

Cyclical Skill-Biased Technological Change*

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

Thijs van Rens
CREI, Universitat Pompeu Fabra, IZA and CEPR
tvanrens@crei.cat

June 2009

Abstract

Over the past two decades, technological progress has been biased towards making skilled labor more productive. What does skill-biased technological change imply for business cycles? To answer this question, we construct a quarterly series for the skill premium from the CPS and use it to identify skill-biased technology shocks in a VAR with long run restrictions. We find that hours worked fall in response to skill-biased, but not in response to skill-neutral improvements in technology. Skill-biased technology shocks are associated with increases in the relative price of investment, indicating that capital and skill are substitutes in aggregate production.

JEL codes: E24 E32 J24 J31

Keywords: skill-biased technology, skill premium, VAR, long-run restrictions, capital-skill complementarity, business cycle

*We thank David Autor, Toni Braun, Paula Bustos, Fabio Canova, Vasco Carvalho, Daniele Coen-Pirani, Juan-Carlos Conesa, Marty Eichenbaum, Jonas Fisher, Jordi Galí, Gino Gancia, Per Krusell, Lars Ljungqvist, Alberto Martin, Monika Merz, Kris Nimark, Franck Portier, Giorgio Primiceri, Valerie Ramey and Harald Uhlig for helpful suggestions, and Matthew Lindquist for making his code available to us. We also gratefully acknowledge financial support from the Spanish Ministry of Education and Science (grants Juan de la Cierva, SEJ2005-01124 and SEJ2006-02235); the Generalitat de Catalunya (DURSI grants Beatriu de Pinós and 2005SGR00490); and the Barcelona GSE Research Network.

1 Introduction

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 marked increase in the skill premium in the US and many other industrialized countries starting in the early 1980s, which coincided with a substantial rise in the average education level of the workforce. This parallel increase in the price and quantity of skill points towards an increase in the demand for skilled workers that exceeded the increase in their supply, suggesting that newly developed production technologies require relatively more educated and fewer uneducated workers (Katz and Murphy (1992); Autor et al. (1998); Acemoglu (2002); Autor et al. (2005) and Autor et al. (2008)).

What are the implications of skill-biased technological change for business cycle fluctuations? To our knowledge, this paper is the first to study this question. The lack of interest in skill-biased technology in the business cycle literature is surprising given the large number of studies dedicated to the effect of this type of technological progress on growth and inequality. Our results show that allowing for skill bias in technological change is important to understand business cycles and in particular speak to two important debates in the macroeconomics literature. First, traditional identifying restrictions, which are justified in models with homogeneous labor, may give a misleading picture of the effect of technology shocks on the economy. In particular, we show that skill-biased improvements in technology lead to a fall in total hours worked, but hours increase in response to skill-neutral technology shocks. Second, we show that the response of the economy to skill biased technology shocks implies restrictions on the production function that are of interest to macroeconomists studying growth as well as business cycles. In particular, we find that skill-biased technological change, unlike skill-neutral technological progress, increases the relative price of investment goods, indicating that capital and skill are substitutes in aggregate production.

Following previous studies on skill-biased technological progress, we identify skill-biased technology shocks from their effect on the skill premium. To this end, we construct a time series for the skill premium, which was so far not available at a quarterly frequency. Using the Current Population Sur-

vey (CPS) outgoing rotation groups, we 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 observable worker characteristics. In combination with comparable measures for the relative supply and employment of skilled workers, these series give a good picture of the high frequency movements in the price and quantity of skill in the US over the 1979:I-2006:II period.

We use a structural vector autoregression (VAR) to estimate the response of the economy to technology shocks, identifying technology shocks using long-run restrictions as in Blanchard and Quah (1989) and Galí (1999). We find that improvements in technology significantly increase the skill premium, providing strong evidence for skill bias in technological change at business cycle frequencies. This finding is novel and somewhat surprising, given that the skill premium is roughly acyclical over our sample period, which seems to suggest that skill-biased technological change is not relevant for business cycle fluctuations.¹ However, in the presence of multiple shocks, unconditional correlations are the result of a mixture of responses, which obscures the effects of changes in technology.² The structural VAR allows us to estimate the response of the economy conditional on technology shocks. This exercise delivers two sets of results.

For our first set of results, described in more detail in section 3, we propose a long-run restriction to separately identify skill-biased technology shocks. We argue that skill-biased technology shocks are the only shocks that affect the skill premium in the long run.³ Following Galí (1999), we identify skill-neutral technology shocks as all remaining shocks that permanently

¹This interpretation seems to be supported by the fact that the skill premium is negatively correlated with the relative supply of skilled labor at business cycle frequencies. For example, Acemoglu (2002) and Autor et al. (2005) argue this observation indicates that fluctuations in the skill premium are driven by fluctuations in the supply of skill rather than its demand.

²Lindquist (2004) reaches a similar conclusion, although from a completely different exercise. Lindquist argues that skill bias in technology shocks, generated by investment-specific technology shocks and capital-skill complementarity in the aggregate production function, explains the cyclical behavior of the skill premium. We discuss his argument in more detail in section 4.3.

³If there are exogenous, permanent changes in the supply of skilled labor, then this restriction is not valid, because increases in the supply of skill would also affect the skill premium in the long run. However, we separately identify skill supply shocks using a short run restriction, assuming that the supply of skilled workers is predetermined, and find that there are no exogenous changes in the supply of skill: skill supply shocks explain a negligible and insignificant fraction of fluctuations in all variables considered.

change labor productivity. We find that skill-biased technology shocks, like skill-neutral technology shocks, increase labor productivity. Skill-biased improvements in technology shocks also cause a large decline in total hours worked. Skill-neutral technology shocks lead to an increase, rather than a decrease in hours. This finding suggests that the fall in hours in response to technology shocks, which has been interpreted as evidence for price rigidities, may instead be driven by compositional changes in the labor force. This result is robust to the precise way to estimate the VAR. We also find very similar results if we construct skill-biased technological changes directly from data on the skill premium and the relative employment of skill using assumptions on the production function rather than a structural VAR.

Our second set of results, described in section 4, concerns the following question: What kind of changes in the aggregate production function best describe the skill-biased improvements in technology we observe over the past two decades? In a production function that takes capital, skilled and unskilled labor as inputs, a change in productivity must be either a change in total factor productivity (TFP) or capital or skilled labor augmenting technological change.⁴ Whereas changes in TFP are always skill-neutral, both capital and skilled labor augmenting technological change may increase the relative demand for skilled labor, depending on the elasticities of substitution between the different inputs. Krusell et al. (2000) argue that capital and skill are complements in the aggregate production function, and that skill-biased technological change is the result of an increase in the relative productivity of the investment-goods producing sector.⁵ We find strong evidence against this hypothesis and argue that over our sample period capital and skill are substitutes rather than complements.

In order to explore the issue of capital-skill substitutability, we include both the skill premium and the relative price of investment goods in the

⁴A change in the productivity of the third input, unskilled labor, cannot be separately identified. For example, a change in technology that makes unskilled labor more productive relative to capital and skilled labor would be the combination of an increase in total TFP and a decrease in capital and skilled labor augmenting productivity.

⁵It 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 sufficiently complementary, 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.

VAR. We use the latter to 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 productivity in the long run, are then 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. Using a simple two-sector real business cycle model that is consistent with our identifying restrictions, we explore what value of the elasticity of substitution between capital and high skilled labor corresponds to these estimates. For different values of the elasticity of substitution, we simulate data from the model and use those to estimate our structural VAR. In order to match the response of the skill premium to investment-specific shocks in the model-simulated data to the response estimated from actual data, we need to assume a very large degree of substitutability between capital and skill.

The remainder of this paper is organized as follows. Section 2 describes our empirical approach. We define the different shocks to the production technology that we consider and 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 present some descriptive statistics on the cyclicity of our quarterly series for the skill premium and the relative supply and employment of skill. In section 3 we describe the properties of skill biased technology shocks using the structural VAR analysis as well as a decomposition using the production function. Section 4 discusses our evidence for capital-skill substitutability in aggregate production. Section 5 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 technological change, discussing various specifications for the aggregate production function. Next, we explain how to identify these

different technology shocks from the data using either the functional form of the production function or 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. The production function 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 are neutral technology shocks, in the sense that they affect the productivity of all inputs in the same proportion. To allow for skill-biased technological change, 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)).

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

Here, A_t is total factor productivity 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 $\sigma > 0$. 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 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 .

There are two ways to interpret skill-biased technology shocks to an aggregate production function as in (1). 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,⁶ a sector-specific technology shock to a skill-intensive sector would 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 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,

$$Y_t = C_t + p_t I_t \quad (2)$$

where the relative price of investment goods p_t reflects technological improvements in the investment goods producing sector. An aggregate production function that allows for capital-skill complementarity is a slightly generalized version of (1).

$$Y_t = A_t \left[\beta \left(\gamma K_t^{\frac{\rho-1}{\rho}} + (1-\gamma) (B_t H_t)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1} \frac{\sigma-1}{\sigma}} + (1-\beta) L_t^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (3)$$

where σ is the elasticity of substitution between skilled and unskilled labor as before, which now 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

⁶This is the case where $\sigma_i = 1$ for all i . In the limit for $\sigma \rightarrow 1$, production function (1) becomes Cobb-Douglas, so that changes in B_t are indistinguishable from changes in A_t .

skilled labor ρ is lower than the elasticity of substitution between capital and unskilled labor σ , i.e. if there is capital-skill complementarity in production.

2.2 Identification and estimation

Under the assumption that workers' wages are proportional to their marginal product, we can calculate the skill premium directly from the production function. Using aggregate production function (1), we get the following expression,

$$\log \left(\frac{w_{H,t}}{w_{L,t}} \right) = \log \left(\frac{\beta}{1-\beta} \right) - \frac{1}{\sigma} \log \left(\frac{H_t}{L_t} \right) + \frac{\sigma-1}{\sigma} \log B_t \quad (4)$$

where $w_{H,t}$ and $w_{L,t}$ are the wages of high and low skilled workers respectively. This equation can be interpreted as a demand curve for skill. The skill premium is decreasing in the relative demand for high skilled workers, $\log(H_t/L_t)$, where the elasticity of demand depends on the elasticity of substitution between high and low skilled workers.

Changes in skill-biased technology B_t represent shifts of the skill demand curve or skill demand shocks. Since the skill premium and the relative quantity of skill are observable, these shocks can be directly retrieved from equation (4), using an estimate for the elasticity of substitution between low and high skilled workers σ .⁷ The estimates for the skill-biased technology shocks obtained this way are identified from the assumption that wages are proportional to marginal products. A sufficient condition for this assumption is that labor markets are perfectly competitive, in which case the wage of all workers equals their marginal product. If there are frictions in the labor market, the weaker assumption that wages are proportional to marginal products still holds approximately. However, if there are frictions in the wage determination process, then wages may deviate from marginal products in the short run. Therefore, we alternatively identify technology shocks using a structural VAR with long-run restrictions, as suggested by Blanchard and Quah (1989) and first used to estimate technology shocks by Galí (1999).

Consistent with equation (4), we identify skill-biased technology shocks as the only shocks that affect the skill premium in the long run, conditional

⁷An estimate for the share parameter β is unnecessary since this parameter affects only the level of B_t and we normalize the mean and variance of the shocks to zero and one respectively.

on the supply of skill. Since the identifying restriction is an assumption on the long-run effects of the structural shocks on the variables in the VAR, it is a weaker assumption than assuming that (4) holds in each period and makes the estimates robust to, for example, wage rigidities. In addition, the long run identification does not depend on the exact functional form of the production function and we no longer need to use an estimate for σ .⁸ Thus, we use long run restrictions in all our estimates, although we compare the results to a direct decomposition using equation (4), see section 3.3, and find that for the simplest estimates the differences are not large.

The estimation of structural shocks using long run restrictions is implemented in two steps. First, we estimate a reduced form VAR in the variables labor productivity, hours worked, the skill premium and in some specifications also the relative price of investment goods. Second, we map the reduced form coefficients and residuals into structural coefficients and shocks normalizing the variance of all structural shocks to one and assuming orthogonality between these shocks, as well as an identifying restriction. The long-run identifying restrictions are incorporated using a Cholesky decomposition of the infinite horizon forecast error variance.⁹

The specific restriction depends on the type of shock we are interested in estimating. 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 discuss the specific identifying restrictions used to identify neutral, skill-biased and investment-specific technology shocks as we describe our results in section 3. The identification of different types of shocks using the Cholesky decomposition is then implemented by simply reordering the variables in the VAR.

Our baseline VAR is estimated on quarterly data from 1979:I to 2000:IV. This period is relatively short because of data limitations, see section 2.3.

⁸Of course the assumption is not valid for all production functions. For example, with capital-skill complementarity, as in (3), any shocks that affect the capital stock also affect the skill premium in the long run. However, the restriction can easily be modified to incorporate this case, see section 4.

⁹The procedure employed here is very similar to the one in Uhlig (2004). We approximate the infinite horizon with 20 years.

All variables are used in first differences in order to allow for unit roots.¹⁰ The reduced form is estimated as a Bayesian VAR with a Minnesota prior, similar to Canova et al. (forthcoming). This prior reflects the belief that the true data generating process is a univariate unit root in each variable. It is implemented as a joint prior that the coefficient matrix of the first lag in the VAR is close to the identity matrix and the coefficients on further lags are close to zero, where the strength of the prior increases with the lag order. We use this prior for two reasons. First, in theory one should employ a VAR with an infinite number of lags in order to correctly identify technology shocks using long run restrictions, see e.g. Chari et al. (2008). The Minnesota prior allows us to generate sensible results for a large number of lags, here we use 8 lags.¹¹ 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 (see Table 4).

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 2006:2.¹² 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

¹⁰In 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. (forthcoming) 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.

¹¹The strength of the Minnesota prior increases with lag length to reflect the belief that the higher order lags are less likely to matter. This is reflected in a harmonic decay of the prior variance on the lag coefficients. We choose a decay parameter of 3. The remaining hyper-parameters are chosen as in the RATS manual such that the Minnesota prior is quite loose.

¹²The 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.

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. 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.

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 using average quantity weights and averaging quantities using average price weights.¹³ The resulting series are adjusted for seasonality using the X-12-ARIMA algorithm of the Census Bureau.

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

¹³For 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.

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.

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. 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 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)).¹⁴

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.¹⁵ The skill premium is basically

¹⁴We thank Jonas Fisher for making his data available to us. The quarterly relative price data runs until 2000, which limits our estimation sample.

¹⁵The sample used to generate these correlations coincides with the estimation sample

acyclical: it is only very mildly positively correlated with output and even less 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. Note that the correlations of the naive measure of the skill premium are quite different than the ones of the baseline measure. Accounting for heterogeneity is important for the cyclical behavior of the skill premium.

3 Skill-biased technology shocks

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’ shocks to total factor productivity. We then assess to what extent these estimates are biased due to the presence of shocks to the supply of skill and find that this bias is negligible. Next, in section 3.3, we propose an identification strategy to separate skill-biased from skill-neutral technology shocks and estimate the response of the economy to each type of shock.

3.1 Skill bias in technology

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

The first row of Figure 3 presents impulse response functions of a VAR as in Galí (1999), extended with the skill premium as a measure of skill bias used in the next section, i.e. 1979:1-2000:4.

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, and the dotted confidence intervals are 68% Bayesian confidence bands of the posterior distribution of the structural impulse-response coefficients. 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. 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.

The estimated technology shocks and their dynamics from the Galí (1999) VAR used here are similar to direct estimates of total-factor productivity from a production function composition in the spirit of a Solow residual by Basu et al. (2006). As a robustness check, we use the quarterly series of the Basu et al. (2006) residuals, constructed by Fernald (2007), instead of labor productivity in the VAR.¹⁶ If the technology shocks identified by the two approaches were identical, then these impulse responses should be the same as those shown in the first row of Figure 3. The results are shown in row two of Figure 3. Indeed, the responses of the ‘purified’ technology measure, hours and the premium are very similar, providing support for the identifying restriction used here.

¹⁶We are grateful to Marty Eichenbaum and Luigi Paciello for drawing our attention to these data and making them available to us.

3.2 Shocks to the supply of skill

In the identification of technology shocks used above, we assumed that technology 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 a standard Cobb-Douglas technology, but either requires high and low skilled labor as separate and imperfectly substitutable inputs, as in equation (1), or output to be produced in multiple sectors with different input shares of skilled labor. In these cases, the identifying assumption of Galí is no longer valid in the presence of exogenous changes in the supply of skill, because such changes may affect labor productivity in the long run.

Suppose a preference shock causes college enrollment to increase permanently. 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.

To assess the importance of this bias, we separately identify shocks to the supply of skill and compare the results from the estimation with and without skill supply shocks. For this purpose, we include a measure of the relative supply of skilled workers in the VAR. We use a short-run restriction to identify shocks to the supply of skill: only skill supply shocks affect the supply of skill in the short run. 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 identifying 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 after a quarter.

There is substantial measurement error in our time series for the relative supply of skill because each observation is based on a relatively small cross-section of individual workers. By construction, this measurement error is independently distributed over time, because the same individual is never in

the outgoing rotation group in two subsequent quarters. In order to prevent that measurement error is identified as a skill supply shock, we implement the short-run restriction after one quarter rather than on the impact effect.

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 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 A.

The third row of Figure 3 documents that controlling for skill supply changes the impulse responses to identified technology shocks very little. The responses of productivity, hours and the skill premium are all very similar to those in the baseline specification in the first row of the Figure. In addition, Table 2 compares the variance decomposition for the identification with and without supply shocks. Supply shocks matter very little for business cycle fluctuations in output, hours and even the skill premium. Moreover, controlling for skill supply shocks does not significantly alter the importance of technology shocks for fluctuations in these three variables. We conclude that the size of the bias induced by supply shocks is small. Since this is true for all of our specifications, we do not report the results controlling for supply shocks in the remainder of the paper.

3.3 Identified skill-biased technology shocks

Our evidence for skill bias in technological change at business cycle frequencies casts doubt on the traditional interpretation of identified technology 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 (4). 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). Precisely, the identifying assumptions are now as follows. We identify skill-biased technology shocks as the only shocks that affect the relative price of skill in the long run. These shocks may or may not affect labor productivity. Skill-neutral technology shocks are all remaining shocks that affect labor productivity in the long run. We implement these assumptions by ordering the respective variables subsequently in the VAR.

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 4 shows the responses of the skill premium, labor productivity and total hours worked to a one-standard deviation skill-biased technology (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 this effect is realized immediately on impact. 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.

In response to a positive SBT shock, hours worked significantly and persistently fall. 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 5 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 dominates the average productivity effect of an SBT shock.

Table 3 shows a decomposition of the forecast error variance of the VAR at business cycle frequencies with periodicities from 8 to 32 quarters. Separating out skill-biased and skill-neutral technology shocks, substantially increases the contribution of technology shocks to fluctuations. Skill-neutral technology shocks alone, explain about 40% of the business cycle variance of output, 5%-points more than the estimated overall contribution of technology shocks in the specification of Table 2. Skill-neutral shocks explain about 8% of the volatility in hours worked, compared to about 6% in Table 2. Skill-biased technology shocks explain an additional 4% of fluctuations in output and over 12% of fluctuations in total hours worked. Finally, fluctuations in the skill premium are almost exclusively due to SBT shocks.

3.4 Robustness

We now explore the robustness of our estimates to changes in the estimation specification and the construction of the data. The results of this exercise are summarized in Table 4. First, we check whether the finding that skill supply shocks are not important carries over to our specification with identified SBT shocks. In this specification, these shocks potentially matter more because of the standard simultaneity problem in estimating demand and

supply equations. 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 find however, that controlling for shocks to skill supply as described in section 3.2, does not significantly affect our estimates.

In our baseline estimates, we impose a Minnesota (Litterman) prior on the decay of the lag coefficients in order to be able to allow for a large number of lags. However, our results are not driven by this prior. The responses of productivity and the skill premium to all shocks are virtually unaltered when we change the number of lags, the strength of the prior, or when we estimate the VAR using ordinary least squares (OLS). The fall in hours worked in response to skill-biased technology shocks is also robust across specifications and is significant if we include at least 4 lags in the VAR. The increase in hours in response to neutral technology shocks is actually stronger in all alternative specifications: whereas in the baseline the positive effect becomes significant only after 3 quarters, in all other specifications it is significant at all horizons.

Next we explore to what extent the way we constructed our measure for the skill premium matters for the results. Using a ‘naive’ measure of the skill premium that does not take into account the heterogeneity over and above two skill types, we would not have observed the fall in hours in response to an SBT shock. As we found in the unconditional moments in section 2.3, accounting for heterogeneity is important for the cyclical behavior of the skill premium.

3.5 Production function decomposition

As a final robustness check, we compare the properties of our identified SBT shocks to a simple decomposition using the production function, see section 2.2. This decomposition is similar in spirit to a Solow residual and requires a value for the elasticity of substitution between high and low skilled workers σ . We use $\sigma = 1.5$, which is the consensus estimate from the literature based on several different data sources (Katz and Murphy (1992), Ciccone and Peri (2006), Teulings and van Rens (2008)). With this value, we can use equation (4) to retrieve changes in skill-biased technology B_t from our data on the skill premium and the relative employment of skill. After demeaning these changes and normalizing their variance to unity, they are comparable to

the identified SBT shocks from the structural VAR. The difference is in the identification. Whereas the identified shocks require assumptions only on the long run behavior of the skill premium, the decomposition requires equation (4) to hold in each period. Figure 6 plots both estimates for the shocks over the sample period. It is encouraging that despite the differences in identification, the resulting estimates for the skill-biased technology shocks look similar. The correlation between the two estimates is 0.53.

To complete the comparison, we compare the response of productivity, hours worked and the skill premium to the identified SBT shocks and the estimated shocks using the decomposition. We regress these variables on lags of the shocks, estimated either from the decomposition using equation (4) or as the residuals from our structural VAR, as suggested by Basu et al. (2006). This is a direct estimate of the moving average representation of the impulse response functions and the results are comparable to the impulse responses in Figure 4. Since the impulse responses in Figure 4 seem to flatten out after about 6 quarters, we use 6 lags of the shocks. The results are presented in the first row of Figure 7. The responses to identified SBT shocks estimated in this way are very similar to those directly calculated from the VAR estimates. We now discuss how the responses to SBT shocks obtained from the decomposition compare to these. The second row of Figure 7 shows the responses to SBT shocks estimated using the decomposition. Generally, the responses are very similar to the responses to the VAR residuals and never significantly different.

4 Capital-skill substitutability

Over our sample period the relative price of investment goods fell substantially. 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 technological change is important not only for long run trends, but also for business cycle fluctuations. Because the increase in the skill premium roughly coincided 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. 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. If this is the case, then investment-specific technology shocks should lead to business cycle fluctuations in the skill premium. In this section, we explore this hypothesis and find evidence against it.

4.1 Skill bias in investment-specific shocks

Consider the alternative aggregate production function (3), as in Krusell et al. (2000), which allows for complementarity or substitutability between capital and skill. Assuming as before that wages are proportional to marginal products in the long run, expression (4) for the skill premium changes to the following.

$$\begin{aligned} \log\left(\frac{w_{H,t}}{w_{L,t}}\right) &= \log\left(\frac{\beta(1-\gamma)}{1-\beta}\right) - \frac{1}{\rho}\log\left(\frac{H_t}{L_t}\right) + \frac{\rho-1}{\rho}B_t \\ &\quad + \frac{\sigma-\rho}{\sigma(\rho-1)}\log\left(\gamma K_t^{\frac{\rho-1}{\rho}} + (1-\gamma)(B_t H_t)^{\frac{\rho-1}{\rho}}\right) \end{aligned} \quad (5)$$

Since investment-specific technological progress raises the long run capital-labor ratio, it is clear that such technological change will also raise the skill premium if $\rho < \sigma$, i.e. if capital and skill are complements rather than substitutes in production. As a result, our identifying restriction that skill-biased technology shocks are the only shocks that affect the skill premium in the long run is no longer valid, and we need to separately control for investment-specific shocks. In addition, it is interesting in itself to assess the skill bias in investment-specific shocks, because it will allow us to assess the degree of capital-skill complementarity in aggregate production.

We follow Fisher (2006) in identifying investment-specific and investment-neutral technology shocks using the relative price of investment goods. We estimate the effect of these shocks on the skill premium in order to evaluate the hypothesis of capital-skill complementarity. We identify investment-specific technology shocks as the only shocks that affect the relative price of investment goods in the long run. Finally, investment-neutral technology shocks are all remaining shocks that drive labor productivity in the long run. For implementation, the relative price of investment and labor productivity are ordered first in the VAR.

Figure 8 shows the responses of the the skill premium, labor productiv-

ity, hours worked and the relative price of investment goods to investment-specific and investment-neutral technology shocks.¹⁷ 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.¹⁸

The skill premium significantly falls 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 increase relative demand for unskilled 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 9, 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 8: 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.¹⁹ In any case, our findings reject

¹⁷Note that again controlling for skill supply shocks changes the results very little.

¹⁸Since productivity increases after an investment-specific technology shock in our specification, we do not need to use an additional assumption on this effect as in Fisher (2006).

¹⁹Lindquist (2004) presents a business cycle with capital-skill complementarity and investment-specific technology shocks and argues that the model can explain fluctuations

the hypothesis that there is a stable aggregate production function with capital-skill complementarity.

4.2 Contribution to business cycle fluctuations

When we allow for investment-specific technology shocks, our estimates replicate the finding in Fisher (2006) that investment-specific shocks are an important source of business cycle fluctuations, whereas investment-neutral technology shocks contribute only a small fraction of fluctuations in output and hours, see Table 5. However, our results suggest that there are at least four different types of technology shocks with distinct implications for the comovement of aggregate variables: skill-neutral, investment-neutral; skill-neutral, investment-specific; skill-biased, investment-neutral; and unskill-biased, investment-specific (or skill-biased, consumption-specific) technology shocks. 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. In this case, skill-biased technology shocks will 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. Similarly, if we identify first the skill-biased shocks, then these shocks will include the skill-biased, consumption-specific shocks.

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, in ordering I, we identify investment-specific technology shocks as all 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

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.

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 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 6 shows the variance decomposition of the forecast error variance in output, hours and the skill premium. The contribution of skill- and investment-neutral technology shocks is very similar in both orderings of the identifying restrictions. Neutral technology shocks explain between 5% and 6% 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, up to 40% of the variation in hours. This finding is consistent with earlier findings in the literature (Fisher (2006), Canova et al. (forthcoming)).

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. Comparing the variance decomposition in Table 6 to that in Table 5, the contribution of investment-neutral technology shocks to fluctuations in output and hours increases, mirroring our results in section 3.3. 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. This finding is consistent with the unconditional moments in Table 1, which show the skill premium to be largely uncorrelated with output.

4.3 Replicating the data with a model

Our finding that the skill premium decreases in response to investment-specific shocks, and the relative price of investment goods increases in response to skill-biased technology shocks suggest that capital and skill are substitutes rather than complements in the aggregate production function.

Yet the estimates by themselves do not give any indication as to how large this effect is. What parameters of production function (3) are consistent with our estimates? To answer this question, we simulate a simple business cycle model with a production function as in (3) and compare the estimated impulse response functions from the actual data to those from simulated data for different values of the substitution parameters. This procedure also allows us to see whether the structural VAR performs well in capturing the conditional moments of the variables in a model that is consistent with our interpretation of the results.

The model is a simple real business cycle model with high and low skilled workers. The model is taken from Lindquist (2004) and combines the two sector model of Greenwood et al. (1997), in which output can be used for consumption or accumulation of capital equipment, with the model of Krusell et al. (2000) with two skill types and capital-skill complementarity. Business cycle fluctuations in the model are driven by shocks to total factor productivity and the relative price of investment goods.

For the calibration of the structural parameters of the model we also follow Lindquist (2004), but we assume that the two productivity shocks are highly persistent and uncorrelated with each other in order to be consistent with the identifying restrictions of our VAR. The substitution parameters in the aggregate production function (3) are $\sigma = 1.67$ and $\rho = 0.67$. These values were estimated by Krusell et al. (2000) to be consistent with the trends in the relative price of investment goods and the skill premium. Since $\rho < \sigma$ in this calibration the aggregate production function exhibits capital-skill complementarity. In alternative calibrations, we keep σ constant, because the value of the elasticity of substitution between high and low skilled workers is well documented, and change ρ to vary the degree of capital skill complementarity. We consider the cases of capital-skill complementarity ($\rho = 0.67$), weak complementarity ($\rho = 1.17$), neither complementarity nor substitutability ($\rho = \sigma = 1.67$), weak substitutability ($\rho = 2.17$), substitutability ($\rho = 2.67$), strong substitutability ($\rho = 3.17$) and very strong substitutability ($\rho = 5$). In each case, we recalibrate the other model parameters to keep the calibration targets constant.

We simulate the model 1000 times for 88 quarters, the same sample length as in our data. In each simulation, the model is first simulated for 200 periods, which are then discarded, in order to remove dependence on

the initial conditions. We then estimate the VAR for each sample of 88 quarters and average the impulse responses across the 1000 simulations. Figure 10 illustrates this for the calibration in which capital and skill are neither complements nor substitutes. For better comparison, the responses are normalized such that they match the responses in the actual data of the investment price and labor productivity to the two technology shocks respectively 10 quarters after the shock has hit. The estimated responses from the simulated data closely match the theoretical ones from the model. This is also the case for other degrees of substitutability of complementarity between capital and high-skilled labor. Most importantly for our purposes, the estimated response of the skill premium to investment-specific shocks is positive if capital and skill are complements, negative if they are substitutes and zero when they are neither substitutes nor complements.

Figure 11 shows the impulse responses of the skill premium to an investment-specific shock from the model simulated for different degrees of capital-skill complementarity/substitutability as well as from the actual data. Comparing the response of the skill premium to investment-specific shocks in the actual data to that in the model, it is clear that our estimates suggest a very large degree of capital-skill substitutability. In fact, the estimates suggest an elasticity of substitution between capital and high skilled labor of around $\rho = 5$, compared to the elasticity of substitution between capital and low skilled labor of $\sigma = 1.67$. These parameters imply that if the capital stock increases by 1%, firms can still produce the same amount of output as before if they fire 1.67% of their low skilled workers or as much as up to 5% of their high skilled workers.

5 Conclusion

In this paper, we explored the implications of skill bias in technological changes for business cycle fluctuations. We constructed a quarterly time series for the skill premium using micro-data from the Current Population Survey (CPS) outgoing rotation groups, and used it to identify skill-biased technology shocks in a structural VAR with long run restrictions. We showed that technology shocks are biased towards skilled labor at all frequencies and documented two main differences between skill-biased and skill-neutral technology shocks. First, the fall in hours in response to improvements

in productivity is driven by skill-biased technology shocks. In response to skill-neutral improvements in technology, hours worked increase. Second, the relative price of investment rises in response to skill-biased improvements in technology, indicating that capital and skill are substitutes in the aggregate production process. Both findings have important implications for the interpretation of well-known results in the literature.

The fall in hours worked in response to technology shocks, as documented by Galí (1999), has typically been interpreted as evidence for price rigidities. Having access to an improved production technology, which reduces marginal costs, a firm would like to reduce prices in order to increase sales. If prices are rigid however, the firm adjusts labor input in order to produce the amount it can sell. Our results cast doubt on this interpretation. We document a sharp drop in hours worked in response to skill-biased, but not in response to skill-neutral technological improvements. This finding suggests that at least part of the fall in hours is driven by a compositional change in labor demand. In response to a skill-biased improvement in technology, firms increase their relative demand for skilled labor. Since high skilled workers are on average more productive than low skilled workers, effective labor input may increase even if total hours worked fall.

Our conclusion that capital and skill are substitutes rather than complements in the aggregate production function, is based on our finding that the relative price of investment goods increases in response to skill-biased technological improvements, or, vice versa, that the skill premium falls in response to investment-specific technology shocks. If capital and skill are complements, as Krusell et al. (2000) argue, we would expect these effects to be reversed. Is it reasonable to think that capital and skill are substitutes rather than complements? Clearly, the answer depends on the type of capital and therefore the time period under consideration. In the industrial revolution, new production technologies often involved machines that could be operated by unskilled workers and replaced skilled laborers.²⁰ Regarding more recent technological developments, Autor et al. (2003) make the point that computer capital complements workers performing nonroutine problem-solving tasks, but substitutes labor in “cognitive and manual

²⁰For example, hand weavers, a skilled profession, opposed the adoption of weaving machinery, “because it threw many of them out of work. Those who got jobs in the factories were obliged to take the same pay as unskilled workers.” (<http://history-world.org/Industrial%20Intro.htm>)

tasks that can be accomplished by following explicit rules.” Since both non-routine and routine tasks may be performed by either skilled or unskilled workers, the aggregate elasticity of substitution between capital and skill may vary with the task composition of the workforce. Our results indicate that over the last 20 years, technological improvements in capital mostly substituted skilled workers. The reason that the skill premium nevertheless increased over this period, is due to investment-neutral technological progress, in which there was strong skill bias.

A Identification with both 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 seek to find a unique transformation matrix A that maps the reduced form residuals v_t into structural shocks e_t . Assuming orthogonality between the structural residuals and normalizing their variance to unity, A therefore satisfies $AA' = \Omega$ where Ω is the variance matrix of the reduced form residuals. In a VAR with n variables, another $n(n - 1)/2$ restrictions are then necessary for exact identification and will come out of the short- and long-run assumptions.

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. 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 neutral or non-technology shocks affect the supply of skill in the short run (on impact). This is equivalent to restricting $a_{12} = a_{13} = \dots = a_{1n} = 0$ (with a_{ij} being elements of A). These zero restrictions in the first row of A , combined with

$$A_{1.} * A'_{.1} = \Omega_1. \quad (6)$$

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. As in the standard long-run assumptions, the subsequent remaining columns should incorporate the effects of the various technology shocks. As before, we therefore use a Cholesky decomposition of the infinite horizon forecast error variance in order to measure the technology shocks. However, we only need to use the lower right block of this matrix, i.e. the part of the forecast variance which remains after the first row and column of A have already been taken into account. The Cholesky decomposition then delivers the remaining elements of A .

References

- Acemoglu, Daron**, “Technical Change, Inequality and the Labor Market,” *Journal of Economic Literature*, March 2002, 40 (1), 7–72.
- Autor, David H., Frank Levy, and Richard J. Murnane**, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, November 2003, 118 (4), 1279–1333.
- , **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*, 1989, 79 (4), 655–673.
- Canova, Fabio, David Lopez-Salido, and Claudio Michelacci**, “The effects of technology shocks on hours and output: A robustness analysis,” *Journal of Applied Econometrics*, forthcoming.
- Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan**, “Are Structural VARs with Long-Run Restrictions Useful in Developing Business-Cycle Theory?,” *Journal of Monetary Economics*, November 2008, 55 (8), 1337–1352.
- 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*, 2002, 5 (2), 243–284.
- 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*, 2006, 114 (3), 413–451.
- Galí, Jordi**, “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?,” *The American Economic Review*, 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*, 1997, 87 (3), 342–362.
- 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 Cyclicalities 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 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*, 2004, *2* (2-3), 361–371.

Table 1: Unconditional business cycle correlations

	Std	Correlation with			
		Output	Hours	Productivity	Price
<i>Baseline measure</i>					
Skill premium	0.0042	0.1597	-0.1266	0.5042	-0.2263
Relative employment	0.0166	-0.4824	-0.3206	-0.3893	0.7344
<i>Naïve measure</i>					
Skill premium	0.0042	0.0445	0.2001	-0.2392	0.1015
Relative employment	0.0144	-0.4853	-0.3990	-0.2684	0.7208
Relative supply	0.0093	0.0460	0.1466	-0.1502	0.4963

Notes: Data series are constructed as explained in section 2.3 and seasonally adjusted using X-12-ARIMA. The series are HP-filtered with $\lambda=1600$.

Table 2: Variance decomposition with identified technology shocks

	without supply shocks			with supply shocks		
Horizon	8	16	32	8	16	32
<i>output</i>						
techn. shock	36.26 (14.6,55.9)	35.78 (14.2,56.0)	35.72 (13.5,56.3)	36.11 (19.6,53.9)	36.42 (20.0,53.9)	36.24 (20.2,53.9)
supply shock				4.38 (0.8,13.2)	3.68 (0.6,12.0)	3.53 (0.5,11.8)
<i>hours</i>						
techn. shock	5.30 (1.0,18.2)	6.05 (0.9,20.2)	6.35 (0.7,21.1)	5.79 (1.1,19.1)	6.62 (1.0,22.0)	6.96 (0.9,23.0)
supply shock				1.76 (0.3,6.6)	1.74 (0.3,6.5)	1.69 (0.2,6.6)
<i>premium</i>						
techn. shock	7.86 (2.3,17.5)	7.75 (1.9,18.7)	7.70 (1.5,19.2)	7.87 (2.4,16.6)	7.28 (1.9,17.0)	7.08 (1.7,17.3)
supply shock				5.09 (2.1,13.0)	4.25 (1.4,12.2)	3.79 (0.9,12.2)

Notes: Numbers are in percents; the contribution of all shocks, including the (omitted) residual shock, adds up to 100% at each horizon. We report posterior medians and 68% Bayesian confidence bands from the posterior distribution.

Table 3: Variance decomposition with identified skill-biased technology shocks

Horizon	8	16	32
<i>output</i>			
SBT shock	4.30 (0.6,15.9)	4.39 (0.5,16.5)	4.48 (0.5,16.6)
neutral shock	40.25 (22.7,56.5)	39.74 (21.6,56.8)	39.84 (21.4,56.5)
<i>hours</i>			
SBT shock	12.43 (1.9,30.6)	12.05 (1.7,30.6)	12.12 (1.6,30.3)
neutral shock	6.66 (1.2,19.9)	7.78 (1.2,22.8)	8.37 (1.2,23.7)
<i>premium</i>			
SBT shock	96.77 (93.0,98.5)	98.25 (96.4,99.2)	99.09 (98.2,99.6)
neutral shock	1.12 (0.5,2.5)	0.61 (0.3,1.4)	0.31 (0.1,0.7)

Notes: Numbers are in percents; the contribution of all shocks, including the (omitted) residual shock, adds up to 100% at each horizon. We report posterior medians and 68% Bayesian confidence bands from the posterior distribution.

Table 4: Robustness of the response of hours to skill-biased and skill-neutral technology shocks

	SBT shock on hours	skill-neutral shock on hours
Baseline specification		
	-, significant at all horizons	+, sign. after 3rd quarter
with supply shocks	-, significant at all horizons	+, sign. after 3rd quarter
Variation of the baseline specification with baseline wage premium		
Minnesota prior with 8 lags changed to		
2 lags	-, not sign. any horizon	+, not sign. any horizon
4 lags	-, significant on impact	+, sign. in long run
12 lags	-, significant all horizons	-, significant on impact
weaker prior ¹	-, significant all horizons	+, sign. after 3rd quarter
Flat prior (OLS equivalent)		
2 lags	-, not sign. any horizon	+, not sign. any horizon
4 lags	-, significant on impact	+, sign. in long run
Baseline specification with different wage premium series		
Naive measure	+, small effect, not sign.	+, sign. in long run
Lindquist measure ²	+, small effect, not sign.	+, sign. in long run

Notes: 1) Decay parameter $d = 1$ instead of $d = 3$ as in the baseline; 2) The Lindquist measure uses data from Lindquist (2004), which are similar but not identical to our naive measure, see section 2.3.

Table 5: Variance decomposition with investment-specific technology shocks

Horizon	8	16	32
<i>output</i>			
i-specific shock	56.77 (40.6,71.5)	58.18 (40.8,72.6)	57.81 (39.4,73.3)
neutral shock	2.82 (0.6,8.4)	2.56 (0.5,8.5)	2.48 (0.4 ,8.9)
<i>hours</i>			
i-specific shock	37.41 (19.5,55.7)	36.34 (17.0,56.6)	33.96 (14.0,55.4)
neutral shock	6.39 (1.5,16.5)	5.31 (1.1,15.3)	4.98 (0.8,15.6)
<i>premium</i>			
i-specific shock	9.82 (4.6,20.2)	13.46 (4.7,29.8)	15.26 (4.2,35.0)
neutral shock	22.61 (14.6,32.7)	22.64 (13.9,34.4)	22.56 (13.2,35.2)

Notes: Numbers are in percents; the contribution of all shocks, including the (omitted) residual shock, adds up to 100% at each horizon. We report posterior medians and 68% Bayesian confidence bands from the posterior distribution.

Table 6: Variance decomposition with identified skill-biased and investment-specific technology shocks

Horizon	8		16		32	
	I	II	I	II	I	II
<i>output</i>						
i-shock (ul,lb)	58.95 (44.4,71.4)	45.84 (31.4,60.2)	62.40 (47.5,75.4)	48.73 (32.9,63.3)	63.75 (48.5,77.5)	50.4 (33.1,64.6)
SBT shock (lb,ub)	0.91 (0.1,4.0)	9.99 (1.8,26.5)	0.88 (0.1,3.9)	10.50 (1.8,28.1)	0.84 (0.1,3.9)	10.8 (1.7,28.6)
neutral shock	6.64 (2.5,12.2)	6.44 (2.4,12.9)	5.66 (2.0,11.2)	5.45 (1.9,11.6)	5.29 (1.7,10.7)	5.16 (1.7,11.1)
<i>hours</i>						
i-shock (ul,lb)	38.89 (22.8,55.6)	21.91 (10.4,36.9)	42.37 (24.5,61.1)	24.24 (11.8,40.3)	43.55 (24.4,62.7)	25.1 (12.1,42.2)
SBT shock (lb,ub)	10.34 (4.2,20.8)	25.13 (10.5,43.4)	9.11 (3.5,19.7)	24.71 (9.7,43.9)	8.84 (3.3,19.1)	24.47 (9.2,43.7)
neutral shock	1.36 (0.5,4.3)	1.53 (0.5,4.7)	0.91 (0.3,3.5)	1.09 (0.3,3.8)	0.78 (0.2,3.2)	0.91 (0.2,3.6)
<i>premium</i>						
i-shock (ul,lb)	8.71 (4.0,17.6)	1.65 (0.7,3.9)	9.39 (3.0,22.4)	0.87 (0.4,2.1)	9.65 (2.3,25.9)	0.4 (0.2,1.1)
SBT shock (lb,ub)	89.57 (80.8,94.4)	96.46 (93.7,98.2)	89.76 (76.6,96.0)	98.17 (96.8,99.0)	89.76 (73.7,97.2)	99.08 (98.4,99.5)
neutral shock	0.52 (0.2,1.1)	0.48 (0.2,1.1)	0.27 (0.1,0.6)	0.25 (0.1,0.6)	0.14 (0.1,0.3)	0.12 (0.1,0.3)

Notes: Numbers are in percents; the contribution of all shocks, including the (omitted) residual shock, adds up to 100% at each horizon. We report posterior medians and 68% Bayesian confidence bands from the posterior distribution.

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: Impulse-responses to identified technology shocks




Notes: Percent responses to a positive one-standard-deviation shock.
 Confidence intervals are 68% Bayesian bands.
 First row: Responses to baseline identification of technology shocks.
 Second row: Responses to using the Basu et al. (2006) measure of TFP instead of labor productivity in the baseline specification.
 Third row: Responses to identified technology shocks controlling for skill supply shocks.

Figure 4: Impulse-responses with identified skill-biased technology shocks




Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 5: Impulse-responses of wages with identified skill-biased technology shocks




Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 6: Comparison of SBT shocks from the VAR and from the production function decomposition



Notes: Black solid line depicts identified SBT shock, red dashed line shows the residual from the production function decomposition.

Figure 7: Impulse-responses to SBT shocks from the VAR
and from the production function decomposition

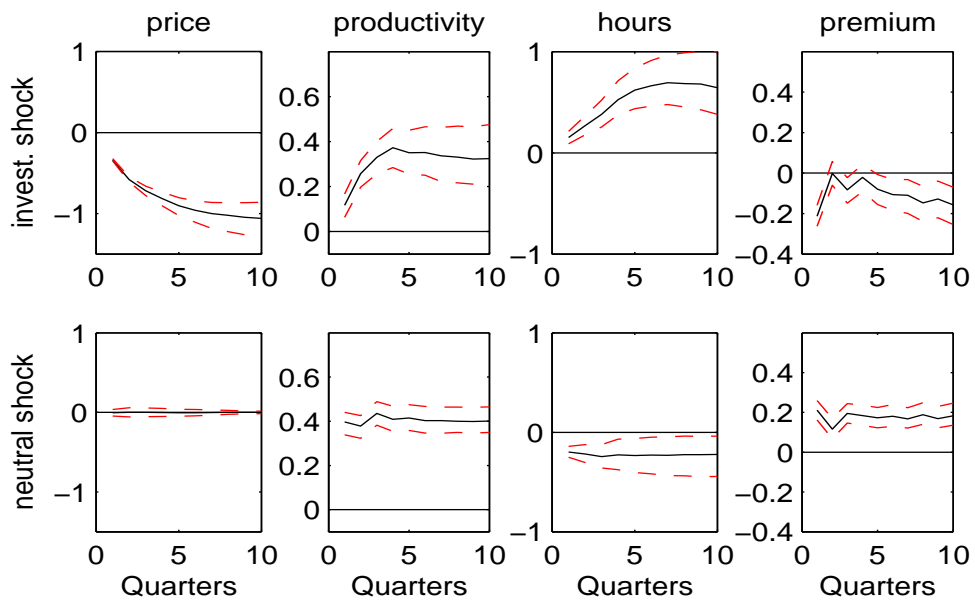


Notes: Percent responses to a positive one-standard-deviation shock.

First row: Impulse-responses from regression of the skill premium, productivity, hours worked and the relative supply on six lags of the identified SBT shock.

Second row: Impulse-responses from regression of the variables on six lags of the residual from the production function decomposition. The black dotted line repeats the estimate from the first row. Confidence intervals are one standard error bands.


Figure 8: Impulse-responses with identified investment-specific technology shocks



Notes: Percent responses to a positive one-standard-deviation shock.

Confidence intervals are 68% Bayesian bands.

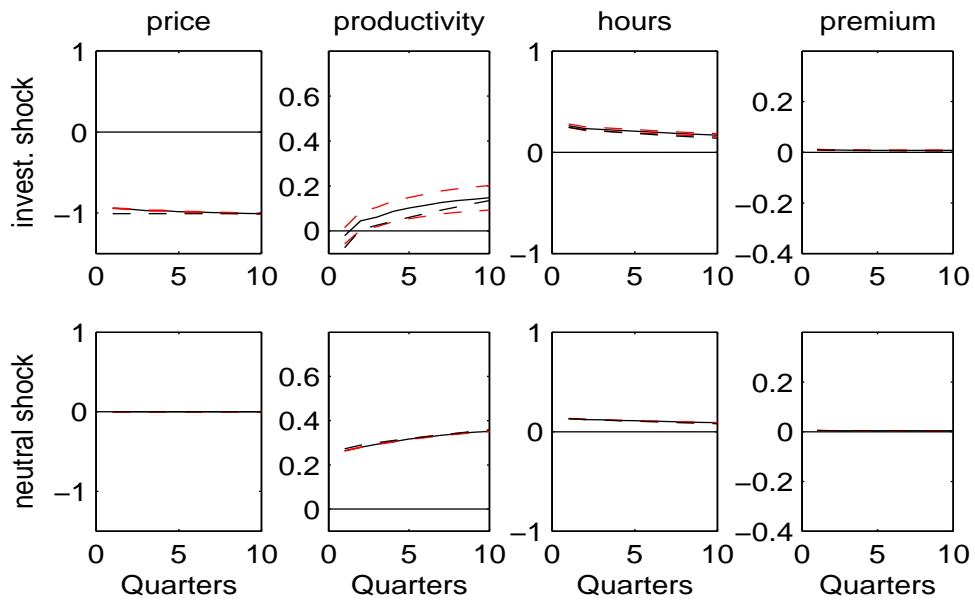
Figure 9: Impulse-responses with identified skill-biased technology shocks including the relative price of investment goods



Notes: Percent responses to a positive one-standard-deviation shock.


Confidence intervals are 68% Bayesian bands.

Figure 10: Impulse-responses from a model with investment-specific shocks and shocks to total factor productivity



Percent responses to a positive one-standard-deviation shock. The dashed lines represent the theoretical responses from the model with $\rho = \sigma = 1.67$. The solid lines are the estimated responses from 1000 simulations of 88 quarters each of the same model. The responses are normalized to match the responses of the investment price and labor productivity in the actual data in the longer run (20 quarters).

Figure 11: Capital-skill substitutability



Notes: Black line depicts response of the premium from the estimated structural VAR with actual data together with the Bayesian 68% confidence bands (red dotted lines). The dashed lines show the responses from the model with $\rho = 0.67$, $\rho = 1.17$, $\rho = 2.17$, $\rho = 2.67$ and $\rho = 5$ respectively.