Skill-Biased Technological Change
and the Business Cycle

Technical Appendix

Not for Publication

May 2012

A Construction of the skill premium

Our measure for the skill premium is the log wage differential between college graduates and high school graduates. The relative hours worked 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.[1] The resulting series are adjusted for seasonality using the X-12-ARIMA algorithm of the Census Bureau.

[1]For the skill premium and relative hours 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.

The way we measure the skill premium and the relative hours 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 hours worked or supply of skill. Suppose, for example, that there is an increase in the number of workers with a masters degree. This represents an increase in the supply of skill. However, a naive measure of the relative supply, which just counts the number of workers with at least a college degree, would not reflect this increase. Moreover, if workers with a masters degree earn on average higher wages than workers with a bachelors degree only, then a naive measure of the skill premium would increase. In our measures, this increase in the supply of skill would leave the skill premium unchanged and increase the relative supply measure. Our data show a pronounced increase in the skill premium since 1980, which seems to slow down mildly towards the end of the 1990s, as documented in previous studies, e.g. Autor et al. (2005). The fluctuations in our measure of the skill premium are similar to those in the Mincerian return to schooling.

B Time-series properties of the variables

This section documents univariate and multivariate time-series properties of the variables used in the baseline and extended specifications in the paper. These time series include the baseline and naive measures of the skill premium and the relative hours worked of skill as well as the relative supply of skill that were constructed for the paper. It also considers other macroeconomic variables that are employed in the specifications such as hours, output, productivity and the relative price of investment.

In order to support our specification in which all variables enter in first differences in the VAR, we check a few issues of concern in this section. First, do the variables exhibit enough autocorrelation to justify the VAR approach? Table 1 supports that this is the case. It shows the partial correlations the series in first differences as in the baseline specification in the paper. Second, is there still interesting autocorrelation in the first differences, even when we control for measurement error via a moving-average component? Table 2 reports ARIMA(4,1) models for the series used in the estimation. The number of AR-lags reported corresponds to best model fit and significance across the various series. The table shows that the MA-component matters for the skill premium and the relative price of investment and that, controlling for this MA component, there is still autocorrelation and hence high frequency dynamics present.

Next, Table 3 examines the unit root properties of the variables employed in the estimation based on a Dickey-Fuller tests which test for the null hypothesis of a unit root and a KPSS with the null hypothesis of stationarity. For all variables, the Dickey-Fuller test does not reject a unit root in levels and rejects stationarity in levels. Note that we also checked for trend stationarity here. Likewise, the KPSS test rejects the null of stationarity in level. For hours, output and productivity the KPSS test does not deliver clear results on stationarity (this has been noted in the hours puzzle debate already). For all variables, both tests support stationarity in first differences in all variables (the investment price showing some exceptions) which supports our specification of the VAR in first differences.

Finally, as all of the variables of interest are integrated, cointegration should be checked in order to justify the VAR in first differences compared to a vector-error correction specification. Table 4 reports the results of pairwise Johansen cointegration tests between the skill premium, skill supply, relative hours worked by skilled workers, the relative price of investment, hours, output and labor productivity. The figures in the table exhibit the rank of the cointegrating vector and hence the number of cointegration relationships. Here, the test is carried out with a constant both in the cointegration relationship and the remaining VAR, referring to μ and γ in the following equation, and with a trend ρt in the cointegration relationship:

$$\triangle y_t = \alpha(\beta y_{t-1} + \mu + \rho t) + \sum_{i=0}^{p} \triangle y_{t-i} + \gamma + \varepsilon.$$

The only incidents of potential cointegration appear in the case of relative hours worked with the investment price and the relative supply of skill. Since the stationarity tests for the relative hours worked and the relative price of investment did not deliver clear results, it is possible that the cointegration relationship between these two variables detects stationarity in one variable rather than a true cointegration relationship.

Apart from the results on the pairwise cointegration relationships shown in the table, the VAR specification with the different combinations of the variables should be checked for cointegration. It can be shown that for the baseline specification with the skill premium, productivity and hours worked and relative hours worked, the Johansen trace test indicates a rank of zero for the cointegration vector. When including the relative price of investment, the rank of this vector is equal to one.

Table 1: Autocorrelation of data series

Series	Partial autocorrelation with lag							
	1	2	3	4	5	6	7	8
Premium baseline	-0.499*	-0.093	-0.040	-0.205*	0.022	0.004	0.162	-0.070
Premium naive	-0.531*	-0.209*	-0.133	-0.181	0.076	0.015	0.160	0.029
Rel. empl. baseline	-0.232*	-0.056	0.069	0.042	-0.089	0.117	0.093	-0.329*
Rel. empl. naive	-0.333*	-0.055	0.123	0.127	-0.114	0.073	0.099	-0.268*
Rel. supply	-0.415*	-0.235*	0.049	0.160	0.012	0.125	0.205*	-0.362*
Price	0.583*	0.053	0.114	0.019	-0.209*	0.063	-0.285*	-0.055
Hours	0.580*	0.170	0.004	-0.082	0.080	-0.023*	-0.017	-0.199
Output	0.335*	0.142	-0.002	0.004	-0.153	0.069	0.005	-0.086
Productivity	0.019	0.147	-0.111	-0.018	-0.068	-0.007	0.107*	-0.043

Notes: Series used are seasonally adjusted series and in log differences. Here, * denotes significance on 5%.

Table 2: ARIMA models

Series	constant	AR lags				MA lags
		1	2	3	4	1
Premium baseline	0.002*	-1.358*	-0.572*	-0.212	-0.176	0.872*
Rel. empl. baseline	0.006*	-0.677	-0.139	0.067	0.095	0.436
Rel. supply	0.006*	-0.098	0.064	0.251*	0.135	-0.443
Price	-0.015*	-0.252	0.429*	0.028	0.234*	0.846*
Hours	0.000	-0.165	0.487*	0.157	-0.091	0.652
Output	0.007*	-0.295	0.298	0.059	0.063	0.584
Productivity	0.004*	0.394	0.139	-0.162	0.008	-0.366

Notes: Here, * denotes significance on 5%.

Table 3: Unit root tests

Series	Dickey-Fuller Test H_0: unit root	KPSS Test H_0: stationarity
Levels		
Premium baseline	not rejected	rejected
Rel. empl. baseline	not rejected*	rejected
Rel. supply	not rejected	rejected
Price	not rejected	rejected
Hours	not rejected	rejected
Output	not rejected	rejected**
Productivity	not rejected	rejected
First Differences		
Premium baseline	rejected	not rejected
Rel. empl. baseline	rejected	not rejected
Rel. supply	rejected	not rejected
Price	rejected	no clear result
Hours	rejected	not rejected
Output	rejected	not rejected
Productivity	rejected	not rejected

Notes: Unit root test for levels done with and without trend (except for hours). Unit root test for first differences done with and without constant. Results based on 1% significance. * DF rejected at 10 %. ** KPSS rejected with constant, with trend rejected on 5% only for less than one lag.

Table 4: Pairwise Johansen cointegration tests

Series	Rank of cointegration vector		
	Premium baseline	Rel. empl. baseline	Rel. supply
Premium baseline	-	0	0
Rel. empl. baseline	0	-	1
Price	0*	1	0
Hours	0	0	0
Output	0	0**	0
Productivity	0	0	0

Notes: Here, * denotes that a rank of zero resulted from the Johansen test with a constant, but rank of one when a trend was included. ** denotes a rank of one when a constant was included and a rank of zero when applying a trend in the cointegration test.

C Identification and Estimation

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 and shocks by normalizing the variance of all structural shocks to one, assuming orthogonality between these shocks and imposing the identifying restrictions. Without sign restrictions, the long-run identifying restrictions are incorporated using a Cholesky decomposition of the infinite horizon forecast error variance. Adding a long-run sign restriction in this setup means that we have to take away a long-run zero restriction. Hence, the forecast variance is no longer lower triangular, since it has one zero less. In order to implement this, we first impose the long-run zero restrictions and then apply the sign restrictions to the variables of interest, here the skill premium, relative hours of skilled workers and labor productivity.[2]

C.1 Standard Long-Run Identification

Identification involves finding a mapping A of the residuals from a reduced form VAR into so-called structural residuals such that these can be interpreted as technology shocks. More precisely, name v_t the residuals from a reduced form VAR with $E[v_t v_t'] = \Omega$. The relationship between the n structural and reduced form residuals is then $e_t = Av_t$ which induces $A\Sigma_e A' = \Omega$. The remaining assumptions in order to pin down A then need to come from restrictions on the matrix of long-run effects. These assumptions can be incorporated as zero restrictions in the matrix of long-run effects $C \equiv \sum_{i=0}^{\infty} \Phi_i A$, where Φ_i are the impulse-response coefficients.

For identification, all identified shocks are assumed to be orthogonal and the variance of the structural residuals is normalized such that $\Sigma_e = I$. We use a lower triangular structure of the matrix C in order to impose our identifying assumptions for the technology shocks. This is easily obtained by decomposing the variance of the ∞-step ahead

[2]This involves rotating and checking the sign restriction in a subsystem of the long-run horizon forecast revision matrix. This works in analogy to the implementation of short-run sign restrictions, see e.g. Peersman (2005), in which case the variance covariance matrix of the residuals (short-run revision matrix) is rotated. Please refer to the Technical Appendix for details on the estimation and identification.

forecast error $\eta_{t,\infty} = X_{t+\infty} - E_t(X_{t+\infty})$ which is equal to

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

with the Cholesky decomposition. We vary the order of the variables in the VAR in order to impose our different identifying restrictions, i.e. order labor productivity first in the Galí identification, order the skill premium first before labor productivity when distinguishing skill-biased and skill-neutral technology shocks and ordering the investment price first before labor productivity when separately identifying investment-specific and investment-neutral technology shocks. The respective variables that the restrictions are imposed upon need to enter in first differences in the VAR for the identification to work[3]. It has to be noted that this procedure uniquely pins down the effect of the identified technology shocks on all variables in the VAR. It can be shown that the result is not affected by the additional (unnecessary) zero restrictions in the matrix of long-run effects. In that sense, this approach is the same as imposing the identifying restrictions within an instrumental-variable setup as suggested by Shapiro and Watson (1988). Francis et al. (2003) document this for the Galí identification.

C.2 Combination of long-run zero and sign restrictions

In order to separate skill-biased from other technology shocks, we combine our previous long-run zero restrictions with sign restrictions on the long-run effects of these two shocks on the skill premium and labor productivity. As before, we assume that $\Sigma_e = I$ and that the reduced form residuals map into structural residuals such that $A\Sigma_e A' = \Omega$. As in section B.1, imposing our long-run restrictions is then equivalent to finding a decomposition L of the long-run forecast revision variance such that $LL' = MSE(\infty)$.

Assume that the skill premium and labor productivity are ordered first in the VAR (This is not necessary, but convenient for comparison with the Cholesky decomposition in section B.1). First, we impose the long-run zero restrictions, i.e. only the two types of technology shocks can affect the skill premium and labor productivity in the long run. This means that $l_{13} = l_{14} = ... = l_{1n} = 0$ and $l_{23} = l_{24} = ... = l_{2n} = 0$ and results in

$$L_{1:2,1:2}L'_{1:2,1:2} = MSE(\infty)_{1:2,1:2}.$$

Next, we implement sign restrictions on this upper left 2-by-2 system in a similar fashion

[3] This can easily be extended to a specification in levels.

as in Peersman (2005). This involves a rotation of $L_{1:2,1:2} = TQ$, using the Cholesky decomposition $TT' = L_{1:2,1:2}$ and an orthonormal matrix Q (i.e. $QQ' = I$):

$$Q = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix},$$

where $\theta \in [0, \pi]$. As in Peersman (and similar to Uhlig (2005)), our VAR is estimated in a Bayesian framework. For each draw of the posterior distribution of the reduced form VAR coefficients, we calculate the long-run forecast revision variance. We then choose θ from a grid of $[0, \pi]$ and use T and Q to calculate the rotation of the upper left elements of the matrix L. If the sign restrictions are satisfied, we keep the draw.

After having implemented the restrictions, we can now proceed to calculate the remaining elements of the matrix L such that this matrix provides a valid decomposition of the long-run variance. For the remaining elements of the first two columns, we use that $L_{3:n,1:2}L'_{1:2,1:2} = MSE(\infty)_{3:n,1:2}$. Last, we can determine the lower right block of L. Note that there are no restrictions imposed on these elements and that they are not related to the shocks of interest. First, we use the information on the first two rows and columns to calculate what is missing to explain the lower right elements of the long-run variance using LL'. This 'remaining' lower right block of the variance is then decomposed using the Cholesky decomposition. Having found all elements of L, we then determine the matrix A via $A = (\sum_{i=0}^{\infty} \Phi_i)^{-1} L$.

Identification of skill supply shocks in addition to skill-biased and other technology shocks adds one more variable, relative hours worked, and one more sign restriction to the system. Here, the relative hours, the skill premium and labor productivity are ordered first in the system. As before, we then impose the long-run zero restrictions, i.e. only skill supply shocks and the two types of technology shocks can affect the relative hours, the skill premium and productivity. This means that $l_{13} = l_{14} = ... = l_{1n} = 0$, $l_{23} = l_{24} = ... = l_{2n} = 0$ and $l_{33} = l_{34} = ... = l_{3n} = 0$ and results in $L_{1:3,1:3}L'_{1:3,1:3} = MSE(\infty)_{1:3,1:3}$. Rotating this 3-by-3 system now involves three rotation matrices:

$$Q_{12} = \begin{pmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{pmatrix}, Q_{13} = \begin{pmatrix} \cos(\psi) & 0 & -\sin(\psi) \\ 0 & 1 & 0 \\ \sin(\psi) & 0 & \cos(\psi) \end{pmatrix}, Q_{23} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\varphi) & -\sin(\varphi) \\ 0 & \sin(\varphi) & \cos(\varphi) \end{pmatrix},$$

with $\theta \in [0, \pi]$, $\psi \in [0, \pi]$ and $\varphi \in [0, \pi]$ and $Q = Q_{12}Q_{13}Q_{23}$. As before, we rotate $L_{1:3,1:3}$ and check whether the sign restrictions are satisfied. We then calculate the remaining elements of L equivalent to above.

In the application, we draw 1000 candidates from the posterior distribution of the reduced from coefficients and divide $[0, \pi]$ into a grid of 100 values of θ for the implementation of the single sign restriction and into a grid of $10^3 = 1000$ values for θ, ψ and φ when identifying supply shocks along with the technology shocks. We compute the impulse responses for all draws that satisfy the sign restrictions and report the median and the 16th and 84th percentile from the resulting distribution. It is straightforward to extend this identification scheme to incorporate more long-run restrictions, e.g. identify investment-specific shocks using long-run restrictions. For this, order the relative price of investment first, and identify investment-specific shocks. Order the skill premium and labor productivity second and third or the relative hours, the skill premium and productivity second to fourth and impose the identification strategy described above on the remaining elements.

C.3 Estimation of the BVAR

Before identification, the reduced form VAR is estimated in a Bayesian framework. More precisely, we obtain 1000 draws of the posterior distribution of the reduced form coefficients and then apply the identification procedure to each of these in order to produce draws of the distribution of the structural coefficients.[4] As is standard in the literature reporting impulse-responses from Bayesian VARs, the point estimates exhibited correspond to the median and the confidence intervals to the 16th and 84th percentiles of the posterior distribution. As these refer to credible sets rather than confidence intervals in the classical sense, note that these cannot directly be compared to one or two standard error bands in the classical sense.

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

The prior variance $\Sigma_{d_0} = \Sigma_{d_0}(\phi)$ of the coefficients depends on three hyper-parameters ϕ_1, ϕ_2 and ϕ_3, that determine the tightness and decay on own lags, other lags and exoge-

[4]This approach goes back to Canova (1991) and Gordon and Leeper (1994) and is feasible if the system is just identified, that is, if there exists a unique mapping between draws of the residual variance covariance matrix and draws of the identification matrix A.

nous variables. Except for the decay, a loose prior is chosen for the hyper-parameters, namely $\phi_1 = 0.2$, $\phi_2 = 0.5$ and $\phi_3 = 10^5$. The decay parameter $d = 3$. The advantage of the structure of the Minnesota prior is exactly this ability to separately deal with the lags of the variables, i.e. own and other lags, as well as exogenous variables. Together with a normal likelihood of the data the Minnesota prior produces a posterior that can be derived analytically. Hence, the estimation does not rely on sampling procedures.

D Long sample results and identifying restrictions

In an earlier version of this paper, we identified skill-biased technology shocks as all shocks that affect the skill premium in the long run, and estimated the responses of these shocks over the 1979:1 to 2000:4 period. In this short sample, the skill premium significantly increases after Galí shocks and hours worked fall after SBT shocks. In this version, we extend the sample period to 1979:1 to 2006:4. Adding 6 years of data significantly changes our results, which are no longer robust over subsamples. This problem is due to a problem with the simple identifying long-run restriction, which does not matter in the short sample but is apparent when extending the sample period.

To understand what drives identification, we plot the first differences of the three variables in our VAR (skill premium, productivity and hours) over time. To facilitate eyeballing this graph, we smooth these first differences using a 9-quarter centered moving average with linearly declining weights. We then plot the identified shocks from our VAR with the long-run zero restriction, similarly smoothed, in the same graph, which is provided in Figure 1.

As is clear from the graph, identified SBT shocks are closely related to changes in the skill premium. Over most of the sample, periods of relatively high growth in the skill premium coincide with periods of relatively high productivity growth, and periods when the skill premium growth is below trend are periods when productivity growth is below trend as well. This explains why we find that SBT shocks increase productivity in the short sample.

Around 2000, however, this relationship breaks down. In the post 2000 period, the rise in the skill premium flattens out, whereas productivity grows strongly. As a result, it seems that SBT shocks estimated on the later part of the sample reduce productivity, calling into question their interpretation as technology shocks. This suggests that there is something wrong about the identification strategy using long-run zero restrictions only.

What type of shock could raise the skill premium permanently but decrease pro-

Figure 1: Skill premium, labor productivity and identified shocks

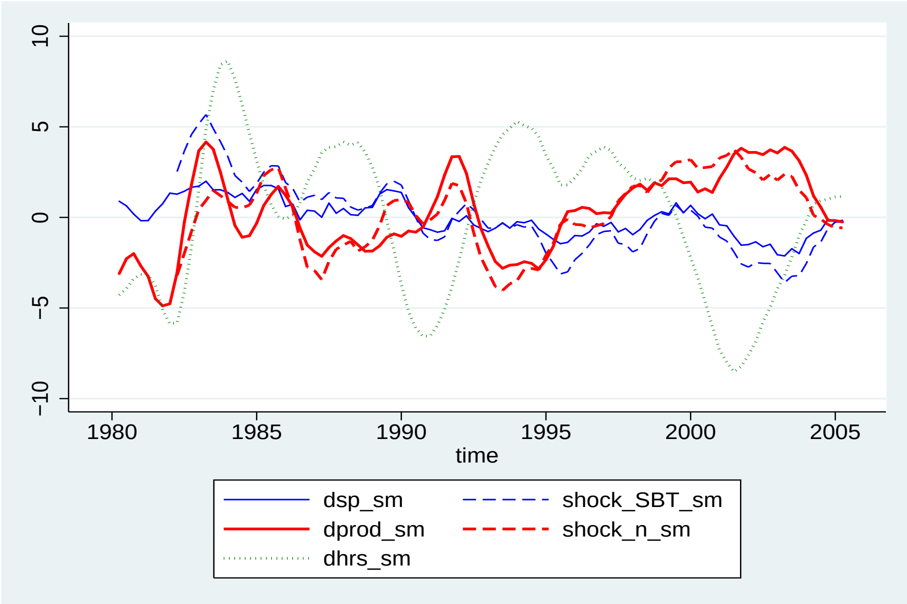

ductivity? In other words: What type of shock raises productivity permanently and decreases the skill premium. If we maintain the assumption that only technology shocks affect productivity in the long run, the obvious candidate is a technology shocks that is biased towards unskilled labor. And of course it is perfectly plausible that such shocks exist in reality. By imposing only the zero restriction that SBT shocks affect the skill premium, we are including technology shocks biased towards unskilled labor as well as towards skilled labor in our estimated shocks. And since SBT and unskill-biased (UBT) shocks that increase the skill premium have opposite effects on productivity (as well as hours worked), the responses of these variables become non-robust and sensitive to the sample period.

The solution to the problem is to identify SBT shocks with a combination of a zero and a sign restriction. SBT shocks are the only shocks that affect the skill premium and affect the premium and labor productivity in the same direction. The zero restriction limits the set of shocks to SBT and UBT shocks. The sign restriction rules out UBT shocks, so that the only shocks that satisfy the combination of both restrictions are SBT shocks. Using this modified restriction, our results are robust to the sample period. This also means that we can no longer use a production function decomposition to replicate the baseline results, because it is not possible to impose the sign restriction using this approach. Therefore, we start with the results from the production function decomposition, show that we can replicate the same results in a VAR with a long-run zero restriction, and then argue why we need the sign restriction and impose that to obtain our baseline results.

E Skill bias in technology, Galí identification

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. Here, we evaluate the skill bias in technology shocks identified in this manner.

Figure 2 presents impulse response functions to technology shocks, identified as in Galí (1999). The first row of responses replicates the estimates in Galí (1999), using data on labor productivity and hours worked over the period 1948:I-1994:IV (postwar sample). These responses are estimated using a VAR with 4 lags and a 'flat' prior with median equal to the OLS point estimate. A positive innovation in technology leads to an almost immediate increase in labor productivity equal to the long-run effect, and a reduction 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 second row shows the responses from the same sample, now estimated with our prior on the declining importance of the higher order lag coefficients, see Section 2 in the paper, and 8 lags as in our baseline specification. Compared to the responses in the first row, the response of hours worked is shifted up slightly in this specification, so that the initial drop in hours is no longer significant. The third row again estimates the same specification, but using data for the 1979:I-2006:II (recent) sample. Over this time period, the fall in hours worked is significant and more persistent than in the postwar sample, although the differences are not significant.

In the fourth row of Figure 2, we add the skill premium and the relative hours worked to the VAR. Introducing this additional regressor leaves the responses of labor productivity and total hours worked virtually unchanged. The skill premium does not respond to a permanent improvement in technology, while relative hours fall significantly. This finding seems inconsistent with the hypothesis of skill-biased technological change. However, since the identifying restriction encompasses both skill-biased and unskill-biased technology shocks and since these two shocks have opposing effects on the skill premium, the skill premium does not react to technology shocks if SBT and UBT shocks are equally important. Table 5 shows that technology shocks only explain about 5% of the business cycle variation of output, but up to 10% of the variation in total and relative hours worked.

Table 5: Variance decomposition with identified technology shocks

Horizon	8	16	32
output			
technology shock	5.85	5.20	5.05
	(0.8, 17.7)	(0.7,17.4)	(0.6,17.6)
total hours			
technology shock	10.29	9.12	8.88
	(1.9, 28.7)	(1.5,27.1)	(1.2,26.6)
premium			
technology shock	1.72	2.17	2.45
	(0.6, 5.1)	(0.6,7.2)	(0.5,8.6)
relative hours			
technology shock	8.94	10.16	10.34
	(2.33,20.2)	(2.26,23.1)	(2.12,24.3)

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

Figure 2: Impulse-responses to technology shocks

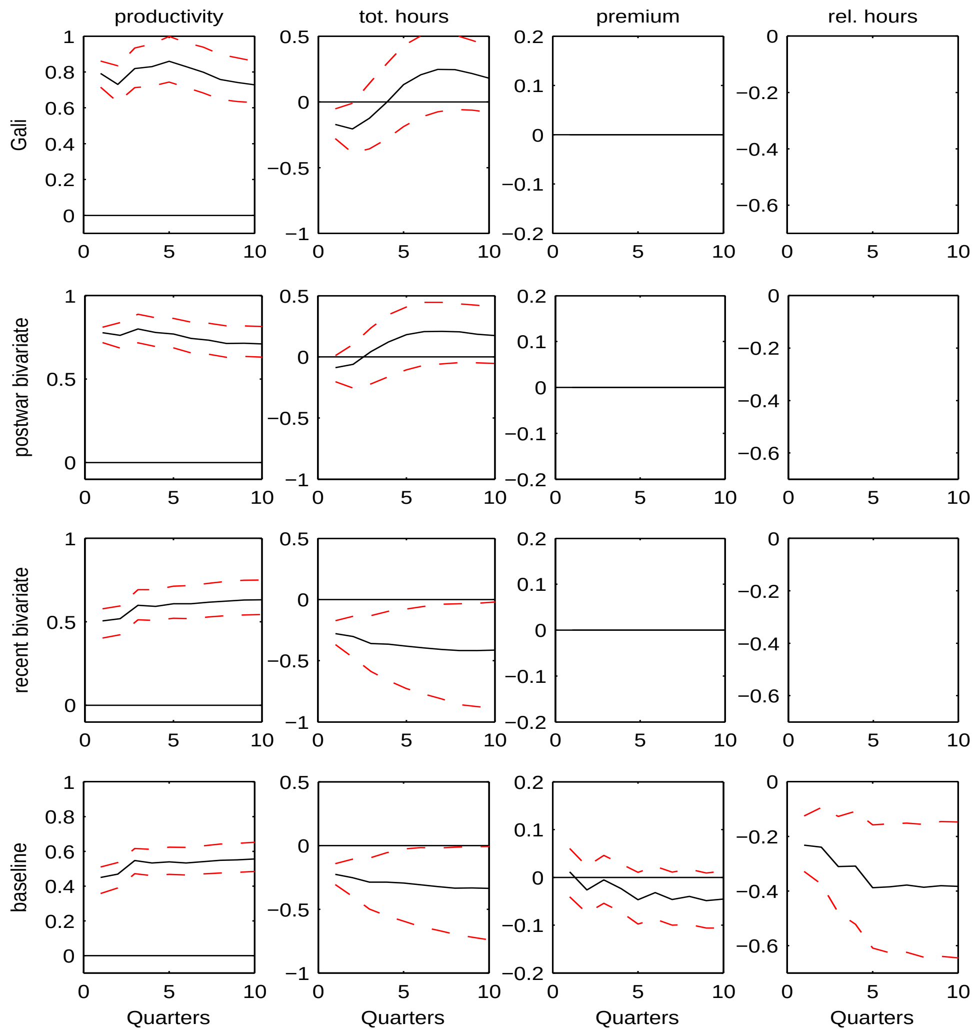

Notes: Responses in percent and quarters to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.
First row: Bivariate VAR with labor productivity and total hours worked, 1948:I-1994:IV, OLS with 4 lags.
Second row: Bivariate VAR, 1948:I-1994:IV, Minnesota prior with 8 lags.
Third row: Bivariate VAR 1979:I-2006:II, with low-frequency adjustment.
Fourth row: VAR with labor productivity, total hours, skill premium and relative hours, 1979:I-2006:II, with low-frequency adjustment.

F Additional Tables and Figures

Figure 3 plots our quarterly series for the log wage premium of college over high school graduates. Our data show a pronounced increase in the skill premium since 1980, which seems to slow down mildly towards the end of the 1990s, as documented in previous studies, e.g. Autor et al. (2005). 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 Mincerian return to schooling, indicating we have adequately controlled for heterogeneity beyond two skill types. Figure 4 shows similar plots for the relative hours worked and the relative supply of skilled labor. Again, there is a substantial difference between our preferred measure and the naive measure of the relative hours 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.

Table 6 shows the business cycle correlation of the skill premium, the relative hours and supply of skill with leads and lags of productivity and output. Here, the series are HP-filtered with $\lambda = 1600$. Note that the skill premium significantly correlated with leads of productivity, but not with output. Relative hours worked of skill are significantly and negatively correlated with output and productivity, while the relative supply of skill is significantly correlated also with leads and lags of productivity, but not output.

Figure 5 shows the responses for skill biased and other technology shocks when including the relative price of investment into the specification. Compared to the original Fisher identification, both investment-neutral and investment-biased technology shocks can now be encompassed by SBT shocks. The Figure documents that the results are virtually the same when including the investment price. Further, both UBT and SBT shocks drive the relative investment price down, but the effect is less strong for SBT shocks. This suggests that the INT and some of the IBT shocks are encompassed by SBT shocks and the rest of the IBT shocks are contained in the UBT shocks now.

Table 6: Unconditional business cycle correlations: Leads and Lags

Series	Correlation with $t + i$ of output and productivity										
	-5	-4	-3	-2	-1	0	1	2	3	4	5
	Output										
Skill premium	-0.0094	-0.0388	0.0599	0.0996	0.0929	0.1131	0.0710	0.0731	0.0260	0.0348	-0.0772
Rel. hours	-0.1420	-0.2427*	-0.3728*	-0.4684*	-0.4961*	-0.4124*	-0.2851*	-0.1294	0.0114	0.1440	0.2217*
Rel. supply	-0.2845*	-0.2250 *	-0.1683*	-0.1788*	-0.1239	-0.0220	0.0723	0.1429	0.2466*	0.3238*	0.3382*
	Productivity										
Skill premium	0.0726	0.0004	0.1341	0.1621	0.1524	0.1763	0.1604*	0.1872*	0.1068	0.1256	-0.0795
Rel. hours	-0.2112*	-0.3477*	-0.3997*	-0.4245*	-0.3734*	-0.2591*	-0.0874	0.0323	0.1584	0.2638*	0.1883*
Rel. supply	-0.0720	-0.0812	-0.0332	-0.0996	-0.0889	-0.0761	0.0057	-0.0190	0.0749	0.1501	0.0765

Notes: Data series are constructed as explained in section 2.3 and seasonally adjusted using X-12-ARIMA. The skill premium and relative hours worked correspond to the baseline series. The series are HP-filtered with λ=1600. The * indicates significance of at least 10%.

Figure 3: Skill premium and Mincer return to schooling in the US

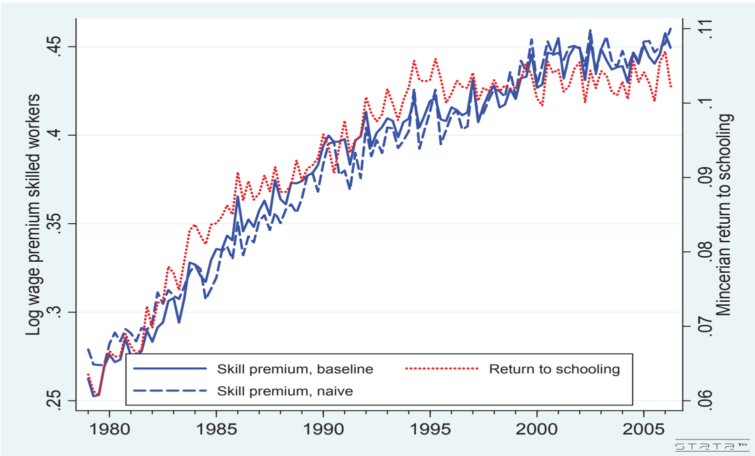

Figure 4: Relative hours worked and relative supply of skill in the US

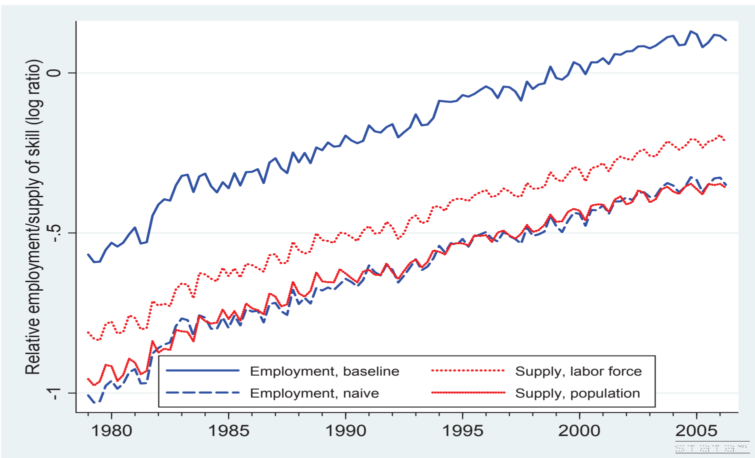

Figure 5: Impulse-responses to skill-biased and other technology shocks including the relative price of investment goods

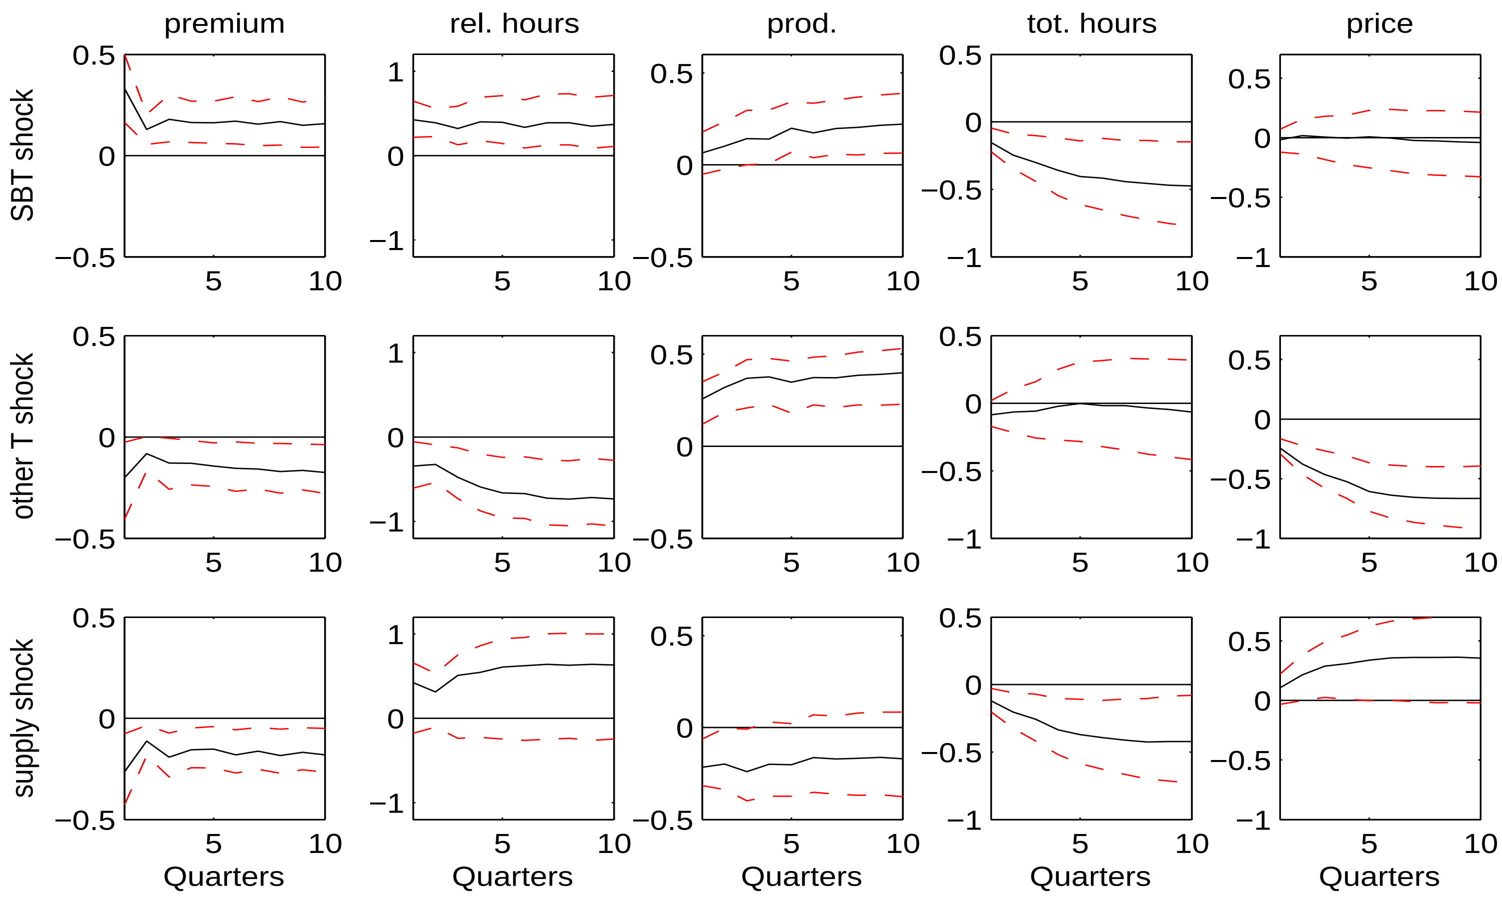

Notes: Percent responses to a positive one-standard-deviation shock. Confidence intervals are 68% Bayesian bands.

References

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney, "Rising Wage Inequality: The Role of Composition and Prices," Discussion Paper 2096, Harvard Institute of Economic Research 2005.

Canova, Fabio, "Source of Financial Crisis: Pre- and Post-Fed Evidence," *International Economic Review*, 1991, *32*, 689–713.

Francis, Neville, Michael T. Owyang, and Athena T. Theodorou, "The use of long-run restrictions for the identification of technology shocks," *The Federal Reserve Bank of St. Louis Review*, 2003, *85* (6), 53–66.

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

Gordon, Robert J. and Eric Leeper, "The Dynamic Impact of Monetary Policy: An Exercise in Tentative Identification," *Journal of Political Economy*, 1994, *102*, 1228–1247.

Peersman, Gert, "What Caused the Early Millenium Slowdown? Evidence Based on Vector Autoregressions," *Journal of Applied Econometrics*, 2005, *20*, 185–207.

Shapiro, M. and M. Watson, "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual*, 1988, pp. 111–148.

Uhlig, Harald, "What are the effects of monetary policy on output? Results from an agnostic identification procedure," *Journal of Monetary Economics*, 2005, *52*, 381–419.