also says that the Francis-Ramey results are dominated by the prewar sample, since that part of the sample contains a lot more variance than the postwar sample. Given that the data for the prewar years is probably not nearly as good as the data for the postwar years, caution may be advised in putting too much weight on these results.

Francis and Ramey do not provide error bands for their impulse responses. When I tried to recalculate their results, using my Bayesian methodology with a "very long" medium-run identification, I found a very wide error band for the prewar sample. Again, this may be more evidence for fragility of the long-run approach.

Finally, it is laudable that Francis and Ramey check, that their technology shocks are not caused by government spending or M2. There is a simpler way of achieving an identification guaranteeing that: one can simply add government spending and M2 to the VAR. It would have been interesting to see the results from that exercise.

4. Theory

Neoclassical growth theory and its cousins asks one to focus on total factor productivity rather than labor productivity alone, and they ask one to consider the accumulation of capital as a key component in the low-frequency movements of labor productivity. Indeed, Chari, Kehoe and McGrattan (2004) have recently shown how leaving away capital in the sort of empirical exercise provided by Francis and Ramey can lead to serious problems from misspecification and to potentially very misleading results, in particular, when using differenced data for hours. They show how a standard real business cycle model (where technology shocks lead to rises in hours worked) would generate data, which would deliver the results found by Francis and Ramey, that hours fall in response to a technology shock identified in a bivariate VAR from long-run restrictions. Chari, Kehoe and McGrattan even go so far as to suggest that one should be generally skeptical about using SVARs. But one can also give their results a constructive reading and add capital into the picture: so let me pursue that.

To construct capital, I start from steady state capital in 1889, calculated private output in 1889 times the average investment-output ratio divided by (1 - (1-delta)/g), where g is the average growth factor of output, i.e. the average of y(t)/y(t-1), and delta is the annual depreciation rate, set to 10 percent. I calculated the investment-output ratio to be 0.25 from NIPA postwar data, where I included durable consumption with gross private domestic investment, and excluded government spending from output. For the same reason, I rescaled the raw investment series from 1889 to 2002 by the factor 1.44 in order to roughly capture the otherwise unmeasured investment in consumer durables (these rescalings do not make much of a difference in the end, since we use the logs of all series in our estimations anyways). I then calculated capital via k(t) = (1-delta) k(t-1) + inv(t), and used a Cobb-Douglas production function with a capital share of one third to calculate total factor productivity. In order to also account for potential policy influences, I now estimate a VAR with six variables, i.e. TFP, adjusted hours, capital, government spending, M2 and the dividend tax rate provided by McGrattan, see figure H and the comment below.

When using the entire sample and medium-run identification, seeking the shock which explains as much as possible of the variance of total factor productivity three years after the shock (see Uhlig, 2004, for details), the results look remarkably reasonable, see figure F. In particular, hours worked move up substantially, following a technology shock identified in this manner, although the initial response is somewhat muted, and the peak response is about two years after the shock. Nonetheless, this looks very much to be in accord with standard real business cycle theory. One item that looks a little puzzling is the rather uncertain response of capital, though.

The response of capital looks much worse, when using long-run identification, though (implemented as medium-run identification applied to the 20-year ahead revision of the forecast for total factor productivity). Now, not only labor, but also capital falls after an initial increase in total factor productivity. Indeed, this is already true in the simple, bivariate VAR of labor productivity and adjusted hours, when adding the capital series: the results are in figure G. I suspect that Francis and Ramey would find the same, if they included capital in their econometric exercise.

This finding seems rather odd at first. It is hard to think of reasonable theories, according to which technological progress would lead to disinvestment. What then, might be going on?

A hint is already in figure E, possibly partly explaining even the muted response of capital in that figure: the capital income tax rate rises substantially in response to a technology shock identified from medium-run restrictions. Indeed, figure H juxtaposes the output-to-capital ratio with that tax rate series, where we added 50% to the tax rate just in order to show both series on the same scale. First of all, that tax rate shows large and very persistent movements. Second, the output-to-capital ratio comoves with this series remarkably closely. Whether it is appropriate to interpret that tax rate as the marginal dividend tax rate faced by entrepreneurs or not, does not even matter much. It is certainly plausible that this tax rate indicates the general investment climate that entrepreneurs are facing. Governments which do not hesitate to impose a high dividend-tax rate probably do not hesitate to expropriate entrepreneurs in other ways as well. So, figure H is very much what one would expect, following some simple neoclassical reasoning. And persistent changes in capital income tax rates can surely lead to persistent changes in the capital-labor ratio and thus to persistent changes in labor productivity, which have nothing whatsoever to do with technology.

Francis and Ramey point to a variety of other references to rule out changes in the tax treatment of dividends as a source of their identified technology shocks. Perhaps, they are right. But given the rather dramatic shape of figure H, caution may be advised. Mountford and Uhlig, for example, argue that tax rate changes often are debated for a long time in the public and in parliament before they are actually implemented: so the statistical finding that technology shocks are not Granger-caused by past changes in the capital-income tax rate may provide a very misleading picture. Case studies and additional research on this topic would help.

But capital tax rate changes may not be the only additional cause of long-run movements in labor productivity. Uhlig (2004) argues that gradual and persistent changes in the social attitude towards the work place lead to similarly persistent distortions in the way labor input is measured and thus to persistent changes in labor productivity. For many white-collar workers, the work place nowadays has become a central part of their social life, with possibilities to surf the Internet or to meet marriage partners. And finally, any endogenous growth theory makes any shock have persistent effects on labor productivity.

Clearly, the main feature of labor productivity is the fact that it is trending upwards, see the figure in Francis and Ramey. There is little doubt that this is due to technological progress. But whether the persistent random fluctuations around this trend are due to random fluctuations in technological progress also should be much in doubt. Only the simplest of theories would lead one to believe that this is so.

5. Conclusions.

The paper by Francis and Ramey is an excellent and careful contribution to a growing literature, investigating whether technology shocks lead to a fall or a rise in hours worked. They provide a long-run data set which is a genuine gift to the profession and which I expect to be used a lot for a variety of purposes: certainly, I have already made ample use of it. Equipped with that data set, they show that hours worked per capita really have not changed all that much during the last 100 or so years: a bold and fascinating claim, which needs to be subjected to further research and scrutiny. Likewise, their data implies that hours and labor productivity are at most mildly correlated over the cycle, which might explain some of their findings and which could lead to a substantial rethinking about business cycles. Next, Francis and Ramey provide a careful investigation of the data, using a variety of specifications for their VAR and taking a variety of recent suggestions and criticisms on board. This is a fine and informative piece, which pushes the frontier forward.

But I do not buy into their conclusions yet. First, their historical accounting of the labor force is interesting and much more careful than what one usually sees. But for a variety of reasons, it does not seem to go far enough. And mismeasurements here imply mismeasurements in the long-run movements of labor productivity, which their methodology seeks to identify.

Second, the long-run restrictions seem to be too fragile an identification device to provide convincing conclusions. The very fact that the first and the second half of the sample look rather different already sheds substantial doubt on any strategy seeking to identify anything from long-run behaviour. Medium-run identification provides more robust results, and may therefore provide a more sensible alternative.

Finally, theory suggests that one should take capital accumulation and the policy variables related to it into account. There is strong comovement in the data over the last 100-something years between the output-capital ratio and the dividend tax rate: this alone might lead to a serious distortions when identifying technology shocks from long-run movements in labor productivity. Furthermore, Chari, Kehoe and McGrattan (2004) have shown, that leaving out capital in the sort of exercise performed by Francis and Ramey can lead to serious problems from misspecification. Finally, theory suggests many reasons why labor productivity fluctuates in the long-run, aside from technology. If the medium-run fluctuations in total factor productivity are more strongly dominated by technology shocks than the long-run, then medium-run identification is more informative about the impact of technology shocks. Figure F shows, that technology shocks identified from medium-run identification seem to lead to rising rather than falling hours worked.

What we need are good theories, explaining the key business cycle facts, like the comovements between labor productivity and output, and the relative volatilities of consumption, output and investment over the cycle. The original real business cycle model has been attacked a lot, and perhaps it is false. But what should replace it? That question still remains unanswered. The evidence provided by Francis and Ramey tells us, that we need to think even harder about the answer, even if that evidence is not conclusive enough either.

References

Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde (2002), "An Estimated Dynamic General Equilibrium Model for Monetary Policy Analysis," working paper, Northwestern University.

Blanchard, Olivier J. and Danny Quah, "The Dynamic Effects of Aggregate Demand and Supply Disturbances," American Economic Review (1989), 79, 655-673.

Chari, V.V., Patrick Kehoe and Ellen R. McGrattan (2004), "Are Structural VARs Useful Guides for Developing Business Cycle Theories?", draft, Federal Reserve Bank of Minneapolis.

Cooley, Thomas F., ed., Frontiers of Business Cycle Research, Princeton University Press, Princeton N.J. (1995)

Christiano, Lawrence J., Martin Eichenbaum and Robert Vigfusson, "The Response of Hours to a Technology Shock: Evidence based on Direct Measures of Technology", Journal of the European Economic Association, April-May 2004, 2 (2-3): 381-395.

Galí, Jordi, David López-Salido and Javier Vallés, "Technology and Monetary Policy: Assessing the Fed's Performance." Journal of Monetary Economics, 50 (2003), 723-743.

Hansen, Gary D. (1985), "Indivisble Labor and the Business Cycle," Journal of Monetary Economics, vol. 16: 309-327.

Kydland, Finn E. and Edward C. Prescott (1982), "Time to Build and Aggregate Fluctuations," Econometrica, vol. 50(6). 1345-1371.

Mountford, Andrew and Harald Uhlig (2004), "What are the effects of fiscal policy shocks?", draft, Humboldt Universität zu Berlin.

Ravn, Morten and Harald Uhlig, "On Adjusting the HP-Filter for the Frequency of Observations," Review of Economics and Statistics, 84(2), May 2002, 371-76.

Sims, Christopher A. and Harald Uhlig, "Understanding Unit Rooters: A Helicopter Tour," Econometrica, vol. 59, no. 6 (1991): 1591-1599.

Uhlig, Harald, "What Macroeconomists Should Know about Unit Roots: A Bayesian Perspective," Econometric Theory, vol. 10., nos 3/4 (1994): 645-671.

Uhlig, Harald, "Do Technology Shocks lead to a Fall in Total Hours Worked? ", Journal of the European Economic Association, April-May (2004), 2 (2-3): 361-371.

Table 1.

Correlations of private sector output, private sector hours, adjusted hours (see Francis-Ramey), private sector productivity, consumption, investment and government spending. The numbers above the diagonal are for the first differences of the logs of the series, whereas the numbers below the diagonal are for the HP-filtered series, using λ=7, see Ravn and Uhlig (2002). All data are from Francis and Ramey. Note the low correlation between labor and labor productivity, possibly explaining some of the Francis-Ramey findings.

Table 1A
Correlations for 1889-2002

	y	n	n adj.	prod.	c	x	g
y		0.82	0.83	0.71	0.73	0.59	0.20
n	0.85		0.92	0.18	0.56	0.55	0.11
n adj.	0.85	0.92		0.28	0.43	0.38	0.39
prod.	0.72	0.25	0.36		0.56	0.34	0.20
c	0.69	0.57	0.40	0.53		0.61	-0.27
x	0.55	0.52	0.31	0.34	0.70		-0.46
g	0.29	0.21	0.50	0.25	-0.29	-0.45	

Table 1B.
Correlations for 1950-2002.

	y	n	n adj.	prod.	c	x	g
y		0.85	0.83	0.53	0.85	0.89	0.18
n	0.90		0.95	0.00	0.64	0.80	0.11
n adj.	0.86	0.98		0.03	0.60	0.72	0.23
prod.	0.50	0.07	0.01		0.59	0.40	0.18
c	0.87	0.73	0.67	0.54		0.75	-0.07
x	0.91	0.81	0.75	0.46	0.84		-0.15
g	0.12	0.11	0.15	0.06	-0.19	-0.20	

Table 1C.
Correlations for 1889-1940.

	y	n	n adj.	prod.	c	x	g
y		0.83	0.83	0.74	0.84	0.79	-0.02
n	0.87		0.97	0.23	0.65	0.75	-0.05
n adj.	0.86	0.98		0.26	0.61	0.72	0.11
prod.	0.77	0.35	0.36		0.67	0.47	0.04
c	0.85	0.69	0.64	0.71		0.57	-0.17
x	0.82	0.76	0.72	0.56	0.65		-0.29
g	0.10	0.11	0.27	0.05	-0.13	-0.17	

Figure A

Death rates for children per 1000, age 1 to 4. Source: AmeriStat, Population Reference Bureau

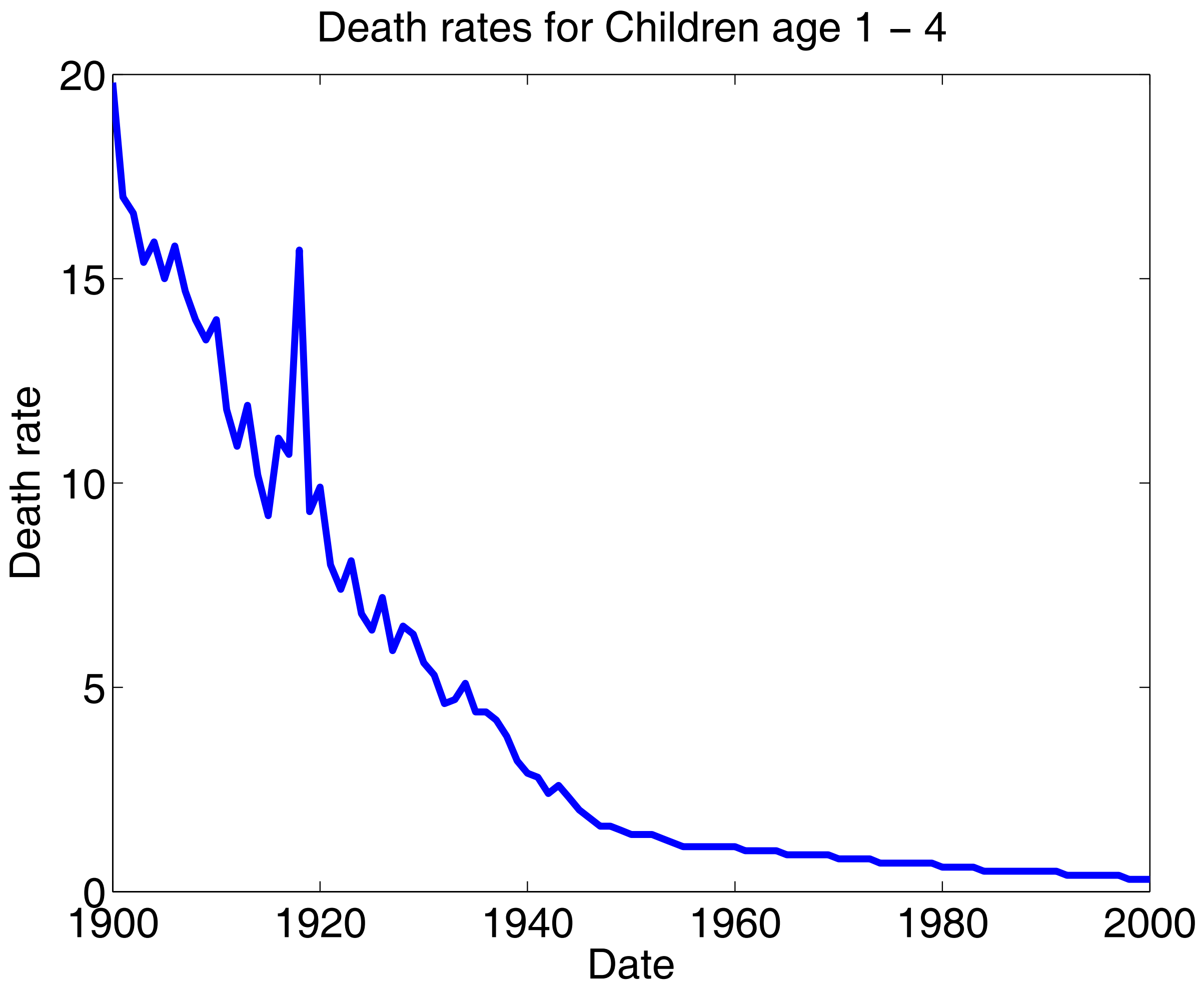

Figure B.

Fraction of children age 5 to 16 in total population.

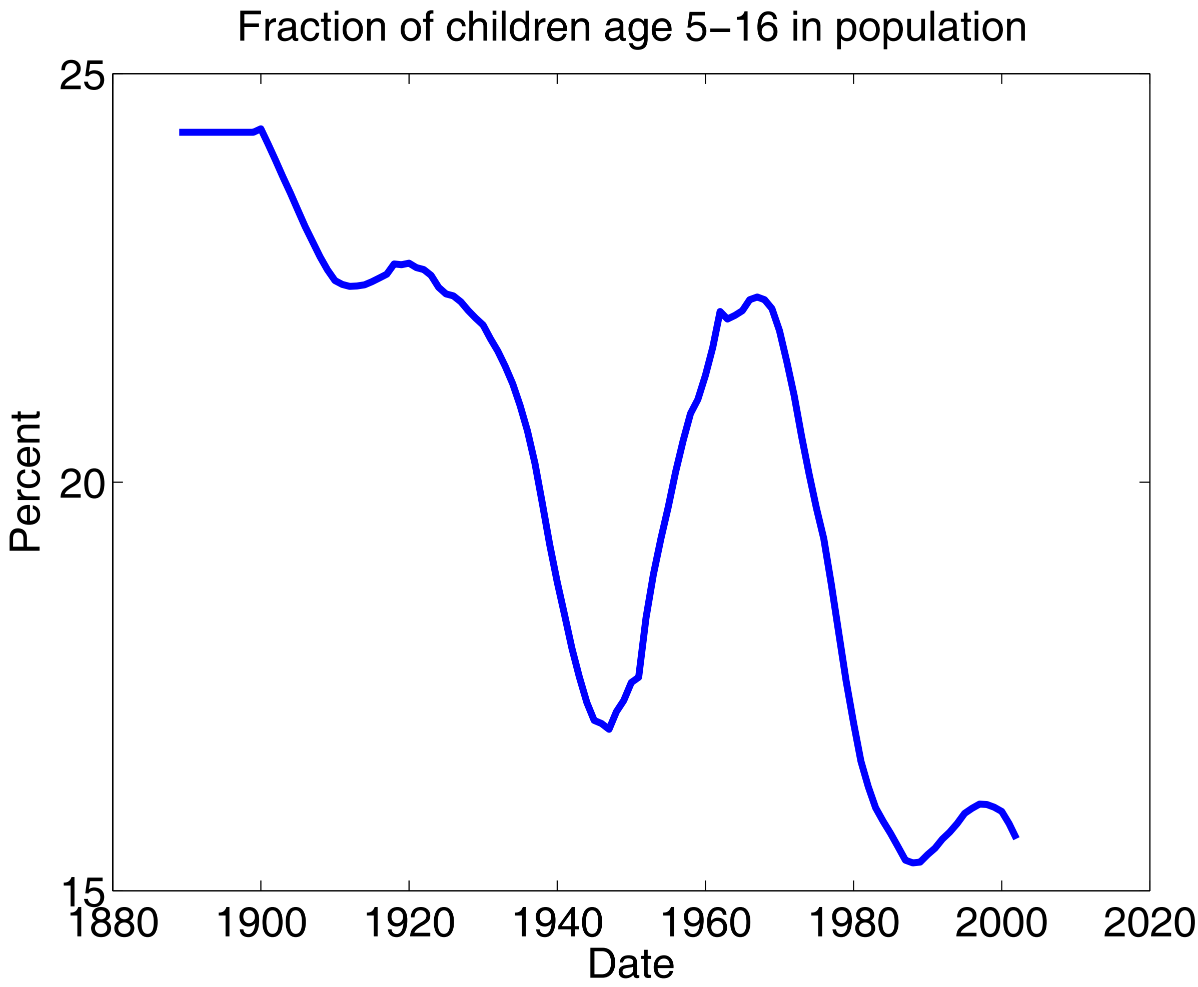

Figure C.

Fraction of children age 5 to 16 in labor force, not accounting for school enrollment, i.e. (pop - pop16 - pop4)/(pop - pop4 - pop65 - govemp)

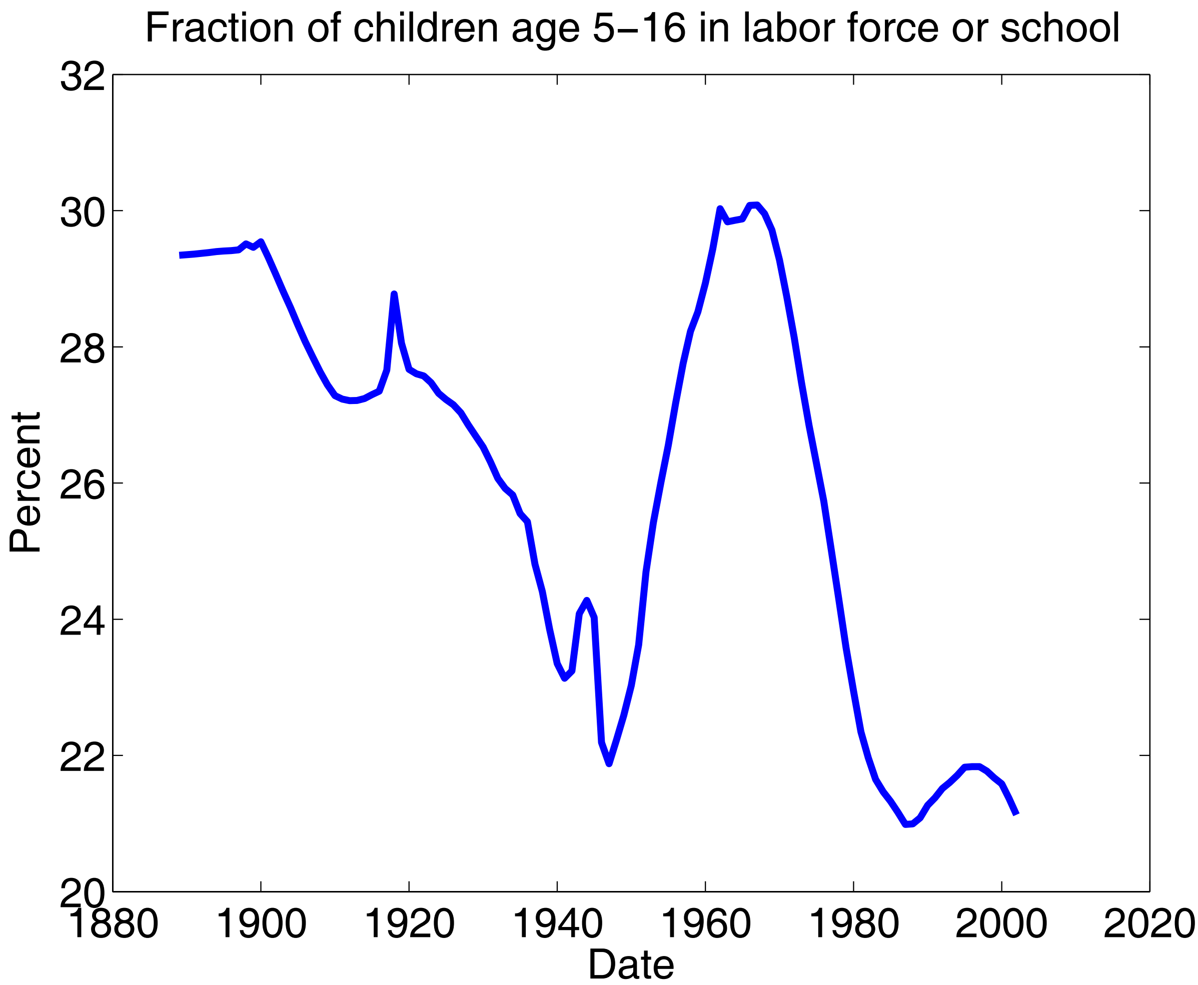

Figure D.

Fraction of children age 5 to 16 in labor force, accounting for school enrollment, i.e. (pop - pop16 - pop4 - school)/(pop - pop4 - pop65 - govemp - school)

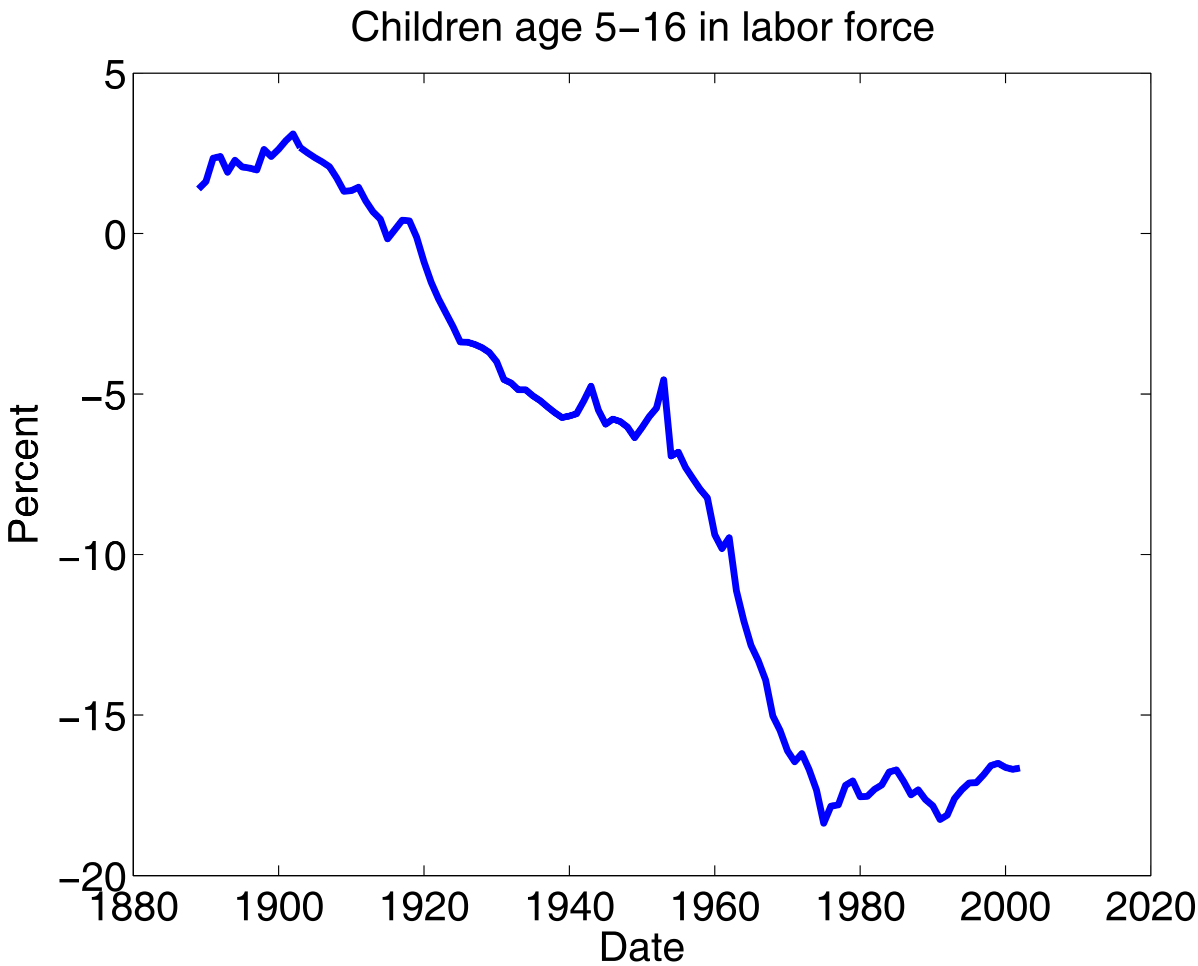

Figure E

Results from medium-run identification, i.e. for the shock that explains as much as possible of the 3-year ahead revision of the forecast for productivity. The horizontal line is the MLE of the one standard deviation of the one-step ahead prediction error for comparison.

	labor	productivity
All data	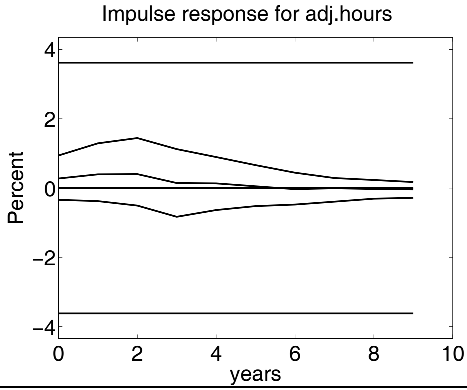	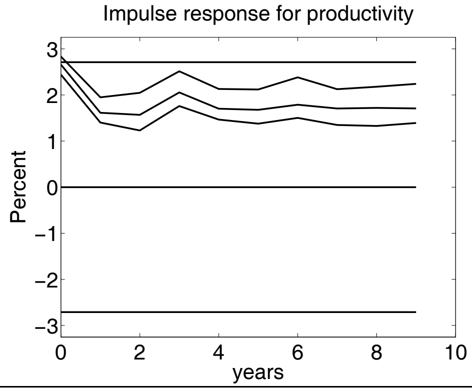
1889-1940	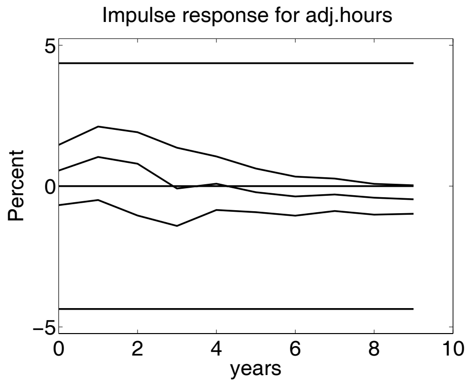	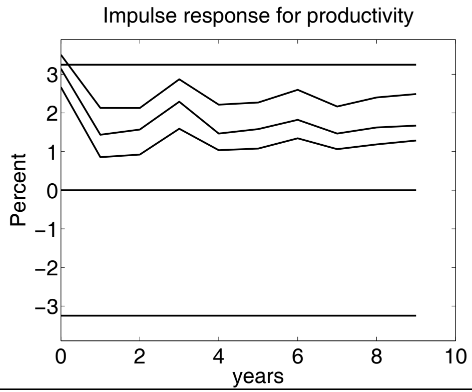
1949-2002	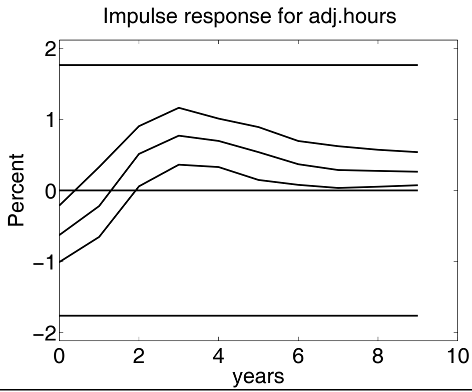	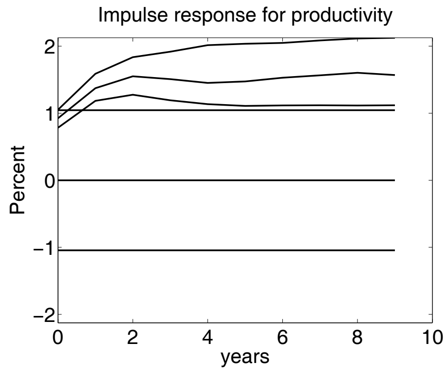

Figure F

Impulse responses to a technology shock identified from medium run identification applied to to the revision of the three-year forecast of total factor productivity, and using a six-variable VAR, including government spending, M2 and capital income tax rates.

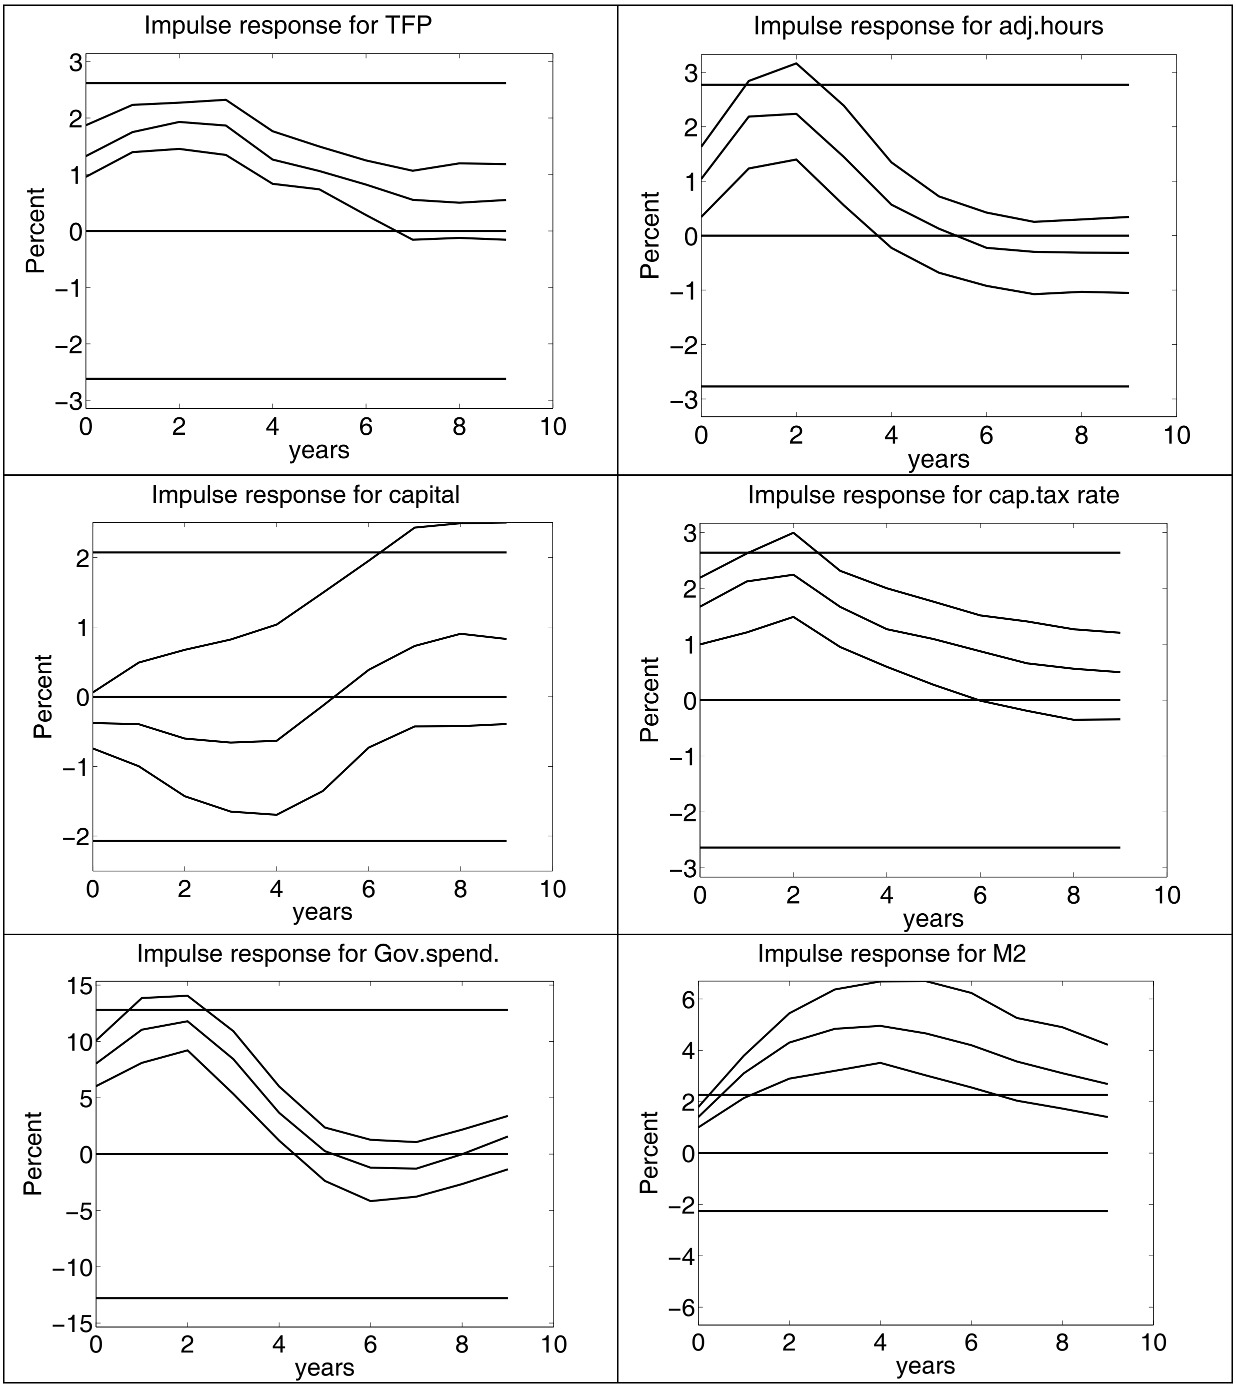

Figure G

Results from long-run identification in a VAR with labor productivity, adjusted hours (in levels) and capital. Note, that capital falls, following a "technology shock" identified this way.

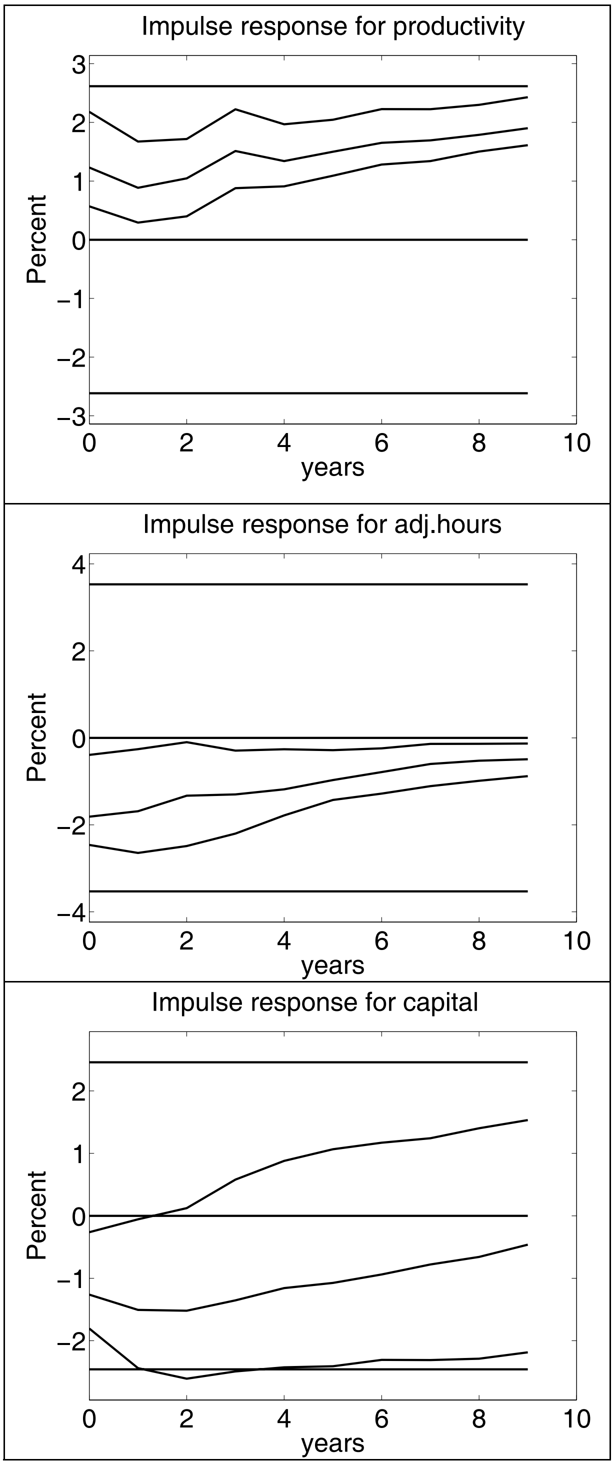

Figure H

Juxtaposing the annual output-capital ratio to capital income tax rates (plus 50% in order to show both on the same scale).

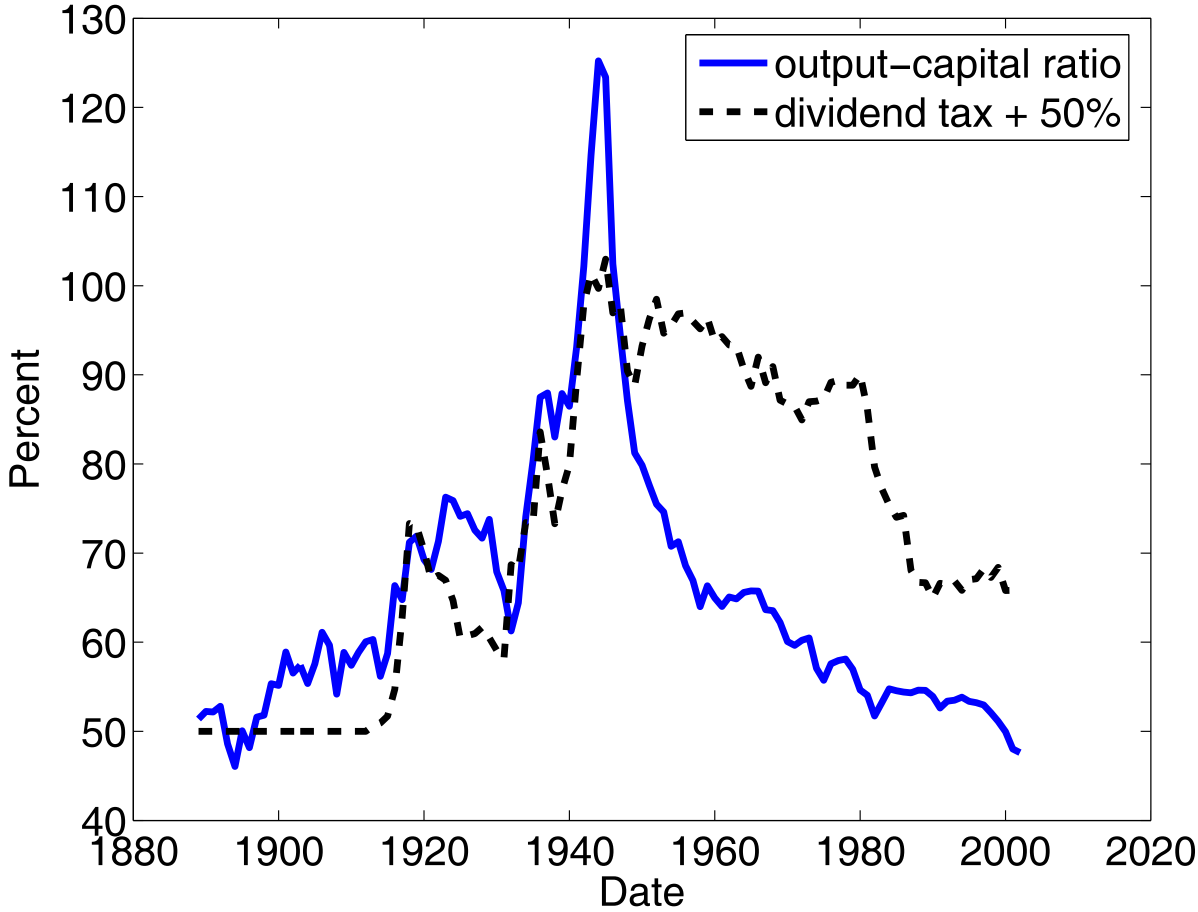

SFB 649 Discussion Paper Series 2006

For a complete list of Discussion Papers published by the SFB 649, please visit http://sfb649.wiwi.hu-berlin.de.

001 "Calibration Risk for Exotic Options" by Kai Detlefsen and Wolfgang K. Härdle, January 2006.
002 "Calibration Design of Implied Volatility Surfaces" by Kai Detlefsen and Wolfgang K. Härdle, January 2006.
003 "On the Appropriateness of Inappropriate VaR Models" by Wolfgang Härdle, Zdeněk Hlávka and Gerhard Stahl, January 2006.
004 "Regional Labor Markets, Network Externalities and Migration: The Case of German Reunification" by Harald Uhlig, January/February 2006.
005 "British Interest Rate Convergence between the US and Europe: A Recursive Cointegration Analysis" by Enzo Weber, January 2006.
006 "A Combined Approach for Segment-Specific Analysis of Market Basket Data" by Yasemin Boztuğ and Thomas Reutterer, January 2006.
007 "Robust utility maximization in a stochastic factor model" by Daniel Hernández–Hernández and Alexander Schied, January 2006.
008 "Economic Growth of Agglomerations and Geographic Concentration of Industries - Evidence for Germany" by Kurt Geppert, Martin Gornig and Axel Werwatz, January 2006.
009 "Institutions, Bargaining Power and Labor Shares" by Benjamin Bental and Dominique Demougin, January 2006.
010 "Common Functional Principal Components" by Michal Benko, Wolfgang Härdle and Alois Kneip, Jauary 2006.
011 "VAR Modeling for Dynamic Semiparametric Factors of Volatility Strings" by Ralf Brüggemann, Wolfgang Härdle, Julius Mungo and Carsten Trenkler, February 2006.
012 "Bootstrapping Systems Cointegration Tests with a Prior Adjustment for Deterministic Terms" by Carsten Trenkler, February 2006.
013 "Penalties and Optimality in Financial Contracts: Taking Stock" by Michel A. Robe, Eva-Maria Steiger and Pierre-Armand Michel, February 2006.
014 "Core Labour Standards and FDI: Friends or Foes? The Case of Child Labour" by Sebastian Braun, February 2006.
015 "Graphical Data Representation in Bankruptcy Analysis" by Wolfgang Härdle, Rouslan Moro and Dorothea Schäfer, February 2006.
016 "Fiscal Policy Effects in the European Union" by Andreas Thams, February 2006.
017 "Estimation with the Nested Logit Model: Specifications and Software Particularities" by Nadja Silberhorn, Yasemin Boztuğ and Lutz Hildebrandt, March 2006.
018 "The Bologna Process: How student mobility affects multi-cultural skills and educational quality" by Lydia Mechtenberg and Roland Strausz, March 2006.
019 "Cheap Talk in the Classroom" by Lydia Mechtenberg, March 2006.
020 "Time Dependent Relative Risk Aversion" by Enzo Giacomini, Michael Handel and Wolfgang Härdle, March 2006.
021 "Finite Sample Properties of Impulse Response Intervals in SVECMs with Long-Run Identifying Restrictions" by Ralf Brüggemann, March 2006.
022 "Barrier Option Hedging under Constraints: A Viscosity Approach" by Imen Bentahar and Bruno Bouchard, March 2006.

SFB 649, Spandauer Straße 1, D-10178 Berlin
http://sfb649.wiwi.hu-berlin.de

This research was supported by the Deutsche Forschungsgemeinschaft through the SFB 649 "Economic Risk".

023 "How Far Are We From The Slippery Slope? The Laffer Curve Revisited" by Mathias Trabandt and Harald Uhlig, April 2006.
024 "e-Learning Statistics – A Selective Review" by Wolfgang Härdle, Sigbert Klinke and Uwe Ziegenhagen, April 2006.
025 "Macroeconomic Regime Switches and Speculative Attacks" by Bartosz Maćkowiak, April 2006.
026 "External Shocks, U.S. Monetary Policy and Macroeconomic Fluctuations in Emerging Markets" by Bartosz Maćkowiak, April 2006.
027 "Institutional Competition, Political Process and Holdup" by Bruno Deffains and Dominique Demougin, April 2006.
028 "Technological Choice under Organizational Diseconomies of Scale" by Dominique Demougin and Anja Schöttner, April 2006.
029 "Tail Conditional Expectation for vector-valued Risks" by Imen Bentahar, April 2006.
030 "Approximate Solutions to Dynamic Models – Linear Methods" by Harald Uhlig, April 2006.
031 "Exploratory Graphics of a Financial Dataset" by Antony Unwin, Martin Theus and Wolfgang Härdle, April 2006.
032 "When did the 2001 recession *really* start?" by Jörg Polzehl, Vladimir Spokoiny and Cătălin Stărică, April 2006.
033 "Varying coefficient GARCH versus local constant volatility modeling. Comparison of the predictive power" by Jörg Polzehl and Vladimir Spokoiny, April 2006.
034 "Spectral calibration of exponential Lévy Models [1]" by Denis Belomestny and Markus Reiß, April 2006.
035 "Spectral calibration of exponential Lévy Models [2]" by Denis Belomestny and Markus Reiß, April 2006.
036 "Spatial aggregation of local likelihood estimates with applications to classification" by Denis Belomestny and Vladimir Spokoiny, April 2006.
037 "A jump-diffusion Libor model and its robust calibration" by Denis Belomestny and John Schoenmakers, April 2006.
038 "Adaptive Simulation Algorithms for Pricing American and Bermudan Options by Local Analysis of Financial Market" by Denis Belomestny and Grigori N. Milstein, April 2006.
039 "Macroeconomic Integration in Asia Pacific: Common Stochastic Trends and Business Cycle Coherence" by Enzo Weber, May 2006.
040 "In Search of Non-Gaussian Components of a High-Dimensional Distribution" by Gilles Blanchard, Motoaki Kawanabe, Masashi Sugiyama, Vladimir Spokoiny and Klaus-Robert Müller, May 2006.
041 "Forward and reverse representations for Markov chains" by Grigori N. Milstein, John G. M. Schoenmakers and Vladimir Spokoiny, May 2006.
042 "Discussion of 'The Source of Historical Economic Fluctuations: An Analysis using Long-Run Restrictions' by Neville Francis and Valerie A. Ramey" by Harald Uhlig, May 2006.

SFB 649, Spandauer Straße 1, D-10178 Berlin
http://sfb649.wiwi.hu-berlin.de

This research was supported by the Deutsche Forschungsgemeinschaft through the SFB 649 "Economic Risk".