Cyclical Skill-Biased Technological Change[*]

Almut Balleer
University of Bonn
aballeer@uni-bonn.de

Thijs van Rens
CREI, Universitat Pompeu Fabra, IZA and CEPR
thijs.vanrens@upf.edu

March 2008

Preliminary and incomplete.
Please do not distribute.
www.crei.cat/~vanrens/skillbias

Abstract

Over the past two decades, technological progress has been biased towards making skilled labor more productive. The evidence for this finding is based on the persistent parallel increase in the skill premium and the supply of skilled workers. What are the implications of skill-biased technological change for the business cycle? To answer this question, we use the CPS outgoing rotation groups to construct quarterly series for the price and quantity of skill. The unconditional correlation of the skill premium with the cycle is zero. However, using a structural VAR with long run restrictions, we find that technology shocks substantially increase the premium. Investment-specific technology shocks are not skill-biased and our findings suggest that capital and skill are (mildly) substitutable in aggregate production.

JEL codes: E24 E32 J24 J31
Keywords: skill-biased technology, skill premium, VAR, long-run restrictions, capital-skill complementarity, business cycle

1 Introduction

The US, as well as many other industrialized countries, have seen a marked increase in the skill premium over the past two decades. Over the same period, the average education level of the workforce also rose substantially. This parallel rise in the price and quantity of skill points towards an increase in the demand for skill that exceeded the increase in the supply of skilled workers. A commonly accepted explanation for this finding is skill bias in technological progress: newly developed technologies require relatively more educated and less uneducated workers (Katz and Murphy (1992); Autor et al. (1998); Acemoglu (2002); Autor et al. (2005) and Autor et al. (2008)).

In this paper, we explore the implications of skill-biased technological change for business cycle fluctuations. Existing studies, including those mentioned above, have focused on slow moving trends in the data. These papers use annual data, constructed from a variety of worker-level data sources. Annual data are not suitable to analyze business cycle fluctuations and we construct a quarterly series for the skill premium and the relative supply of skill over the 1979:I-2006:II period, using the Current Population Survey (CPS) outgoing rotation groups. Every month, about one fourth of workers in the CPS is in an outgoing rotation group, meaning they are being interviewed for the fourth month in a row and are therefore being rotated out of the sample. These workers are asked about earnings and hours as well as education and other personal characteristics. We use this information to calculate the skill premium as the log ratio of wages of college graduate equivalent workers over high school graduate equivalents, controlling for experience and other standard Mincer controls.[1]

The skill premium is close to acyclical over our sample period. If we think of business cycles as being driven by technology shocks, one might conclude from this observation that most of the higher frequency movements in the skill premium are driven by fluctuations in the supply of skill rather than its demand. Acemoglu (2002) and Autor et al. (2005) reach this conclusion, although from a different observation: once we detrend the skill premium and the relative supply of skill, the two series are negatively rather than positively correlated.[2] Our estimates confirm that shocks to the supply of skill are an important determinant of fluctuations in the skill premium. However, we also find significant effects of technology shocks on the premium.

Unconditional correlations are the result of a variety of shocks to the

---

[1]Lindquist (2004) also construct a quarterly series for the skill premium from the CPS outgoing rotation groups, but does not control for multiple education levels and other sources of worker heterogeneity, see section 2.3.

[2]Acemoglu (2002) regresses the skill premium on the relative supply of skill controlling for a linear trend and finds a coefficient of -0.74 (table 2, column 1). Autor et al. (2005) detrend the time series and show graphically that there is strong comovement in both series, but they move in opposite directions (figure 7, panel A).

economy, which may obscure the effects of changes in technology. We use a structural vector autoregression (VAR) to estimate the conditional response of the skill premium and the relative supply of skill to technology shocks. In order to control for fluctuations in the supply of skill, we separately identify skill supply shocks using a short run restriction, assuming that the supply of skilled workers is predetermined. We then identify technology shocks using long-run restrictions as in Blanchard and Quah (1989).

We start by assessing the overall skill bias in technology shocks, identified following Galí (1999) as the only shocks that affect labor productivity in the long run. Improvements in technology significantly increase the skill premium. This effect is realized in full within a year, providing evidence in favor of skill-biased technological change and its potential importance for business cycle fluctuations.

Next, we ask the question whether all technological changes are skill-biased or whether there is a difference between skill-biased and skill-neutral technology shocks. We propose a strategy to identify skill-biased technology (SBT) shocks from a long-run restriction, arguing that SBT shocks are the only shocks that affect the skill premium in the long run. Skill-neutral technology shocks are all remaining sources of permanent changes in labor productivity.

Skill-biased technology shocks are similar to skill-neutral technological changes in that they increase labor productivity. However, they have different implications for other aggregate variables. In particular, a positive SBT shock leads to a much larger reduction in total hours worked than a skill-neutral technology shock. In addition, SBT shocks increase the supply of skill in the long run, as we would expect, whereas skill-neutral shocks lead to reduced supply of skill.

Finally, we attempt to better understand what drives skill-biased technological change. In particular, we evaluate the hypothesis, put forward by Krusell et al. (2000), that skill-biased technological change is the result of an increase in the relative productivity of the investment-goods producing sector. It is a well-documented fact that, over the same period that the skill premium has risen, the relative price of investment goods (software, equipment structures) has fallen substantially, providing evidence for investment-specific technological change (Gordon (1990); Greenwood et al. (1997); Cummins and Violante (2002)). Krusell et al. (2000) show that if capital and skilled labor are complements in the aggregate production function, investment-specific technological progress can explain the increasing trend in the skill premium, because the increase in the capital-labor ratio makes skilled labor relatively more productive.

We identify investment-specific technology shocks, following Fisher (2006), as the only shocks that affect the relative price of investment in the long run. An investment-specific improvement in technology lowers the relative price of investment goods. The remaining shocks that affect labor produc-

tivity in the long run, are investment-neutral technology shocks. We find that investment-specific technology shocks have a significant, but negative effect on the skill premium, while investment-neutral technology shocks have a positive effect on this variable. Conversely, skill-biased technology shocks, identified as described above, raise the relative price of investment goods. This evidence is in direct contradiction with the hypothesis of capital-skill complementarity, suggesting instead that capital and skill are (to some degree) substitutes in the aggregate production process.

The remainder of this paper is organized as follows. Section 2 describes our empirical approach. First we define the different shocks to the production technology that we consider, then we discuss how to identify the effects of these shocks using long-run restrictions. We also describe the data that are necessary to estimate these effects and show some descriptive statistics on the cyclicality of our quarterly series for the skill premium and the relative supply of skill. In section 3, we discuss our results from the structural VAR analysis. Section 4 concludes.

2 Empirical Approach

In this section, we outline our approach to estimate the implications of skill-biased technological progress for the business cycle. We start by defining different types of of technological change, discussing various specifications for the aggregate production function. Next, we explain how to identify these different technology shocks from the data using a VAR with long-run restrictions. Finally, we describe the data needed for the identification, including quarterly series for the skill premium and the relative supply and employment of skilled labor, which we construct from micro data.

2.1 Shocks to the production technology

Consider an aggregate production function for output Y_t that takes capital K_t, high skilled labor H_t and low skilled labor L_t as inputs.

$$Y_t = A_t f(K_t, H_t, L_t) \tag{1}$$

The function f satisfies the standard conditions: it is increasing and concave in all its arguments and homogenous of degree one so that there are constant returns to scale. Shocks to total factor productivity A_t are neutral technology shocks, in the sense that they affect the productivity of all inputs in the same proportion.

To allow for skill-biased technology shocks, the literature has typically assumed an aggregate production function of the following form (see e.g.

Katz and Murphy (1992), Katz and Autor (1999), Autor et al. (2008)).

$$Y_t = A_t f(K_t, B_t H_t, L_t)$$

$$= A_t K_t^{\alpha} \left[\beta \left(B_t H_t \right)^{\frac{\sigma-1}{\sigma}} + (1 - \beta) L_t^{\frac{\sigma-1}{\sigma}} \right]^{\frac{(1-\alpha)\sigma}{\sigma-1}} \tag{2}$$

Here, A_t is neutral technology and B_t is skilled labor augmenting technology. An increase in B_t can be skill or unskill biased, depending on the elasticity of substitution between skilled and unskilled labor $\sigma > 0$. Under the assumption that workers' wages are proportional to their marginal product,[3] we can calculate the skill premium directly from production function (2).

$$\log \left(\frac{w_{H,t}}{w_{L,t}} \right) = \log \left(\frac{B_t f_2(K_t, B_t H_t, L_t)}{f_3(K_t, B_t H_t, L_t)} \right)$$

$$= \log \left(\frac{\beta}{1-\beta} \right) + \frac{\sigma-1}{\sigma} \log B_t - \frac{1}{\sigma} \log \left(\frac{H_t}{L_t} \right) \tag{3}$$

where $w_{H,t}$ and $w_{L,t}$ are the wages of high and low skilled workers respectively. If high and low skilled labor are substitutes rather than complements ($\sigma > 1$), the substitution effect of improvements in skilled labor augmenting technology dominates the income effect so that an increase in B_t increases the demand for skill and therefore the skill premium (assuming that the supply curve for skill is downward sloping). The consensus estimate for σ is around 1.5 (see Katz and Murphy (1992), Ciccone and Peri (2006), Teulings and van Rens (2008)), so that we can think of skill-biased technology shocks as increases in B_t, which increase the skill premium.

There are two ways to interpret skill-biased technology shocks to an aggregate production function as in (2). If the production function for all goods in the economy is the same, then we can think of an increase in B_t as a technological development that makes skilled labor more productive in all sectors. Alternatively, we may think that the production in different sectors i requires skilled labor in different proportions β_i of total labor input. In this case, even if skilled and unskilled labor are neither substitutes nor complements within each sector,[4] a sector-specific technology shock to a skill-intensive sector could still increase the skill premium.

A particularly interesting case is an economy that consists of a consumption goods producing sector and an investment goods producing sector. In this economy there are two mechanisms, by which sector-specific shocks may affect the skill premium. First, the input shares for skill might be different

---

[3]A sufficient condition is that labor markets are perfectly competitive, in which case the wage of all workers equals their marginal product. However, even if there are frictions on the labor market, the weaker assumption that wages are proportional to marginal products might still hold.

[4]This is the case where $\sigma_i = 1$ for all i. In the limit for $\sigma \to 1$, production function (2) becomes Cobb-Douglas, so that changes in B_t are indistinguishable from changes in A_t.

across the two sectors as explained above. Because investment goods are used to build up capital, which is an input in the production process, sector-specific shocks may also affect the capital-labor ratio used in production. If capital and skill are complements, as argued by Krusell et al. (2000), then a higher capital labor ratio increases the relative demand for skilled labor and therefore the skill premium.

Suppose the two sectors have identical production functions except for a difference in total factor productivity. In this case, as shown among others by Fisher (2006) and Krusell et al. (2000), the economy can be aggregated to a one-sector economy, where total output is divided between consumption and investment,

$$Y_t = C_t + p_t I_t$$

where the relative price of investment goods p_t reflects technological improvements in the investment goods producing sector. The aggregate production function in this economy, allowing for capital-skill complementarity, is a slightly generalized version of (2).

$$Y_t = A_t \left[\beta \left(\gamma K_t^{\frac{\rho-1}{\rho}} + (1-\gamma)(B_t H_t)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}\frac{\sigma-1}{\sigma}} + (1-\beta) L_t^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (4)$$

where σ is the elasticity of substitution between skilled and unskilled labor as before, except that now it also measures the elasticity of substitution between capital and unskilled labor, ρ is the elasticity of substitution between capital and skilled labor and β and γ are share parameters. As shown by Krusell et al. (2000), improvements in investment-specific technology increase the skill premium if and only if the elasticity of substitution between capital and skilled labor ρ is lower than the elasticity of substitution between capital and unskilled labor σ, i.e. if the production technology displays capital-skill complementarity.

2.2   Identification and Estimation

As explained above, skill-biased technology shocks are shocks to the production technology that affect the skill premium, investment-specific technology shocks change the relative price of investment goods and in the presence of capital-skill complementarity technology shocks may be both investment-specific and skill-biased. Neutral technology shocks increase productivity but do not affect either the relative price or the skill premium. We use data on the skill premium, the relative price of investment goods and productivity to identify these technology shocks in a structural vector-autoregression (VAR).

We identify technology shocks using long-run restrictions, as suggested by Blanchard and Quah (1989). In this procedure, unobserved structural shocks to the economy are identified in two steps. First, we estimate a

reduced form VAR. Second, we map the reduced form coefficients and residuals into structural coefficients and shocks assuming orthogonality between the structural shocks and an identifying restriction. The details of the identification procedure are in appendix B.

The identifying restriction is an assumption on the long-run effects of the structural shocks on the variables in the VAR. For example, (neutral) shocks to total factor productivity are identified as the only shocks that affect labor productivity in the long run, as in Galí (1999). We discuss the specific identifying restrictions used to identify skill-biased and investment-specific technology shock as we describe our results in section 3.

In our basic VAR, we include labor productivity and hours worked, which are required to identify technology shocks. We add the skill premium in order to assess the skill bias and to identify skill-biased technology shocks. Depending on the specification, we also include other variables if these are needed for identification or if their impulse response is of interest. All variables are used in first differences in order to allow for unit roots.[5] We use quarterly data from 1979:I to 2000:IV. This period is relatively short because of data limitations, see section 2.3.

The reduced form is estimated as a Bayesian VAR with a Minnesota prior, similar to Canova et al. (2006), which is a prior on the decay of the coefficients on higher order lags and pushes towards a unit root (in levels). We use this prior for two reasons. First, critics of VARs (e.g. Chari et al. (2005)) have stated that in theory one should employ a VAR with an infinite number of lags in order to correctly identify technology shocks using long run restrictions. The Minnesota prior allows us to generate sensible results for a large number of lags. We use 8 lags and a decay parameter of 3. Second, the prior makes our estimation results more stable in the presence of high frequency variation in the skill premium that is due to measurement error. The prior does not affect the long-run restrictions in any way and we show that our results are robust to the strength of the prior and to estimating the reduced form VAR using ordinary least squares.

For the purposes of this paper, identifying technology shocks using a structural VAR is preferable over alternatives such as constructing Solow residuals. The approach allows us to identify full dynamic impulse responses to the different types of technology shocks by imposing only very little theory (the identifying restrictions). In the case of neutral technology shocks, the structural VAR estimates of the shocks are very similar to Solow residuals, once the latter have been properly corrected for non-technological effects

---

[5]In the context of the identification of neutral technology shocks, there has been a debate in the literature whether hours worked should be included in levels (Christiano et al. (2003)) or in first differences (Galí and Rabanal (2004)). Canova et al. (2006) show that once the very low frequencies are purged out from the data, the results of Galí (1999) are robust to using hours worked in levels. In all specifications, we verified that our results are also robust to this choice.

such as varying utilization of capital and labor, nonconstant returns, imperfect competition and aggregation effects (Basu et al. (2006)). While the estimation of other types of technology shocks using production function estimation would require taking a strong stance on the production structure, the structural VAR allows to identify these shocks using transparent and easy-to-understand assumptions that can be justified in a wide range of macroeconomic models.

2.3 Data

We construct quarterly series for the skill premium and the relative employment and supply of skill using individual-level wage and education data from the CPS outgoing rotation groups. This survey has been administered every month since 1979 so that our series runs from 1979:1 to 2006:2.[6] Wages are usual hourly earnings (weekly earnings divided by usual weekly hours for weekly workers) and are corrected for top-coding and outliers. We limit our sample to wage and salary workers between 16 and 64 years old in the private, non-farm business sector and weight average wages by the CPS-ORG sampling weights as well hours worked in order to replicate aggregate wages as close as possible, see Haefke et al. (2007) for a more detailed discussion. Education is measured in five categories (less than high school, high school degree, some college, college degree, more than college) and made consistent over the full sample period following Jaeger (1997). In an average quarter, we have wage and education data for about 35,000 workers. A more detailed description of the data may be found in appendix A.

Our measure for the skill premium is the log wage differential between college graduates and high school graduates. The relative employment and supply of skill are defined as the log ratio of the number of college graduates over the number of high school graduates in the population and the workforce respectively. Following Autor et al. (2005), we map the five education levels in the data to college and high school graduate equivalents and control for changes in experience, gender, race, ethnicity and marital status. To do this, we first estimate a standard Mincerian earnings function for log wages. The predicted values from this regression for males and females at 5 education levels in 5 ten-year experience groups yield average wages for 50 education-gender-experience cohorts keeping constant the other control variables. We then calculate the number of workers in each cell as a fraction of the workforce or population. Dividing by a reference category, this procedure gives us relative the prices and quantities of skill for 50 skill categories. Finally, we aggregate to two skill types by averaging relative prices

---

[6]The BLS started asking questions about earnings in the outgoing rotation group (ORG) surveys in 1979. The March supplement goes back much further (till 1963), but does not allow to construct wage series at higher frequencies than annual. The same is true for the May supplement, the predecessor of the earnings questions in the ORG survey.

using average quantity weights and averaging quantities using average price weights.[7]

The way we measure the skill premium and the relative employment and supply of skill allows easy comparison to models with workers of only two skill levels. Yet, the measures do justice to the greater degree of heterogeneity in the data. This is necessary to ensure that changes in the price of skill are correctly attributed to changes in the skill premium and changes in the quantity of skill to the relative employment or supply of skill. Suppose, for example, that there is an increase in the number of workers with a masters degree. This represents an increase in the supply of skill. However, a naive measure of the relative supply, which just counts the number of workers with at least a college degree, would not reflect this increase. Moreover, if workers with a masters degree earn on average higher wages than workers with a bachelors degree only, then a naive measure of the skill premium would increase. In our measures, this increase in the supply of skill would leave the skill premium unchanged and increase the relative supply measure. In section 3, we explore the robustness of our results to alternative ways to construct these series.

Figure 1 plots our quarterly series for the log wage premium of college over high school graduates. As documented in previous studies, the data show a pronounced increase in the skill premium since 1980, which seems to slow down mildly towards the end of the 1990s. For comparison, the figure also shows a naive measure of the skill premium (the log wage difference between workers with at least a college degree and those with at most a high school degree) and the Mincerian return to schooling. The trend and fluctuations in our measure of the skill premium are similar to those in the Mincer return, indicating we have adequately controlled for heterogeneity beyond two skill types.[8]

Figure 2 shows similar plots for the relative employment and the relative supply of skilled labor. Again, there is a substantial difference between our preferred measure and the naive measure of the relative employment of skill. The increase in the employment and the supply of skill was roughly similar over the last two decades, but the higher frequency fluctuations differ markedly as we document below.

The other data series we use in our analysis are the following. Output is non-farm business output of all persons from the national income and product accounts (NIPA). Hours are total hours of non-supervisory workers from

---

[7]For the skill premium and relative employment series, we calculate average prices and quantities weighting individual workers in each cell by hours worked. For the relative supply series this is not possible since we do not observe hours worked for non-employed workers. For this series, we weight averages only by the CPS-ORG sample weights.

[8]Note that all our series exhibit large high frequency movements. These fluctuations are not seasonal effects but reflect measurement error. In the estimation, we smooth the impulse responses using the Minnesota prior.

the Current Employment Statistics establishment survey. Labor productivity is output per hour. All three series are available from the Bureau of Labor Statistics (BLS) productivity and cost program. As the relative price of investment goods, we use a quarterly intrapolation as in Fisher (2006) of the quality adjusted NIPA deflator for producer durable equipment over the consumption deflator (Gordon (1990); Cummins and Violante (2002)).[9]

Table 1 shows the business cycle correlations of the skill premium and the relative employment and supply of skill with output, hours, productivity and the relative price of investment goods.[10] The skill premium is basically acyclical: it is only very mildly positively correlated with output and even less correlated with hours worked. This finding is consistent with previous studies (Keane and Prasad (1993); Lindquist (2004)). The relative supply of skill is acyclical as well, but the relative employment of skill is higher in recessions than in booms, indicating the presence of a composition bias in employment as argued by Solon et al. (1994). The correlation of the skill premium with the relative investment-price is weak and negative. This is a first indication that capital-skill complementarity does not seem an important feature of the data at business cycle frequencies.

3 Results

In this section, we present our results for the effects of technology shocks on aggregate variables. We start by assessing the degree of skill bias in 'traditional' neutral technology or total factor productivity shocks. We then discuss how exogenous shocks to the supply of skill may bias these estimates and how we can control for these skill supply shocks. Next, in section 3.3, we propose a strategy for separating skill-biased technology shocks from skill-neutral shocks. In section 3.4, we address the issue of capital-skill complementarity and evaluate the hypothesis that it is investment-specific technological progress that produces the skill-bias observed in the data. Finally, in section 3.5, we jointly estimate all three types of technology shocks and evaluate their importance for business cycle fluctuations in various aggregates.

3.1 Skill bias in neutral technology shocks

Galí (1999) identifies permanent technology shocks as the only source of long-run movements in labor productivity. In a wide range of models, closed-

---

[9]We thank Jonas Fisher for making his data available to us. Fisher's data runs until 2000:4, which limits our estimation sample. Riccardo DiCecio has updated the quarterly time series for the relative price to 2004 (DiCecio (2005)). We will incorporate the extended data in the next draft of this paper.

[10]The sample used to generate these correlations coincides with the estimation sample used in the next section, i.e. 1979:1-2000:4.

economy, stationary, one-sector RBC models as well as models of the new Keynesian variety, shocks to total factor productivity are the only shocks that satisfy this identifying restriction. The remaining disturbances in the structural VAR are non-technology or 'demand' shocks, an amalgam of other possible shocks in the model: government expenditure shocks, preference shocks, or shocks to price or wage markups. As a first pass at our data, we evaluate the skill bias in technology shocks identified in this manner.

Figure 3 presents impulse response functions of a VAR as in Galí (1999), extended with the skill premium as a measure of skill bias in addition to labor productivity and hours worked, and estimated on our smaller sample. Here, as in all graphs that will follow, the point estimate is the median from the posterior distribution of the structural impulse-response coefficients. The dotted confidence intervals are one-standard error bands from the same distribution.

Introducing the price of skill as an additional regressor and using a different estimation sample leaves the responses of labor productivity and total hours worked almost unchanged compared to Galí (1999). As in his estimates, a positive innovation in technology leads to an almost immediate increase in labor productivity equal to the long run effect, and an initial reduction and a subsequent increase in hours worked. The first finding is supportive of the interpretation of the identified shock as a permanent improvement in technology. The second finding has typically been interpreted as evidence in favor of price rigidities, which dampen the substitution effect on impact and thus make the income effect of higher productivity that increases the demand for leisure dominant in the short run. Note that the skill premium increases in response to a permanent improvement in technology. The effect is permanent and is almost fully realized after two quarters. This finding is consistent with the hypothesis of skill-biased technological change, suggesting that the improved technology increased the demand for high-skilled labor.

When we include the wages and hours of high and low skilled workers separately, the wage of high skilled workers increases as expected, see Figure 4. The wage of low skilled workers stays roughly constant and initially even decreases a bit. Apparently, the skill bias in the technology shocks is so large that the relative price effect dominates the average price effect on the wage of low skilled workers. A different picture emerges for hours worked. Here, hours worked by high skilled workers decrease, while they increase for low skilled workers. This result is somewhat counter-intuitive, since we would have expected the relative quantity of skilled labor to increase. Since we have not properly identified skill-biased technology shocks here, this result could, however, obscure different kinds of disturbances such as different types of technology shocks or skill supply shocks.

The estimated technology shocks and their dynamics from the Galí (1999) VAR used here, are similar to the direct estimates of total-factor

productivity by Basu et al. (2006). As a robustness check, we use the quar-
terly series of the Basu et al. (2006) residuals, constructed by Fernald (2007),
instead of labor productivity in the VAR.[11] If the technology shocks iden-
tified by the two approaches were identical, then these impulse responses
should be the same as those shown in Figures 3 and 4. The results are
shown in Figures 5 and 6. Indeed, the responses of the 'purified' technology
measure, hours and the premium are very similar, providing support for
the identifying restriction used here. Interestingly, the increase in the wage
premium stems from a fall in the wage of low skilled workers rather than an
increase in high skilled wages, however.

3.2   Shocks to the supply of skill

In the identification of technology shocks used above, we assumed that tech-
nology shocks are the only shocks that drive productivity in the long run.
We showed that these shocks have asymmetric effects on the demand for
high and low skilled labor. Thus, production does not use the standard
technology as in (1), but either requires high and low skilled labor as sepa-
rate and imperfectly substitutable inputs, as in (2), or output is produced in
multiple sectors with different input shares of skilled labor. In these cases,
the identifying assumption of Galí is no longer valid because shocks to the
supply of skill may affect labor productivity in the long run.

Suppose a preference shock causes college enrollment to increase perma-
nently. When the new, larger cohort of college graduates enters the labor
market, the supply of skill exogenously increases. The resulting lower skill
premium leads firms to employ relatively more skilled workers. Since skilled
workers are more productive, this raises average labor productivity. Thus,
this shock to the supply of skill satisfies the identifying restriction for a
technology shock, even though technology has not changed at all.

We separately identify shocks to the supply of skill in order to avoid
biasing the estimated technology shocks. For this purpose, we include a
measure of the relative supply of skilled workers in our VAR. We use a
short-run restriction to identify shocks to the supply of skill: only skill
supply shocks affect the supply of skill within a quarter. This restriction is
equivalent to assuming that the supply of skill is predetermined.

Of course there are many other shocks that may increase the supply of
skill endogenously, through an increase in the skill premium. Skill-biased
technology shocks are just one example. However, the intuition for the iden-
tifying restriction is that in order to increase the supply of skill in response
to an increase in its price, workers need to obtain more education, which
lasts at least a year. It seems unlikely therefore, that other shocks would
affect the supply of skill within a quarter.

---

[11]We are grateful to Marty Eichenbaum and Luigi Paciello for drawing our attention to
these data and making them available to us.

It is crucial for our identification that we use a measure of the relative *supply* of skill, not the relative employment. It is reasonable to assume that the supply of skill is predetermined, but the same is not true for the employment of skill. If low and high skilled workers are imperfect substitutes, then firms may hire relatively more skilled workers in recessions, when the unemployment pool is larger and these workers are more abundantly available. This composition bias has been documented by Solon et al. (1994). We measure the relative supply of skill as the ratio of skilled workers to low skilled workers in the workforce, whereas the relative employment is the the equivalent ratio among employed workers, see section 2.3.

The strategy to identify technology shocks conditional on skill supply shocks is recursive. We first identify skill supply shocks with the short-run restriction and next use the same long run restriction discussed in the previous subsection to identify technology shocks. Thus, skill supply shocks are allowed to have a long run effect on productivity. Having identified fluctuations in productivity (as well as other variables in our VAR) that are due to skill supply shocks, technology shocks are the only *remaining* shocks that affect labor productivity in the long run. The details on the implementation of this combination of short and long run restrictions can be found in Appendix B.

Figure 7 shows the impulse response functions for this identification scheme. The lower row shows the responses to a one-standard deviation skill supply shock. By construction, the supply of skill increases immediately in response to this shocks. The estimates indicate that the effect is permanent: the supply of skill remains high in subsequent quarters. Somewhat counter-intuitively, labor productivity falls after a positive skill supply shock, hours jump up on impact and continue to increase and the skill premium is almost unaffected.

Controlling for skill supply shocks affects the impulse responses to technology shocks very little. The responses of productivity, hours and the skill premium are all very similar to the estimates without controlling for skill supply shocks. The response of productivity is a bit stronger and the response of the skill premium a bit weaker than before. The supply of skill falls moderately, but significantly, in response to a positive technology shock. We conclude that, while the direction of the bias is as expected, its size seems to be small. Nevertheless, we will control for shocks to the supply of skill in all specifications in the rest of the paper.

3.3   Skill-biased technology shocks

While the response of the skill premium is consistent with skill-biased technological change, it casts doubt on the traditional interpretation of these shocks. If these were truly shocks to total factor productivity, as in equation (1), the demand for skilled and unskilled labor should increase in equal

proportions and the relative demand should be unaffected. Here, we propose an alternative identification strategy to directly identify skill-biased technology shocks in addition to skill-neutral shocks to productivity.

In sections 3.1 and 3.2 above, we interpreted the increase in the skill premium in response to a technology shock as a measure of skill bias in technology. Here, we formalize that interpretation as an identifying restriction, identifying skill-biased technology shocks as those shocks that affect the relative price of skill in the long run, see equation (3). This restriction is similar in spirit to the identification of investment-specific technology shocks as shocks that affect the relative price of investment goods proposed by Fisher (2006). Controlling for shocks to the supply of skill is particularly important in this context, because of the standard simultaneity problem in estimation of demand or supply elasticities. An exogenous, permanent increase in the supply of skill would permanently reduce the price of skill and thus satisfies our identifying restriction for skill-biased technology shocks. We control for skill supply as described above in section 3.2.

Precisely, the identifying assumptions are now as follows. First, we identify skill supply shocks as the only shocks that affect the supply of skill contemporaneously. Next, we identify skill-biased technology shocks as the only remaining shocks that affect the relative price of skill in the long run. Both types of shocks could potentially affect labor productivity. Finally, skill-neutral technology shocks are all remaining shocks that affect labor productivity in the long run.

This identification scheme strictly speaking is not a *decomposition* of technology shocks as in Galí (1999) into skill-biased and skill-neutral shocks. In principle, there might be shocks that affect the skill premium but not labor productivity in the long run. However, as explained in section 2.1, it is hard to imagine non-technology shocks other than skill supply shocks to affect the skill premium in the long run. Moreover, our estimates indicate that the shocks we identify as skill-biased technology shocks increase labor productivity, supporting our interpretation of these shocks as a specific type of technology shock.

Figure 8 shows the responses of the skill premium, the supply of skill, labor productivity and total hours worked to a one-standard deviation skill-biased technology shock (SBT shock) and skill neutral technology shock. By assumption, a positive SBT shock drives the skill premium up in the long run. The estimates indicate that half of this effect is realized immediately and the rest within a year. A skill-neutral technology shock has no significant effect on the wage premium on impact and by assumption there is no long run effect either. SBT shocks increase the supply of skill in the long run, as should be expected with a higher skill premium, but this effect is small.

In response to a positive SBT shock, hours worked significantly and persistently fall. Interestingly, skill-neutral technology shocks barely decrease hours on impact and significantly and substantially increase hours worked

less than a year after impact. This finding suggests that at least part of the
fall in hours worked in response to technology shocks, as in Galí (1999) and
in the estimates in section 3.1, is related to the skill bias in these shocks.
If high skilled workers are much more productive than low skilled workers,
then it is possible that by substituting low skilled for high skilled workers in
response to an SBT shock, firms may increase effective labor input in their
production process, while reducing total hours or employment. Figure 9
confirms this interpretation: in response to an SBT shock, the wage of high
skilled workers increases substantially, but the wage of low skilled workers
actually falls. In contrast, the wages of both types of workers are affected
identically by a skill-neutral technology shock. These findings indicate that
for low skilled workers the *relative* productivity effect dominate the average
productivity effect of an SBT shock.

3.4   Investment-specific technology shocks and capital-skill complementarity

Over our sample period the relative price of investment goods fell substan-
tially. This finding has been interpreted to mean that technological progress
has been faster in investment goods producing sectors than in consumption
goods producing sectors (Greenwood et al. (1997), Cummins and Violante
(2002)). Fisher (2006) has argued that such investment-specific technologi-
cal change is important not only for long run trends, but also for business
cycle fluctuations. Because the increase in the skill premium roughly coin-
cided with the decrease in the relative price of investment goods, Krusell
et al. (2000) argue that investment-specific and skill-biased technological
change might be one and the same. They show that if capital and skill are
complements in the aggregate production function, technological innovation
in the investment-sector will necessarily lead to an increase in the demand
for skill (see section 2.1 and equation 4). If this is the case, then investment-
specific technology shocks should lead to business cycle fluctuations in the
skill premium.

  In this section, we follow Fisher (2006) in identifying investment-specific
and investment-neutral technology shocks using the relative price of invest-
ment goods. We estimate the effect of these shocks on the skill premium in
order to evaluate the hypothesis of capital-skill complementarity. As before,
we control for skill supply shocks, so that the exact identifying restrictions
are as follows. First, we identify skill supply shocks using a short run restric-
tion as described above. Then, we identify investment-specific technology
shocks as the only remaining shocks that affect the relative price of invest-
ment goods in the long run. Finally, investment-neutral technology shocks
are all remaining shocks that drive labor productivity in the long run.

  Figure 10 shows the responses of the the skill premium, the relative
supply of high skilled workers, labor productivity, hours worked and the

relative price of investment goods to investment-specific and investment-neutral technology shocks. Note that controlling for skill supply shocks changes the original results in Fisher (2006) very little. After an improvement in investment-specific technology, the relative price of investment falls, productivity increases and hours worked increase as well. An investment-neutral technology shock, has no effect on the relative price of investment, increases productivity and leads to a fall in hours worked.

The skill premium and the supply of skill significantly fall after an improvement in investment-specific technology. While there is certainly evidence for a relation between skill bias and investment-specific technical change, these estimates point towards capital-skill substitutability rather than complementarity: investment-specific shocks favor demand for unskilled labor rather than skilled labor. Because we have already documented that technology shocks are skill biased, it should not be surprising that investment-neutral technology shocks increase the skill premium, suggesting these shocks increase the demand for skilled labor.

The same finding can be documented in an alternative way. In Figure 11, we present impulse responses of the relative price of investment goods to skill-biased and skill-neutral technology shocks, identified as in section 3.3. The graphs provide the mirror image to those in Figure 10: skill-biased technology shocks increase the relative price of investment goods significantly, suggesting these shocks are 'consumption-specific' or capital and skill are substitutes in production.

Our findings are in striking contradiction with the argument in Krusell et al. (2000). What explains the difference is that Krusell et al. (2000) base their argument on a correlation in the long run trends in the skill premium and the relative price of investment goods. In our approach, the identifying variation are comovement between those two series at all frequencies *except* the trends, which are captured by the constant term in the VAR. It is possible that the comovement in the trends in both relative prices is a spurious correlation between two integrated series. It is also possible that the model needed to explain long run growth trends is different from the model that describes higher frequency fluctuations.[12] In any case, our findings reject the hypothesis that there is a stable aggregate production function with capital-skill complementarity.

---

[12]Lindquist (2004) presents a business cycle with capital-skill complementarity and investment-specific technology shocks and argues that the model can explain fluctuations in the skill premium and the capital-skill ratio. However, he evaluates the model based on the unconditional correlations of the skill premium with output and does not consider the correlation of the skill premium with the investment price.

3.5   Technology Shocks and Business Cycle Fluctuations

Our results suggest that there are at least four different types of technology shocks with distinct implications for the comovement of aggregate variables: un-skill-biased, investment-specific; skill-neutral, investment-specific; skill-biased, un-investment-specific; and skill-neutral, investment-neutral technology shocks. In this section, we assess the relative importance of these shocks for the business cycle fluctuations in output, hours worked and the skill premium and try to reconcile our findings with the unconditional correlations between these variables. We do this using a variance decomposition of the forecast error of a VAR.

With the identifying restrictions discussed above, it is not possible to separately identify all four different shocks simultaneously. Recall that both investment-specific and investment-neutral technology shocks affect the skill premium. Conversely, both skill-biased and skill-neutral technology shocks affect the relative price of investment goods. Hence, if we use a recursive identification scheme, identifying first investment-specific technology shocks, then these shocks will include the unskill-biased, investment-specific shocks. Skill-biased technology shocks will then be identified as all *remaining* shocks that affect the skill premium in the long run and will exclude shocks that affect both the relative price of investment and the skill premium.

Our solution to this problem is to estimate both orderings and use the estimates as a lower and upper bound for the contribution of the various shocks. To be more precise, we always identify supply shocks first as above. Then, in ordering I, we identify investment-specific technology shocks as all remaining shocks that affect the relative price of investment goods. These shocks are allowed to affect the skill premium. Skill-biased technology shocks are identified as all remaining shocks that affect the skill premium in the long run. The estimates of this VAR provide an upper bound for the contribution of investment-specific shocks and a lower bound for the contribution of skill-biased technology shocks. In ordering II, we identify skill-biased technology shocks as all shocks that affect the skill premium in the long run (conditional on skill supply shocks) and investment-specific shocks as the remaining shocks that affect the relative price in the long run. This ordering provides an upper bound for the contribution of skill-biased shocks and a lower bound for the contribution of investment-specific shocks. In both cases, the remaining shocks affecting labor productivity are neutral technology shocks.

Table 2 shows the variance decomposition of the forecast error in output, hours and the skill premium. The contribution of skill supply shocks and neutral technology shocks is very similar in both orderings of the identifying restrictions. This illustrates that we identify the same shocks in both orderings. Neutral technology shocks explain less than 5% of business cycle fluctuations in output and play virtually no role for fluctuations in hours

and the skill premium. Investment-specific technology shocks explain up to two thirds of the volatility in output at business cycle frequencies, about one third of the variation in hours. This finding is consistent with earlier findings in this literature (Fisher (2006), Canova et al. (2006)).

Skill-biased technology shocks explain almost all of the entire business cycle variation in the skill premium. These shocks are important for fluctuations in output and (especially) hours as well, but only insofar as they also affect the relative price of investment goods. Investment-specific, skill-neutral technological progress is important for fluctuations in output, but does not have much of an effect on the skill premium. These results suggest that shocks that drive fluctuations in the skill premium are largely unrelated to other variables in the economy. These is consistent with the unconditional moments in table 1, which show the skill premium to be largely uncorrelated with output. Note however, that while the supply of skill is acyclical, shocks to the supply of skill contribute substantially to business cycle fluctuations in hours worked.

4   Conclusions

[To be added]

A   Data description

[To be added]

B   Details on the specification

B.1   Long-run identification

As explained in section 2.2, structural identification involves finding a mapping from the residuals v_t of a reduced form VAR,

$$X_t = \sum_{j=1}^{p} D_j X_{t-j} + v_t$$

into structural residuals that can be interpreted as technology shocks. Here, $X_t = (X_{it}, ..., X_{nt})$ is a vector of n variables and p the number of lags included in the VAR. Let $\Sigma_v \equiv E[v_t v_t'] = \Omega$ denote the variance-covariance matrix of the reduced form residuals. Throughout, the identification will be exact, i.e. there exists a unique mapping from the reduced to the structural form of the VAR. The structural residuals e_t are assumed to be orthogonal and their variance is normalized such that $\Sigma_e \equiv E[e_t e_t'] = I$. Then, the relationship between the structural and reduced form residuals can be described by $e_t = A v_t$ which implies $A \Sigma_e A' = \Omega$, delivering $n(n + 1)/2$ of the necessary restrictions in order to pin down all n^2 elements of the matrix A.

The remaining $n(n - 1)/2$ assumptions stem from restrictions on the matrix of long-run effects and can be incorporated as zero restrictions in the matrix of long-run effects $C \equiv \sum_{i=0}^{\infty} \Phi_i A$. Here, the Φ_i are the impulse-response coefficients from the reduced form VAR, namely $\Phi_0 = I_n$ and $\Phi_s = \sum_{j=1}^{s} \Phi_{s-j} D_j$. One can now re-order the $n(n - 1)/2$ zero restrictions such that C has a lower-triangular structure. The matrix C, and consequently A, is then obtained by decomposing the variance of the ∞-step ahead forecast error with the Cholesky decomposition.[13] For this, note that the k-step ahead forecast error is equal to $\eta_{t,k} = X_{t+k} - E_t(X_{t+k})$ resulting in the following variance

$$MSE(k) = (\sum_{i=0}^{k} \Phi_i) \Omega (\sum_{i=0}^{k} \Phi_i)'$$

In the application, $k = \infty$ has to be approximated by some large value, here 20 years. Note that the above implies that the variables that are relevant for the identification are specified in first differences in the VAR.

The lower triangular structure of the matrix of long-run effects reflects the long-run restrictions of the various specifications discussed above if the

---

[13]This procedure is for example also used in Uhlig (2004).

variables in the VAR are ordered conveniently. Implementing the assumptions by Galí (1999), labor productivity is ordered first. In the identification of skill-biased technology shocks, the skill premium is ordered first, followed by labor productivity. The assumptions in Fisher (2006) result in the investment price ordered first in the VAR and labor productivity second. In the joint identification, ordering I, the investment price is followed by the skill premium and labor productivity. This results in the identification of investment-specific technology shocks, investment-neutral skill-biased technology shocks and investment-neutral skill-neutral technology shocks. In the joint identification, ordering II, the skill premium is followed by the investment price and productivity. Hence, we identify skill-biased technology shocks, skill-neutral investment-specific technology shocks and skill-neutral investment-neutral technology shocks. This procedure uniquely pins down the effects of the identified technology shocks on all variables in the VAR and the results are not affected by additional (superfluous) zero restrictions in the matrix of long-run effects.

B.2   Combination of short- and long-run restrictions

To implement the short-run restriction, which identifies skill supply shocks, together with the long-run restrictions for the various technology shocks, we need to exactly identify the transformation matrix A that maps reduced form into structural coefficients. Under exact identification, we can then proceed with the estimation of the reduced form VAR and the structural mapping as before. As explained above, A satisfies $AA' = \Omega$ and we then need another $n(n-1)/2$ restrictions for exact identification. Similar to before, we can formulate the problem in a triangular structure when the variables are conveniently ordered. This means ordering the supply of skill first in the VAR and then ordering the other variables according to the respective specification (see above).

   The identification then works as follows. First, one identifies the supply shock through its short-run effect. More precisely, in order to identify supply shocks we assume that neither i-shocks, nor SBT-shocks nor skill-neutral or investment-neutral technology shocks affect the supply of skill in the short run (on impact). This is equivalent to restricting $a_{12} = a_{13} = a_{14} = 0$ (with a_{ij} being elements of A). These zero restrictions in the first row of A, combined with

$$A_{1.} * A'_{.1} = \Omega_{1.}$$

pin down the first column of A. The first column uniquely determines the effects of the supply shocks on the system of variables.

   Second, we need to determine all other elements of the matrix A except for the first row and column. Here, we apply long-run restrictions in order to attribute the missing values to particular structural shocks. The remaining

lower right block of A, or the respective elements of the matrix of long-run effects $C_\infty = (\sum_{i=0}^{\infty})A$ excluding the first row and columns is lower triangular as in the standard long-run restrictions applied before. We obtain the elements in this block by applying the Cholesky decomposition to the 'updated' lower right forecast revision variance, i.e. for which the already known elements of A have to be taken into account.

B.3   Estimation of the reduced form VAR

We estimate the reduced form VAR for all specifications in a Bayesian framework with a Minnesota prior. The Minnesota prior consists of a normal prior for the VAR coefficients and a fixed and diagonal residual variance. The prior mean d_0 is restricted such that it represents a random walk structure on the VAR coefficients, i.e. in the standard case, the prior mean on the first lag is set to unity and the prior mean on the other lags (remaining parameters) is set to zero. Here, this is reflected by the fact that all variables enter the VAR in first differences resulting in a zero mean for all lags.

The prior variance $\Sigma_{d_0} = \Sigma_{d_0}(\phi)$ of the coefficients depends on three hyper-parameters ϕ_1, ϕ_2 and ϕ_3, that determine the tightness on own lags, other lags and exogenous variables. More precisely, for the coefficient γ of variable j with lag l in equation i:

$$
\begin{aligned}
Var(\gamma_{ijl}) &= \frac{\phi_1}{h(l)} \text{ for own lags} \\
&= \frac{\phi_1 \phi_2}{h(l)} \frac{\sigma_i^2}{\sigma_j^2} \text{ for lags on variable } j \neq i \\
&= \phi_1 \phi_3 \text{ for exogenous variables,}
\end{aligned}
$$

Here, $h(l) = l^d$ measures the harmonic decay on the lags and σ_i are elements from the residual variance-covariance matrix from the OLS regression. Except for the decay, a loose prior is chosen for the hyper-parameters, namely $\phi_1 = 0.2$, $\phi_2 = 0.5$ and $\phi_3 = 10^5$. The decay parameter $d = 3$. The advantage of the structure of the Minnesota prior is exactly this ability to separately deal with the lags of the variables, i.e. own and other lags, as well as exogenous variables. Together with a normal likelihood of the data the Minnesota prior produces a posterior that can be derived analytically. Hence, the estimation does not rely on sampling procedures.

References

Acemoglu, Daron, "Technical Change, Inequality and the Labor Market," *Journal of Economic Literature*, March 2002, *40* (1), 7–72.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger, "Computing Inequality: Have Computers Changed the Labor Market?," *The Quarterly Journal of Economics*, November 1998, *113* (4), 1169–1213.

_ , _ , **and Melissa S. Kearney**, "Rising Wage Inequality: The Role of Composition and Prices," Discussion Paper 2096, Harvard Institute of Economic Research 2005.

_ , _ , **and** _ , "Trends in the U.S. Wage Inequality: Re-Assessing the Revisionists," *Review of Economics and Statistics*, forthcoming 2008.

Basu, Susanto, John Fernald, and Miles Kimball, "Are Technology Improvements Contractionary?," *American Economic Review*, December 2006, *96* (5), 1418–1448.

Blanchard, Olivier and Danny Quah, "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *The American Economic Review*, September 1989, *79* (4), 655–673.

Canova, Fabio, David Lopez-Salido, and Claudio Michelacci, "On the robust effects of technology shocks on hours worked and output," Working Paper, 2006.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan, "A Critique of Structural VARs Using Business Cycle Theory," Working Paper, 2005.

Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson, "What Happens After a Technology Shock," *NBER Working Papers*, 2003, (10254).

Ciccone, Antonio and Giovanni Peri, "Identifying Human-Capital Externalities: Theory with Applications," *Review of Economic Studies*, 2006, *73*, 381–412.

Cummins, Jason G. and Giovanni Luca Violante, "Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics*, April 2002, *5* (2), 243–284.

DiCecio, Riccardo, "Comovement: It's Not A Puzzle," Working Paper 035B, Federal Reserve Bank of St. Louis, 2005.

Fernald, John, "A Quarterly, Utilization-Corrected Series on Total Factor Productivity," Working Paper, 2007.

Fisher, Jonas D.M, "The Dynamics Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy*, June 2006, *114* (3), 413–451.

Galí, Jordi, "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *The American Economic Review*, March 1999, *89* (1), 249–271.

_ **and Pau Rabanal**, "Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar U.S. Data?," *NBER Macroeconomics Annual*, 2004.

Gordon, Robert J., *The Measurement of Durable Goods Prices*, Chicago: University of Chicago Press, 1990.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, June 1997, *87* (3), 342–362.

Haefke, Christian, Thijs van Rens, and Marcus Sonntag, "Wage Rigidity and Job Creation," Working Paper, 2007.

Jaeger, David A., "Reconciling the Old and New Census Bureau Education Questions: Recommendations for Researchers," *Journal of Business and Economics Statistics*, July 1997, *15* (3), 300–309.

Katz, Lawrence F. and David H. Autor, *Changes in the Wage Structure and Earnings Inequality*, 1 ed., Vol. 3a of *Handbook of Labor Economics*, Amsterdam, North Holland: Orley Ashenfelter and David Card, June

_ **and Kevin M. Murphy**, "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, February 1992, *107* (1), 35–78.

Keane, Michael P. and Eswar S. Prasad, "The Relation Between Skill Levels and the Cyclical Variability of Employment, Hours and Wages," IMF Staff Papers 50 (3), 1993.

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante, "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, September 2000, *68* (5), 1029–1053.

Lindquist, Matthew J., "Capital-skill complementarity and inequality over the business cycle," *Review of Economic Dynamics*, 2004, *7*, 519–540.

Solon, Gary, Robert Barsky, and Jonathan A. Parker, "Measuring the Cyclicality of Real Wages: How Important is Composition Bias," *The Quarterly Journal of Economics*, February 1994, *109* (1), 1–25.

Teulings, Coen and Thijs van Rens, "Education, Growth and Income Inequality," *Review of Economics and and Statistics*, February 2008, *90* (1), 89–104.

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: Unconditional business cycle correlations*

	Std	Correlation with			
		Output	Hours	Productivity	Price
Baseline measure					
Skill premium	0.0077	0.1017	-0.0598	0.2874	-0.1486
Relative employment	0.0248	-0.3529	-0.2372	-0.2805	0.5123
Naive measure					
Skill premium	0.0086	0.0199	0.0788	-0.0898	0.0236
Relative employment	0.0232	-0.3153	-0.265	-0.165	0.4724
Relative supply	0.0114	0.0213	0.0759	-0.0824	0.2430

*Series are HP-filtered with λ=1600.

Table 2: Variance decomposition from joint identification

Horizon	8		16		32	
	I	II	I	II	I	II
output						
supply shock	5.3	5.9	10.0	10.8	12.3	13.1
invest. shock (ub,lb)	63.9	54.8	60.6	50.7	57.3	48.7
SBT shock (lb,ub)	2.5	9.1	1.9	9.7	1.9	8.9
neutral shock	4.2	4.9	4.3	5.0	4.6	5.2
hours						
supply shock	20.6	21.3	30.2	30.7	35.9	36.0
invest. shock (ub,lb)	46.0	26.6	38.8	22.1	31.8	18.7
SBT shock (lb,ub)	1.0	19.4	1.1	17.8	1.1	15.3
neutral shock	1.3	1.1	0.7	0.6	0.6	0.4
premium						
supply shock	1.7	1.5	2.0	1.8	2.4	2.2
invest. shock (ub,lb)	11.2	5.4	21.5	2.2	25.2	1.0
SBT shock (lb,ub)	86.0	92.2	76.0	95.6	72.2	96.6
neutral shock	0.4	0.2	0.1	0.1	0.1	0.0

Figure 1: Skill premium and Mincer return to schooling in the US
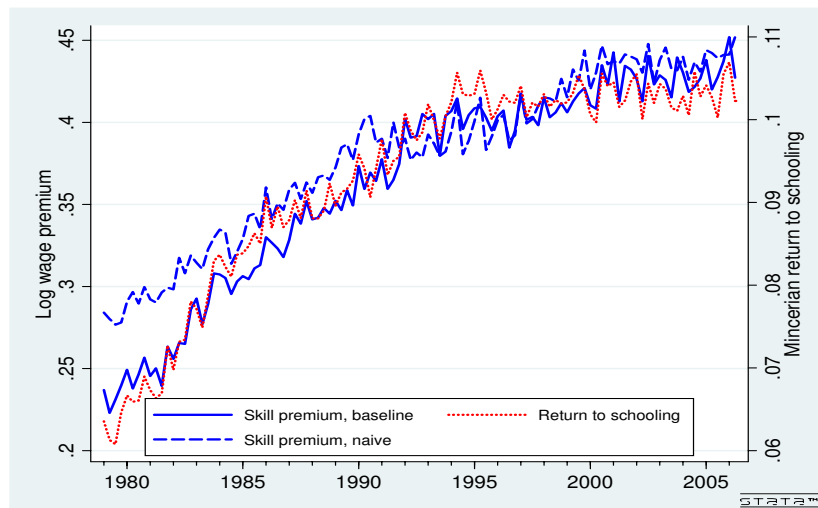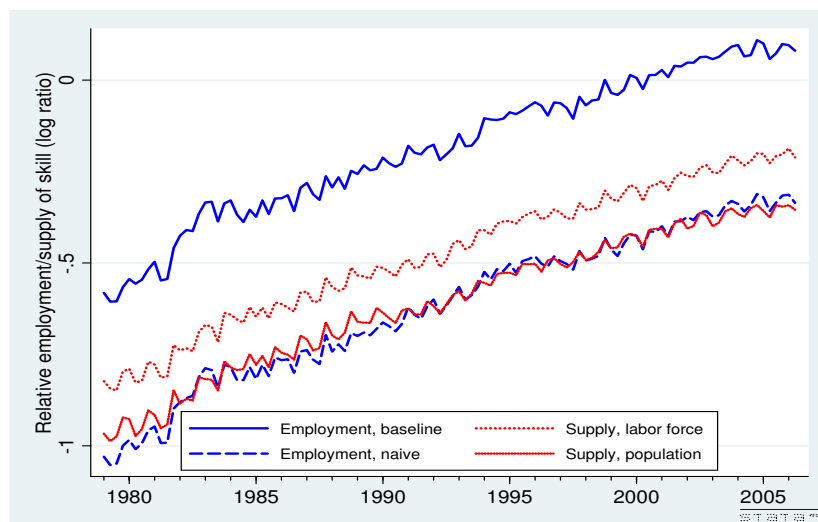

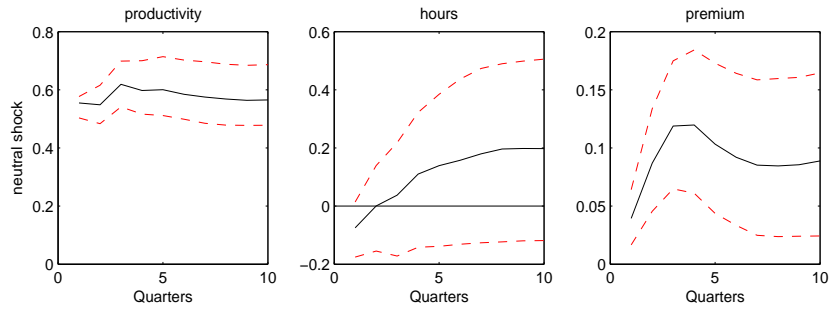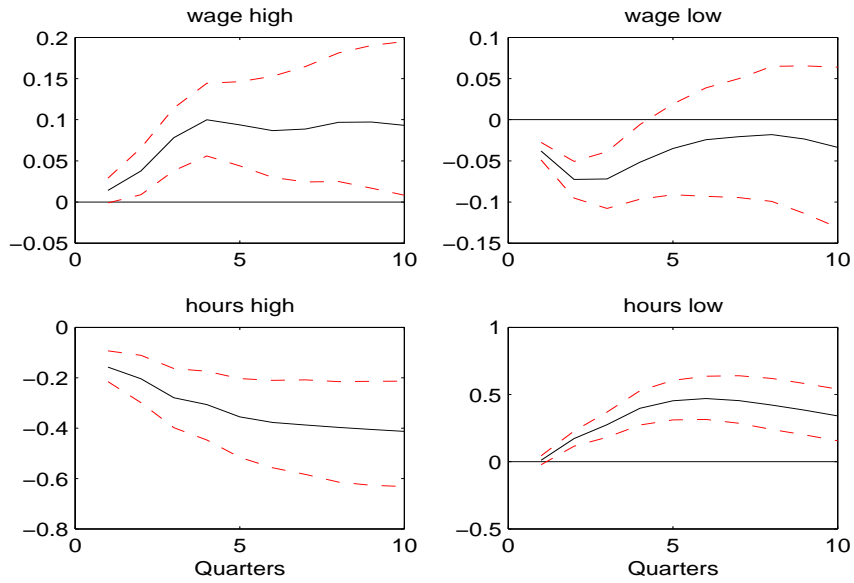Figure 2: Relative employment and relative supply of skill in the US
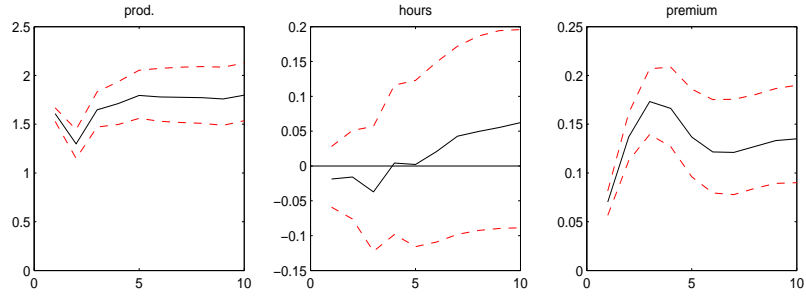
Figure 3: Galí identification with skill premium*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.
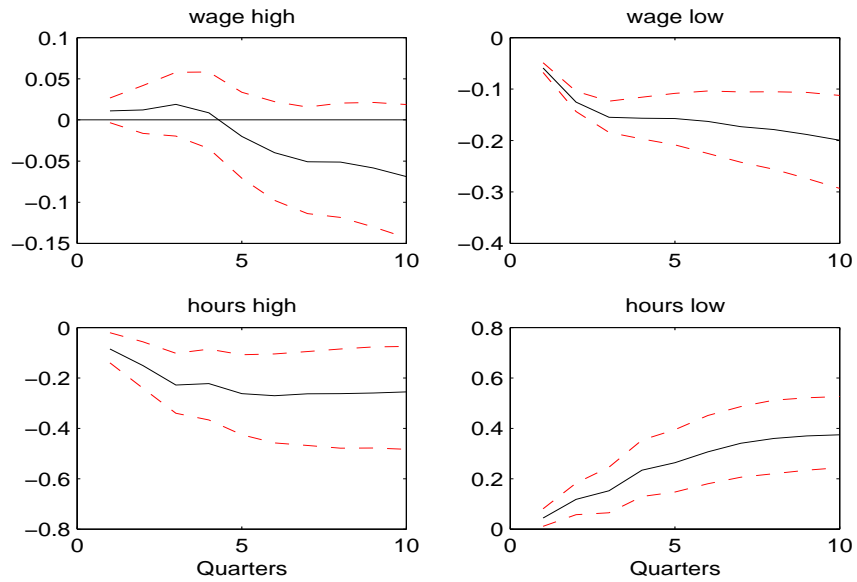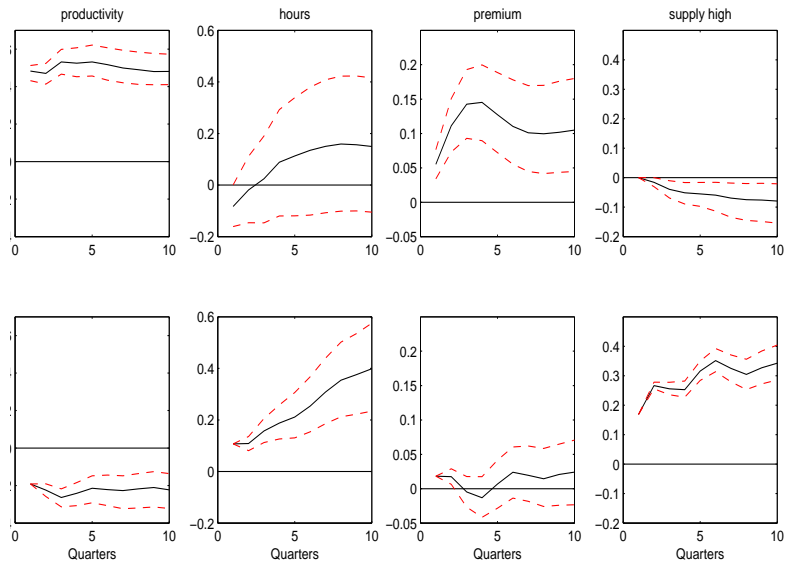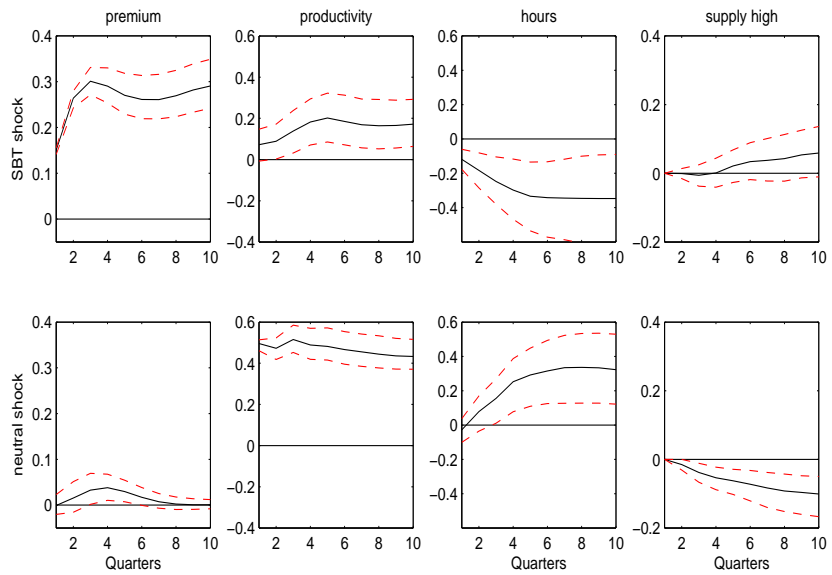
Figure 4: Galí identification - additional variables*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.

Figure 5: Galí identification with TFP measure*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.

Figure 6: Galí with TFP measure and additional variables*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.
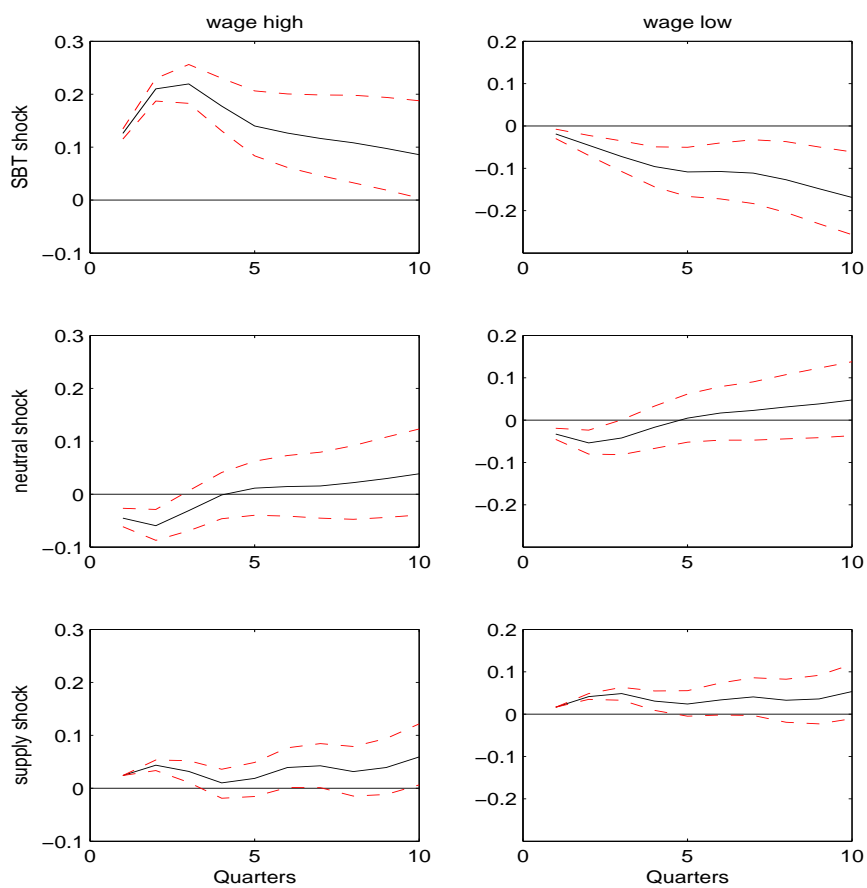
Figure 7: Galí identification with skill supply shocks*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.
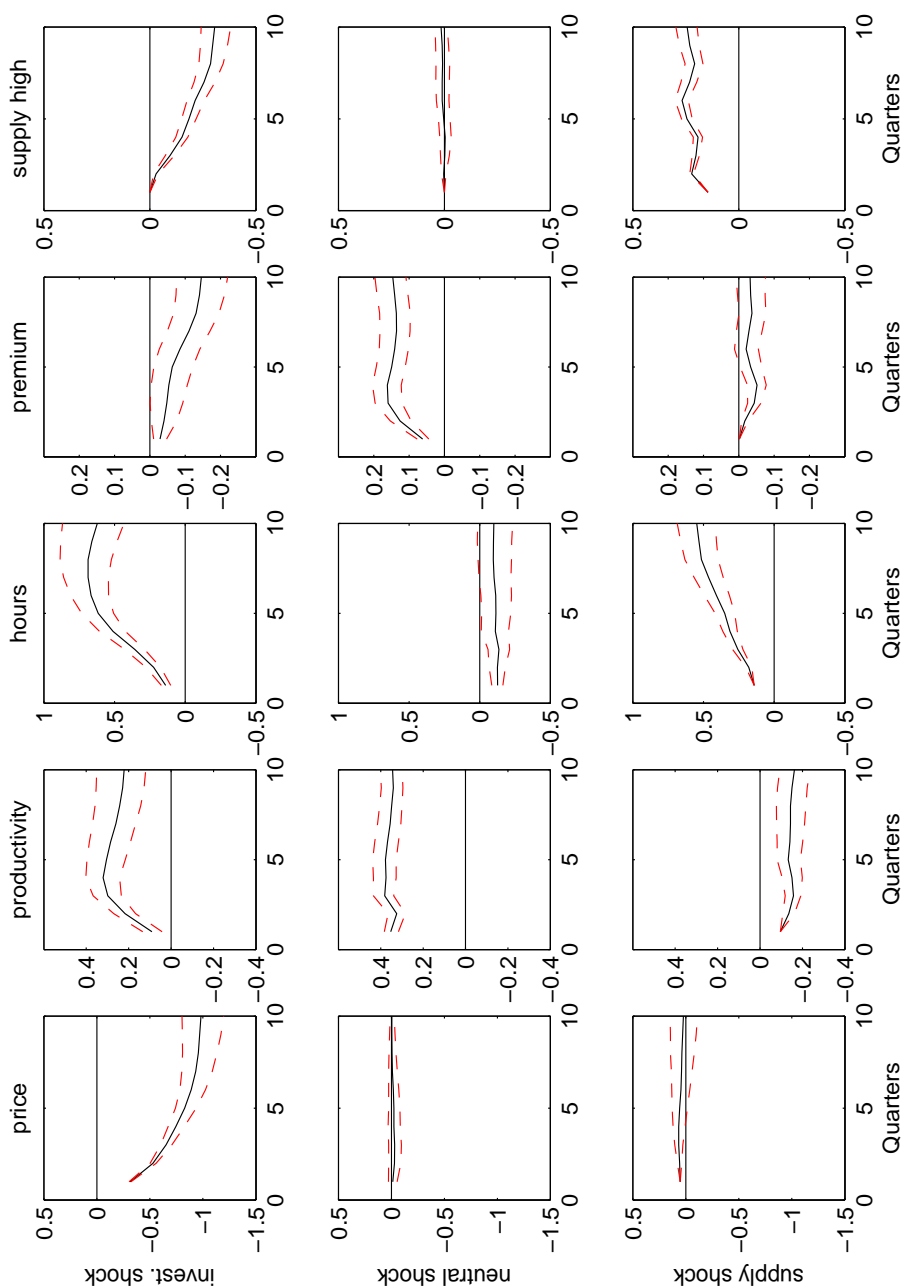
Figure 8: SBT identification*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.
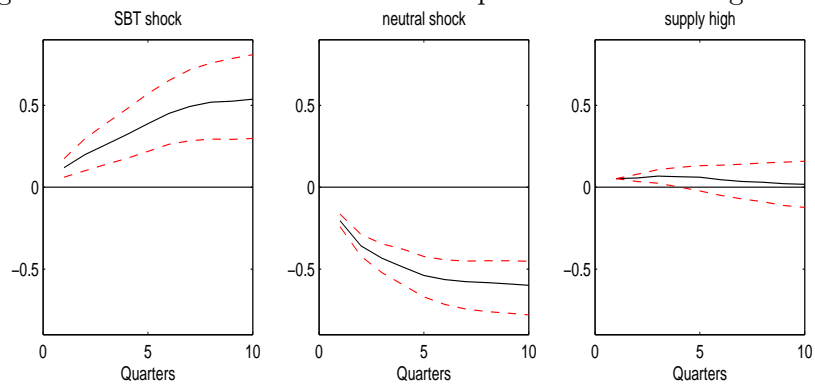
Figure 9: SBT identification - additional variables*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.

Figure 10: Fisher identification with skill supply shocks*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.

Figure 11: SBT identification - relative price of investment goods*

*Responses to a one-standard deviation shock. The dotted confidence intervals are one-standard error bands. All values are in percent.