# On the Implications of Unobserved Technology and Preference Shifts for Aggregate Labor Demand and Supply

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
Almut Balleer
aus Göttingen

Bonn 2009

Dekan: Prof. Dr. Christian Hillgruber
Erstreferent: Prof. Monika Merz, PhD
Zweitreferent: Prof. Dr. Jörg Breitung

Tag der mündlichen Prüfung: 28.07.2009

Diese Dissertation ist auf dem Hochschulschriftenserver der ULB Bonn
(http://hss.ulb.uni-bonn.de/diss_online) elektronisch publiziert.

to my parents

# Acknowledgments

I owe much gratitude to my supervisor Monika Merz for her guidance throughout the dissertation, her fruitful advice during many discussions and her extensive and valuable support on the job market. I would also like to thank Thijs van Rens for many useful comments and our pleasant and instructive cooperation. His continuous support during the dissertation and the job market cannot be taken for granted.

I have greatly benefitted from productive comments in- and outside the seminars at the University of Bonn and the Universitat Pompeu Fabra in Barcelona. In particular, I would like to thank Jörg Breitung in Bonn as well as Jordi Galí and Fabio Canova in Barcelona for helpful suggestions and constructive comments. Many thanks also to Jarkko Turunen for the support during my time at the European Central Bank and our joint project. I greatly enjoyed both the academic and recreational interaction with my fellow graduate students. In particular, I would like to thank Michael Evers for interesting and fruitful discussions, his continuous help and a lot of fun. Special thanks also go to Zeno Enders, Andrea Felfe, Katharina Greulich, Stefan Niemann and Markus Poschke.

Financial support from the Bonn Graduate School of Economics, Jörg Breitung and the Marie Curie Fellowship at the Universitat Pompeu Fabra is gratefully acknowledged. Many thanks go also to Urs Schweizer and Jürgen von Hagen for managing the Bonn Graduate School of Economics.

Last but not least, I am deeply indebted to Malte, my friends and my family for their patience, their unconditional emotional support and encouragement as well as their enduring belief in me.

# Contents

# List of Figures

# List of Tables

# Introduction

The macroeconomic approach to the labor market aims at explaining aggregate labor market phenomena that have been present in many industrialized countries such as the United States or the western European countries in the postwar period. These phenomena include the permanent and simultaneous presence of unemployment and unfilled job vacancies as well as the fact that aggregate hours worked, employment and unemployment strongly co-move with the business cycle. At the same time, the amount and composition of labor supplied to the market and employed in production have substantially changed over the last decades. Understanding the driving forces and economic mechanisms that lead to the outcome in the labor market is of high importance to grasp the business-cycle fluctuations as well as the evolution of long-run growth of the economy as a whole.

A large strand in the macroeconomic labor literature builds on the seminal work of Dale Mortensen and Christopher Pissarides (see Mortensen and Pissarides (1994) and Pissarides (2000)) who employ search frictions in the labor market in order to explain the parallel existence of unemployment and unfilled vacancies in equilibrium. In their model, posting a vacancy is costly for firms, and matching in the labor market takes time depending on the tightness of the labor market, i.e., the ratio of unemployed workers seeking employment and open vacancies required to be filled. Firms and workers who meet in the labor market bargain over the wage given their economic conditions such as labor productivity or unemployment benefits.

The dynamic version of the Mortensen-Pissarides model aims at replicating the procyclical fluctuations of employment and countercyclical fluctuations of unemployment respectively. Here, shocks to labor productivity increase the incentive for firms to post vacancies and therefore decrease unemployment. This means that the fluctuations in the labor market are prominently driven by labor demand. Building on this baseline model, a variety of issues have been raised in recent years in order to improve the empirical performance of the model, i.e., its ability to replicate the employment and unemployment fluctuations it aims to explain. For example, the introduction of rigid wages is conjectured to improve the model with respect to the volatility of the key variables (see Shimer (2005a) or Hall (2005)). Another issue of importance is the question whether the data support the assumption of a constant job separation rate, i.e., the probability with which an employed worker is separated from an existing employment relationship (see Fujita and Ramey (forthcoming) or Ramey (2008)).

A related part of the literature incorporates search-and-matching in the labor market as in Mortensen and Pissarides into a general macroeconomic real-business-cycle (RBC) setup (see Merz (1995) and Andolfatto (1996) and also denHaan et al. (2000)). This research was initially motivated to improve upon the empirical performance of the RBC model. In contrast to the standard model, it provides a more sophisticated framework which links the developments in the labor market to the key macroeconomic aggregates such as output, capital investment and consumption. Fluctuations in the basic RBC model are (mainly) driven by shocks to technological progress or more generally aggregate supply. As labor productivity shocks in the more general setup of the Mortensen-Pissarides model are widely interpreted as technology shocks (Shimer (2005a)), the RBC framework indeed provides the natural framework to link the cyclical movements of the labor market variables to the induced driving force of the business-cycle. However, the RBC paradigm about the cyclical movements of the variables and the sources of the business cycle is heavily disputed in the literature (see for example Galí (1999)). In the tradition of this dispute, Chapter 1 challenges the conventional view that cyclical labor market dynamics are mainly driven by technology shocks by highlighting the role of shocks to the aggregate demand (here preference shocks) as a complementary source of fluctuations in employment and unemployment. Based on this analysis, Chapter 1 also contributes to many of the issues raised with respect to the Mortensen-Pissarides model such as its empirical performance, or the endogeneity of job separations.

Another part of the macroeconomic labor literature has focused on the long run developments of labor demand. Also here, technological progress constitutes the main determinant of the labor demand of firms in the medium to long run. A wide literature has documented that this demand has shifted in favor of higher rather than lower educated workers in recent decades, so-called skill-biased technological progress (for example Katz and Murphy (1992) or Autor et al. (1998)). Analyzing these shifts in the composition of inputs to aggregate production via a growth decomposition framework reveals not only the skill-bias of technological progress but also the degree of substitutability or complementarity between high and low skilled labor as well as capital in aggregate production (Krusell et al. (2000)). Building on aspects of the above mentioned RBC literature as well as results and methods of Chapter 1, Chapter 2 investigates the implications of skill-biased technological change for the business cycle. Over and above linking the conventional driving forces of business-cycles and long-run growth to the developments in the labor market, Chapter 2 therefore attempts to identify sources of cycles and growth that originate in the labor market itself.

Apart from labor demand, aggregate labor market supply evolves as a complex result of individual economic choices which are potentially affected by wages, labor market policies such as unemployment benefits or labor taxes, but also preferences for education and skills or social norms and culture. Shifts in the magnitude or composition of labor supply and consequently labor input are an important factor of actual and potential economic growth. In fact, this highlights the importance of the underlying heterogeneity between different workers in the labor force for the aggregate developments. Like labor demand, labor supply has shifted towards a higher ratio of high skilled workers in the recent decades (see again Katz and Murphy (1992) or Autor et al. (1998)). This may be due to the aforementioned increase in skill demand, but is also attributable to an increased preference for education. Together with compositional changes with respect to skill, the age composition of aggregate labor supply has also changed. As more and more young persons tend to stay in education longer, they enter the labor force later in life, but are also better educated. Further, changes in the age composition of the population towards a higher share of older persons also affect the age composition of the labor and work force.

In addition to growth, skill supply shifts have also importantly affected the business-cycle fluctuations of employment as shown in Chapter 2. Apart from this, the labor force in the aforementioned literature (and the first two chapters of this dissertation) was assumed to be constant. However, fluctuations in labor market participation behavior are present in the data, affect the aggregate labor supply and play an important role for the labor market business-cycle fluctuations (for a recent study on this, see Veracierto (2008)). In the presence of labor market frictions as assumed in Chapter 1, a rise in labor market participation initially increases the pool of unemployed and therefore decreases the probability to find a job. Participation may increase over the cycle in response to rising wages, i.e., higher opportunity costs of staying out of the labor force. This effect strongly depends on the underlying wage elasticity of the labor supply and is often associated with the so-called "added worker" effect of persons, often females, in a liaison with a working partner.

In addition to cyclical fluctuations, labor market participation has substantially increased over time, especially for females. This rise in participation has often been associated with cohort effects. Cohort effects generally encompass any factor associated with a particular birth year and constitute unobserved sources of labor supply shifts between different generations of workers. These shifts are often reflected in choices made early on in life such as marriage, maternity leave and/or education and are related to changes in the underlying preferences or cultural factors. Chapter 3 attempts to disentangle shifts in labor market participation that are due to movements of the business cycle or a change in labor market policies from those that are due to unobserved cohort effects.

In general, the subsequent analysis is positive and policy applications will be referred to, but not addressed explicitly. However, as this dissertation aims at achieving a better understanding of the movements of hours worked, employment, unemployment and labor market participation over the business cycle and along their long-run trends, its results are highly relevant for both the conduct and the evaluation of labor market policy, and also fiscal and monetary policy. The following chapter summaries will provide more detailed insights into the particular questions addressed and the results obtained in this dissertation and will also illustrate the differences and connections between the respective chapters.

**Chapter 1**[1]. The standard workhorse model to study business-cycle movements into and out of employment and unemployment is based on a dynamic version of the Mortensen and Pissarides (1994) model with search-and-matching in the labor market. The question whether this model is able to replicate the business-cycle fluctuations in U.S. time series data has been one of the most controversially discussed issues in the recent macro-labor literature. Shimer (2005a), one of the most important contributors to this debate, has documented that the model does not mirror the high volatility of the job finding rate and unemployment that is observed in the data.

Chapter 1 re-addresses the empirical performance of the model by relating the business-cycle fluctuations in the labor market variables to the sources of the cycle. In particular,

[1]This chapter is based on the working paper "New Evidence, Old Puzzles: Technology Shocks and Labor Market Dynamics" (Balleer (2009))

the dynamics in the standard frictional labor market model stem from fluctuations in labor productivity and "a change in labor productivity is most easily interpreted as a technology or supply shock" (Shimer (2005a)). This points towards a representation of the labor market dynamics within a real-business-cycle (RBC) and growth model in which technology shocks are the main driving forces of labor productivity. However, other disturbances such as demand shocks may affect labor productivity as well. In Chapter 1, I disentangle different types of technology (supply) and non-technology (demand) shocks using a structural VAR with long-run restrictions similar to Galí (1999). I then assess the empirical performance of the standard model based on second moments that are conditional on technology shocks rather than on overall unconditional moments.

I document that a baseline model that is driven by technology shocks and is conventionally calibrated performs well to replicate the standard deviations in the key labor market variables that are conditional on technology shocks. This means that much of the criticism that has been uttered with respect to the performance of these models based on overall unconditional moments does in fact not apply. However, the model is not able to replicate the positive correlation between unemployment and labor productivity that is conditional on technology shocks. This result displays a new puzzle in the macro labor literature that is in many ways parallel to the existing "hours puzzle" with respect to the RBC model without labor market frictions (Galí (1999)). In addition, I show that non-technology shocks are necessary to explain both the overall volatility in the labor market and the co-movement of the labor market variables with labor productivity. I test the effect of demand shocks in the form of shocks to the marginal-rate of substitution between consumption and leisure, generally referred to as preference shocks, in the macroeconomic labor market model. I document that these shocks can help to explain the overall volatility in the labor market variables, but are not suitable to analyze the conditional correlations in the data.

Conventional technology shocks commonly referred to as a source of the business cycle may either be factor-neutral or biased towards new investment rather than consumption goods (investment-specific technology shocks). Motivated by observed variation in the data, I propose a new identification scheme to separate different types of technology shocks that are neither neutral nor investment-specific in Chapter 1. I show that these remaining technology shocks play an important role for the business-cycle variance of the labor market variables. While it is difficult to interpret these technology shocks in the theoretical framework of this chapter, when assuming a Cobb-Douglas production technology, Chapter

2 provides a more general context in this respect in order to further explore these shocks.

**Chapter 2**[2]. Like the preceding chapter, Chapter 2 builds on the evidence that technology shocks have been an important driving source of the US business cycle in the last two decades. At the same time, the US as well as many other industrialized countries, have seen a strong parallel increase in the price and quantity of skill. This fact has been taken as evidence for an increase in skill demand driven by technological progress that is biased towards making skilled labor more productive (see Katz and Murphy (1992) or Autor et al. (1998)). This means that newly developed technologies require relatively more educated and fewer uneducated workers. Chapter 2 attempts to relate these two phenomena by exploring the implications of skill-biased technological change for the business cycle.

Existing studies on skill-biased technological change have focused on slow moving trends in annual data. As annual data are not suitable to analyze business cycle fluctuations, we construct a quarterly series for the skill premium and the supply of skilled workers for the US from the Current Population Survey outgoing rotation groups. Measuring neutral and investment-specific technology shocks using long-run restrictions in a structural VAR similar to Chapter 1 then allows us to assess the relationship between technology and the skill premium, and hence skill-bias, over the business cycle. Over and above this assessment, Chapter 2 adds a different angle as it proposes a structural VAR to identify skill-biased, and complementary also skill-neutral, technological change directly. This strategy allows to explore potential sources of the business cycle that originate in the labor market.

We document that cyclical improvements in technology significantly increase the skill premium. This effect is realized in full within a year, providing evidence in favor of skill-biased technological change and its potential importance for business cycle fluctuations. Further, skill-biased and skill-neutral technology shocks have different implications for other aggregate variables. In particular, a positive skill-biased technology shock leads to a much larger reduction in total hours worked than a skill-neutral technology shock and may therefore provide an explanation to the "hours puzzle" already mentioned in Chapter 1. Apart from this, the results from Chapter 2 are instructive with respect to the relationship between the inputs to aggregate production. In particular, high and low skilled labor are not perfect substitutes in production. It is documented that investment-specific technology shocks have a significant negative effect on the skill premium and, conversely,

[2]This chapter draws on joint work with Thijs van Rens (Balleer and van Rens (2008)).

skill-biased technology shocks raise the relative price of investment goods. This evidence directly contradicts the hypothesis of capital-skill complementarity, suggesting instead that capital and skill are in fact substitutes in the aggregate production process.

Chapter 2 highlights the role of non-technology shocks for the business-cycle dynamics in the labor market. These shocks are different in nature than the non-technology shocks investigated in Chapter 1. Here, shifts in the ratio of high to low skilled workers in the aggregate labor supply affect the composition of the pool of workers available to firms. If the two groups of workers differ in their productivity, these shifts affect the composition of aggregate employment, aggregate productivity and the wage premium. This means that shocks to the relative supply of high skilled workers may mistakenly be measured as technology shocks. In Chapter 2, we propose a strategy to separately gauge these skill supply shocks in the structural VAR. This cleans the measure of technology shocks from the potential influence of the skill supply shocks. We furthermore document that skill supply shocks play an important role for the business cycle movements of hours worked as they account for around 30% of the business cycle variance of this variable.

**Chapter 3**[3]. Similar to the US, the euro area labor force participation rate, defined as the ratio between the labor force and the working age population, has increased from below 65% in the early 1980s to 70.9% in 2007. The participation rate of females in the euro area has increased by more than 15 percentage points over this time period, to 63.3% in 2007, compared to the participation rate of 78.6% for males. The large increase in the propensity of the euro area population to work or to search for and be available for jobs has been one of the main driving forces of the substantial increase in euro area labor supply that has accelerated since the mid-1990s. This strong increase has significantly reduced the gap in the use of labor input between the euro area and the United States, and has substantially contributed to output growth and welfare in the euro area. A number of determinants could have factored into this rise in participation, including reforms in the labor market, changes in cultural attitudes towards work (particularly for women), as well as demographic factors. As demographic factors will become less favorable with population ageing increasing in the future, positive participation trends within age and gender groups will be important for sustaining potential growth in the euro area.

[3]This chapter is based on the working paper "Labor Force Participation in the Euro Area: A Cohort-Based Analysis" which is joint work together with Jarkko Turunen and Ramon Gomez-Salvador (Balleer et al. (2009))

Chapter 3 uses a cohort based model of labor force participation to analyze determinants of participation for disaggregated groups of workers in the euro area and the five largest euro area countries (Germany, France, Italy, Spain and the Netherlands). The model is used to decompose the evolution of time-series of age-specific participation rates into the impact of the business cycle, observed structural determinants of participation and other unobserved determinants captured by fixed effects that are specific to ages and cohorts.

Chapter 3 documents that analyzing participation behavior both between (age and gender effects) and within (cohort effects) detailed age and gender groups is particularly useful for modelling trends in euro area aggregate participation rates and projecting them forward. The results suggest that age and cohort effects can explain a substantial part of the recent increase in labor force participation rates in the euro area, although not the surge since early 2000s. Cohort effects are particularly relevant for women, with those born in the 1920s and 1930s less likely and those born in the late 1960s and early 1970s more likely to participate in the labor market over the life-cycle. There is substantial variation in cohort effects across the five largest euro area countries that are analyzed. Depending on the country, the estimated cohort profiles suggest an increase of 10 to 30 percentage points in female participation rates. In addition, a number of observed determinants, such as labor taxes, union density, unemployment benefits and the average number of children have had an impact on labor force participation rates, although the specific impact varies across age and gender groups and countries.

We use the results from the cohort model in order to consider different scenarios of future labor market participation. As cohort effects have increased for those born before the mid 1970's, but have then stabilized at a constant level, they continue to positively affect future participation as long as these groups still remain in the labor force. Looking forward, demographic factors will negatively affect participation as population ageing increases. We document that positive cohort effects are not large enough to compensate for the downward impact of population ageing on future labor force participation rates in the euro area.

# Chapter 1

# On the Implications of Technology and Non-Technology Shocks for Aggregate Labor Demand

## 1.1 Introduction

U.S. business cycles are characterized by large movements into and out of employment. The standard framework commonly used to study these movements comprises search-and-matching in the labor market as first presented by Mortensen and Pissarides (1994). In the dynamic version of this model, business-cycle fluctuations of labor market variables originate in fluctuations of labor productivity. These dynamics can be characterized by gross worker flows, i.e., the flow of unemployed workers filling an open job vacancy and employed workers separating from an existing employment relationship. The question whether the standard model is able to replicate the business-cycle fluctuations in U.S. time series data has been one of the most controversially discussed issues in the recent macro-labor literature.

Shimer (2005a) has fuelled the debate by criticizing the standard model with respect to its empirical performance. His criticism was based on comparing second moments generated from the model to second moments in worker flow data calculated from the U.S. Current Population Survey (CPS). He showed that the model did not mirror the high volatility of the job finding rate and unemployment that is observed in the data. In addition, the correlation between the job finding rate and the unemployment rate with labor productivity is much too high in the model.

While the dynamics in the standard frictional labor market model stem from fluctuations in labor productivity, "a change in labor productivity is most easily interpreted as a technology or supply shock" (Shimer (2005a), p. 25). Hence, labor market dynamics can be represented within a real-business-cycle (RBC) and growth model as in Merz (1995), or Andolfatto (1996). In these models technology shocks are the main driving forces of labor productivity. However, other disturbances such as demand shocks may affect labor productivity as well. Within this context, Galí (1999) demonstrated how to separately identify technology and non-technology shocks in time series data via restricting their long-run effects in structural vector-autoregressions (SVARs).

Against this background, this paper re-addresses the empirical performance of the standard search-and-matching model of the labor market in which fluctuations are driven by technology shocks. The empirical performance of the model is assessed based on second moments that are *conditional* on technology shocks rather than on overall unconditional moments.[1] Since conditional and unconditional moments substantially differ in this case, the judgement of the model that is based on unconditional moments may be very misleading. The results provide answers to various issues of importance to the standard labor market model. First, one can gain important insights into the failure of the model to generate sufficient volatility on the unconditional level as documented by Shimer. Second, in addition to the moments conditional on technology shocks, this analysis provides information about the importance of non-technology shocks and the dynamics induced by these shocks. Put differently, unconditional dynamics may encompass various different dynamics on the conditional level. Since the identified shocks are structural, the results deliver a meaningful guidance for the formal modelling of the labor market dynamics. Third, if the identified shocks are in fact shocks to the business cycle, their effect on the rate of job separations sheds light on the validity of assuming a constant job separation rate in a business-cycle model.

Two main findings emerge. With respect to volatility, the standard deviations of the job finding rate and the unemployment rate that are conditional on technology shocks are much lower than the unconditional ones. In addition, these standard deviations are, in fact, close to the standard deviations that are generated within a commonly calibrated version of the

[1] I am not the first to address conditional moments with respect to labor market dynamics. Michelacci and Lopez-Salido (2007), Ravn and Simonelli (2006), Fujita (2009) and many others all have also used SVARs in order to investigate the effect of different shocks on worker and job flows. I will refer to differences in the focus as well as methods and results below.

standard model that is driven by technology shocks. Consequently, the Shimer critique of the model with respect to its lack of volatility does not apply when the empirical performance is based on conditional moments. Since the technology shocks generate only a part of the overall volatility in the data, non-technology shocks play a substantial role for this volatility as well. In order to replicate the unconditional moments in the data, the standard model should therefore be augmented by additional non-technological sources of fluctuations rather than with respect to a better propagation of technology shocks as suggested in the literature. I show that shocks to the marginal rate of substitution between consumption and leisure, so-called preference shocks, may work in this respect. Further, job separations significantly move after both types of estimated shocks. This means that it is not reasonable to assume the job separation rate to be constant over the cycle.

With respect to the conditional correlations, the co-movement of the job finding rate with labor productivity that is conditional on technology shocks is negative, while the conditional correlation of unemployment with productivity is positive. Put differently, job finding falls after a positive technological innovation while unemployment increases. In the standard labor market model, a positive technology shock of the same size leads to an increase in labor productivity and, hence, to an increase in the job finding rate and a fall in unemployment. This result constitutes a "job finding puzzle" from the viewpoint of the standard model that is comparable to the so-called "hours puzzle" documented in Galí (1999). Since technology shocks play a considerable role for the business cycle variance of the job finding rate and unemployment, this result is a much more serious challenge to the empirical performance of the standard model than the Shimer volatility in unemployment puzzle. Hence, this result supports models which are able to incorporate these effects. Since the correlations of these two variables with productivity that are conditional on technology shocks are of opposite sign as the respective unconditional moments, non-technology shocks are necessary again to fully describe the overall dynamics in the data. However, I show that preference shocks are not suitable to explain the remaining variation in the data.

This paper presents results for different types of technology shocks and different types of measures for the labor market dynamics. Based on Galí (1999), technology shocks are the only shocks that have a long-run effect on labor productivity. This assumption holds in the RBC framework with frictional labor markets that is presented in Section 1.2. The identification of these standard Galí technology shocks within a structural VAR as well as their conditional moments that are estimated including the Shimer worker flow data are presented in Section 1.3. In addition, Fisher (2006) has motivated the separate identification of factor-neutral and investment-specific (or capital-embodied) technology shocks from the data. In the model, both of these shocks positively affect labor productivity in the long-run, while investment-specific technology shocks have a negative long-run effect on price of investment goods relative to consumption goods in addition. Section 1.4 presents the identification of these two shocks based on assumptions derived from the model and documents the results. Note that the identification employed uses an additional assumption on the effect of investment-specific technology shocks on labor productivity that goes back to Fisher. This assumption has an important effect on the results and has been neglected by many other authors in similar studies (such as Canova et al. (2007) and Ravn and Simonelli (2006)). Here, even though investment-specific technology shocks provide an additional source of volatility in job finding and unemployment, they are not large enough to explain the high volatility in the data. Further, investment-specific and neutral technology shocks generate very similar dynamics in the worker flow data and hence support the findings from the Galí identification.

Moments conditional on neutral and investment-specific shocks from the Fisher identification are presented for job flow data in Section 1.5.2. Data on job flows are generally viewed as an alternative to worker flows in order to assess the empirical performance of a model with a frictional labor market. Using recent data collected by Davis et al. (2006), the volatility result outlined above prevails. The job finding puzzle vanished however when incorporating job flows rather than worker flows in the estimation. Again, non-technological disturbances are necessary in order to fully understand the overall dynamics in the data.

Complementary to the Galí and Fisher identification, Section 1.6 proposes a new and alternative identification strategy for technology shocks which attempts to shed light on a few issues that arise from the estimation of technology shocks and their potential impact on the results. First, I document that the identified standard Galí technology shocks have a positive and significant effect on the relative price of investment. This means that the Galí technology shocks are neither truly neutral technology shocks nor are they investment-specific technology shocks. Rather, these shocks are negatively biased towards new investment. Neither the Galí nor the Fisher identification accommodates this variation in the data. Second, the Fisher identification of technology shocks employs an assumption which fixes the effect of the investment-specific technology shock on labor productivity and consequently the correlation between this shock and the neutral technology shock.

I propose a mixture of long-run zero and sign restrictions to distinguish positive productivity shocks with positive from positive productivity shocks with negative effects on the investment price. On the one hand, this provides an identification of investment-specific technology shocks alternative to the Fisher identification. Thereby I can test the critical Fisher restriction for its validity. On the other hand, I identify a new kind of technology shocks, namely positive technology shocks that are negatively biased towards investment. These shocks have so far not been taken into account in the literature as it is not clear how to interpret them. However, they are shown to play a significant role for the dynamics of the labor market variables. For both types of technology shocks following from this identification, the general results with respect to the empirical performance of the standard model based on moments conditional on these shocks continue to hold.

## 1.2 A standard labor market model

### 1.2.1 The model

The standard labor market framework referred to in the following nests search-and-matching in the labor market within a real-business-cycle (RBC) and growth model as in Merz (1995). The model comprises the subsequent equations:

$$\max_{\{C_t, N_{t+1}, V_t, K_{t+1}\}_{t=0}^{\infty}} \mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \left( \chi \ln(C_t) - \frac{N_t^{1+\phi}}{1+\phi} \right)$$

subject to

$$\begin{aligned}
A_t K_t^{\alpha} N_t^{1-\alpha} &\geq C_t + X_t + aV_t Z_t \\
K_{t+1} &\leq (1-\delta)K_t + I_t X_t \\
N_{t+1} &= (1-\psi)N_t + \mu V_t^{1-\eta}(1-N_t)^{\eta} \\
A_t &= \exp(\gamma + \varepsilon_{at})A_{t-1} \\
I_t &= \exp(\nu + \varepsilon_{it})I_{t-1}.
\end{aligned}$$

The posting of vacancies $V_t$ creates a cost $a$ and thereby search frictions. Employment next period is determined by those jobs that remain after exogenous separation $\psi$ and the new job matches that are formed in this period via a commonly used Cobb-Douglas matching function with matching elasticity $\eta$. The labor force is assumed to be constant, so that unemployment in period $t$ can be measured by $1 - N_t$. Job finding per period can be described by $F_t = \mu(\frac{V_t}{1-N_t})^{1-\eta}$ and thus co-moves with labor market tightness, defined as the ratio of vacancies to unemployment. The social planner representation can be derived from a decentralized problem in which workers and firms bargain over the wage. In order to meet the Hosios condition, the bargaining weight is implicitly set equal to the matching elasticity in this setup.

As in Fisher (2006), growth is exogenously generated by two types of technological progress. $A_t$ represents general purpose technology in the production function and will be called neutral technology in the following. $I_t$ is referred to as investment-specific technology as makes new investment goods relatively cheaper than consumption goods and hence drives the real price of new investments down.[2] Through the capital accumulation equation it favors new investments, leads to new capital formation and hence positively affects output and labor productivity. As in Fisher, output, consumption, investment and labor productivity grow with the rate $\frac{\alpha\nu+\gamma}{1-\alpha}$ along a balanced growth path, while the capital stock grows at rate $\frac{\nu+\gamma}{1-\alpha}$. Employment, unemployment and vacancies are stationary[3]. Shocks to these two types of technology generate business cycle fluctuations in the model. Note that each one of these technology shocks also constitutes a labor productivity shock. Through its positive effect on labor productivity, job finding rises after a positive technology shock, while unemployment falls. Following from the two laws of motion for technology, the investment-specific technology shock has a permanent effect on the relative price of investment, and both technology shocks have permanent effects on labor productivity. These two properties will serve as identifying restrictions in the estimation and hence, this framework serves as the suitable setup for the subsequent empirical investigation.[4]

[2]This can also be described as $\frac{1}{P_t}$. Greenwood et al. (2000) derive this one-sector representation of the model from a two-sector version with a consumption and an investment sector. Empirically, investment-specific technological progress is believed to be responsible for the persistent fall in the real price of equipment goods from 1955 until 2000 as measured by Cummins and Violante (2002) among others.

[3]Hence, vacancies are multiplied by $Z_t = A_t^{\frac{1}{1-\alpha}} I_t^{\frac{\alpha}{1-\alpha}}$ in the budget constraint.

[4]Note that DeBock (2006) also presents a search-and-matching model with investment-specific technology shocks. However, the shocks are transitory in his framework and therefore not in line with our identification of technology shocks applied later. Michelacci and Lopez-Salido (2007) describe a search-and-matching model with permanent neutral and investment-specific technology shocks. Their model is much more complicated than the standard model here and aims at describing different results in the data.

The labor market model outlined above differs in many respects from the standard Mortensen and Pissarides (1994) model that provides the basis for the Shimer model. Utility is not linear, but follows the standard assumptions in the RBC literature. In addition, due to the explicit modelling of capital and capital accumulation (i.e. savings) as well as output fluctuations, the RBC setting aims much more at imitating real fluctuations outside the labor market. This will be important for potential extensions in order to augment the performance of the model with respect to other variables and to other shocks. However, as in Shimer, this study focusses on the second moments of the central variables that this model wants to explain, that is the dynamics in the job finding, job separation and unemployment rate.

Both the Shimer model and the model outlined above lack many features that have been shown to be important to replicate overall dynamics in the data such as nominal or real rigidities outside the labor market. The standard labor market model serves as a baseline model in order to contrast its empirical performance based on unconditional moments with moments conditional on labor productivity shocks, that is, technology shocks. It is straightforward to add any other non-technological source of variation on productivity, e.g. demand shocks. As long as extensions of the model do not affect the validity of the identification, the empirical results documented below remain equally valid. In Section 1.3, I will consider preference shocks which move the marginal rate of substitution between consumption and leisure. In the model, this means that the parameter $\chi$ will be replaced by a stochastic process of the form $\ln(x_t) = \rho_x \ln(x_{t-1}) + \varepsilon_{xt}$.

### 1.2.2 Empirical performance based on neutral shocks

Due to the difference to the Shimer model, I re-consider the empirical performance of the model outlined above. To keep the framework as simple as possible, I start with considering neutral shocks as the only source of variation in the model. For this, I calibrate the model and generate artificial time series from the model, compute the respective second moments and compare them to the unconditional ones in the data. I choose a set of standard parameters for the calibration: a capital share in production of $\alpha = \frac{1}{3}$, the time discount factor of $\beta = 0.99$ and capital depreciation of $\delta = 0.02$. The Frisch labor supply elasticity is pinned down by $\phi = 1$ and $\chi = 1$. In line with Mortensen and Nagypal (2007), the elasticity of the matching function with respect to unemployment is set to $\eta = 0.46$. The constant of the matching function ($\mu = 1.5$) and the cost of posting vacancies ($a = 0.02$)

are calibrated such that the steady state labor market tightness is equal to one and the respective job finding rate equals the mean quarterly job finding rate of 1.5 in the worker flow data used later in the estimation. The same data delivers the mean quarterly job separation rate of $\psi = 0.09$.[5]

The first and second column of Table 1.1 compare the second moments in the data to those that are generated from the model driven by neutral shocks only. Hence, $\varepsilon_{it} = 0$. The growth rate and standard deviation of the neutral technology shock $\varepsilon_{at}$ are then calibrated to match the standard deviation of labor productivity which results in $\gamma = 0.0035$. Both the artificial and the data series are detrended with a very smooth HP-filter ($\lambda = 10^5$) as in Shimer in order to relate my results directly to his. In the actual data, the job finding rate and unemployment are a lot more volatile than the job separation rate. From this, Shimer concludes that unemployment fluctuations are mainly driven by fluctuations in the job finding rather than the job separation rate. Furthermore, the standard deviation of the job finding rate and unemployment are about ten times as large as the one in labor productivity. All series are highly autocorrelated in the first lag.

The comparison with the model moments mirrors the Shimer volatility in unemployment puzzle. First, the standard deviations of job finding and unemployment generated in the model are very small compared to the ones in the data. Second, the correlation of unemployment and job finding with productivity is too high in the model compared to the data.[6] Shimer concludes that there exists no internal propagation mechanism of labor productivity shocks in the model, since the real wage strongly reacts to labor productivity shocks and hence weakens the incentives for firms to post vacancies. In order to improve its empirical performance, Shimer and also Hall (2005) have therefore proposed to introduce rigid wages into the standard framework.

Hagedorn and Manovskii (2008) and many other authors have argued that Shimer's volatility in unemployment puzzle disappears for a different calibration of the model, more precisely for a different calibration of the outside option of the workers in the wage bargaining. This parameter is not considered here. Within the framework used above, the parameters are chosen such that the volatility in the job finding rate and unemployment is as high

[5] For more details on the data and the sample, see Section 1.3.1.

[6] Table 1.7 shows that these result do not depend on the choice of the smoothing parameter in the HP-Filter.

as possible[7]. Put differently, the aim of this study is not to find a calibration such that the model driven by technology shocks matches the unconditional moments in the data. Rather, the output from this model in the standard calibration is to be compared to the moments that are conditional on technology shocks.

## 1.3 Moments conditional on technology shocks

In the model, business cycle fluctuations of labor productivity, job finding and unemployment originate in movements of technological progress. It is therefore straightforward to evaluate the empirical performance of the model based on second moments conditional on technology shocks rather than on unconditional moments. In the data, shocks other than technology shocks play a role for the overall fluctuations as well. Thus disentangling the technology shocks from other shocks potentially serves three purposes. First, I can investigate the dynamic relationships (correlations and impulse responses) between the variables of interest that are conditional on technology shocks. Second, since these may be different from the unconditional ones it may therefore be possible explain the failure of the model on the unconditional level. Third, it is possible to assess the importance of technology shocks for the unconditional data dynamics.

### 1.3.1 Identification and estimation

The effects of technology shocks on labor market variables can be investigated within a structural VAR framework with long-run restrictions based on Blanchard and Quah (1989). The main idea is to find a mapping that transforms the residuals from a reduced form VAR into structural residuals such that the latter can be interpreted as certain types of shocks such as technology shocks. These mappings typically involve assumptions on the variance-covariance matrix of the structural shocks as well as restrictions on the effects of these shocks on the variables in the VAR.

Based on Galí (1999), technology shocks are identified via the central assumption that they are the only shocks that positively affect labor productivity in the long-run. In addition, the technology shocks are orthogonal to each of the non-technology shocks estimated.

[7] Investigating sensitivity of this result to the choice of parameter values, it is possible, for example, to increase the matching elasticity with respect to unemployment to the value proposed by Shimer of $\lambda = 0.72$ which clearly decreases the volatility of job finding and unemployment.

Table 1.1: Historical decomposition of Galí identification

| | Uncond. Sample | Model I | Model II | Conditional Moments Technology | Conditional Moments Residual |
|---|---|---|---|---|---|
| **A: Standard Deviations** | | | | | |
| JFinding | 0.1542 | 0.0536 | 0.0417 | 0.0548 | 0.1229 |
| | | | | (0.04,0.08) | (0.10,0.14) |
| JSeparation | 0.062 | | | 0.0503 | 0.056 |
| | | | | (0.04,0.06) | (0.05,0.06) |
| Unemployment | 0.1786 | 0.0519 | 0.0404 | 0.0881 | 0.1409 |
| | | | | (0.06,0.12) | (0.12,0.16) |
| Productivity | 0.0156 | 0.0156 | 0.0116 | 0.0116 | 0.0166 |
| | | | | (0.01,0.02) | (0.01,0.02) |
| **B: Autocorrelations** | | | | | |
| JFinding | 0.9128 | 0.9071 | 0.9061 | 0.9189 | 0.8869 |
| | | | | (0.86,0.95) | (0.86,0.90) |
| JSeparation | 0.6336 | | | 0.9256 | 0.6158 |
| | | | | (0.89,0.95) | (0.59,0.66) |
| Unemployment | 0.9218 | 0.845 | 0.8443 | 0.9131 | 0.9109 |
| | | | | (0.88,0.93) | (0.90,0.92) |
| Productivity | 0.8507 | 0.8701 | 0.868 | 0.8927 | 0.9206 |
| | | | | (0.86,0.92) | (0.90,0.94) |
| **C: Cross-Correlations** | | | | | |
| JFind.,Prod. | 0.0567 | 0.8625 | 0.8522 | -0.436 | 0.6739 |
| | | | | (-0.66,-0.10) | (0.52,0.77) |
| JSep.,Prod. | -0.4392 | | | 0.3544 | -0.6703 |
| | | | | (0.11,0.48) | (-0.74,-0.59) |
| Unemp.,Prod. | -0.1858 | -0.7776 | -0.7668 | 0.4613 | -0.8014 |
| | | | | (0.17,0.63) | (-0.88,-0.70) |
| JFind.,Unemp. | -0.9558 | -0.9272 | -0.9266 | -0.9041 | -0.9359 |
| | | | | (-0.96,-0.75) | (-0.95,-0.91) |
| JSep.,Unemp. | 0.6845 | | | 0.885 | 0.6302 |
| | | | | (0.80,0.92) | (0.56,0.69) |
| JFind.,JSep. | -0.4404 | | | -0.596 | -0.3167 |
| | | | | (-0.76,-0.19) | (-0.40,-0.19) |

Notes: All series are detrended with the smooth HP-filter as in Shimer (2005a). For the conditional moments, the series are simulated with the respective shock operating only. The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution. Calibration I of the model matches the unconditional standard deviation of labor productivity, calibration II matches the same moment, conditional on technology shocks.

These assumptions are implemented by including labor productivity in first differences and ordered first in the VAR and then applying a Cholesky decomposition to the long-run horizon forecast revision variance[8]. It has to be noted that many structural disturbances other than technological innovations can affect labor productivity in the short and medium run, but that technology shocks can be distinguished from non-technology shocks with respect to their long-run effects on this variable. With this approach, I do not exactly estimate the model outlined above. Rather the conditional moments obtained should hold for a broad class of different model specifications that fulfill the identifying assumptions. The long-run assumption about the nature of technology shocks holds in the model presented as well as in many other models, such as the neoclassical growth model or the New Keynesian model[9].

All identification alternatives presented in the following are based on the same reduced-form VAR which contains labor productivity, the job finding and separation rate. For later comparison with alternative identification schemes, the relative price of investment is added to the VAR. The reduced-form VAR is estimated within a Bayesian framework with a Minnesota prior, similar to Canova et al. (2007). The Minnesota prior incorporates a unit root in the levels of the variables included in the VAR and a fixed residual variance which determines the tightness on own lags, other lags and potential exogenous variables as well as the decay of the lags. Using the latter parameter, this prior allows us to generate sensible results for a large number of lags, as Canova et al. outline. This addresses an often cited criticism on the VAR approach (e.g. by Chari et al. (2008)) which states that in theory one should employ a VAR with an infinite number of lags (here eight lags will be employed) in order to correctly identify technology shocks using long run restrictions. Except for the decay, I will use a relatively loose prior in the estimation[10]. Further, the VAR is estimated with a trend as suggested by Canova et al. (2006). Here, the trend is a dummy that is deterministically broken at 1973:2 and 1997:1. These dates have been considered as break points in the growth literature and replicate the turning points in the job separation rate and unemployment series.[11]

[8]See the Appendix A.1 for further details.

[9]It does not hold in endogenous growth frameworks.

[10]The prior variance of the coefficients depends on three hyper-parameters $\phi_1 = 0.2$, $\phi_2 = 0.5$ and $\phi_3 = 10^5$, that determine the tightness and decay on own lags, other lags and exogenous variables. The decay parameter is set to $d = 7$.

[11]See Fernald (2007b) for empirical evidence on the trend breaks. Section 1.4.3 presents robustness checks to this specification along various dimensions including different priors, different break points for

The baseline specification is estimated using quarterly time series data for the U.S. over the sample 1955:1-2004:4. The job finding and separation rates are taken from the worker flow data produced by Robert Shimer[12]. Labor productivity (output per hours of all persons) is the standard non-farm business measure provided by the U.S. Bureau of Labor Statistics. The real price of investment consists of a price index for equipment and software and a consumption price deflator that is chain weighted from nondurable, service and government consumption. The standard data from the National Income and Product Accounts (NIPA) have been criticized not to take into account the price-per-quality change in the investment goods of interest (see Gordon (1990)). I use the quarterly series generated by Fisher (2006) that is based on the measure of Cummins and Violante (2002) and that takes these flaws into account[13]. Labor productivity and the relative price of investment are included in growth rates in the VAR, while the job finding and separation rates are included in levels.

Under the assumption of homogenous workers and a constant labor force, the unemployment rate can be approximated by the steady state unemployment rate $\tilde{u} = \frac{js}{js+jf}$. Linearizing this relationship, one can also deduct the impulse-response of unemployment from the responses of the job finding and the job separation rates. Shimer's assumption that the job separation rate does not move over the cycle and, therefore, does not play a role for the fluctuations of unemployment has been criticized by Fujita and Ramey (forthcoming) among others. In fact, the job separation rate is more strongly correlated with labor productivity than the job finding rate as can be seen from the first column in Table 1.1. I include the job separation rate in the VAR in order to test this criticism.

### 1.3.2 The Shimer puzzle

Table 1.1 depicts the historical decomposition of the actual time series into the technology and non-technology (or residual) components. These component series are generated assuming the exclusive presence of the respective shock and using information on the first lags in the sample. Detrending the resulting series with the smooth HP-filter as in Shimer then delivers the business cycle components of interest. The historical decomposition documents the ability of the single shocks to replicate exactly those moments in the data that

the trend and no trend as well as different lag lengths in the VAR.

[12] This is the worker flow data officially posted on the website of Robert Shimer and documented in Shimer (2005b). For additional details, see http://home.uchicago.edu/∼shimer/data/flows.

[13] The series by Jonas Fisher was extended by Ricardo DiCecio. I thank both for making their data available to me.

have been used for judging the empirical performance of the model.[14]

Volatility is measured by the standard deviation in panel A. The standard deviations of the component series of the job finding rate and unemployment that are driven by technology shocks are less than half of the overall sample volatility. In fact, if the model is calibrated to match the standard deviation of labor productivity that is conditional on technology shocks (calibration II in column 3 of Table 1.1), the standard deviation of the job finding rate generated in the model is close to and lies within the confidence bands of the standard deviation that is conditional on technology shocks.

The model assumes a constant job separation rate over the business cycle. The estimated standard deviation of the job separation rate that is conditional on both technology shocks and non-technology shocks is, however, significantly positive. If business cycles are driven by technology shocks, this result undermines the assumption of a constant separation rate over the cycle. Instead, this result favors a theoretical context with endogenous rather than exogenously fixed job separation as in denHaan et al. (2000).

Addressing the empirical performance of the model with constant job separation nevertheless, one should therefore consider the volatility of unemployment that is driven by the job finding rate only, setting the job finding rate to its mean value throughout the sample period. The unconditional standard deviation of 0.1525 is then contrasted with the 0.0548 conditional on technology shocks and 0.1237 conditional on non-technology shocks (see first row in Table 1.2). The standard deviation in unemployment that is generated by the model therefore lies within the confidence bands conditional on technology shocks. As a result, the technology-shock driven model works well to replicate the volatilities in the job finding rate and unemployment that are conditional on technology shocks in the data. As a consequence, the Shimer critique does not apply.

While the model works well to generate the volatility that is conditional on technology shocks, it, however, still fails to explain the overall volatility in the sample. In fact, a large part of the volatility still remains to be unexplained in the "residual" disturbances as depicted in the last column of Table 1.1[15]. In order to replicate the dynamics in the

[14]Note that the second moments resulting from these series do not add up to the unconditional moment. Note also that all results discussed also hold for HP-filtered data using the standard parameter of $\lambda = 1600$ as can be seen in Table 1.7 in Section 3.4.

[15]In a parallel developed paper, Barnichon (2008) also shows the importance of non-technology shocks for worker flows. He argues that these remaining shocks are monetary policy shocks.

overall data, the standard search-and-matching model should consequently be augmented by additional non-technology sources of volatility, generally referred to as demand shocks. Hall (1997) has proposed a candidate for these residual shocks, namely preference shocks or shocks to the marginal rate of substitution between consumption and leisure.[16] As mentioned in Section 1.2, it is easy to incorporate these kinds of shocks into the model. After a positive preference shock, agents in the economy want to consume and work more, hence they are willing to accept a lower wage in order to become employed which increases the incentive for firms to post vacancies and decreases unemployment. Panel A and B of Table 1.2 depict the unconditional and conditional moments in the data (assuming a constant job separation rate) as well as the moments from the model that is driven by preference shocks only. The model is calibrated to match the standard deviation of labor productivity that is conditional on the non-technology shocks which involves $\rho_x = 0.5$ and $\sigma_x = 0.2$. Preference shocks are suitable to generate high volatility in these two variables as suggested by Hall.

### 1.3.3 The "job finding puzzle"

The autocorrelations conditional on technology shocks are close to the unconditional ones. The model lacks some persistence with respect to the job finding rate as the autocorrelation is a bit too low compared to the one in the data. Generally however, the model performs well in replicating the conditional and unconditional autocorrelations. The conditional co-movement of the variables is depicted in panel C of Table 1.1 and also in the impulse-responses to a one-standard deviation technology shock in Figure 1.1[17]. Most prominently, job finding falls after a positive technology shock and the conditional correlation between job finding and productivity is negative. Regardless of the job separation rate, unemployment increases after the fall in job finding and the correlation of unemployment and productivity is positive. These two effects are opposite to those in the overall sample and the exact contrary to what the standard model proposes. Hence, this result

[16]Hall decomposes macroeconomic variables into fluctuations that originate in technology, government spending and preference shocks. He bases his decomposition on equations derived from a standard RBC-model, he does not use structural VAR techniques for his analysis. He shows that preference shocks account for most of the fluctuations in hours worked. His results are therefore similar to the results documented here.

[17]The response of unemployment is calculated from the linearized relationship between the approximated unemployment rate and the responses of the job finding and separation rates according to $\hat{u}_t = \frac{f}{(s+f)^2}\hat{s}_t - \frac{s}{(s+f)^2}\hat{f}_t$, where $s$ and $f$ are the mean values of the two rates respectively.

Table 1.2: The role of job separation and preference shocks

| | Unconditional Sample | Model Pref. Shocks | Conditional Moments Technology | Residual |
|---|---|---|---|---|
| A: Standard Deviations | | | | |
| JFind. and Unemp. | 0.1526 | 0.1314 | 0.0548 | 0.1238 |
| | | | (0.04,0.08) | (0.10,0.14) |
| Productivity | 0.0156 | 0.0165 | 0.0116 | 0.0165 |
| | | | (0.01,0.02) | (0.01,0.02) |
| B: Autocorrelations | | | | |
| JFind. and Unemp. | 0.9128 | 0.832 | 0.9207 | 0.8873 |
| | | | (0.85,0.95) | (0.86,0.90) |
| Productivity | 0.8507 | 0.9184 | 0.8902 | 0.9208 |
| | | | (0.86,0.92) | (0.90,0.94) |
| C: Cross-Correlations | | | | |
| JFind.,Prod. | 0.0489 | -0.7702 | -0.4347 | 0.662 |
| | | | (-0.64,-0.07) | (0.53,0.76) |
| Unemp.,Prod. | -0.0489 | 0.892 | 0.4332 | -0.6626 |
| | | | (0.07,0.64) | (-0.76,-0.53) |

Notes: All series are detrended with the smooth HP-filter as in Shimer (2005a). Unemployment is calculated with a job separation rate that is constant and set equal to its mean value over the sample. For the condi tional moments, the series are simulated with the respective shock operating only. The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution. The model is driven by preference shocks only and is calibrated such that it matches the conditional standard deviation of labor productivity.

challenges the conventional dynamics in the standard search-and-matching model in a similar fashion as the results in Galí (1999), known as the "hours puzzle", have challenged the RBC paradigm with frictionless labor markets.[18]

A variance decomposition adds up the impulse-response coefficients from the estimation to a certain conventional business cycle horizon. This statistic reports the respective contribution of each shock to the overall variance and therefore also highlights the importance of the shocks relative to each other. Decomposing the business cycle variance of the Galí identification into the contribution of technology and non-technology shocks, technology shocks explain up to 17% of the business cycle variance of job finding and over 20% of the variance of unemployment. Hence, an appropriate model should take these dynamics into account.

[18]Researchers have questioned that the identified shocks can in fact be interpreted as technology shocks. Section 1.3.4 shows robustness for this finding using an alternative measure of technology derived by Basu et al. (2006).

Figure 1.1: Impulse-responses to Galí technology shocks

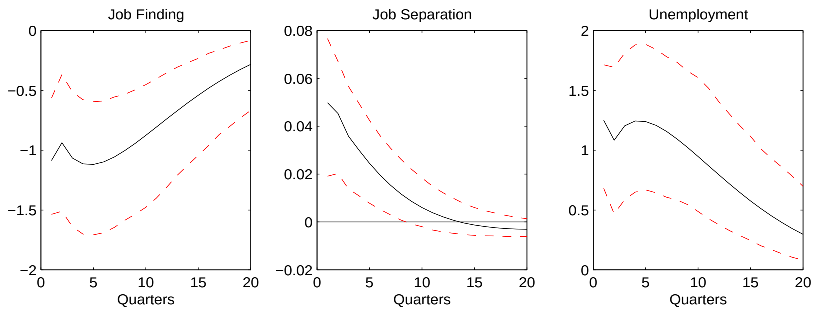


Notes: Responses in percentage points to a positive one-standard-deviation shock. Confidence intervals are 68% Bayesian bands.

Galí has explained the drop in hours worked within a sticky price New Keynesian framework. Can the natural extension of this framework including search-and-matching in the labor market equally explain the drop in the job finding rate? In the case of hours, fixed demand in the short run leads firms to adjust hours worked after a positive technology shock. Since it is much more costly to adjust employment rather than hours worked, it is not clear that the same mechanism works equally well in this context. In their specification with real rigid wages, Blanchard and Galí (2006) document that unemployment increases after a positive productivity shock. Here, labor market tightness and hence the job finding rate move together with unemployment replicating the dynamics documented above. Barnichon (2008) uses a similar reasoning to generate the fall in labor market tightness which he documents in a similar SVAR-framework as the one presented here. However, as conjectured, his model is not able to generate the large fall in labor market tightness and strong increase in unemployment that we see in the data.[19]

There exist explanations for this empirical finding different from a New Keynesian setup. Chapter 2 documents that the shocks that have been identified as neutral technology shocks in the Galí identification are in fact positively biased towards new skills (as they have a positive effect on the wage premium of high to low skilled workers). Consider a framework in which two types of workers are used in production and are to some degree substitutable. After a positive skill-biased technology shock, high-skilled workers become

[19]In contrast, Krause and Lubik (2007) present a framework in which job finding falls after a positive productivity shock mentioning that the resulting dynamics are counterfactual. This is no longer true based on conditional moments. In Christoffel et al. (2006), vacancies fall and unemployment increases after a positive productivity shock, resulting in an fall of labor market tightness and the job finding rate.

more productive than low-skilled workers and overall labor productivity increases. Low-skilled workers will then be substituted out of employment. The job finding rate for low-skilled workers will drop, while it will potentially increase for high-skilled workers. If the negative effect on low-skilled is larger than the positive effect on high-skilled workers, the overall job finding rate drops and unemployment increases.

Regardless of the mechanism, a model driven by technology shocks is again not suitable to explain the overall dynamics in the data. Rather, non-technology shocks are needed in order to model the unconditional dynamics in the data. Reconsidering the preference shocks from above, these kinds of shocks have been popular in the RBC-literature in order to explain the empirical correlation of labor productivity with hours[20]. Table 1.2 documents that the correlations of the job finding rate and unemployment with productivity that are generated by preference shocks in the model are opposed to the ones conditional on non-technology shocks in the data, however. After a positive preference shock, agents want to consume more and hence decrease investments. Capital falls and, after an initial increase, output falls as a consequence. Due to the increase in employment, labor productivity falls which induces a negative correlation of this variable with the job finding rate and a positive one with unemployment. Hence, preference shocks are not suitable to explain the conditional correlations within this setup. It has to be noted that in a New Keynesian setup, the induced correlations are different and preference shocks could replicate the empirical dynamics. A distinction between skill-biased and skill-neutral shocks could also provide two shocks that match the conditional correlations in the data.

As exhibited in Figure 1.1, job separation significantly increases after a positive technology shock contributing to an even larger increase in unemployment. A rise in job separation after a positive innovation in technology might be due to the fact that not all of the existing job matches can freely use this new technology. Hence, technological innovation is embodied in new jobs, or specific to existing vintages. Canova et al. (2007) employ a vintage human capital in order to model the "Schumpeterian creative destruction" after a neutral technology shock. As is documented in greater detail in Section 1.4.3, the effect of job separation is not robust neither when considering different sub-samples nor to the in- or exclusion of a trend in the estimation.

[20]See for example Bencivenga (1992) on the Dunlop-Tarshis observation.

Figure 1.2: Impulse-responses to BFK technology shocks

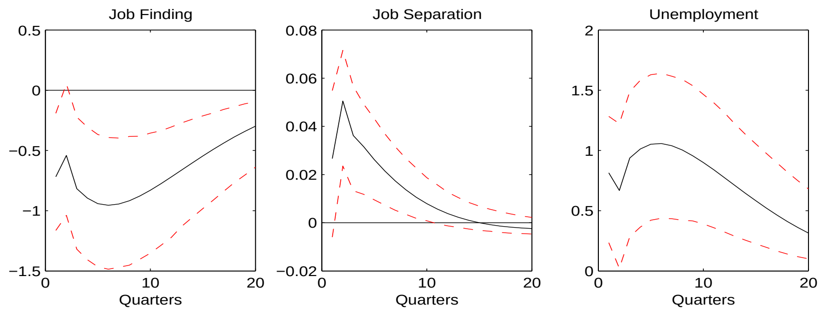


Notes: Responses in percentage points to a positive one-standard-deviation shock. Confidence intervals are 68% Bayesian bands.

### 1.3.4 Are the estimated shocks really technology shocks?

Many researchers have questioned that the structural residuals that are identified from a Galí-style VAR are in fact estimates of technological progress. Supporting the findings from Galí, a recent piece of evidence from Basu et al. (2006) has documented that their measure of technological progress, derived as a “sophisticated” Solow residual from a very different exercise, also induces a contractionary effect on hours worked. Here, I use this measure in order to support the effect of technology on the job finding and separation rate from my estimation in two different ways. First, I incorporate “true” total-factor-productivity (TFP) instead of labor productivity into my SVAR with long-run restrictions. Neutral technology shocks are then the only shocks that move TFP in the long run. As depicted in Figure 1.2, the effects of these shocks on the job finding rate, the job separation rate and unemployment are very similar to the ones from the estimation with labor productivity. Second, as suggested by Basu et al. (2006), I regress four lags of their technology measure (dz) on job finding and job separation. Here, I detrend the two rates as in the VAR by regressing them on a dummy trend broken at 1973:2 and 1997:1. Table 1.3 shows the results, for impulse-responses, one could simply add the estimated coefficients. Here, TFP has a negative effect on the job finding rate. The effect on the job separation rate is also negative, but since this effect is small (and insignificant), unemployment still increases after a shock to TFP.

Table 1.3: Regression on BFK measure

| Dependent variable | Regressor | | | | |
|---|---|---|---|---|---|
| | dz | dz(-1) | dz(-2) | dz(-3) | dz(-4) |
| JFinding | -0.6250* | -0.3429 | -0.4441* | -0.5339* | -0.3447 |
| JSeparation | -0.1473 | 0.0305 | -0.0835 | -0.1753 | -0.1848 |

Notes: The star * denotes significance based on one standard error bands.

## 1.4 Different shocks: Fisher identification

Fisher (2006) based on Greenwood et al. (1997) has addressed the issue that fluctuations in labor productivity might be generated not only by factor-neutral technological progress, but also by investment-specific technological innovations. Consequently, investment-specific technological progress satisfies the identifying assumption for the Galí technology shocks and hence invalidates the interpretation of these shocks to be factor-neutral. Fisher proposes a strategy to separately estimate neutral and investment-specific technology shocks and documents that the two shocks might have different effects on macroeconomic variables. Further, investment-specific technological progress contributes to a larger extend to growth and cyclical fluctuations of macroeconomic variables (in particular of output and hours worked) than neutral technology. Investment-specific technological progress thus provides a potential additional source of variation in the job finding rate and unemployment.

In the original Shimer framework, it is not possible to distinguish between these two sources of variation in labor productivity, while the model in Section 1.2 does differentiate between these two shocks. As mentioned before, the labor market dynamics that are induced by the two technology shocks are actually very similar, i.e., job finding increases and unemployment falls after both technology shocks. However, since the formation of capital takes time, productivity increases with a lag in response to investment-specific technological progress. This increases the overall standard deviation of the job finding rate and unemployment in the model in which both types of technology shocks operate (see second column of Table 1.8 in the Appendix to Chapter 1). Further, the correlation between the job finding rate and productivity is smaller than in the model with neutral shocks only. However, these effects are not large enough to replicate the unconditional data moments, hence the Shimer critique still holds.[21]

### 1.4.1 Identification

In order to identify the two types of technology shocks, Fisher imposes the assumption that investment-specific technology shocks are the only shocks that (negatively) affect the relative price of investment in the long-run and that are additionally allowed to affect labor productivity in the long-run. (Investment-)neutral technology shocks are then the only remaining shocks that affect labor productivity in the long run. Note that this assumption is true in the model outlined in Section 1.2.

It is easy to implement these two assumptions ordering the first differences of the relative investment price and labor productivity first in the reduced-form VAR and applying a Cholesky decomposition to the long-run forecast revision variance. However, the effect of the investment-specific shocks on labor productivity is estimated to be negative in our baseline specification. This means that all or at least a part of the identified investment-specific shocks are not technology shocks according to the Galí definition and more importantly not positive shocks to labor productivity as the ones in the model and referred to by Shimer. Fisher addresses this problem by introducing the additional assumption that positive investment-specific shocks increase labor productivity by a fixed proportion to their effect on the investment price. Derived from the production function in the model this proportion is set to $\frac{\alpha}{(1-\alpha)}$. This additional assumption comes at a cost as it not only strongly restricts the long-run productivity effect of investment-specific shocks to a certain value but also implies a positive and fixed correlation between the investment-specific and neutral technology shocks.[22]

There exist several a few studies that consider the responses of worker flows to both neutral and investment-specific technology shocks based on the Fisher identification. The work by Canova et al. (2006) is closely related to the analysis in this section of the paper. The estimation of the reduced form VAR in a Bayesian framework with a Minnesota prior is

[21] In this simulation of the model, the growth rates and standard deviations of the two types of technology shocks are calibrated to match the moments of labor productivity and the investment price which results in $\gamma = 0.0074$ and $\nu = -0.0117$ for our sample. The mean growth rate of labor productivity then equals $\frac{1}{1-\alpha}\gamma + \frac{\alpha}{1-\alpha}\nu$.

[22] See Figure 1.6 for a comparison of the responses of the restricted and the unrestricted Fisher identification. See the Appendix A.3 for more details and the implementation of this identification scheme. Parallel to the model calibration I use $\alpha = \frac{1}{3}$.

taken directly from them. However, Canova et al. employ the Fisher identification without the additional third restriction. Equally, Ravn and Simonelli (2006) identify technology shocks without the third restriction in a framework which also incorporates fiscal and monetary policy shocks. Adding the third restriction delivers quite different dynamics induced by the investment-specific technology shock. I will discuss this issue further in Section 1.6 in which I also propose a test for the third restriction. Complementary to these studies, there exist many contributions in the literature that estimate medium or large scale DSGE models which incorporate search-and-matching in the labor market. Here, technology shocks are usually identified based on a combination of short-run sign restrictions as in Fujita (2009) or Braun et al. (2006). While these shocks should generally depict the same dynamics as the technology shocks identified in this paper, this is not always the case and depends on the fact that the co-movement between labor input and productivity in the short run is explicitly used for identification.

### 1.4.2 Results

The historical decomposition of the standard deviation supplements the results from the Galí identification, see Table 1.8 in the Appendix to Chapter 1. Both types of technology shocks, as well as both technology shocks taken together, generate standard deviations in the job finding rate and unemployment that are much smaller than the unconditional standard deviations, but quite close to the ones produced from the model. Again, sources other than technology are necessary to understand the unconditional volatility in the data.[23]

With respect to the conditional dynamics, Figure 1.3 depicts the responses of the job finding and separation rate as well as unemployment to positive one standard deviation technology shocks from the Fisher identification. Note that the responses to the neutral shock are very similar to the responses derived from the Galí identification. Job finding drops after both types of technology. This effect is stronger and more persistent after a neutral technology shock than after an investment-specific shock. The job separation rate does not significantly react to an investment-specific technology shock. The falling job finding rate positively affects the unemployment rate, but the effect is again not as strong as for the neutral technology shock. Consequently, the contrast between the conditional

[23]Note that here, the two technology shocks are not orthogonal. Hence, the historical decomposition is not truly a decomposition. Technology shocks and the residual disturbances are orthogonal, however.

Figure 1.3: Impulse-responses to Fisher technology shocks

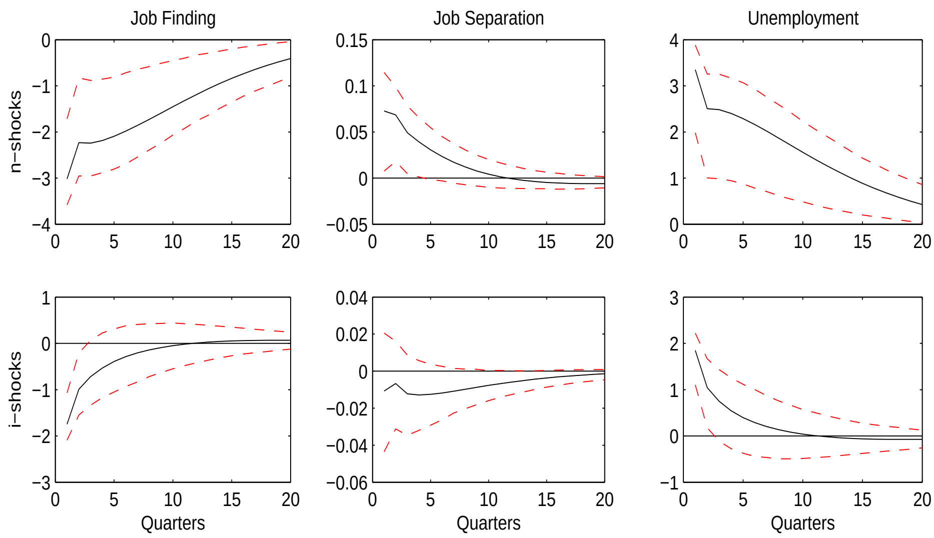


Notes: Percentage responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

dynamics in the data versus the ones in the model still exists, but is weaker in case of the investment-specific shocks. This is also reflected in the conditional correlations in panel C in Table 1.8. The conditional correlation of job finding and productivity is much lower than the one conditional on a neutral shock, the correlation of unemployment with productivity has the same sign as the unconditional one, both of these figures are insignificant. The investment-specific technology shock therefore moderates the effect of the neutral shock. Both technology shocks taken together however still generate dynamics that are opposite to the unconditional dynamics and that are not replicated in the model.

Table 1.6 in the Appendix to this chapter exhibits the contribution of the shocks to the forecast error variance of the variables in this small VAR. The neutral shock is much more important for the variances of the labor market variables than the investment-specific shock. This highlights again the importance to replicate the dynamics of this shock in an appropriate model. Together, the technology shocks explain between 45% to 60% of the variance of job finding and unemployment.[24]

[24] This result is similar to Canova et al. (2007) who, in spite of an alternative identification of investment-specific technology shocks, document that employment effects can mainly be attributed to neutral technology shocks.

### 1.4.3 Robustness

This section investigates the robustness of the main results from the Fisher identification. As documented above, the neutral shocks from the Fisher identification and the Galí identification are very similar in fact. The robustness analysis focusses on the two main results: The low standard deviation conditional on neutral and invest-ment-specific technology shocks in job finding and unemployment and the drop in the job finding rate after positive innovations of both types of technology. Table 1.4 summarizes the results.

The first set of robustness checks deals with the prior and the lag length in the estimation of the reduced form VAR. Clearly, the baseline specification with the Minnesota prior is different from a standard OLS specification with 2 to 4 lags in the VAR. In the Minnesotay prior, a high decay parameter is necessary for a large number of lags to generate both significant and sensible results. Using a smaller number of lags together with a smaller decay on these lags, or similarly a flat prior (OLS equivalent) for the estimation of the reduced form VAR, qualitatively supports the findings in the baseline specification, but is not significant, however. Further, the results are robust to relaxing the assumption of a fixed residual variance within a Normal-Wishart prior structure. The prior suggested by Kadiyala and Karlsson (1997) employs the same mean for the coefficients as the Minnesota prior and generalizes the Minnesota prior in terms of a non-diagonal, unknown residual variance. Compared to the Minnesota prior, the coefficient variance additionally weights the effect of the exogenous variables on a variable with its respective variance and fixes $\phi_1 = 1$.

The baseline specification includes a broken dummy-trend into the specification which is not uncontroversial. In fact, the question of whether or not to include a trend into the specification is closely related to the debate on how to specify hours worked in a similar structural VAR. Here, it has been shown that if specified in first differences or HP-filtered, hours worked fall after a positive Galí-type technology shock, while they increase after the same type of shock if specified in levels (see Galí (1999) and Christiano et al. (2003) respectively). The fall in hours worked after a positive technology shock contradicts the standard RBC paradigm and has become famous as the “hours puzzle” in the literature. In fact, a trend as the one applied here takes out slow-moving components from the series and is therefore related to taking first differences of the labor market variables. Canova et al. (2006) argue that if the variables are specified in levels, long-run restrictions may pick up the slowly moving components of the variables, even though they aim at explaining

Table 1.4: Robustness of the Fisher identification

| | Conditional Standard Deviation | | | | Impulse Response | |
|---|---|---|---|---|---|---|
| | Job Finding | | Unemployment | | Job Finding | |
| | i-shock | n-shock | i-shock | n-shock | i-shock** | n-shock |
| Baseline | 0.0627 | 0.0667 | 0.0692 | 0.0972 | -,sign. | -,sign. |
| Baseline specification with Minnesota prior changed to | | | | | | |
| 4 lags, decay 7 | 0.0651 | 0.071 | 0.0808 | 0.1129 | -,sign. | -,sign. |
| 12 lags, decay 7 | 0.069 | 0.0702 | 0.847 | 0.1053 | -,sign. | -,sign. |
| 8 lags, decay 4 | 0.579 | 0.0477 | 0.0745 | 0.0689 | -;+,not sign. | -,not sign. |
| 3 lags, decay 1 | 0.0533 | 0.0567 | 0.0706 | 0.0809 | -,not sign. | -,not sign. |
| Flat prior (OLS equivalent) with | | | | | | |
| 2 lags | 0.0511 | 0.0609 | 0.727 | 0.0971 | -,not sign. | -,not sign. |
| 3 lags | 0.0533 | 0.0649 | 0.0737 | 0.0899 | -;+,not sign. | -,not sign. |
| K and K prior* | 0.651 | 0.0738 | 0.689 | 0.1037 | -,sign. | -,sign. |
| Trend specification | | | | | | |
| no break | 0.0667 | 0.0595 | 0.058 | 0.0494 | -,sign. | -,sign. |
| Fisher subsamples without break | | | | | | |
| 1955:I-1979:II | 0.0828 | 0.0853 | 0.0784 | 0.0895 | -,sign. | -,sign. |
| 1982:III-2004:IV | 0.0352 | 0.059 | 0.0777 | 0.0402 | -;+,sign. | -,sign. |
| Fujita and Ramey subsample without break | | | | | | |
| 1976:III-2004:IV | 0.0424 | 0.0699 | 0.0622 | 0.0528 | -;+,sign. | -,sign. |

Notes: **Describes the effect on impact. Here, -;+ indicates initial drop, then hump-shaped increase.
*Kadiyala and Karlsson prior with Minnesota structure, same parameters as in baseline specification.

business cycles fluctuations.

Figure 1.11 shows the results for the baseline specification without the dummy breaks. The job finding rate still decreases after positive innovations of both technology shocks. This means that the "job finding" puzzle is is robust to including a trend or not in the specification. Note further that job separation now falls significantly after both shocks. In fact, it falls by such a large extend that the unemployment rate falls in the longer horizon which reflects the result from the hours debate. In addition, the results from the entire sample are compared to results for subsamples suggested by Fisher (2006). Here, no trend is incorporated into the specification, the results are robust to an inclusion of trend breaks as in the baseline specification, however. In the latter sample, investment-specific technology shocks induce an initial fall in the job finding rate and a subsequent, (borderline) significant increase. Job separation does not react to a neutral shock, but decreases significantly after an investment-specific technology shock. Hence, these shocks do generate dynamics different from the neutral shocks in this sample.

## 1.5 Alternative variables

### 1.5.1 Alternative worker flows

The worker flow data of Shimer and the respective business-cycle facts are not uncontroversial in the literature. Fujita and Ramey (forthcoming) have also calculated worker flows from the CPS. The Fujita and Ramey dataset does not encompass the same sample as the one by Shimer; it ranges from 1976:3 to 2004:4[25]. As stated by the authors, the standard deviation of the job separation rate is higher and the one of job finding is lower in their data series compared to Shimer. This suggests a larger role for the first series in the dynamics of unemployment. Job separation is also more persistent. The correlations of the job finding and separation rates with productivity are much lower than in the Shimer series. Figure 1.4 shows that the responses in both datasets are quite similar. Note that job separation decreases after a positive technology shock. However, this is mainly due to the subsample rather than the difference in the measurement of the data. In fact, results for the job separation rate are not robust to subsample choices or different specifications.

[25] I thank Shigeru Fujita for making the data available to me.

Figure 1.4: Shimer versus Fujita-Ramey

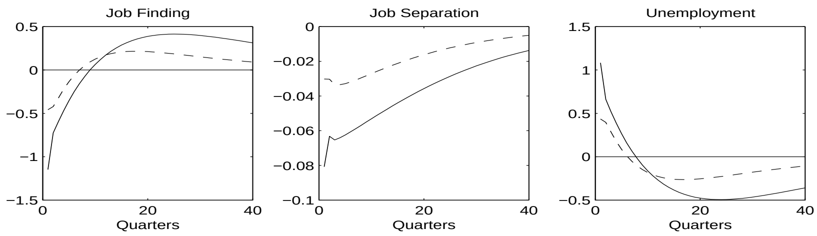


Notes: Solid lines depict Shimer data, broken lines show Fujita and Ramey data.
Responses in percentage points to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

### 1.5.2 Job flows

Instead of worker flows, so-called job flow data have often been used to assess the empirical validity of the standard labor market model (similar to Cole and Rogerson (1999) and Davis et al. (1998)). Note that from the perspective of the standard model job flows and worker flows are indistinguishable, i.e., when a worker moves into or out of a job, the job match is automatically created or destroyed. In the data, these two concepts show quite different unconditional business-cycle moments however, and hence it is interesting to consider conditional moments in job flows complementary to the above.

Here, I use data from Faberman (2006) which encompasses the fluctuations of jobs defined as small size units ("plants") that are created and destroyed within the U.S. manufacturing sector[26]. The resulting rates are usually referred to as job creation and destruction rates and both are measured in percent of employment. Unemployment dynamics are approximated by unemployment growth which results from taking the difference between the job destruction and creation rate. In the following, the same exercise as in the Fisher identification in Section 1.4 is repeated by using job flows rather than worker flows. Table 1.5 presents the conditional and unconditional moments from this set of data together with the familiar moments from the model.

Note that in this sample, that job destruction is about twice as volatile as job creation. Both series are less persistent than the worker flows, while the cross-correlations between the variables are qualitatively similar, but quite different in value from the ones in the

[26]The data is also described in Davis et al. (2006). I thank Jason Faberman for making the data available to me.

Table 1.5: Historical decomposition from Fisher identification - Job flows

| | Uncond. Sample | Model | Conditional Moments Inv. Tech. | Neu. Tech. | All Tech. | Residual |
|---|---|---|---|---|---|---|
| A: Standard Deviations | | | | | | |
| Creat. | 0.0765 | 0.0775 | 0.0455 | 0.0336 | 0.041 | 0.0732 |
| | | (0.0717) | (0.04,0.06) | (0.03,0.04) | (0.03,0.05) | (0.07,0.08) |
| Dest. | 0.1311 | | 0.0547 | 0.0473 | 0.0583 | 0.1214 |
| | | | (0.04,0.08) | (0.03,0.07) | (0.04,0.07) | (0.11,0.13) |
| Unemp. | 1.0612 | 0.0708 | 0.3921 | 0.2604 | 1.7027 | 3.9574 |
| | | (0.0657) | (0.25,0.57) | (0.17,0.40) | (1.27,2.17) | (3.78,4.16) |
| Prod. | 0.0156 | 0.0156 | 0.0174 | 0.0191 | 0.013 | 0.01 |
| | | (0.0129) | (0.01,0.02) | (0.02,0.02) | (0.01,0.01) | (0.01,0.01) |
| B: Autocorrelations | | | | | | |
| Creat. | 0.6177 | 0.8655 | 0.8254 | 0.9051 | 0.8226 | 0.6383 |
| | | (0.8671) | (0.75,0.90) | (0.81,0.96) | (0.76,0.89) | (0.60,0.67) |
| Dest. | 0.7222 | | 0.8247 | 0.6189 | 0.7751 | 0.7146 |
| | | | (0.63,0.88) | (0.45,0.82) | (0.65,0.86) | (0.70,0.74) |
| Unemp. | 0.6683 | 0.8607 | 0.8133 | 0.5611 | 0.9457 | 0.9455 |
| | | (0.8632) | (0.69,0.86) | (0.33,0.77) | (0.93,0.96) | (0.94,0.95) |
| Prod. | 0.8507 | 0.8482 | 0.8563 | 0.8141 | 0.864 | 0.8514 |
| | | (0.855) | (0.80,0.90) | (0.77,0.85) | (0.85,0.88) | (0.79,0.90) |
| C: Cross-Correlations | | | | | | |
| JC,P | 0.1545 | 0.5141 | 0.4224 | 0.2328 | 0.2206 | 0.0636 |
| | | (0.4087) | (0.21,0.57) | (-0.12,0.46) | (0.08,0.34) | (-0.08,0.24) |
| JD,P | -0.4733 | | -0.4225 | 0.2207 | -0.0073 | -0.6159 |
| | | | (-0.65,-0.02) | (-0.31,0.43) | (-0.22,0.23) | (-0.76,-0.45) |
| U,P | -0.4449 | -0.4427 | -0.5901 | 0.0538 | 0.1561 | -0.0678 |
| | | (-0.3506) | (-0.72,-0.35) | (-0.47,0.37) | (-0.13,0.40) | (-0.14,0.02) |
| JC,U | -0.7176 | -0.8718 | -0.6134 | -0.3034 | 0.2599 | 0.114 |
| | | (-0.8749) | (-0.79,-0.31) | (-0.56,-0.05) | (0.11,0.43) | (0.04,0.19) |
| JD,U | 0.9242 | | 0.7912 | 0.782 | 0.3496 | 0.1383 |
| | | | (0.59,0.90) | (0.58,0.89) | (0.25,0.45) | (0.10,0.18) |
| JC,JD | -0.4187 | | 0.0523 | 0.4158 | -0.0813 | -0.3764 |
| | | | (-0.33,0.49) | (0.07,0.65) | (-0.32,0.22) | (-0.42,-0.34) |

Notes: All series are detrended with the smooth HP-filter as in Shimer (2005a). The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution. The model is calibrated to match the unconditional standard deviation of labor productivity and the same figure that is conditional on both technology shocks (in brackets).

worker flow series. With regard to the empirical performance of the standard model based on unconditional second moments, this means that while the model now replicates the standard deviation of job creation (in fact the standard deviation is a little too high in the model), it does not mirror the volatility of the job destruction rate and hence unemployment. A natural extension of this model would include endogenous job destruction as in Mortensen and Pissarides (1998) or denHaan et al. (2000) in order to account for fluctuations in this variable. The model does not aim at explaining the positive correlation between productivity and job creation with unemployment.

Conditional on investment-specific and neutral technology shocks, the standard deviation in job creation is even smaller than the unconditional one. More importantly, the two technology shocks generate a standard deviation of job destruction and unemployment that is only about a third of the one in the data. Hence, job destruction does move after technology shocks, but most of its volatility stems from non-technological disturbances. This means that endogenous job destruction alone cannot realign the moments from the model with the unconditional moments. Complementary to this result, technology shocks explain only up to 17% of the business cycle variance of job creation and destruction as is exhibited in Table 1.9 in the Appendix to Chapter 1. My result supports the findings from the previous sections that an additional non-technological disturbance is needed in order to explain the fluctuations observed unconditionally.

Panel C of Table 1.5 depicts the conditional cross-correlations of the labor market variables with each other and productivity. Figure 1.7 in the Appendix to this chapter also visualizes the dynamics induced by the two technology shocks. Most importantly, job creation and labor productivity are positively correlated after both technology shocks. As a consequence, the "job finding puzzle" after a neutral technology innovation from before disappears. Unemployment still increases after a positive neutral shock, due to the strong increase in job destruction (This is also reflected in the positive co-movement of these variables with productivity). Even though insignificant, in a model with endogenous job destruction and vintage technologies, job destruction may increase after a positive shock to technology if it can only be used in newly formed jobs rendering many existing job matches technologically obsolete. Then, these effects provide a valid and easy explanation to the rise in unemployment or parallel the fall in hours after a technology shocks and, hence, to the hours puzzle documented by Galí (1999). Strikingly, investment-specific technology shocks induce dynamics that are different from the ones generated by neutral technology shocks and that are similar to those expected from the standard model: Job creation goes up and job destruction falls after a positive innovation in investment-specific technology. As a consequence, unemployment decreases before converging back to zero. The responses after the investment-specific shocks exhibit greater persistence than the ones after a neutral shock.[27] However, investment-specific technology shocks are not important enough to explain the unconditional moments. Again, an additional source of fluctuations is necessary here.

Are the results from the Fisher identification with worker and with job flows are truly comparable? Plotting the structural shocks from the two estimations and calculating their correlation, it is possible to see that the investment-specific shocks are almost identical in both specifications. The neutral shocks from both estimations are positively correlated (the correlation coefficient is about 0.6), but not identical. Alternatively, both job and worker flow data can be included into one common specification. This is also important in the light of the joint dynamics of these two data concepts which has been an issue in the literature. The results show that the effects of the neutral shock on job creation and job destruction hardly change[28]. To summarize, since the two data concepts not only generate quite different unconditional statistics, but also react differently to the estimated shocks, it seems reasonable to try to distinguish the different concepts and model the empirical dynamics of these two sets of data in a theoretical framework as well.

## 1.6 Alternative identification

### 1.6.1 Motivation and identification

This section investigates to which extend the results outlined above in sections 1.3 and 1.4.2 strongly rely on the imposed identification assumption for the technology shocks, or whether they are robust to an alternative identification scheme as well. To motivate, let us briefly return to the Galí identification of technology shocks. In fact, the identified Galí shocks have a significant and positive effect on the relative price of investment. These shocks are therefore negatively biased towards new investment and mistakenly labelled

[27]Michelacci and Lopez-Salido (2007) do a similar empirical exercise with job flow data. They document similar responses after a neutral technology shock, but different responses after an investment-specific technology shock due a different identification.

[28]Job creation drops on impact after a positive neutral technology shock, but then rises with a hump-shape above zero.

factor-neutral, see Figure 1.8 in the Appendix to Chapter 1[29].

The Fisher identification separates technology shocks that have an effect on the relative price of investment from technology shocks that do not have an effect on the relative price of investment and hence are truly investment-neutral. However, the Fisher identification disregards those shocks that have a positive effect on both productivity and the price. When estimated without the third restriction on the productivity effect of investment-specific shocks, these shocks are incorporated into the investment-specific technology shocks in the Fisher identification. The difference between the results from the Fisher identification with and without the third restriction documents that these shocks may play an important role in the overall dynamics of these two variables. More precisely, labor productivity falls in response to these unrestricted investment-specific technology shocks (see discussion in Section 1.4). Additionally, these unrestricted shocks produce labor market dynamics that are quite different from the ones generated by the restricted shocks. Namely, job finding increases in a hump-shape after a positive investment-specific technology shock and job separation falls. As a result, unemployment decreases.[30] The unrestricted shocks also play a much larger role for the business cycle variance of the labor market variables than the restricted shocks.

Against this background, I propose an alternative identification of technology shocks which separates investment-specific technology shocks from those other shocks. The identification strategy imposes the following assumptions:

1. Technology shocks are assumed to be the only shocks that affect the relative price of investment and labor productivity in the long run.
2. Out of these shocks, investment-specific technology shocks are those shocks that affect labor productivity positively and the relative price of investment negatively in the long run.
3. Out of these shocks, the remaining shocks may affect labor productivity positively and the relative price of investment positively in the long run.

These assumptions are implemented with a mixture of long-run zero and sign restrictions similar to the Galí and Fisher identifications. I order the relative price of investment and

[29]Chapter 2 documents that these shocks are not only biased negatively towards investment, but also towards skilled labor.

[30]See Figure 1.9 in the Appendix to this chapter.

labor productivity first in the VAR and impose zero restrictions on the long-run effects of all but the first two shocks on these variables. Sign restrictions similar as in Peersman (2005) are then applied to the upper left 2-by-2 system of the long-run horizon forecast revision matrix according to the restrictions outlined above. The remaining elements of the long-run effects can then be calculated subsequently.[31].

Figure 1.10 in the Appendix to Chapter 1 visualizes the assumed responses of price and productivity to the two newly identified shocks. Not surprisingly, the new shocks turn out to be negatively biased towards investment and may consequently called investment-unspecific technology shocks. Note that the Galí, Fisher and the alternative identification strategies all offer an alternative decomposition of the long-run variance of the investment price and productivity[32]. The Fisher and Galí identification each impose an extra zero restriction on this system. This means that by construction the Fisher identification does not deliver shocks that induce the same effect on the price and productivity as the Galí identification. Thus, the Fisher identification does not provide a decomposition of the Galí technology shocks. My alternative identification is more closely related to the Galí identification as this scheme decomposes Galí's productivity shocks into investment-specific and -unspecific shocks. I can now test Fisher's third identifying assumption based on the effect of the first shock in a more general context in which all shocks are in fact orthogonal. Further, I can assess the importance of those shocks that resulting from the unrestricted Fisher identification might have been labelled investment-specific technology shocks by mistake and can explore their properties. However, it is no longer possible to distinguish between investment-specific and investment-neutral shocks in this setup.

What are technology shocks that drive the relative price of investment up? In the model outlined in Section 1.2, shocks that have a positive effect on the relative price of investment negatively affect labor productivity and, hence, are not technology shocks. As a consequence, the model outlined above does not accommodate these shocks and it is therefore not clear how to interpret them in this context. Chapter 2 suggests to identify technology shocks which originate in the labor market. More precisely, it is documented that technology shocks that are biased towards skilled labor have a positive effect on the relative price of investment and could therefore capture the variation of the data documented

[31]For further details of the implementation of the long-run sign restrictions are contained in Appendix A.4.

[32]This is true if the price is ordered second in the Galí identification. The remaining elements of the first two rows of this matrix are always zero.

here.[33] Once more, this points to the use of a more complex production function with which it is possible to distinguish between low and high skilled labor in order to replicate the empirical dynamics.

### 1.6.2 Results

Table 1.11 in the Appendix to this chapter exhibits the historical decompositions for this identification scheme. Regarding volatility, the standard deviations conditional on investment-specific technology shocks are very close to the results from the Fisher identification. The two identified technology shocks together generate a conditional standard deviation that is again less than half of the unconditional standard deviation in job finding, separation and unemployment. This is not surprising, since the alternative identification is just a different decomposition of the technology shocks from the other identification schemes.

More interesting in this respect are the labor market dynamics induced by the two new shocks documented in Figure 1.5 and Table 1.11. For both types of shocks, job finding drops and unemployment increases supporting the findings of the Fisher and Galí identification. There are significant differences between the responses of the two shocks however. After an investment-specific productivity shock job separation does not move significantly. Note that the dynamics of this shock are very similar to the ones I have documented for the restricted Fisher investment-specific technology shocks. Indeed, the estimated relationship between the effect of this shock on the price and productivity is very close to the one imposed via the third restriction. After an investment-unspecific shock job finding does not react on impact and subsequently decreases in a hump-shape, job separation significantly rises and the rising unemployment inherits the hump-shape from the effects on the job finding rate[34].

The variance decomposition in Table 1.10 in the Appendix to this chapter sheds light on the

[33] The identification of these shocks originates in the effect of technological progress on the skill premium in a model which allows for both skilled and unskilled labor in production. The fact that the investment price increases in responses to these shocks provides evidence for capital-skill substitutability in the data.

[34] Note that the inverse of this shock is an investment-specific technology shock with a negative effect on productivity. The resulting dynamics are strikingly close the the ones from the unrestricted Fisher identification, see Figure 1.9 in the Appendix to Chapter 1 or Canova et al. (2007). This means that the major part of the unrestricted investment-specific technology shocks consists of shocks that do not positively affect labor productivity and are consequently not in line with our model.

Figure 1.5: Productivity shocks from sign restrictions

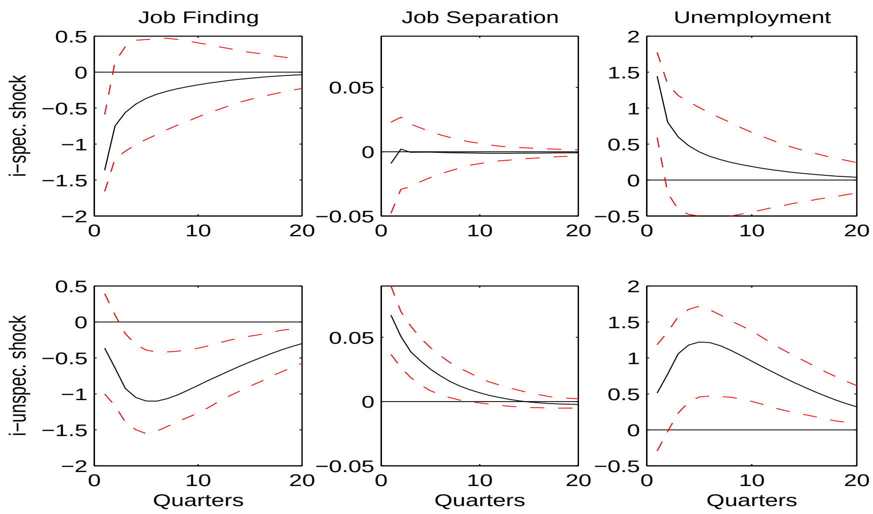


Notes: Responses in percentage points to a one-standard deviation shock.
Confidence intervals are 68% Bayesian bands.

relative importance of investment-specific -unspecific technology shocks. The investment-unspecific technology shock is more important for the business cycle variance of labor productivity than the investment-specific technology shock. The investment-specific technology shock explains more of the variance of the relative investment price in the first two horizons, while the investment-unspecific shock is more important in the longer run. This means that a substantial part of the dynamics in the unrestricted investment-specific shocks are not driven by positive productivity shocks and this highlights the importance of distinguishing between the two types of shocks. The investment-unspecific shock explains a substantial fraction of the job finding and separation rate and consequently unemployment. This shock is generally more important for the business cycle variation of the labor market variables than the investment-specific technology shock. Together, both shocks explain about 20% of the business cycle variation in job finding and unemployment.

Investment-unspecific technology shocks have not been identified so far. The reason clearly lies in the fact that they are difficult to interpret in the context of a standard model as the one outlined in Section 1.2. Here I have shown that they carry some weight with respect to the dynamics in the labor market. As argued above, these shocks reflect skill-biased technology shocks as identified in Chapter 2. Skill-biased technology shocks have a negative effect on total hours worked and thus induce similar dynamics to the shocks identified here.

## 1.7 Conclusion

Starting from the recent ongoing debate on the empirical performance of the Mortensen-Pissarides search-and-matching model, this study provides an important contribution to the debate as it judges the empirical performance of the model on basis of moments conditional on technology shocks rather than on unconditional moments. My analysis breaks down the second moments of labor productivity, the job finding, job separation and unemployment rate into the contribution of technology and non-technology shocks. These shocks are identified within a SVAR framework with conventional long-run restrictions and a combination of long-run zero and sign restrictions.

I find that technology shocks cannot be the source of the high volatility in the job finding rate and unemployment present in the data. As a result, the standard deviation of these variables that is generated from a standard model replicates the volatility conditional on technology shocks. A large part of the volatility remains unexplained in the residual from the structural estimation. This residual might be called non-technology or demand shock. In order to mirror the overall volatility in the data, the model should be augmented with an additional non-technological source of volatility rather than with respect to the propagation of technology shocks as proposed by Shimer. Ravn and Simonelli (2006) identify government spending shocks in a similar SVAR. Their shocks indeed mirror the dynamics of our "residual" disturbances as they drive labor productivity and labor market tightness up and unemployment down. Barnichon (2008) argues that these shocks are shocks to monetary policy. Here, I investigate an idea by Hall (1997) that preference shocks in the form of shocks to the marginal rate of substitution between consumption and leisure are important for labor market dynamics. These shocks in fact add a lot of volatility to the model.

Technology shocks induce a negative co-movement between job finding and productivity and a positive co-movement between unemployment and productivity, while the respective figures in the overall sample are directly the opposite. Put differently, job finding falls and importantly contributes to an increase in unemployment after a positive technology shock. This result contradicts the effects generated in the standard search-and-matching model. Chapter 2 contains evidence that these effects may be explained through a distinction between high- and low-skilled labor in production. Since the identified technology shocks are (possibly) biased towards the productivity of high-skilled labor, low-skilled labor gets substituted out of production. Further results in the following chapter show that the "job finding puzzle" vanishes when job flow data rather than worker flow data are employed in the specification. In any case, additional non-technological disturbances are needed in order to replicate the unconditional correlation between productivity, the job finding rate and unemployment.

In different specifications, I distinguish technology shocks that are factor-neutral or investment-specific as in Galí (1999) and Fisher (2006). I document that the two main results are robust to these extensions. The role of technology shocks for labor market dynamics is further assessed through a distinction of positive productivity shocks that have either a negative or a positive effect on the relative price of investment. The latter my be called investment-unspecific technology shocks. First, this identification tests and verifies a critical assumption in the Fisher identification on the effect of investment-specific technology shocks on labor productivity. Second, this procedure investigates the relationship between constrained and unconstrained investment-specific technology shocks. I find that investment-unspecific technology shocks might by mistakenly labelled investment-specific in the unconstrained identification. In addition, these shocks play a significant role for labor market fluctuations. However, these shocks cannot be interpreted in the context of the standard model. It will be shown in chapter 2 that it is reasonable to assume that these shocks are the same as skill-biased technology shocks in their paper. Technology shocks that are skill-biased induce similar dynamics in the investment price and the labor market as the shocks identified here. This result again provides empirical foundation to allowing for a more sophisticated production function in this class of model in which low- and high-skilled labor are substitutable in production.

# Appendix to Chapter 1: Additional Tables and Graphs

Table 1.6: Variance decomposition in Fisher identification

| | Investment-specific Shock | | | | Neutral Shock | | | |
|---|---|---|---|---|---|---|---|---|
| Quarters | 1 | 8 | 16 | 32 | 1 | 8 | 16 | 32 |
| Price | 59.27 | 78.02 | 86.83 | 93.30 | 6.55 | 4.47 | 3.07 | 1.58 |
| | (31,79) | (52,91) | (68,95) | (83,98) | (1,27) | (1,20) | (0,14) | (0,7) |
| Productivity | 13.50 | 14.12 | 12.94 | 11.46 | 68.33 | 76.50 | 82.49 | 86.43 |
| | (8,19) | (11,18) | (11,16) | (10,13) | (50,78) | (67,82) | (77,86) | (84,88) |
| JFinding | 15.92 | 6.73 | 6.23 | 6.28 | 46.86 | 42.34 | 42.98 | 42.70 |
| | (8,23) | (4,12) | (3,11) | (3,11) | (28,59) | (18,58) | (19,58) | (19,58) |
| JSeparation | 1.87 | 3.02 | 3.46 | 3.62 | 19.27 | 21.26 | 21.15 | 21.59 |
| | (0,9) | (1,11) | (1,11) | (1,11) | (3,41) | (4,43) | (5,43) | (5,43) |
| Unemployment | 15.19 | 6.38 | 6.00 | 6.02 | 49.44 | 43.54 | 43.86 | 43.48 |
| | (8,22) | (3,12) | (3,11) | (3,10) | (31,61) | (19,59) | (20,59) | (19,59) |

Notes: The values for the investment-specific shock, the neutral shock and the (omitted) residual disturbances add up to 100 for each variable at each time horizon. The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution.
All numbers are percent.

Table 1.7: Galí identification with standard detrending

| | Uncond. Sample | Model I | Model II | Conditional Moments Technology | Conditional Moments Residual |
|---|---|---|---|---|---|
| A: Standard Deviations | | | | | |
| JFinding | 0.1019 | 0.0407 | 0.0251 | 0.0327 | 0.0835 |
| | | | | (0.02,0.05) | (0.07,0.10) |
| JSeparation | 0.0497 | | | 0.0255 | 0.0444 |
| | | | | (0.02,0.03) | (0.04,0.05) |
| Unemployment | 0.1181 | 0.041 | 0.0252 | 0.0479 | 0.0928 |
| | | | | (0.03,0.07) | (0.08,0.11) |
| Productivity | 0.0105 | 0.0105 | 0.0066 | 0.0066 | 0.009 |
| | | | | (0.00,0.01) | (0.00,0.01) |
| B: Autocorrelations | | | | | |
| JFinding | 0.8137 | 0.8008 | 0.8031 | 0.7939 | 0.7688 |
| | | | | (0.68,0.87) | (0.71,0.80) |
| JSeparation | 0.4409 | | | 0.7408 | 0.3913 |
| | | | | (0.66,0.83) | (0.32,0.44) |
| Unemployment | 0.8345 | 0.6784 | 0.6791 | 0.7325 | 0.8136 |
| | | | | (0.64,0.79) | (0.79,0.83) |
| Productivity | 0.6881 | 0.6651 | 0.6651 | 0.7161 | 0.7286 |
| | | | | (0.64,0.80) | (0.70,0.76) |
| C: Cross-Correlations | | | | | |
| JFind.,Prod. | 0.1443 | 0.9522 | 0.9532 | -0.7619 | 0.5986 |
| | | | | (-0.87,-0.48) | (0.45,0.72) |
| JSep.,Prod. | -0.4826 | | | 0.4837 | -0.6975 |
| | | | | (0.22,0.63) | (-0.80,-0.60) |
| Unemp.,Prod. | -0.3051 | -0.6943 | -0.696 | 0.7441 | -0.8329 |
| | | | | (0.55,0.84) | (-0.89,-0.72) |
| JFind.,Unemp. | -0.9254 | -0.8405 | -0.8408 | -0.908 | -0.8984 |
| | | | | (-0.96,-0.75) | (-0.92,-0.86) |
| JSep.,Unemp. | 0.6346 | | | 0.8453 | 0.5455 |
| | | | | (0.68,0.91) | (0.44,0.61) |
| JFind.,JSep. | -0.2947 | | | -0.5102 | -0.1169 |
| | | | | (-0.71,-0.04) | (-0.22,0.05) |

Notes: All series are detrended with the HP-Filter with $\lambda = 1600$. The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution. Calibration I of the model matches the unconditional standard deviation of labor productivity, calibration II matches the same moment, conditional on technology shocks.

Table 1.8: Historical decomposition of Fisher identification

| | Uncond. Sample | Model | Conditional Moments Inv. Tech. | Neu. Tech. | All Tech. | Residual |
|---|---|---|---|---|---|---|
| A: Standard Deviations | | | | | | |
| Find. | 0.1542 | 0.0775 | 0.0684 | 0.0741 | 0.0671 | 0.1283 |
| | | (0.0717) | (0.05,0.09) | (0.05,0.11) | (0.05,0.09) | (0.11,0.15) |
| Sep. | 0.062 | | 0.0401 | 0.048 | 0.0512 | 0.0543 |
| | | | (0.03,0.05) | (0.04,0.06) | (0.04,0.06) | (0.05,0.06) |
| Unemp. | 0.1786 | 0.0708 | 0.0658 | 0.0996 | 0.088 | 0.1434 |
| | | (0.0657) | (0.05,0.09) | (0.06,0.14) | (0.07,0.12) | (0.12,0.17) |
| Prod. | 0.0156 | 0.0156 | 0.0185 | 0.0184 | 0.0129 | 0.016 |
| | | (0.0129) | (0.01,0.02) | (0.02,0.02) | (0.01,0.01) | (0.01,0.02) |
| B: Autocorrelations | | | | | | |
| Find. | 0.9128 | 0.8655 | 0.7116 | 0.8182 | 0.8771 | 0.9009 |
| | | (0.8671) | (0.62,0.82) | (0.69,0.89) | (0.81,0.92) | (0.87,0.92) |
| Sep. | 0.6336 | | 0.9245 | 0.8984 | 0.8757 | 0.6389 |
| | | | (0.85,0.96) | (0.83,0.95) | (0.82,0.92) | (0.59,0.70) |
| Unemp. | 0.9218 | 0.8607 | 0.7692 | 0.8326 | 0.9045 | 0.9143 |
| | | (0.8632) | (0.67,0.88) | (0.74,0.88) | (0.87,0.93) | (0.90,0.92) |
| Prod. | 0.8507 | 0.8482 | 0.9055 | 0.8597 | 0.8909 | 0.9253 |
| | | (0.855) | (0.85,0.95) | (0.80,0.91) | (0.87,0.92) | (0.91,0.94) |
| C: Cross-Correlations | | | | | | |
| JF,P | 0.0567 | 0.5141 | -0.1674 | -0.5569 | -0.3274 | 0.6979 |
| | | (0.4087) | (-0.38,0.11) | (-0.29,-0.70) | (-0.55,0.01) | (0.57,0.79) |
| JS,P | -0.4392 | | -0.4355 | 0.2757 | 0.2059 | -0.6298 |
| | | | (-0.61,-0.21) | (0.03,0.46) | (-0.02,0.38) | (-0.73,-0.53) |
| U,P | -0.1858 | -0.4427 | -0.0838 | 0.5323 | 0.3431 | -0.821 |
| | | (-0.3506) | (-0.44,0.19) | (0.27,0.67) | (0.03,0.55) | (-0.89,-0.72) |
| JF,U | -0.9558 | -0.8718 | -0.8394 | -0.9147 | -0.8606 | -0.9409 |
| | | (-0.8749) | (-0.92,-0.72) | (-0.79,-0.97) | (-0.94,-0.75) | (-0.91,-0.95) |
| JS,U | 0.6845 | | 0.3897 | 0.794 | 0.7584 | 0.5997 |
| | | | (0.06,0.65) | (0.60,0.88) | (0.58,0.85) | (0.51,0.66) |
| JF,JS | -0.4404 | | 0.2296 | -0.4877 | -0.3075 | -0.2893 |
| | | | (-0.12,0.52) | (-0.17,-0.69) | (-0.58,0.11) | (-0.17,-0.38) |

Notes: All series are detrended with the smooth HP-filter as in Shimer (2005a). The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution. The model is calibrated to match the unconditional standard deviation of labor productivity and the same figure that is conditional on both technology shocks (in brackets).

Table 1.9: Variance decomposition in Fisher identification - Job flows

| | Investment-specific Shock | | | | Neutral Shock | | | |
|---|---|---|---|---|---|---|---|---|
| Quarters | 1 | 8 | 16 | 32 | 1 | 8 | 16 | 32 |
| Price | 76.39 | 92.80 | 96.60 | 98.39 | 4.44 | 0.91 | 0.42 | 0.20 |
| | (54,90) | (82,98) | (91,99) | (96,100) | (0,19) | (0,5) | (0,2) | (0,1) |
| Productivity | 12.15 | 11.94 | 11.01 | 10.50 | 80.46 | 85.85 | 87.87 | 88.94 |
| | (9,15) | (11,13) | (10,12) | (10,11) | (73,85) | (84,88) | (87,89) | (88,89) |
| JCreation | 6.32 | 6.84 | 7.04 | 7.05 | 3.93 | 10.19 | 10.45 | 10.45 |
| | (1,14) | (3,13) | (3,12) | (3,12) | (0,15) | (3,24) | (3,24) | (3,24) |
| JDestruction | 1.37 | 4.60 | 4.66 | 4.66 | 15.77 | 11.79 | 11.81 | 11.81 |
| | (0,5) | (2,12) | (2,12) | (2,12) | (2,40) | (4,31) | (4,31) | (4,31) |
| Unemployment | 1.35 | 6.12 | 6.12 | 6.12 | 8.20 | 9.11 | 9.28 | 9.27 |
| | (0,6) | (2,13) | (2,13) | (2,13) | (1,26) | (3,21) | (3,22) | (3,22) |

Notes: The values for the investment-specific shock, the neutral shock and the (omitted) residual disturbances add up to 100 for each variable at each time horizon. The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution. All numbers in percent.

Table 1.10: Variance decomposition in sign identification

| | Investment-specific Shock | | | | Investment-unspecific Shock | | | |
|---|---|---|---|---|---|---|---|---|
| Quarters | 1 | 8 | 16 | 32 | 1 | 8 | 16 | 32 |
| Productivity | 24.66 | 28.25 | 29.67 | 31.68 | 46.85 | 59.75 | 63.89 | 65.10 |
| | (2,59) | (3,67) | (3,70) | (3,74) | (18,77) | (23,86) | (24,91) | (24,93) |
| Price | 27.75 | 35.71 | 38.80 | 37.82 | 11.45 | 24.53 | 39.06 | 51.75 |
| | (7,52) | (10,61) | (9,68) | (8,74) | (1,36) | (4,53) | (9,69) | (16,82) |
| JFinding | 16.86 | 6.44 | 6.00 | 5.90 | 3.54 | 9.88 | 12.93 | 13.42 |
| | (4,33) | (2,18) | (2,17) | (2,17) | (0,14) | (3,28) | (4,31) | (4,31) |
| JSeparation | 2.69 | 2.98 | 3.06 | 3.10 | 17.87 | 15.26 | 14.50 | 14.51 |
| | (0,13) | (1,12) | (1,12) | (1,12) | (6,36) | (5,34) | (5,33) | (5,32) |
| Unemployment | 16.38 | 6.35 | 5.91 | 5.83 | 4.09 | 10.62 | 13.61 | 14.00 |
| | (4,34) | (2,18) | (2,17) | (2,17) | (0,16) | (3,30) | (4,32) | (4,32) |

Notes: The values for the investment-specific shock, the investment-unspecific shock and the (omitted) residual disturbances add up to 100 for each variable at each time horizon. The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution.
All numbers are percent.

Table 1.11: Historical decomposition of sign identification

| | Uncond. Sample | Conditional Moments I-Specific | I-Unspecific | Both Shocks | Residual |
|---|---|---|---|---|---|
| A: Standard Deviations | | | | | |
| Find. | 0.1542 | 0.0456 | 0.051 | 0.0643 | 0.1242 |
| | | (0.04,0.07) | (0.04,0.07) | (0.05,0.09) | (0.10,0.15) |
| Sep. | 0.062 | 0.0408 | 0.0499 | 0.0527 | 0.0535 |
| | | (0.03,0.05) | (0.04,0.06) | (0.04,0.06) | (0.05,0.06) |
| Unemp. | 0.1786 | 0.0538 | 0.0742 | 0.088 | 0.139 |
| | | (0.04,0.08) | (0.05,0.10) | (0.07,0.11) | (0.12,0.16) |
| Prod. | 0.0156 | 0.0122 | 0.0109 | 0.0127 | 0.0156 |
| | | (0.01,0.01) | (0.01,0.01) | (0.01,0.01) | (0.01,0.02) |
| B: Autocorrelations | | | | | |
| Find. | 0.9128 | 0.8091 | 0.9436 | 0.8653 | 0.9028 |
| | | (0.70,0.91) | (0.90,0.96) | (0.79,0.91) | (0.87,0.92) |
| Sep. | 0.6336 | 0.9374 | 0.8886 | 0.8634 | 0.6507 |
| | | (0.88,0.96) | (0.83,0.94) | (0.80,0.92) | (0.59,0.71) |
| Unemp. | 0.9218 | 0.897 | 0.9185 | 0.8992 | 0.9137 |
| | | (0.83,0.95) | (0.89,0.95) | (0.87,0.92) | (0.90,0.92) |
| Prod. | 0.8507 | 0.92 | 0.9381 | 0.8929 | 0.9225 |
| | | (0.88,0.97) | (0.89,0.98) | (0.87,0.92) | (0.90,0.94) |
| C: Cross-Correlations | | | | | |
| JF,P | 0.0567 | 0.003 | -0.0897 | -0.3597 | 0.7118 |
| | | (-0.46,0.29) | (-0.47,0.21) | (-0.53,-0.04) | (0.58,0.80) |
| JS,P | -0.4392 | -0.1501 | -0.1297 | 0.235 | -0.6269 |
| | | (-0.58,0.33) | (-0.51,0.30) | (-0.02,0.40) | (-0.73,-0.54) |
| U,P | -0.1858 | -0.1624 | -0.0406 | 0.3822 | -0.8218 |
| | | (-0.65,0.46) | (-0.48,0.44) | (0.05,0.56) | (-0.91,-0.76) |
| JF,U | -0.9558 | -0.7386 | -0.8048 | -0.8396 | -0.9408 |
| | | (-0.90,-0.55) | (-0.92,-0.61) | (-0.93,-0.70) | (-0.95,-0.91) |
| JS,U | 0.6845 | 0.6339 | 0.7937 | 0.7583 | 0.5913 |
| | | (0.33,0.86) | (0.64,0.88) | (0.64,0.86) | (0.51,0.66) |
| JF,JS | -0.4404 | 0.1652 | -0.2492 | -0.2512 | -0.2781 |
| | | (-0.46,0.47) | (-0.59,0.19) | (-0.57,0.04) | (-0.38,-0.15) |

Notes: All series are detrended with the smooth HP-filter as in Shimer (2005a). The point estimate is the median, the confidence intervals are 68% Bayesian bands from the posterior distribution.

Figure 1.6: Restricted and unrestricted Fisher identification

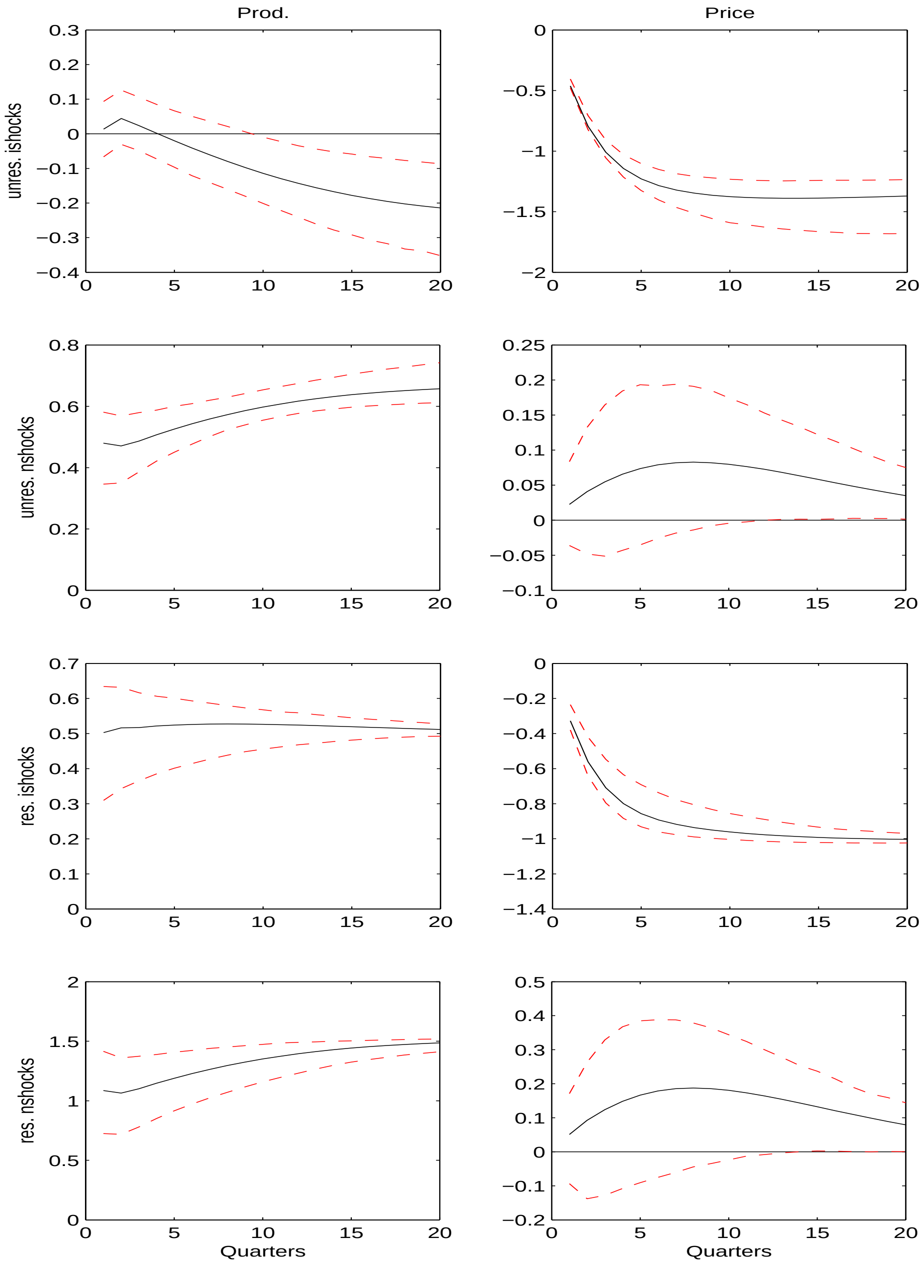


Notes: Responses in percent to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 1.7: Job flow responses to Fisher technology shocks

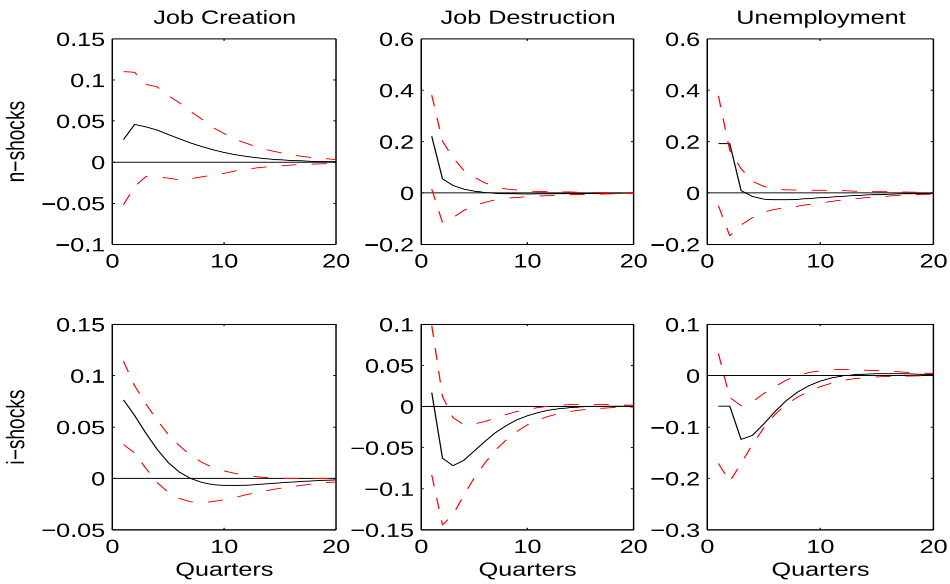


Notes: Percentage responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 1.8: Galí identification - price and productivity

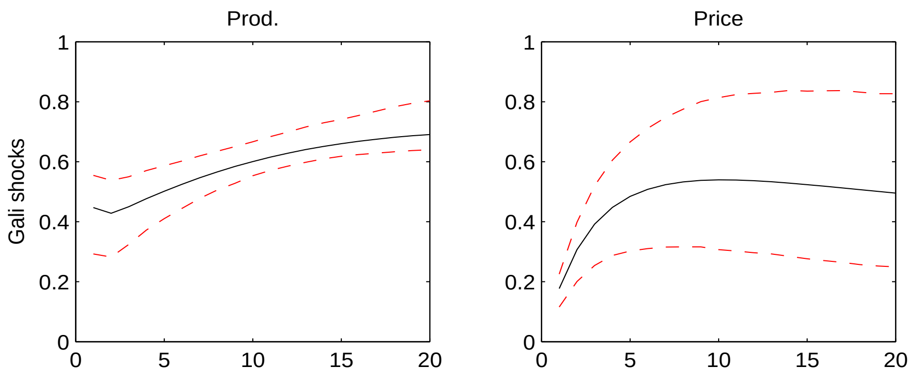


Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 1.9: Unrestricted Fisher technology shocks

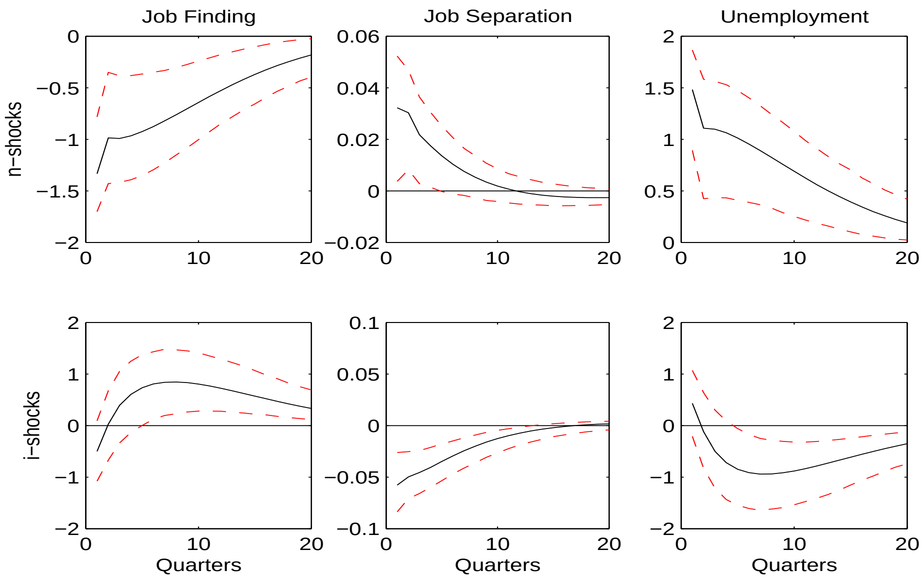


Notes: Percentage responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 1.10: Sign identification - price and productivity

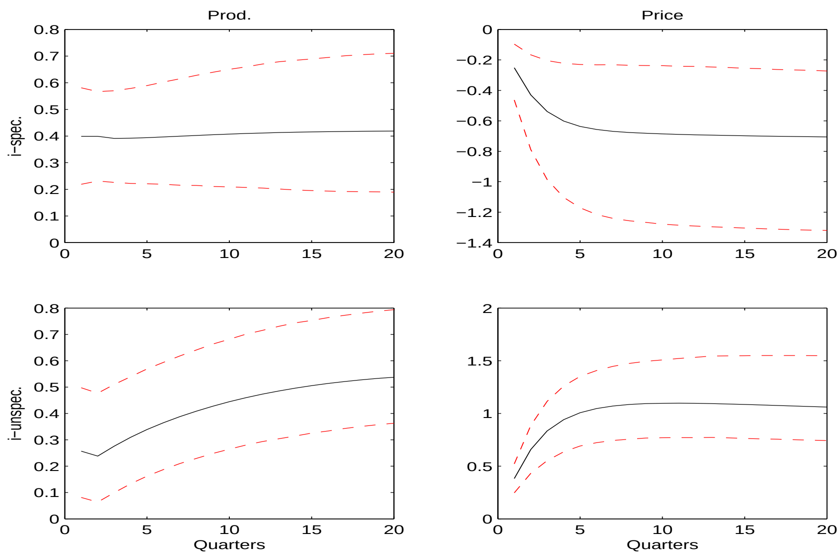


Notes: Responses in percentage points to a one-standard deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 1.11: Fisher technology shocks - no trend

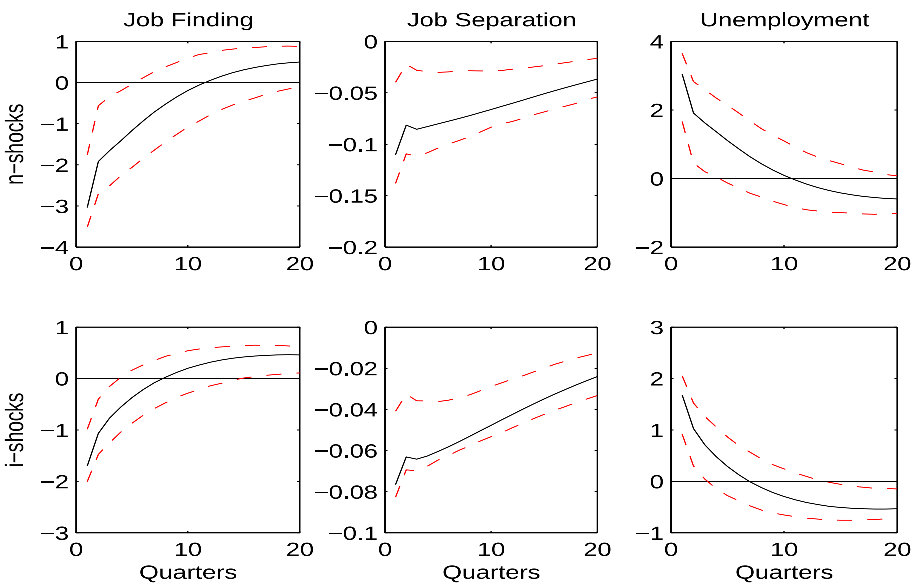


Notes: Percentage responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

# Chapter 2

# On the Implications of Skill-Biased Technological Progress for the Business Cycle

## 2.1 Introduction

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

At the same time, shocks to technological progress have been attributed to be an important driving force of US business cycles. Conventional technology shocks commonly referred to as a source of the business cycle may either be factor-neutral or biased towards new investments (investment-specific technology shocks). In this paper, we relate these two phenomena by exploring the implications of skill-biased technological change for business cycle fluctuations. We will consider this issue from two angles: First, we investigate the relationship between technology and the skill premium, and hence skill-bias, over the business cycle. Second, we propose an empirical strategy to identify skill-biased, and complementary also skill-neutral, technological change directly. Over and above linking the conventional driving forces of business-cycles to the developments in the labor market, this paper therefore attempts to identify sources of cycles that originate in the labor market itself. Our approach allows us to address many important questions in this respect: Does skill-biased technological progress play an important role for the business cycle? What does skill-biased technological progress imply for the business-cycle dynamics of macroeconomic aggregates? How are the production inputs capital, high and low skilled labor related over the business cycle?

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

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

[1]Lindquist (2004) also construct a quarterly series for the skill premium from the CPS outgoing rotation groups, but does not control for multiple education levels and other sources of worker heterogeneity, see section 2.2.3.

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

shocks on the premium.

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

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

This result rises the question whether all technological changes are skill-biased or whether there is a difference between skill-biased and skill-neutral technology shocks. We propose a strategy to identify skill-biased technology (SBT) shocks from a long-run restriction, arguing that SBT shocks are the only shocks that affect the skill premium in the long run. Skill-neutral technology shocks are all remaining sources of permanent changes in labor productivity. Skill-biased technology shocks are similar to skill-neutral technological changes in that they increase labor productivity. However, they have different implications for other aggregate variables. In particular, a positive SBT shock leads to a much larger reduction in total hours worked than a skill-neutral technology shock. In addition, SBT shocks increase the supply of skill in the long run, as we would expect, whereas skill-neutral shocks lead to reduced supply of skill. For robustness, we show that the impulse-responses that result from a decomposition of a production function that allows for inputs of high and low skilled labor as well as capital are similar to the ones from our estimated SBT shocks.

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

We identify investment-specific technology shocks, following Fisher (2006), as the only shocks that affect the relative price of investment in the long run. An investment-specific improvement in technology lowers the relative price of investment goods. The remaining shocks that affect labor productivity in the long run, are then investment-neutral technology shocks. We find that investment-specific technology shocks have a significant, but negative effect on the skill premium, while investment-neutral technology shocks have a positive effect on this variable. Conversely, skill-biased technology shocks, identified as described above, raise the relative price of investment goods. This evidence is in direct contradiction with the hypothesis of capital-skill complementarity, suggesting instead that capital and skill are (to some degree) substitutes in the aggregate production process. We support this result by simulating data from a model with different degrees on complementarity and substitutability between capital and skilled labor and estimating shocks and responses from these data with our structural VAR.

The remainder of this paper is organized as follows. Section 2.2 describes our empirical approach. First we define the different shocks to the production technology that we consider, then we discuss how to identify the effects of these shocks using long-run restrictions. We also describe the data that are necessary to estimate these effects and show some descriptive statistics on the cyclicality of our quarterly series for the skill premium and the relative supply of skill. In Section 2.3 we discuss skill-biased in technology shocks based on the structural VAR analysis and the production function decomposition. Section 2.4 deals with investment-specific technology shocks and capital-skill complementarity. Section 2.5 concludes.

## 2.2 Empirical approach

In this Section, we outline our approach to estimate the implications of skill-biased technological progress for the business cycle. We start by defining different types of of technological change, discussing various specifications for the aggregate production function. Next, we explain how to identify these different technology shocks from the data using 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.

### 2.2.1 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 technology shocks, the literature has typically assumed an aggregate production function of the following form (see e.g. Katz and Murphy (1992), Katz and Autor (1999), Autor et al. (2008)).

$$Y_t = A_t K_t^{\alpha} \left[ \beta \left( B_t H_t \right)^{\frac{\sigma-1}{\sigma}} + (1-\beta) L_t^{\frac{\sigma-1}{\sigma}} \right]^{\frac{(1-\alpha)\sigma}{\sigma-1}} \tag{2.1}$$

Here, $A_t$ is total factor productivity and $B_t$ is skilled labor augmenting technology. An increase in $B_t$ can be skill or unskill biased, depending on the elasticity of substitution between skilled and unskilled labor $\sigma > 0$. If high and low skilled labor are substitutes rather than complements ($\sigma > 1$), the substitution effect of improvements in skilled labor augmenting technology dominates the income effect so that an increase in $B_t$ increases the demand for skill and therefore the skill premium (assuming the supply curve for skill is downward sloping). The consensus estimate for $\sigma$ is around 1.5 (see Katz and Murphy (1992), Ciccone and Peri (2006), Teulings and van Rens (2008)), so that we can think of skill-biased technology shocks as increases in $B_t$.

There are two ways to interpret skill-biased technology shocks to an aggregate production function as in (2.1). If the production function for all goods in the economy is the same, then we can think of an increase in $B_t$ as a technological development that makes skilled labor more productive in all sectors. Alternatively, we may think that the production in different sectors $i$ requires skilled labor in different proportions $\beta_i$ of total labor input. In this case, even if skilled and unskilled labor are neither substitutes nor complements within each sector,[3] a sector-specific technology shock to a skill-intensive sector could still increase the skill premium.

A particularly interesting case is an economy that consists of a consumption goods producing sector and an investment goods producing sector. In this economy there are two mechanisms, by which sector-specific shocks may affect the skill premium. First, the input shares for skill might be different across the two sectors as explained above. Because investment goods are used to build up capital, which is an input in the production process, sector-specific shocks may also affect the capital-labor ratio used in production. If capital and skill are complements, as argued by Krusell et al. (2000), then a higher capital labor ratio increases the relative demand for skilled labor and therefore the skill premium.

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

$$Y_t = C_t + p_t I_t \tag{2.2}$$

where the relative price of investment goods $p_t$ reflects technological improvements in the investment goods producing sector. An aggregate production function that allows for capital-skill complementarity is a slightly generalized version of (2.1).

$$Y_t = A_t \left[ \beta \left( \gamma K_t^{\frac{\rho-1}{\rho}} + (1-\gamma)(B_t H_t)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}\frac{\sigma-1}{\sigma}} + (1-\beta) L_t^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{2.3}$$

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

[3] This is the case where $\sigma_i = 1$ for all $i$. In the limit for $\sigma \to 1$, production function (2.1) becomes Cobb-Douglas, so that changes in $B_t$ are indistinguishable from changes in $A_t$.

### 2.2.2 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 (2.1), we get the following expression,

$$\log\left(\frac{w_{H,t}}{w_{L,t}}\right) = \log\left(\frac{\beta}{1-\beta}\right) - \frac{1}{\sigma}\log\left(\frac{H_t}{L_t}\right) + \frac{\sigma-1}{\sigma}\log B_t \tag{2.4}$$

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

Changes in skill-biased technology $B_t$ represent shifts of the skill demand curve or skill demand shocks. Since the skill premium and the relative quantity of skill are observable, these shocks can be directly retrieved from equation (2.4), using an estimate for the elasticity of substitution between low and high skilled workers $\sigma$.[4] The estimates for the skill-biased technology shocks obtained this way are identified from the assumption that wages are proportional to marginal products. A sufficient condition for this assumption is that labor markets are perfectly competitive, in which case the wage of all workers equals their marginal product. If there are frictions in the labor market, the weaker assumption that wages are proportional to marginal products still holds approximately. However, if there are frictions in the wage determination process, then wages may deviate from marginal products in the short run. Therefore, we alternatively identify technology shocks using a structural VAR with long-run restrictions, as suggested by Blanchard and Quah (1989) and first used to estimate technology shocks by Galí (1999).

Consistent with equation (2.4), we identify skill-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 (2.4) holds in each period and makes the estimates robust to wage rigidity for example. 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 $\sigma$.[5] Thus, we use long run restrictions

[4] An estimate for the share parameter $\beta$ is unnecessary since this parameter affects only the level of $B_t$ and we normalize the mean and variance of the shocks to zero and one respectively.

[5] Of course the assumption is not valid for all production functions. For example, with capital-skill

in all our estimates, although we compare the results to a direct decomposition using equation (2.4), see Section 2.3.3, and find that for the simplest estimates the differences are not large.

The estimation of structural shocks using long run restrictions is implemented in two steps. First, we estimate a reduced form VAR in the variables labor productivity, hours worked, the skill premium and in some specifications also the relative price of investment goods. Second, we map the reduced form coefficients and residuals into structural coefficients and shocks normalizing the variance of all structural shocks to one and assuming orthogonality between these shocks, as well as an identifying restriction. The long-run identifying restrictions are incorporated using a Cholesky decomposition of the infinite horizon forecast error variance.[6]

The specific restriction depends on the type of shock we are interested in estimating. Skill-biased technology shocks are shocks to the production technology that affect the skill premium, investment-specific technology shocks change the relative price of investment goods and in the presence of capital-skill complementarity technology shocks may be both investment-specific and skill-biased. Neutral technology shocks increase productivity but do not affect either the relative price or the skill premium. We discuss the specific identifying restrictions used to identify neutral, skill-biased and investment-specific technology shocks as we describe our results in Section 2.3. The identification of different types of shocks using the Cholesky decomposition is then implemented by simply reordering the variables in the VAR.

Our baseline VAR is estimated on quarterly data from 1979:I to 2000:IV. This period is relatively short because of data limitations, see Section 2.2.3. All variables are used in first differences in order to allow for unit roots.[7] The reduced form is estimated as a Bayesian VAR with a Minnesota prior, similar to Canova et al. (2006). The prior

complementarity, as in (2.3), 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 2.4.

[6]The procedure employed here is very similar to the one in Uhlig (2004). We approximate the infinite horizon with 20 years. The procedure uniquely pins down the effects of the identified shocks on all variables in the VAR and the results are not affected by additional (superfluous) long-run zero restrictions.

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

mean pushes towards a unit root (in levels), the prior variance affects the tightness of the lags of the autoregressive variables and of exogenous variables. We use this prior for two reasons. First, in theory one should employ a VAR with an infinite number of lags in order to correctly identify technology shocks using long run restrictions, see e.g. Chari et al. (2008). The Minnesota prior allows us to generate sensible results for a large number of lags simultaneously adjusting the importance (decay) of these additional lags for the estimation. Here, we use 8 lags and a decay parameter of 3.[8] Second, the prior makes our estimation results more stable in the presence of high frequency variation in the skill premium that is due to measurement error. The prior does not affect the long-run restrictions in any way and we show that our results are robust to the strength of the prior and to estimating the reduced form VAR using ordinary least squares (see Table 2.4).

### 2.2.3 Data

We construct quarterly series for the skill premium and the relative employment and supply of skill using individual-level wage and education data from the CPS outgoing rotation groups. This survey has been administered every month since 1979 so that our series runs from 1979:1 to 2006:2.[9] 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). In an average quarter, we have wage and education data for about 35,000 workers.

Our measure for the skill premium is the log wage differential between college graduates and high school graduates. The relative employment and supply of skill are defined as the

[8] The remaining hyper-parameters are chosen as in the RATS manual such that the Minnesota prior is quite loose: $\phi_1 = 0.2$ for the tightness on own lags of a variable, $\phi_2 = 0.5$ for the tightness on lags of other variables and $\phi_3 = 10^5$ for the tightness on exogenous variables.

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

Figure 2.1: Skill premium and Mincer return to schooling in the US

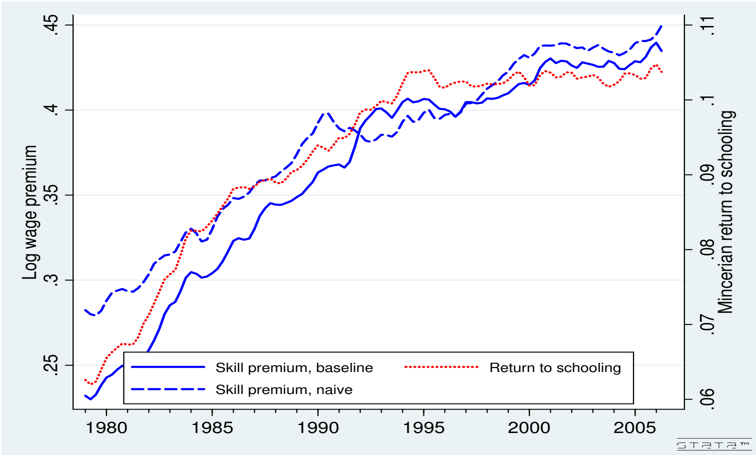


Figure 2.2: Relative employment and relative supply of skill in the US

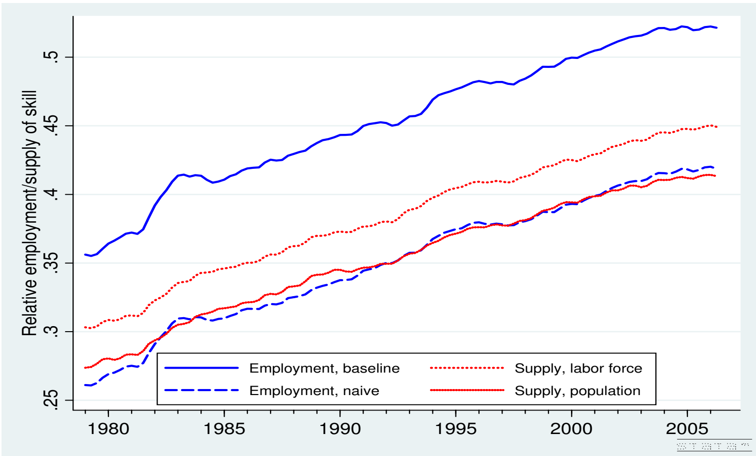


log ratio of the number of college graduates over the number of high school graduates in the population and the workforce respectively. Following Autor et al. (2005), we map the five education levels in the data to college and high school graduate equivalents and control for changes in experience, gender, 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.[10]

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

Figure 2.1 plots our quarterly series for the log wage premium of college over high school graduates.[11] As documented in previous studies, the data show a pronounced increase in the skill premium since 1980, which seems to slow down mildly towards the end of the 1990s. For comparison, the figure also shows a naive measure of the skill premium (the log wage difference between workers with at least a college degree and those with at most a high school degree) and the Mincerian return to schooling. The trend and fluctuations in our measure of the skill premium are similar to those in the Mincer return, indicating we have adequately controlled for heterogeneity beyond two skill types.

Figure 2.2 shows similar plots for the relative employment and the relative supply of skilled labor. Again, there is a substantial difference between our preferred measure and the naive measure of the relative employment of skill. The increase in the employment and the supply of skill was roughly similar over the last two decades, but the higher frequency

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

[11] Note that all our original data series exhibit large high frequency movements. These fluctuations are not seasonal effects but reflect measurement error (sampling error). In a first attempt to get rid of this measurement error the series, as exhibited in figure 2.1 and 2.2, are smoothed using an HP-filter with a very small smoothing parameter, here $\lambda = 1$. The impulse responses are further smoothed by the Minnesota prior.

Table 2.1: Unconditional business cycle correlations

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

Notes: Series are HP-filtered with $\lambda$=1600.

fluctuations differ markedly as we document below.

The other data series we use in our analysis are the following. Output is non-farm business output of all persons from the national income and product accounts (NIPA). Hours are total hours of non-supervisory workers from the Current Employment Statistics establishment survey. Labor productivity is output per hour. All three series are available from the Bureau of Labor Statistics (BLS) productivity and cost program. As the relative price of investment goods, we use a quarterly intrapolation as in Fisher (2006) of the quality adjusted NIPA deflator for producer durable equipment over the consumption deflator (Gordon (1990); Cummins and Violante (2002)).[12]

Table 2.1 shows the business cycle correlations of the skill premium and the relative employment and supply of skill with output, hours, productivity and the relative price of investment goods.[13] The skill premium is basically acyclical: it is only very mildly positively correlated with output and even less correlated with hours worked. This finding is consistent with previous studies (Keane and Prasad (1993); Lindquist (2004)). The relative supply of skill is acyclical as well, but the relative employment of skill is higher in recessions than in booms, indicating the presence of a composition bias in employment as argued by Solon et al. (1994). The correlation of the skill premium with the relative investment-price is weak and negative. This is a first indication that capital-skill comple-

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

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

Figure 2.3: Galí identification with skill premium

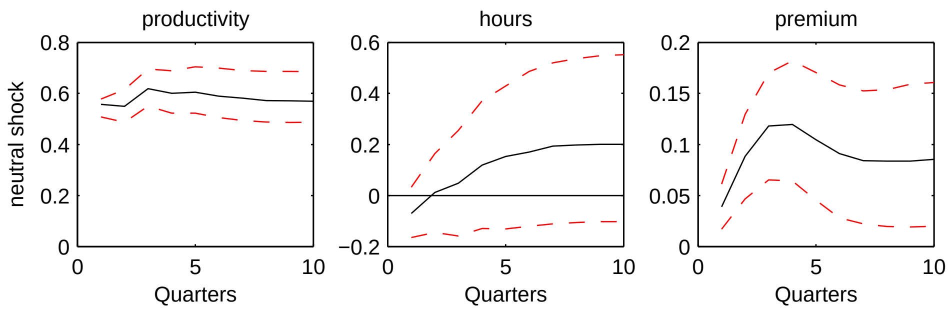


Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

mentarity does not seem an important feature of the data at business cycle frequencies.

## 2.3 Skill-biased technology shocks

In this section, we present our results for the effects of technology shocks on aggregate variables. We start by assessing the degree of skill bias in 'traditional' neutral technology or total factor productivity shocks. We then discuss how exogenous shocks to the supply of skill may bias these estimates and how we can control for these skill supply shocks. Next, in Section 2.3.3, we propose a strategy for separating skill-biased technology shocks from skill-neutral shocks. In Section 2.4, we address the issue of capital-skill complementarity and evaluate the hypothesis that it is investment-specific technological progress that produces the skill-bias observed in the data. Finally, in Section 2.4.2, we jointly estimate all three types of technology shocks and evaluate their importance for business cycle fluctuations in various aggregates.

### 2.3.1 Skill bias in 'neutral' technology shocks

Galí (1999) identifies permanent technology shocks as the only source of long-run movements in labor productivity. In a wide range of models, closed-economy, stationary, one-sector RBC models as well as models of the new Keynesian variety, shocks to total factor productivity are the only shocks that satisfy this identifying restriction. The remaining disturbances in the structural VAR are non-technology or 'demand' shocks, an amalgum of other possible shocks in the model: government expenditure shocks, preference shocks, or shocks to price or wage markups. As a first pass at our data, we evaluate the skill bias in technology shocks identified in this manner.

Figure 2.3 presents impulse response functions of a VAR as in Galí (1999), extended with the skill premium as a measure of skill bias in addition to labor productivity and hours worked, and estimated on our smaller sample. Here, as in all graphs that will follow, the point estimate is the median from and the dotted confidence intervals are 68 % Bayesian bands from the posterior distribution of the structural impulse-response coefficients. Introducing the price of skill as an additional regressor and using a different estimation sample leaves the responses of labor productivity and total hours worked almost unchanged compared to Galí (1999).

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

When we include the wages and hours of high and low skilled workers separately in the VAR, the wage of high skilled workers increases as expected, see Figure 2.8 in the appendix to this chapter. The wage of low skilled workers stays roughly constant and initially even decreases a bit. Apparently, the skill bias in the technology shocks is so large that the relative price effect dominates the average price effect on the wage of low skilled workers. A different picture emerges for hours worked. Here, hours worked by high skilled workers decrease, while they increase for low skilled workers. This result is somewhat counter-intuitive, since we would have expected the relative quantity of skilled labor to increase. Since we have not properly identified skill-biased technology shocks here, this result could, however, obscure different kinds of disturbances such as different types of technology shocks or skill supply shocks.

The estimated technology shocks and their dynamics from the Galí (1999) VAR used here, are similar to the direct estimates of total-factor productivity by Basu et al. (2006). As a robustness check, we use the quarterly series of the Basu et al. (2006) residuals, constructed by Fernald (2007a), instead of labor productivity in the VAR.[14] If the technology shocks identified by the two approaches were identical, then these impulse responses should be the same as those shown in Figures 2.3 and 2.8. The results are shown in Figures 2.9 and 2.10 in the appendix to this chapter. Indeed, the responses of the 'purified' technology measure, hours and the premium are very similar, providing support for the identifying restriction used here. Interestingly, the increase in the wage premium stems from a fall in the wage of low skilled workers rather than an increase in high skilled wages, however.

### 2.3.2 Shocks to the supply of skill

In the identification of technology shocks used above, we assumed that technology shocks are the only shocks that drive productivity in the long run. We showed that these shocks have asymmetric effects on the demand for high and low skilled labor. Thus, production does not use a standard Cobb-Douglas technology, but either requires high and low skilled labor as separate and imperfectly substitutable inputs, as in equation (2.1), or output to be produced in multiple sectors with different input shares of skilled labor. In these cases, the identifying assumption of Galí is no longer valid because shocks to the supply of skill may affect labor productivity in the long run.

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

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

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

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

It is crucial for our identification that we use a measure of the relative *supply* of skill, not the relative employment. It is reasonable to assume that the supply of skill is predetermined, but the same is not true for the employment of skill. If low and high skilled workers are imperfect substitutes, then firms may hire relatively more skilled workers in recessions, when the unemployment pool is larger and these workers are more abundantly available. This composition bias has been documented by Solon et al. (1994). We measure the relative supply of skill as the ratio of skilled workers to low skilled workers in the workforce, whereas the relative employment is the the equivalent ratio among employed workers, see Section 2.2.3.

The strategy to identify technology shocks conditional on skill supply shocks is recursive. We first identify skill supply shocks with the short-run restriction and next use the same long run restriction discussed in the previous subsection to identify technology shocks. Thus, skill supply shocks are allowed to have a long run effect on productivity. Having identified fluctuations in productivity (as well as other variables in our VAR) that are due to skill supply shocks, technology shocks are the only *remaining* shocks that affect labor productivity in the long run. The details on the implementation of this combination of short and long run restrictions can be found in Appendix A to Chapter 1 and Chapter 2.

Figure 2.11 in the appendix to this chapter shows the impulse response functions for this identification scheme. The lower row shows the responses to a one-standard deviation skill supply shock. By construction, the supply of skill increases immediately in response to this shocks. The estimates indicate that the effect is permanent: the supply of skill remains high in subsequent quarters. Somewhat counter-intuitively, labor productivity falls after a positive skill supply shock, hours jump up on impact and continue to increase and the skill premium is almost unaffected.

Controlling for skill supply shocks affects the impulse responses to technology shocks very little. The responses of productivity, hours and the skill premium are all very similar to the estimates without controlling for skill supply shocks. The response of productivity is

Figure 2.4: SBT identification

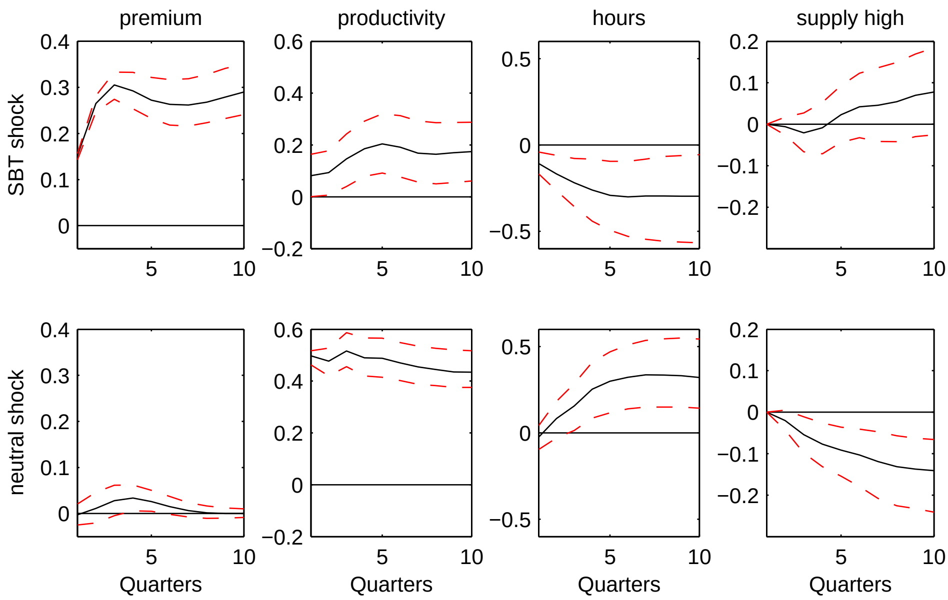


Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

a bit stronger and the response of the skill premium a bit weaker than before. The supply of skill falls moderately, but significantly, in response to a positive technology shock. We conclude that, while the direction of the bias is as expected, its size seems to be small. Nevertheless, we will control for shocks to the supply of skill in all specifications in the rest of the paper.

### 2.3.3 Identified skill-biased technology shocks

While the response of the skill premium is consistent with skill-biased technological change, it casts doubt on the traditional interpretation of these shocks. If these were truly shocks to total factor productivity, as in equation (2.1), the demand for skilled and unskilled labor should increase in equal proportions and the relative demand should be unaffected. Here, we propose an alternative identification strategy to directly identify skill-biased technology shocks in addition to skill-neutral shocks to productivity.

In Sections 2.3.1 and 2.3.2 above, we interpreted the increase in the skill premium in response to a technology shock as a measure of skill bias in technology. Here, we formalize that interpretation as an identifying restriction, identifying skill-biased technology shocks as those shocks that affect the relative price of skill in the long run, see equation (2.4). This restriction is similar in spirit to the identification of investment-specific technology shocks as shocks that affect the relative price of investment goods proposed by Fisher (2006). Controlling for shocks to the supply of skill is particularly important in this context, because of the standard simultaneity problem in estimation of demand or supply elasticities. An exogenous, permanent increase in the supply of skill would permanently reduce the price of skill and thus satisfies our identifying restriction for skill-biased technology shocks. We control for skill supply as described above in Section 2.3.2.

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

This identification scheme strictly speaking is not a *decomposition* of technology shocks as in Galí (1999) into skill-biased and skill-neutral shocks. In principle, there might be shocks that affect the skill premium but not labor productivity in the long run. However, as explained in Section 2.2.1, it is hard to imagine non-technology shocks other than skill supply shocks to affect the skill premium in the long run. Moreover, our estimates indicate that the shocks we identify as skill-biased technology shocks increase labor productivity, supporting our interpretation of these shocks as a specific type of technology shock.

Figure 2.4 shows the responses of the skill premium, the supply of skill, labor productivity and total hours worked to a one-standard deviation skill-biased technology shock (SBT shock) and skill neutral technology shock. By assumption, a positive SBT shock drives the skill premium up in the long run. The estimates indicate that half of this effect is realized immediately and the rest within a year. A skill-neutral technology shock has no significant effect on the wage premium on impact and by assumption there is no long run effect either. SBT shocks increase the supply of skill in the long run, as should be expected with a higher skill premium, but this effect is small.

In response to a positive SBT shock, hours worked significantly and persistently fall. Interestingly, skill-neutral technology shocks barely decrease hours on impact and significantly and substantially increase hours worked less than a year after impact. This finding suggests that at least part of the fall in hours worked in response to technology shocks, as in Galí (1999) and in the estimates in Section 2.3.1, is related to the skill bias in these shocks. If high skilled workers are much more productive than low skilled workers, then it is possible that by substituting low skilled for high skilled workers in response to an SBT shock, firms may increase effective labor input in their production process, while reducing total hours or employment. Figure 2.12 in the appendix to this chapter confirms this interpretation: In response to an SBT shock, the wage of high skilled workers increases substantially, but the wage of low skilled workers actually falls. In contrast, the wages of both types of workers are affected identically by a skill-neutral technology shock. These findings indicate that for low skilled workers the *relative* productivity effect dominates the average productivity effect of an SBT shock.

Table 2.3 shows a decomposition of the forecast error of the VAR at various horizons. At business cycle frequencies with periodicities from 8 to 32 quarters, SBT shocks explain a little over 3% of fluctuations in output, which seems unimportant compared to the 45% of fluctuations explained by skill-neutral technology shocks. Fluctuations in the skill premium are almost exclusively due to SBT shocks, with skill supply shocks and neutral technology shocks combined explaining less than 2% of the variance. Thus, it seems that fluctuations in the skill premium are largely driven by shocks that are unrelated to output fluctuations. However, to understand business cycles in the labor market, it is important to allow for skillbias in technology shocks. Skill-biased shocks are responsible for about 10% of fluctuations in hours worked, slightly more than neutral technology shocks.

### 2.3.4 Robustness

In our baseline estimates, we impose a Minnesota (Litterman) prior on the decay of the lag coefficients in order to be able to allow for a large number of lags. However, our results are not driven by this prior. The responses of productivity and the skill premium to all shocks are virtually unaltered when we change the number of lags, the strength of the prior, or when we estimate the VAR using ordinary least squares (OLS). The fall in hours worked in response to skill-biased technology shocks is also robust across specifications and is significant if we include at least 4 lags in the VAR. The increase in hours in response to neutral technology shocks is actually stronger in all alternative specifications: whereas in the baseline the positive effect becomes significant only after 3 quarters, in all other specifications it is significant at all horizons. These results are summarized in the first panel of Table 2.4.

Next we explore to what extent the way we constructed our measure for the skill premium matters for the results. Using a 'naive' measure of the skill premium that does not take into account the heterogeneity over and above two skill types, we would not have observed the fall in hours in response to an SBT shock. The likely reason is that this measure picks up some changes in the quantity of skill as changes in the premium. Thus, the SBT shock identified off changes in the naive measure would include some skill supply shocks. And since skill supply shocks would be expected to increase hours worked, this would mitigate the fall in hours in response to the identified SBT shocks.

Finally, we compare the properties of our identified SBT shocks to a simple decomposition using equation (2.4), see Section 2.2.2. 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 $\sigma$. 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 (2.4) to retrieve changes in skill-biased technology $B_t$ from our data on the skill premium and the relative employment of skill. After demeaning these changes and normalizing their variance to unity, they are comparable to the identified SBT shocks from the structural VAR. The difference is in the identification. Whereas the identified shocks require assumptions only on the long run behavior of the skill premium, the decomposition requires equation (2.4) to hold in each period. Figure 2.13 plots both estimates for the shocks over the sample period. It is encouraging that despite the differences in identification, the resulting estimates for the skill-biased technology shocks look similar, except at the beginning of the sample. The correlation between the two estimates is 0.48. Moreover, the decomposition is robust to the value of $\sigma$ chosen. In fact, the estimates for the SBT shocks using the decomposition are similar to the first difference of the skill premium.

To complete the comparison, we compare the response of productivity, hours worked and the skill premium to the identified SBT shocks and the estimated shocks using the decomposition. We regress these variables on lags of the shocks, estimated either from the decomposition using equation (2.4) or as the residuals from our structural VAR, as suggested by Basu et al. (2006). This is a direct estimate of the moving average representation of the impulse response functions and the results are comparable to the impulse responses in Figure 2.4. Since the impulse responses in Figure 2.4 seem to flatten out after about 6

Figure 2.5: Impulse-responses to Solow residual

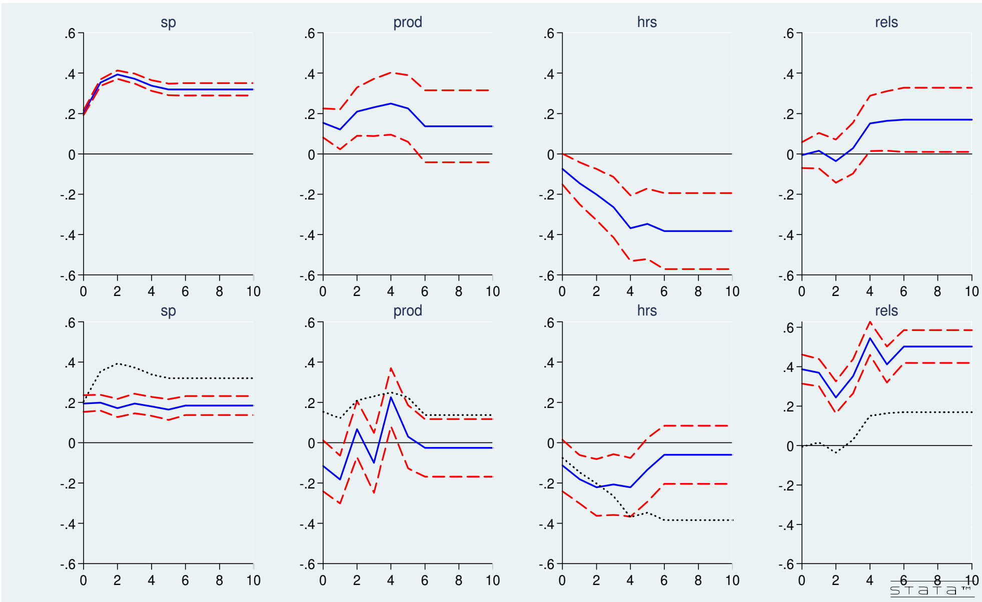


Notes: Percent responses in quarters to a positive one-standard-deviation shock.
First row: Impulse-responses from regression of the skill premium, productivity, hours worked and the relative supply on six lags of the identified SBT shock.
Second row: Impulse-responses from regression of the variables on six lags of the residual from the production function decomposition. The black dotted line repeats the estimate from the first row. Confidence intervals are one standard errors.

quarters, we use 6 lags of the shocks. The results are presented in the first row of Figure 2.5. The responses to identified SBT shocks estimated in this way are quite similar, especially in sign, to those directly calculated from the VAR estimates. We now discuss how the responses to SBT shocks obtained from the decomposition compare to these.

The second row of Figure 2.5 shows the responses to SBT shocks estimated using the decomposition. Generally, the responses are different from the responses to the VAR residuals but throughout equal in sign. The responses of the skill premium and hours worked to the VAR residuals are larger, the respective response of the relative supply of skill is smaller in absolute value than the responses to the decomposition residuals. The largest difference in the responses to SBT shocks estimated in the two different ways is the response of labor productivity. Labor productivity falls below zero on impact and does overall not respond significantly to a shock to the decomposition. Given the difference in the identifying assumption underlying both sets of estimates, the difference must be due to short run deviations of the skill premium from equation (2.4), for example because of wage rigidities. We have discussed above that we find some evidence for other shocks affect both the skill premium and/or labor productivity in the long run, for example neutral technology shocks or skill supply shocks. These shocks are not adequately filtered out by the decomposition, which justifies using long run restrictions.

## 2.4 Investment-specific shocks and the skill premium

Over our sample period the relative price of investment goods fell substantially. This finding has been interpreted to mean that technological progress has been faster in investment goods producing sectors than in consumption goods producing sectors (Greenwood et al. (1997), Cummins and Violante (2002)). Fisher (2006) has argued that such investment-specific technological change is important not only for long run trends, but also for business cycle fluctuations. Because the increase in the skill premium roughly coincided with the decrease in the relative price of investment goods, Krusell et al. (2000) argue that investment-specific and skill-biased technological change might be one and the same. If capital and skill are complements in the aggregate production function, technological innovation in the investment-sector will necessarily lead to an increase in the demand for skill. If this is the case, then investment-specific technology shocks should lead to business cycle fluctuations in the skill premium. In this section, we explore this hypothesis and find no evidence for capital-skill complementarity.

### 2.4.1 Skill bias in investment-specific shocks

Consider the alternative aggregate production function (2.3), as in Krusell et al. (2000), which allows for complementarity or substitutability between capital and skill. Assuming as before that wages are proportional to marginal products in the long run, expression (2.4) for the skill premium changes to the following.

$$\log\left(\frac{w_{H,t}}{w_{L,t}}\right) = \log\left(\frac{\beta(1-\gamma)}{1-\beta}\right) - \frac{1}{\rho}\log\left(\frac{H_t}{L_t}\right) + \frac{\rho-1}{\rho}B_t + \frac{\sigma-\rho}{\sigma(\rho-1)}\log\left(\gamma K_t^{\frac{\rho-1}{\rho}} + (1-\gamma)(B_tH_t)^{\frac{\rho-1}{\rho}}\right) \tag{2.5}$$

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

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

Figure 2.6 shows the responses of the the skill premium, the relative supply of high skilled workers, labor productivity, hours worked and the relative price of investment goods to investment-specific and investment-neutral technology shocks. Note that controlling for skill supply shocks changes the results in Fisher (2006) very little. After an improvement in

Figure 2.6: Fisher identification with skill supply shocks

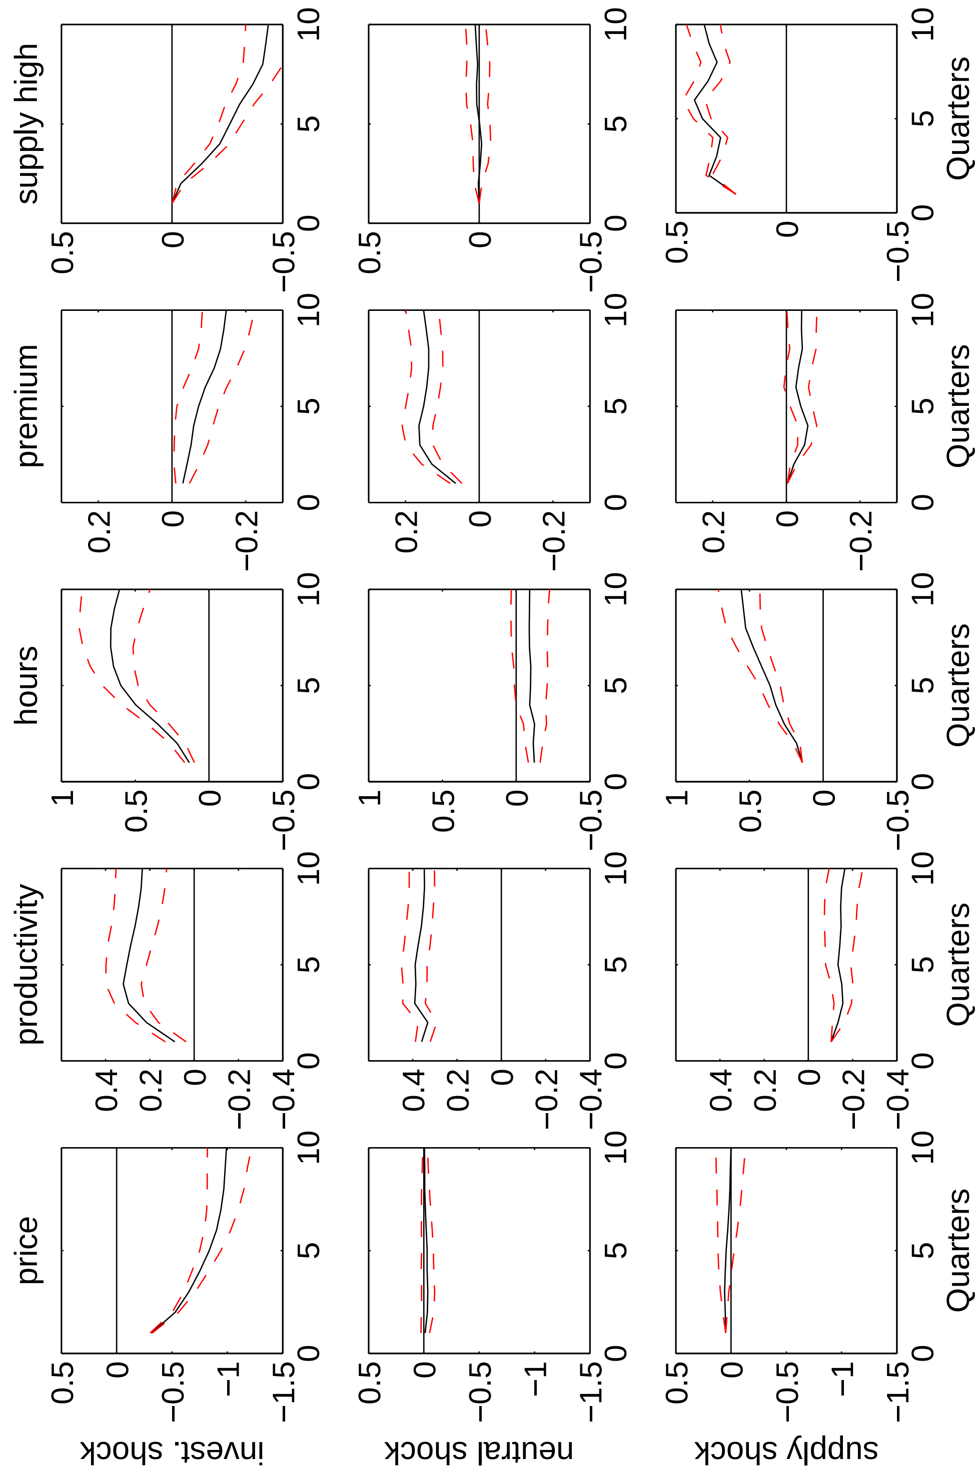


Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

investment-specific technology, the relative price of investment falls, productivity increases and hours worked increase as well. An investment-neutral technology shock, has no effect on the relative price of investment, increases productivity and leads to a fall in hours worked[15].

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

The same finding can be documented in an alternative way. In Figure 2.14 in the appendix to this chapter, we present impulse responses of the relative price of investment goods to skill-biased and skill-neutral technology shocks, identified as in Section 2.3.3. The graphs provide the mirror image to those in Figure 2.6: Skill-biased technology shocks increase the relative price of investment goods significantly, suggesting these shocks are "consumption-specific" or capital and skill are substitutes in production. Note that these shocks share the same features as the investment-unspecific technology shocks identified in Chapter 1.

Our findings are in striking contradiction with the argument in Krusell et al. (2000). What explains the difference is that Krusell et al. (2000) base their argument on a correlation in the long run trends in the skill premium and the relative price of investment goods. In our approach, the identifying variation are comovement between those two series at all frequencies *except* the trends, which are captured by the constant term in the VAR. It is possible that the comovement in the trends in both relative prices is a spurious correlation between two integrated series. It is also possible that the model needed to explain long run growth trends is different from the model that describes higher frequency fluctuations.[16] In any case, our findings reject the hypothesis that there is a stable aggregate production

[15] Since 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).

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

function with capital-skill complementarity.

### 2.4.2 Contribution to business cycle fluctuations

Our results suggest that there are at least four different types of technology shocks with distinct implications for the comovement of aggregate variables: skill-neutral, investment-neutral; skill-neutral, investment-specific; skill-biased, investment-neutral; and unskill-biased, investment-specific (or skill-biased, consumption-specific) technology shocks. With the identifying restrictions discussed above, it is not possible to separately identify all four different shocks simultaneously. Recall that both investment-specific and investment-neutral technology shocks affect the skill premium. Conversely, both skill-biased and skill-neutral technology shocks affect the relative price of investment goods. Hence, if we use a recursive identification scheme, identifying first investment-specific technology shocks, then these shocks will include the unskill-biased, investment-specific shocks. In this case, skill-biased technology shocks will be identified as all *remaining* shocks that affect the skill premium in the long run and will exclude shocks that affect both the relative price of investment and the skill premium. Similarly, if we identify first the skill-biased shocks, then these shocks will include the skill-biased, consumption-specific shocks.

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

Table 2.2 shows the variance decomposition of the forecast error in output, hours and the skill premium. The contribution of skill supply shocks and neutral technology shocks

Table 2.2: Variance decomposition from joint identification

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

Notes: Numbers are in percent. The values for the shocks and the (omitted) residual disturbances add up to 100 for each horizon. The point estimate is the median, the confidence intervals are 68 % Bayesian bands from the posterior distribution.

is very similar in both orderings of the identifying restrictions. This illustrates that we identify the same shocks in both orderings. Neutral technology shocks explain less than 5% of business cycle fluctuations in output and play virtually no role for fluctuations in hours and the skill premium. Investment-specific technology shocks explain up to two thirds of the volatility in output at business cycle frequencies, about one third of the variation in hours. This finding is consistent with earlier findings in this literature (Fisher (2006), Canova et al. (2006)).

Skill-biased technology shocks explain almost all of the entire business cycle variation in the skill premium. These shocks are important for fluctuations in output and (especially) hours as well, but only insofar as they also affect the relative price of investment goods. Investment-specific, skill-neutral technological progress is important for fluctuations in output, but does not have much of an effect on the skill premium. These results suggest that shocks that drive fluctuations in the skill premium are largely unrelated to other variables in the economy. This finding is consistent with the unconditional moments in Table 2.1, which show the skill premium to be largely uncorrelated with output.

### 2.4.3 Capital-skill complementarity

Our finding that the skill premium decreases in response to investment-specific shocks, and the relative price of investment goods increases in response to skill-biased technology shocks suggest that capital and skill are substitutes rather than complements in the aggregate production function. Yet the estimates by themselves do not give any indication as to how large this effect is. What parameters of production function (2.3) are consistent with our estimates? To answer this question, we simulate a simple business cycle model with a production function as in (2.3) and compare the estimated impulse response functions from the actual data to those from simulated data for different values of the substitution parameters. This procedure also allows us to see whether the structural VAR performs well in capturing the conditional moments of the variables in a model that is consistent with our interpretation of the results.

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

For the calibration of the structural parameters of the model we also follow Lindquist (2004), but we assume the two productivity shocks are highly persistent and uncorrelated with each other in order to be consistent with the identifying restrictions of our VAR. The substitution parameters in the aggregate production function (2.3) are $\sigma = 1.67$ and $\rho = 0.67$. These values were estimated by (2.3) 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 $\sigma$ constant, because the value of the elasticity of substitution between high and low skilled workers is well documented, and change $\rho$ 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 ($\rho = 2.67$) and strong substitutability ($\rho = 3.17$) and also an extreme case of substitutability ($\rho = 5$). In each case, we recalibrate the other model parameters if necessary to keep the calibration targets constant.

We simulate the model 1000 times for 88 quarters, the same sample length as in our data. In each simulation, the model is first simulated for 200 periods, which are then discarded, in order to remove dependence on the initial conditions. We then estimate the VAR for each sample of 88 quarters and average the impulse responses across the 1000 simulations. Figure 2.15 in the appendix to this chapter illustrates this for the calibration in which capital and skill are neither complements nor substitutes. For better comparison, the responses are normalized such that they match the responses in the actual data of the investment price and labor productivity to the two technology shocks respectively. Nicely, the estimated responses from the simulated data closely match the theoretical ones from the model.[17] This is also the case for other degrees of substitutability of complementarity between capital and high-skilled labor. Most importantly for our purposes, the estimated response of the skill premium to investment-specific shocks is positive if capital and skill are complements, negative if they are substitutes and zero when they are neither substitutes nor complements.

Figure 2.7 shows the impulse responses of the skill premium to an investment-specific

[17] The small differences may be due to many reasons: technology shocks in the model are persistent but not permanent, the prior smoothes the estimated responses, the finite lag length in the VAR, etc.

Figure 2.7: Capital-skill substitutability

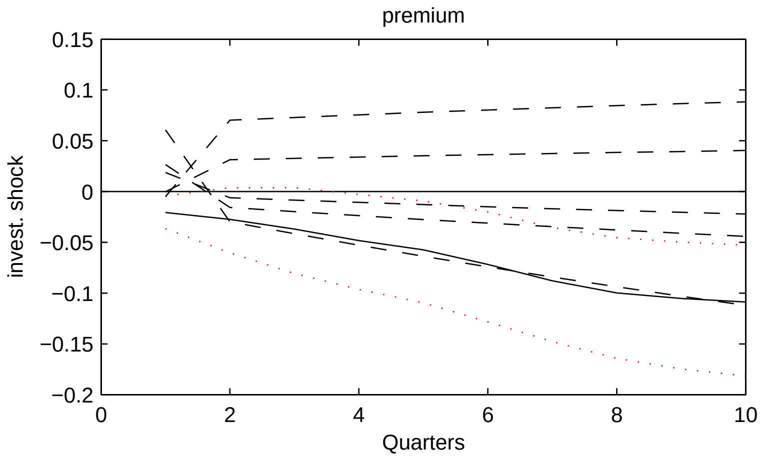


Notes: Black line depicts response of the premium from the estimated structural VAR with actual data together with the Bayesian 68% confidence bands (red dotted lines). The dashed lines show the responses from the model with $\rho = 0.67$, $\rho = 1.17$, $\rho = 2.17$, $\rho = 2.67$ and $\rho = 5$ respectively.

shock from the model simulated for different degrees of capital-skill complementarity/substitutability as well as from the actual data. Comparing the response of the skill premium to investment-specific shocks in the actual data to that in the model, it is clear that our estimates suggest a fairly large degree of capital-skill substitutability. In fact, the estimates suggest an elasticity of substitution between capital and high skilled labor of around $\rho = 5$, whereas the elasticity of substitution between capital and low skilled labor is $\sigma = 1.67$. These parameters imply that if the capital stock increases by 1%, firms can still produce the same amount of output as before if they fire 1.67% of their low skilled workers or as much as up to 5% of their high skilled workers. One should recall here that the estimated response to the skill premium is particularly low for the baseline specification of the estimation. Other estimates will provide lower values of substitutability between skilled labor and capital.

## 2.5 Conclusions

This paper has investigated the implications of skill-biased technological change for the business cycle. In order to address this issue we have constructed a quarterly series of the skill premium and skill supply from the CPS outgoing rotation groups. We have identified conventional neutral and investment-specific technology shocks from structural VARs with long-run restrictions using quarterly U.S. data. In addition, we have proposed a strategy to identify skill-biased technology shocks through reshuffling the variables in the same VAR that we have used for the identification of the conventional technology shocks. Skill-biased technology shocks are those technology shocks that drive the skill premium up and may affect productivity in the long-run. Since they potentially bias the results, we have additionally controlled for shocks to the supply of skill using a short-run restriction.

We have investigated the effect of neutral and investment-specific technology shocks on the skill premium and documented that technology shocks are skill-biased at all business cycle frequencies. Further, there exists no evidence for complementarity between capital and skill over the business cycle as investment-specific technology shocks do not significantly drive up the skill premium. Rather, capital and skill are substitutes in production. Moreover, we find that skill-biased technology shocks lead to a fall in hours worked and may thus be suitable to explain what Galí (1999) has documented as the 'hours puzzle'. As a consequence, skill-biased technology shocks are important to understand the business cycle fluctuations.

We have addressed a great variety of robustness checks for our results. We have considered different measures for the skill premium and skill supply as well as various specifications of estimating the baseline VAR. In addition, we have constructed a measure of skill biased technology from a production function decomposition, similar to a Solow residual. We have furthermore simulated artificial data from a model with different degrees of capital-skill substitutability and complementarity in order to test our identification procedure.

There are nevertheless still some caveats and issues that will need further attention in future research. For example, the induced dynamics to shocks to the supply of skill do not fully agree with what they were expected to ex ante. With respect to the capital-skill substitutability, the biggest open question still lies in the discrepancy of our results with the evidence of complementarity between these two production inputs in the trends (zero frequency). The joint coincidence of these two results points to the existence of two different production functions in the different frequencies. Last, it would be insightful to study the degree of capital-skill substitutability or complementarity in different sectors in order to see how this relates to or causes the aggregate substitutability that we have documented above.

# Appendix to Chapter 2: Additional Tables and Graphs

Table 2.3: Variance decomposition SBT identification

| Horizon | 1 | 8 | 16 | 32 |
|---|---|---|---|---|
| *output* | | | | |
| supply shock | 1.39 | 1.30 | 2.44 | 2.97 |
| | (1.3,1.4) | (0.6,3.2) | (0.7,8.1) | (0.5,10.7) |
| SBT shock | 1.66 | 3.12 | 3.44 | 3.66 |
| | (0.1,7.5) | (0.5,11.4) | (0.6,13.6) | (0.5,14.4) |
| neutral shock | 47.55 | 45.10 | 44.46 | 44.23 |
| | (29.5,64.2) | (27.7,60.4) | (26.5,60.2) | (26.2,60.6) |
| *hours* | | | | |
| supply shock | 7.83 | 8.21 | 14.45 | 17.22 |
| | (7.8,7.9) | (3.8,14.7) | (6.1,25.8) | (7.3,31.7) |
| SBT shock | 7.87 | 10.28 | 9.90 | 9.56 |
| | (1.6,18.1) | (1.9,24.7) | (1.7,25.3) | (1.5,25.4) |
| neutral shock | 1.75 | 9.30 | 9.51 | 9.16 |
| | (0.2,6.6) | (2.2,22.2) | (2.1,23.3) | (1.8,23.3) |
| *premium* | | | | |
| supply shock | 1.07 | 0.79 | 1.04 | 1.18 |
| | (1.0,1.1) | (0.3,2.6) | (0.3,3.9) | (0.2,4.7) |
| SBT shock | 90.98 | 94.74 | 96.43 | 97.40 |
| | (80.7,96.7) | (89.3,97.9) | (92.5,98.5) | (93.8,99.1) |
| neutral shock | 0.88 | 0.84 | 0.41 | 0.19 |
| | (0.1,4.2) | (0.3,2.6) | (0.1,1.3) | (0.1,0.6) |

Notes: Numbers are in percent. The values for the shocks and the (omitted) residual disturbances add up to 100 for each horizon. The point estimate is the median, the confidence intervals are 68 % Bayesian bands from the posterior distribution.

Table 2.4: Robustness of SBT Identification

| | SBT shock on hours | skill-neutral shock on hours |
|---|---|---|
| Baseline specification | | |
| | -, sign. | +, sign. after 3rd quarter |
| Variation of the baseline specification with baseline wage premium | | |
| Minnesota prior with 8 lags changed to | | |
| 2 lags | -, not sign. | +, sign. |
| 4 lags | -, sign. | +, sign. |
| 12 lags | -, sign. | -, sign. on impact, |
| | | +, sign. after 3rd quarter |
| weaker prior* | -, sign. | +, sign. |
| Flat prior (OLS equivalent) | | |
| 2 lags | -, not sign. | +, not sign. |
| 4 lags | -, sign. | +, sign. |
| Baseline specification with different wage premium series | | |
| Naive measure | +, small effect, | +, sign. |
| | sign. on impact* | |
| Lindquist measure | +, small effect, | +, sign. |
| | sign. on impact* | +, sign. after 3rd quarter |

Notes: *Changed decay parameter from $d = 3$ to $d = 1$.

Table 2.5: Variance decomposition Fisher identification

| Horizon | 1 | 8 | 16 | 32 |
|---|---|---|---|---|
| *output* | | | | |
| supply shock | 0.47 | 5.48 | 9.86 | 11.86 |
| | (0.4,0.5) | (2.1,9.9) | (3.5,17.9) | (3.9,21.7) |
| i-shock | 15.21 | 60.21 | 56.95 | 54.21 |
| | (7.7,25.2) | (47.8,70.5) | (42.5,69.4) | (37.7,69.1) |
| neutral shock | 17.17 | 6.65 | 6.29 | 6.20 |
| | (8.0,29.6) | (2.5,13.4) | (2.1,13.2) | (1.7,14.3) |
| *hours* | | | | |
| supply shock | 17.25 | 20.43 | 28.87 | 33.22 |
| | (17.2,17.3) | (14.7,27.6) | (19.3,39.9) | (22.2,46.1) |
| i-shock | 15.44 | 41.95 | 35.79 | 29.74 |
| | (8.1,23.9) | (29.9,54.2) | (21.2,51.3) | (15.5,48.8) |
| neutral shock | 13.65 | 2.08 | 1.69 | 1.53 |
| | (5.9,23.3) | (0.6,6.8) | (0.5,6.1) | (0.3,6.1) |
| *premium* | | | | |
| supply shock | 0.05 | 2.70 | 3.17 | 3.75 |
| | (0.0,0.1) | (0.9,6.3) | (0.8,8.7) | (0.7,10.5) |
| i-shock | 3.98 | 10.39 | 20.02 | 23.12 |
| | (0.7,10.5) | (2.8,22.9) | (6.3,37.7) | (6.5,43.6) |
| neutral shock | 21.13 | 31.85 | 28.59 | 27.37 |
| | (11.7,31.4) | (20.3,44.1) | (17.5,41.4) | (16.0,40.6) |

Notes: Numbers are in percent. The values for the shocks and the (omitted) residual disturbances add up to 100 for each horizon. The point estimate is the median, the confidence intervals are 68 % Bayesian bands from the posterior distribution.

Figure 2.8: Galí identification - additional variables

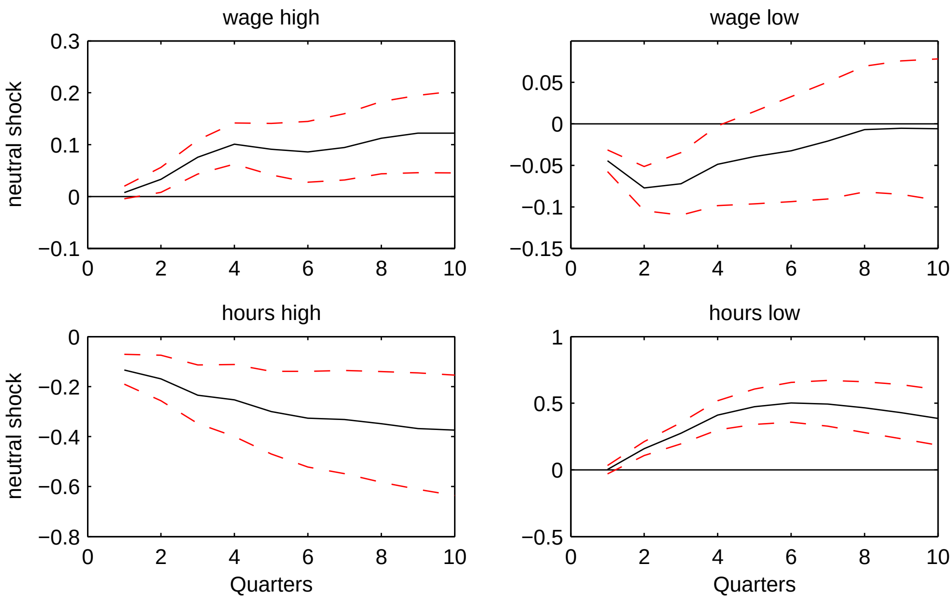


Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 2.9: Galí identification with TFP measure

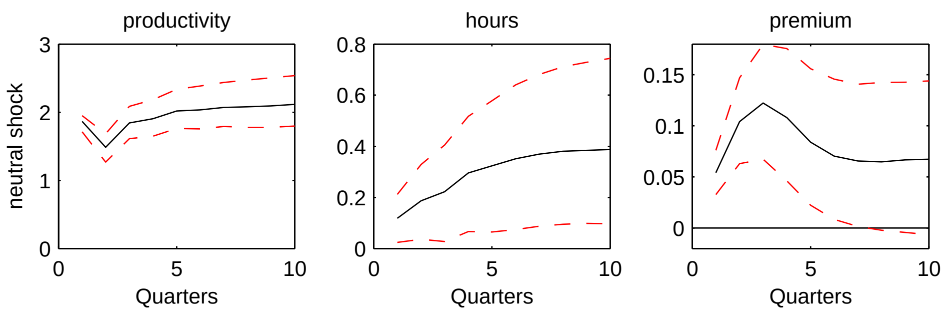


Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 2.10: Galí with TFP measure and additional variables

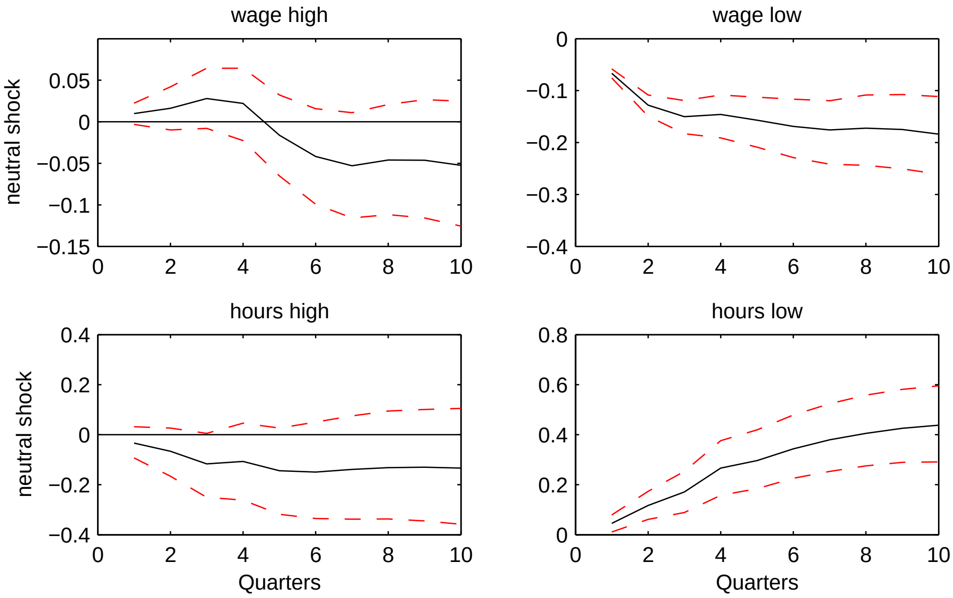


Notes: Percent responses to a positive one-standard-deviation shock. Confidence intervals are 68% Bayesian bands.

Figure 2.11: Galí identification with skill supply shocks

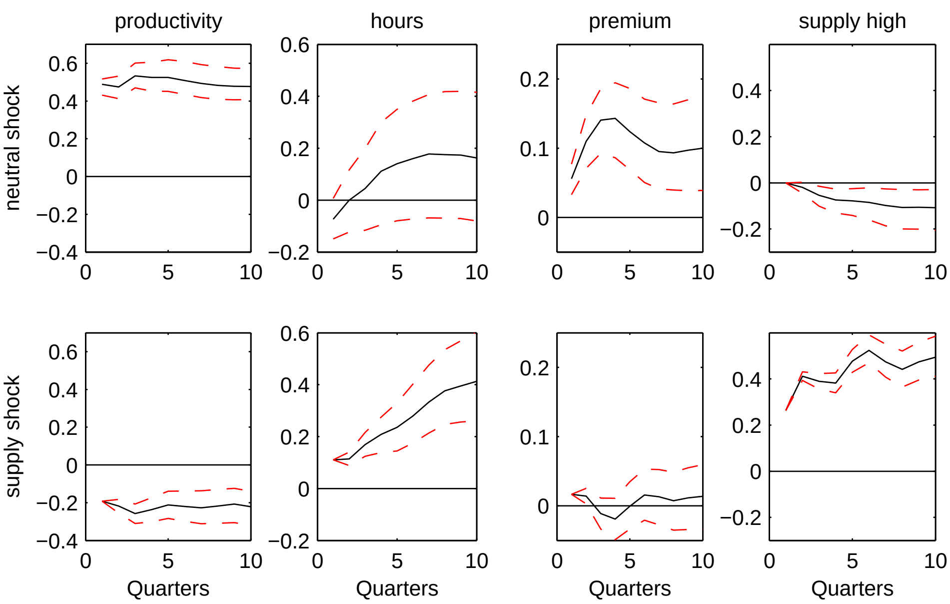


Notes: Percent responses to a positive one-standard-deviation shock. Confidence intervals are 68% Bayesian bands.

Figure 2.12: SBT identification - additional variables

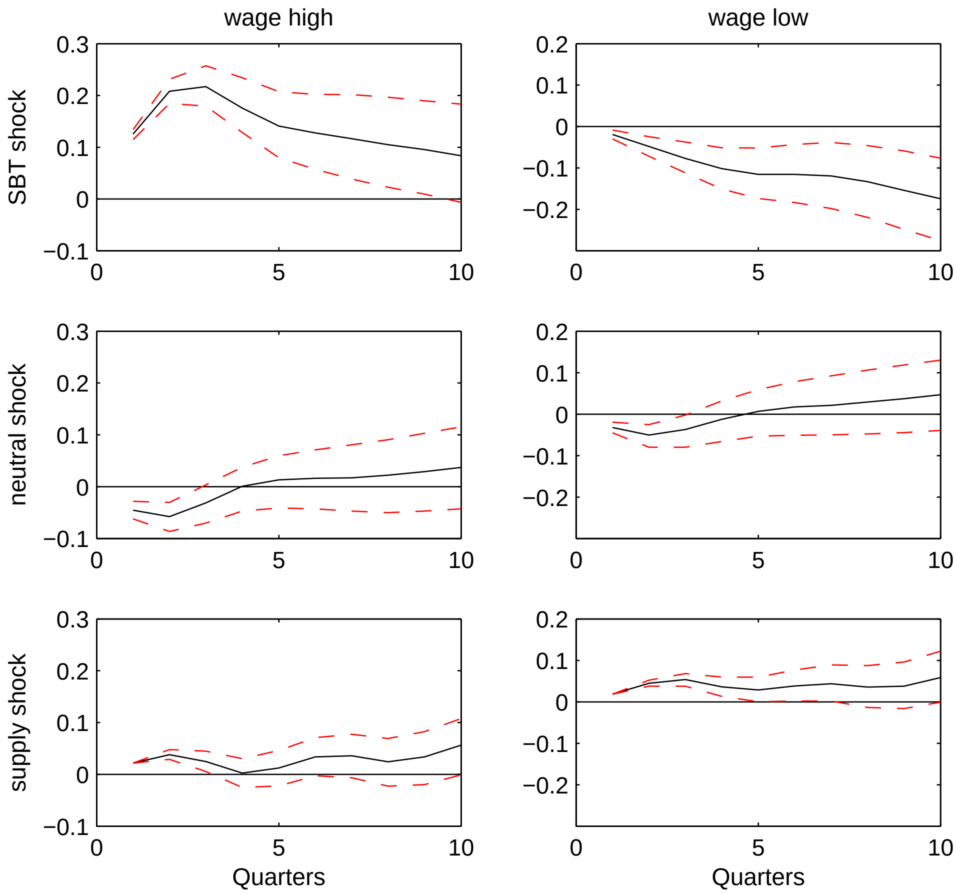


Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 2.13: Comparison of SBT shock and decomposition

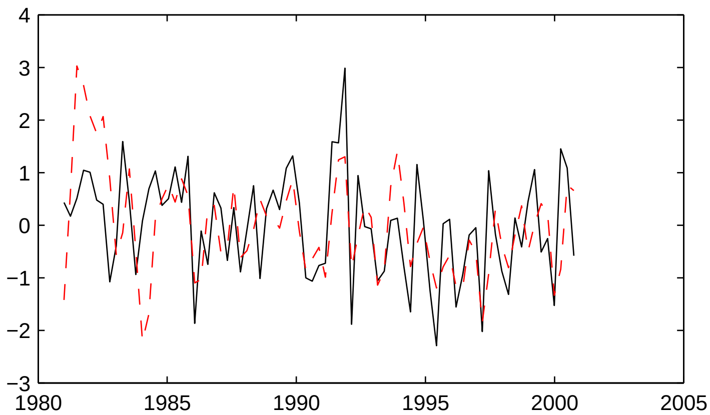


Notes: Black solid line depicts identified SBT shock, red dashed line shows the residual from the production function decomposition.

Figure 2.14: SBT identification - relative price of investment goods

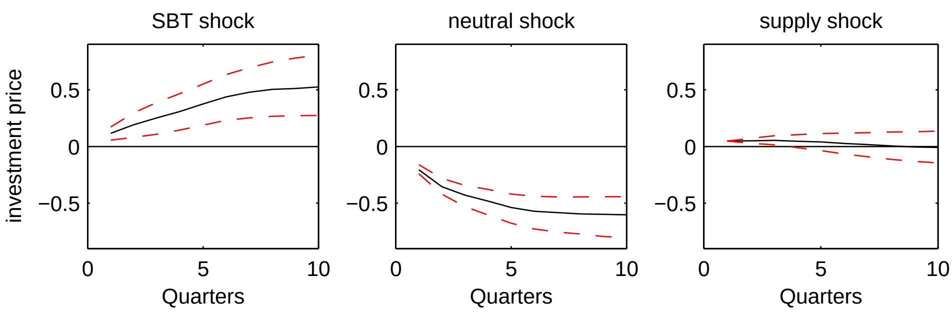


Notes: Percent responses to a positive one-standard-deviation shock.
Confidence intervals are 68% Bayesian bands.

Figure 2.15: Impulse-responses from model and simulated data

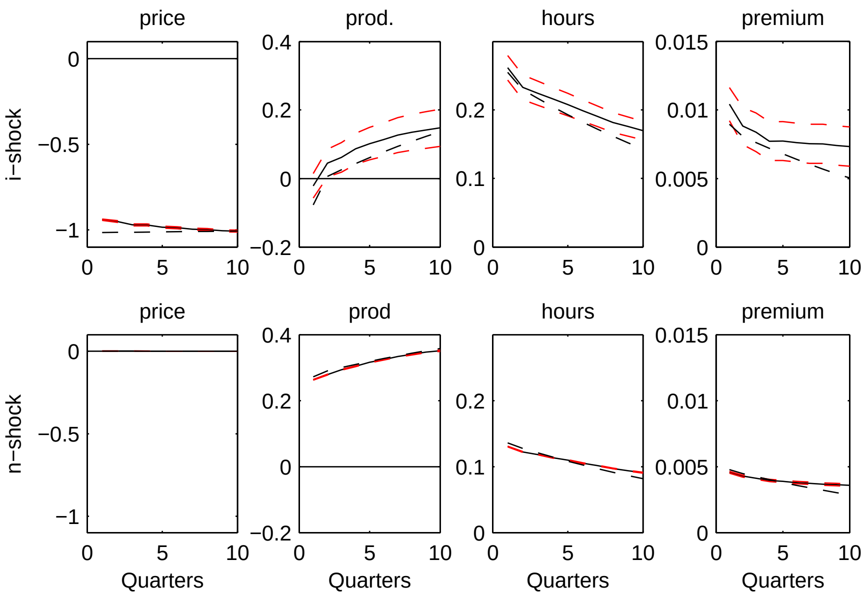


Notes: Percent responses to a positive one-standard-deviation shock. Responses from the model with $\rho = \sigma = 1.67$ are dashed lines. Solid lines with 68% Bayesian confidence bands are estimated from 1000 simulations from the same model. The responses are normalized to match the responses of the investment price and labor productivity in the actual data in the longer run (20 quarters).

# Chapter 3

# On the Implications of Unobserved Age and Cohort Effects for Aggregate Labor Supply

## 3.1 Introduction

The euro area labor force participation rate, defined as the ratio between the labor force and the working age population, has increased from below 65% in the early 1980s to 70.9% in 2007. The increase in the propensity to work or to search for and to be available for jobs has been the main driver of the substantial increase in euro area labor supply that has accelerated since the mid-1990s. However, the overall increase reflects substantial heterogeneity in the evolution of participation behavior across population groups and across euro area countries. The participation rate of females in the euro area has increased by more than 15 percentage points over this time period, to 63.3% in 2007, compared to the participation rate of 78.6% for males (see upper panel of Figure 3.1). The participation rate of the young (15-24 years old) declined markedly until the mid-90s and has stabilized to around 45% in the last decade, whereas, following a long period of stable participation rates, the participation rate of those 55-64 years old increased markedly in the last few years (see lower panel of Figure 3.1). Also, the extent of the increase in participation and its composition across worker groups varies across euro area countries, suggesting an important role for cross-country heterogeneity in the underlying factors that determine individual labor supply decisions. A number of factors could have contributed to the overall increase in participation: robust, employment intensive economic growth (in particular from the mid-1990s onwards), reforms targeted at groups with lower attachment to the labor market, changes in cultural attitudes towards work (particularly for women), as well as demographic factors, such as the larger share of the population in prime working age.

We use harmonized data from the EU Labour Force Survey (LFS) and a cohort-based model to analyze determinants of labor market participation in the euro area and the five largest euro area countries (Germany, France, Italy, Spain and the Netherlands) over the last few decades. We refer to two euro area aggregates. The first (EA12) consists of the euro area 12 countries before Slovenia, and later Malta, Cyprus and Slovakia, entered the euro area and the second (EA5) an aggregation of the five largest euro area countries. The cohort-based model is used to decompose time-series of age-specific participation rates in euro area countries into the impact of the business cycle, observed structural determinants of participation (such as labor market institutions) and most importantly unobserved determinants captured by age and birth-year specific (i.e. cohort) effects.

The age and cohort effects are derived from the evolution of the age-participation profile over time. The propensity to participate evolves over the life-cycle, as reflected in an inverted u-shape age-participation profile. The age effects in the model capture this feature of the underlying age-participation profile. At the same time, the age-participation profile is continuously shifting. The cohort based model captures parallel shifts in the profile that are specific to a birth-year through the unobserved cohort effects. While cohort effects generally encompass any factor associated with a particular birth year, they are likely to reflect the impact of individual participation choices made early on in life (for example choices relating to starting a family, maternity leave and/or education) that persist throughout the life-cycle. They may also reflect crowding-out effects or slowly evolving preferences, cultural factors or institutions.[1] Controlling for business-cycle effects, we first estimate age and cohort effects for the euro area aggregate (EA12) and individual euro area countries.

While cohort effects explain shifts in the age-participation profile, potential changes in the shape of this profile are captured through observed time-varying determinants, such as demographic trends and changing labor market institutions. We use the model with

[1] For example, Fernandez (2007) builds a model of female participation that is based on culture and learning. She argues that cultural factors can explain the increase (and the S-shaped time series pattern) of female participation rates in the United States. Antecol (2000) finds that the home country plays an important role in participation decisions of first generation female immigrants in the US, suggesting that culture matters for participation behavior.

Figure 3.1: Participation rates by worker groups in the euro area (EA12)

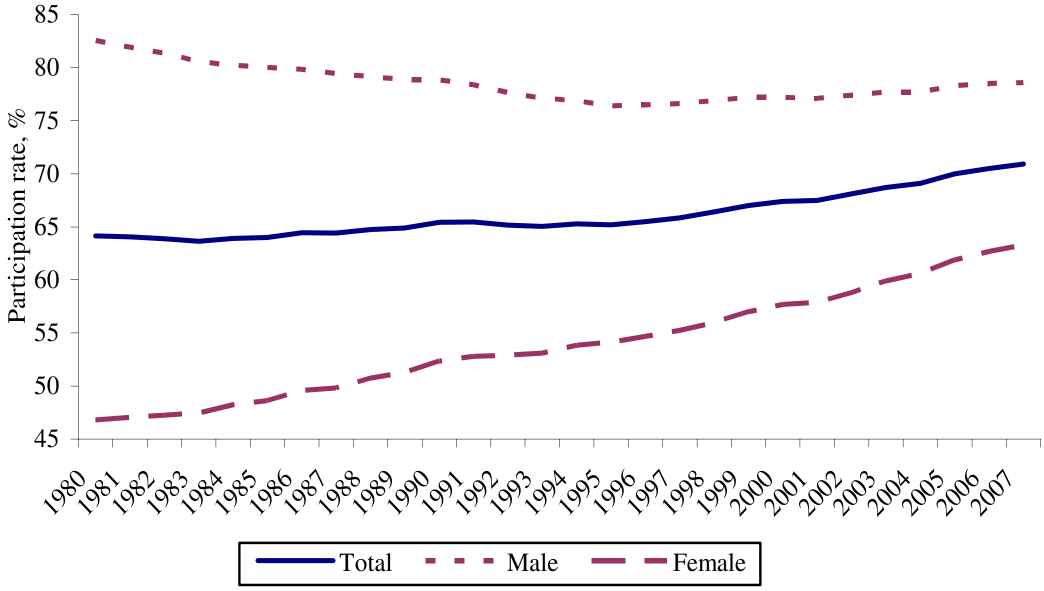


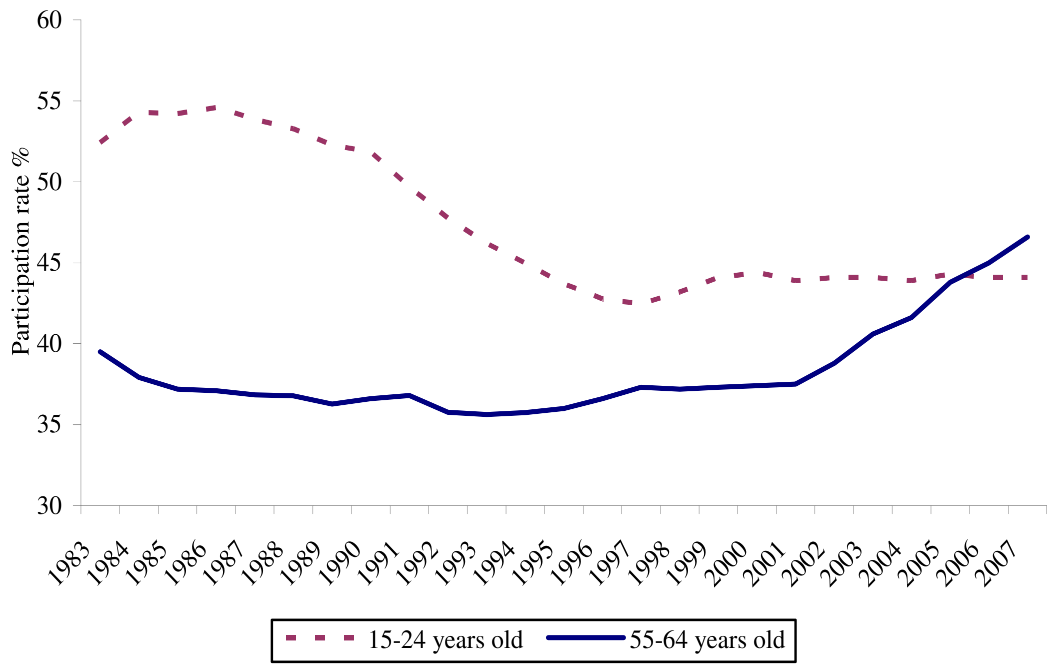


Sources: EU LFS (Eurostat), OECD and own calculations.

observed determinants to explain changes in trend participation rates over time and also to project them forward in the five largest euro area countries. We then aggregate these country trends and projections for the euro area (EA5). Projections that take age and cohort effects and the changing population structure into account provide a useful benchmark scenario for future labor supply in the euro area. In particular, a cohort based model takes into account the extent of the pass-through of participation behavior from the young cohorts to the oldest cohorts. Looking forward, demographic factors will become less favorable with population ageing increasing the importance of positive participation trends within age and gender groups in sustaining potential growth in the euro area. As we estimate the model separately for individual euro area countries and aggregate the results, the results for the euro area also fully incorporate heterogeneity across countries.

Our paper is related to two main strands of literature. First, our main focus is on accurately estimating trends in participation based on both observed determinants and the unobserved age and cohort effects in the euro area. For this purpose we use a modified version of the cohort-based model presented in Fallick and Pingle (2007) and applied in Aaronson et al. (2006) to data for the United States. By simultaneously estimating participation equations for single ages for each gender and taking advantage of cross-equation restrictions the model provides a detailed account of the role of age and cohort effects in explaining movements in the aggregate participation rate. Fallick and Pingle find that these effects provide additional insights compared to time series based trend/cycle decompositions. For example, they find that the levelling off of the increase in the propensity to participate at cohorts born around 1950 suggests that increased labor market attachment is less likely to support an increase in the participation of females in the United States. We are not aware of a cross-country study of European participation rates that accounts for these features. Euwals et al. (2007) find using micro-data that cohort effects have played an important role in explaining the increase in female participation from 1992 to 2004 in the Netherlands. Fitzenberger et al. (2004) use an alternative age, cohort, and period accounting model to study participation and employment in Germany and find significant cohort effects for females.[2]

Second, a number of studies have documented the impact of labor market institutions on unemployment and employment in European countries (for a recent contribution and

[2]Other studies that use closely related methods include Beaudry and Lemieux (1999) for Canada and Fukuda (2006) for Japan. In addition, Carone (2005) and Burniaux et al. (2004) take advantage of cohort effects to project participation rates for EU and OECD countries respectively.

review of the literature, see Bassanini and Duval (2006) and Bertola et al. (2007) for age-group specific analysis). Participation decisions have received less attention in this context. Blöndal and Scarpetta (1999) and Duval (2006) focus on older workers and their retirement decisions and Jaumotte (2003) on females. Genre et al. (2005) and (2008) focus on group specific participation rates in European countries. Using annual data for a panel of European Union countries, they estimate participation equations for age and gender groups in order to identify the impact of institutions in participation decisions. They find that labor market institutions indeed matter for labor supply: higher union density, more employment protection and more generous unemployment benefits lower participation rates. Genre et al. (2008) also find using lagged participation rate as a proxy, that a common (across countries) cohort effect is an important element for understanding participation rates of older women (those between 55 and 64) in European countries. We add to these studies by considering disaggregated groups and by evaluating age and cohort effects and possible observed determinants of participation in the same model. Instead of the cross-country focus of most previous studies, we exploit the time series dimension of the data and incorporate the impact of a broader set of factors through the unobserved age and cohort effects.

We find that analyzing participation behavior both between (age and gender effects) and within (cohort effects) detailed age and gender groups is particularly useful for modelling trends in euro area aggregate participation rates and projecting them forward. Our results suggest that age and cohort effects can explain a substantial part of the recent increase in labor force participation rates in the euro area, although not the surge since early 2000s. Cohort effects are particularly relevant for women, with those born in the 1920s and 1930s less likely and those born in the late 1960s and early 1970s more likely to participate in the labor market over the life-cycle. There is substantial variation in cohort effects across the five largest euro area countries that we analyze. Depending on the country, the estimated cohort profiles suggest an increase of 10 to 30 percentage points in female participation rates. We also find that a number of observed determinants, such as labor taxes, union density, unemployment benefits and the average number of children have had an impact on labor force participation rates, although the specific impact varies across age and gender groups and countries. Looking forward, while they continue to provide some upward support to participation rates of women in the euro area, positive cohort effects are not large enough to compensate for the downward impact of population ageing on labor force participation rates in the euro area.

The rest of the paper is organized as follows. In Section 3.2 we describe sources and characteristics of the data employed and the cohort based model of participation. In Section 3.3 we present results from the model in three parts. We first illustrate the role of estimated age and cohort effects in determining participation. Second, we analyze the impact time-varying observed determinants of participation within a full model. Third, we present projections for participation rates up to 2030 based on the model and compare them with alternative scenarios. Finally, we summarize our results and conclude in Section 3.4.

## 3.2 Data and methodology

Participation behavior and its determinants vary systematically by age and gender and changes in group-specific participation rates translate into the aggregate through an evolving population structure.[3] As a result, analyzing participation behavior of detailed age and gender groups is essential for understanding aggregate participation developments.

The source for data on population, employment and unemployment for detailed age and gender groups for euro area countries is the EU Labour Force Survey (LFS) compiled by Eurostat.[4] The same LFS data are used by Eurostat to calculate official statistics on participation and unemployment for EU countries. The LFS data are harmonized across countries and therefore particularly well-suited for cross-country comparative analysis. The annual data from 1983 to 2007 are based on the spring (second quarter) results. Data are available for ages from 15 to those over 70.[5]

Constructing consistent data over time requires some adjustments. In the case of Germany, data prior to 1991 have been extrapolated backwards on the basis of the developments in

[3]Naturally, participation behavior varies also across other personal characteristics, such as education and skills, immigrant status etc. We focus on age and gender for reasons of data availability: in particular, LFS data by education categories is only available from the early 1990s onwards. That data shows that more educated workers tend to have higher participation rates and that an increase in overall educational attainment over time has coincided with an increase in participation rates, particularly for women.

[4]A detailed description of the sampling methods and adjustment procedures used in the LFS can be found in "The European Union Labour Force Survey - Methods and Definitions, 2001", the available variables are listed and described in the "EU Labour Force Survey database - User guide". The change from annual to quarterly periods by Eurostat has resulted in breaks in the LFS survey in many euro area countries. Therefore we rely on the more consistent spring (second quarter) data throughout the sample period, except for France and Austria (first quarter).

[5]Except for Spain where data are available for those above 16.

West Germany. We refer to two euro area aggregates in the paper. The first consists of the euro area 12 countries before Slovenia, and later Malta, Cyprus and Slovakia, entered the euro area. While there is no information available in the LFS for the euro area countries before they joined the European Union (i.e., for Spain and Portugal prior to 1986, for Austria and Finland prior to 1995), this has been taken into account in the calculation of the euro area 12 aggregate. In particular, data for the euro area 12 aggregate prior to 1996 have been obtained on the basis of the growth rate of the largest aggregate available (i.e., 12 countries in 1995 to 2006, 10 countries between 1986 and 1995 and 8 countries before 1986). The second aggregate (euro area 5) contains only the largest 5 countries of the euro area, namely Germany, France, Italy, Spain and the Netherlands. This aggregate is calculated from the actual and estimated participation rates of the single countries, weighted with their respective population shares.

Labor supply and participation rates evolve substantially over the life-cycle, tracing a well-known overall inverted u-shape profile of participation rates that peaks around the prime working age. Figure 3.8 in the Appendix to this chapter illustrates these profiles for the EA12 aggregate for males and females in 2007. The participation rates of younger workers (those below 25) and older workers (those above 50) change substantially from one age group to another, whereas the substantially higher participation rates for those in prime working age show a relatively flat profile between the ages of 30 and 50. The age participation profile for females is always below, peaking earlier than the profile for males. The gap to the male profile is smaller at younger than at older ages. At the country level, we can observe somewhat different participation profiles in 2007, pointing to important heterogeneity in participation behavior. For instance, the gap between male and female participation rates is more substantial in Italy and Spain than in France, Germany and the Netherlands, especially for those in prime age. While the participation rates of the youngest age groups are comparable between most countries (at levels around 10-30%), they are substantially higher in the Netherlands (at around 60%). Finally, for the oldest age groups (60-64 years old), differences are mostly concentrated in female participation, which varies from 10% in Italy to around 30% in Germany, while for males, participation rates are generally between 30-50%, with the only exception of France (below 20%). The age effects in the model for the euro area and the different countries will capture these features of the underlying age-participation profiles for men and women.

These age-participation profiles for males and females are continuously evolving as a result

Figure 3.2: Changes in participation rates by age for females in the euro area (EA12)

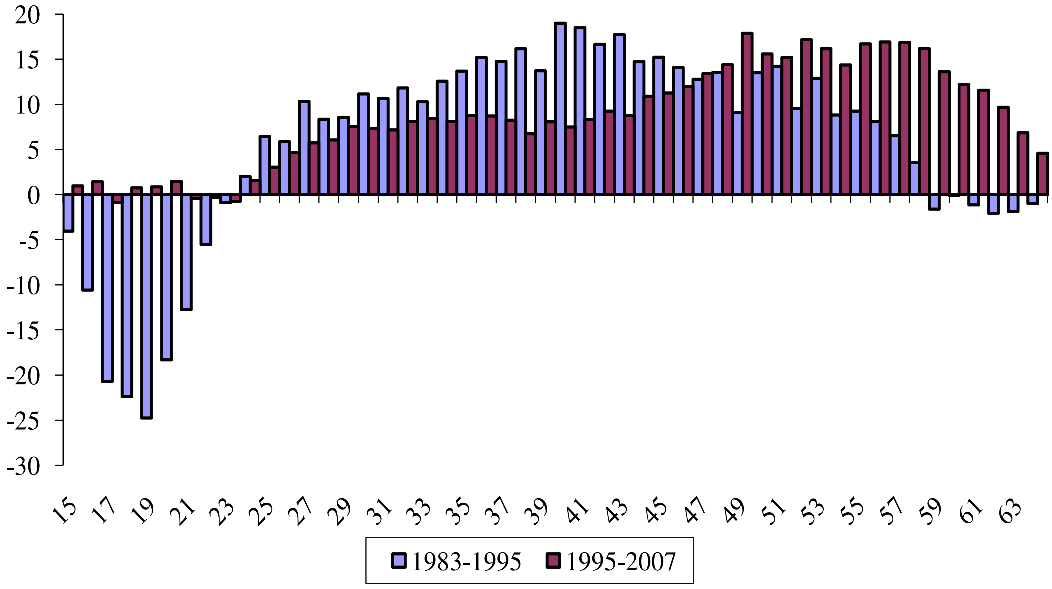


Sources: EU LFS (Eurostat)and own calculations.

of changes both between and within age groups. Figure 3.2 plots the overall change in participation rates for each single age for females in two time periods, 1983-1995 and 1995-2007. These two periods are comparable both in terms of length and in terms of economic developments (i.e. the business cycle). Overall, since 1983 the female profile has been lifted up for those between 25 and 58 years old, in particular for older women who increasingly stay in the labor market after child-bearing. At the same time, the participation rates for the youngest women have declined. Since the mid-1990s the latter effect has decreased, while the hump-shaped pattern of an increase in participation for those between 25 and 58 years old has shifted towards older age groups. This effect is reminiscent of cohort specific participation effects, i.e., female participation behavior for a particular cohort persists over time. In terms of the age-participation profile, the estimated cohort effects in the model describe upward shifts in the profile that are specific to a particular birth-year.

In contrast, participation rates for prime-age males have not changed significantly in the entire period. However, the participation rates of those between 15-24 years old have slightly declined, while those of the oldest, between 55 to 64 years old, have increased a bit. For both males and females, the increases in participation rates for the youngest and oldest workers may be related to the impact of labor market reforms that have focused on groups with a weaker attachment to the labor market. While the cohort effects capture parallel shifts in the age-participation profile, changes in the shape of the age-participation profile over time will be captured by time-varying institutions and other explanatory variables such as the share of youth in education.

We use the output gap to measure the business cycle. The output gap is calculated as a deviation of real GDP from an Hodrick-Prescott filtered trend. In line with Uhlig and Ravn (2002), the smoothing parameter in the HP-filter for annual data is set at $\lambda = 6.25$. The real GDP data for both the euro area 12 and the single countries is taken from the AMECO database. The full model specifications include a number of indicators for key time-varying institutions. We include OECD indicators for union density, labor taxes, implicit tax for older workers, the unemployment benefit replacement rate, the share of youth in education and average number of children also used in Bassanini and Duval (2006) and a measure of life expectancy from Eurostat. When missing, data for the last few years has been extrapolated based on past trends. As noted before, we include institutions to control for changes in the shape of the age-participation profile. This means that we rely on time-series variation of institutions within a single country to identify the impact of institutions. Therefore, several important institutional determinants of labor supply that do not generally vary over time, such as the mandatory retirement age, are excluded from this analysis.

The estimation strategy is based on the cohort-based model presented in Fallick and Pingle (2007).[6] Specifically, we estimate a system of constrained least squares regressions for single ages 15 to 70 and over, separately for men and women:

$$\ln\left(\frac{LFPR_{g,t}}{1-LFPR_{g,t}}\right) = \alpha_g + \sum_{b=1917}^{1992} C_{g,b,t}\beta_b + \lambda_g X_{g,t} + \varepsilon_{g,t}$$

The dependent variable is the logistic transformation of the participation rate for males or females. We use the logistic transformation to ensure that predicted participation rates remain bounded between 0 and 100 and undo the transformation after estimation.

[6] Closely related models based on age, cohort and period accounting have a long tradition in sociological and demographic research and have been recently applied to analyze labor supply in Beaudry and Lemieux (1999) and Fitzenberger et al. (2004). Articles in Mason and Fienberg (1985) provide an early discussion of basic accounting models and applications that rely on functional form assumptions. From an economic perspective the pure age, cohort and period accounting approach seems rather ad hoc in nature. The current model is therefore an attempt to move beyond a pure statistical decomposition by including observable variables that capture underlying factors that determine participation rates. See also Euwals et al. (2007) for a discussion and comparison of different modelling strategies.

The coefficient $\alpha$ represents an age effect that is constant over time and measures the average propensity to participate in the labor market at a certain age. The $\alpha$ for all ages trace an underlying fixed age-participation profile. The coefficients $C_{g,b,t}$ represent dummies for the different birth years and are equal to one if the birth cohort $b$ appears in age $g$ at time $t$. Within each gender group and country, the coefficient $\beta$ is constrained to be the same across equations. This allows an identification of cohort effects separately from the age and business cycle effects. As a consequence, the coefficients $\beta$ represent cohort effects that are constant over time and may be interpreted as the average propensity to participate in the labor force when born in a particular year. The cohort effect shifts the underlying age participation profile up and down, depending on the propensity of the birth year cohort to participate in the labor market throughout their working lives. We include all cohorts in the estimation which results in considering persons born between the years 1917 and 1992. However, as the most recent birth cohorts are only observed when they are very young, we estimate the model without the last eight cohorts.[7] Later, we assign a cohort effect to these cohorts after estimation by setting it equal to the last estimated cohort effect (equal to the cohort effect for those born in 1984).

Finally, $X$ contains other variables that have explanatory power for participation rates of particular age groups. In the baseline specification, this encompasses business cycle effects represented by the contemporaneous value and two preceding lags of the output gap. In addition, both the estimated age and cohort effects are potentially influenced by time-varying institutions. In the full model therefore $X$ includes also a set of indicators of observed determinants. Note that the institutions do not vary across ages, although some institutions are included only in the equations for young (youth in education), female (number of children) or older workers (implicit tax and life expectancy). The coefficients of the observed determinants vary freely across ages and therefore allow the underlying age-participation profile to tilt.

The total system is estimated based on 1400 age-year observations, with 56 equations, resulting in 56 estimated age and 168 estimated business cycle parameters each and 72 (constrained) cohort parameters. The unconstrained model results in a regressor matrix that is of reduced rank. With the help of the restrictions on the cohort effects, the estimation is nevertheless possible as shown in Greene and Seaks (1991). Significance tests

[7] We do this by replacing the values of the participation rate and the other explanatory variables of the ages affected with means from the rest of the sample. We also restrict the cohort effects of the last eight cohorts to equal the average of the remaining cohorts for the respective age.

Table 3.1: Contribution of population composition to changes in participation rates

| | 1983-1995 | 1995-2007 | 2007-2015 | 2007-2030 |
|---|---|---|---|---|
| 15-19 | -1.4 | -0.2 | -0.1 | 0.0 |
| 20-24 | -0.4 | -0.8 | -0.3 | -0.1 |
| 25-34 | 1.9 | -2.4 | -1.4 | -1.9 |
| 35-44 | 1.1 | 1.9 | -1.9 | -3.5 |
| 45-54 | -0.3 | 1.7 | 2.2 | 0.4 |
| 55-64 | 0.4 | 0.2 | 1.0 | 2.8 |
| Total | 1.3 | 0.4 | -0.6 | -2.4 |
| Change in PR | 1.5 | 5.7 | – | – |

Notes: Numbers in percentage points. Sources: EU LFS (Eurostat) and own calculations.

are based on robust (White-corrected) standard errors. The cross-equation constraints identify the cohort effects only up to a scale factor. As in Fallick and Pingle (2007), we therefore normalize the coefficient estimates by setting the parameter of one cohort (here 1969) to one.

Based on the results for labor market participation of the single ages for males and females, we then construct aggregate participation trends using population weights. In addition to changes in group-specific participation behavior, Table 3.1motivates how changes in the composition of the euro area population have affected participation rates over time.[8] In particular, the positive total effect of the population composition of 1.3 percentage points observed in 1983-1995 declined significantly in 1995-2007 mainly resulting from the decreasing share of the prime-age population in favor of older groups with lower participation rates. Moreover, the supportive contribution of the 25-34 year-old to aggregate participation in 1983-1995 turned negative in 1995-2007. Note that even though the population effect declined, the overall participation rate increased much more in the second period than in 1983-1995.

Finally, in order to construct a scenario for future labor supply, we use population projections from the New Cronos database by Eurostat (EUROPOP2008). EUROPOP2008 contains statistical information on population projections with reference to projected 1st

[8]The impact of the change in composition can be measured by applying the change in the population composition between the two periods to the participation rates of the first period, by age and gender groups.

of January population by sex and single year of age, projected vital events (births and deaths) and assumptions concerning fertility, life expectancy at birth by sex and international migration. In the projections, we have made use of two variants: the baseline projection includes migration, while an alternative scenario captures the population developments without migration. Looking forward, population ageing implies that the older age groups within the working age population gain more weight: those above 55 years old that are expected to be around 20% of the working age population in 2015, compared with 17.6% in 2007. In contrast, the weight of the ones below 24 years is expected to decline by 1 percentage point over the same period to 16.5%; and the weight of the group between 35 and 44 years old, i.e., those most attached to the labor market, is expected to decline by more than 2 percentage points to 21.4%. The mechanical decomposition depicted in Table 3.1 therefore suggests a substantial decline (by 0.6 percentage points) in the aggregate participation rate, putting downward pressure on total labor supply and potential growth in the euro area. This downward pressure intensifies significantly (a decline of 2.4 percentage points in the aggregate participation rate) if the horizon is extended up to 2030 when the oldest group (those between 55 to 64 years) is expected to account for one fourth of the working age population.

Figure 3.3: Estimated age-participation profiles in the EA12

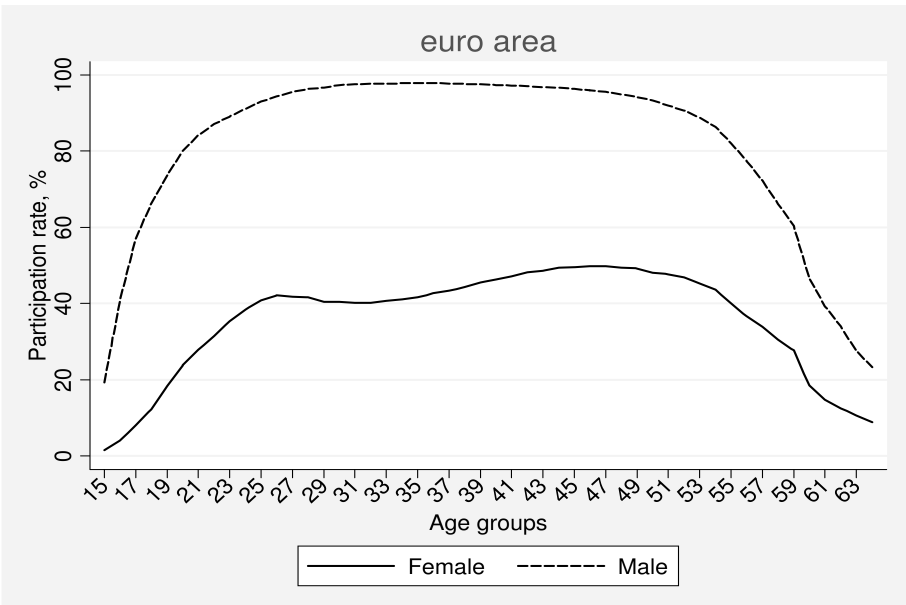

## 3.3 Results

We present results in three parts. We first illustrate the role of age and cohort effects using a basic decomposition of participation rates into age, cohort and business cycle effects. Second, we add a number of indicators of time-varying observed determinants of participation behavior, such as labor market institutions, in the model. Finally, we present projections for male and female participation rates up to 2030. While the focus is on the euro area, we use country results to illustrate and to account for cross-country heterogeneity in participation behavior. In particular, the full model with time-varying observed determinants is estimated separately for the five largest euro area countries (Germany, France, Italy, Spain and the Netherlands) and the results, in terms of trends and projections, are then aggregated to form a euro area 5 aggregate.

### 3.3.1 Basic model

For the basic model, we will investigate the results from the decomposition to age and cohort effects for the euro area (EA12) and the five largest euro area countries. The results suggest substantial and highly statistically significant age effects that show the familiar hump-shaped pattern. For males, underlying euro area participation rates increase until age 30, remain stable until age 50, before gradually declining again (see Figure 3.3). The overall level of the underlying euro area age-participation profile is lower for females. While male participation rates are close to 100% in prime-age, female participation is highest at ages 40-50 at around 45%. In addition, for females the estimated age participation profile features a pronounced dip around early 30s. The dip suggests that a number of women leave the labor market temporarily to start a family, returning to work in their late 30s. The overall pattern of the age effects is similar across euro area countries. However, both the level and shape of the underlying female age-participation profiles differ substantially across countries (see Figure 3.9 in the Appendix to this chapter). While female participation rates peak at close to 70% in France, they do not exceed 40% in Spain and the Netherlands. The post child-bearing-age increase in participation is visible in all countries except Italy. These differences point to significant underlying differences in the participation behavior of European women that are likely to reflect a combination of time-invariant cultural and institutional factors. The results also point to the need to model the female participation rate using a flexible functional form that varies across

Figure 3.4: Estimated cohort effects in the EA12

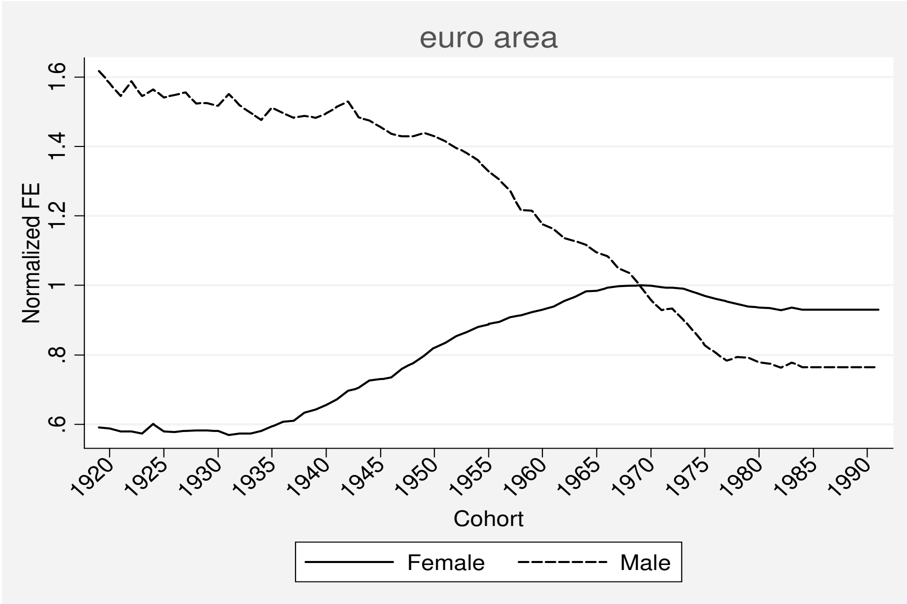


countries. While the male age-participation profile could be characterized by a second order polynomial of age, female age-participation profiles are more complex and cannot be adequately captured by simple polynomials.

The results from the basic model also confirm that cohort effects are statistically significant and robust to age and period effects as measured by an indicator of the business cycle. In line with the descriptive evidence, cohort effects appear more significant in size for females than males. The normalized cohort fixed effects for the euro area are plotted in Figure 3.4 for both males and females. The results show a broadly declining profile for males and an increasing profile for females. The overall pattern of declining cohort effects for men and decreasing cohort effects for women appears similar to that observed in the United States (see Figure 8 in Fallick and Pingle (2007)). This mix of a positive cohort effect for the middle female cohorts and a negative effect for the younger female cohorts has a large impact on overall labor supply and, as demographic change shifts the weight between birth cohorts, turns out to be a relevant factor for future euro area labor supply. Again country results broadly confirm the overall pattern of estimated cohort effects (see Figure 3.10 in the Appendix to this chapter). The relative decline in cohort effects for men varies most across countries, with a substantial decline in Italy contrasted with an increase throughout in the Netherlands.

Combining the age and cohort effects, and excluding the business cycle and the error term, provides a measure of trend participation rates for each age group. For both females and males the actual and trend participation rates show a decline in participation of the younger age groups (up to 20-24 for females and 25-29 for males). For those in prime age and for older workers the trends diverge somewhat. For females, actual and trend participation increase for those in prime working age, and post mid 1990s also for older workers. For males, actual and trend participation rates are either stable or declining for those in prime age, whereas a more recent increase in participation rates for older workers results in a mild u-shaped pattern. Estimated trend participation rates from the simple decomposition capture actual developments reasonably well for most detailed age groups. The model does particularly well in explaining the increasing trend of female participation and the recent increase in the participation of older workers for both males and females (with different timing across genders). In comparison, the results for some age groups suggest that the simple model misses important determinants of participation. Aggregating results for both males and females shows that beyond the broad trends of increasing participation of females and decreasing participation of males, important medium term developments are not fully captured by the simple model. For both males and females this includes a mild slump in participation in the 1990s and the most recent increase beginning around 2004. For males, actual participation rates were also above trend rates as captured by the model in the early 1980s. Overall, while the simple decomposition does well in explaining broad trends in participation, in particular for females, for some groups, age and cohort effects alone are not sufficient to capture trend participation patterns in the euro area.

### 3.3.2 Model with observed determinants

Going beyond the basic model, it is likely that other factors, such as time-varying labor market institutions, may have influenced participation trends in the medium term. Therefore, in a second step, we estimate the cohort model for the five largest euro area countries with a number of indicators of observed determinants that may matter for participation decision. We include union density, labor taxes, implicit tax on retirement for older workers, unemployment benefit replacement rate, the share of highly educated in the youth population, average number of children and life expectancy. The list of indicators is suggested by previous empirical analysis on the impact of institutions on labor force participation (see Bassanini and Duval (2006) and Genre et al. (2005) and 2008) and theoretical considerations. In addition, availability of comparable indicators with sufficient time variation limits the list of relevant institutional factors that are considered (excluding, for example, indicators of employment protection legislation or the retirement age).

A number of hypotheses about the likely impact of these institutional factors can be put forward. First, we expect that declining union density in a number of euro area countries may have contributed to increase participation through its positive impact on expectations about the availability of jobs to those that have been previously inactive. As unions tend to compress the wage distribution, the decline in unionization may have more of an impact on those at the lower part of the wage distribution (more likely to be younger and older workers). Second, an increase in labor taxes (observed in a number of euro area countries) over time is also expected to result in lower labor participation by making leisure relatively less expensive. However, from a household labor supply perspective an increase in labor taxes for the head of the household may also result in an increased propensity to participate for other members of the household (more likely to be women). For older workers the implicit tax on continued work, a summary measure of retirement incentives, is likely to be more relevant than the overall labor tax. A higher implicit tax rate is expected to lower incentives to retire early (for the ages 55-64 considered here) and therefore to increase the participation rate of older workers. Third, observed declines in the generosity of the unemployment benefits system, as measured by the replacement rate, in a number of euro area countries is likely to lower the incentive to participate in the labor market by lowering alternative income when unemployed relative to inactivity. By contrast, unemployment benefits may also have a positive impact on participation via wage bargaining, with lower generosity leading to weakening of the insider's position in the labor market relative to the outsiders, or as a proxy for the overall generosity of the welfare system. Fourth, longer life expectancy is likely to lead to higher participation for older workers as they remain active and may also anticipate a longer period of retirement. Fifth, the higher share of young in education relative to older workers is expected to lower participation of young workers. Finally, the number of children is expected to influence female participation rates, with more children lowering participation rates of women around the typical age for starting a family. While union density, unemployment benefits and labor taxes are included in the equations for all age and gender groups (in working age), variables relating to education are included only for the youngest workers, life expectancy and the implicit tax on continued work for the oldest workers and, finally, the number of children for females only.

Tables 3.2 and 3.3 show the aggregated coefficient estimates and their t-statistics of the observed determinants of participation for three main age groups: young (15-24), prime-aged (25-54) and older (55-64), for all five countries. To simplify comparisons of coefficient estimates across groups and indicators, the data on observed determinants has been standardized. Note that the identification of the impact of institutions here relies only on available within-country time variation, which is often limited for the indicators of labor market institutions considered here. As a result, relatively few indicators turn out to be statistically significant. With this caveat in mind, a number of institutional indicators seem to matter, although the magnitude, and in some cases the sign, varies across countries and age groups. Higher labor taxes tend to lower participation rates (as reflected in 16 out 19 statistically significant coefficients). This impact is estimated more consistently for males in all countries. Higher union density (in 11 out of 15 statistically significant coefficients) and more generous unemployment benefits (17 out of 24 statistically significant coefficients) also tend to lower participation rates. The negative impact of unemployment benefits is consistent with the interpretation that unemployment benefits impact participation rates either via their impact on bargaining (with increased power for insiders leading to higher bargained wages and lower participation rates for outsiders) or via their role as a proxy of the overall generosity of the welfare state (more generous benefits tend to coincide with more generous welfare benefits for financing non-participation, lowering participation rates). Exceptions to this result occur mainly for young people, whose participation rates in some countries are positively associated with unemployment benefits. The results also suggest that unemployment benefits increase participation of all males in Germany. While not conclusive, these results are suggestive of negative incentive effects for the unemployed stemming from generous unemployment benefits that are also of relatively long duration. In this case, a decline in benefits over time would lead some unemployed workers (who may have not been actively looking for jobs) to leave the labor force altogether. Overall, the results for union density and unemployment benefits are broadly in line with panel regression results in Genre et al. (2005) and 2008, who also find that higher union density and more generous unemployment benefits lower participation rates.

For females, with the exception of young females in France and the Netherlands, higher number of children tends to lower participation. The decline in number of children in most euro area countries is therefore associated with an increase in female participation rates.

Table 3.2: Impact of observed determinants: males

| | LT | UD | UB | TR | LE | YE |
|---|---|---|---|---|---|---|
| Germany: | | | | | | |
| Young | 0.00 | -0.01 | 0.10 | | | 0.05 |
| | (-0.22) | (-0.33) | (5.07) | | | (3.68) |
| Prime-aged | -0.07 | -0.03 | 0.03 | | | |
| | (-7.03) | (-1.43) | (1.97) | | | |
| Older | -0.12 | -0.04 | 0.07 | 0.03 | 0.40 | |
| | (-7.89) | (-1.02) | (3.34) | (0.54) | (7.29) | |
| France: | | | | | | |
| Young | -0.16 | 0.46 | 0.05 | | | -0.27 |
| | (-4.86) | (5.14) | (2.54) | | | (-5.00) |
| Prime-aged | -0.07 | -0.04 | -0.02 | | | |
| | (-5.10) | (-2.10) | (-2.28) | | | |
| Older | 0.02 | -0.09 | -0.06 | 0.02 | -0.09 | |
| | (0.86) | (-2.06) | (-5.49) | (1.44) | (-2.47) | |
| Italy: | | | | | | |
| Young | -0.14 | -0.07 | -0.32 | | | 0.20 |
| | (-6.10) | (-2.00) | (-9.00) | | | (2.95) |
| Prime-aged | -0.11 | 0.00 | -0.10 | | | |
| | (-4.83) | (-0.15) | (-3.85) | | | |
| Older | 0.00 | -0.15 | 0.06 | 0.04 | 0.21 | |
| | (0.08) | (-4.43) | (1.54) | (3.12) | (3.06) | |
| Spain: | | | | | | |
| Young | -0.03 | -0.04 | -0.12 | | | 0.05 |
| | (-3.18) | (-2.24) | (-4.27) | | | (0.85) |
| Prime-aged | -0.08 | -0.01 | -0.11 | | | |
| | (-6.52) | (-0.50) | (-7.65) | | | |
| Older | -0.03 | 0.01 | -0.07 | -0.03 | 0.05 | |
| | (-1.78) | (0.62) | (-2.33) | (-0.59) | (1.49) | |
| Netherlands: | | | | | | |
| Young | -0.08 | -0.34 | -0.08 | | | -0.02 |
| | (-1.91) | (-5.21) | (-2.97) | | | (-0.25) |
| Prime-aged | -0.03 | -0.18 | -0.08 | | | |
| | (-1.37) | (-4.42) | (-4.21) | | | |
| Older | -0.11 | 0.05 | 0.04 | 0.42 | 0.14 | |
| | (-2.66) | (0.71) | (1.30) | (4.49) | (2.20) | |

Note: T-statistics based on robust standard errors in parenthesis. LT is labor taxes, UD is union density, UB is unemployment benefits, TR is tax on retirement, LE is life expectancy, YE is youth education. For each age group, the coefficients and their standard errors have been aggregated from single ages using labor force weights in 2007.

Table 3.3: Impact of observed determinants: females

| | LT | UD | UB | TR | LE | YE | NC |
|---|---|---|---|---|---|---|---|
| Germany: | | | | | | | |
| Young | 0.00 | 0.13 | 0.06 | | | 0.03 | -0.12 |
| | (0.35) | (4.70) | (4.22) | | | (3.26) | (-10.04) |
| Prime-aged | -0.04 | 0.02 | 0.01 | | | | -0.01 |
| | (-8.08) | (1.70) | (1.01) | | | | (-1.51) |
| Older | -0.04 | 0.00 | -0.05 | 0.06 | -0.03 | | -0.02 |
| | (-3.78) | (-0.10) | (-3.98) | (2.12) | (-0.95) | | (-2.07) |
| France: | | | | | | | |
| Young | -0.14 | 0.51 | 0.07 | | | -0.29 | 0.17 |
| | (-3.96) | (4.18) | (3.36) | | | (-5.11) | (1.89) |
| Prime-aged | 0.03 | -0.02 | -0.01 | | | | 0.00 |
| | (2.47) | (-1.07) | (-2.27) | | | | (-0.11) |
| Older 0.06 | -0.19 | -0.02 | -0.03 | -0.09 | | 0.08 | |
| | (1.99) | (-3.64) | (-1.49) | (-1.50) | (-2.22) | | (0.91) |
| Italy: | | | | | | | |
| Young | -0.11 | -0.09 | -0.22 | | | 0.59 | -0.04 |
| | (-5.79) | (-2.14) | (-6.94) | | | (8.85) | (-1.95) |
| Prime-aged | -0.07 | -0.01 | -0.10 | | | | -0.02 |
| | (-6.18) | (-0.84) | (-8.17) | | | | (-2.55) |
| Older | -0.01 | -0.04 | -0.06 | 0.01 | 0.06 | | 0.00 |
| | (-0.28) | (-1.29) | (-1.75) | (0.86) | (0.94) | | (-0.13) |
| Spain: | | | | | | | |
| Young | 0.03 | 0.05 | -0.09 | | | -0.01 | 0.15 |
| | (2.31) | (1.42) | (-3.02) | | | (-0.05) | (1.03) |
| Prime-aged | 0.01 | 0.03 | -0.06 | | | | -0.10 |
| | (1.28) | (2.88) | (-4.02) | | | | (-1.29) |
| Older | 0.00 | -0.04 | -0.05 | -0.01 | 0.01 | | 0.06 |
| | (0.04) | (-1.11) | (-1.61) | (-0.08) | (0.35) | | (0.19) |
| Netherlands: | | | | | | | |
| Young | 0.03 | -0.15 | 0.05 | | | 0.04 | 0.15 |
| | (0.71) | (-2.41) | (2.72) | | | (0.58) | (4.73) |
| Prime-aged | -0.01 | -0.07 | -0.04 | | | | -0.04 |
| | (-0.31) | (-3.39) | (-2.67) | | | | (-2.07) |
| Older | 0.02 | -0.11 | -0.02 | 0.20 | -0.03 | | -0.05 |
| | (0.32) | (-1.67) | (-0.42) | (2.20) | (-0.52) | | (-0.76) |

Note: T-statistics based on robust standard errors in parenthesis. LT is labor taxes, UD is union density, UB is unemployment benefits, TR is tax on retirement, LE is life expectancy, YE is youth education and NC is number of children. For each age group, the coefficients and their standard errors have been aggregated from single ages using labor force weights in 2007.

This is also in line with the Genre et al. (2008) finding that the fertility rate is negatively associated with participation rate of prime-aged females.[9] Other group specific variables appear to be estimated less consistently, with both the sign and statistical significance changing across age groups and countries. Higher implicit tax on retirement, in the few cases when it is statistically significant, increases participation of older workers. With few counterintuitive exceptions (older people in France) increased life expectancy also increases participation of older workers. Both the sign and statistical significance of the share of youth in education varies across countries, suggesting that investment in human capital may not be well captured in the model.[10]

As regards the business cycle, we find that the sum of the coefficients of current and two lags of the output gap for worker groups are often not statically significant (not shown). In addition, for a number of groups we find a negative business cycle effect. For some groups, such as young people and females, this result could reflect "added worker" effects. For example, for individuals in families with a main bread-winner, in good times labor income from the rest of the family members may not be needed, whereas additional income from a second job is needed in bad times.[11] We tried other indicators of the business cycle (unemployment and employment gap measures) with similar results. We therefore conclude that the business cycle has little influence on participation decisions in these countries, in line with results that show that European unemployment and employment rates are mainly influenced by structural factors or interactions of structural factors and shocks (e.g. Bassanini and Duval (2006) and Blanchard and Wolfers (2000)).

In order to illustrate the size of the total impact of observed determinants we compare the total change in trend participation rates as measured by our model with a scenario of keeping the observed determinants at their 1995 values. The scenario reflects the view that the acceleration of labor market reforms from the second half of the 1990s onwards has

[9] We also experimented with other determinants of female participation, in particular, the tax rate on second earners and marriage rate. Previous literature has suggested that both are potentially important determinants (see Jaumotte (2003). The tax rate on second earners was usually not statistically significant for these countries. We found some (counterintuitive) indication that the marriage rate is positively associated with participation. Both variables where therefore excluded from the final model. Jaumotte (2003) and Genre et al. (2008) exploit cross-section variation to establish other potential determinants of female participation (such as maternity leave) that we do not consider here.

[10] These effects could be captured better by changes in returns to education. However, we are not aware of comparable estimates of returns to education with a sufficiently long time-series that we could use.

[11] Prieto-Rodriguez and Rodriguez-Gutierrez (2000) find these effects to be relevant for women in Spain, in line with our finding of negative business cycle effects for women of all ages.

Figure 3.5: Total impact of observed determinants

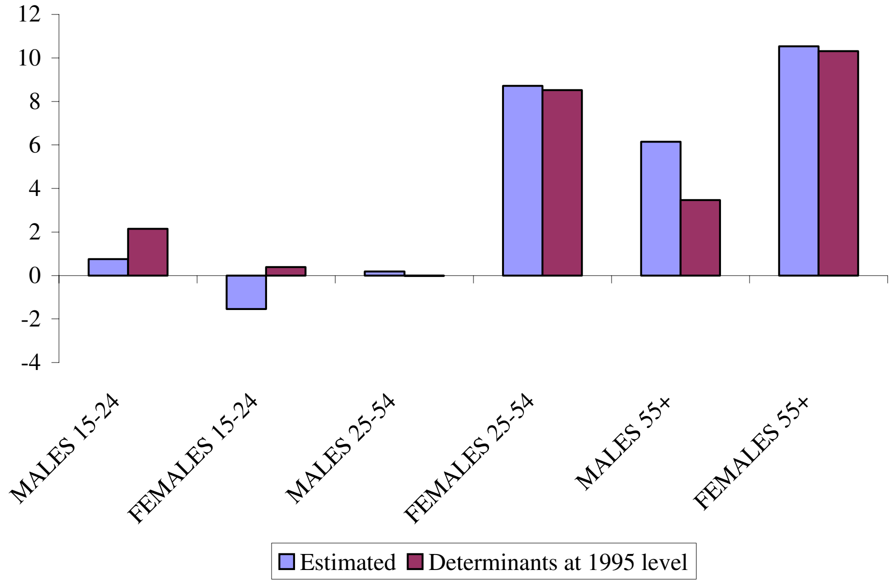


contributed positively to participation rates (see for example, Masuch and Force (2008)). The results are shown in Figure 3.5. The positive impact of observed determinants on participation is most evident for older males, as reflected in the large gap between the two bars. The most relevant variable in this respect appears to be life expectancy. The increase in life expectancy since 1995 has had a positive impact in the participation rate of older males. Overall, the observed determinants have resulted only in small increase in participation rates for females, with most of the increase over this time period attributed to age and cohort effects instead. For young people, the impact of observed determinant has been to dampen participation rates. This is partly explained by the increase in the proportion of young people in education. At the country level, it is worth mentioning that the impact of the change in the institutional framework is broadly based for prime-age age males and females, and for females aged 55 and over. In contrast, for the young, the developments are strongly influenced by the results for Italy and France, and for males aged 55 and over by the results for Netherlands and Germany.

Both age and cohort effects remain jointly statistically significant in all models even after including business cycle indicators and other time-varying determinants of participation. These coefficients can be thought of as capturing the impact of other time-invariant cultural or institutional factors (for the age coefficients) or slowly changing impact of factors

Figure 3.6: Estimated cohort profiles in the EA5

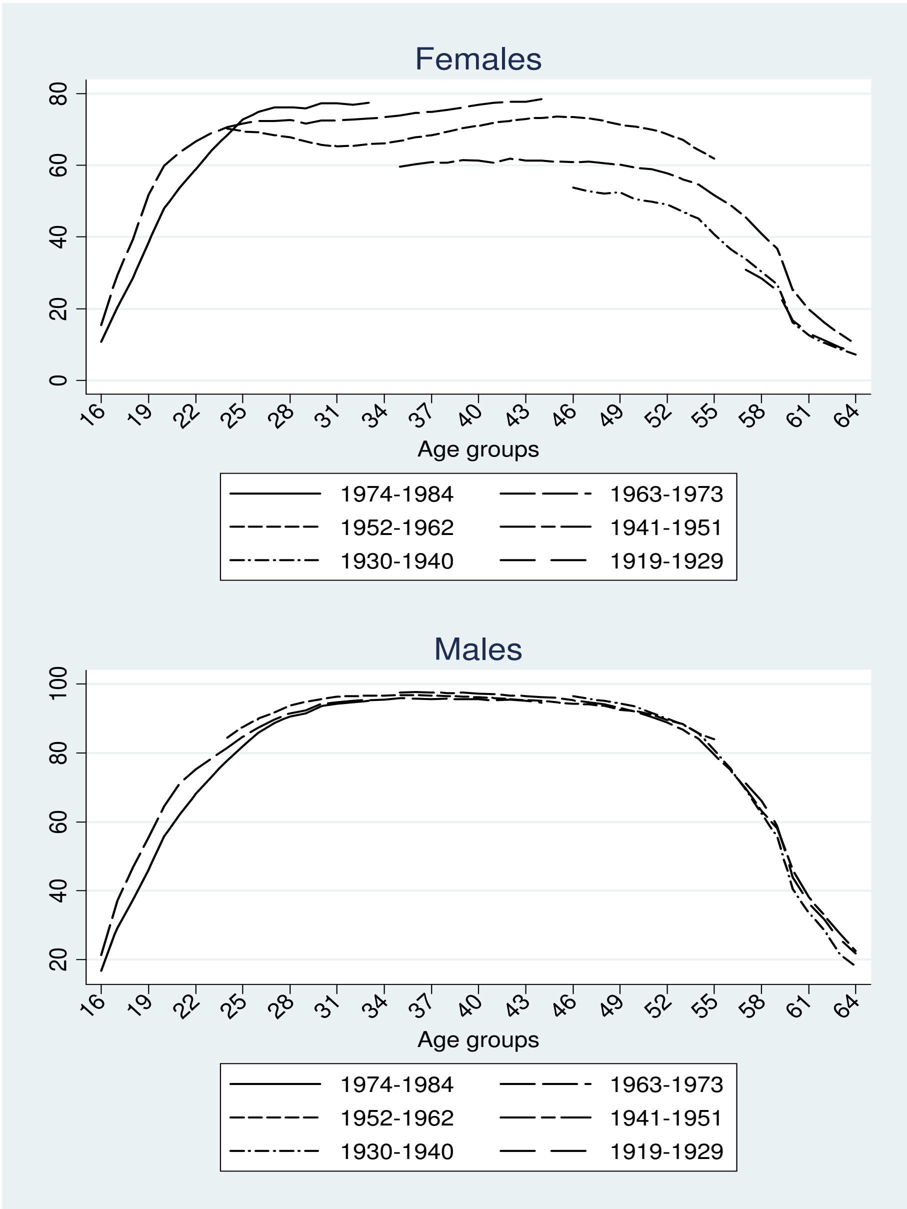

that are specific to birth years (for the cohort coefficients). The latter may include factors such as cultural attitudes towards labor market participation (for women in particular) or institutional factors and reforms that are not captured by the observed determinants. Figure 3.6 plots the estimated cohort profiles based on the trend participation rates from the model, aggregated to the euro area five (EA5) level. For males, the lines indicating participation rates for specific cohorts are mostly overlapping. There is some indication that most recent cohorts enter the labor market later, reflecting the substantial increase in the average number of years spent in education. The same impact is visible also for the youngest female cohort. However, in addition, the cohort profile for females suggests a substantial shifting up of the age participation profile over time. For prime-aged women, those in their mid 30s and 40s, the participation rate has increased by more than 20 percentage points. Furthermore, while the cohort profile for those born between 1953 and 1962 shows a pronounced dip at child bearing age, this dip is not visible for the next cohort (those born between 1963 and 1973). The higher propensity to participate of females born in the late 1960s and early 1970s has therefore contributed to the increase in female participation in the euro area.

Country results show that participation behavior differs across the largest five euro area countries and that again this is most evident for women. For women in their 30s and 40s, the estimated cohort profiles show that participation rates of most recent cohorts has increased most, by more than 30 percentage points, in the Netherlands and Spain and the least in France, with roughly 10 percentage points (see Figure 3.11 in the Appendix to this chapter). The disappearance of the dip at child bearing age is most pronounced in the Netherlands: while the participation rates of women in the late 20s and early 30s for those born in 1950s dropped by as much as 20 percentage points, the more recent cohorts appear to have stayed in the labor market through the child bearing years.

### 3.3.3 Projections

In a third and final step we use the model results to project participation rates forward until 2030. The results of both trend and projection for the euro area (EA5) are obtained by aggregating the full model estimates for Germany, France, Italy, Spain and the Netherlands weighting the countries with their respective populations. We assume that age and cohort effects are fixed throughout the sample and keep observed determinants at their 2007 values. In addition, for the young cohorts, i.e., the last eight cohorts of our sample and

Figure 3.7: Trend and participation and projections in the EA5 by gender, 1986-2030

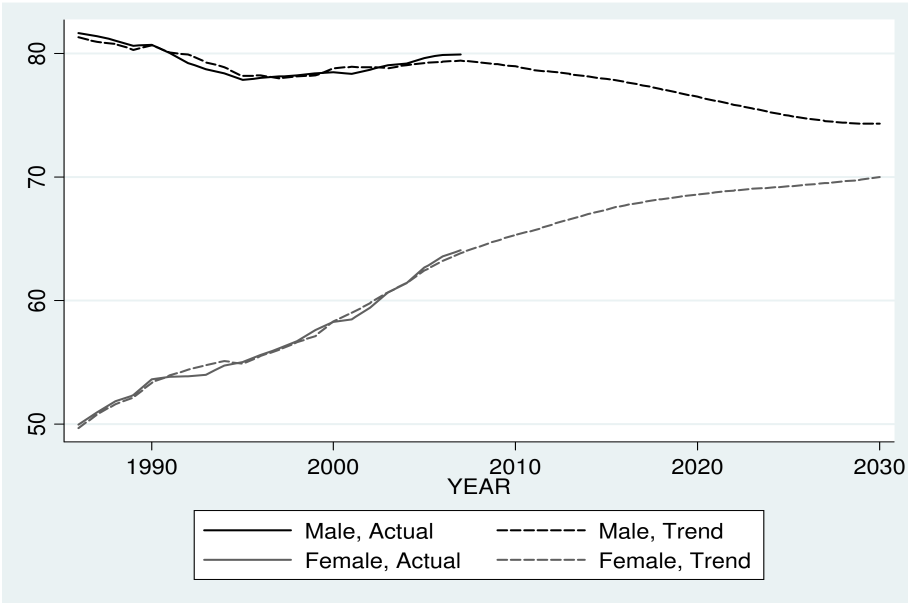


those that enter the labor market after 2007, we fix their cohort effects at the level of the last cohort effect we estimate, namely those born in 1984. Figure 3.7 shows the trend from the estimation of the full model for males and females together with the actual participation rates for the euro area (EA5). The results clearly show that within the sample period the full model captures both trends and medium-term developments well. This is confirmed by results for individual ages shown in Figures 3.12 to 3.17 in the Appendix to this chapter.

The projected euro area participation rate decreases for males throughout the projection period. In contrast, the euro area participation rate for females increases before stabilizing at about 70% in 2030. This pattern is in line with the waning impact of the positive cohort effects for females that continue to support participation rates looking forward. While the gap between male and female participation rates is expected to decline substantially, at the end of the projection horizon male participation rate remains 4.6 percentage points above the female participation rate. Overall, towards the end of the sample the negative impact of population ageing shifting the larger share of the population to older age groups with lower participation rates begins to dominate and dampen the overall participation rate. As a result, the overall participation rate is anticipated to increase slightly up to 2015, by 1 percentage point, but to decline thereafter. However, in 2030 it is still expected to remains at just above the 2007 level (see Table 3.4). The underlying country results from the baseline model are shown in Table 3.5. The results for all countries point to an ongoing increase in female participation and a decline in male participation. Indeed, in the Netherlands and France the gap between male and female participation closes by 2030. Reflecting the continued positive trend in female participation, the overall participation rate is expected to continue increasing in all countries except Germany.

In order to explore these results further and to evaluate robustness we calculate three additional scenarios for the euro area (see Table 3.4). First, we compare our results with a scenario that keeps participation rates by age and gender groups unchanged at their 2007 level, i.e., accounting only for population effects. The model based results imply a more positive outlook for participation than a scenario based on unchanged participation behavior. Indeed, in the latter scenario, the overall participation rate declines already in 2015, with a gap of 3.9 percentage points in 2030 between the baseline model results and the alternative scenario. Second, we calculate a scenario that accounts for the impact of migration through the population structure. Specifically we compare the baseline model results with a scenario that assumes no migration and find that the impact of migration through the population structure is positive. On average, migrants tend to be younger and therefore to have higher participation rates than the native population. The impact is relatively small, but its relevance grows over time - the gap in the participation rate between the baseline migration and non-migration scenario is 0.1 percentage point in 2015, but reaches 0.8 percentage point in 2030. A comparison of the scenarios by gender shows that migration is only relevant for the male participation rate, while the impact on the female participation rate is negligible. Finally, we compare our results with those derived from the participation rate projections at the country level published by the European Commission (EuropeanCommission (2008)). We find that there is a significant gap between our baseline model results and the European Commission projections. According to the European Commission, the overall participation rate is expected to increase somewhat more, by 2.2 percentage points in 2015 and 3 percentage points in 2030. This gap reflects a substantially more positive outlook for male participation in the European Commission projections - for males the gap in participation is 2.4 and 5.8 percentage points in 2015 and 2030 respectively. While it is not straightforward to decompose the difference in terms of underlying determinants, the European Commission projections appear to incorporate more inertia from recent participation trends for males. Note that we keep the effect of institutional variables unchanged in the model based projections. Therefore, recent changes in observed determinants that persist or have lagged effects are not reflected in our model based scenario. In contrast, reflecting the important role of cohort effects in explaining past participation trends, our results suggest a somewhat more positive outlook for female participation.

## 3.4 Conclusion

We use a cohort based model of labor force participation to analyze determinants of participation for disaggregated groups of workers in European countries, with a focus on the euro area. The model identifies significant age and cohort effects for detailed worker groups as indicators of (unobserved) structural determinants. We use observed structural determinants and age and cohort effects to construct trend measures of labor supply and to disentangle the impact of structural and business cycle factors on labor force participation rates.

Our results suggest that age and cohort effects can explain a substantial part of the recent increase in labor force participation rates in the euro area, although not the surge since early 2000s. Cohort effects are particularly relevant for women, with those born in the 1920s and 1930s less likely and those born in the late 1960s and early 1970s more likely to participate in the labor market over the life-cycle. There is substantial variation in cohort effects across the five largest euro area countries that we analyze. While cohort effects generally encompass any factor associated with a particular birth year, we speculate that the cohort effects that we observe reflect evolving preferences or social norms that vary across countries. Depending on the country, the estimated cohort profiles suggest an increase of 10 to 30 percentage points in female participation rates. We control for a number of observed time-varying institutions, such as labor taxes, union density, unemployment benefits and the average number of children and find that they have had an impact on labor force participation rates, although the specific impact varies across age and gender groups and countries. Looking forward, while they continue to provide some upward support to participation rates of women in the euro area, positive cohort effects are not large enough to compensate for the downward impact of population ageing on labor force participation rates in the euro area.

# Appendix to Chapter 3: Additional Tables and Graphs

Table 3.4: Alternative scenarios for future participation rates (EA5)

| | 2007 | 2015 | 2020 | 2025 | 2030 |
|---|---|---|---|---|---|
| Total participation rate | | | | | |
| PR (baseline model) | 72.0 | 73.0 | 72.9 | 72.5 | 72.5 |
| PR (model - no migration) | 72.0 | 72.9 | 72.6 | 71.9 | 71.8 |
| PR (2007 level) | 72.0 | 71.3 | 70.3 | 69.1 | 68.6 |
| PR (EC) | 72.0 | 74.2 | 74.6 | 74.6 | 75.0 |
| Females participation rate | | | | | |
| PR (baseline model) | 64.1 | 67.6 | 68.8 | 69.5 | 70.2 |
| PR (model - no migration) | 64.1 | 67.5 | 68.7 | 69.2 | 69.9 |
| PR (2007 level) | 64.1 | 63.1 | 62.1 | 60.9 | 60.5 |
| PR (EC) | 64.1 | 67.6 | 68.5 | 68.8 | 69.5 |
| Males participation rate | | | | | |
| PR (baseline model) | 79.9 | 78.4 | 76.9 | 75.4 | 74.8 |
| PR (model - no migration) | 79.9 | 78.2 | 76.5 | 74.6 | 73.8 |
| PR (2007 level) | 79.9 | 79.4 | 78.4 | 77.1 | 76.5 |
| PR (EC) | 79.9 | 80.8 | 80.8 | 75.9 | 80.6 |

Notes: Euro area obtained as the aggregation of Germany, Italy, France, Spain and the Netherlands. PR (2007 level) refers to a scenario based on unchanged participation rates at the 2007 level. PR (EC) refers to a scenario derived from European Commission (2008); it has been re-based to the 2007 level derived from the EU-LFS. Sources: EU LFS (Eurostat) and own calculations.

Table 3.5: Country projections

| | 1987 | 1997 | 2007 | 2015 | 2020 | 2025 | 2030 |
|---|---|---|---|---|---|---|---|
| Germany | | | | | | | |
| Total | 69.6 | 70.4 | 75.5 | 75.1 | 74.3 | 73.2 | 72.7 |
| Females | 57.6 | 61.8 | 70.0 | 71.4 | 71.5 | 71.1 | 70.8 |
| Males | 81.8 | 78.8 | 81.0 | 78.7 | 77.0 | 75.2 | 74.5 |
| France | | | | | | | |
| Total | 67.1 | 68.0 | 68.6 | 68.8 | 68.8 | 68.7 | 68.8 |
| Females | 57.0 | 61.1 | 64.1 | 66.4 | 67.6 | 68.5 | 69.3 |
| Males | 77.6 | 75.2 | 73.3 | 71.2 | 70.0 | 68.8 | 68.4 |
| Italy | | | | | | | |
| Total | 59.7 | 58.7 | 62.7 | 66.3 | 67.0 | 67.3 | 68.3 |
| Females | 41.9 | 44.0 | 50.9 | 57.7 | 60.3 | 62.1 | 64.1 |
| Males | 78.3 | 73.6 | 74.5 | 74.8 | 73.7 | 72.5 | 72.5 |
| Spain | | | | | | | |
| Total | 58.3 | 62.8 | 72.5 | 74.2 | 74.0 | 73.3 | 72.8 |
| Females | 36.9 | 47.7 | 61.9 | 66.3 | 67.5 | 68.0 | 68.5 |
| Males | 80.3 | 78.0 | 82.7 | 82.0 | 80.3 | 78.4 | 76.9 |
| Netherlands | | | | | | | |
| Total | 63.9 | 71.4 | 78.2 | 79.6 | 79.9 | 80.3 | 81.0 |
| Females | 48.4 | 61.2 | 71.8 | 76.7 | 78.7 | 80.4 | 81.9 |
| Males | 79.1 | 81.4 | 84.4 | 82.4 | 81.1 | 80.2 | 80.0 |

EU LFS (Eurostat) and own calculations.

Figure 3.8: Age-participation profiles by gender in the euro area (EA12), 2007

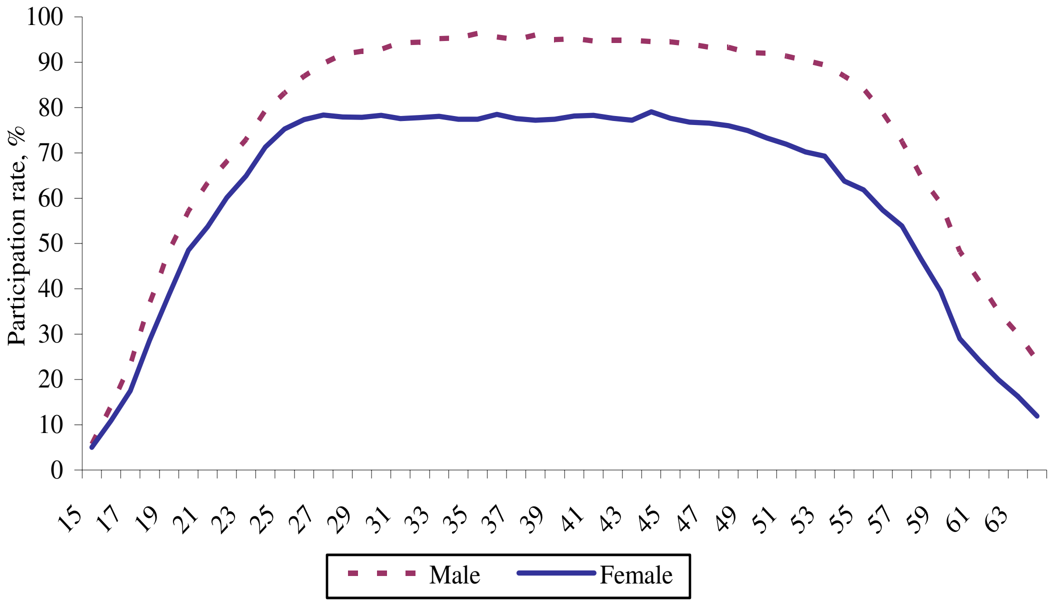


Sources: EU LFS (Eurostat)and own calculations.

Figure 3.9: Estimated age-participation profiles by country

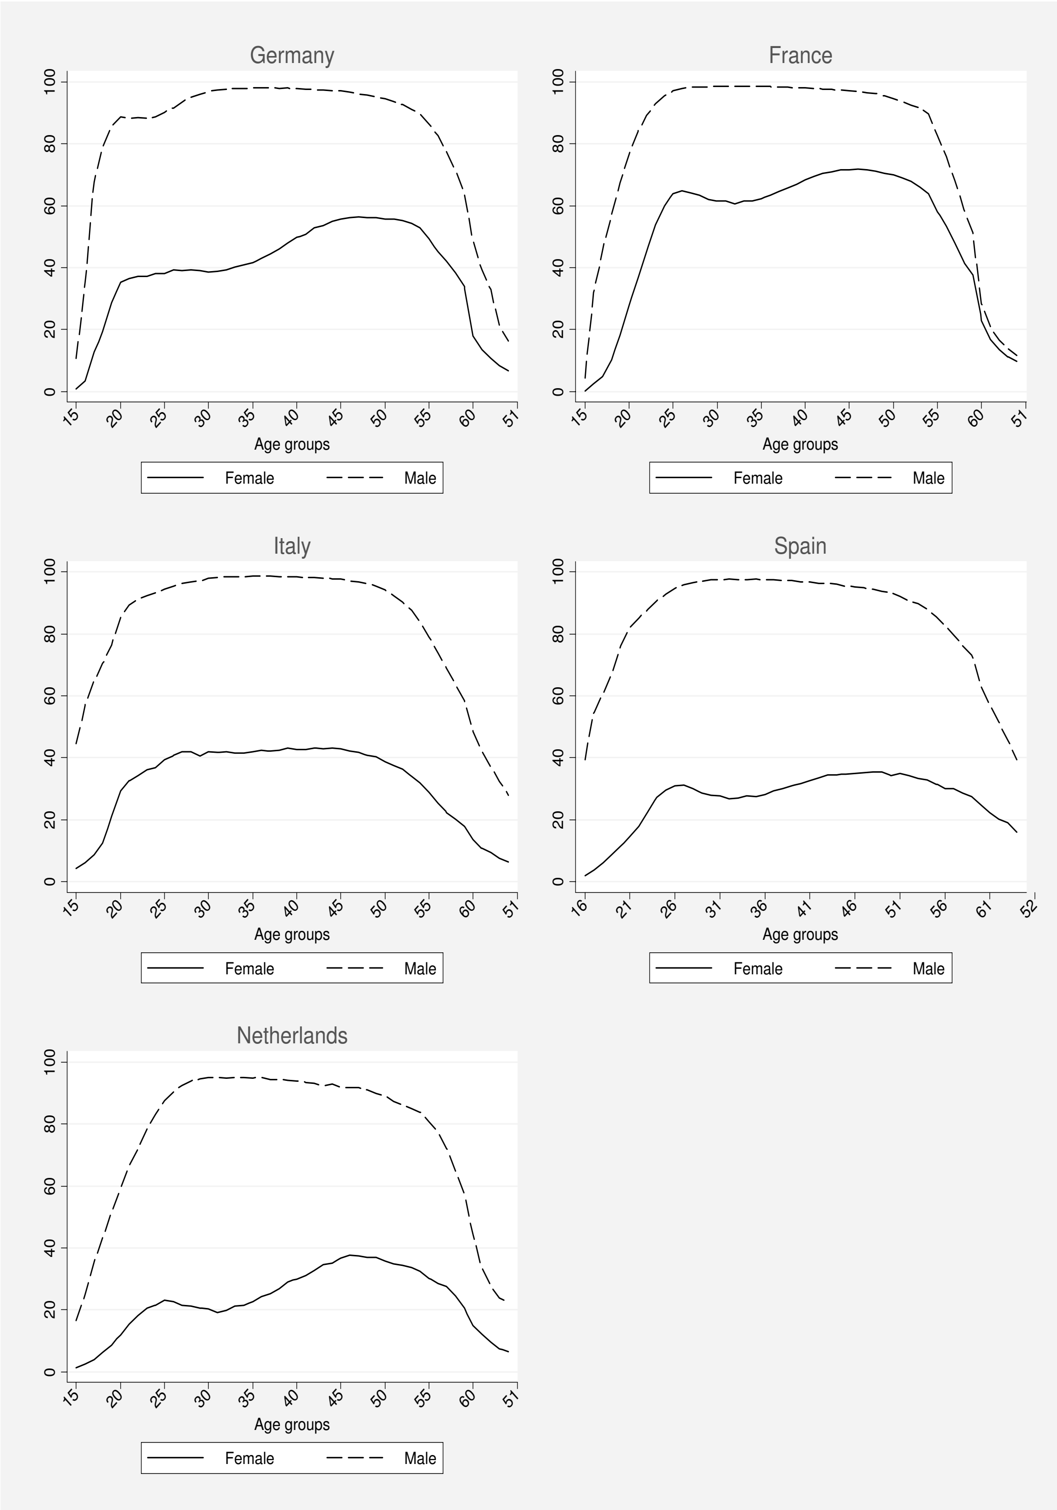

Figure 3.10: Estimated cohort effects by country

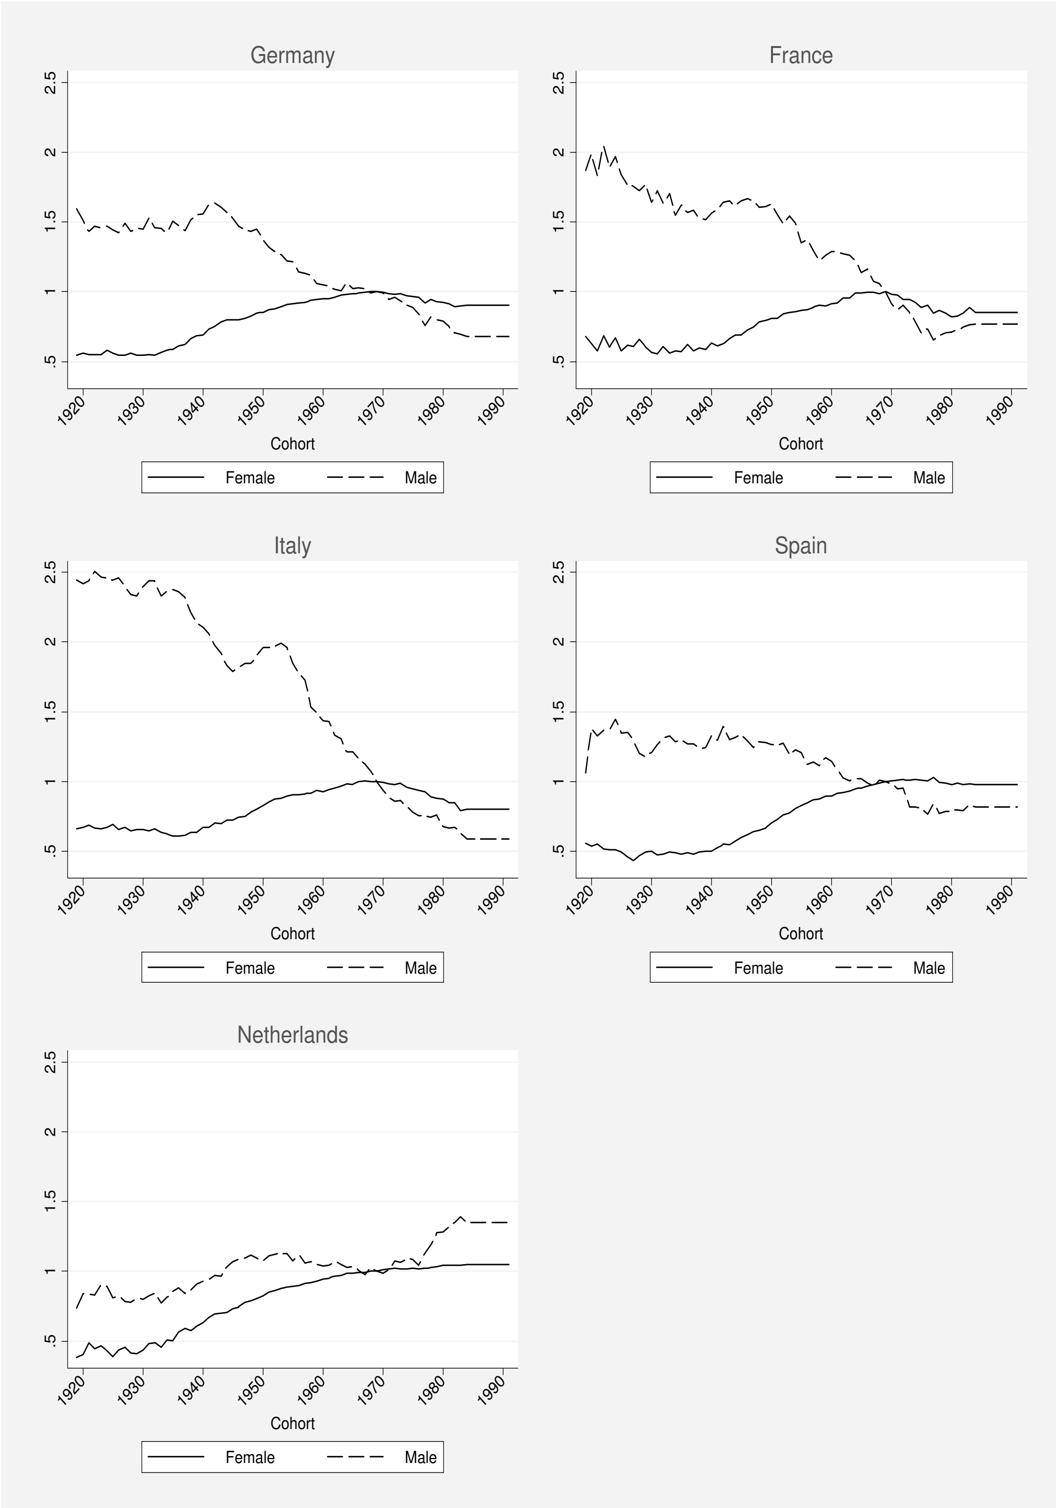

Figure 3.11: Estimated cohort profiles by country, females

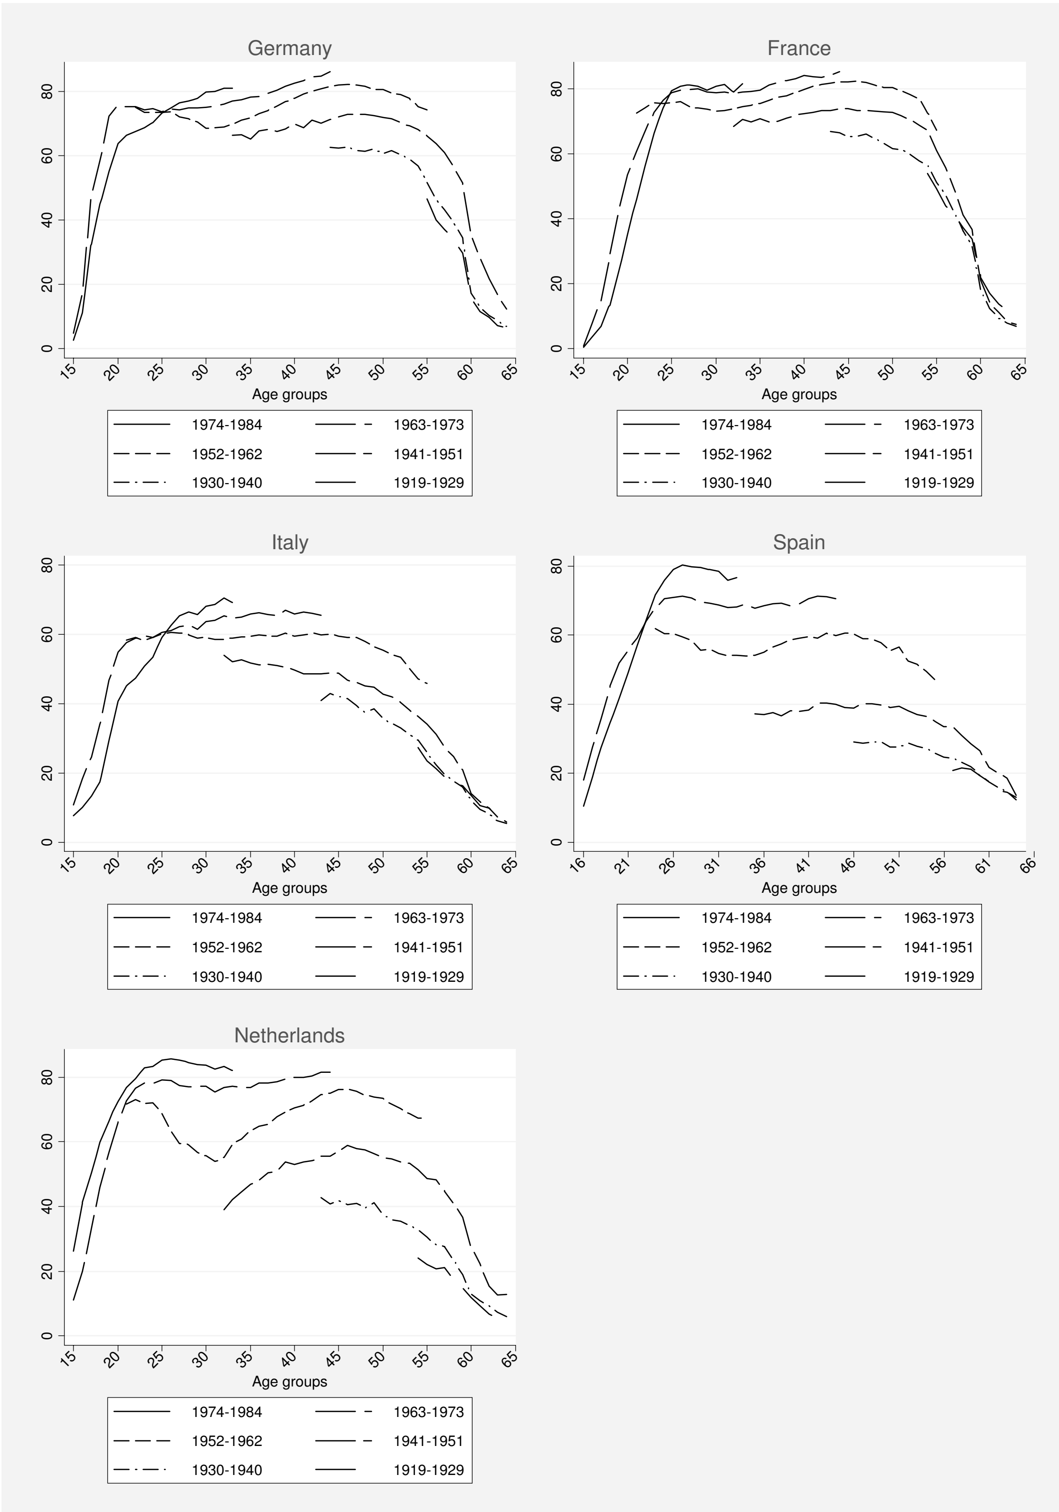

Figure 3.12: Trend and participation rates in the EA5: young females

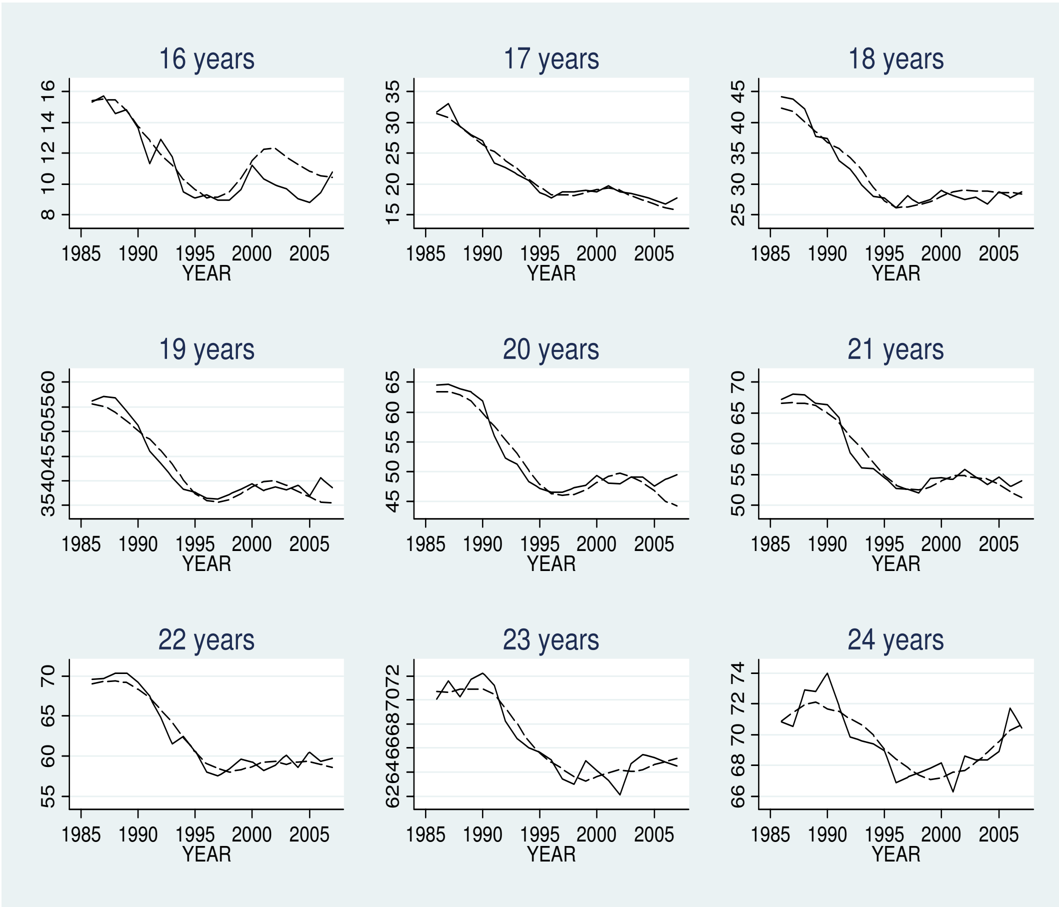


Notes: Trend is a three-year moving average of the estimated trend participation rate.

Figure 3.13: Trend and participation rates in the EA5: prime-age females

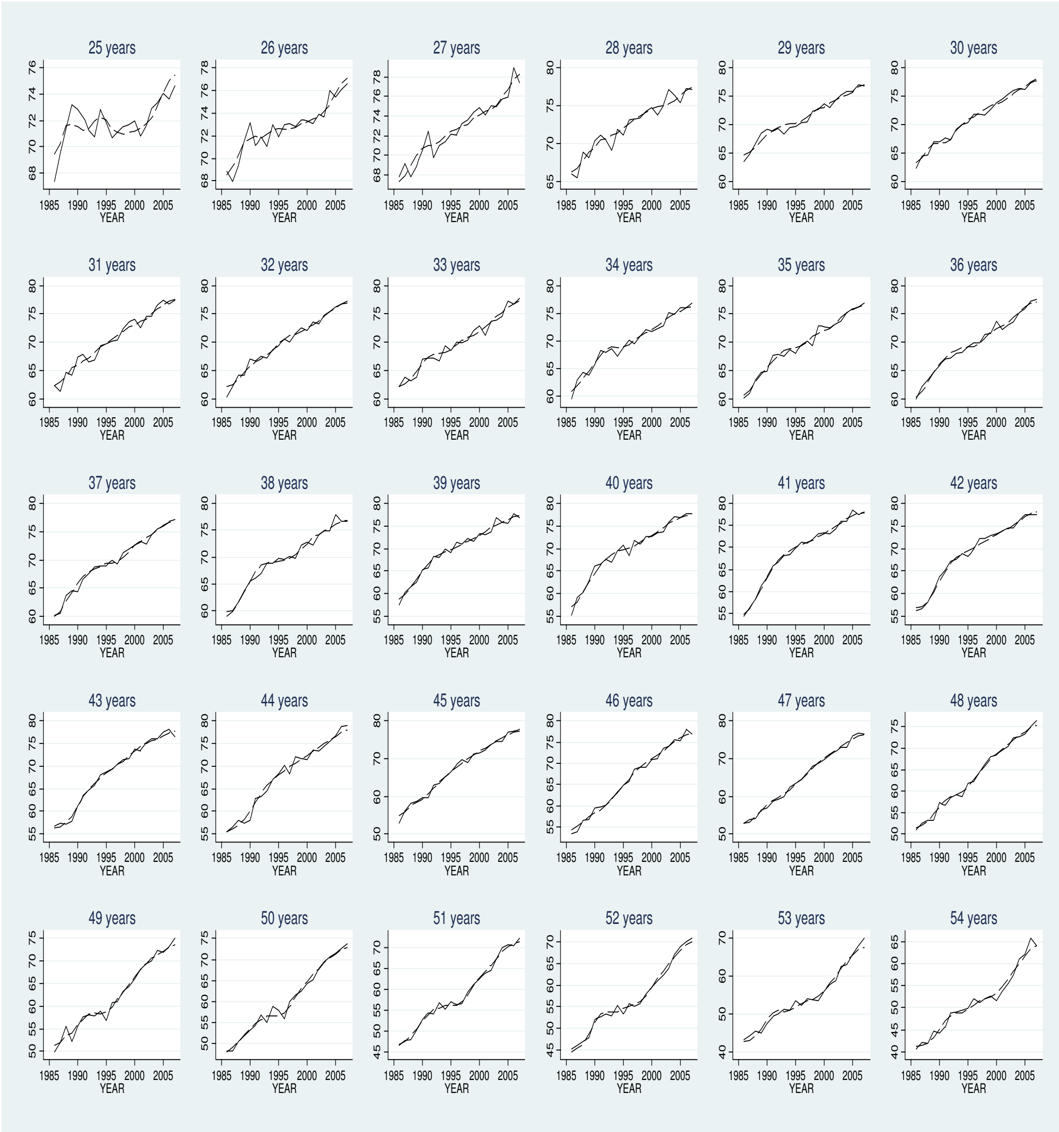


Notes: Trend is a three-year moving average of the estimated trend participation rate.

Figure 3.14: Trend and participation rates in the EA5: older females

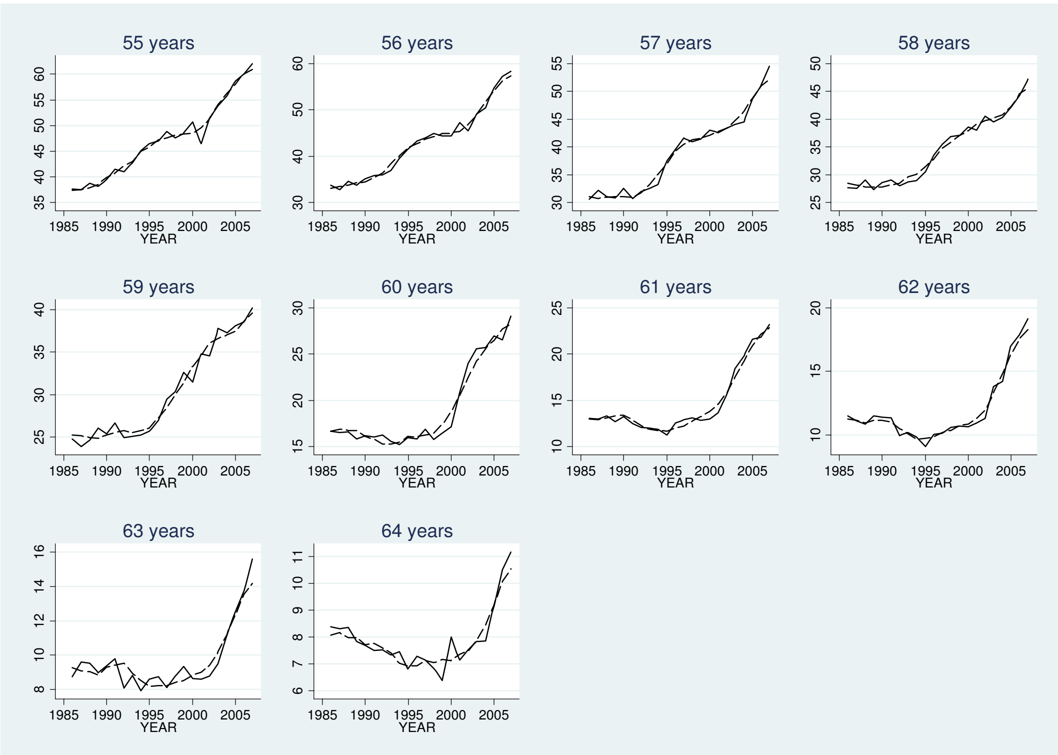


Notes: Trend is a three-year moving average of the estimated trend participation rate.

Figure 3.15: Trend and participation rates in the EA5: young males

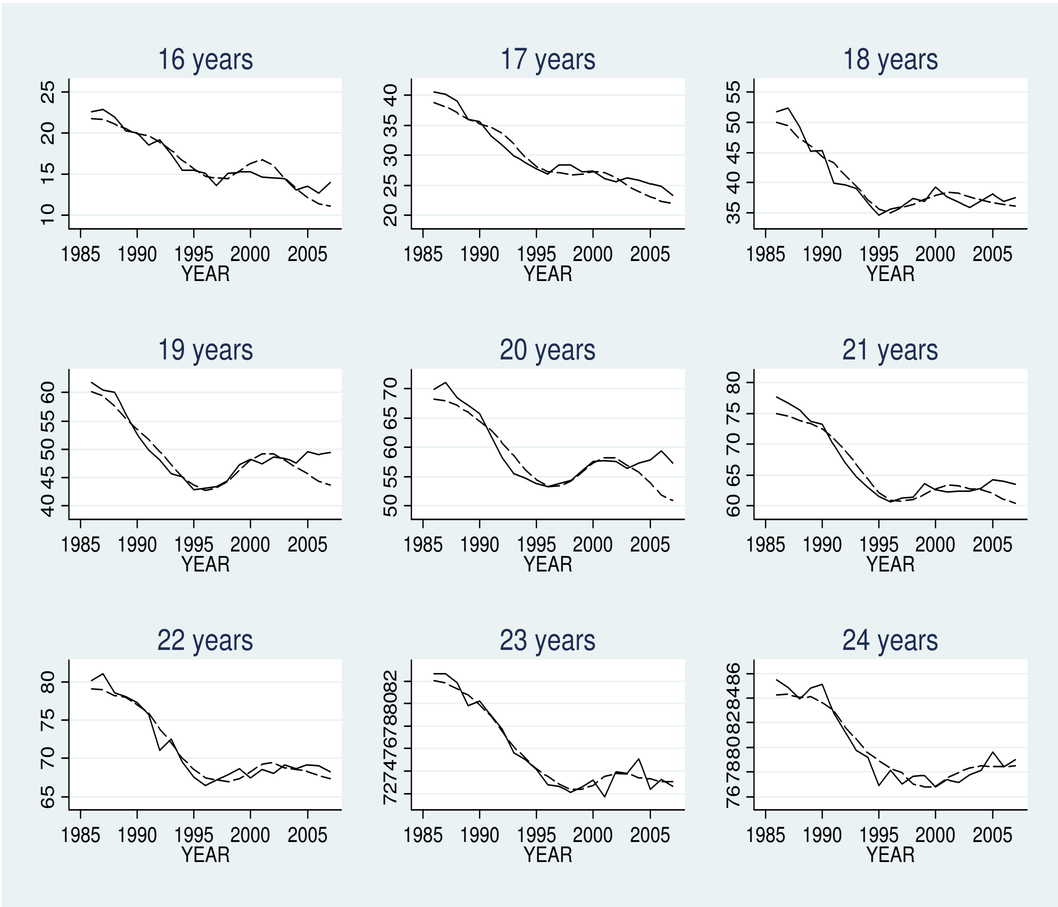


Notes: Trend is a three-year moving average of the estimated trend participation rate.

Figure 3.16: Trend and participation rates in the EA5: prime-age males

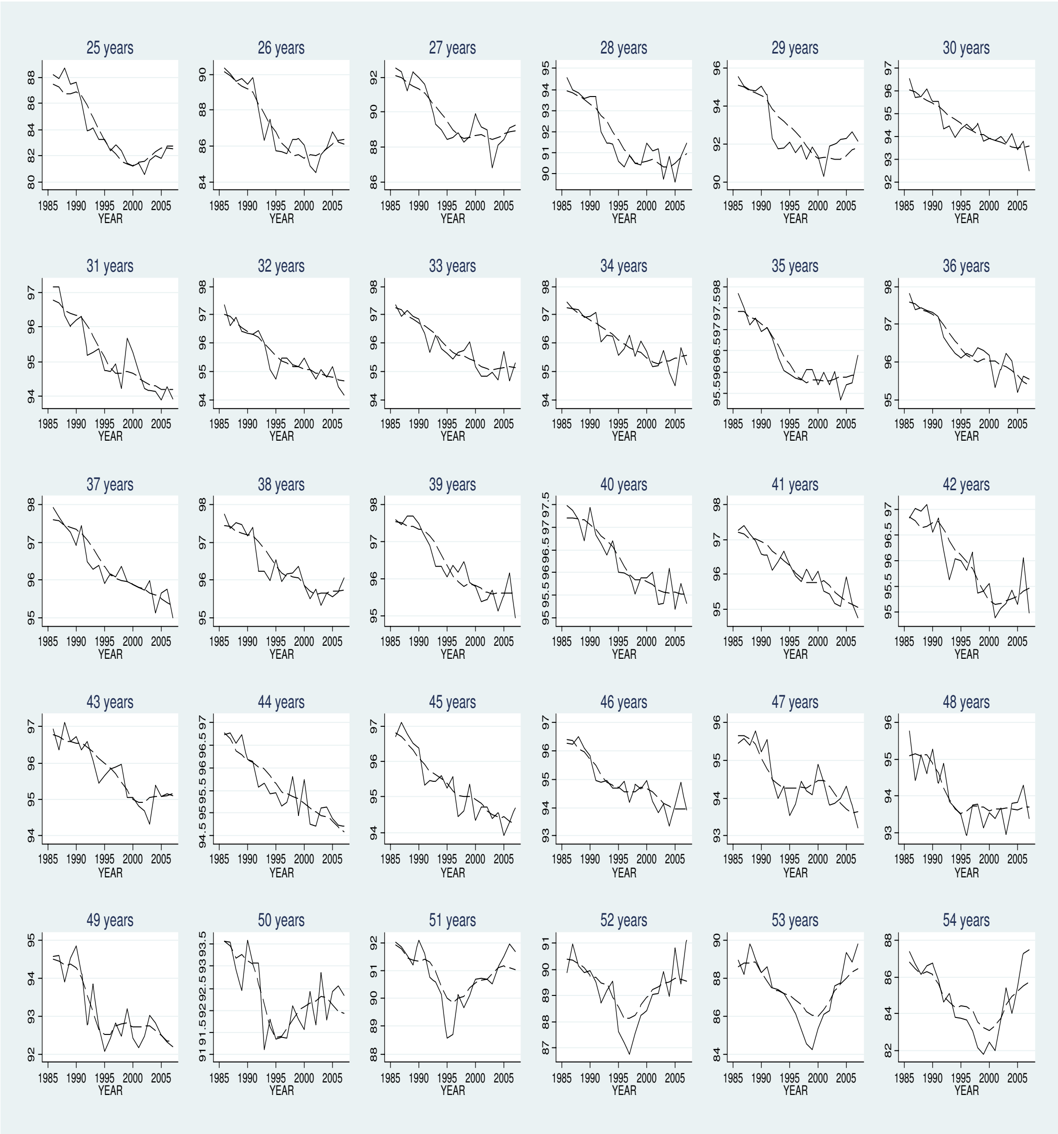


Notes: Trend is a three-year moving average of the estimated trend participation rate.

Figure 3.17: Trend and participation rates in the EA5: older males

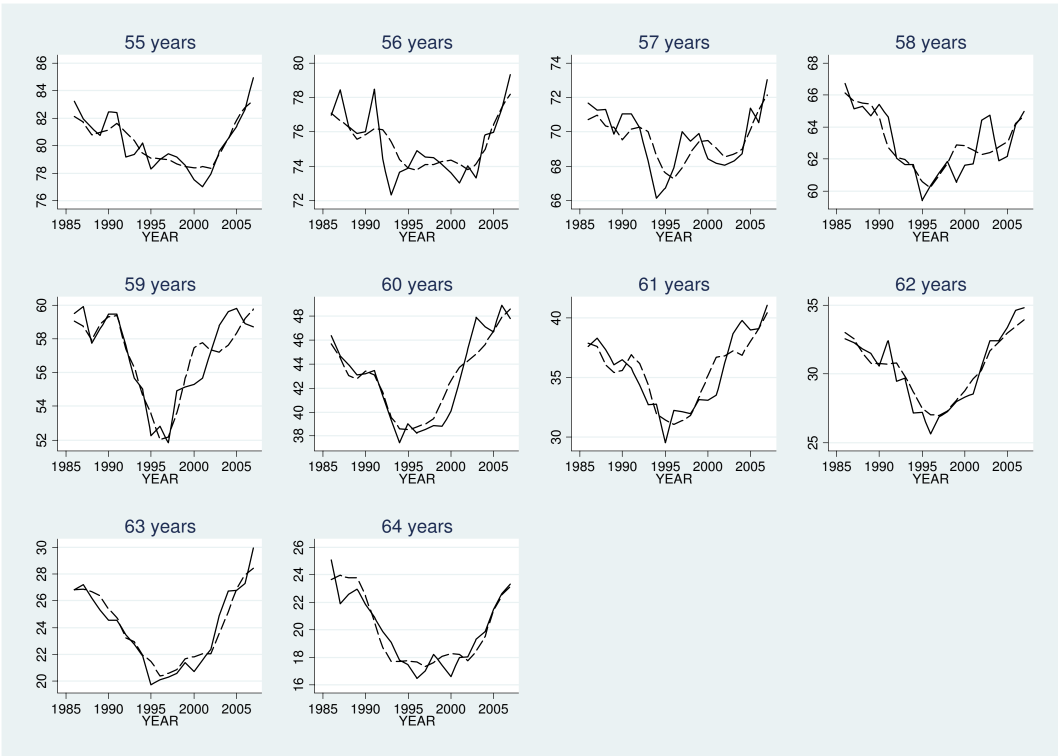


Notes: Trend is a three-year moving average of the estimated trend participation rate.

# Conclusion

This dissertation has investigated different aspects of business-cycle and long-run movements of aggregate labor supply and demand. In particular, the three chapters have addressed the implications of unobserved shocks for the evolution of these two variables and the resulting equilibrium in the labor market. While Chapter 1 has considered the importance of technology versus non-technology shocks for the cyclical fluctuations in labor demand, Chapter 2 has investigated the role of skill-biased technological change for these fluctuations. In addition, Chapter 2 has highlighted the role of these shocks as well as of shocks to the supply of skill (which may be interpreted as preference shifts in favor of higher education) for both the business-cycle and economic growth. Chapter 3 in turn has disentangled business-cycle from cohort effects (which may be interpreted as preference shifts in favor of participation in the labor market) as determinants for the overall increase in aggregate labor supply.

Contributing to the recent ongoing debate on the empirical performance of the Mortensen-Pissarides search-and-matching model, Chapter 1 has judged the empirical performance of the model on basis of moments conditional on technology shocks and non-technology shocks rather than on overall unconditional moments. These shocks were identified within a SVAR framework with conventional long-run restrictions and a combination of long-run zero and sign restrictions. Chapter 1 has documented that technology shocks cannot be the source of the high volatility in the job finding rate and unemployment present in the data. Opposite to the figures in the overall sample, technology shocks induce a negative co-movement between job finding and productivity and a positive co-movement between unemployment and productivity. Instead, additional non-technological disturbances are needed in order to replicate the unconditional volatility and correlations. Chapter 1 has documented that preference shocks which have been suggested to be important in the literature can only partly account for this.

Chapter 2 has investigated the implications of skill-biased technological change for the business cycle. Based on a new quarterly series of the skill premium and skill supply from the CPS outgoing rotation groups, this chapter has assessed the effect of technology shocks identified from a SVAR with long-run restrictions on the skill premium. Here, skill supply shocks have been controlled for using a short-run restriction. In addition, the chapter has proposed a strategy to identify skill-biased technology shocks through simple reshuffling of the variables in the same VAR that was used for the identification of the conventional technology shocks. Chapter 2 has documented that technology shocks are skill-biased at all business cycle frequencies. Further no evidence was provided for complementarity between capital and skill over the business cycle as investment-specific technology shocks do not significantly drive up the skill premium. Rather, capital and skill are substitutes in production.

Even though Chapter 2 abstracts from search-and-matching on the labor market, it contains valuable results that may be used in order to enhance the empirical performance of the models addresses in Chapter 1 in future research. More precisely, Chapter 2 has provided empirical foundation for a (cyclical, i.e. short- to medium-run) production function in which low- and high-skilled labor as well as high-skilled labor and capital are substitutes in production. In this framework, skill-biased technology shocks lead to a fall in hours worked and may thus be suitable to explain the "hours puzzle". Intuitively, as these technology shocks make high skilled workers more productive, low skilled workers are substituted out of production and overall hours fall. Further, skill-biased technology shocks mirror the dynamics of investment-unspecific technology shocks that have been identified in Chapter 1 in order to assess shortcomings of the conventional Fisher identification for investment-specific technology shocks. As a consequence, skill-biased technology shocks are important to understand the overall business cycle fluctuations that we observe in the data.

Based on the evidence in Chapters 1 and 2, it is potentially worthwhile to consider business-cycle labor market dynamics within a New Keynesian model rather than a RBC framework. New Keynesian models may well incorporate the dynamics driven by both non-technology and technology shocks that are documented in Chapter 1. These models usually investigate the effect of labor market frictions on optimal monetary policy. But it can be interesting to turn this question around: How do real and nominal rigidities in prices and wages influence the dynamics on the labor market? The answer to this question is not clear a priori. To give an example, in the presence of rigid nominal wages and flexible nominal prices and given a certain wage bargaining setup, firms may strategically affect the real wage through the setting of the nominal price. Since this potentially has an important effect on the labor market dynamics different from those in a setup in which nominal wages are fully flexible, it is possibly worthwhile to investigate this issue further in future research. This assessment becomes even more interesting when allowing for more heterogeneity on the labor market, for example with respect to skill.

Chapter 3 also highlights the role of heterogeneity in order to understand the dynamics on the labor market, in particular those of labor supply. Chapter 3 has used a cohort based model of labor force participation in order to analyze determinants of participation for disaggregated groups of workers in European countries, with a focus on the euro area. The model identifies significant age and cohort effects for detailed worker groups as indicators of (unobserved) structural determinants and disentangles these effects from other structural and business cycle factors on labor force participation rates. Indeed, it has been documented that age and cohort effects can explain a substantial part of the recent increase in labor force participation rates in the euro area, although not the surge since early 1990s. Cohort effects are particularly relevant for women, with those born in the 1920s and 1930s less likely and those born in the late 1960s and early 1970s more likely to participate in the labor market over the life-cycle. Looking forward, while they continue to provide some upward support to participation rates of women, positive cohort effects are not large enough to compensate for the downward impact of population ageing on labor force participation rates in the euro area.

Against this background, participation decisions of women may also have an important impact on the labor market dynamics over the business cycle and are worthwhile to be considered further in future research. In fact, two partners of a household interact with respect to their choices about market work, home production and leisure. This underlying heterogeneity plays a role for the individual allocation of time and hence the elasticity of labor supply. Allowing for this kind of heterogeneity in the aggregate potentially allows new insights into the determinants of aggregate labor supply and may have important implications for labor market or fiscal policy.

# Appendix

# Appendix A

# Identification and estimation in Chapters 1 and 2

## A.1 Standard long-run identification

Chapter 1 and 2 both employ structural identification in a VAR with long-run restrictions in order to estimate different types of technology shocks and disentangle them from non-technology shocks. 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 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 $\Phi_i$ are the impulse-response coefficients.

In the case of the Galí identification, all identified shocks, i.e. the neutral technology shock plus the remaining $n-1$ non-technology shocks, are assumed to be orthogonal. In addition, the variance of the structural residuals is normalized such that $\Sigma_e = I$. If labor productivity is ordered first in the VAR, a lower triangular structure of the matrix $C$ satisfies Galí's assumption that only neutral technology shocks drive labor productivity in the long run. This is easily obtained by decomposing the variance of the $k$-step ahead forecast error $\eta_{t,k} = X_{t+k} - E_t(X_{t+k})$ which is equal to

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

with the Cholesky decomposition[1]. In the application, $k = \infty$ has to be approximated by some large value, here $k$ is 80 quarters. It has to be noted that this procedure uniquely pins down the effect of the neutral technology shock on all variables in the VAR and that the result is not affected by the additional (unnecessary) zero restrictions in the matrix of long-run effects.

The reduced form VAR for all baseline specifications is estimated in a Bayesian framework in the main application. More precisely, I 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.[2] The point estimates exhibited then correspond to the median and the confidence intervals to the 16th and 84th percentiles of the posterior distribution (this is equivalent to one standard error).

## A.2 Estimation of the BVAR

All baseline results in chapter 1 and 2 are based on the reduced form VAR that is 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 $\phi_1$, $\phi_2$ and $\phi_3$, that determine the tightness and decay on own lags, other lags and exogenous 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 in chapter 1 is $d = 7$, in chapter 2 it is $d = 3$. 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

[1]See for example Uhlig (2004). Note that the variables important for identification, here labor productivity, need to enter in first differences in the VAR for this equation to hold.

[2]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 $A$.

posterior that can be derived analytically. Hence, the estimation does not rely on sampling procedures.

## A.3 Restricted Fisher identification

In the restricted Fisher identification in chapter 1, I implement the identifying assumption for neutral and investment-specific technology shocks as in Fisher (2006). Here, the restricted Fisher identification is very similar to the Galí identification, apart from a few issues.[3] First, if the the real investment price is ordered first and labor productivity second in the VAR, the matrix of long-run effects may be lower triangular in order to impose the first two restrictions, namely that only investment-specific technology shocks affect the investment price in the long run and only technology shocks may be sources of long-run fluctuations in labor productivity. In addition, the third restriction implies that $\frac{c_{21}}{c_{11}} = \frac{\alpha}{1-\alpha}$, where $c_{ii}$ are the respective elements of the matrix of long-run effects $C$. Since the lower triangular structure already imposes the number of conditions necessary for the identification of $A$, I need to relax one of the other assumptions in order to maintain exact identification. Here, the third restriction results in a positive correlation between neutral and investment-specific technology shocks. Hence, $\Sigma_e$ is no longer diagonal, but rather

$$\Sigma_e = \begin{bmatrix} 1 & \rho & O \\ \rho & 1 & O \\ O & O & I \end{bmatrix}.$$

Naming $\Lambda$ the lower triangular Cholesky factor from the decomposition of the $k$-step ahead forecast error, the identification matrix is then $A = FB$ with $F = (\sum_{i=0}^{k} \Phi_i)^{-1}$ and

$$B = \begin{bmatrix} 1 & 0 & O \\ b & \sqrt(1+b^2) & O \\ O & O & I \end{bmatrix}.$$

With $b = \frac{\frac{\alpha}{1-\alpha}\lambda_{11} - \lambda_{21}}{\lambda_{22}}$, with $\lambda_{ii}$ being the elements of $\Lambda$, the correlation between the two technology shocks is pinned down as $\rho = \frac{-b}{\sqrt(1+b)^2}$.

[3] Note that Fisher imposes his restrictions in an instrumental variable framework similar to Shapiro and Watson (1988). I thank Fabio Canova for the solution of the implementation of the Fisher restrictions as explained above.

## A.4 Alternative identification

The alternative identification in chapter 1 combines long-run zero restrictions as in the Fisher identification with sign restrictions on the long-run effects of the structural shocks on the real investment price and labor productivity. Assuming $\Sigma_e = I$, the $n$ elements of the matrix $A$ that maps the reduced form into structural residuals have to be determined such that $AA' = \Omega$ and our long-run restrictions are fulfilled. Note that this is equivalent to finding a decomposition $L$ of the long-run forecast revision variance such that

$$LL' = \Sigma^{\infty} = (\sum_{i=0}^{\infty} \Phi_i)\Omega(\sum_{i=0}^{\infty} \Phi_i)'.$$

Consider the same order of variables as in the Fisher identification, i.e. the real price of investment and labor productivity are ordered first in the VAR. First, I assume that only the two types of productivity shocks can affect the real investment price 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

$$L_{1:2,1:2}L'_{1:2,1:2} = \Sigma^{\infty}_{1:2,1:2}.$$

Next, I implement sign restrictions on this upper left 2-by-2 system in a similar fashion as in Peersman (2005). This involves a rotation of $L_{1:2,1:2}$ using an orthonormal matrix $Q$ (i.e. $QQ' = I$):

$$Q = \left[\begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}\right].$$

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, I calculate the long-run forecast revision variance. I then randomly draw $\theta$ from a uniform distribution $[0, \pi]$, use $Q$ to calculate the upper left elements of the matrix $L$ and check whether our sign restrictions are satisfied. In the application, I draw 100 candidates from the posterior distribution of the reduced from coefficients and another 100 values of $\theta$ for the rotation. I compute the impulse responses for all draws that satisfy the sign restriction and report the median and the 16th and 84th percentile from the resulting distribution. On average over a third of the draws satisfies the sign restrictions.

After having implemented the restrictions, I 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, I use that

$L_{3:n,1:2}L'_{1:2,1:2} = (\Sigma^{\infty}_{1:2,3:n})'$. Now I still need to determine the lower right elements of $L$. Note that these elements do not impose any of the restrictions nor are they related to the shocks of interest. I use the information on the first two rows and columns in order to adjust the lower right elements of the long-run variance. This 'remaining' block of the variance is then decomposed using the Cholesky decomposition. Having found all elements of $L$, I can then determine the matrix $A$ via $A = (\sum_{i=0}^{\infty} \Phi_i)^{-1} L$.

## A.5 VAR identification with short- and long-run restrictions

This section describes the method to implement the combination of short-run and long-run restrictions in chapter 2. To implement the short-run restriction which identifies skill supply shocks together with the long-run restrictions for the various technology shocks, we seek to find a unique transformation matrix $A$ that maps the reduced form residuals $v_t$ into structural shocks $e_t$. Assuming orthogonality between the structural residuals and normalizing their variance to unity, $A$ therefore satisfies $AA' = \Omega$ where $Omega$ is the variance matrix of the reduced from residuals. In a VAR with $n$ variables, another $n(n-1)/2$ restrictions are then necessary for exact identification and will come out of the short- and long-run assumptions.

Similar to before, we can formulate the problem in a triangular structure when the variables are conveniently ordered. This means ordering the supply of skill first in the VAR and then ordering the other variables according to the respective specification. The identification then works as follows. First, one identifies the supply shock through its short-run effect. More precisely, in order to identify supply shocks we assume that neither i-shocks, nor SBT-shocks nor neutral or non-technology shocks affect the supply of skill in the short run (on impact). This is equivalent to restricting $a_{12} = a_{13} = ... = a_{1n} = 0$ (with $a_{ij}$ being elements of $A$). These zero restrictions in the first row of $A$, combined with

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

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

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

# Bibliography

**Aaronson, Stephanie, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William Wascher**, "The Recent Decline in the Labour Force Participation Rate and Its Implications for Potential Labour Supply," *Brookings Papers of Economic Activity*, March 2006.

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

**Andolfatto, David**, "Business Cycles and Labor-Market Search," *The American Economic Review*, 1996, *86* (1), 112–132.

**Antecol, Heather**, "An examination of cross-country differences in the gender gap in labor force participation rates," *Labour Economics*, 2000, *7* (4), 409–426.

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

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

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

**Balleer, Almut**, "New Evidence, Old Puzzles: Technology Shocks and Labor Market Dynamics," Working Paper 1500, Kiel Institute for the Wold Economy March 2009.

**_ and Thijs van Rens**, "Cyclical Skill-Biased Technological Change," Working Paper, 2008.

**_ , Ramon Gomez-Salvador, and Jarkko Turunen**, "Labor Force Participation in the Euro Area: A Cohort-Based Analysis," ECB Working Paper 1049 May 2009.

**Barnichon, Regis**, "The Shimer Puzzle and the Identification of Productivity Shocks," Working Paper 2008.

**Bassanini, Andrea and Romain Duval**, "Employment Patterns in OECD Countries: Reassessing the Role of Policies and Institutions," Economics Department Working Papers 486, OECD 2006.

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

**Beaudry, Paul and Thomas Lemieux**, "Evolution of the Female Labour Force Participation Rate in Canada, 1976-1994," Working Papers 99-02, CIRANO 1999.

**Bencivenga, Valerie R.**, "An Econometric Study of Hours and Output Variation with Preference Shocks," *International Economic Review*, May 1992, *33* (2), 449–471.

**Bertola, Giuseppe, Francine D. Blau, and Lawrence M. Kahn**, "Labor Market Institutions and Demographic Employment Patterns," *Journal of Population Economics*, 2007, *20*, 833–867.

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

_ **and Jordi Galí**, "Labor Markets and Monetary Policy: A New Keynsian Model with Unemployment," *NBER Working Paper*, 2006, (13897).

_ **and Justin Wolfers**, "The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence," *Economic Journal*, 2000, *110* (462), C1–33.

**Blöndal, Sveinbjörn and Stefano Scarpetta**, "The Retirement Decision in OECD Countries," Economics Department Working Papers 202, OECD 1999.

**Braun, Helge, Reinout DeBock, and Riccardo DiCecio**, "Aggregate Shocks and Labor Market Fluctuations," Working Paper, Federal Reserve Bank of St. Louis 2006.

**Burniaux, Jean-Marc, Romain Duval, and Florence Jaumotte**, "Coping with Ageing: A Dynamic Approach to Quantify the Impact of Alternative Policy Options on Future Labour Supply in OECD Countries," Economics Department Working Papers 371, OECD 2004.

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

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

_ **,** _ **, and** _ , "Schumpeterian Technology Shocks," Working Paper, 2007.

**Carone, Guiseppe**, "Long-Term Labour Force Projections for the 25 EU Member States: A Set of Data for Assessing the Economic Impact of Ageing," Economic Papers 235, DG ECFIN 2005.

**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).

**Christoffel, Kai, Keith Kuester, and Tobias Linzert**, "Identifying the Role of Labor Markets for Monetary Policy in an Estimated DSGE Model," Working Paper Series 635, ECB June 2006.

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

**Cole, Harold L. and Richard Rogerson**, "Can the Mortensen-Pissarides Matching Model Match the Business Cycle Facts?," *International Economic Review*, 1999, *40* (4), 933–959.

**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.

**Davis, Steven J., John C. Haltiwanger, and Scott Schuh**, *Job Creation and Destruction*, Cambridge: MIT Press, 1998.

_ **, R. Jason Faberman, and John Haltiwanger**, "The Flow Aproach to Labor Markets: New Data Sources and Micro-Macro Links," *NBER Working Paper*, 2006, (12167).

**DeBock, Reinout**, "Investment-Specific Technology Shocks and Labor Market Frictions," Working Paper, 2006.

**denHaan, Wouter J., Garey Ramey, and Joel Watson**, "Job Destruction and the Propagation of Shocks," *The American Economic Review*, June 2000, *90* (3), 482–498.

**Duval, Romain**, *Pension Systems, Social Transfer Programmes and the Retirement Decision in OECD Countries* in Messina, J., C. Michelacci, J. Turunen G. and Zoega (eds), Labour Market Adjustments in Europe, Cheltanham: Edward Elgar, 2006.

**EuropeanCommission**, *The 2009 Ageing Report: Underlying Assumptions and Projection Methodologies* European Economy 7 2008.

**Euwals, Rob, Marike Knoef, and Daniel van Vuuren**, "The Trend in Female Labour Force Participation: What Can be Expected for the Future?," Discussion Paper 3225, IZA 2007.

**Faberman, Jason R.**, "Job Flows and the Recent Business Cycle: Not All 'Recoveries' Are Created Equal," Working Paper 391, Bureau of Labor Statistics 2006.

**Fallick, Bruce and Jonathan Pingle**, "A Cohort Based Model for Labour Force Participation," Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System 2007.

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

**Fernald, John G.**, "Trend breaks, long-run restrictions, and contractionary technology improvements," *Journal of Monetary Economics*, 2007, *54*, 2467–2485.

**Fernandez, Raquel**, "Culture as Learning: The Evolution of Female Labor Force Participation over a Century," Working Paper, 2007.

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

**Fitzenberger, Bernd, Reinhold Schnabel, and Gaby Wunderlich**, "The Gender Gap in Labor Market Participation and Employment: A Cohort Analysis for West Germany," *Journal of Population Economics*, 2004, *17*, 83–116.

**Fujita, Shigeru**, "Dynamics of Worker Flows and Vacancies: Evidence from the Sign Restriction Approach," *Journal of Applied Econometrics*, forthcoming 2009.

_ **and Garey Ramey**, "The Cyclicality of Job Loss and Hiring," *International Economic Review*, forthcoming.

**Fukuda, Kosei**, "A Cohort Analysis of Female Labor Force Participation Rates in the U.S. and Japan," *Review of Economics of the Household*, 2006, *4*, 379–393.

**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.

**Genre, Veronique, Ramon Gomez-Salvador, and Ana Lamo**, *The Determinants of Labour Force Participation in the EU* in R. Gómez-Salvador, A. Lamo, B. Petrongolo, M. Ward, and E. Wasmer (eds.), Labour Supply and Incentives to Work in Europe, Cheltenham: Edward Elgar, 2005.

_ **,** _ **, and** _ , "European Women: Why Do(n't) They Work," *Applied Economics*, forthcoming 2008.

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

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

**Greene, William and Terry Seaks**, "The Restricted Least Squares Estimator: A Pedagogical Note," *Review of Economics and Statistics*, 1991, *73* (3), 563–567.

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

_ **,** _ **, and** _ , "The role of investment-specific technological change in the business cycle," *European Economic Review*, 2000, *44*, 91–115.

**Hagedorn, Marcus and Iourii Manovskii**, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, September 2008, *98* (4), 1692–1706.

**Hall, Robert E.**, "Macroeconomic Fluctuations and the Allocation of Time," *Journal of Labor Economics*, 1997, *15* (1), S223–S250.

— , "Employment Fluctuations with Employment Wage Stickiness," *American Economic Review*, 2005, *95* (1), 53–69.

**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.

**Jaumotte, Florence**, "Female Labour Force Participation: Past Trends and Main Determinants in OECD Countries," Economics Department Working Papers, OECD, 376 2003.

**Kadiyala, K. Rao and Sune Karlsson**, "Numerical Methods for Estimation and Inference in Bayesian VAR-Models," *Journal of Applied Econometrics*, 1997, *12* (2), 99–132.

**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.

**Krause, Michael and Thomas Lubik**, "The (ir)relevance of real wage rigidity in the New Keynesian model with search frictions," *Journal of Monetary Economics*, 2007, *54*, 706–727.

**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.

**Mason, William and Stephen Fienberg**, *Cohort Analysis in Social Research: Beyond the Identification Problem*, New York: Springer, 1985.

**Masuch, Klaus and ECB Task Force**, "Labour Supply and Employment in the Euro Area Countries: Developments and Challenges," Occasional Paper 87, ECB 2008.

**Merz, Monika**, "Search in the Labor Market and the Real Business Cycle," *Journal of Monetary Economics*, 1995, *36*, 269–300.

**Michelacci, Claudio and David Lopez-Salido**, "Technology Shocks and Job Flows," *Review of Economic Studies*, 2007, *74*, 1195–1227.

**Mortensen, Dale T. and Christopher A. Pissarides**, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 1994, *61* (3), 397–415.

**Mortensen, Dale T and Christopher A. Pissarides**, "Technological Progress, Job Creation, and Job Destruction," *Review of Economic Dynamics*, 1998, *1*, 733–753.

**Mortensen, Dale T. and Eva Nagypal**, "More on Unemployment and Vacancy Fluctuations," *Review of Economic Dynamics*, 2007, *10*, 327–347.

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

**Pissarides, Christopher A.**, *Equilibrium Unemployment Theory*, Vol. 2, Camebridge, MA: MIT Press, 2000.

**Prieto-Rodriguez, Juan and Cesar Rodriguez-Gutierrez**, "The Added Worker Effect in the Spanish Case," *Applied Economics*, 2000, *32*, 1917–1925.

**Ramey, Garey**, "Exogenous vs. Endogenous Separation," Working Paper, 2008.

**Ravn, Morten O. and Saverio Simonelli**, "Labor Market Dynamics and the Business Cycle: Structural Evidence for the United States," Working Paper, 2006.

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

**Shimer, Robert**, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 2005, *95* (1), 25–49.

— , "Reassessing the Ins and Outs of Unemployment," Working Paper, 2005.

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

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

**Uhlig, Harald**, "Do Technology Shocks Lead to a Fall in Total Hours Worked," *Journal of the European Economic Association*, 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.

— **and Morten Ravn**, "On Adjusting the HP-Filter for the Frequency of Observations," *Review of Economics and Statistics*, 2002, *84* (2), 371–376.

**Veracierto, Marcelo**, "On the Cyclical Behavior of Employment, Unemployment and Labor Force Participation," *Journal of Monetary Economics*, September 2008, *55* (6), 1143–1157.