



Kiel

Working Papers

**Kiel Institute
for the World Economy**



Skill-Biased Technological Change and the Business Cycle

by Almut Balleer and Thijs van Rens

No. 1775 | June 2012

Kiel Working Paper No. 1775 | June 2012

Skill-Biased Technological Change and the Business Cycle*

Almut Balleer and Thijs van Rens

Abstract:

Over the past two decades, technological progress in the United States has been biased towards skilled labor. What does this imply for business cycles? We construct a quarterly skill premium from the CPS and use it to identify skill-biased technology shocks in a VAR with long-run restrictions. Hours fall in response to skill-biased technology shocks, indicating that at least part of the technology-induced fall in total hours is due to a compositional shift in labor demand. Investment-specific technology shocks reduce the skill premium, indicating that capital and skill are not complementary in aggregate production.

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

JEL classification: E24 E32 J24 J31

Almut Balleer

IIES-Institute for International Economic
Studies, Stockholm University
Universitetsvägen 10
10405 Stockholm, Sweden
Telephone: +46 8 16 22 68
E-mail: almut.balleer@iies.su.se

Thijs van Rens

CREI, Universitat Pompeu Fabra, Barcelona
GSE, CEPR, IZA and Kiel Institute
Ramon Trias Fargas 25
08005 Barcelona, Spain
Telephone: +34 93 542 25 98
E-mail: tvanrens@crei.cat

* We thank David Autor, Toni Braun, Paula Bustos, Fabio Canova, Vasco Carvalho, Daniele Coen-Pirani, Juan-Carlos Conesa, Marco Del Negro, Marty Eichenbaum, Jonas Fisher, Jordi Gali, Gino Gancia, Ferre de Graeve, Per Krusell, Lars Ljungqvist, Alberto Martin, Monika Merz, Kris Nimark, Franck Portier, Giorgio Primiceri, Valerie Ramey, Barbara Rossi, Harald Uhlig and Fabrizio Zilibotti 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, SEJ2006-02235 and ECO2008-01665); the Generalitat de Catalunya (DIUE grants Beatriu de Pinos and 2009SGR-1157); and the Barcelona GSE Research Network.

The responsibility for the contents of the working papers rests with the author, not the Institute. Since working papers are of a preliminary nature, it may be useful to contact the author of a particular working paper about results or caveats before referring to, or quoting, a paper. Any comments on working papers should be sent directly to the author.

Coverphoto: uni_com on photocase.com

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)).

This paper documents a set of stylized facts about the implications of skill-biased technological change for business cycle fluctuations. To our knowledge, this paper is the first to undertake this task. 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 the fall in hours in response to improvements in technology is due at least in part to a compositional shift in labor demand from unskilled to skilled workers. Second, we show that the response of the economy to skill-biased technology shocks implies restrictions on the production technology that are of interest to macroeconomists studying growth as well as business cycles. In particular, we find that investment-specific technological change reduces the skill premium. These results reject the hypothesis that there is capital-skill complementarity in the aggregate production function.

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 Survey (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 hours 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 zero and sign restrictions. We find evidence for substantial 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. Part of this restriction is a long-run zero restriction, as in Blanchard and Quah (1989): we argue that skill-biased technology shocks are the only shocks that affect the skill premium in the long

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.

run. We complement this zero restriction with two sign restrictions. First, we require that skill-biased technology shocks, which are shocks to the demand for skill, affect the skill premium and the relative hours of skilled labor in the same direction. This rules out shocks to the supply of skilled labor, which also affect the skill premium in the long run. Second, we require that skill-biased technology shocks affect the skill premium and productivity in the same direction. This rules out technology shocks biased towards unskilled labor, which increase productivity but decrease the skill premium (or vice versa). We identify all other technology shocks as the remaining shocks that permanently change labor productivity, following Galí (1999). These other technology shocks include skill-neutral as well as unskill-biased shocks. We find that skill-biased improvements in technology cause a decline in total hours worked. This finding suggests that the fall in hours in response to technology shocks, which has been interpreted as evidence for price rigidities, is due at least in part to a compositional shift in labor demand towards skilled workers.

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, skilled labor or unskilled labor augmenting technological change. 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.³ Our results cast doubt on this hypothesis.

³It 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

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 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 reduce the skill premium, while investment-neutral technology shocks have a positive effect on this variable. 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. We obtain the best match of the response of the skill premium to investment-specific shocks in the model-simulated data to the response estimated from actual data, if we assume capital and skill are substitutable.

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. Section 4 discusses our evidence against capital-skill complementarity in aggregate production. Section 5 concludes.

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.

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)).

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

There are three technology parameters in this production function: A_t is neutral, $B_{H,t}$ skilled labor augmenting and $B_{L,t}$ unskilled labor augmenting technology. Increases in A_t are improvements in skill-neutral technology (SNT). Increases in $B_{H,t}$ and $B_{L,t}$ 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 $B_{H,t}$ increases the demand for skill and therefore the skill premium (assuming the supply curve for skill is downward sloping) and an increase in $B_{L,t}$ decreases the skill premium. 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 (SBT) shocks as changes in $B_{H,t}$ and unskill-biased technology (UBT) shocks as changes in $B_{L,t}$. Note that both positive SBT shocks and negative UBT shocks increase the skill premium, but the two shocks are conceptually different. Positive technology shocks, whether to A_t , $B_{H,t}$ or $B_{L,t}$ represent improvements in technology that raise total factor productivity, whereas negative technology shocks of all types reduce productivity.

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 $B_{H,t}$ 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 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. Second, because investment goods are used to build up capital, which is an input in the production process, sector-specific shocks 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.

⁴This 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 $B_{H,t}$ and $B_{L,t}$ are indistinguishable from changes in A_t .

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 decreases in the relative price of investment goods p_t reflect technological improvements in the investment goods producing sector or investment-biased technology (IBT) shocks. An aggregate production function that allows for capital-skill complementarity is a slightly generalized version of (1), where A_t now denotes not only skill-neutral but also investment-neutral technology (INT).

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

The elasticity of substitution between skilled and unskilled labor σ now also measures the elasticity of substitution between capital and unskilled labor, whereas ρ is the elasticity of substitution between capital and skilled labor. As shown by Krusell et al. (2000), improvements in investment-specific technology (positive IBT shocks) 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 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,

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

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 $B_{H,t}$ or unskill-biased technology $B_{L,t}$ represent shifts of the skill demand curve or skill demand shocks.

The first, and easiest, way to estimate shocks to the relative demand for skill is a type of production function decomposition. Since the skill premium and the relative quantity of skill are observable, these shocks can be calculated directly from equation (4), using an estimate for the elasticity of substitution between low and high skilled workers σ .⁵ This approach, the results of which are described in Section 3.1, has two disadvantages. First, we cannot separately identify skill-biased and unskill-biased improvements in technology. This is problematic, because the effects of an improvement in skill-biased technology and a deterioration in unskill-biased technology on the economy are likely to be quite different, even though both lead to an increase in the skill premium. Second, the estimates for the skill-biased technology shocks obtained this way are identified from the assumption that wages are proportional to marginal products. This assumption is not problematic if labor markets (and product markets) are perfectly competitive and 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. In order to address the second issue, we will use only long-run effects to identify skill-biased technological

⁵An estimate for the share parameter β is unnecessary since this parameter affects only the level of $B_{H,t}$ and $B_{L,t}$, and we normalize the mean and variance of the shocks to zero and one respectively.

progress. To address the first issue, we use the additional restriction that improvements in technology increase productivity.

We implement our identification strategy using a structural VAR with long-run restrictions. Consistent with equation (4), we identify skill-biased and unskill-biased technology shocks as the only shocks that affect the skill premium in the long run, conditional 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 .⁶ In Section 3.2, we compare the estimated technology shocks using long-run restrictions to the estimated shocks using the production function decomposition described above and find that the differences are small.

To separately identify skill-biased and unskill-biased technology shocks, we use a long-run sign restriction on the response of productivity in addition to the long-run zero restriction on the response of the skill premium.⁷ In addition, we use a second long-run sign restriction to separate skill-biased technology shocks from exogenous changes in the supply of skilled labor, which may also affect the skill premium in the long run. We implement these additional restrictions, which are described in more detail in Sections 3.3 and 3.4, only in the structural VAR, because it would be impossible to do in the context of a production function decomposition. The results, described in Section 3.5, show that the sign restrictions affect the results substantially, indicating that there were skill-biased as well as unskill-biased technology shocks and skill supply shocks over the

⁶Of 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.

⁷Alternatively, we could have exclusively used sign restrictions, imposed on a broader range of frequencies, as in Uhlig (2005) and Dedola and Neri (2007). We opt for a long-run zero restrictions in combination with a sign restriction because we believe that any assumption on the short run behavior of the skill premium would be more problematic than the assumption that only skill-biased technology shocks affect the premium in the long run.

sample period.

The estimation of structural shocks using long-run zero and sign restrictions is implemented in two steps. First, we estimate a reduced form VAR in the variables labor productivity, total hours worked, the skill premium, relative hours of skilled workers and in some specifications also the relative price of investment goods. Second, we map the reduced form coefficients and residuals into structural coefficients by means of our identifying restrictions.⁸ Our baseline VAR includes 8 lags and is estimated on quarterly data from 1979:I to 2006:II. All variables are used in first differences in order to allow for unit roots.⁹ The baseline specification includes labor productivity, hours worked, skill premium and relative hours of skilled workers. In Section 3.7, we show that our results are robust to adding other variables, such as consumption and investment. In order to identify investment-specific technology shocks, we further include the relative price of investment goods into the VAR in Section 4.

We use a Bayesian VAR to estimate the reduced form and employ a prior on the coefficients of the lagged variables, similar to Canova et al. (2010). This prior is a type of Minnesota or Litterman prior, reflecting the belief that the true data generating process for each variable is a univariate unit root, so that in first differences the variables are serially uncorrelated. It is implemented as a prior that the lag coefficients in the VAR are close to zero, where the strength of the prior increases with the lag order.¹⁰ 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

⁸Please refer to the Technical Appendix for details on the implementation of these restrictions.

⁹In 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. (2010) 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. We show in section 3.7 that our results are robust to using hours worked in levels even if we do not filter out the low frequencies.

¹⁰The strength of the prior increases with lag length to reflect the belief that the higher order lags are less likely to matter. This is imposed in form of a harmonic decay of the prior variance on the lag coefficients. Apart from the decay, the prior employed is quite loose. The Technical Appendix provides more information on the specification.

the strength of the prior and to estimating the reduced form VAR using ordinary least squares.

2.3 Data

We construct quarterly series for the skill premium and the relative hours worked 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:I to 2006:II.¹¹ 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. 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).¹² We use these data to construct the skill premium as the log wage differential between college graduates and high school graduates, controlling for other sources of observable heterogeneity. In an average quarter, we have wage and education data for about 35,000 workers.¹³

The other data series we use in our analysis are the following. Output is non-farm business output per capita of all persons from the national income and product accounts (NIPA). Hours per capita are hours of non-supervisory workers in the non-farm business sector from the Current Employment Statistics establishment survey, corrected

¹¹The 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.

¹²The most important change in the education question occurs in 1992. Until 1991, educational attainment is coded as years of schooling. From 1992, the coding is based on the degree. As a result, there are some jumps in the fractions of workers in each educational category in this year. We correct these jumps by imposing that the fractions of workers in each of the 5 categories do not change from 1991:IV to 1992:I beyond what may be expected based on seasonal effects and a slow-moving trend.

¹³Please refer to the Technical Appendix for details on the construction of the skill premium.

to be representative for the entire workforce including supervisors. 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 of the quality adjusted NIPA deflator for producer durable equipment over the consumption deflator. These data were constructed by DiCecio (2009), extending the series by Fisher (2006) and based on the annual data constructed by Gordon (1990) and Cummins and Violante (2002).¹⁴

Table 1 shows business cycle statistics for the skill premium, the relative hours worked and supply of skill, output, hours, productivity and the relative price of investment goods for our estimation sample 1979:I to 2006:II. The skill premium is basically acyclical: it is only very mildly positively correlated with output and productivity. 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 hours of skill are 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 negative and insignificant. This is a first indication that capital-skill complementarity does not seem an important feature of the data at business cycle frequencies.

3 Skill-biased technology shocks

In this section, we present our results for the effects of technology shocks on aggregate variables. We start with a simple production function decomposition, which allows estimating shocks that affect the skill premium, but not for separately identifying skill demand shocks (technology shocks) versus skill supply shocks, and skill-biased versus unskill-biased technology shocks. We then estimate the same shocks again, this time

¹⁴We thank Ricardo DiCecio for making these data available to us. The Technical Appendix describes the time-series properties of the data that are relevant for the specification of the VAR such as autocorrelations, integration and cointegration properties.

using a structural VAR with long-run restrictions and find that the estimates are very similar. The structural VAR framework allows for additional restrictions to separate out skill supply shocks and unskill-biased technology shocks. Imposing these restrictions gives rise to our baseline estimates in Section 3.5. Finally, we report responses to skill-biased technology shocks for some additional variables and explore the robustness of our results.

3.1 Production function decomposition

Our first estimates of skill and unskill-biased shocks are constructed using a simple decomposition based on the production function, as described in 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 and unskill-biased technology $\log(B_{H,t}/B_{L,t})$ from our data on the skill premium and the relative hours of skill. For comparability with the identified shocks from a structural VAR in the continuation of this section, we demean these changes and normalize their variance to unity.

In order to obtain impulse responses of aggregate variables to skill and unskill-biased technology shocks, we regress these variables on lags of the estimated shocks, 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 from a VAR. Since the impulse responses 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 1. Not surprisingly, skill and unskill-biased technology shocks estimated in this manner increase the skill premium. On average, these shocks seem to have little effect on productivity and hours worked.

The finding that skill and unskill-biased shocks do not increase labor productivity

seems inconsistent with our interpretation of these shocks as technology shocks. However, it is important to remember that there are two types of technology shocks that affect the skill premium. Positive skill-biased technology shocks, i.e. increases in $B_{H,t}$, and negative unskill-biased technology shocks, i.e. decreases in $B_{L,t}$, both increase the skill premium, but have opposite effects on productivity. Our estimates indicate that over the sample period, both types of shocks were present, and the increase in productivity driven by positive SBT shocks was compensated by a decrease in productivity driven by UBT shocks.¹⁵ In Section 3.3, we describe an additional restriction to separately identify skill-biased technology shocks $B_{H,t}$. Since this restriction can only be imposed in the context of a structural VAR, we first make sure we can replicate the results from the production function decomposition in a VAR with long-run restrictions.

3.2 Long-run restrictions

Identifying skill and unskill-biased technology shocks from their long-run effects on the relative price of skill is consistent with skill demand equation (4), but more general because we do not require this equation to hold true in each period. The long-run restriction we use is similar in spirit to the identification of investment-biased technology shocks as shocks that affect the relative price of investment goods proposed by Fisher (2006). Skill and unskill-biased technology shocks identified in this manner may or may not affect labor productivity. Skill-neutral technology shocks, following Galí (1999), 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. Our identification scheme is strictly speaking is not a *decomposition* of technology shocks into skill and unskill-biased versus 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 changes

¹⁵The Technical Appendix discusses this issue in more detail and argues that UBT shocks were particularly important in the post-2000 period.

in the skill supply that affect the skill premium in the long run (we address the issue of skill supply shocks in Section 3.4).

For comparability with the results from the production function decomposition, we first regress the skill premium, labor productivity and total hours worked on 6 lags of the identified skill and unskill-biased shocks from the structural VAR. The responses obtained in this manner are presented in row 2 of Figure 1. By the identifying assumption, positive skill-biased and negative unskill-biased shocks drive the skill premium up in the long run. The estimates indicate that this effect is realized immediately on impact. The response of labor productivity to skill and unskill-biased shocks identified using long-run restrictions is very similar to its response to the same shocks estimated from a production function decomposition as well. This result indicates that equation (4) is a good description of skill demand at all frequencies, not only in the long run. It also shows that the structural VAR identifies the same skill and unskill-biased shocks as a decomposition using the production function, mirroring a similar equivalence result for neutral technology shocks in Basu et al. (2006).

Rows 3 and 4 in Figure 1 show the responses of the skill premium, productivity and hours worked to one-standard deviation skill and unskill-biased and skill-neutral technology shocks, calculated directly from the coefficient estimates of the VAR. Here, as in all graphs that will follow, we present the median as well as the 16th and 84th percentiles of the posterior distribution of the structural impulse-response coefficients, following Uhlig (2004). The responses to skill and unskill-biased shocks estimated in this manner are again very similar to the responses to shocks from production function decomposition. Skill-neutral technology shocks have no significant effect on the skill premium at any horizon (by assumption, there is no effect in the long run). The response of productivity and hours to these shocks looks very similar to the response of these variables to identified technology shocks from a VAR without skill-biased shocks, as in Galí (1999).¹⁶

¹⁶In the Technical Appendix, we replicate the estimates in Galí (1999) and show that the results are

3.3 Unskill-biased technology shocks

How to separately identify skill-biased from unskill-biased technology shocks? From skill demand equation (4) it is clear that improvements in skill-biased technology $B_{H,t}$ and deteriorations in unskill-biased technology $B_{L,t}$ raise the price of skill in the same way. It is not possible, therefore, to separately identify each type of technological change using data on the skill premium alone. However, using data on labor productivity, the sign of technology shocks is observable: positive technology shocks increase and negative technology shocks decrease productivity, see production function (1). We implement this observation as a sign restriction. For the same reasons as set out in Section 3.2, we impose this sign restriction only on the long-run response of labor productivity, although the results change very little if we impose the restriction at other horizons as well. We identify skill-biased technology shocks as those shocks that affect the skill premium in the same direction as labor productivity in the long run. Other technology shocks, which affect labor productivity in the long run, may be unskill-biased or skill-neutral.

Because skill-biased, unskill-biased and skill-neutral technology shocks are linearly dependent, it is not possible to identify all three shocks separately. For example, an increase in unskill-biased technology $B_{L,t}$, which increases productivity and decreases the skill premium, is observationally equivalent to the combination of a decrease in skill-biased technology $B_{H,t}$, which decreases the skill premium, and an increase in skill-neutral technology A_t , which increases productivity. Therefore, we separately identify skill-biased technology shocks and refer to the remaining technology shocks as ‘other’ technology shocks, which include both unskill-biased and skill-neutral shocks.

We implement this identification scheme by assuming that only technology shocks may affect labor productivity and the wage premium in the long-run (usual long-run zero restrictions). We then impose sign restrictions on the long-run variance between these two variables in order to separate skill-biased technology from other technology

robust to changing the sample to our time period and to adding the skill premium as an additional variable, and that they differ very little from the responses presented here.

shocks. The Technical Appendix provides more details on this procedure.

3.4 Shocks to the supply of skill

In the identification of technology shocks, we have so far assumed that skill-biased and unskill-biased technology shocks are the only shocks that affect the skill premium, and technology shocks are the only shocks that affect productivity in the long run. Neither assumption is valid in the presence of exogenous changes in the relative supply of skill, $\log(H_t/L_t)$, because of the standard simultaneity problem in estimating demand and supply equations, see equation (4).

Suppose a preference shock causes college enrollment to increase permanently or cheaper child care makes market work more attractive for highly educated parents. In both cases, the supply of skill increases for reasons unrelated to the production technology. The increase in skill supply must decrease the skill premium because skill demand is not affected. Productivity may be affected as well, although the direction of the effect is ambiguous. On the one hand, the lower skill premium leads firms to employ relatively more skilled workers, which tends to raise average labor productivity if skilled workers are more productive than unskilled workers. On the other hand, diminishing returns in skilled labor push down productivity.

To make sure that our estimates for skill-biased technology shocks do not include shocks to the supply of skill, we separately identify skill supply shocks using a second sign restriction. As opposed to skill-biased technology shocks, which are skill demand shocks, skill supply shocks affect the price and the quantity of skill in opposite directions. We exploit this property to identify these shocks. For this purpose, we include a measure of the relative hours worked of skilled workers in the VAR.

3.5 Identified skill-biased technology shocks

Figure 2 shows the responses of the skill premium, relative hours worked of skilled workers, labor productivity and total hours worked per capita to skill-biased and other technology shocks from our baseline specification using long-run zero and sign restrictions. Rows 1 and 2 present the responses if we separately identify skill-biased and unskill-biased technology shocks, as described in Section 3.3, but ignore shocks to the supply of skill. In rows 3 to 5 we separately identify these shocks as well, as in Section 3.4. By the identifying assumption, skill-biased technology shocks raise the skill premium, the relative hours of skilled workers as well as productivity in the long run. In both cases, the increase is significant and fully realized on impact, indicating that all technology shocks are close to permanent. The effect of skill-biased and other technology shocks on productivity is roughly of the same magnitude. Other technology shocks reduce the skill premium, indicating that these shocks now include improvements in unskill-biased technology.

Hours worked fall strongly and significantly in response to skill-biased improvements in technology, but not in response to other technological improvements.¹⁷ The finding that hours fall after a positive technology shocks, first documented by Galí (1999), is typically interpreted as evidence for price rigidities. Rigid prices dampen the substitution effect on impact and thus make the income effect of higher productivity, which increases the demand for leisure, dominant in the short run.¹⁸ Our results suggest, however, that part of the fall in hours may be 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. In Section 3.6, we explore this mechanism in more detail.

¹⁷This finding is qualitatively unchanged and quantitatively stronger if we separately identify investment-specific technology shocks as in Section 4.4.

¹⁸An alternative explanation that has been suggested is the combination of habit formation in consumption and adjustment costs in investment, see Francis and Ramey (2005).

The finding that hours fall in response to skill-biased but not in response to other technology shocks only becomes apparent when we separate skill supply shocks from skill-biased technology shocks. The reason is that a decrease in the supply of skill, which raises the skill premium and therefore satisfies the identifying restriction of skill-biased technology shocks, shifts employment towards unskilled workers. Since unskilled workers are on average less productive than skilled workers, firms need to increase overall hours worked in order to achieve the desired level of production.

Table 2 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 other technology shocks increases slightly the overall contribution of technology shocks to fluctuations. Skill-biased and other technology shocks together explain about 11% of the business cycle variance of output, compared to about 5% if we identify technology shocks as in Galí (1999).¹⁹ Technology shocks explain about 18% of the volatility in hours worked, compared to about 9% in the Galí (1999) specification. Skill-biased technology shocks are relatively more important for hours worked, but unskill-biased and skill-neutral shocks explain a larger fraction of fluctuations in output. Fluctuations in the skill premium are due to skill-biased technology shocks, unskill-biased technology shocks and skill supply shocks in roughly equal proportions. Skill supply shocks also explain a sizable share of the variance of output and total hours worked. Overall, we find strong evidence that skill-bias plays an important role in technological change at business cycle frequencies.

3.6 Wages and hours of high and low skilled workers

By adding additional variables to the VAR, we can evaluate their response to skill-biased technology shocks. Here, we explore the response of the wages and hours worked of skilled and unskilled workers separately, in order to provide supportive evidence for our interpretation for the fall in total hours worked in response to skill-biased technology

¹⁹These results are available in the Technical Appendix.

shocks as a compositional shift in labor demand. Figure 3 presents these responses.

The first row in Figure 3 replicates the increase in the skill premium in response to skill-biased improvements in technology as in Figure 2, and decomposes this increase into the responses of wages of skilled and unskilled workers. In response to an SBT shock, the wage of skilled workers increases substantially, while the wage of unskilled workers almost does not change. According to the point estimate, unskilled workers still benefit a little from a skill-biased improvement in technology, different from what we would expect if unskilled and skilled workers are more substitutable than complementary. However, the increase in the wage of unskilled workers is small and insignificant, consistent with our interpretation of a compositional shift in labor demand.

In the second row of Figure 3, we look directly at the quantity of labor of each type that is employed in equilibrium. By our identifying restriction, the relative hours of skilled workers with respect to unskilled workers increase in response to skill-biased improvements in technology. This increase is driven by a strong fall in hours of unskilled workers. Hours of skilled workers respond very little to a skill-biased improvement in technology. To understand this result, we argue that skill-biased improvements in technology lead to a shift in the composition of labor demand towards skilled workers. This compositional shift tends to increase hours of skilled workers and decrease hours of unskilled workers. In addition, skill-biased technology shocks have the same effect as other improvements in technology, which is a mild decrease in hours worked. The combination of the two effects is a sharp drop in hours of unskilled workers, and virtually no effect on hours of skilled workers.

3.7 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 3. The fall in hours after a skill-biased technology shock is robust across specifications. In response to other technology shocks hours worked sometimes rise and sometimes fall,

but always less than in response to skill-biased technology shocks, consistent with our interpretation that part of the fall in hours is due to a compositional shift in labor demand.

To explore whether our results may be driven by low frequency movements in the data that may not be well described by the model we have in mind, we try various ways of detrending hours worked per capita. We employ a dummy broken at 1997:I, as suggested by Fernald (2007), filter the series with a low-pass filter excluding frequencies above 52 quarters, as in Canova et al. (2010), and include a deterministic polynomial trend (up to a third order polynomial) into the equation for hours worked. We also check robustness to including hours worked in levels and to using a shorter sample ending in 2000:IV. In all of these cases, the results are qualitatively unchanged and largely quantitatively unchanged.

In our baseline estimates, we impose a prior on the decay of the lag coefficients, see Section 2.2, in order to be able to include a larger number of lags. However, our results are not driven by this prior. The responses of productivity and the skill premium to skill-biased technology shocks are virtually unaltered when we vary the strength of the prior or when we estimate the VAR using ordinary least squares (OLS).

We also re-estimated the VAR using total hours worked instead of total hours per capita, and using hours worked from the CPS rather than the usual series from the establishment survey. The CPS series is much noisier than the baseline series because the underlying micro-data sample is much smaller, but it is more consistent with our skill premium series. All results are robust to these alternative series for hours worked.

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 weakens the results but does not change them qualitatively.

Finally, as shown by Fernandez-Villaverde et al. (2007), it is important to include a proper set of variables in order to have a mapping between the VAR and the underlying

DSGE model.²⁰ Therefore, we try including additional and potentially omitted variables in the VAR: the relative price of investment goods; consumption, measured as real personal consumption expenditures from the NIPA; investment, measured consistent with the series for the relative price; and the interest rate, measured as the return on a 3-month T-bill as in Fisher (2006). Including these variables does not significantly alter any of our results.

4 Capital-skill substitutability

The relative price of investment goods fell substantially over our sample period. 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) argued that this investment-specific or investment-biased 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-biased technology shocks should lead to business cycle fluctuations in the skill premium. In this section, we explore this hypothesis and find no evidence for it.

4.1 Skill bias in investment-specific technology

Consider the alternative aggregate production function (3), as in Krusell et al. (2000), which allows for complementarity or substitutability between capital and skill. Assuming

²⁰Including additional variables may also alleviate the problem of finite lag length, see Erceg et al. (2005).

as before that wages are proportional to marginal products in the long run, expression (4) for the skill premium changes to the following.

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

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 $\sigma > \rho$, 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 unless we control for investment-specific shocks. In addition, it is interesting in itself to assess the skill bias in investment-specific shocks, because it allows us to assess the degree of capital-skill complementarity in aggregate production.

We follow Fisher (2006) in identifying investment-biased and investment-neutral technology shocks using the relative price of investment goods. Investment-biased technology shocks are the only shocks that affect the relative price of investment goods in the long run. Investment-neutral technology shocks are all remaining shocks that drive labor productivity in the long run. For implementation, we include the relative price of investment in the VAR, ordering it first, before labor productivity.

Figure 4 shows the responses of the the skill premium, labor productivity, hours worked and the relative price of investment goods to investment-biased and investment-neutral technology shocks. After an improvement in investment-specific technology, the relative price of investment falls, productivity rises and hours worked increase. An investment-neutral technology shock, has no effect on the relative price of investment, increases productivity and leads to a fall in hours worked.²¹

²¹Since 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).

The skill premium falls in response to an improvement in investment-specific technology. Thus, we find no evidence for complementarity between capital and skill in the production technology. If anything, capital and skilled labor seem to be more substitutable than capital and unskilled labor. Investment-neutral technology shocks increase the skill premium, suggesting that these shocks are more often skill-biased than unskill-biased. We further decompose these shocks into skill-biased and other technology shocks in Section 4.4 below. There, we also show that the relative hours of skilled workers decrease after an investment-specific improvement in technology, again consistent with capital-skill substitutability rather than complementarity. Total hours per capita still fall after skill-biased technology shocks when we include the relative price of investment into the specification.

4.2 Relation to previous literature

Our findings are in striking contradiction with the argument in Krusell et al. (2000). To explain the difference, we need to understand what drives identification in their paper and in ours. From the skill demand equation (5), we see that the skill premium depends directly on skill-biased and unskill-biased technology as well as on the capital-labor ratio, which in turn depends on investment-biased technology. Let $\pi_t = \log(w_{H,t}/w_{L,t})$ denote the skill premium, $b_t = \log(B_{H,t}/B_{L,t})$ the combination of skill-biased and unskill-biased technology shocks and d_t investment-biased technology. Then, using a hat to denote the deviation of a variable from its mean, this relation can be expressed in log-linear approximation as $\hat{\pi}_t = \omega_b \hat{b}_t + \omega_d \hat{d}_t$, where $\omega_b = (\sigma - 1)/\sigma > 0$ and $\omega_d > 0$ if $\sigma > \rho$, i.e. if the production function exhibits capital-skill complementarity. Suppose we were to estimate this relation as a univariate regression. Technological change \hat{b}_t and \hat{d}_t is of course unobserved, but investment-biased technology is closely related the relative price of investment goods, $\hat{d}_t = -\hat{p}_t$, so we could regress $\hat{\pi}_t = \beta_1 \hat{p}_t + \beta_2 t + \varepsilon_t$. In this regression, $\beta_1 = -\omega_d < 0$ indicates capital-skill complementarity and the remaining time trend and short-term fluctuations measure the effect of skill-biased but investment-

neutral technological change, $\beta_2 t + \varepsilon_t = \omega_b \hat{b}_t$.

The baseline estimates in Krusell et al. (2000) do not allow for skill-biased but investment-neutral technological change. Imposing $\beta_2 = 0$, the coefficient β_1 is identified largely off the trends in \hat{p}_t and $\hat{\pi}_t$. Since we know that over the sample period \hat{p}_t trends down while $\hat{\pi}_t$ trends up, we estimate $\beta_1 < 0$ and conclude there is evidence for capital-skill complementarity in this case, as Krusell et al. (2000) do. Of course, the particular estimation method is not important for this result: whether we estimate the comovement in \hat{p}_t and $\hat{\pi}_t$ in a univariate regression, in a structural model, as in Krusell et al. (2000), or in a structural VAR, as in this paper, it is clear that $\beta_1 < 0$ is necessary to match the trends in these variables if we impose $\beta_2 = 0$. However, if we do not impose this restriction and allow for skill-biased but investment-neutral technological change, then β_1 is no longer identified off the trends but off the higher frequency comovement. Since after detrending \hat{p}_t and $\hat{\pi}_t$ are positively rather than negatively correlated, we would now find $\beta_1 > 0$ and conclude that there is evidence for capital-skill substitutability instead of complementarity, as we do in this paper.²²

Lindquist (2004) argues that capital-skill complementarity explains not only the trends, but also the business cycle fluctuations in the skill premium. He develops a business cycle model with neutral and investment-specific technology shocks and evaluates this model by comparing its predictions for the (unconditional) moments to the data, in particular the fact that the skill premium is volatile but acyclical. Lindquist argues that strong capital-skill complementarity is necessary to explain these facts. In his model, investment-specific technological improvements increase the skill premium whereas neutral improvements in technology decrease the skill premium, in both cases because of capital-skill complementarity. Since business cycles are driven by both types of shocks in the model, this makes the skill premium volatile, but roughly acyclical.

²²Krusell et al. (2000) estimate a version of their model that allows for a trend in \hat{b}_t and find that this model can also “account for changes in the skill premium”. They argue, however, that the implied difference in the growth rate of productivity of skilled versus unskilled workers of 11%-points per year makes this interpretation of the data “less compelling than that of capital-skill complementarity” (p.1047).

It is crucial for Lindquist's argument that the model has at least two shocks, the effects of which on the skill premium roughly cancel out against each other. If business cycles were driven exclusively by investment-specific shocks, the skill premium would be strongly procyclical in his model. Although Lindquist presents impulse responses of the premium to each shock separately from the model, he does not compare the conditional moments to the data. Our estimated impulse responses show that his model implies the wrong response of the skill premium to investment-specific shocks. Another way to say that is that although the model with capital-skill complementarity captures the volatility of the skill premium, the implied correlation of the premium with the relative price of investment goods has the wrong sign. In Section 4.3, we show that Lindquist's model can replicate the empirical response of the skill premium to investment-specific shocks if we recalibrate the model such that capital and skill are mild substitutes rather than strong complements in the production function.

4.3 A model with capital-skill substitutability

Our finding that the skill premium falls in response to investment-specific shocks suggest that capital and skill are substitutes rather than complements in the aggregate production function. But these responses are measured with error. This raises the question what range of 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.²³ 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

²³Alternatively we could estimate the model, which would provide a more precise estimate of the degree of complementarity or substitutability between capital and skill in the production function. However, in order to do this we would have to make additional assumption about parts of the economy that are unrelated to the production function. Our test for capital-skill complementarity would then be a joint test together with these auxiliary assumptions. Therefore, we prefer to focus on the impulse response that is likely to be most informative about the degree of capital-skill complementarity and estimate this response with minimal assumptions on the structure of the rest of the economy.

interpretation of the results.²⁴

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 technology shocks are highly persistent and uncorrelated with each other.²⁵ 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 110 quarters, the same sample length as in our

²⁴In particular, it allows us to check whether our VAR includes sufficiently many lags to properly identify the true model impulse responses, addressing the potential problem with the VAR approach pointed out by Chari et al. (2008).

²⁵We assume shocks in the model are uncorrelated in order to be consistent with the identifying assumptions of our VAR. In addition, we are not sure how to interpret the predictions of a structural model with correlated shocks, which introduce comovement outside of the model. Similar to Uhlig (2004) we assume persistent, but not permanent, autoregressive processes for the shocks because the production function does not imply balanced growth. These changes in the calibration with respect to Lindquist change the simulated data very little, and none of our conclusions change if we follow Lindquist's calibration exactly.

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 add measurement error to the simulated variables as we seek to identify two shocks out of four variables in the VAR. We then estimate the VAR for each sample of 110 quarters and average the impulse responses across the 1000 simulations. Figure 5 illustrates this for the case that 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.

Figure 6 shows the impulse responses of the skill premium to an investment-biased shock according to the model for different degrees of capital-skill complementarity or substitutability, as well as the response estimated from the actual data. Comparing the response of the skill premium to investment-specific shocks in the actual data to the responses in the model, we find that our estimates are consistent with capital-skill substitutability or with capital and skill being neither substitutes nor complements. However, we can reject even weak capital-skill complementarity. Our point estimate for the long run response of the skill premium suggests an elasticity of substitution between capital and high skilled labor ρ between 2.67 and 5 which corresponds to strong substitutability between the two inputs in production.

4.4 Contribution to business cycle fluctuations

When we allow for investment-biased technology shocks, our estimates replicate the finding in Fisher (2006) that investment-specific technology is an important source of business cycle fluctuations, whereas investment-neutral technology shocks contribute only a small fraction of fluctuations in output and hours. However, our results indicate that investment-neutral shocks include technology shocks of different types, with distinct implications for the comovement of aggregate variables: skill-biased, unskill-biased and

skill-neutral technology shocks. In addition, we emphasize the potential importance of shocks to the supply of skilled labor. With the identifying restrictions discussed above, it is not possible to separately identify all these different shocks simultaneously. Recall that both investment-biased and investment-neutral technology shocks may affect the skill premium. Similarly, both skill-biased and skill-neutral or unskill-biased technology shocks may affect the relative price of investment goods.

To separately identify as many shocks as we can, we use a recursive identification scheme, identifying first investment-biased technology shocks as all shocks that affect the relative price of investment goods. Then, skill supply shocks, skill-biased and other technology shocks are identified as all remaining shocks that affect the skill premium in the long run and satisfy their respective sign restrictions, excluding shocks that affect both the relative price of investment and the skill premium. Therefore, the estimated fraction of the variance in aggregate variables that is due to skill supply shocks and skill-biased technology shocks should be interpreted as a lower bound on the actual contribution of these shocks to business cycle fluctuations. Figure 7 shows the impulse responses to the various technology shocks from this joint identification strategy. Comparing the results in this figure to the responses to skill-biased and other technology shocks in Figure 2, it is clear that the results do not change much, with hours worked still falling significantly in response to investment-neutral, skill-biased technology shocks.

Table 4 shows the variance decomposition of the forecast error variance in output, hours, the skill premium and relative hours of skilled workers. Investment-specific technology shocks explain between 30 and 40% of the volatility in output at business cycle frequencies, consistent with earlier findings in the literature (Fisher (2006), Canova et al. (2010)). Although investment-neutral technology shocks explain only a small fraction of about 5% of the forecast variance of output, these shocks are important for fluctuations in hours worked, explaining about 23% of the volatility of hours per capita, about two thirds of which is due to skill-biased technology shocks. Skill supply shocks also explain a non-negligible part of about 10% of fluctuations in hours. Investment-

specific technology shocks explain only about 8% of the volatility of hours, less than skill-biased but investment-neutral shocks. Investment-specific shocks play virtually no role for fluctuations in the skill premium, with skill-biased technology shocks and skill supply shocks together explaining over 60% of the variance in that variable and other technology shocks (including unskill-biased shocks) about 10%.

5 Conclusion

In this paper, we explored the implications of skill bias in technological change 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 zero and sign restrictions. We documented two main differences between skill-biased and other technology shocks. First, the fall in hours in response to investment-neutral improvements in productivity is driven at least in part by the skill-bias in these shocks. Second, investment-specific improvements in technology are biased towards unskilled labor, indicating that capital and skill are substitutes rather than complements 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 drop in hours worked in response to skill-biased 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 in the aggregate production function, is based on our finding that the skill premium falls in response to investment-biased technology shocks. If capital and skill are complements, as Krusell et al. (2000) argue, we would expect the demand for skill and therefore the skill premium to increase in response to investment-biased improvements in technology.

Is it reasonable to think that capital and skill are complements, substitutes or neither? 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.²⁶ 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 tasks that can be accomplished by following explicit rules.” Since both nonroutine 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 substituted skilled workers more than unskilled workers. The reason that the skill premium nevertheless increased over this period, is due to investment-neutral technological progress, which was biased towards skilled labor.

Finally, it is important to emphasize that we use a broad interpretation of what constitutes ‘technological’ change. For example, Philippon and Reshef (2010) show that financial deregulation dramatically increased the demand for skilled labor in the financial sector over the 1980-2010 period. In our estimates, this change in regulation is indistinguishable from a skill-biased change in technology. Moreover, if deregulation

²⁶For example, hand weavers, a skilled profession, opposed the adoption of weaving machinery, going so far as destroying these machines, because many of them lost their jobs and the others were forced to accept lower wages (Noble et al. (2002), p.701).

affected not only the type, but also the total amount of services provided by the financial sector, then it may even look like there is capital-skill substitutability in aggregate production (assuming the financial sector uses less capital than rest of the economy), because the relative demand for skilled labor rises while the relative demand for capital goods falls.

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*, 2010, 25, 755–773.
- 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.
- Dedola, Luca and Stefano Neri**, “What does a technology shock do? A VAR analysis with model-based sign restrictions,” *Journal of Monetary Economics*, 2007, 54, 512–549.
- DiCecio, Ricardo**, “Sticky wages and sectoral labor comovement,” *Journal of Economic Dynamics and Control*, 2009, 33, 538–553.

- Erceg, Christopher J., Luca Guerrieri, and Christopher Gust**, “Can long-run restrictions identify technology shocks?,” *Journal of the European Economic Association*, 2005, 3 (6), 1237–1278.
- Fernald, John G.**, “Trend breaks, long-run restrictions, and contractionary technology improvements,” *Journal of Monetary Economics*, 2007, 54, 2467–2485.
- Fernandez-Villaverde, Jesus, Juan F. Rubio-Ramirez, and Thomas J. Sargent**, “A,B,C’s (and D’s) for Understanding VARs,” *The American Economic Review*, 2007, 97 (3), 1021–1026.
- Fisher, Jonas D.M.**, “The Dynamics Effects of Neutral and Investment-Specific Technology Shocks,” *Journal of Political Economy*, 2006, 114 (3), 413–451.
- Francis, Neville and Valerie A. Ramey**, “Is the technology-driven real business cycle theory dead? Shocks and aggregate fluctuations revisited,” *Journal of Monetary Economics*, 2005, 52 (8), 1379–1399.
- 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.
- Noble, Thomas F.X., William Cohen, Barry S. Strauss, Duane J. Osheim, and Kristen Neuschel**, *Western civilization, the continuing experiment*, Boston, Massachusetts: Cengage Learning, 2002.
- Philippon, Thomas and Ariel Reshef**, “Wages and Human Capital in the U.S. Financial Industry: 1909-2006,” Working Paper, 2010.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker**, “Measuring the Cyclical-ity 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.
- , “What are the effects of monetary policy on output? Results from an agnostic identification procedure,” *Journal of Monetary Economics*, 2005, *52*, 381–419.

Table 1: Business cycle moments

	Std	Correlation with			
		Output	Hours	Productivity	Price
<i>Baseline measure</i>					
Skill premium	.0066	0.1131	0.0114	0.1763	-0.1535
Relative hours	.0183	-0.4124*	-0.2917*	-0.2591*	0.5838*
<i>Naive measure</i>					
Skill premium	.0071	0.0962	0.1989*	-0.1853	0.0852
Relative hours	.0161	-0.4242*	-0.3533*	-0.1688	0.5418*
Relative supply	.0111	-0.0220	0.0400	-0.0761	0.2756*

Notes: Data series are constructed as explained in section 2.3 and seasonally adjusted using X-12-ARIMA. The series are HP-filtered with $\lambda=1600$. The * indicates significance of at least 5%.

Table 2: Variance decomposition with skill-biased and other technology shocks

Horizon	8	16	32
<i>output</i>			
SBT shock	3.09 (0.5,15.8)	3.09 (0.4,16.2)	3.09 (0.4,16.3)
other T shock	8.29 (0.8,29.1)	7.51 (0.8,27.8)	7.38 (0.7,27.5)
supply shock	23.94 (7.2,6.8)	23.09 (6.8,44.3)	22.84 (6.6,44.2)
<i>total hours</i>			
SBT shock	14.11 (2.3,34.6)	12.83 (1.9,32.9)	12.24 (1.8,32.5)
other T shock	5.83 (0.9,19.8)	5.41 (0.7,18.8)	5.32 (0.6,18.6)
supply shock	19.62 (5.2,39.5)	19.50 (5.1,39.8)	19.44 (5.0,39.8)
<i>premium</i>			
SBT shock	24.69 (3.6,71.6)	22.38 (2.3,69.2)	20.81 (1.5,67.7)
other T shock	17.80 (1.8,58.8)	19.47 (1.6,60.4)	20.09 (1.3,61.9)
supply shock	32.93 (3.1,66.0)	36.32 (3.3,70.6)	38.29 (3.2,73.1)
<i>relative hours</i>			
SBT shock	10.82 (1.8,44.3)	10.35 (1.3,45.0)	10.20 (1.0,46.5)
other shock	32.83 (3.5,73.5)	34.17 (3.3,77.7)	34.26 (3.0,79.9)
supply shock	32.02 (4.9,74.3)	35.07 (5.4,79.1)	36.77 (5.5,81.3)

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.

Table 3: Robustness of the response of hours to skill-biased and other technology shocks

	SBT shock	other technology shocks
Baseline specification		
with supply shocks	-, insignificant -, significant	-, sign. before 5th quarter -, insignificant
Variation of the baseline specification with supply shocks		
Taking into account low frequency variation in hours		
dummy ¹	-, significant	- on impact, insign.
low-pass trend	-, significant	- on impact, insign.
polyn. trend	-, significant	- on impact, insign.
hours in levels	-, insignificant	+, significant
subsample stability		
1979:I-2000:IV	-, significant	+, sign. after 5th quarter
Minnesota prior with 8 lags changed to		
2 lags	-, significant	- on impact, insign.
4 lags	-, significant	-, insignificant
12 lags	-, significant	-, significant
weaker prior ²	-, significant	-, insignificant
Flat prior (OLS equivalent)		
2 lags	-, significant	-, insignificant
4 lags	-, insignificant	-, insignificant
Alternative and additional variables		
CPS hours	-, sign. on impact	-, insignificant
total hours	-, significant	-, insignificant
Naive wage premium	-, significant	-, insignificant
Baseline + invest. price	-, significant	-, insignificant
the above + investment	-, significant	- on impact, insign.
the above + consumption	-, significant	- on impact, insign.
the above + interest rate	-, sign. on impact	-, significant

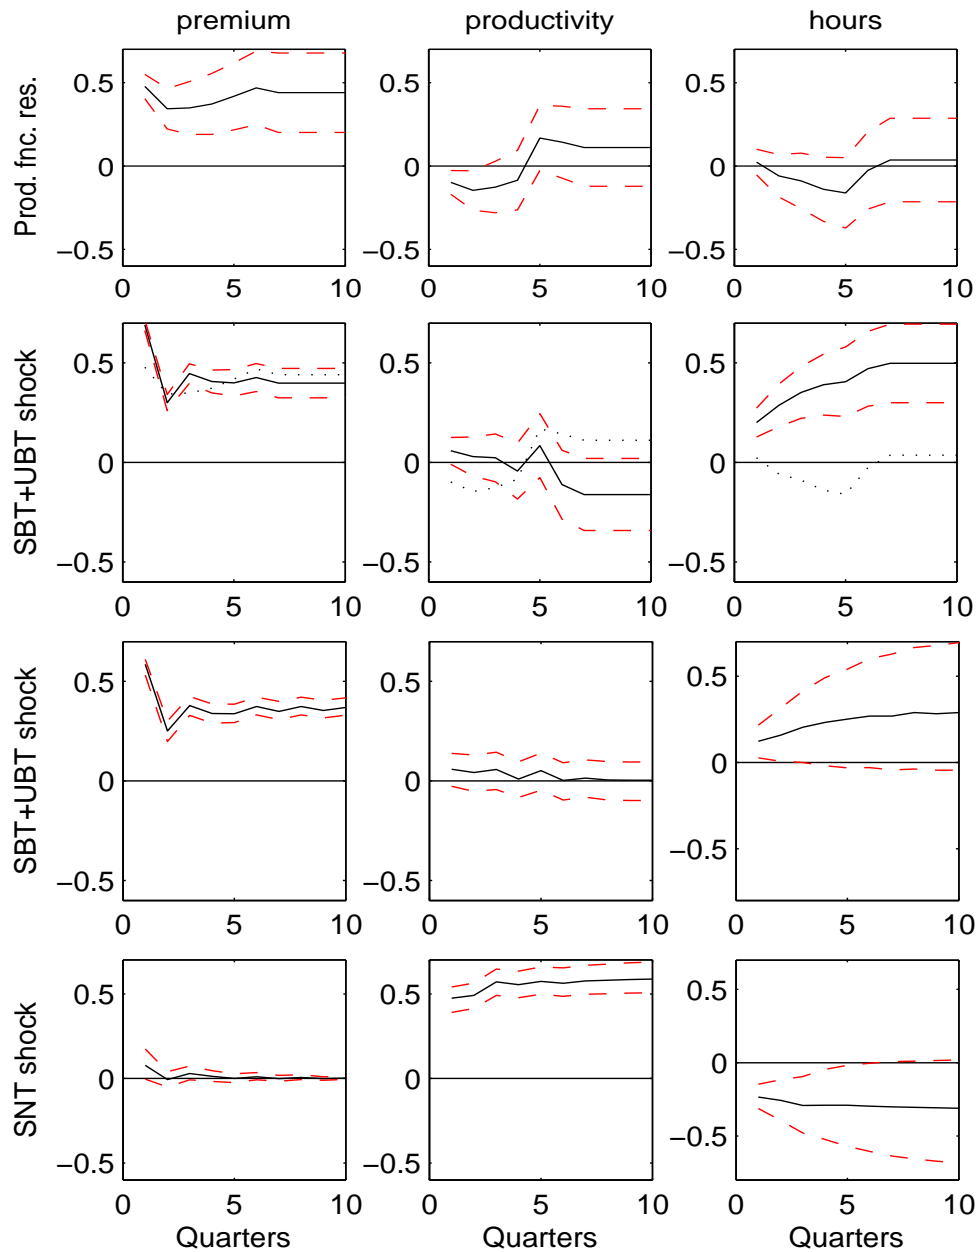
Notes: 1) dummy break at 1997:I; 2) Decay parameter $d = 1$ instead of $d = 3$ as in the baseline.

Table 4: Variance decomposition with investment-biased, skill-biased and other technology shocks

Horizon	8	16	32
<i>output</i>			
IBT shock	41.57 (21.6,57.5)	33.94 (15.9,51.4)	28.97 (10.6,47.6)
SBT shock	2.8 (0.4,11.1)	3.4 (0.5,13.9)	3.8 (0.5,15.9)
other T shock	2.0 (0.4,7.5)	2.3 (2.3,8.6)	2.5 (0.3,9.6)
supply shock	8.09 (2.0,19.3)	8.22 (1.9,20.6)	8.75 (1.9,22.3)
<i>total hours</i>			
IBT shock	11.66 (2.4,30.2)	8.26 (2.9,18.7)	6.85 (2.4,17.5)
SBT shock	14.2 (1.7,34.3)	14.6 (1.8,34.6)	14.9 (1.8,35.2)
other T shock	8.6 (1.2,25.2)	8.0 (1.0,24.7)	7.9 (0.9,25.1)
supply shock	10.50 (1.5,29.1)	10.46 (1.4,29.2)	10.54 (1.4,29.8)
<i>premium</i>			
IBT shock	2.41 (0.9,7.3)	4.93 (1.3,14.3)	8.00 (1.1,19.3)
SBT shock	37.37 (8.44,81.55)	32.7 (6.10,77.70)	27.8 (4.29,74.80)
other T shock	9.34 (0.95,50.55)	9.97 (0.87,51.59)	10.06 (0.69,51.69)
supply shock	27.52 (2.0,55.5)	29.72 (2.3,56.5)	31.01 (2.5,56.9)
<i>relative hours</i>			
IBT shock	38.15 (26.48,50.3)	48.84 (33.02,59.5)	50.72 (33.52,63.0)
SBT shock	8.04 (1.18,33.3)	7.52 (1.01,29.7)	7.52 (0.98,29.5)
other T shock	15.02 (1.96,42.2)	12.04 (1.43,36.9)	10.58 (1.16,35.5)
supply shock	18.22 (3.16,43.1)	16.56 (2.96,38.9)	16.14 (2.93,38.1)

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

Figure 1: Impulse-responses to technology shocks from the production function decomposition and from a VAR with long-run restrictions

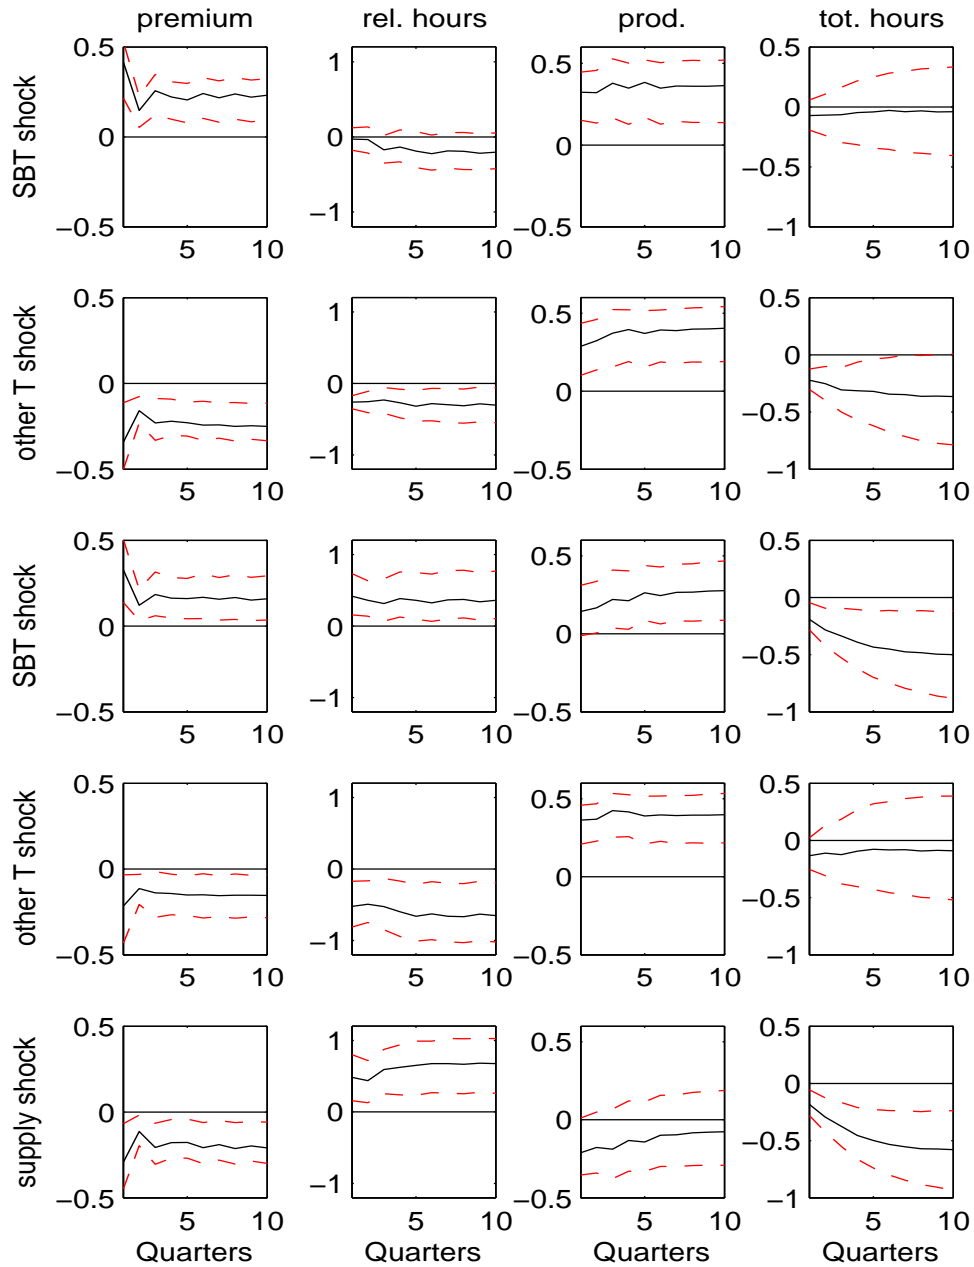


Notes: All responses are in percent to a positive one-standard-deviation shock.

The first two rows show impulse-responses from regressing the variables on six lags of the production function residual and the SBT and UBT shock from the SVAR. The black dotted line repeats the estimate from the first row. Confidence intervals are one standard error bands.

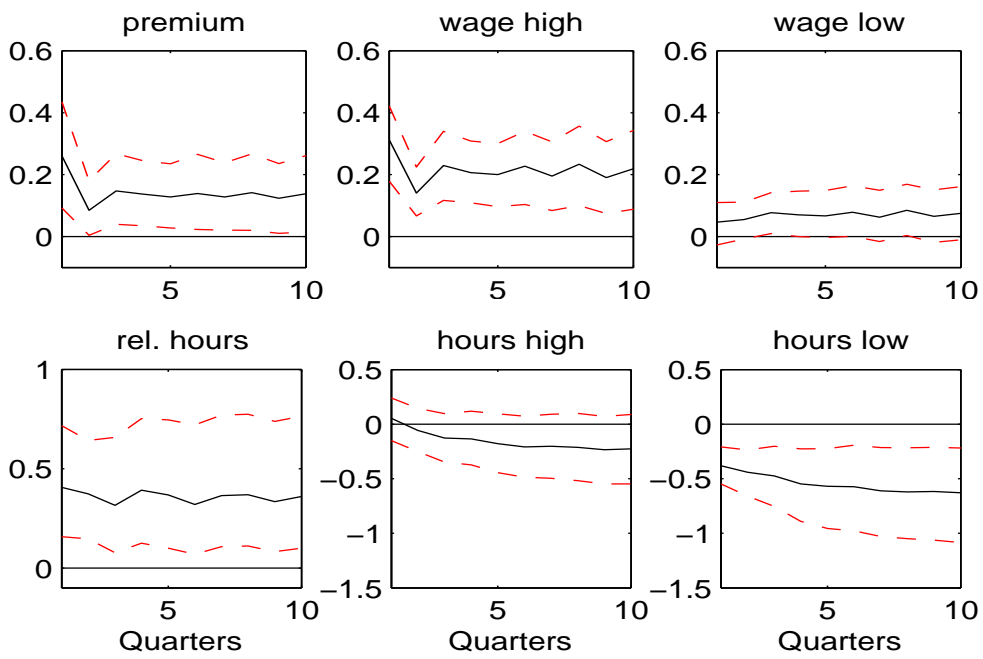
The third and fourth row show the responses to SBT and UBT as well as SNT shocks estimated within the SVAR. Here, confidence intervals are 68% Bayesian bands.

Figure 2: Impulse-responses to technology shocks from a VAR with long-run zero and sign restrictions



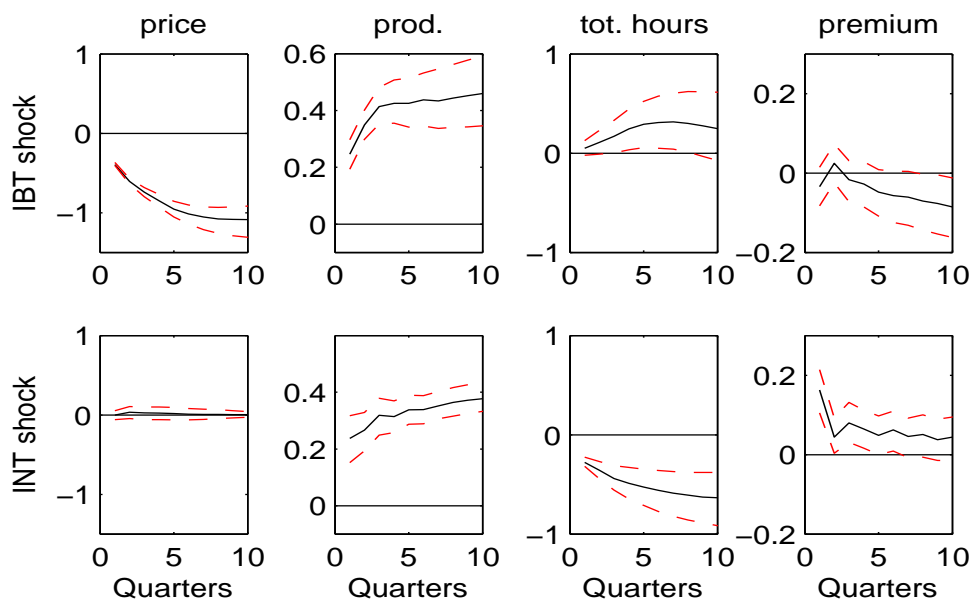
Notes: Percent responses to a positive one-standard-deviation shock. Confidence intervals are 68% Bayesian bands. The first two rows show the results from the identification without supply shocks, the last three rows from the identification with supply shocks

Figure 3: Impulse-responses to skill-biased technology shocks for additional variables



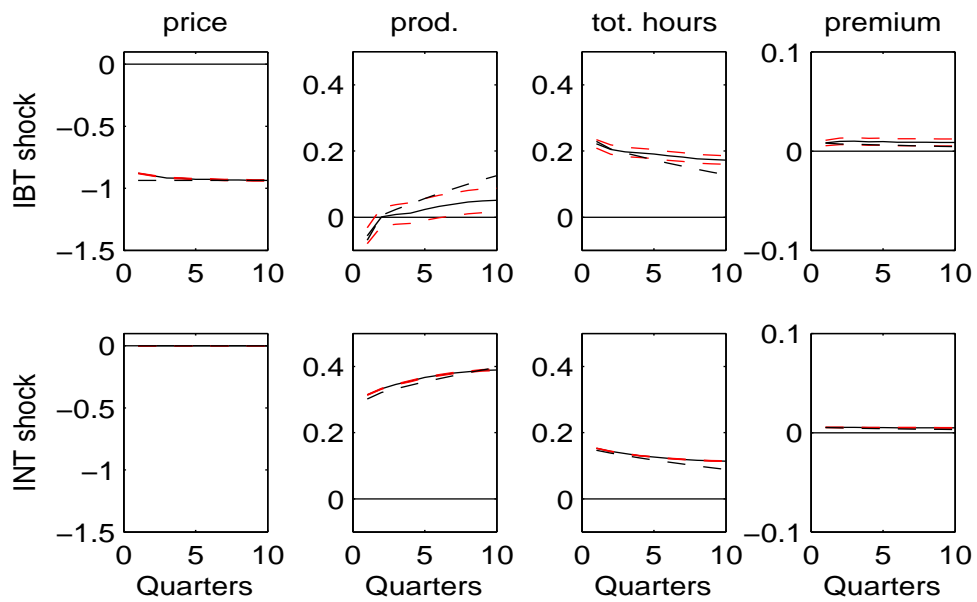
Notes: Percent response to a positive one-standard-deviation shock. Confidence intervals are 68% Bayesian bands.

Figure 4: Impulse-responses to investment-biased technology shocks



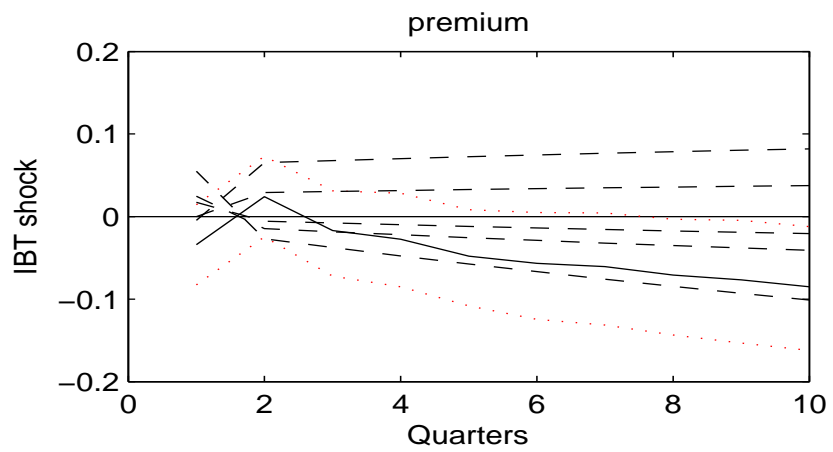
Notes: Percent responses to a positive one-standard-deviation shock. Confidence intervals are 68% Bayesian bands.

Figure 5: Impulse-responses from the model



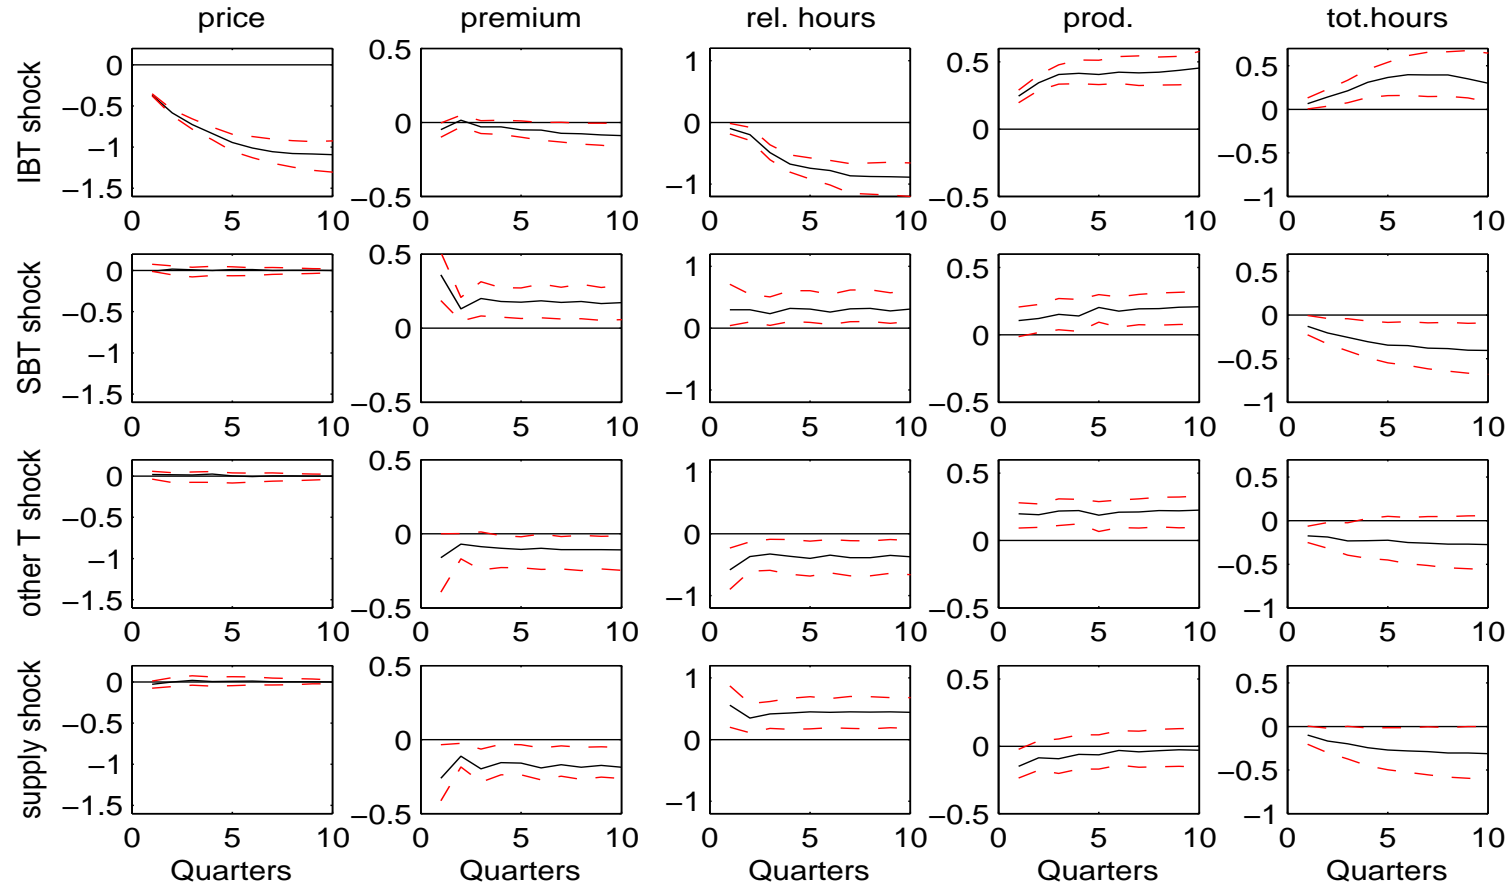
Notes: Percent responses to a positive one-standard-deviation shock. The dashed lines represent the theoretical responses from the model with $\rho = \sigma = 1.67$. The solid lines are the estimated responses from 1000 simulations of 110 quarters each of the same model. The responses are normalized to match the responses of the investment price and labor productivity in the actual data in the longer run (20 quarters).

Figure 6: 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.

Figure 7: Impulse-responses to investment- and skill-biased and other technology shocks



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