Essays on labor market dynamics with worker heterogeneity

Carlo Pizzinelli
Keble College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Economics

Hilary Term 2018

In memory of Giulio (1988-2016):
a friend and a scholar.

Supervisors: Prof. Francesco Zanetti and Dr. Michalis Rousakis.

Word count: Approximately 65,000 words, taking p.10 as a typical page.

Statement of authorship: The work in this thesis is based on research carried out by me as a student at the University of Oxford between October 2014 and March 2018. No part of this thesis has been submitted for a degree at the University of Oxford or any other university. Chapter 1 of this work is single-authored and is fully my own work. Chapter 2 is co-authored with Dr. Bradley Speigner (Bank of England). This chapter developed from the research project I undertook in collaboration with Dr. Speigner during my internship at the Bank of England in the summer of 2016. The topic was proposed by Dr. Speigner with the aim of contributing to the broader research agenda of his analytical division. As the project engaged with a literature familiar to both of us, we jointly devised the empirical strategy and particularly relied on Dr. Speigner's previous work for the estimation of "matching efficiency" through a state space model. During the internship and in the following months, I designed, carried out, and wrote the majority of the work in all sections, regularly discussing the results with Dr. Speigner and relying on his guidance to keep a policy-relevant focus throughout the analysis. An earlier version of this work has been published as a Bank of England Staff Working Paper (No. 667) with the title "Matching Efficiency and Labour Market Heterogeneity in the United Kingdom." Chapter 3 is co-authored with Prof. Francesco Zanetti. We started this work during the first year of my DPhil, as Prof. Zanetti encouraged me to explore in greater depth the non-linearities in search and matching models that I had started investigating when writing my MPhil thesis. While we jointly developed the paper's research question, I undertook most of the technical work individually. Prof. Zanetti provided guidance and expertise in writing the work.

Statement of funding: I acknowledge the generous financial support of the Economics and Social Research Council (for the academic years 2014/2015 and 2015/2016) and of the David Walton Distinguished Doctoral Scholarship (for the academic year 2016/2017).

Disclaimer: The views expressed in Chapter 2 are solely those of the authors and should not in any way be interpreted as reflecting those of the Bank of England. The authors alone are responsible for any possible errors.

Contents

Abstracts

Housing, borrowing constraints, and labor supply over the life cycle

Leverage-based borrowing constraints are important determinants of labor supply and homeownership over the life cycle. In this paper, I develop a life cycle model of a two-worker household with female labor supply and housing, where leverage constraints are formulated as upper limits of the Loan-To-Value (LTV) and Loan-To-Income (LTI) ratios. I find that female employment exhibits opposite cross-sectional relations with these two measures of leverage. The driver of this result is the interaction between the constraints and the household's composition of earnings. Furthermore, the model has two key implications. First, credit policies enacted through changes in the LTV or LTI limit affect households' labor supply decisions differently, with the latter exhibiting a greater spillover onto aggregate employment. Second, leverage constraints restrict households' ability to buffer income fluctuations and generate large heterogeneity in females' labor supply response to income shocks. Finally, using micro-data from the British Household Panel Survey, I find evidence for the model's predictions on the relationship between leverage and the employment of households' secondary earners.

JEL Classification: D91, J22, R21.
Keywords: Life Cycle Models, Labor Supply, Housing Demand, Leverage.

Labor force composition and the aggregate matching function in the United Kingdom (with Bradley Speigner)

This paper investigates how compositional changes in the labor market affect the matching process between vacancies and job seekers in the UK. We augment a state space representation of the aggregate matching function with a measure of job seekers' "search intensity" recovered from micro-data on unemployment-to-employment transitions. The baseline results show a worsening of the overall frictions in the labor market, which we call "matching efficiency", defined as the linear coefficient in the matching function. We find that matching efficiency declined by 15 percent between 1995 and 2010 but recovered by about 5 percent in the following 6 years. Prior to 2008, compositional changes in the labor force that improved aggregate search intensity offset the decline in matching efficiency. Considering broader definitions of job seekers that include marginally attached workers and on-the-job searchers exacerbates the registered decline in matching efficiency. Furthermore, changes in "recruiting intensity" and the share of vacancies posted by different industries provide a potential explanation for the initial fall in matching efficiency but not for the decline that preceded the 2008 recession. Finally, we quantitatively analyze the role that heterogeneity in labor force composition and changes in matching efficiency play for the shape and location of the UK Beveridge Curve.

JEL Classification: E24, E32, J64, J82.
Keywords: Unemployment, Labor Heterogeneity, Matching Function, Beveridge Curve.

State dependence in labor market fluctuations: Evidence, theory, and policy implications (with Francesco Zanetti)

This paper documents a novel fact: the volatility of the unemployment rate and the job separation rate is larger in periods with low aggregate productivity. A Diamond-Mortensen-Pissarides model with endogenous job separation and on-the-job search replicates these empirical regularities well. Endogenous job separation embeds powerful state dependence: fluctuations in the separation rate are larger in periods of low aggregate productivity and in response to contractionary shocks. Similar dynamics are in turn acquired by the unemployment rate. We study the implications of this asymmetry for structural reforms of labor market institutions. State dependence implies that the effect of labor market reforms is different across phases of the business cycle. A permanent removal of layoff taxes is welfare-enhancing in the long run, but it involves distinct short-run costs depending on the initial state of the economy. The welfare gain of a tax removal implemented in a low-productivity state is 3.5 percent larger than the same reform enacted in a state with high aggregate productivity.

JEL Classification: E24, E32, J64.

Keywords: Search and Matching Models, State Dependence, Business Cycles.

Acknowledgements

I am extremely grateful to my supervisors, Francesco Zanetti and Michalis Rousakis, for their sound guidance and constant support throughout my doctoral research. I would also like to thank other professors who have given me valuable advice on specific chapters, in particular Árpád Ábrahám from the European University Institute and Jesús Fernández-Villaverde from the University of Pennsylvania. Within Oxford, my work has benefited from interactions with other professors and post-docs of the macroeconomics research group, including Guido Ascari, Paul Beaudry, Martin Ellison, Andrea Ferrero, Jean Flemming, Agnes Kovacs, Michael MacMahon, and John Muellbauer. I am also grateful to Charles Gottlieb, from the University of St. Gallen, and Howard Smith, my advisor at Keble College, for their mentorship and friendship. Special thanks go to Bradley Speginer for being an enthusiastic co-author and a great mentor during my internship at the Bank of England.

I would like to acknowledge the Economic and Social Research Council and the David Walton Scholarship for their generous financial support, which has allowed me to pursue my DPhil studies without any financial worry. I am grateful to Keble College for providing me with a place I could call home for more than five years, and to the European University Institute for their hospitality during my four-month academic visit in 2016. I should note the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out the computational work related to Chapter 1.

For their comments and good humour, I thank other DPhil students at Oxford: Andrea Bernini, Manuel Fernández Sierra, Kris Hammarbäck, Sam Jindani, Michael Koelle, Thomas McGregor, Karan Nagpal, Simone Pedemonte, Viviana Perego, Chris Roth, Federico Torracchi, and Belinda Tracey. On a personal and non-academic level, I would like to acknowledge all my friends in Oxford who have been a source of joy and companionship, including my housemates at 2FK (Libby, Pat, Shea, Tim, and Valerio) as well as Alejandro Betancourt, Edoardo Borgomeo, Eduardo Dos Santos, Suzanne Jones, Luis Porangaba, Stefan Saftescu, and the wonderful people of the Keble College MCR. On the "Italian front," I should also thank my old friends who have been there for me during my years abroad, in particular Michela Golia, Alessandro Laterza, Margherita Maspero, Lorenzo Riva, Filippo Salmoiraghi, and my godmother Giovanna Majno.

No words can describe my heartfelt gratitude towards my parents, my brother, my grandfather, and Laura for their enduring love and support throughout these formative years. Whether nearby or far away, they have been by my side during both the happy and the challenging patches of this journey. I could not have done it without them.

Finally, Although I had the privilege to meet many brilliant and kind people during these four years, my thoughts also go to those dear to me whom I have lost along the way: Giulio, Serafino, and Antonia.

Introduction

This thesis is comprised of three chapters which discuss topics related to labor market dynamics from a macroeconomic perspective. Although each chapter is self-standing in terms of research question and methodology, they are united by a common interest for the macroeconomic implications of worker heterogeneity. The chapters vary with respect to the time horizon over which they study aggregate dynamics, covering business cycle frequency, the economy's long run steady state, and households' life cycle. Furthermore, they develop the concept of heterogeneity across different dimensions: stages of the life cycle, households' income and wealth, observed worker characteristics, and worker-firm productivity levels. The overall purpose of this thesis is therefore to contribute to the study of labor markets and labor policies through a multi-faceted approach.

Chapter 1 studies how borrowing constraints originating from the housing market affect the labor supply decisions of households' secondary earners over the life cycle. This question addresses two topics that have been at the center of the policy debate in the UK following the 2008 recession: homeownership and employment. First, housing is the main item in households' balance sheet with respect to both assets (through houses' market value) and liabilities (in the form of mortgages and collateralized debt). As households' ability to access debt is assessed through their financial leverage, housing becomes a crucial determinant of borrowing constraints. The two main measures of financial leverage with respect to housing are the Loan-to-Value and Loan-to-Income ratios, which are applied by banks to assess their customers' financial position and ability to re-finance their mortgages. Moreover, after the 2008 financial crisis, LTV and LTI limits are often directly targeted by policymakers to address macroprudential stability or increase homeownership, possibly with contrasting effects. For instance, in 2013 the UK government launched the "Help-to-Buy" scheme, which allowed first-time homebuyers to access low-interest mortgages with a LTV limit of 0.95. Concurrently, to restrain the exposure of the financial system to the real estate sector, in 2014 the Bank of England published guidelines for lending institutions which capped the issuance of new mortgages with an LTI ratio above 4.5. Second, labor earnings are the main driver of idiosyncratic income risk for households, and labor supply is a key margin of adjustment in the presence of financing constraints. It is therefore important to study the connection of these two key aspects of households' decisions.

A review of the literature, as well as preliminary empirical analysis on UK data, convinced me that the most appropriate approach to study the connection between leverage limits and labor supply is (*i*) through the lens of a life cycle model and (*ii*) focusing on households' secondary earners. Young households often face high leverage levels, as they purchase their homes through mortgages with small downpayments. Consequently, young households are the most likely ones to be affected by leverage limits, both in their decision to purchase a house and in their ability to smooth consumption when owning one. Meanwhile, a vast body of research has focused on labor supply choices of females, which have traditionally been identified as more likely to adjust working hours in response to fertility and childcare costs. In particular, several empirical works establish a connection between female labor supply, on the one hand, and house prices and mortgage leverage, on the other (Fortin, 1995; Del Boca and Lusardi, 2003). Human capital accumulation, fertility choices, and childcare costs imply marked life cycle variation in households' returns and costs from supplying labor. Therefore, a household simultaneously makes choices regarding both leverage and labor supply taking into account expectations over the life cycle.

To answer the proposed question, I develop a life cycle model of a two-worker household, calibrated to the UK. I assume that the secondary earner, calibrated to match average female earnings and employment, can adjust hours worked at the extensive and intensive margins while facing age-varying labor supply costs. Households can access debt in order to finance housing and smooth consumption subject to borrowing constraints, which are formulated in terms of LTV and LTI limits. The model uncovers the important interaction between leverage choices and labor supply, which is crucially shaped by the presence of leverage limits and of secondary earners. LTV and LTI ratios are reflective of different motives within the household. Depending on the composition of male and female earnings, either limit may be the binding one for a given household. However, only the LTI constraint can be directly relaxed by the female's choice to work. Households with high male earnings can substitute away from costly female employment and still purchase housing by committing future income to the repayment of debt. These households will have high LTV ratios and low female employment. When male earnings are temporarily or permanently low, female labor acts not only as a means of consumption insurance, but also as a way to access more debt by relaxing the LTI constraint. Using data from the British Household Panel Survey, I verify that these opposite relations of employment of secondary earners with the different leverage ratios broadly hold empirically.

While these relations generally apply to all age groups, they are particularly relevant for young households. Through a set of counterfactual calibrations of the leverage limits, I show that the homeownership rate of young households is the most sensitive to tighter financial constraints. The key difference between the LTV and the LTI limit, however, is that a tightening of the latter entails a fall in aggregate female labor supply that persists throughout the life cycle, and therefore lasts longer than the fall in homeownership. These results are relevant for understanding the spillovers of macroprudential policies on the real economy. Restricting credit by tightening either measure of leverage has distinct effects for two reasons. First, the two limits interact differently with agents' decisions. Second, the households affected by either margin are not the same ones.

Overall, this work approaches the theme of heterogeneity in several dimensions. First, on average, households of different ages make different decisions because they are at different stages of the life cycle. Second, at a given age, households' decisions vary based on the earnings of the two earners, leverage, and net worth. Masked by aggregate statistics, heterogeneity over the life cycle is of high importance for understanding today's macroeconomy. The homeownership rate in the UK has fallen by less than 10 percentage points since the early 2000's. However, it has actually dropped by more than 20 points for households below the age of 35, while remaining almost constant for older ones.[1] The life cycle framework highlights the intrinsically dynamic nature of agents' decisions and shifts the analysis towards a long-run scope: borrowing constraints directly affect households when they are young, but they ultimately also determine their wealth and consumption choices in old ages. For instance, a household that is borrowing constrained at a young age will also have greater net

1 The UK Office of National Statistics provides some highly insightful graphics on homeownership at https://visual.ons.gov.uk/uk-perspectives-2016-housing-and-home-ownership-in-the-uk/.

wealth later on in life and will adjust its behavior accordingly. This mechanism implies that leverage-based borrowing constraints may have long-lasting effects on the macroeconomy.

Chapter 2, coauthored with Bradley Speigner, studies how the observable characteristics of the pool of job seekers affect the process of matching between workers and recruiting firms in the UK. To this end, we estimate an aggregate matching function *à la* Diamond-Mortensen-Pissarides (DMP) augmented by a term for job seeker composition. The objective of this work is to understand the main drivers of unemployment dynamics besides labor demand and labor supply. In particular, we ask whether composition matters for long-run developments or business cycle fluctuations, and whether there have been changes in the unobserved matching frictions in the last two decades.

Certain demographic characteristics are strongly correlated with an individual's likelihood to find a job while unemployed. For instance, young and high-skill job seekers spend on average less time in unemployment than old and low-skilled ones. Similarly, unemployment duration is linked to lower chances of re-employment. Assuming that these differences in relative unemployment exit probabilities remain constant, compositional changes in the pool of job seekers entail that the job finding probability of the "average" searcher, as represented by the matching function, also changes. Furthermore, unemployed individuals are not the only job seekers in the economy. A significant fraction of employed workers search for a new job. Concurrently, some individuals out of the labor force still consider themselves willing an able to start a job although they are not actively searching for one. At any point in time, on-the-job search (OJS) and "marginally attached" (MA) workers constitute a substantial fraction of all flows into employment.

Thanks to its simplicity and adaptability, the DMP matching function has become the workhorse model to capture in reduced form the frictions operating in the labor market. We adapt and expand the framework proposed by Barnichon and Figura (2015) for the US, which relies on a two-step procedure, to account for the composition of job seekers in the matching function. In the first stage, the relative hazard rates of different worker types are estimated from micro-data on individual transitions into employment from the Labour Force Survey. The estimated coefficients and the labor force data are then combined to create an aggregate time series of "search intensity," which reflects the changes in the aggregate job finding rates accounted for by the composition of the searcher pool. This series is then included in the estimation of the aggregate matching function through a state space model. This macro-level step allows us to recover both the parameters of the function and the estimated path of the "matching efficiency", which captures the inherent frictions of the matching process.

We find that shifts in demographic characteristics (age and education) account for a secular increase in search intensity over the long run, while unemployment duration underpins its cyclical fluctuations. Finally, being driven by individual choices, the levels of OJS and MA job seekers also exhibit cyclical fluctuations which affect search intensity. In both the baseline estimation and the extended models, the estimated path of matching efficiency evinces a marked downward path from the late 1990's until the 2008 recession. Over the first part of this period, however, this worsening in matching efficiency was offset by improvements in search intensity. Moreover, using the composition-adjusted matching function we show how labor force composition crucially affects the shape and the position of the Beveridge Curve, which traces the steady-state relationship between unemployment and vacancies.

Ultimately, this project directly assesses the quantitative importance of (observable) worker heterogeneity for the dynamics of the unemployment rate. The framework disentangles long-run demographic drivers from the sources of cyclical fluctuations in search intensity. With regards to the estimated path of matching efficiency, the main result is that its decline started about a decade before the Great Recession. This finding is in line with several works focusing on the US, which detect a worsening in efficiency preceding the 2008 downturn (Sedláček, 2016; Hall and Schulhofer-Wohl, 2018). Being the residual term in the matching function, efficiency really represents the inherent functioning of the labor market that cannot be explained by labor demand, supply, and composition. Hence, there remains ample scope for future work to more deeply investigate its determinants.

Chapter 3, coauthored with Francesco Zanetti, starts from a stylized empirical fact regarding cyclical labor market dynamics in the US: the job destruction rate and the unemployment rate exhibit a larger volatility in periods of low aggregate productivity compared to periods of high productivity. We rationalize this finding through a DMP search model with endogenous job destruction and on-the-job search.

Within the model, the driver of the state-dependent volatility is the interaction between the distribution of workers' individual productivity levels and firms' threshold for efficient job destruction. When a job seeker and a vacancy are matched, they receive an idiosyncratic productivity value that augments the exogenous aggregate productivity level. In order for the employment contract to be efficient, the match-specific productivity must be sufficiently high. Hence, firms impose a reservation level below which matches do not turn into employment. If, over time, the match-specific productivity falls below this threshold, the worker is laid off. Moreover, the separation threshold fluctuates with aggregate productivity. In times of low aggregate productivity, the minimum level of individual productivity required for a match to be efficient is higher. Under conventional assumptions regarding the distribution of match-specific productivity, a higher level of the threshold implies that its subsequent movements lead to larger fluctuations in the separation rate.

The generalized Impulse Response Functions (IRFs) of the model show that the job destruction rate responds more pronouncedly to productivity innovations in periods of low aggregate productivity. Furthermore, job separations also respond more strongly to negative than to positive productivity shocks. The unemployment rate acquires the same asymmetries of the separation rate, and therefore exhibits a larger volatility in periods of low productivity.

Given the importance of the job separation rate for these cyclical asymmetries, we use the model to study the effect of a labor market policy that directly interferes with firms' job destruction choices. We thus introduce labor protection in the model through a "wasteful" tax on layoffs. We find that the overall effect of employment protection on the long-run unemployment rate crucially depends on the endogenous response of the workers who decide to search on the job. In the baseline model, OJS is undertaken by low-productivity workers, who are close to the separation threshold and try to find a better match. Since this group is comprised of workers who would be laid off in case of a negative aggregate shock, OJS becomes a way to avoid separation. Employment protection depresses labor demand, as firms internalize the future cost of the tax and lower their vacancy postings. Facing lower chances of finding a new job, fewer workers undertake OJS and eventually more workers are vulnerable to being separated.

Finally, we consider the transition dynamics of the unemployment rate ensuing from a permanent removal of the layoff tax. To further highlight the relevance of state-dependence in the labor market, we compare the case in which the tax is removed during a period of low aggregate productivity with the case in which the reform happens at a time of high productivity. While the reform is beneficial in the long run, the short-run costs, due to the higher matching frictions, are larger in periods of high productivity.

This chapter deals with the heterogeneity from the viewpoint of the distribution of idiosyncratic productivity of different workers, which drives the cyclical dynamics of unemployment and explains its state-dependent volatility. In studying the impact of the layoff tax, heterogeneity is once again relevant because it underpins the fraction of workers who decide to search on the job.

The following three chapters contain the essays outlined above. These works are presented as self-contained journal-style articles. Each chapter includes a discussion of its individual contribution to the relevant academic literature.

1 Housing, borrowing constraints, and labor supply over the life cycle [1]

1.1 Introduction

Households' ability to mitigate unexpected shocks and smooth consumption over the life cycle is crucially affected by the composition of their balance sheets (Mian et al., 2013; Kaplan and Violante, 2014). To ensure aggregate stability, policymakers therefore pay close attention to the mortgage market and constrain households' access to debt by setting limits to their financial leverage. Most households face high levels of leverage at young ages, when they purchase houses through mortgages, and extinguish their debts over their working lives. In this process, labor supply provides a key margin of adjustment to finance house buying, repay mortgage debt, and smooth consumption. In fact, several empirical studies establish that households adjust the labor supply of secondary earners in response to mortgage market reforms (Del Boca and Lusardi, 2003) and in reaction to changes in their ability to repay housing debt (Fortin, 1995). The aim of this paper is to develop a life cycle model that isolates the key channel for the interaction between leverage-based borrowing constraints and labor supply decisions, drawing important conclusions for credit policies.

I build a life cycle model *à la* Attanasio et al. (2012) with uninsurable income risk and housing preferences, calibrated to the United Kingdom (UK), which I extend to a two-worker household with female labor supply.[2] I assume that females function as secondary earners and face empirically relevant, age-varying labor supply costs. Moreover, exogenous borrowing constraints are formulated as upper limits of two measures of mortgage leverage: the Loan-to-Value (LTV) and the Loan-to-Income

1 I would like to thank Árpád Ábrahám, Charles Gottlieb, Jesús Fernández-Villaverde, Christopher Roth, Michalis Rousakis, and Francesco Zanetti, as well as seminar participants at the European University Institute, the University of St. Gallen, the University of Oxford, and the Workshop on Dynamic Macroeconomics (Vigo, Spain) for useful comments. I would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work (http://dx.doi.org/10.5281/zenodo.22558).

2 Similar to previous studies, I model two-earner households as composed of one male and one female. Henceforth, the term "couple" is employed assuming a household where the two earners are of opposite sex.

(LTI) ratios. The LTV represents the ratio of a household's outstanding mortgage to the value of housing assets, while the LTI is the ratio of the outstanding mortgage to yearly income. These ratios are the most common measures of housing leverage and are used by banks to set limits on their customers' ability to borrow. This feature of the model captures the institutional framework of the UK as well as other developed economies.

The structure of mortgage markets, together with long-standing empirical evidence on intra-household risk-sharing, make the life cycle dimension the most informative lens to study the relationship between housing debt and female labor supply. Mortgage contracts can entail repayment periods of up to 35 years, hence affecting agents' choices throughout a significant segment of their lives. Meanwhile, a large literature shows that, within multi-earner households, female labor supply is the principal margin of adjustment to insure consumption against income shocks to the primary earner, a mechanism known as the "added worker effect" (Lundberg, 1985; Stephens, 2002).

I show that leverage-based borrowing constraints crucially affect the household's decisions on debt levels and labor supply. The composition of earnings determines whether either leverage limit (or neither) is the binding one for a given household, consequently determining choices with respect to debt and hours worked. This result has two important implications. First, the quantitative impact of changes in the credit constraints on homeownership depends on their interaction with households' labor supply choices and varies over the life cycle. Second, the response of female labor supply to unexpected income shocks exhibits large heterogeneity across different levels of leverage.

In the model, employment of the secondary earner is negatively correlated with the LTV ratio and positively correlated with the LTI computed using the primary earner's income - which I denote as p-LTI. The earnings of the two workers, and their expected growth over the life cycle, underpin these cross-sectional relationships. In particular, the LTV- and LTI-based constraints interact with two opposing channels. On the one hand, when labor supply costs are high, a household has a motive to avoid female employment and instead finance consumption and housing through borrowing. On the other hand, if male earnings are low, female labor supply can provide a key source of income and relax the LTI-based debt constraint. Intuitively, young households anticipate growing income and falling labor supply costs later in life. When purchasing a house, and in the following years, they have an incentive to avoid female work by accumulating debt, which increases their LTV ratio. However, even for low levels of debt, if the male's earnings are sufficiently low, the LTI limit becomes binding first unless the female also works. Hence, it is only households with high male earnings that can obtain a high LTV ratio. Among these, only those with low female potential earnings choose to accumulate debt instead of supplying costly labor. Meanwhile, a high p-LTI reflects the need for female labor to sustain debt repayments and relax a credit constraint when the male's income is low, leading to high employment.

Leverage-based borrowing constraints are crucial to establish a direct link from the credit market to households' decisions. Macroeconomic policies often directly target LTV and LTI limits to address homeownership and financial stability objectives.[3] However, credit policies enacted through LTV and LTI limits crucially differ with respect to their spillover onto the labor market. First, the two limits interact differently with households' labor decisions, as only the LTI limit directly affects the returns from work. Second, the households who are "at the margin" of either leverage constraint are not the same.

To evaluate these predictions and their policy implications, I consider counterfactual scenarios with alternative values of the LTI and the LTV limit, respectively. I show that comparable restrictions of the LTI and LTV limits have different impacts on aggregate labor supply. A tightening of the LTI limit, enacted to reduce the homeownership rate of the youngest cohort by 10 percentage points, also implies a fall in aggregate female employment between 1 and 2 percentage points which persists through the life cycle. The decrease in labor supply originates from those households who rely on female labor supply to access the debt needed for homebuying. A comparable contraction of the LTV limit does not directly interfere with the returns to labor supply and does not lead to a fall in employment at the aggregate level. These results also imply large heterogeneity in the welfare losses suffered by households because of the credit restrictions. However, both limits are particularly relevant for young households, which are the most likely to be affected by the binding constraints.

A critical policy question is whether the response of female labor supply to per-

3 For instance, in 2013 the UK government launched the "Help-to-Buy" policy to allow new homebuyers to obtain subsidized mortgages with a maximum LTV of 0.95. Meanwhile, in 2014, following the advice of the Bank of England, the Financial Conduct Authority issued rules on the maximum share of high-LTI mortgages that banks should issue (FCA, 2014).

manent income shocks varies with households' leverage levels. I find that high-LTV households greatly adjust their labor supply after income shocks, compared to low-LTV ones. The former group includes those households with low female wages and high male earnings, who can afford to borrow up to the LTV limit. They thus exhibit the largest fall in employment after a negative shock to female wages as they can easily substitute away from costly female work. However, they also experience a large rise in employment after a drop in male earnings because the fall in primary income impairs their main channel for financing consumption.

To assess empirically the predictions of the model, I use data from the 2001-2006 British Household Panel Survey (BHPS). I follow the survey's definition of "head of household" to identify couples' secondary earners, which in 90 percent of the cases are female. In line with the model, I show through a linear probability regression that the probability of the secondary earner being employed is negatively related to the household's LTV ratio and positively related to the p-LTI. The results are robust to alternative sample selections and to an extended time interval. Furthermore, I find that falls in primary earnings account for a large fraction of cases of high p-LTI, indicating circumstances in which secondary labor acts as a means of insurance.

This paper is related to the large literature that uses life cycle models to separately study housing demand or female labor supply. I contribute to this literature by developing a model that shows the strong interaction between the two channels, which is relevant for both fields. On housing demand, seminal works include Iacoviello (2008), Yang (2009), Fernández-Villaverde and Krueger (2011), Attanasio et al. (2012), Bajari et al. (2013), and Iacoviello and Pavan (2013). While these works consider leverage constraints, I extend their results by showing that the endogeneity of leverage choices to the labor supply of secondary earners entails key dynamics, which are not otherwise captured. Studies on two-earner households include Low (2005), Attanasio et al. (2005, 2008, 2015), Blundell et al. (2016), and Wu and Krueger (2016). These papers highlight the role played by female labor as a means of consumption insurance, which is particularly relevant in the presence of borrowing constraints. However, by abstracting from housing, they only consider constraints on net wealth. As I show, leverage-based debt limits imply that households need not be poor to be close to a borrowing constraint. Consequently, some households with high wealth may exhibit a large added worker effect in response to income shocks.[4]

To the best of my knowledge, the only work that combines female labor supply and housing demand within a life cycle framework is Bottazzi et al. (2007). While adopting a similar approach, my paper traces a direct link between life cycle expectations and leverage-based constraints to explain the empirical relationships between female labor supply and different measures of leverage. I further extend the analysis by assessing the relevance of the labor supply channel for the effect of changing leverage limits on homeownership.

The rest of the paper is structured as follows. Section 1.2 provides empirical evidence on the life cycle profiles of homeownership, employment, and leverage. Sections 1.3 and 1.4 outline the model and the calibration, respectively. Section 1.5 presents the main results and discusses the model's key dynamics. Section 1.6 evaluates the importance of the labor supply channel with respect to changes in leverage limits. Section 1.7 analyzes the implications for the response of female labor supply to income shocks. Using the BHPS, Section 1.8 finds empirical support for the model's main predictions. Section 1.9 concludes.

1.2 Housing, employment, and leverage in the BHPS

In this section, I provide stylized empirical evidence to motivate the joint analysis of homeownership, leverage, and labor supply. Using the British Household Panel Survey (BHPS), I show that these variables exhibit both marked life cycle patterns and large cross-sectional variation for two-earner households.

Mortgage repayments are a form of "committed consumption", implying that part of the household's budget cannot be adjusted in response to income fluctuations. Additionally, homeowners with an outstanding mortgage have lower housing equity and hence a reduced ability to further borrow against their housing stock to smooth consumption. The size of the mortgage relative to the flow of income is also indicative of how difficult it may be for the household to extinguish its debt over the years. Reflecting these concepts, two conventional measures of household leverage are the Loan-to-Value and the Loan-to-Income ratios. Higher values of these two metrics indicate a higher leverage. However, the measures are distinctively informative in the

4 Using the definition of Kaplan and Violante (2014), this study explores the labor supply dimension of "wealthy hand-to-mouth" households.

sense that they compare a household's outstanding debt relative to a stock of wealth and a flow of income.

The BHPS is comprised of yearly individual-level observations with a longitudinal dimension from 1991 to 2008. For the analysis, I use the survey waves for the years 2001 to 2006. I choose this time interval because, despite the sustained house price growth, it was a period of general macroeconomic stability in the UK, marked by a stable unemployment rate just above 5 percent and moderate real earnings growth.[5] I consider all households formed by couples, either married or in a cohabiting relationship, where the secondary earner is between 23 and 65 years old.[6] A large literature has showed that the secondary earner's extensive margin is the main margin of labor supply adjustment for most households.[7] I define as secondary earner the member of the couple who is not classified as the head of household by the BHPS in each wave.[8] Based on this selection criterion, in the 2001 wave of the BHPS roughly 90 percent of secondary earners in prime-age couples are female.[9]

Using this sample, the first row of panels in Figure 1 shows the life cycle profiles of housing assets and the employment rate of females and males within couples.[10] Keeping in mind that these profiles do not account for endogenous household formation and separation decisions, clear age trends are visible in all series. The homeownership rate starts just below 60 percent and rises until age 45, after which it stabilizes around 85 percent. As the average log value of the primary residence shows (Panel 3), older homeowners also tend to accumulate larger amounts of housing assets.[11] As shown in the second panel of Figure 1, female employment evinces a more pronounced life cycle profile than male employment. The percent of females in employment gradually rises from 70 percent at age 25 to its peak just above 80 percent in the late 40's, and

5 House prices began to fall in 2007 and the unemployment rate rose from 5 percent to 8 percent in 2008, while average earnings fell. The results of the following analysis, however, are robust to including the years 2007 and 2008 in the sample.

6 For comparability with the model presented below, I exclude same-sex couples in the baseline sample. However, sensitivity analysis shows that the empirical results hold equally when including them.

7 See for instance Lundberg (1985), Stephens (2002), and Mankart and Oikonomou (2017).

8 In the BHPS, the head of the household is identified as the member who is legally and financially responsible for accommodation or the elder of two people who share the responsibility.

9 Table 1.A.1 reports the sex of the household head and the secondary earners. The table also shows that the secondary earner is more likely to be involved in family care or other non-employment activities. Furthermore, for couples with children, the secondary earner is more likely to be solely responsible for childcare activities.

10 The employment panel plots the employment rate against the age of the member of the respective sex. The other panels use the age of the secondary earner.

11 For the entire analysis in this paper, variables in British pounds are deflated to 2001 levels using the UK Consumer Price Index, unless otherwise stated.

Figure 1: Life cycle profiles (3-year moving averages) for homeownership, employment, average (log) house value, percentage of mortgagors, average LTV, LTI, and p-LTI using the 2001-2006 waves of the BHPS.

Note. Source: BHPS waves 2001-2006. The sample includes married couples and those in permanent partnerships aged 23 to 65. All series are reported as 3-year centered moving averages. The solid line in Panel 6 reports the LTI based on the full household income, while the dashed line reports the average value using all household income except for the secondary worker's earnings. In Panel 3, house values are deflated to 2001 prices using the UKHPI series.

subsequently falls until age 65. Meanwhile, male employment is higher at all ages and almost constant until age 50, when transitions into retirement begin.

The lower row of Figure 1 focuses on debt and leverage. Panel 4 shows that the proportion of households with an outstanding mortgage has a hump-shaped profile over age, implying that almost all young homeowners have a mortgage, which they progressively repay throughout their working years.[12] The last two quadrants of Figure 1 focus on households' financial position in terms of LTV and LTI ratios for those with outstanding mortgages. Both measures of leverage show a clear downward path, indicating that young mortgagors have on average a larger portion of their house used as collateral and their outstanding mortgage is high relative to their flow of income. The last panel plots the LTI computed with the full household income

12 Since the BHPS does not specifically ask about available Home Equity Lines of Credit (HELOCs), this measure may underestimate the proportion of households with outstanding debt secured against their house.

and based only on primary labor earnings and other non-labor income (i.e. the p-LTI). The gap between the two series highlights the importance of labor income from the secondary earner to lower the overall ratio, which may be necessary to avoid a leverage limit.[13]

The 6-year period used for the analysis featured sustained growth in house prices, which may be the main driver of younger households' leverage decisions. As Figure 1.A.1 in the Appendix shows, rising prices in the decade preceding the 2008 recession imply that younger cohorts of homeowners had higher LTI ratios than older ones.[14] I therefore check whether the age profiles in the period of interest are not the result of spurious year or cohort effects. Figure 1.A.2 reports the age effects obtained from the regression method proposed by Deaton and Paxson (1993), which confirm the trends from Figure 1.

Besides differences across age groups, homebuyers make different choices on the leverage they take up at the time of purchase and sustain over the following years. Figure 2 shows the large variation in the LTV and (p-)LTI values chosen by new homeowners in the year of purchase. This variation originates from the value of the house, the size of the mortgage, and income, which are all chosen by households based on their future prospects. In the next section, I develop a life cycle framework to establish how expectations of earnings and labor supply costs determine choices of leverage and work.

1.3 Model

This section describes the theoretical model, while the details of the computational solution are contained in Appendix 1.B.

To assess the stylized facts from Section 1.2, I develop a model of a unitary household with two earners: a male and a female. The main features of the model follow the lines of Bottazzi et al. (2007) and Attanasio et al. (2012). Given that 90 percent of secondary earners in the BHPS are females, I assume for explanatory purposes that the female acts as secondary earner. The household faces idiosyncratic income shocks to both male earnings and female wages over a finite lifetime. Households accumulate

13 Interestingly, the rise in primary earner's LTI after age 55 may be underpinned by selection in the type of households who have an outstanding mortgage late in their working lives, and the small sample size for mortgagors at old ages.

14 To provide a clearer picture of the long-run differences across cohorts, this figure is produced using the full 1991-2008 BHPS sample.

Figure 2: Distribution of LTV, p-LTI, and LTI among new homeowners in the BHPS waves 2001-2006.

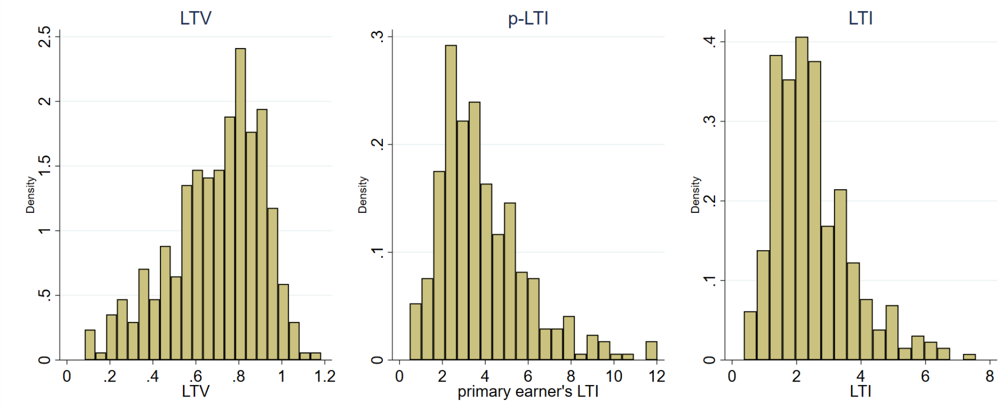

Note. Source: BHPS waves 2001-2006. The sample includes new homeowners where the secondary earner is aged 23 to 65.

assets to smooth consumption and purchase discrete units of housing. Those who do not own any housing must pay rent. The purchase of housing can be partly financed through a mortgage, hence holding negative liquid assets. Homeowners can also borrow against the value of their house, up to a limit, after the period of purchase. After a finite working life, agents retire and receive a fixed income stream. Death occurs at the end of the retirement period. The model is in partial equilibrium: house prices and interest rates are exogenous and deterministic.[15]

The life cycle lasts J_D periods of one year each. From period $j = 1$ to period J_{R-1} households are active in the labor force and receive an income y_j for the male and a wage w_j^f for the female, both of which exhibit stochastic fluctuations around a deterministic age trend. For the male, income shocks include the possibility of temporary involuntary unemployment. From age J_r to J_{R-1}, females face the possibility of retiring permanently. From J_R to J_D, both earners in the household are retired and receive pensions b and b^f, respectively. Upon reaching period J_D, households cease to live and must repay all their debt.[16]

Except for involuntary unemployment, the male always works. The female decides

15 The absence of these two features may be a limitation for an exhaustive analysis of housing, but neither of them is essential for the key mechanisms studied in this work. A following section of the paper discusses the effect of unanticipated price shocks.

16 I assume they derive no utility from leaving bequests. I also choose not to model household destruction through early death or divorce for simplicity. The possibility of both events has important repercussions for savings decisions and for intra-household risk sharing but are beyond the direct scope of this work. If empirically relevant, the interaction of these channels with the relationship studied in this paper would constitute an important direction for future research.

whether to work ($0 < n^f \leq 1$) or not ($n^f = 0$) based on her preference for leisure and on the additional age-varying cost $\xi(j)$ incurred if employed.

Housing is a discrete variable $h_j \in \{0, ..., \overline{h}\}$. When $h_j = 0$, the household is renting at a cost q. Housing units have a price p and their trading is subject to proportional transaction costs F_b and F_s for each unit bought and sold, respectively. I use a specification similar to that of Attanasio et al. (2012) in setting $\overline{h} = 4$, where $h = 1$ equals a small house and $h = 4$ represents the largest house size. It is worth noting that h is interpreted broadly as "housing services", which embodies all features that give greater value to a housing unit. It therefore represent not just house size but also location and any other factor increasing the price.

Net financial assets a_j are a continuous variable receiving a per-period interest r and can be traded at no cost. A negative value of a_j signifies an outstanding mortgage. Debt for the following period is constrained by a borrowing limit $\phi(h_j, y_j, w_j^f, n_j^f)$ or by the current outstanding debt. Details of this constraint are provided further below.

Household preferences are a function of consumption c_j, housing h_j, female labor n_j^f, and age. Future periods are discounted with a constant factor $0 < \beta < 1$. Given my interest in how the "lumpiness" of housing affects other household choices, I adopt a specification of preferences where such discreteness enters directly into the agents' utility. Households derive utility from consumption, housing, and leisure as follows:

$$U(c, n^f, h, j) = \frac{c^{1-\sigma}}{1-\sigma} + \theta_f \frac{(1-n^f)^{1-\psi}}{1-\psi} + \mathbb{I}\{h > 0\}\mu(h, j),$$

where $\mathbb{I}\{h > 0\}$ is an indicator function for homeownership and the preferences for housing are represented by the function $\mu(h, j)$, which has the following form:

$$\mu(h, j) = \chi(j) \left(\mu_h + \phi_h \frac{h-1}{\overline{h}-1} \right).$$

The first term $\chi(j)$ is a deterministic age-varying weight that implies a change in the preference for housing over the life cycle. The constant μ_h represents the preference for a small house, while the slope coefficient ϕ_h represents the marginal preference for larger houses.

For $j < J_R$, male earnings follow the process

$$\log y_j = \alpha_0 + \alpha_1 j + \alpha_2 j^2 + \log z_j \ ,$$

$$\log z_j = \rho_z \log z_{j-1} + \epsilon_j, \quad \epsilon_j \sim \mathrm{N}(-\frac{\sigma_\epsilon^2}{2}, \sigma_\epsilon^2).$$

The first component of earnings is a deterministic quadratic function of age, the second one is a stochastic process.[17] Furthermore, with some probability π_u, each period the male may be jobless and receive unemployment insurance y^u. Earnings become constant once the male retires, so that $y_j = b$, for $J_R \leq j \leq J_D$. Retirement income is assumed to be a constant proportion of the earnings received in the final period of work. Female wages follow a similarly structured process:

$$\log w_j^f = \alpha_0^f + \alpha_1^f j + \alpha_2^f j^2 + \log z_j^f \ ,$$
$$\log z_j^f = \rho_z^f \log z_{j-1}^f + \epsilon_j^f, \quad \epsilon_j^f \sim \mathrm{N}(-\frac{\sigma_{\epsilon^f}^2}{2}, \sigma_{\epsilon^f}^2).$$

Females face no involuntary unemployment risk but incur a per-period retirement probability π_r between J_r and $J_R - 1$, or retire with certainty at J_R if still active.[18] At retirement, they receive income b^f equal to a fraction of their final wage times the average hours worked by females in the economy.

The household's age-j value function V_j depends on whether the female is active or retired. Denoting $X_j = [a_j, h_{j-1}, y_j, w_j^f]$ as the vector of relevant states, the problem for a couple with an active female expressed in terms of the value $V_j^A(X_j)$ for all $j < J_r$ is

$$V_j^A(X_j) = \max_{c_j, h_j, a_{j+1}, n_j^f} \Big\{ U(c_j, h_j, n_j^f, j) + \beta \mathbb{E}_j V_{j+1}^A(X_{j+1}) \Big\} \tag{1}$$

subject to:

$$h_j \in \{0, ..., \overline{h}\}$$
$$a_{j+1} + c_j + p h_j + \Phi(h_j, h_{j-1}) + q\mathbb{I}\{h_j = 0\}$$
$$= (1+r)a_j + y_j + w_j^f n_j^f - \xi(j)\mathbb{I}\{n_j^f > 0\} + p h_{j-1}$$
$$a_{j+1} \geq \min\{a_j, \phi(h_j, y_j, w_j^f, n_j^f)\}$$
$$n_j^f \in [0,1],$$

where the indicator function $\mathbb{I}\{h_j = 0\}$ is equal to one when the household rents and $\mathbb{I}\{n_j^f > 0\}$ equals one when the female works. Expectations for the future are taken with respect to income and wages.

17 I approximate the stochastic component with a finite vector of states $z \in [\underline{z}, ..., \overline{z}]$ and a set of transition probabilities $\pi_{z_j|z_i}$.

18 Although a thorough analysis of retirement decisions is beyond the scope of this work, modeling heterogeneity in retirement timing and pension income, even if in a stochastic way, is important to match the life cycle profile of earnings and labor supply. For recent work on the relationship between housing wealth, labor supply, and retirement decisions, see Zhao (2018)

Using a standard approach, as done in Borella et al. (forthcoming), I assume that female labor supply costs follow a quadratic function in age:

$$\xi(j) = \xi_1 j + \xi_2 j^2.$$

Denoting $X_j^R = (a_j, h_{j-1}, y_j, b^f)$ as the relevant state vector, the problem for a couple with a retired female is

$$V_j^R(X_j^R) = \max_{c_j, h_j, a_{j+1}} \Big\{ U(c_j, h_j, 0, j) + \beta \mathbb{E}_j V_{j+1}^R(X_{j+1}^R) \Big\} \tag{2}$$

subject to:

$h_j \in \{0, ..., \overline{h}\}$

$a_{j+1} + c_j + ph_j + \Phi(h_j, h_{j-1}) + q\mathbb{I}\{h_j = 0\} = (1+r)a_j + y_j + b^f + ph_{j-1}$

$a_{j+1} \geq \min\{a_j, 0\}.$

where male earnings y_j are stochastic until $J_R - 1$ but become constant afterwards.

Between ages J_r and $J_R - 1$ female retirement can happen with some probability between two periods, so that the relevant objective function is

$$V_j^A(x) = \max_{c_j, h_j, a_{j+1}, n_j^f} \Big\{ U(c_j, h_j, n_j^f, j) + \beta \Big[(1-\pi_r)\mathbb{E}_j V_{j+1}^A(X_{j+1}) + \pi_r \mathbb{E}_j V_{j+1}^R(X_{j+1}^R) \Big] \Big\} \tag{3}$$

subject to the same constraints as (1).

In the last period of life, the household must extinguish all debts and consume all remaining income and wealth.

The transaction costs function $\Phi(h_j, h_{j-1})$ is asymmetric with respect to selling and buying:

$$\Phi(h_j, h_{j-1}) = \begin{cases} ph_{j-1}F_s + ph_jF_b & \text{if } h_j \neq h_{j-1} \\ 0 \quad \text{if } h_j = h_{j-1} \end{cases}$$

The borrowing constraint function $\phi(h_j, y_j, w_j^f, n_j^f)$ is a key component of the model. Its specification builds on that of Attanasio et al. (2012), by assuming it depends on both the value of real estate holdings, previous financial assets, and current income. To illustrate its functioning, it is useful to express the household problem in terms of debt, $d_j = -a_j$. The borrowing constraint then becomes

$$d_{j+1} \leq \max\{d_j, \hat{\phi}(h_j, y_j, w_j^f, n_j^f)\},$$

$$\hat{\phi}(h_j, y_j, w_j^f, n_j^f) = \min\Big\{ \underbrace{\lambda_h p h_j}_{\text{LTV limit}}, \underbrace{(\lambda_y y_j + w^f \bar{n} \mathbb{I}\{n^f > 0\})}_{\text{LTI limit}} \Big\} \tag{4}$$

Renters are not allowed to have debt. Homeowners can hold debt subject to a set of collateral constraints. Specifically, buyers can finance the purchase of a housing unit through debt up to the minimum between the LTV limit $\lambda_h p h_j$ and the LTI limit $\lambda_y y_j + w^f \bar{n} \mathbb{I}\{n^f > 0\}$, where $\bar{n}$ is the average hours of full-time work. The indicator function $\mathbb{I}\{n^f > 0\}$ implies that, by supplying labor, females can increase the LTI limit. Moreover, households can use their house as collateral to borrow at any time, as long as their current level of debt already satisfies the two leverage ceilings. If not, they are unable to increase their debt. However, they are not forced to immediately satisfy the leverage limits whenever these are violated, although they must at least repay interests on the outstanding debt. Similarly, I assume retirees are allowed to pay off their debt gradually but not to borrow any further.

A few further comments regarding the modeling of debt and leverage limits are in order. Representing net liquid savings as a single continuous variable is motivated by the need to model mortgage-related borrowing constraints without having a decoupled choice of both deposits and an individual mortgage contract within the household's balance sheet. In reality, most households hold both positive liquid assets and a mortgage contract. However, modeling these two separately would entail doubling the continuous dimension of the state-space for assets, which is computationally cumbersome.[19] Allowing the household to choose any value of a_{j+1} subject to some constraint, rather than imposing a fixed debt repayment schedule, is equivalent to a yearly renegotiation of mortgage terms. This assumption is tenable given the UK's flexible institutional framework and the availability of "interest-only mortgages."[20]

The specification of the leverage limits allows for homeowners to suddenly become borrowing-constrained after an income shock, even without a change in assets. For instance, a household may obtain a mortgage at the time of purchase that satisfies both the LTV and LTI constraints. However, an income fall in the following period may suddenly imply that debt is above the LTI limit. The main implication is that a household with high net worth but also high debt can become borrowing-constrained.

19 See Druedahl (2015) for an analysis of this scenario.

20 Also known as "endowment mortgages", these require debtors to make regular interest payments while also accumulating savings in a separate endowment fund to repay the principal at maturity. This set-up gives households flexibility over the timing of debt repayment.

The LTI limit also entails that female employment can relax the borrowing constraint, although the secondary income carries less weight in the function. As discussed in the calibration section, this formulation captures the main features of mortgage contracts.

Figure 3 provides graphic intuition for the mechanics of the constraints. The solid black line represents the relationship between debt (x-axis) and the ratio of debt to house value (y-axis), for a given quantity of housing assets that a household owns (or purchases). The horizontal red dashed line at λ_h is the fixed LTV limit. The vertical blue dashed lines represent the LTI limit, which depends on the male's and female's incomes and on the female's labor decision.

Panel (a) shows the effect of higher male earnings in the case of a female not working. Assuming $y_m < y'_m < y''_m$, the LTI limit shifts to the right, allowing the household to access more debt. The change from y_m to y'_m moves the household debt limit to $\lambda_y y'_m$. However, since $\lambda_h ph < \lambda_y y''_m$, the LTV constraint becomes the binding limit for y'', and the higher income does not translate one-to-one into greater borrowing capacity. Panel (b) shows how female employment relaxes the borrowing constraint. In the first case, the constraint is represented by the LTI limit $\lambda_y y_m + w_f n$. In the second case, female wages are so high that employment allows the household to raise the debt ceiling only up to the LTV ceiling, which binds before the LTI limit. Finally, in Panel (c) female employment has no impact on the borrowing constraint since male earnings are very high and the LTV limit already binds.

1.4 Calibration

The majority of the parameters are calibrated externally, while the parameters concerning housing preferences, leisure, and childcare costs are chosen through an internal calibration.

1.4.1 External calibration

The period length is one year. I assume agents start working at age J_1 =25 and die at $J_D = 80$. Females start retiring with some probability at age $J_r = 50$, and mandatory retirement begins at $J_R = 65$ for both members.

Following the convention of many life cycle models, I assume that the stochastic parts of both male log earnings and female log wages are unit roots. The variances of the innovations, as well as the coefficients of the quadratic trends are estimated on

Figure 3: Examples of LTV and LTI constraints for a given level of housing.

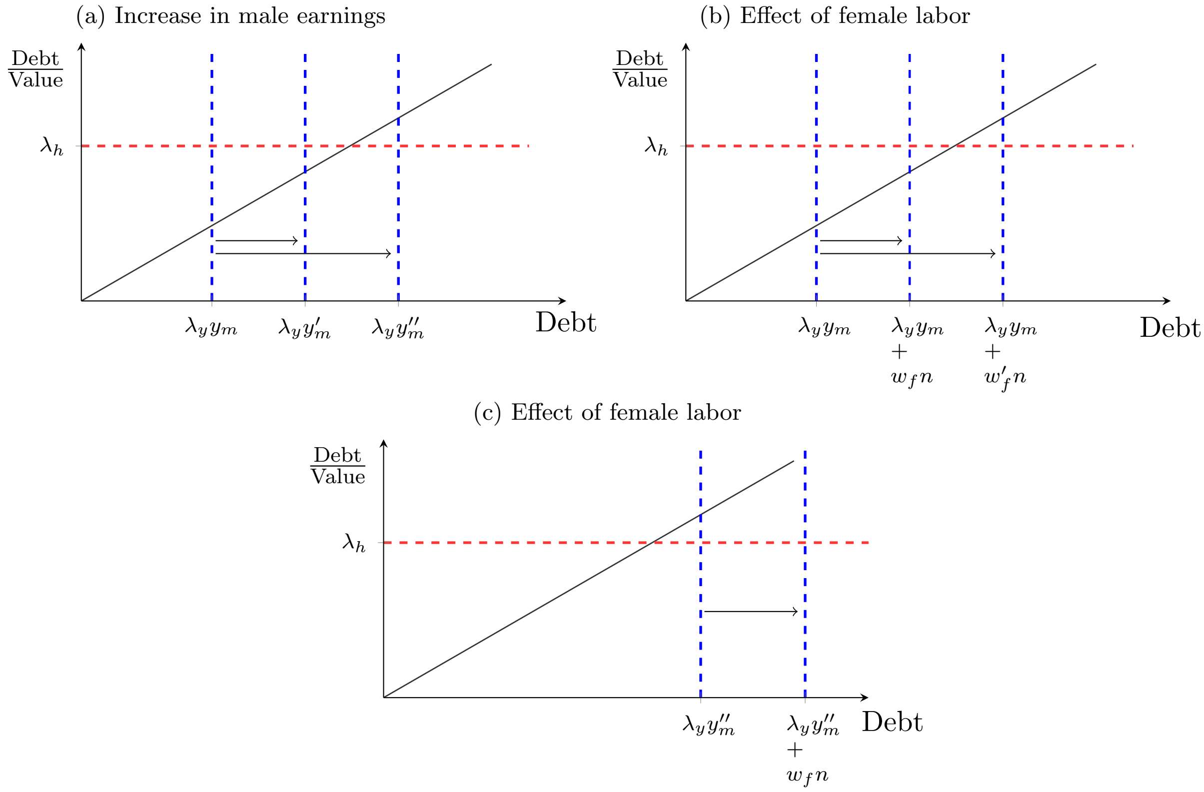

BHPS data following the approach of Blundell et al. (2008), as described in Appendix 1.A.[21] Average male earnings are normalized to 1. For males, I also assume that the probability of being unemployed in any period is 0.055, which approximates the average unemployment rate for the UK in the period 2001-2006.[22] I set the average earnings of females compared to males to 0.85. In other words, assuming a labor supply of 35 hours ($\bar{n} = 0.35$), the average wage for females is 85 percent of average male earnings.[23] I also assume that unemployment insurance yields 30 percent of average male earnings: i.e. $y^u = 0.3$.

The issues of assortative matching in household formation and intra-household income correlation are potential concerns for the calibration. A way to exogenously

21 Note that modeling earnings as a random walk implies that all shocks are permanent. Furthermore, while the distribution of the innovations remains constant, the overall variance of earnings increases with age.

22 This assumption abstracts from the important issue of persistence in labor market status. However, it increases clarity in terms of distinguishing the effect of temporary versus permanent income risk.

23 There is a large literature on the estimation of the gender gap, mainly attempting to correct for selection bias in observed wages. This issue is beyond the scope of the current work. The value of 0.85 is in line with estimates of the female pay gap in gross hourly earnings in the early 2000s from the UK's Annual Survey of Hours and Earnings, administered by the Office of National Statistics. The gap has been following a downward trend and is closer to 0.9 in recent years (ONS, 2016).

model assortative matching is to impose correlation in the earnings of the male and the female. Several studies model households where each member is either "high-skill" or "low-skill", and account for assortative forces by choosing the probability that both members belong to the same skill group. A related issue concerns the idiosyncratic earnings shocks to the two members, which may also be due to assortative matching. Instead of explicitly modeling permanent skills, I address the issue by drawing the starting earnings of the two members from a joint distribution with positive correlation. Specifically, I set the initial correlation in earnings to 0.2, in line with the findings of Lise and Seitz (2007) for the intra-household correlation of income in the UK in 2000. As the earnings processes are unit roots, this approach creates some persistence in earnings correlation over the life cycle. Furthermore, I allow the shocks to the two processes to be correlated. I choose the correlation to be 0.25, as used by Attanasio et al. (2015), a value estimated by Hyslop (2001). Although this value is estimated on US data, it is a tenable assumption to apply it to the UK, a similar economy, on a similar time period. Additionally, this value of the correlation in the *shocks* results in a correlation in the *level* of earnings within each age group of 0.2 on average, which is consistent with the initial target.[24]

The probability of retirement for females, π_r, is set based on the retirement rates from the BHPS for ages 50 and 65. Only 1 percent of females is retired at age 50, while 72 percent are retired by age 65. Hence $\pi_r = 1 - ((1-0.72)/.99)^{(1/15)} = 0.082$. This approximation matches the general retirement trend from age 50 in a linear way. The replacement rate of retirement earnings is 0.5, implying that workers receive retirement income equal to half of their final-year income for males and half of the full-time equivalent income (i.e. assuming $\bar{n} = 0.35$) for females.

Following Bajari et al. (2013), I let the the interest rate depend on whether the household's assets are positive or negative, reflecting the fact that returns on deposits are usually lower than interest rates for mortgages. I set a 3 percent return on savings ($r_s = 0.03$) and a 7 percent interest rate on debt ($r_d = 0.07$). The latter is close to the historical average nominal interest rate on new mortgages in the early 2000s in the UK (FSA, 2009).

I allow for four distinct amounts of housing assets: i.e. $h \in \{0, 1, 2, 3, 4\}$, where 0 represents renting. With some abuse of notation, instead of setting a unit price for

24 Sensitivity analysis, however, shows that the main results of the model are not strongly dependent on the level of intra-household earnings correlation.

housing assets, I set increasing prices for the four house sizes $p = [p_1,\ p_2,\ p_3,\ p_4]$. This formulation, although effectively equivalent, allows for a more intuitive internal calibration, as shown below. Using the sample of couples where the secondary earner is aged 23 to 65 in the 2001 wave of the BHPS, I set the prices based on different percentiles of the distribution of house values, normalized by the mean yearly income of working males in the 2001 wave of the BHPS, which is 18,448 GBP. I set $p_1 = 3.2$ to represent the 25^{th} percentile, $p_2 = 4.34$ for the 45^{th}, $p_3 = 6.17$ for the 65^{th}, and $p_4 = 9.7$ for the 85^{th}. I assume no price growth over time. Renting for one period costs one percent of a large house ($q = 0.097$). Following Yang (2009), when buying a house, 2.5 percent of the value has to be paid in costs ($F_b = 0.025$), while there is a cost equal to 7 percent of the value ($F_s = 0.07$) for selling.

The values of the leverage-based borrowing constraints are set to replicate in a parsimonious way the main features of the UK institutional environment in the mid-2000s, in line with Bottazzi et al. (2007) and Attanasio et al. (2012). I set $\lambda_h = 0.9$, so that a minimum downpayment of 10 percent of the house price is required for a purchase. This value is close to the typical maximum LTV ratio for the UK in the early 2000s.[25] For the LTI limit there is greater variation across lending institutions, and the maximum LTI often depends on whether the loan is undersigned by only one member of the household or both. The FSA's 2004 *Guide to Mortgages* states: "Typically, the maximum mortgage a lender offers is three times the main earner's income plus one times any second earner's income, or two-and-a-half times your joint income" (FSA, 2004). I therefore set $\lambda_y = 3$. Given the model's assumption that the male is the primary earner, a household's joint LTI limit is three times the male income plus one times the female's full-time equivalent income (if she works).

I set the CRRA parameter of consumption utility $\sigma = 2$ and the discount factor $\beta = 0.95$, which are standard values in the literature. The parameter ψ of leisure preferences is set to obtain a Frisch elasticity of labor supply of 0.3. In this case, defining $l = 1 - n^f$, $\epsilon_n^f = \frac{U_l}{U_{ll} * n^f} = \frac{1-n^f}{\psi n^f} = 0.3$, which implies $\psi = 6.19$ for full-time hours worked $\bar{n} = 0.35$.

The time-varying component of housing preferences $\chi(j)$ is computed based on the average number of children by age of the secondary earner in the BHPS, adjusted by the OECD equivalization scale. Details are contained in Appendix 1.A.3. Finally, I

25 However, in the years preceding the 2008 recession, mortgage rules were looser and higher LTV's were frequent.

calibrate the initial distribution of financial assets for the simulations on the empirical distribution of net worth for couples in the 2000 BHPS where the head is aged 23 to 27, leaving the details to Appendix 1.A.4.[26]

Table 1 reports all the externally calibrated parameters and their values.

Table 1: Externally calibrated parameters.

Parameter	Value	Description	Target/Source
Preferences			
β	0.95	discount factor	
σ	2	CRRA parameter	
ψ	6.19	leisure elasticity parameter	Frisch elasticity =.3
$\chi(j)$	see App. 1.A.3	equivalization coefficient	average household size
Life cycle and earnings			
J_1	25	starting age	
J_r	50	starting age of early retirement	
J_R	65	age of mandatory/male retirement	
J_D	80	final age	
ρ_z	1	persistence parameter of earnings	Attanasio et al. (2012)
σ^2_ϵ	.0133	variance of income shock - males	BHPS (Appendix 1.A)
$\sigma^2_{\epsilon^f}$	.0148	variance of income shock - females	BHPS (Appendix 1.A)
$\text{corr}(\epsilon, \epsilon^f)$	0.25	correlation of income shocks	Hyslop (2001)
$\alpha_1,\ \alpha_2$	.0576, -0.000834	income profile coefficients - males	BHPS (Appendix 1.A)
$\alpha^f_1,\ \alpha^f_2$	.0384, -0.000468	income profile coefficients - females	BHPS (Appendix 1.A)
b	$0.5 * y_{J_{R-1}}$	retirement income	
ω_f	0.8	gender earnings gap	
π_u	0.055	probability of male unemployment	UK unemployment rate
y^u	0.3	unemployment insurance	
π_r	0.082	retirement probability, $j = 50, .., 65$	BHPS
Housing market			
$F_b,\ F_s$	0.025, 0.07	buying / selling transaction costs	Yang (2009)
λ_h	0.9	LTV borrowing limit	Attanasio et al. (2012)
λ_y	3	p-LTI limit	FSA (2004)
r_s	0.03	interest rate on savings	
r_b	0.07	interest rate on debt	FSA (2009)
$p_1,\ p_2,\ p_3,\ p_4$	3.2, 4.24, 6.17, 9.7	price of housing units	BHPS
q	0.097	rental cost	

1.4.2 Internal calibration

The two housing preference parameters ϕ_h and μ_h are internally calibrated to match the empirical average homeownership rate of 82 percent for households between the ages of 25 and 50 and to obtain an average value of housing assets of 5.4 in the same age range.

The parameters of the labor supply cost function ξ_1 and ξ_2 are calibrated to match the employment rate of females in the age groups 25-29, 30-34, 35-39, 40-44, 45-49. The resulting function, displayed in Figure 1.D.1, has a hump-shaped path

26 The BHPS only provides information on savings and unsecured debt every five years. I use the 2000 wave as it is the closest to the beginning of the sample.

with a negative slope after age 30. The linear coefficient of the leisure preferences θ^f is calibrated to match an average of 35 hours worked for employed females (i.e. $n^f = 0.35$), equivalent to standard full-time contracts in the UK.

I focus the internal calibration on aggregate moments up to age 50 because I do not aim to explain dynamics specific to retirement decisions. Furthermore, I explicitly calibrate moments that are not related to housing debt levels and leverage. Table 2 reports the internally calibrated parameters. Table 1.D.1 reports the targeted moments and the corresponding values obtained from 10,000 simulations of the exogenous earnings processes and an equal number of draws from the distribution of initial wealth.

Table 2: Internally calibrated parameters.

Parameter	**Value**	**Description**
μ_H	.25	Housing preference constant
ϕ_H	.35	Housing size preference
ξ_1	0.026	Labor cost, linear coefficient
ξ_2	0.00045	Labor cost, quadratic coefficient
θ^f	0.17	Linear coefficient for leisure

1.5 Results

This section presents the baseline results of the model. First, I show how the simulations match the aggregate profiles from the BHPS. I then examine the cross-sectional relationship between leverage and employment, discussing how earnings composition within the household underpins the results.

1.5.1 Aggregate life cycle profiles

The upper plots of Figure 4 report the average profiles for homeownership, female employment, and house value from the 10,000 simulations. Overall, the model captures the main life cycle dynamics of the BHPS. The homeownership profile starts lower in the model than in the data and eventually overshoots, leading to an overall steeper profile but matching the shape of the path well. Female employment peaks at 50 in the model, when the retirement transitions start, while in the data it peaks slightly earlier. The quantity of housing has a very similar upward trend, although it levels off in the data.

Figure 4: Average simulated life cycle profiles for homeownership rate, home size, female employment, net worth, LTV, LTI, and p-LTI.

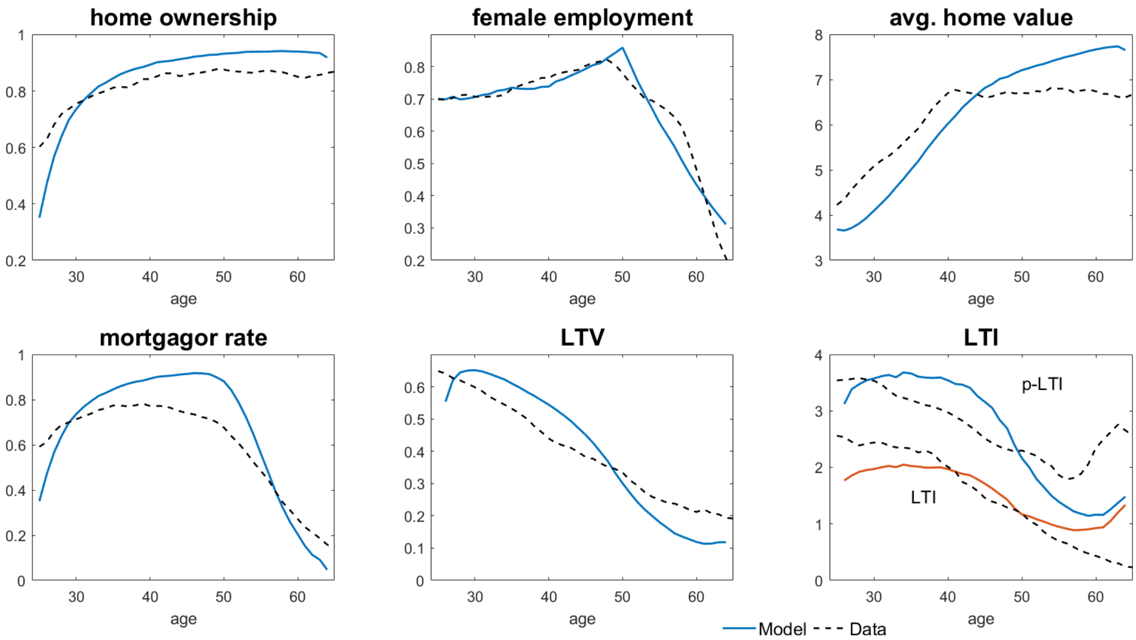

Note. The solid lines represent the average values from the simulated model. The dashed lines represent the moving average of the empirical life cycle profiles from the BHPS. For the LTI panel the upper blue solid line reports the p-LTI, i.e. the LTI computed only using the male's earnings, which is compared against the empirical LTI computed using earnings from the primary earner and other income sources. The lower red line represents the LTI computed with the entire household income, which is compared against its corresponding series from the BHPS. The simulated profiles are computed as averages from 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth.

The lower plots of Figure 4 focus on households' leverage, which are not targeted by the internal calibration. The percentage of mortgagors follows a similar path to the empirical one from the BHPS, although its peak occurs later than in the data and exceeds it in value. The average LTV tracks the empirical series closely. The p-LTI (the higher series in the last subplot) overall tracks the empirical counterpart well, initially exceeding it but later falling below it in value. The household's joint LTI (the lower red series) starts slightly below the empirical one but, because of its flatter slope, eventually overshoots. Overall, the life cycle paths and the average levels of the ratios are well captured by the model.

1.5.2 Female employment and leverage ratios

To elaborate on the cross-sectional relationship between leverage and labor, Figure 5 plots the employment rate of females across the distributions of LTV and p-LTI over 10-year age groups. The plots show the relative employment rate of households at a given quintile of the LTV or p-LTI distribution relative to households in the first quintile. The left panel shows that, except for the first age group, there is a negative relationship between LTV and employment. For the age groups 35-44 and 45-54, the employment rates of households in the fifth quintiles are 40 and 17 percentage points lower, respectively, compared to that of households of the same age group in the first quintile. For the 25-34 age group, however, the plot is hump-shaped and positive, implying a negative relationship only from the third quintile onwards. The right panel shows an opposite relation for the p-LTI. Here, across all age groups the plot is upward sloped, indicating a higher employment rate for households whose p-LTI is high. The relationship is quantitatively very large, especially for the youngest group. In the appendix, Figures 1.D.2 and 1.D.3 break down the plot by house size to show that the relationship between employment and leverage holds across levels of wealth. In particular, Figure 1.D.4 shows that between two households with the same net worth but different house sizes, the one with the higher LTV has on average a lower labor supply.[27]

It is worth noting that, for a given LTV (p-LTI) value, households will have different p-LTI (LTV) ratios. To render the relationship graphically, in the appendix Figure 1.D.5 displays a surface plot of the employment rate over the LTV and p-LTI ratios from the simulations using the same 10-year age groups. The nonlinear relationship between the three variables emerges from the plots. For instance, the negative relationship between employment and LTV is steeper at lower levels of the p-LTI while it is almost absent for high p-LTI's. Similarly, the positive slope of the relationship with the p-LTI is steeper for high LTV values. This result is consistent with the intuition explained above. For example, a household with a high LTV and a low p-LTI has a greater ability to substitute future consumption for current leisure than one with a high LTV and a high p-LTI.

To better understand the characteristics of households with high leverage, Table 3 provides summary statistics divided by levels of leverage and across 10-year age groups. The upper panel divides households into those in the bottom 80 and the top

27 For a given level of wealth, the owner of a small house has a lower LTV compared to the owner of a larger one. The figure shows that, in the regions of the net wealth axis where smaller and larger owners overlap, the latter have a lower employment rate. Therefore, the negative correlation between LTV and employment holds at all levels of housing, implying that leverage is the main driver of employment decisions rather than net wealth.

Figure 5: Employment rate of females in different 10-year age groups over quintiles of the the LTV and p-LTI distribution relative to those in the first quintile.

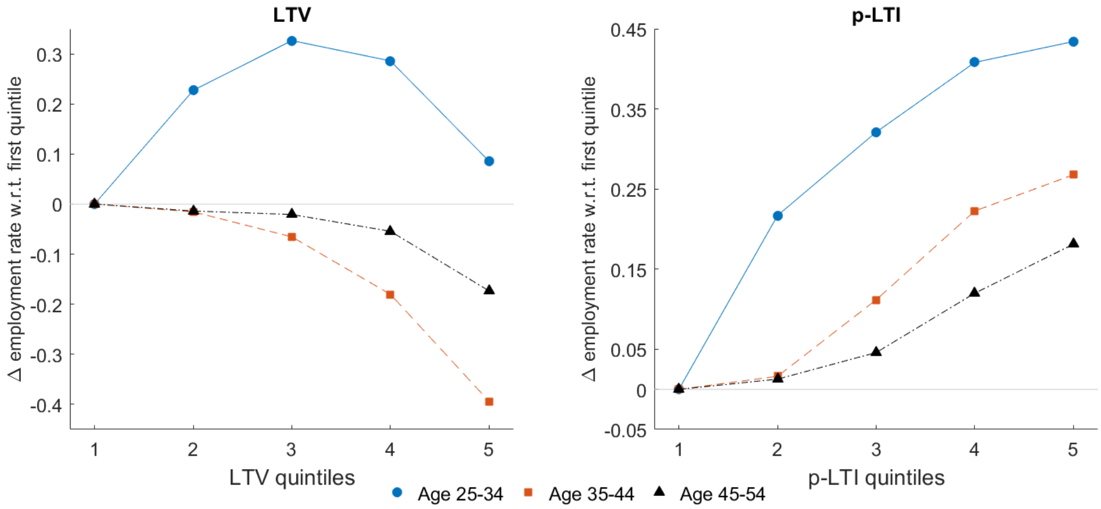

Note. Each plot is produced by grouping households into 10-year age groups. For each group, each point represents the employment rate (y-axis) for a quintile of the LTV (x-axis) and an age group. In this way, the plot shows how a variable tends to change across the simulated distribution of the x-variable. The plots are produced using 10,000 simulations of individual income shocks starting from the calibrated initial distribution of net worth.

20 percent of the LTV distribution for each age group, while the lower panel similarly divides households based on the p-LTI.

Consistent with Figure 5, the upper panel shows that high-LTV households have a lower employment rate than the low-LTV group at all ages, with the difference being particularly large for young and mid-life ones. The interaction of the leverage constraints with the workers' earnings is central to the results. Those in the high-LTV group also have higher average male earnings and lower female wages. Given the unit-root nature of earnings, high (low) current earnings imply high (low) expected future earnings. For a given level of housing, accessing debt up to the LTV limit is only possible if the LTI limit does not bind first. Hence, households with high primary earnings are able to access higher LTV's without the need for female labor. Furthermore, within this group, those with low potential female wages are those who choose to avoid costly female labor in exchange for debt. The high-LTV group therefore includes those who have a high motive to substitute away from female labor and a high capacity to borrow at early ages thanks to high expected primary earnings.

As the lower panel of Table 3 shows, households in the top 20 percent of the p-LTI distribution have a higher employment rate than those with a low p-LTI. On average,

Table 3: Summary statistics by LTV and p-LTI levels across 10-year age groups.

	LTV					
	Age 25-34		**Age 35-44**		**Age 45-54**	
	Bottom 80%	**Top 20%**	**Bottom 80%**	**Top 20%**	**Bottom 80%**	**Top 20%**
Employment	0.82	0.62	0.85	0.56	0.79	0.74
Avg. $\boldsymbol{w^f}$	0.73	0.56	0.92	0.52	1.02	0.58
Avg. $\boldsymbol{y^m}$	0.78	0.86	1.06	0.98	1.20	1.13
Avg. $\boldsymbol{\%\Delta y^m}$	0.04	0.02	0.03	0.01	0.03	0.00
% Male Unemployed	0.06	0.05	0.05	0.06	0.05	0.06

	p-LTI					
	Age 25-34		**Age 35-44**		**Age 45-54**	
	Bottom 80%	**Top 20%**	**Bottom 80%**	**Top 20%**	**Bottom 80%**	**Top 20%**
Employment	0.74	0.92	0.75	0.93	0.75	0.90
Avg. $\boldsymbol{w^f}$	0.68	0.74	0.83	0.90	0.92	0.96
Avg. $\boldsymbol{y^m}$	0.87	0.50	1.17	0.54	1.33	0.63
Avg. % $\boldsymbol{\Delta y^m}$	0.09	-0.17	0.08	-0.18	0.06	-0.14
% Male Unemployed	0.00	0.26	0.00	0.27	0.02	0.21

Note. All results are produced using the same set of 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth.

they also have higher female wages and much lower male earnings. The difference in earnings across the two groups reflects the two main mechanisms driving high levels of p-LTI. First, households with low male income can use female labor income to relax the LTI constraint and access more debt. Second, by construction, a fall in male income leads to a rise in the p-LTI. As reported in the fourth and fifth rows of the lower panel, negative income shocks to the primary earner underlie high levels of the p-LTI: households in the top 20 percent have a negative average income growth and high unemployment. Female labor supply thus acts as a means of insurance against the adverse shocks to primary earnings that lead to high p-LTI's.

1.6 The effect of changes in the constraints

In this section, I discuss the propagation of leverage limits from the housing market to labor supply choices. I consider how changes in the LTV and LTI constraints may exert different effects on homeownership and labor supply over the life cycle. Not only do the two constraints interact differently with households' decisions, but also households who are close to either constraint differ with respect to their earnings composition and will therefore have distinct responses. To answer this question, I solve the model under alternative values for the leverage limits, compare the resulting life cycle profiles to the baseline, and consider the welfare implications.

1.6.1 Changes in the LTI constraint

I consider the case where the weight of female earnings within the LTI constraint is tightened (relaxed) enough to produce a fall (rise) in the homeownership rate of 10 percentage points for households at age 25 (i.e. the youngest age). Specifically, I allow the debt limit to be equal to λ_y times the male's income plus λ_y^f times the female's full time earnings (if she works). In other words:

$$\text{LTI limit} = \lambda_y y_j + \lambda_y^f w_j^f \bar{n} \mathbb{I}\{n_j^f > 0\},$$

where $\lambda_y^f < 1$ in case of a tightening and $\lambda_y^f > 1$ for a loosening compared with the baseline calibration (i.e. $\lambda_y = 1$). The two values leading to 10-percentage point falls and rises in homeownership in the first age group are 0.18 and 1.55, respectively.

A change in the limit alters credit availability in two ways. First, holding female labor supply constant, it lowers or raises the LTI-based debt limit for households where the female is employed. Second, it may alter the labor supply decisions of some households. For instance, a LTI loosening induces some females who previously were not employed to work in order to obtain credit and purchase a house. Similarly, a lower λ_y^f decreases the returns to work for some females and induces a decrease in labor supply at the extensive margin.

The left panel of Figure 6 shows the alternative life cycle profiles of homeowenrship. In both cases, the difference with the baseline profile at age 25 is 10 percentage points. However, the gaps closes up between ages 35 and 40, implying that the leverage limit has a strong impact on young households but not on middle-aged and old ones. The right panel shows the profiles of labor supply. Female employment rises slightly in the case of the relaxed limit and falls under the tightened scenario. Both changes in the initial age are between 1 and 2 percentage points. However, the difference also persist farther into the life cycle than for homewonership. The persistence is due to the implications of the leverage limit for mortgagors' debt level: higher ability to access credit at younger ages implies a larger debt level for older households and in turn a higher labor supply. Furthermore, at all ages there are households experiencing negative shocks to primary earnings who thus rely on female work to access more debt.

Figure 6: Homeownership and female employment in the baseline model and in the model with the tightened/relaxed LTI constraint.

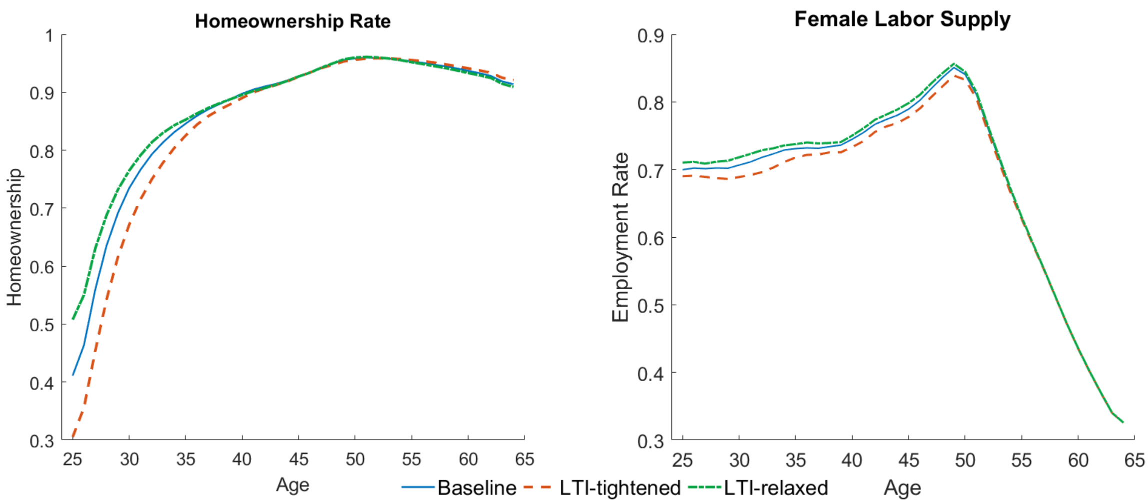

Note. All results are produced using the same set of 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth and the policy functions for the baseline model and the alternative calibrations.

To focus on the marginal households affected by the change, I divide the simulations into three groups based on the tenure status at age 25. The first group includes those who own a house in both the baseline and the counterfactual-LTI model, the second group is formed by those who change their tenure status at age 25. For the LTI tightening case, these are households who owned a house in the baseline case and become renters in the counterfactual scenario (i.e. "new renters"). For the LTI loosening case, these are renters who become owners (i.e. "new owners"). The third group includes those who are renters in both cases.

Table 4 reports summary statistics by tenure group. The "new renters" and the "new owners" groups, within each respective exercise, are indicative of the interaction between the extensive margin of labor supply and the LTI limit. Both groups are characterized by high female wages, high male earnings, and low initial assets, indicating the need to borrow to purchase a house, and with a high potential to do so through female earnings. In fact, in both cases, these "marginal" group accounts for the main change in aggregate employment. For the tightening, the employment rate of new renters falls by almost 14 percentage points. Meanwhile, under the looser LTI limit, the employment rate of the new owners rises by 17 percentage points.

It should be noted that the relationship between the credit limit and aggregate labor supply evinced by this exercise is distinct from the canonical relationship present in standard life cycle models. In particular, in models where the borrowing constraint is specified in terms of net wealth or of the ratio of assets to liabilities, a tightening of the borrowing limit usually induces a higher labor supply via the precautionary savings channel. The LTI ceiling, however, introduces an additional labor supply channel which is relevant for a fraction of households.

Table 4: Summary statistics for the three tenure groups with respect to the baseline model and the model with tightened/relaxed LTI model.

LTI Tightening

	% of HH's	**Δ Empl. (p.p.)**	**Avg. w_j^f**	**Avg. y_j**	**p-LTI in baseline case**	**LTV in baseline case**	**Assets**
Owners in both cases	0.26	0.66	0.55	0.62	2.81	0.47	1.97
New Renters	0.10	-13.79	0.62	0.66	3.96	0.79	0.34
Renters in both cases	0.65	-0.05	0.50	0.52	0.00	0.00	0.13
Aggregate	1.00	-1.17	0.52	0.56	3.12	0.55	0.62

LTI Loosening

	% of HH's	**Δ Empl. (p.p.)**	**Avg. w_j^f**	**Avg. y_j**	**p-LTI in reform case**	**LTV in reform case**	**Assets**
Owners in both cases	0.35	-0.26	0.57	0.63	3.35	0.57	1.53
New Owners	0.10	17.34	0.60	0.62	4.67	0.85	0.19
Renters in both cases	0.55	-0.02	0.48	0.50	0.00	0.00	0.12
Aggregate	1.00	1.67	0.52	0.56	3.65	0.63	0.62

Note. All results are produced using the same set of 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth and the policy functions for the baseline model and the alternative calibrations. The employment change is reported in percentage points.

1.6.2 Tightening the LTV constraint

The second counterfactual exercise considers the effect of changes in the LTV limit on homeownership and labor supply. For brevity, I focus on the case of a tightening of the LTV constraint only.[28] I choose an alternative value of λ_h equal to 0.675, which

28 I do not present the case of LTV loosening for two reasons. First, under the current calibration there is no alternative value for the LTV ceiling that yields a rise in homeownership of 10 percentage points in the youngest cohort. The reason is that for LTV values above 0.9, the LTI limit becomes the binding constraint for all households. Furthermore, LTV values above 1 would imply households with negative net worth at the time of purchase and would introduce the incentive for bankruptcy, which I do not model.

yields a fall of 10 percentage points in the homeownership rate of households at age 25.[29]

The left panel of Figure 7 shows that the tightening of the LTV limit lowers homeownership in the early part of the life cycle. Compared with the LTI tightening, however, the gap is less persistent and closes up before age 35. The right panel shows that the lower LTV ceiling does not have an effect on the aggregate female labor supply.

As above, I divide the households into three groups based on housing tenure at age 25: owners in both periods, "new renters", and renters in both periods. The middle group accounts for the fall in homeownership. Table 5 shows that the "new renters" are characterized by high female wages, highly above-average male earnings, and low initial wealth. They thus comprise those households with a strong motive to borrow against future income (both male and female) in order to access housing at young ages. Intuitively, because they have very high earnings, the LTV limit is the relevant borrowing constraint for this group. The ability to borrow is impaired by the more restrictive LTV limit, which cannot be affected by females moving into employment. In fact, in the baseline case, their average LTV is 0.85, which is above the new limit.

Interestingly, in the low-LTV scenario, labor supply for the "new renters" falls by 4.8 percentage points while it rises by roughly 1 percentage point or less for the other two groups. Some of the "new renters" switch out of employment because labor supply is no longer needed to access credit. However, the larger downpayment needed for homebuying also induces greater savings motives for the other two groups. The constant aggregate rate therefore masks changes within the tenure groups.

1.6.3 Welfare benefits and costs of leverage limit changes.

Because leverage limits alter the decisions of households on the margin of the respective borrowing constraint, it is worth analyzing whether changes in leverage limits also have a heterogeneous impact on welfare. To this end, I compute the consumption equivalent welfare benefit/cost for every household under each of the above counterfactual scenarios. While the details of the computation are left to Appendix 1.C, this metric can be interpreted as the constant proportional change in consump-

29 The magnitude of the tightenting in the LTV limit, 0.225, is large but realistically relevant. For instance, the UK government's 2013 Help-to-Buy scheme provides interest-free mortgages on 20 percent of the house price.

Figure 7: Homeownership and female employment in the baseline model and in the model with the tightened LTV constraint.

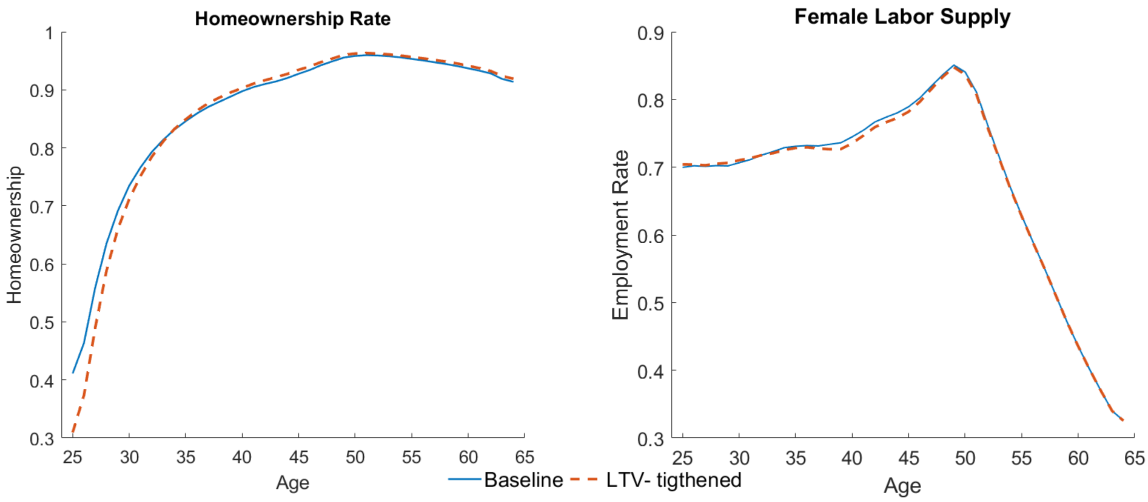

Note. All results are produced using the same set of 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth and the policy functions for the baseline model and the tightened-LTV model.

Table 5: Summary statistics for the three tenure groups with respect to the baseline model and the tightened-LTV model.

	% of HH's	Δ Empl. (p.p.)	Avg. w_j^f	Avg. y_j	p-LTI in baseline case	LTV in baseline case	Assets
Owners in both cases	0.25	0.85	0.55	0.58	2.82	0.43	2.12
New Renters	0.10	-4.77	0.63	0.75	3.84	0.85	0.11
Renters in both cases	0.65	1.45	0.50	0.52	0.00	0.00	0.13
Aggregate	1.00	0.68	0.52	0.56	3.12	0.55	0.62

Note. All results are produced using the same set of 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth and the policy functions for the baseline model and the tightened-LTV model. The employment change is reported in percentage points.

tion in each period that would raise/lower a household's *ex post* lifetime discounted utility to the point where they are indifferent between the baseline economy and the counterfactual scenario. A positive value therefore indicates that the counterfactual scenario entails a higher lifetime welfare for the household.

Table 6 reports the average welfare benefit or cost by tenure group and for the aggregate economy in each counterfactual exercise. Intuitively, greater credit availability via a looser LTI limit implies a welfare increase for all households, with the average value just above one percent of per-period consumption. The benefit is almost double for the new homeowners, who can immediately purchase a house in the first year of life under the relaxed limit. The two scenarios with tighter limits entail an aggregate welfare cost. The LTI tightening shows a substantially higher aggregate cost, which is consistent with the fact that the fall in homeownership is more persistent than in the case of a stricter LTV. The other key difference is that the former case implies a fall in welfare of more than one percent even for those households who are homeowners in the first period of life. Finally, in both cases the new renters experience the largest fall in welfare, of 3.2 and 2.6 percent, respectively. Overall, this exercise points to large heterogeneity in the welfare impact of leverage constraints across households and depending on whether a comparable restriction of credit occurs through the LTI or the LTV limit.

Table 6: Welfare benefits and costs of the counterfactual policy exercises expressed in terms of percent consumption compensation equivalent.

	Average % consumption compensation equivalent		
	LTI loosening	LTI tightening	LTV tigthening
Homeownwers in both cases	0.83	-1.27	-0.15
New Owners	2.03		
New Renters		-3.21	-2.59
Renters in both cases	1.05	-1.84	-1.15
Aggregate	1.07	-1.83	-1.05

1.7 Response to income shocks

The model can be used to understand how the response of female labor supply to unexpected permanent income shocks varies with leverage. Because the shocks change the permanent component of earnings, they alter households' choices not just at the age in which they occur but also in the following years. I thus solve counterfactual models in which either all female or male earnings suffer a one-standard deviation drop at given ages and compare the ensuing average life cycle profiles to the baseline simulation.

Figures 8 and 9 plot the responses of female labor supply, homeownership, and consumption at ages 30, 40, and 50 to unexpected one-standard deviation permanent falls in female and male earnings, respectively. Both shocks cause a permanent fall in consumption as well as a persistent fall in homeownership. While the percentage fall in consumption is similar in the three ages considered, the fall in homeownership

Figure 8: Response to an unexpected permanent fall in female wages at different ages.

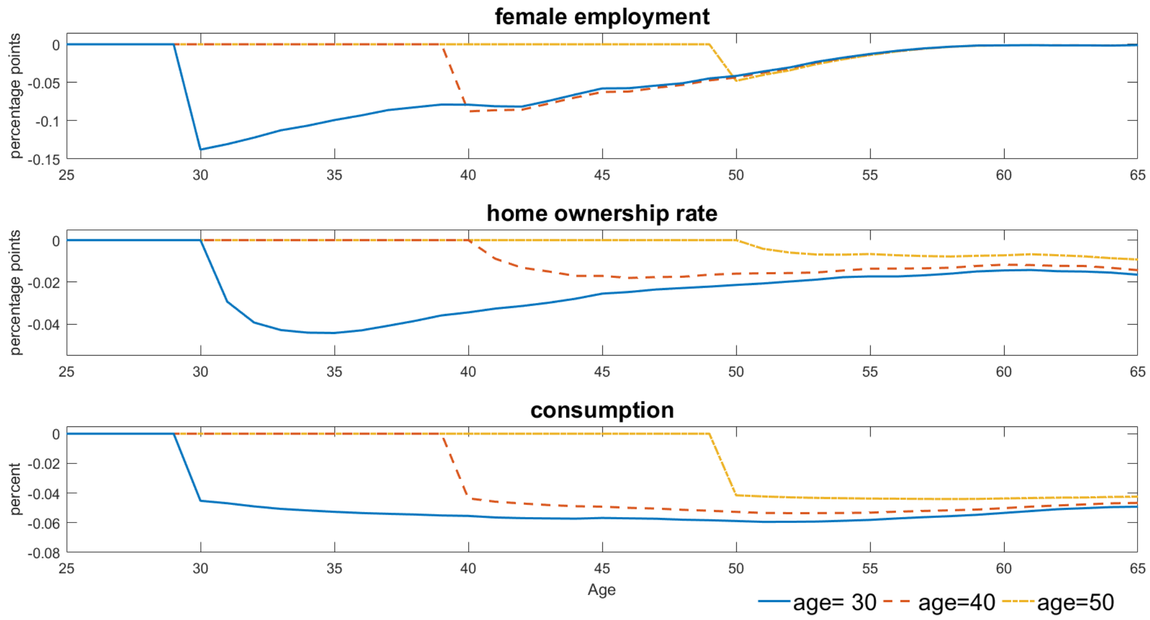

Note. Each plot reports the response of the respective variable in percentage points or percent to an unexpected one-standard deviation permanent fall in female wages occurring at ages 30, 40, and 50. All plots are produced with the same set of 10,000 individual simulations.

is larger at younger ages.[30]

Labor supply responds negatively to the fall in female wages. For young households, the initial response is very large but progressively winds down over the years. The magnitude of the initial fall decreases with age. As labor supply costs fall with age while wages rise, the drop in earnings becomes a less determining factor for the extensive labor margin.

The fall in male income triggers a persistent added worker effect between 2 and 4 percentage points. The initial jump in employment does not vary strongly with the age at which the shock occurs. Furthermore, the rise in employment starts to wind down only after age 50, when females begin to retire. Despite its presence throughout the work years, the average added worker effect is of modest magnitude compared to the response to a fall in female wages.

30 As homeownership rate for young households is lower than for older ones, the main margin for the fall in the aggregate rate is the postponing of home purchases rather than an increase in home sales.

Figure 9: Response to an unexpected permanent fall in male earnings at different ages.

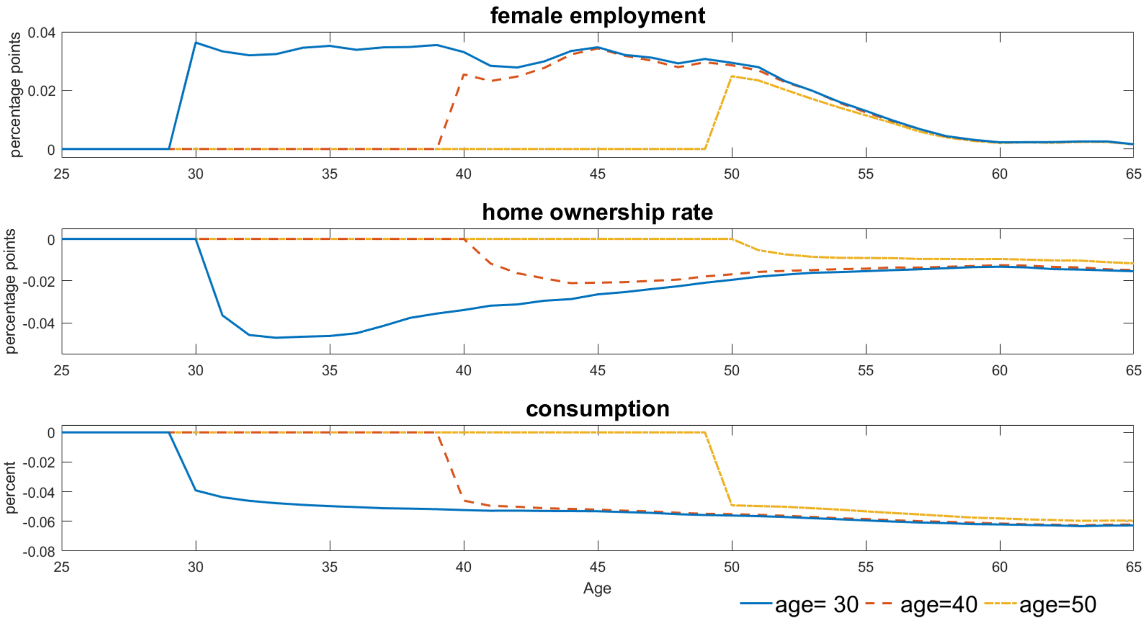

Note. Each plot reports the response of the respective variable in percentage points or percent to an unexpected one-standard deviation permanent fall in male earnings occurring at ages 30, 40, and 50. All plots are produced with the same set of 10,000 individual simulations.

1.7.1 Decomposing the aggregate response

How much heterogeneity do the average responses in labor supply mask? To answer this question, I break down the average response to the income shock along a chosen dimension of leverage for all ages from 30 to 55. I focus on the difference between high-LTV and low-LTV homeowners, choosing as the threshold the value of 0.65, which is close to the 80^{th} percentile of the LTV distribution in the simulations. A similar analysis by high and low p-LTI is left to Appendix 1.E. It must be noted that the composition of these groups and their size change with age: highly leveraged households at age 30 are not the same who are leveraged at 55.

Figure 10 provides a visual representation of how the immediate response of labor supply to an earnings fall differs across groups and across ages. Each point represents the contemporaneous fall of the employment rate in percentage points for a given age group belonging to households with a high or low LTV, as well as for all homeowners.[31]

The left panel decomposes the immediate employment change after a fall in female wages. The immediate fall in labor supply is larger for high-LTV households.

31 As similar picture with the breakdown by p-LTI is provided by Figure 1.E.1.

Figure 10: Contemporaneous labor supply response to female and male earnings shocks for low-LTV and high-LTV homeowners across age groups.

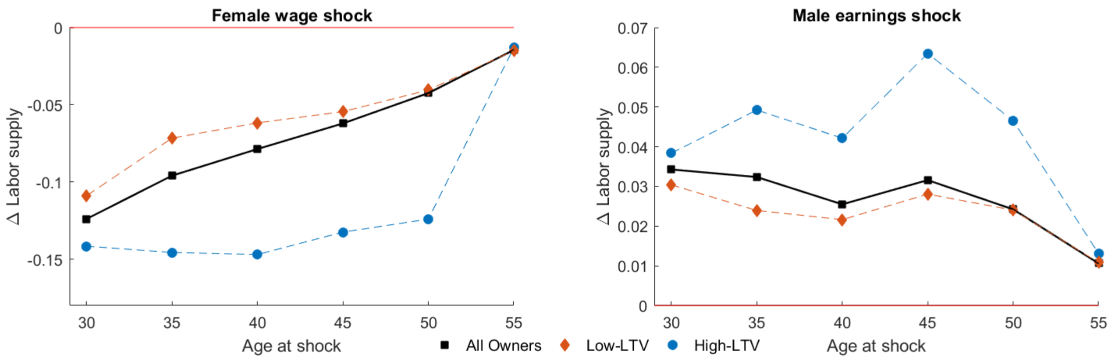

Note. Each point reports the response of the extensive margin of female labor supply in percentage points (i.e. the employment rate) to an unexpected one-standard deviation permanent fall in female wages or male earnings occurring at a given age. Households are divided into low-LTV (LTV<0.65), and high-LTV homeowners LTV (≥0.65). The black squares report the response for all homeowners. All plots are produced with the same set of 10,000 individual simulations.

Additionally, the response of low-LTV households has a more pronounced life cycle profile. The fall for low-LTV households is 10.5 percentage points at age 30 and only 4.5 at 50. Meanwhile, the corresponding values for high-LTV owners are 14 and 13. It is only after retirement transitions begin that the response of the two groups is comparable. Because the proportion of high-LTV households falls with age (see Figure 1.D.6 in the Appendix), the aggregate response of all homeowners in later stages of the life cycle is progressively closer to that of low-LTV ones.

The result that high-LTV females are the most responsive to wage shocks stems from the intuition outlined above. This group includes households with high male income and low female earnings, and hence with a large propensity to exchange current leisure with future commitments to repay debt.

The right panel reports the case of a shock to male earnings. High-LTV households have a more elastic extensive margin: the added worker effect (in the first period) is always larger for this group than for low-LTV households. This result is also consistent with the previously discussed mechanisms. High-LTV households are those who rely the most on male earnings to avoid costly female work. The fall in male income impairs this channel by lowering lifetime earnings and the ability to amortize outstanding debt. The shock also entails an immediate increase of the p-LTI, which pushes more households above the LTI ceiling and leads more females to work to relax the borrowing constraint. Figure 1.D.7 in the Appendix shows the increase in

Table 7: Welfare losses from shocks expressed in % consumption compensation equivalent by age and tenure type.

Shock	Age	Welfare costs (consumption equivalent)					
		LTV		p-LTI			
		Low	High	Low	High	Renter	Aggregate
Female Earnings	30	-7.3	-7.3	-5.5	-8.3	-7.1	-7.2
	40	-7.0	-4.9	-5.0	-7.3	-5.9	-6.3
	50	-5.2	-3.7	-4.6	-6.6	-5.1	-5.1
Male Earnings	30	-9.9	-11.5	-11.9	-10.3	-11.1	-11.0
	40	-8.6	-11.5	-10.7	-8.6	-9.4	-9.5
	50	-7.4	-10.7	-8.0	-6.4	-6.5	-7.5

Notes: the table reports the average consumption equivalent compensation for the welfare loss to the respective group. The compensation corresponds to the proportional increase in consumption each period for the counter-factual scenario that would make the household indifferent *ex post* between the compensated scenario and the baseline case. Appendix 1.C describes the computational details.

the percentage of households with a high p-LTI (i.e. greater than 3.5) after the fall in male income. Among the high-LTV group there is a greater rise in the fraction of households who are also facing a high p-LTI, providing a further channel for their larger response.

1.7.2 Welfare costs of income shocks

Heterogeneity in the response of labor supply to income shocks also entails heterogeneity in the resulting welfare costs based on the household's ability to adjust consumption and housing. I thus use the consumption equivalent measure to gauge the variation in welfare costs arising from falls in female potential wages or male earnings. Using the same set of simulations from the previous sections, I compute the consumption equivalent compensation for each household and then take the average across ages and tenure groups (i.e. high- and low-LTV owners), which are reported in Table 7.[32]

Several trends emerge from the table. First, welfare losses decrease with age for all types of households and hence for the aggregate economy. Second, comparing high- and low-LTV households, the former incur lower welfare losses from a female wage shock than the latter but a higher loss from a male earnings shock. This result is consistent with the difference between the two groups in the labor supply

32 Details of the computation are left to Appendix 1.C.

responses discussed in Section 1.7.1. High-LTV households substantially decrease female labor supply after a female wage shock. Since this group tends to have low female wages to start with, the loss of earnings from moving out of employment is smaller. Moreover, as they have high expected male income, they have a greater capacity to borrow to substitute away from costly labor. Hence, they are better able to buffer the diminished female productivity. For a similar reason, when a fall in male earnings occurs, high-LTV households have to increase their low-productivity female labor supply the most to make up for the lost income. Third, focusing on the breakdown by p-LTI, high-p-LTI households at all ages are more damaged by female earnings shocks and less harmed by male earnings falls than low-p-LTI ones. Given the mechanisms highlighted by the model, this result is also intuitive. High-p-LTI households use female labor supply to support consumption and debt repayment, while relying less on the male's income.

1.7.3 Adding house price shocks

So far, the analysis has maintained house prices constant, implying that households do not face exogenous wealth volatility or unexpected changes in their ability to use housing assets as collateral. This is a tenable assumption for normal times, when aggregate prices exhibit only small fluctuations. However, the occurrence of often-unexpected large price falls, like that of 2008-2009, calls for an assessment of the interaction between price drops and recessionary periods.

Using the same framework as above, in this section I consider the effect of a large and persistent price fall on the responses to earning shocks. In particular, my interest lies in whether the price shock affects the female labor supply dynamics ensuing from income drops.

A change in house prices has several effects on households depending on their financial position. For renters, it makes housing cheaper to purchase (assuming no changes in the ability to access debt). For current homeowners, it has a wealth effect that realizes only in the case of a sale. Hence, the fall in wealth only affects those who plan to downscale or upscale their housing assets. The latter group, however, also benefit from the fall in the cost of the future housing they plan to purchase. Finally, the fall in house prices implies a tightening of the LTV limit, diminishing the borrowing ability of households with already high leverage. Given the results of the

Figure 11: Contemporaneous response of female labor supply to earnings shocks with and without the addition of a fall in house prices, across high- and low-LTV groups and ages.

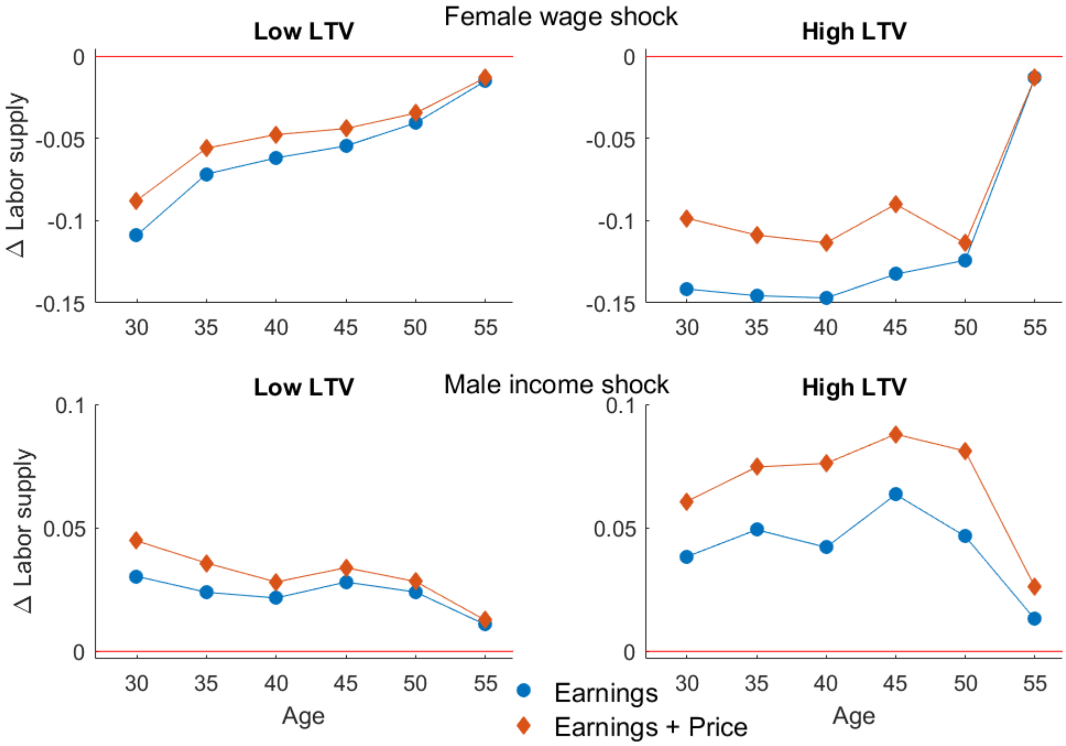

Note. Each circle point reports the response of the extensive margin of female labor supply in percentage points (i.e. the employment rate) to an unexpected one-standard deviation permanent fall in female wages (top row) or male earnings (bottom row) occurring at a given age. The diamond points report the case in which the shock is combined with a persistent fall in house prices of 10 percent, as described in Section 1.7.3. The left plots consider homeowners with a LTV ratio below 0.65, and the right ones consider households whose LTV is equal or greater than 0.65. All plots are produced with the same set of 10,000 individual simulations.

previous section, the effect of this case on labor supply is worthy of discussion.

To investigate this issue, I simulate an additional counterfactual case in which the earning shocks are combined with a persistent change in house prices. The simulated price fall is 10 percent, a large value which is approximative of the fall in the real UK House Price Index (UKHPI) in the Great Recession. The ensuing path is computed using an AR(1) autoregressive process with a coefficient of 0.94, as calculated by Attanasio et al. (2012) from the detrended deflated UKHPI series. Because in the baseline model prices are taken as deterministic and constant, the shock is modelled as an unexpected fall in prices at a given age. After the change, the path of prices back to the original level is taken as deterministic by the households.

The top panels of Figure 11 report the first-period (i.e. immediate) response of the female employment rate to the wage shock broken down by the LTV position for each age group. The circle points represent the case in which prices do not change and are equal to the plots from Figure 10. The diamond points represent the case in which house prices fall. The response of employment for low-LTV households is almost unchanged at all ages between the two scenarios, with only a minor attenuation. For high-LTV households, however, the difference is substantial. The fall in employment for this group is attenuated by 2 to 3 percentage points at most ages. This result implies that high-LTV households are more strongly affected by the change in house prices. While both groups receive a negative wealth shock, the sudden tightening of the LTV limit only affects the highly leveraged ones.

The same reasoning underpins the results for the case of a shock to male earnings, shown in the bottom panels of Figure 11. The addition of the price drop only modestly amplifies the response of low-LTV households but greatly amplifies the increase in employment for high-LTV females throughout the female's active years. Once again this result originates from the heterogeneous effect of the tightening of the LTV caused by the fall in house prices: only high-LTV households experience a significant reduction in their ability to use housing as collateral.

1.8 Empirical Analysis

In this section, I assess whether stylized facts from the data are consistent with the model's main predictions on the relationship between household leverage and labor supply of secondary earners.

I exploit the panel dimension of the BHPS and use a linear probability model to analyze in reduced form the relationship between the secondary earner's labor decision and the household's financial position, controlling for other variables and age trends. To isolate the role of secondary labor, I focus on the primary earner's LTI (i.e. the p-LTI) rather then the combined income LTI. Following the approach of Del Río and Young (2008), I create a set of dummy variables to classify the household's leverage position as either not having any debt or as belonging to a given quintile from the distribution of the LTV and p-LTI ratios, respectively.[33]

Similar to Disney and Gathergood (2013), I run a linear probability model where a dummy for the secondary worker's employment is regressed on the set of LTV and p-LTI quintile dummies, as well as individual, year, and region fixed effects. Other

33 Employing a set of dummies over the distribution is useful to allow for the effect of the two debt ratios to be nonlinear without imposing any functional form.

controls include a quadratic in age, number of children, educational level of both the secondary earner and the head, a dummy for being a renter, and the log of primary household income.[34] The linear equation for household i at time t has the following form:

$$Emp_{it} = \alpha + \beta' X_{it} + \sum_{j=1}^{5} \gamma_j \, q_{ijt}^{LTV} + \sum_{j=1}^{5} \delta_j \, q_{ijt}^{p-LTI} + \eta_i + \zeta_t + \epsilon_{it},$$

where Emp_{it} is a dummy variable for being employed, X_{it} is a vector of household controls, q_{ijt}^{LTV} and q_{ijt}^{p-LTI} are dummies for the LTV and p-LTI quintiles, η_i and ζ_t are household and wave fixed effects, and ϵ_{it} is an error term. The coefficients $\{\gamma_j\}_{j=1}^{5}$ and $\{\delta_j\}_{j=1}^{5}$ can be interpreted as the difference in the employment rate of workers in the j^{th} quintile of LTV and p-LTI, respectively, compared to outright homeowners. Furthermore, because in a given period the amount of outstanding mortgage is an endogenous decision, I use the lag of the mortgage value to compute the LTV and p-LTI dummies. The use of lags can be interpreted as beginning-of-period state variables before decisions on consumption and labor are made. The timing is thus consistent with the fact that all information on income and labor status relates to the period between two waves of the survey.[35] I estimate the regression using the same waves as Secton 1.2, which include the years 2001-2006.[36]

The first column of Table 8 shows the estimated coefficients of the baseline specification. The coefficients on the dummy variables for the LTV quintiles are moderately declining in value from the first to the last. The coefficient for the fourth quintile is statistically significant at the 10 percent level and the last quntile, with a value of -0.077, is significant at the 5 percent level. Meanwhile, the dummies for the p-LTI are positive and increase monotonically from the first to the last quintile, achieving statistical significance from the second quintile onwards. Everything else constant, the likelihood of employment for a secondary worker in a household where the main earner's LTI is in the top quintile is almost 10 percentage points higher than in a household with no outstanding mortgage. Figure 12 plots the coefficients from

34 Primary household income includes income from primary labor and all other non-labor earnings.

35 Unfortunately, the BHPS provides information on non-mortgage debt and on financial savings only for the years 2000 and 2005. Using the 2000 wave of the survey shows that non-housing debt is infrequent and does not constitute a large part of total liabilities. Financial savings, however, do constitute an important part of households' portfolio and, if available, it would be important to control for them.

36 Since I am using the lag of the outstanding mortgage in the right-hand side, the 2001 wave is excluded.

Column (1) to provide graphic evidence for the different relations of the LTV and the p-LTI with the probability of secondary earners being employed.

Table 8: Regression of employment of secondary earner on quintiles of LTV and primary earner's LTI, controlling for household characteristics.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
LTV							
LTV < 20th pctile	-0.008	-0.006	-0.012	0.003	0.004	-0.009	-0.026
	(0.024)	(0.024)	(0.024)	(0.027)	(0.027)	(0.024)	(0.024)
20th < LTV < 40th pctile	-0.018	-0.015	-0.023	-0.015	-0.015	-0.020	-0.023
	(0.026)	(0.026)	(0.026)	(0.029)	(0.029)	(0.026)	(0.027)
40th < LTV < 60th pctile	-0.035	-0.032	-0.034	-0.029	-0.028	-0.034	-0.037
	(0.028)	(0.028)	(0.029)	(0.031)	(0.031)	(0.028)	(0.029)
60th < LTV < 80th pctile	-0.053*	-0.049	-0.051	-0.048	-0.047	-0.056*	-0.060*
	(0.030)	(0.030)	(0.031)	(0.034)	(0.034)	(0.030)	(0.032)
80th < LTV	-0.077**	-0.073**	-0.071**	-0.070**	-0.070*	-0.080**	-0.079**
	(0.032)	(0.032)	(0.033)	(0.036)	(0.036)	(0.033)	(0.035)
p-LTI							
p-LTI < 20th pctile	0.038	0.040*	0.039	0.030	0.031	0.037	0.033
	(0.024)	(0.024)	(0.024)	(0.026)	(0.027)	(0.024)	(0.024)
20th < p-LTI < 40th pctile	0.065**	0.066**	0.066**	0.058**	0.059**	0.067**	0.049*
	(0.025)	(0.026)	(0.026)	(0.029)	(0.029)	(0.026)	(0.026)
40th < p-LTI < 60th pctile	0.076**	0.075**	0.083**	0.068**	0.069**	0.076**	0.053*
	(0.028)	(0.029)	(0.029)	(0.032)	(0.032)	(0.029)	(0.029)
60th < p-LTI < 80th pctile	0.093**	0.092**	0.094**	0.090**	0.090**	0.100**	0.068**
	(0.031)	(0.031)	(0.032)	(0.035)	(0.035)	(0.032)	(0.033)
80th < p-LTI	0.101**	0.098**	0.110**	0.091**	0.091**	0.105**	0.081**
	(0.033)	(0.033)	(0.034)	(0.037)	(0.038)	(0.035)	(0.036)
Renter	0.023	0.026	0.031	0.019	0.019	0.026	0.030
	(0.028)	(0.026)	(0.028)	(0.032)	(0.032)	(0.029)	(0.039)
No. children in the HH	-0.086***	-0.087***	-0.083***	-0.090***	-0.092***	-0.097***	-0.093***
	(0.014)	(0.014)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Log other HH income	-0.006	-0.007	-0.005	-0.003	-0.004	-0.004	-0.004
	(0.005)	(0.005)	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)
Constant	0.027	0.323	-0.369	-0.093	-0.003	-0.157	-0.096
	(0.695)	(0.662)	(0.853)	(0.680)	(0.927)	(0.689)	(0.736)
Observations	8,416	8,416	8,133	6,952	6,885	7,789	7,063
No. of Households	2,425	2,425	2,394	2,077	2,048	2,159	1,976
R^2	0.033	0.027	0.033	0.030	0.030	0.034	0.032
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	Yes		Yes	Yes	Yes	Yes	Yes
Age quadratic	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes		Yes	Yes	Yes	Yes	Yes
Local house prices			Yes				
Sample				Head in lab. force	Head in lab. force, no LTU	Females	Married

Note. The sample includes secondary earners from married and cohabiting couples aged 23 to 65 for the years 2001-2006.
*** 1 percent ** 5 percent * 10 percent significance level.
Robust standard errors in parenthesis.

The following columns of Table 8 assess the sensitivity of the results to the inclusion of further controls and restrictions on the sample. Column (2) uses a more parsimonious set of controls, which excludes education dummies and region fixed effects, as both have very low variation over time within households. Column (3) controls for average house prices at the local authority level. Columns (4) and (5) focus on households in which the primary earner is in the labor force and is not in long-term unemployment (LTU), defined as a year or more without a job. Columns (6) and (7) restrict the sample to females and married couples, respectively. The main result on the LTV and p-LTI coefficients holds through all these specifications.

In Appendix 1.F I also check the sensitivity of the results to different selection criteria for the secondary earner. First, the definition of household head earner may depend on exogenous circumstances, such as permanent job loss or adverse health conditions. Therefore, I consider an alternative sample in which the primary worker is identified based on the household head in 2001 and hence at the beginning of the time period used for the regression. Second, since females account for 90 percent of the secondary earner group, I estimate the regression on the sample of all females regardless of whether they are classified as head or not.[37] Tables 1.F.1 and 1.F.2, show that the results are robust to the first change in sample. In the female-only sample the LTV dummies, despite showing a downward trend, do not achieve statistical significance. This result suggests that there is a difference in the behavior of female workers with respect to leverage depending on whether they act as secondary workers or heads of household. Finally, the results hold in an extended time period by using the years 1998-2008 of the BHPS (Table 1.F.3).[38]

Overall, this regression analysis shows that the main prediction of the model regarding the relationship between employment and the two different measures of leverage is present empirically. The relationships, however, are qualitatively milder in the data. In particular, the relationship with the LTV is tenuous until the top quintile. The strength of the relationship in the model comes from the simplified assumptions that abstract from other important factors that may be present in the data. A primary candidate for explaining the milder empirical relationship of the LTV is house price risk, which exposes high-LTV households to fluctuations in the LTV-based borrowing constraints. This source of uncertainty over the ability to borrow would induce a stronger precautionary behavior and in turn motivate a higher labor supply.

Table 1.F.4 in Appendix 1.F explores whether this reduced-form association acts

37 Note that these samples differ in terms of how the age restrictions are satisfied and the latter subsample excludes couples formed after 2001. Hence the sample size differs compared to the baseline specification.

38 Although not reported in this paper, the results are also robust to using labor force participation, rather than employment, as the dependent variable.

Figure 12: Regression coefficients of female employment on LTV and p-LTI quintiles.

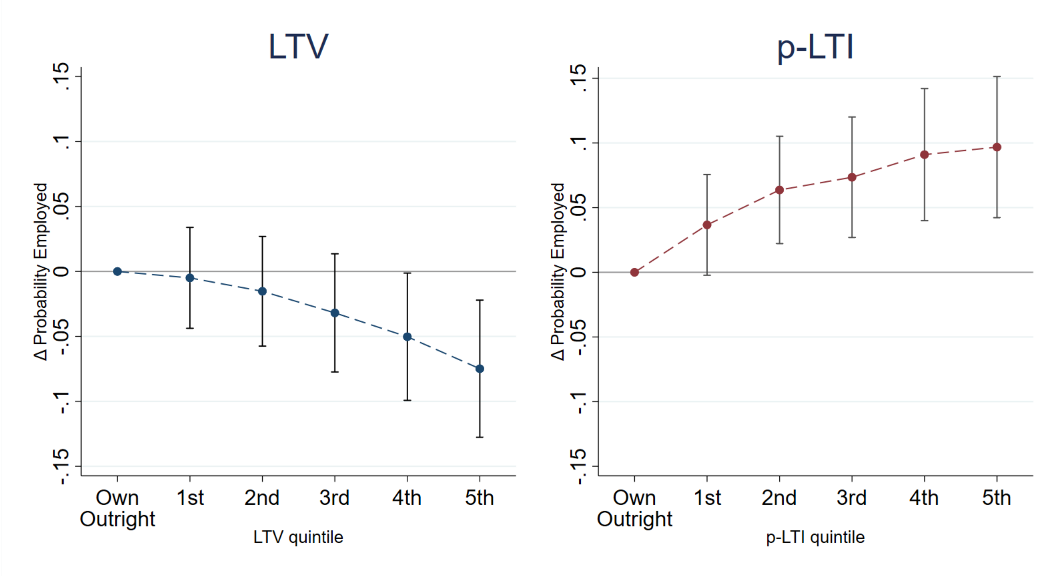

Note. The sample includes secondary earners from married and cohabiting couples aged 23 to 65 for the years 2001-2006. The bars represent the 90 percent confidence interval.

only through the extensive margin (i.e. employment) or also the intensive one (i.e. number of hours worked conditional on being employed).[39] I inspect the hours margin through a set of alternative specifications. Using the same variables as the baseline model, Column (1) uses the number of usual weekly hours worked as the dependent variable, zero for non-employed individuals. In this case, the result from the first specification holds: the p-LTI quintile dummies are positive, significant, and increasing in value while the LTV ones are negative and show a generally decreasing trend, although insignificant. Column (2) focuses only on employed females and uses the log of hours worked as the dependent variable. Keeping in mind the potential bias due to unobserved selection through the endogenous employment decision, this specification shows no significant relation between hours worked and LTI and LTV ratios. In Columns (3) and (4), the use of the inverse hyperbolic sine transformation, as suggested by Burbidge et al. (1988), confirms that the main channel of action is the extensive one.[40] An alternative approach, proposed by Blau and Kahn (2007), is to employ a Tobit model where hours are considered to be left-censored at zero.

39 Specifying a latent-variable selection model, such as a Heckman-type correction, would be the most appropriate method to study the latter. However, this approach is not implementable in the context of panel data due to the incidental parameter problem. See Wooldridge (2010) for a detailed explanation.

40 The inverse hyperbolic sign transformation has very similar properties to the natural log in terms of its approximative interpreation of percentage changes. Unlike the natural log, however, this function has a defined value at 0, which allows to include workers who are not employed in the regression.

The results from this specification, presented in Column (5), also show no significant mechanism through the intensive margin.[41]

A second prediction of the model is that high levels of p-LTI are caused by falls in primary earnings, either temporary or permanent. I take a fully nonparametric approach to check this prediction in the data. Figure 13 shows the kernel density distribution of year-on-year percentage changes in household income excluding secondary labor earnings.[42] The left panel focuses on the LTV ratio. The dashed red line represents the nonparametric distribution of income growth for households with a LTV in the bottom 80 percent of the distribution. The solid blue line reports the distribution for households with a LTV in the top 20 percent. The two distributions do not show any particular difference, indicating that households with a high LTV are not more likely than the rest to have experienced negative or positive income changes. The right panel replicates the same exercise for the p-LTI. The difference in the income growth distribution of households in the bottom 80 and top 20 percent is very pronounced. The latter is negatively skewed, with a large mass below zero, implying that many households with high p-LTI's have experienced a fall in either main earnings or non-labor income in the current year. While only 31 percent of households in the first four quintiles had a negative income change, the fraction rises to 50 percent for those in the top quintile. When considering large falls of 15 percent or more, the comparison is 11 and 32. As income is the denominator of the p-LTI ratio, by construction a fall in income yields a rise in the p-LTI, which is the reason for many p-LTI's in the top quintile. This result suggests that the "added worker effect" is to a large extent the driver of high secondary worker employment for the top quintile.[43]

Overall, the above empirical analysis shows that stylized facts from the data are consistent with the model's main predictions on the relationship between leverage and the labor supply of secondary earners. In particular, the labor supply of secondary earners shows a negative association with LTV levels and a positive one with the p-LTI ratio. The magnitude of this relationship, however, is more moderate for the former measure. Moreover, high p-LTI values are partly accounted for by falls in the

41 The estimation uses the "trimmed" least squares method devised by Honoré (1992) to adjust the Tobit model to panel data with a short time dimension.

42 More precisely, these are wave-on-wave changes. For most households, the interviews in the two waves take place exactly 12 months apart. However, in some cases there are a few months of discrepancy.

43 Interestingly, Figure 1.A.3 in the appendix suggests that some of these falls are temporary, as the distribution of income growth for the following wave is positively skewed.

Figure 13: Nonparametric distribution of primary income growth by LTV and p-LTI.

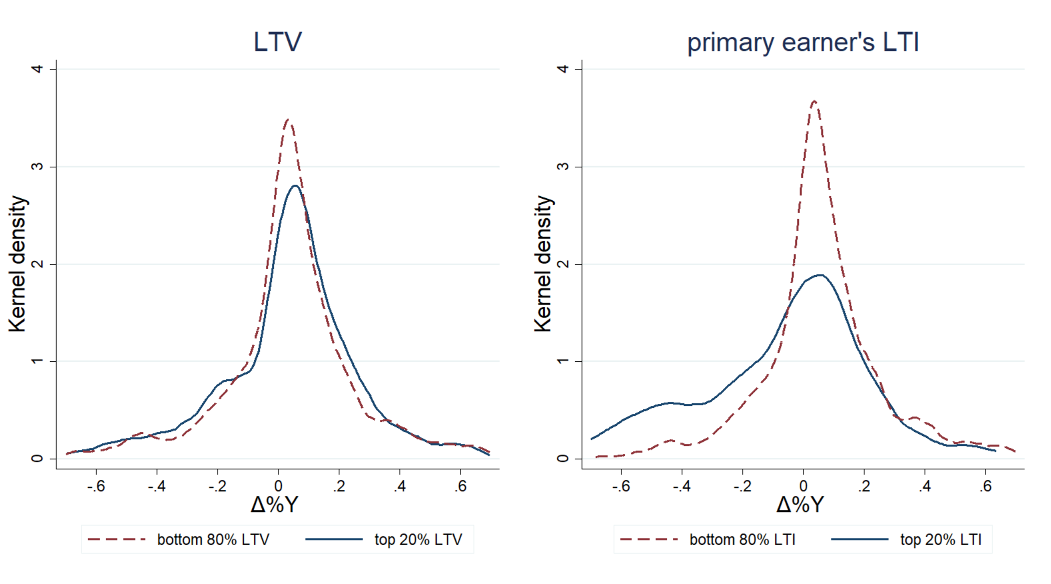

Note. The sample includes secondary earners from married and cohabiting couples aged 23 to 65 for the years 2001-2006. The non-parametric densities are computed using the Epanechnikov kernel on 50 points between -0.7 and 0.7.

earnings of the primary earner.

1.9 Conclusion

This paper analyzed the relationship between housing debt, leverage-based borrowing constraints, and labor supply. Through a life cycle model of a two-worker household, I showed how leverage limits on debt are important determinants of labor supply decisions at the extensive margin. The model implies that the labor supply of the secondary earner (which is assumed to be female) is negatively correlated with the household's LTV and positively correlated with the primary earner's LTI. The results originate from the interaction of the leverage limits with two competing motives within the household: the incentive to substitute away from costly female labor supply -which is reflected by high LTV ratios- and the need to use female labor to repay debt when male earnings are low -which is associated to high p-LTI ratios.

The interaction between labor choices and leverage constraints is important because the latter are used by policymakers for macroprudential reasons or to influence homeownership rates. Using the model, I showed that changes in the LTV and LTI constraints interact with labor supply decisions very differently and have heterogeneous welfare effects across households. The LTI limit has a larger spillover onto the aggregate profile of labor supply over the life cycle, as it more directly interacts with the returns to work for some households.

The LTV- and LTI-based constraints entail large heterogeneity in the response of labor supply to income shocks depending on the household's leverage position. At all ages, high-LTV households have a more elastic extensive margin of labor supply compared to households with low leverage. High-LTV households decrease employment more pronouncedly after a fall in female wages and increase it the most after a fall in male earnings. Finally, unexpected house price drops also play an important role for those who are highly leveraged. The fall in prices tightens the LTV limit, forcing females in high-LTV households to supply more labor than desired.

The analysis in this paper leaves scope for further research in two main directions. First, a more sophisticated life cycle framework should consider three key features: house price risk, house price growth expectations, endogenous fertility choices, and bankruptcy. The first two topics deal with housing as an investment good but ultimately also affect labor supply decisions. Endogenous fertility choices would deal with the empirical correlation in fertility and housing preferences: high utility from housing may be due to intentions to have children and would in turn be associated with high labor supply costs. Regarding bankruptcy and foreclosure decisions, recent research has shown that regulatory frameworks are crucial for both macroeconomic fluctuations (Mitman, 2016) and for the relationship between housing and the labor market (Hsu et al., 2017). The second possible direction would simplify the model's dynamics to incorporate them into a general equilibrium framework for both house prices and wages. Computing the stationary distributions needed to obtain equilibrium prices is extremely challenging under the strong nonlinearity and multiple discontinuities of the policy functions. However, a general equilibrium model could address important questions such as the role played by high housing debt in the sustained period of low wages and high employment experienced by the UK after the Great Recession.

1.A Additional descriptive analysis on the BHPS

1.A.1 Additional descriptive empirical analysis

Table 1.A.1: Descriptive statistics for household head and secondary earner in the 2001 wave of the BHPS for couples where the secondary earner is aged 23-40.

	Head of Household	**Not Head of Household**
Sex		
% Female	12.6	87.4
Labor Force Status		
% Employed	90.2	75.6
% Unemployed	2.1	2.3
% Maternity leave	0.1	2.4
% Family care	3	15.9
% Other LF status	4.6	3.7
Child Care*		
% Sole responsible for child care	8.3	58
% Sole carer for ill children	14.6	69

Note. Summary statistics produced using couples in the year-2001 wave of the BHPS where the secondary earner is between the ages of 23 and 40.
* Child care questions refer only to the subset of couples with one or more children under the age of 12.

Figure 1.A.1: Life cycle profiles by 5-year cohorts for the years 1991-2008.

Note. Source: BHPS waves 1991-2008. The sample includes married couples and those in cohabiting partnerships where the secondary earner is aged 23 to 65. Households are divided into 5-year cohorts based on the secondary earner's year of birth. In the rightmost panel of the first row, house values are deflated to 2001 prices using the UKHPI series.

Figure 1.A.2: Age effects for homeownership, employment, average house value, percentage of mortgagors, average LTV, LTI, and p-LTI using the regression approach by Deaton and Paxson (1993).

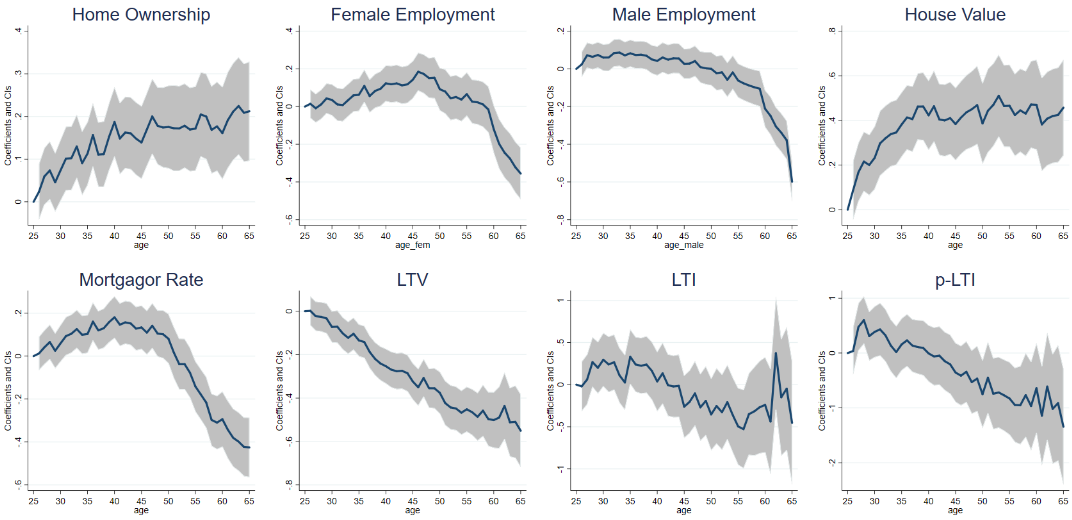

Note. Source: BHPS waves 2001-2006. Sample includes married couples and those in permanent partnerships where the secondary earner is aged 23 to 65. The age effects are obtained using a linear regression on age dummies, 10-year cohort dummies, and wave dummies. These last are constrained to sum to 0 as in Deaton and Paxson (1993).

1.A.2 Income profiles and idiosyncratic shocks

I follow Bottazzi et al. (2007) in using the estimation suggested by Blundell et al. (2008). The identification assumes that log wages for individual i at period j follow the process:

$$\log W_{i,j} = H_{i,j} + Z_{i,j} + u_{i,j},$$

where $H_{i,j}$ is a deterministic component based on observable characteristics, $Z_{i,j}$ is a persistent idiosyncratic component, and $u_{i,j}$ is a transitory one-period component. The idiosyncratic process is a martingale: $Z_{i,j} = Z_{i,j-1} + \epsilon_{i,j}$. Letting $w_{i,j} = \log W_{i,j} - H_{i,j}$ be the unexplained part of earnings, then the unexplained earnings growth is given by $\Delta w_{i,j} = \Delta u_{i,j} + \epsilon_{i,j}$. Hence:

$$cov(\Delta w_{i,j}, \Delta w_{i,j+s}) = \begin{cases} var(\epsilon_{i,j}) + var(\Delta u_{i,j}) & \text{if } s = 0 \\ cov(\Delta u_{i,j}, \Delta u_{i,j+s}) & \text{if } s \neq 0 \end{cases}$$

The assumptions of serially uncorrelated transitory shocks, in particular, means that the variance of the innovations to the idiosyncratic component can be identified by the single moment

$$var(\epsilon_{i,j}) = cov(\Delta w_{i,j}, \Delta w_{i,j-1} + \Delta w_{i,j} + \Delta w_{i,j+1}).$$

First, I obtain the deterministic component of wages by regressing the log hourly wage on a set of control variables. The chosen variables for males are an age quadratic, dummies for government administrative region, education level of the head of the household, five-year cohort dummies, household size, and number of children. For

Figure 1.A.3: Nonparametric distribution of income growth in wave $t + 1$ by LTV and p-LTI for wave t.

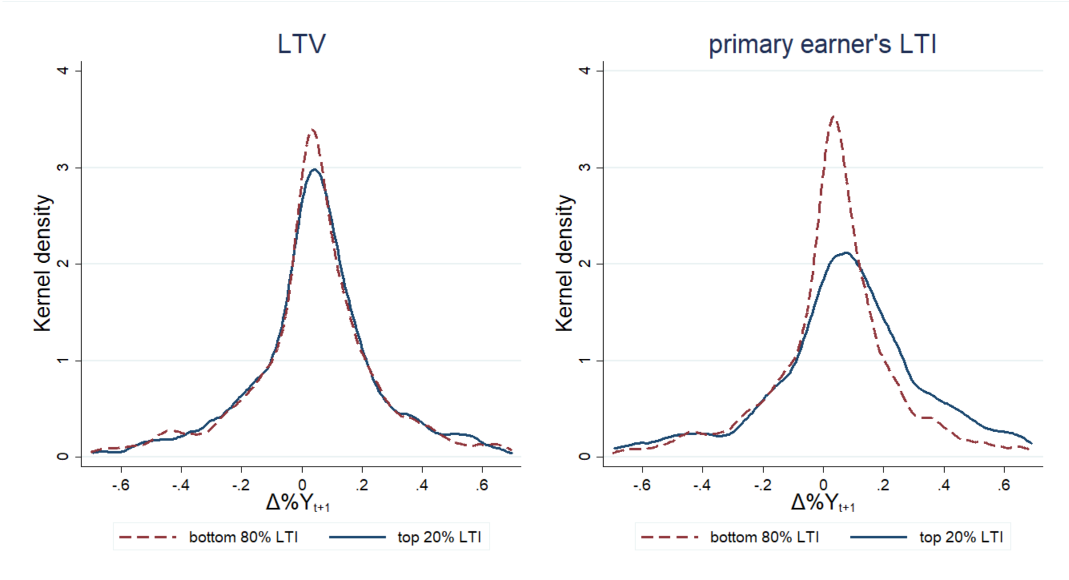

Note. The non-parametric densities are computed using the Epanechnikov kernel on 50 points between -0.7 and 0.7.

females, I account for selection in observing wages based on the decision to work. I therefore use a Heckman-correction 2-step estimation where the selection is based on "other household income", being married, and number of children. The age quadratic is used to calibrate the deterministic component of income and wages for each sex. The residuals from the regressions are used to estimate the parameters of the idiosyncratic income process.

Because the parameter is identified through three periods, at least three years of data are required. The sample of interest for the model is 2001-2006. However, to improve the estimation by having more data and by spanning a time period that includes at least one recession and one recovery, I use all the BHPS years from 1991 to 2008.

1.A.3 Number of children and equivalization coefficient

Figure 1.A.4 shows the paths for the age-varying coefficient $\chi(j)$, which is computed through a cubic polynomial for average household size adjusted by the OECD equivalization coefficient. First, I estimate a cubic polynomial for the average number of children by age of the secondary earner using couples in the BHPS. The equivalization coefficient is based on the OECD scale, where the first adult is given a weight of 1, the second of .7 and each child a weight of 0.5. Hence, $\chi(j)$ is equal to 1.7 plus 0.5 times the function for the average number of children at each age from the BHPS.

1.A.4 Initial distribution of assets

In the years 1995, 2000, and 2005 the BHPS included questions regarding savings and investments. I use the 2000 wave to obtain the empirical distribution of net worth for couples within the ages of 23 and 27.

Figure 1.A.4: Housing equivalization coefficient based on avergae household size from the BHPS.

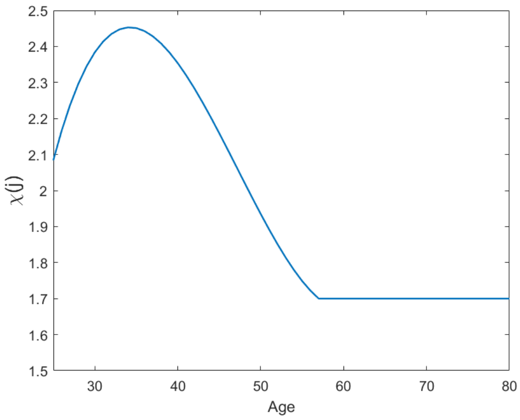

The sample includes all individuals who report being married or living as a couple. Those for which the spouse is present in the survey are matched with their partner to compute all combined measures of savings and investments. Net worth is computed as (joint) savings and investments plus the value of their house minus the outstanding mortgage balance for homeowners and outstanding unsecured debt. As discussed by Banks et al. (2002), the savings and investment variables in the BHPS present some important limitations at the household level. In many cases the two members of a couple report having joint holdings that are inconsistent with each other, or report that only part of their savings are jointly held but do not specify the precise amount. Furthermore, many households report no savings, or reply that they do not know their savings level. Given that these cases account for a relatively small fraction of my sample, I do not implement the kind of "hot-deck" imputation and case-by-case assumptions that are discussed by Banks et al. (2002). For joint savings I take the average of the values reported by the two members of the couple.

It is important to note that about half of the couples in the sample report no savings. However, a significant portion of these still own a house or have a mortgage. It is therefore important to use these variables, and not just liquid savings to calibrate the initial distribution of assets. I assume that all households start with no housing. However, given that initial assets are calibrated to include housing assets, several households purchase a house within the first period of the model. A large portion of households report negative values of total net worth, although of very small size in the majority of cases. Since in the model households cannot have negative net worth, I censor the initial distribution of assets in the simulation at 0, assigning this value to the empirical fraction of households that have zero or negative net worth.

1.B Computational solution and simulation

The nature of the borrowing constraint function and the discontinuity of the labor policy decision create several points of non-differentiability in the value function. Computationally fast algorithms such as the Endogenous Gridpoint Method adapted for discrete-continuous choices are therefore unfeasible. I hence opt for the slower but more robust method of computing the value function over a fixed discretized grid within the state space. The computation starts by solving the problem for the last period and then moving backwards one period at the time. In order to solve the model over a square space I follow Bajari et al. (2013) in transforming the state space for financial assets a_t. Specifically, I define $m_j(h_{j-1}) = a_j + \lambda_h p h_{j-1}$, $m_{j+1}(h_j) = a_{j+1} + \lambda_h p h_j$, so that the state spaces for the variables m_t and m_{t+1} always have 0 as a lower bound. Hence, for every combination of the finite values of h_{j-1} and h_j, the maximization over $m_{j+1}|m_j$ can be solved on a fixed two-dimensional grid. The corresponding values of a_j and a_{j+1} can then be recovered. In each period, for a given h_{j-1} I first solve the optimal level of consumption, savings, and work conditional on each level of h_j. I then maximize over the possible choices for h_j. The optimal choice of labor, conditional on h_{j-1} and h_j, is found through two-state budgeting. Conditional on $[a_t, a_{t+1}, h_{t-1}, h_t, y_t, w_t^f]$, n_t^f and c_t need to be solved simultaneously through the budget constraint and the intra-temporal first-order condition that equalizes the Marginal Rate of Substitution of leisure and consumption to the wage:

$$c_j + n_j^f w_j^f = (1+r)a_j - a_{j+1} + p(h_{j-1} - h_j) - \xi(n_j^f, j) - \Phi(h_{j-1}, h_j) + y_j \tag{5}$$

$$\frac{\partial U(c_j, h_j, n_j^f, j)}{\partial n_j^f} = w_j^f \frac{\partial U(c_j, h_j, n_j^f, j)}{\partial c_j} \tag{6}$$

In order to do this in a computationally efficient manner, for a given w_j^f I first numerically approximate c and n^f as functions of the non-labor resources through a tenth-order Chebyshev polynomial. Then, during the value function computation, I use the linear approximation to compute the values of c and n^f for a given value of the right-hand-side of (5). In this way, I obtain the optimal solution conditional on the female being in the labor force. I then also compute the case in which the female does not work and take the maximum of the two options for each point in the state space.

For the computation, I use 200 grid points to discretize the grid of financial assets. I use the method described inTauchen (1986) for the discretization of the earnings process for men and women. Since the earnings processes are age-dependent, the discretization needs to be done separately for each age. I use 10 points for the male's permanent earnings and 15 for female wages. For the male, the grid space is then doubled to include the possibility of unemployment.

The simulations to compute the life cycle averages are done by simulating the joint Markov process for male and female earnings 10,000 times and then computing the resulting paths for the choice variables. The starting distribution for financial assets is based on the empirical distribution of net worth in the 2000 wave of the BHPS for households aged 23-27.

The labor supply responses with respect to shocks to female wage and male income for a given age are computed by solving the policy function once more imposing an unexpected permanent one-standard deviation fall in all earnings at the respective age. The difference between the two average life cycle profiles constitutes the IRF.

1.C Details of the computation of welfare losses

This section describes the computation of welfare costs for income shocks. However, the same principle applies to the counterfactual calibrations of the LTV and LTI limits. Let j_s be the age at which a given shock occurs. For a household i, given its initial assets a_0^i and the earnings processes $\{y_j, w_j^f\}_{j=0}^T$, the resulting paths of consumption, housing, and labor supply are enough to compute its discounted lifetime utility from age j^s: $\mathbb{U}_j^i = \sum_{j^s}^T \beta^{j-j^s} U(c_j^i, n_j^i, h_j^i)$. A similar value $\overline{\mathbb{U}}_j^i = \sum_{j^s}^T \beta^{j-j^s} U(\overline{c}_j^i, \overline{n}_j^i, \overline{h}_j^i)$ can be calculated for the counterfactual scenario with income shocks. The method aims to find the value of τ^i such that $\sum_{j^s}^T \beta^{j-j^s} U(c_j^i(1+\tau^i), n_j^i, h_j^i) = \overline{\mathbb{U}}_j^i$. Since the compensation is found using *ex post* values of the endogenous variables specific to each household, the method does not require an iterative procedure as in Attanasio et al. (2012).

Using the same set of 10,000 simulations and their counter-factual cases for the various shocks and ages, I compute the average τ for the aggregate economy and the various tenure groups: renters, low-LTV, high-LTV, low-p-LTI, and high-p-LTI.

The similar intuition applies to the counterfactual calibrations of the LTV and LTI limits. The relevant "age of the shock" is the first period of the life cycle.

1.D Additional tables and figures from model

Figure 1.D.1: Calibrated labor supply cost

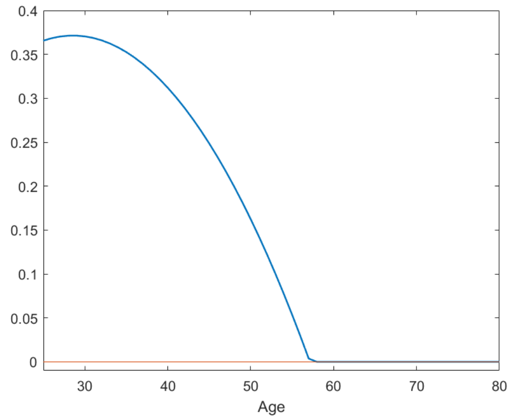

Note. The figure shows the quadratic function $\xi(j)$, where the coefficients ξ_1 and ξ_2 are calibrated as explained in Section 1.4.

Table 1.D.1: Empirical targets and simulated moments.

Moment	**Target**	**Model**
Average homeownership rate, age 30-55	.82	.82
Average home value, age 30-55	5.4	5.4
Female employment rate, age 25-29	.71	.70
Female employment rate, age 30-34	.71	.71
Female employment rate, age 35-39	.75	.72
Female employment rate, age 40-44	.78	.76
Female employment rate, age 45-49	.82	.82
Female average hours 30-54	35	35

Note. The table reports the empirical moments that are targeted by the internal calibration and the corresponding moments from the model obtained from 10,000 simulations of the individual income processes starting from the calibrated initial distribution of wealth.

Figure 1.D.2: Employment rate by LTV ratio across age groups and p-LTI above or below median.

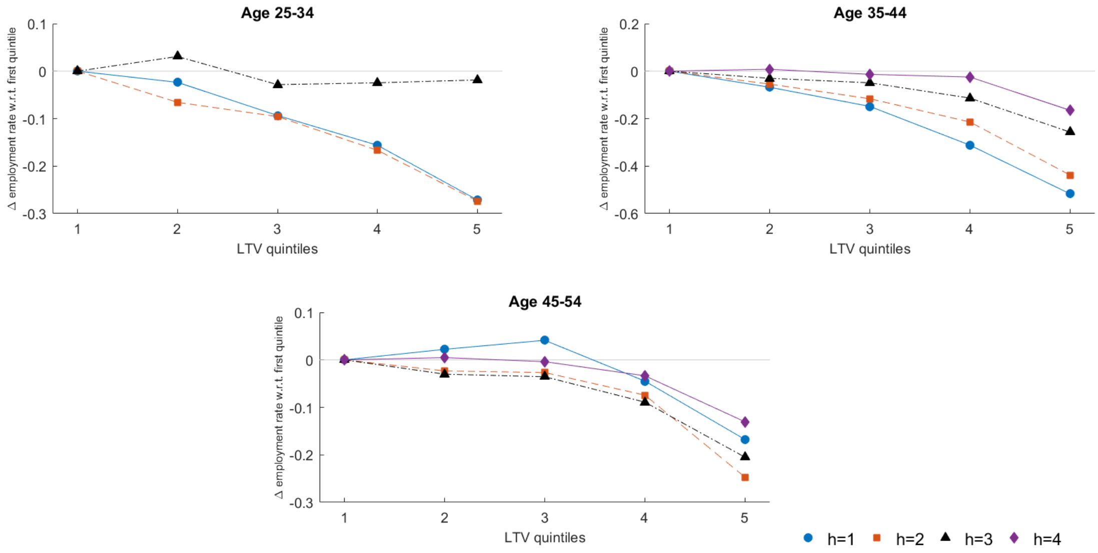

Note. Each plot is produced by grouping households into 10-year age groups. For each group, each point represents the employment rate (y-axis) for a quintile of the LTV (x-axis), an age group, and a size of housing assets. In this way, the plot shows how a variable tends to change across the simulated distribution of the x-variable. The plots are produced using 10,000 simulations of individual income shocks starting from the calibrated initial distribution of net worth.

Figure 1.D.3: Employment rate by p-LTI ratio across age groups and LTV above or below median.

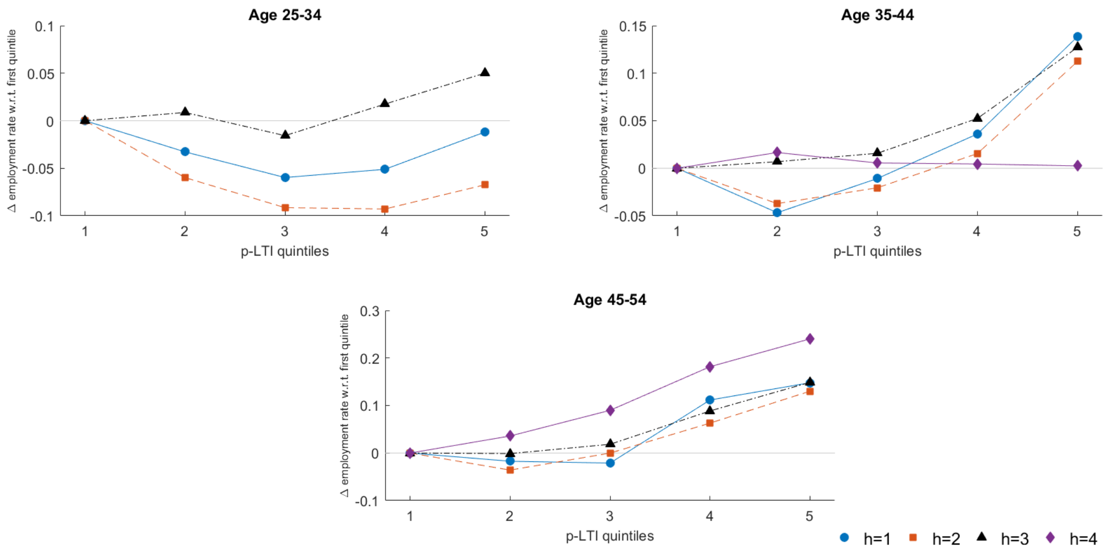

Note. Each plot is produced by grouping households into 10-year age groups. For each group, each point represents the employment rate (y-axis) for a quintile of the p-LTI (x-axis), an age group, and a size of housing assets. In this way, the plot shows how a variable tends to change across the simulated distribution of the x-variable. . In this way, the plot shows how a variable tends to change across the simulated distribution of the x-variable. The plots are produced using 10,000 simulations of individual income processes starting from the calibrated initial distribution of net worth.

Figure 1.D.4: Employment rate by net worth across age groups and house size.

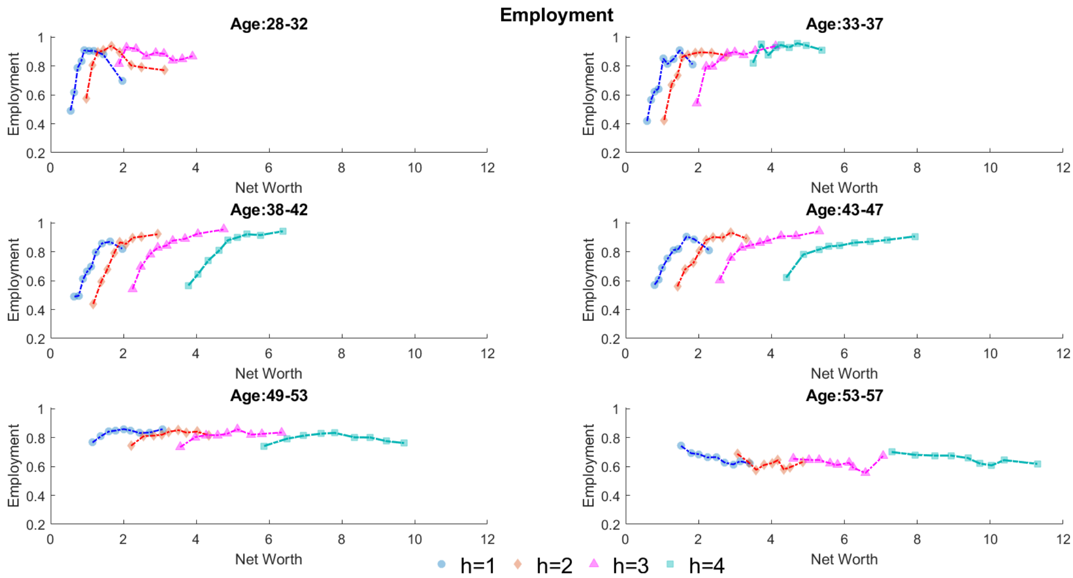

Note. Each plot is produced by grouping together households in five-year group and divide them into owners of each house size. For each group, each point represents the employment rate on the y-axis for a given decile of the variable on the x-axis. In this way, the plot shows how a variable tends to change across the simulated distribution of the x-variable. The plots are produced using 10,000 simulations of individual income processes starting from the calibrated initial distribution of net worth.

Figure 1.D.5: Surface plots of employment, LTV, and p-LTI across 10-year age groups.

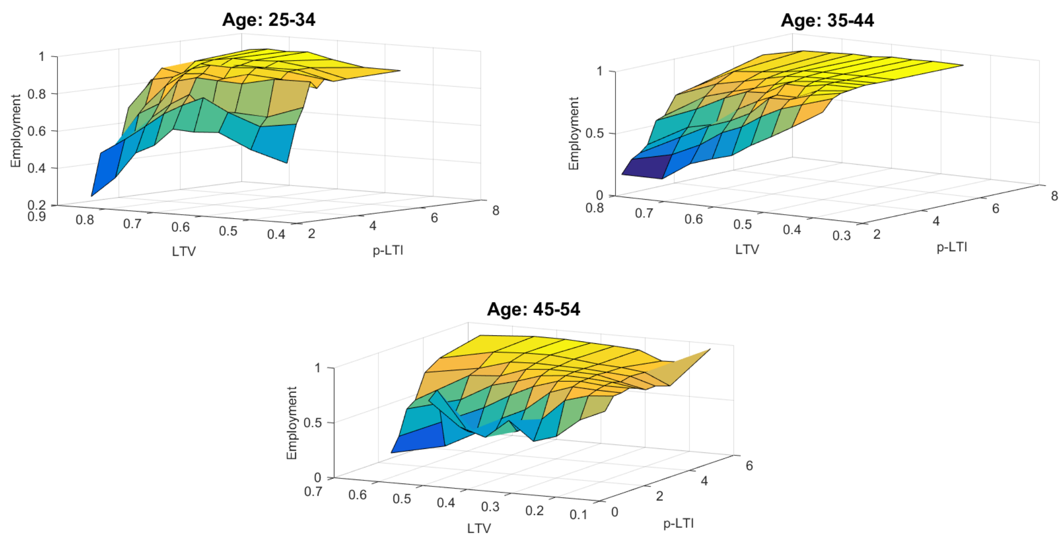

Note. All results are produced using the same set of 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth.

Figure 1.D.6: Fraction of high-LTV homeowners at different ages in the baseline model.

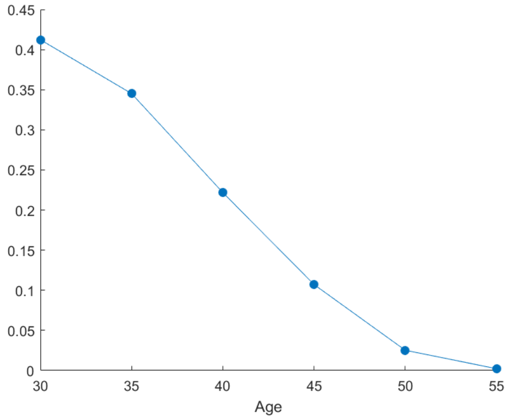

Note. The points correspond to the ages at which the initial female employment response is computed in Figure 10. The series is computed using 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth.

Figure 1.D.7: Change in the fraction of high p-LTI homeowners at different ages and by low- and high-LTV group after a one-standard deviation permanent fall in male earnings.

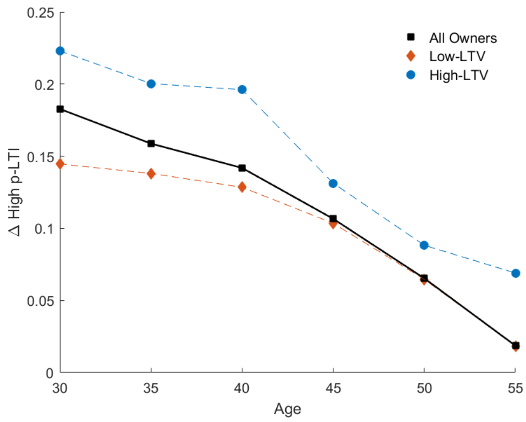

Note. The points correspond to the ages at which the initial female employment response is computed in Figure 10. The series is computed using 10,000 simulations of individual income processes and draws from the calibrated initial distribution of net worth.

1.E Decomposition of labor supply response to income shocks by the male's LTI.

Figure 1.E.1 reports the initial response of female employment to a one-standard deviation fall in either female potential wages or male earnings, broken down by p-LTI level. I choose the value of 3.5 as the threshold for high-p-LTI households, which is close ot the 80$^{\text{th}}$ percentile of the unconditional distribution for all mortgagors. At this value the household is borrowing-constrained unless the female is employed. The fall in employment after a negative shock to female earnings is slightly lower for high-p-LTI households, as they rely more heavily on female labor for debt repayment. The added worker effect is larger for the low-p-LTI group, as it includes "marginal" households who become constrained as a result of the fall in male earnings. Meanwhile, the high-p-LTI includes households who already rely on female labor and therefore their decisions are not strongly affected.

Figure 1.E.1: Contemporaneous labor supply response to female and male earnings shocks between low-p-LTI and high-p-LTI homeowners and ages.

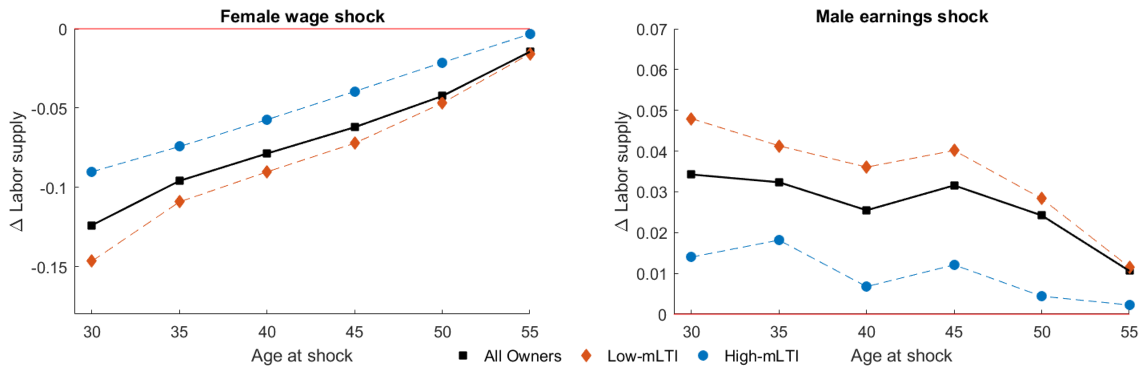

Note. Each point reports the response of the extensive margin female labor supply in percentage point (i.e. employment rate) to an unexpected one-standard deviation permanent fall in female wages or male earnings occurring a given age. Households are divided into low-p-LTI homeowners (p-LTI<3), and high-p-LTI homeowners (p-LTI≥3). The black squares report the response for all homeowners. All plots are produced using 10,000 individual simulations.

1.F Empirical robustness checks

Table 1.F.1: Regression of employment of the secondary earner on quintiles of LTV and primary earner's LTI, controlling for household characteristics and using an alternative sample in which secondary earners are selected in the year 2001 only.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
LTV							
LTV < 20th pctile	-0.014	-0.013	-0.016	-0.004	-0.003	-0.010	-0.021
	(0.023)	(0.023)	(0.023)	(0.027)	(0.027)	(0.024)	(0.023)
20th < LTV < 40th pctile	-0.030	-0.029	-0.033	-0.028	-0.029	-0.026	-0.031
	(0.025)	(0.025)	(0.026)	(0.029)	(0.030)	(0.027)	(0.026)
40th < LTV < 60th pctile	-0.037	-0.035	-0.033	-0.035	-0.035	-0.032	-0.037
	(0.026)	(0.026)	(0.027)	(0.031)	(0.031)	(0.029)	(0.027)
60th < LTV < 80th pctile	-0.061**	-0.057**	-0.057*	-0.059*	-0.059*	-0.061*	-0.062**
	(0.029)	(0.028)	(0.029)	(0.033)	(0.033)	(0.031)	(0.030)
80th < LTV	-0.073**	-0.070**	-0.070**	-0.072**	-0.072**	-0.076**	-0.078**
	(0.030)	(0.030)	(0.032)	(0.035)	(0.035)	(0.034)	(0.033)
p-LTI							
p-LTI < 20th pctile	0.030	0.031	0.026	0.020	0.021	0.028	0.028
	(0.024)	(0.023)	(0.024)	(0.026)	(0.026)	(0.025)	(0.024)
20th < p-LTI < 40th pctile	0.060**	0.060**	0.056**	0.051*	0.051*	0.060**	0.055**
	(0.024)	(0.024)	(0.024)	(0.027)	(0.028)	(0.026)	(0.025)
40th < p-LTI < 60th pctile	0.051*	0.047*	0.051*	0.041	0.041	0.048*	0.036
	(0.027)	(0.027)	(0.027)	(0.030)	(0.031)	(0.029)	(0.028)
60th < p-LTI < 80th pctile	0.072**	0.070**	0.067**	0.072**	0.073**	0.076**	0.071**
	(0.030)	(0.030)	(0.030)	(0.033)	(0.033)	(0.032)	(0.032)
80th < p-LTI	0.082**	0.078**	0.084**	0.076**	0.076**	0.083**	0.078**
	(0.031)	(0.031)	(0.031)	(0.035)	(0.035)	(0.034)	(0.033)
Renter	0.054*	0.056*	0.064**	0.043	0.048	0.052	0.078*
	(0.029)	(0.030)	(0.030)	(0.033)	(0.034)	(0.036)	(0.040)
No. children in the HH	-0.065***	-0.065***	-0.063***	-0.072***	-0.074***	-0.082***	-0.071***
	(0.013)	(0.013)	(0.013)	(0.014)	(0.014)	(0.015)	(0.014)
Log other HH income	-0.005	-0.005	-0.004	-0.002	-0.002	-0.004	-0.003
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)
Constant	0.295	0.276	0.173	0.269	0.439	0.306	0.230
	(0.697)	(0.679)	(0.829)	(0.674)	(0.938)	(0.733)	(0.753)
Observations	8,069	8,069	7,775	6,567	6,510	7,194	6,969
No. of Households	2,050	2,050	2,028	1,754	1,734	1,817	1,783
R^2	0.029	0.025	0.029	0.023	0.024	0.031	0.030
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	Yes		Yes	Yes	Yes	Yes	Yes
Age quadratic	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes		Yes	Yes	Yes	Yes	Yes
Local house prices			Yes				
Sample				Head in lab. force	Head in lab. force, no LTU	Females	Married

Sample includes secondary earners from married and cohabiting couples aged 23 to 65 selected in the year 2001.
*** 1 percent ** 5 percent * 10 percent significance level.
Robust standard errors in parenthesis.

Table 1.F.2: Regression of employment of females on quintiles of LTV and male earner's LTI, controlling for household characteristics.

	(1)	(2)	(3)	(4)	(5)	(6)
LTV						
LTV < 20th pctile	-0.007	-0.007	-0.008	0.002	0.002	-0.022
	(0.023)	(0.022)	(0.023)	(0.026)	(0.027)	(0.024)
20th < LTV < 40th pctile	-0.009	-0.009	-0.009	-0.007	-0.007	-0.013
	(0.025)	(0.025)	(0.025)	(0.029)	(0.029)	(0.027)
40th < LTV < 60th pctile	-0.020	-0.019	-0.015	-0.016	-0.016	-0.022
	(0.027)	(0.027)	(0.027)	(0.031)	(0.031)	(0.029)
60th < LTV < 80th pctile	-0.031	-0.029	-0.026	-0.026	-0.027	-0.039
	(0.029)	(0.028)	(0.030)	(0.032)	(0.033)	(0.031)
80th < LTV	-0.048	-0.046	-0.040	-0.044	-0.045	-0.049
	(0.031)	(0.031)	(0.032)	(0.035)	(0.035)	(0.034)
p-LTI						
p-LTI < 20th pctile	0.036	0.038	0.032	0.031	0.032	0.039
	(0.024)	(0.023)	(0.024)	(0.027)	(0.027)	(0.025)
20th < p-LTI < 40th pctile	0.066**	0.067**	0.062**	0.060**	0.061**	0.058**
	(0.024)	(0.025)	(0.025)	(0.028)	(0.028)	(0.025)
40th < p-LTI < 60th pctile	0.061**	0.061**	0.061**	0.056*	0.057*	0.048*
	(0.027)	(0.027)	(0.028)	(0.031)	(0.031)	(0.029)
60th < p-LTI < 80th pctile	0.095**	0.093**	0.089**	0.091**	0.092**	0.077**
	(0.030)	(0.029)	(0.030)	(0.034)	(0.034)	(0.032)
80th < p-LTI	0.100**	0.100**	0.101**	0.091**	0.092**	0.091**
	(0.031)	(0.031)	(0.032)	(0.036)	(0.036)	(0.035)
Renter	0.000	-0.002	0.004	-0.010	-0.007	0.016
	(0.037)	(0.039)	(0.039)	(0.040)	(0.040)	(0.045)
No. children in the HH	-0.085***	-0.086***	-0.084***	-0.089***	-0.093***	-0.086***
	(0.014)	(0.014)	(0.015)	(0.015)	(0.015)	(0.015)
Log other HH income	-0.007	-0.006	-0.006	-0.004	-0.005	-0.004
	(0.005)	(0.005)	(0.005)	(0.004)	(0.005)	(0.004)
Constant	0.169	0.344	-0.557	0.067	0.136	-0.150
	(0.668)	(0.651)	(0.830)	(0.663)	(0.888)	(0.741)
Observations	8,879	8,879	8,569	7,479	7,389	7,421
No. of Households	2,382	2,382	2,356	2,095	2,063	1,966
R^2	0.033	0.027	0.033	0.028	0.030	0.032
Household FE	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	Yes		Yes	Yes	Yes	Yes
Age quadratic	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes		Yes	Yes	Yes	Yes
Local house prices			Yes			
Sample				Head in lab. force	Head in lab. force, no LTU	Married

Sample includes females from married and cohabiting couples aged 23 to 65 for the years 2001-2006.
*** 1 percent ** 5 percent * 10 percent significance level.
Robust standard errors in parenthesis.

Table 1.F.3: Regression of employment of the secondary earner on quintiles of LTV and primary earner's LTI, controlling for household characteristics, and using the extended time sample 1998-2008.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
LTV							
LTV < 20th pctile	-0.002	-0.003	-0.007	-0.006	0.002	-0.002	-0.008
	(0.019)	(0.018)	(0.019)	(0.020)	(0.029)	(0.019)	(0.020)
20th < LTV < 40th pctile	-0.013	-0.014	-0.016	-0.015	-0.020	-0.014	-0.017
	(0.020)	(0.020)	(0.020)	(0.021)	(0.033)	(0.020)	(0.022)
40th < LTV < 60th pctile	-0.047**	-0.049**	-0.045**	-0.041*	-0.066*	-0.046**	-0.045**
	(0.021)	(0.021)	(0.021)	(0.022)	(0.035)	(0.021)	(0.022)
60th < LTV < 80th pctile	-0.074***	-0.076***	-0.072**	-0.066**	-0.076*	-0.073**	-0.075**
	(0.022)	(0.023)	(0.023)	(0.024)	(0.040)	(0.023)	(0.025)
80th < LTV	-0.049**	-0.053**	-0.050**	-0.034	-0.065	-0.047*	-0.042
	(0.024)	(0.024)	(0.024)	(0.025)	(0.041)	(0.024)	(0.026)
LTI							
p-LTI < 20th pctile	0.032*	0.035*	0.041**	0.037*	0.032	0.033*	0.034*
	(0.019)	(0.019)	(0.019)	(0.021)	(0.032)	(0.019)	(0.020)
20th < p-LTI < 40th pctile	0.057**	0.059**	0.064**	0.057**	0.048	0.061**	0.059**
	(0.020)	(0.020)	(0.021)	(0.022)	(0.035)	(0.021)	(0.022)
40th < p-LTI < 60th pctile	0.056**	0.058**	0.064**	0.048**	0.055	0.059**	0.054**
	(0.021)	(0.021)	(0.021)	(0.022)	(0.038)	(0.021)	(0.023)
60th < p-LTI < 80th pctile	0.080***	0.083***	0.083***	0.077***	0.087**	0.083***	0.077***
	(0.021)	(0.021)	(0.021)	(0.022)	(0.037)	(0.021)	(0.023)
80th < p-LTI	0.105***	0.111***	0.111***	0.092***	0.078*	0.108***	0.117***
	(0.023)	(0.023)	(0.023)	(0.024)	(0.042)	(0.024)	(0.025)
Renter	0.014	0.017	0.015	0.008	0.041	0.022	0.023
	(0.023)	(0.024)	(0.024)	(0.022)	(0.039)	(0.024)	(0.028)
No. children in the HH	-0.083***	-0.084***	-0.083***	-0.080***	-0.087***	-0.089***	-0.081***
	(0.008)	(0.009)	(0.009)	(0.009)	(0.015)	(0.009)	(0.009)
Log other HH income	-0.010**	-0.011**	-0.010*	-0.012**	-0.009	-0.010*	-0.011*
	(0.005)	(0.005)	(0.005)	(0.006)	(0.008)	(0.006)	(0.006)
Constant	-0.502	-0.423	-0.569	-0.056	-1.183	-0.627	-0.824
	(0.479)	(0.472)	(0.678)	(0.568)	(1.215)	(0.465)	(0.504)
Observations	16,748	16,748	16,034	13,710	6,071	15,557	14,180
No. of Households	3,085	3,085	3,022	2,625	650	2,683	2,485
R^2	0.058	0.049	0.058	0.042	0.072	0.059	0.060
Household FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region FE	Yes		Yes	Yes	Yes	Yes	Yes
Age quadratic	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes		Yes	Yes	Yes	Yes	Yes
Local house prices			Yes				
Sample				Male in lab. force no male LTU	Balanced	Females	Married

Sample includes secondary earners from married and cohabiting couples aged 23 to 65 for the years 1998-2008.
*** 1 percent ** 5 percent * 10 percent significance level.
Robust standard errors in parenthesis.

Table 1.F.4: Inspecting the extensive versus the extensive margin channel. Fixed effects regression of secondary earner employment (OLS, Probit, Logit), and hours (unconditional and conditional on working) on quintiles of primary earner's LTI and LTV, controlling for household characteristics.

	(1)	(2)	(3)	(4)	(5)
	Hours	log(hours) hours >0	IHS(Hours)	IHS(hours) hours >0	Tobit
LTV					
LTI < 20th pctile	0.471	0.029	0.035	0.029	2.045
	(0.705)	(0.027)	(0.089)	(0.027)	(1.387)
20th < LTV < 40th pctile	-0.671	-0.012	-0.038	-0.012	3.027**
	(0.792)	(0.028)	(0.095)	(0.028)	(1.180)
40th < LTV < 60th pctile	-1.506*	-0.044	-0.130	-0.044	1.246
	(0.816)	(0.028)	(0.098)	(0.028)	(0.890)
60th < LTV < 80th pctile	-0.955	-0.017	-0.100	-0.017	-0.077
	(0.889)	(0.029)	(0.108)	(0.029)	(0.758)
80th < LTV	-1.458	-0.007	-0.165	-0.007	0.661
	(0.982)	(0.032)	(0.116)	(0.032)	(0.610)
LTI					
p-LTI < 20th pctile	1.530**	-0.002	0.188**	-0.002	-3.249**
	(0.725)	(0.027)	(0.089)	(0.027)	(1.353)
20th < p-LTI < 40th pctile	1.693**	0.007	0.208**	0.007	-1.495
	(0.763)	(0.026)	(0.097)	(0.026)	(1.145)
40th < p-LTI < 60th pctile	1.907**	0.014	0.220**	0.014	-1.374
	(0.850)	(0.027)	(0.105)	(0.027)	(0.934)
60th < p-LTI < 80th pctile	2.168**	0.009	0.256**	0.008	-0.974
	(0.911)	(0.029)	(0.114)	(0.028)	(0.849)
80th < p-LTI	2.608**	0.024	0.271**	0.024	-0.624
	(0.992)	(0.031)	(0.121)	(0.030)	(0.747)
Renter	-0.269	-0.003	0.028	-0.003	-0.430
	(0.987)	(0.038)	(0.112)	(0.038)	(1.289)
No. children in the HH	-3.539***	-0.110***	-0.330***	-0.109***	-5.174***
	(0.434)	(0.015)	(0.050)	(0.015)	(0.595)
Log other HH income	-0.175	-0.007	-0.023	-0.007	-0.225
	(0.171)	(0.007)	(0.019)	(0.007)	(0.216)
Constant	18.261	2.358**	2.878	3.055***	
	(20.023)	(0.867)	(2.422)	(0.866)	
Observations	8,416	5,419	8,416	5,419	8,416
No. of Households	2,425	1,752	2,425	1,752	
R^2	0.047	0.039	0.039	0.039	
Household FE	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes
Region FE	Yes	Yes	Yes	Yes	Yes
Age quadratic	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes

Sample includes secondary earners from married and cohabiting couples aged 23 to 65 for the years 2001-2006. IHS(hours) stands for the inverse hyperbolic sine function of hours worked.
*** 1 percent ** 5 percent * 10 percent significance level.
Robust standard errors in parenthesis.

2 Labor force composition and the aggregate matching function in the United Kingdom [1]

2.1 Introduction

The rate at which job seekers are matched with vacancies is a key labor market indicator to understand unemployment dynamics over both the business cycle and the longer term. The job finding rate is determined by numerous factors besides labor demand and supply, including demographic trends, cyclical shifts in labor force composition, and structural changes in the matching process between workers and firms. In this work, we investigate the quantitative relevance of these channels for the job finding rate and the steady-state unemployment rate in the United Kingdom (UK). To this end, we undertake an analysis of the job matching process using micro-level data on individual worker transitions from unemployment to employment. We then estimate an aggregate matching function, a key building block of the Diamond-Mortensen-Pissarides (DMP) model, that is generalized to incorporate time-varying matching efficiency and heterogeneity in individual characteristics.

Empirically, employment prospects covary strongly with observable demographic characteristics such as age, sex, education, and duration of one's unemployment spell. The propensity to find a job of the "average worker", which we call "search intensity", therefore depends on the composition of the pool of job seekers and follows its fluctuations over time. Compositional changes in the characteristics of the searcher pool can arise from cyclical job separation shocks disproportionately affecting workers of certain types or from long-run demographic trends. Furthermore, average unemployment duration also endogenously fluctuates over the business cycle in response to changes in labor demand. It is worth remarking at the outset that we broadly

1 This chapter is co-authored with Bradley Speigner (Bank of England). I would like to thank Will Abel for help with data preparation and seminar participants at the Bank of England for useful comments.
Disclaimer: the views expressed in this work express solely those of the authors and should not in any way be interpreted as reflecting those of the Bank of England. The authors alone are responsible for any possible errors.

interpret search intensity to encompass exogenous differences in unemployment exit probabilities that are associated with observable characteristics. We hence do not attempt to distinguish between effort and opportunity, as long as they both drive cross-sectional variation in job finding prospects.[2]

Within the DMP framework, an additional determinant of the job finding rate is the "matching efficiency" of the labor market, which captures the inherent frictions in the search process. In reduced form, the effect that aggregate search intensity has on the labor market is similar to an improvement in the efficiency of the aggregate matching technology: i.e. for given levels of unemployment and vacancies, the total number of new matches rises. The primary objective of this work is therefore to explain the dynamics of the job finding rate in the UK by disentangling the separate effects of labor demand -proxied by vacancies-, search intensity -driven by composition-, and matching efficiency -which captures all factors not explained by composition or vacancies. Our analysis encompasses the period 1995-2016, with a particular focus on the decade spanning the Great Recession of 2008.

To disentangle the effects of search intensity and matching efficiency, we adopt a two-step estimation strategy based on Barnichon and Figura (2015) with some extensions. In essence, we augment the aggregate matching function of the canonical DMP model with a term representing the average search intensity of job seekers, thereby accounting for observed heterogeneity. In the first stage, we use micro-level data from the UK Labour Force Survey (LFS) to estimate via Maximum Likelihood the relative hazard rates of unemployed workers with different characteristics. Using the estimated coefficients and changes in the shares of worker types, we then construct a time series for aggregate search intensity. The second stage of the estimation process casts the aggregate matching function in state space form, treating matching efficiency as an unobserved time-varying state variable and modeling search intensity as an observed exogenous state. Maximum Likelihood Estimation (MLE) of the state space model generates estimates of the key parameters of the matching function, in particular the elasticity of matches with respect to vacancies. Kalman smoothing techniques also provide an estimate of the time-varying path of matching efficiency, which explains changes in the job finding probability that are unaccounted for by fluc-

2 We recognise that using the word "intensity" may be susceptible to misinterpretation, and sometimes "employability" may be a better descriptor. The estimation process remains agnostic about the causes of such differences in employment prospects.

tuations in labor demand and aggregate search intensity. Intuitively, by controlling for compositional changes in the observable individual characteristics of the searcher pool, an attempt is made to filter out a "purified" path of matching efficiency that is not distorted by changes in search intensity.

Our main findings are summarised as follows. First, matching efficiency commenced a downward trajectory in the late 1990s, which reached its trough in 2009. Although the downward trend has no longer been apparent since then, there has only been a partial recovery in the level. Second, we do not find a lasting deterioration in matching efficiency that is associated with the 2008 recession. There was also a pronounced rise in search intensity in the first part of the sample, during the late 1990s, peaking in 2005 and later falling during the Great Recession. This pro-cyclicality is accounted for by changes in mean unemployment duration. Over the longer term, education is the main factor underpinning a secular upward trend in search intensity, while age also has some effect. Changes in education and age composition account for an increase in aggregate search intensity of 8 percent between 1995 and 2016. Gender is found to have very little effect on employment transitions.

We consider extensions to the baseline specification in two directions. First, rather than focusing narrowly on unemployment, we broaden the definition of job seekers to include individuals who are "marginally attached" to the labor force and employed workers that are searching for another job. Our main conclusions are robust to this extension. In fact, the estimated path of matching efficiency evinces an even sharper fall prior to the Great Recession. Second, in recognition of the fact that the baseline model accounts for heterogeneity solely on the labor supply side, we also make an attempt to model heterogeneity in the demand for labor by controlling for variation in the "recruiting intensity" of different industries. Due to the lack of firm-level micro data, this approach is not as comprehensive as our modelling of the supply side. However, we do uncover suggestive evidence that sectoral shifts in vacancies may help explain the fall in matching efficiency in the period 1995-2002 that is detected by the baseline model.

Our final exercise is to map our results on changing labor force composition and matching efficiency into an analysis of the Beveridge Curve, which traces the steady state relationship between vacancies and unemployment. In the standard DMP model, purely cyclical dynamics in unemployment are generated by fluctuations in the level of labor demand (represented by vacancies) around a stable Beveridge Curve. We show that labor force heterogeneity implies that the composition of the pool of job seekers changes *along* the curve for different steady-state levels of labor demand. Unemployment spell duration is the main characteristic that changes along the curve. Relative to the canonical model with homogeneous job seekers, this channel causes the Beveridge Curve to pivot relative to the homogeneous version. Furthermore, long-run shifts in the individual characteristics of labor force participants, as well as structural aggregate conditions, affect the performance of the labor market in terms of matching job seekers to vacant jobs, thereby shifting the Beveridge Curve. We show that improvements in matching efficiency and search intensity both cause inward movements of the curve and we quantify how the Beveridge Curve may have shifted over time as a result. Using the mean level of vacancies from the 2001-2006 period, the Beveridge Curve computed using the matching efficiency and labor force composition from 2016 implies a steady state unemployment rate that is 0.6 percentage points lower than the 2005 curve.

This work contributes to the vast literature identifying the main drivers of unemployment dynamics. The idea that fluctuations in the aggregate job finding rate are driven by compositional changes in the pool of job seekers dates back to at least Darby et al. (1985, 1986), who note that "cyclical unemployment is concentrated in groups with low [...] exit probabilities". This claim was challenged by Shimer (2012), whose reassessment of the "ins and outs" of unemployment led to the finding that job seeker heterogeneity was not a major factor driving fluctuations in the aggregate unemployment exit rate. Numerous subsequent studies, often expanding on the DMP framework, have reached contrasting conclusions regarding key drivers of fluctuations.[3] While most of the analyses have focused on the US, several works also investigate the dynamics of UK labor flows (Smith, 2011; Gomes, 2012; Sutton, 2013) and the resulting compositional changes in the pool of job seekers (Elsby et al., 2011; Singleton, 2017).[4]. We contribute to this literature by providing a quantitative appraisal of the importance of observed worker characteristics for the historical

3 Key lines of analysis include general demographic characteristics (Barnichon and Figura, 2015; Kroft et al., 2016; Bachmann and Sinning, 2016), long-term unemployment and duration-dependence of employment probability (Krueger et al., 2014; Kroft et al., 2016), unobserved heterogeneity (Ahn and Hamilton, 2016; Morchio, 2016), the labor force participation margin (Elsby et al., 2015), job-to-job transitions (Sedláček, 2016), firm size (Gavazza et al., 2016), variable search effort (Hornstein et al., 2015), firm hiring standards (Sedláček, 2014), and sectoral segmentation (Şahin et al., 2014).

4 For studies on German data, see Kohlbrecher and Merkl (2016) and Klinger and Weber (2016)

fluctuations in the job finding rate and for the position of the UK Beveridge Curve.

Our work also relates to a more select number of studies on the path of matching efficiency, with a particular focus on the outward shift in the empirical Beveridge Curve after 2008 observed in several advanced economies (see Hobijn and Şahin, 2012, for an overview). A key question is whether this shift was driven by a change in the composition of the unemployed towards groups of workers with an inherently lower propensity to find jobs (i.e. search intensity) or a worsening of the overall process of matching (i.e. efficiency). Barnichon and Figura (2015) incorporate worker heterogeneity in a generalized matching function and find that shifts in the average characteristics of the workforce explain much of the residual in standard matching rate regressions for the US. Considering broader definitions of job seekers that include inactive and/or on-the-job searchers, Sedláček (2016) and Hall and Schulhofer-Wohl (2018) also reveal the presence of downward trends in matching efficiency in the US that began before the Great Recession. While our estimation methodology closely resembles that of Barnichon and Figura (2015), we also consider broader types of job seekers as in the latter works. Our results show that, similarly to the US, the UK also evinces a pre-crisis trend in falling efficiency. This finding contrasts that of Patterson et al. (2016), who find that occupational "mismatch" rose sharply over the Great Recession but without any long-term trend.[5]

Finally, we also make a first attempt, to our knowledge, at considering the role of shifts in the composition of labor demand across industries, an issue which has not yet been investigated in the UK. The majority of the works elsewhere in the literature focus on labor supply heterogeneity as opposed to heterogeneity in labor demand, probably due to lack of data. Notable exceptions include Davis and Haltiwanger (2014) and Gavazza et al. (2016) who look in detail at the macro implications of firm-side recruiting behavior.

The rest of the paper is structured as follows. Section 2.2 describes the data for the estimation, discusses unemployment dynamics in the UK, and presents stylized facts on the composition of the unemployed labor force. Section 1 describes the estimation procedure. Section 2.4 presents the baseline results and Section 2.5 discusses the two extensions. In Section 2.6 we report the implications of our matching function

5 The divergence, however, could be explained by the substantially different angle of analysis. We do not account for occupational and geographic segmentation, but we consider heterogeneity in worker characteristics.

analysis for the location of the Beveridge Curve. Section 2.7 concludes.

2.2 Transition rates and unemployment composition in the UK

This section describes the main data sources we employ and provides a descriptive analysis of the transition rates and the composition of the unemployed labor force in the UK.

2.2.1 The data

We restrict our attention to the sample period from 1994q1 to 2016q2. This is the period for which it is possible to measure labor force transitions at quarterly frequency using the two-quarter longitudinal UK Labour Force Survey (LFS). For each quarter, individual observations contain information regarding demographic characteristics as well as labor force status for the previous and current quarters. The linked data allow us to observe transitions from unemployment to employment as well as continued spells of unemployment. The discrete-time job finding probability F_t is defined as the fraction of workers transitioning into employment from unemployment in a given quarter. We apply the recommended survey weights to make the sample representative of the UK population. The LFS is also used to construct the aggregate unemployment level, U_t, needed to compute labor market tightness θ_t.

The vacancies data come from two sources. For the more recent part of our sample, covering 2002 to 2016, a national survey of job vacancies is available from the ONS Vacancy Survey. For the years 1995-2001, data on the stock of vacancies is obtained from vacancies at job centres, available on the NOMIS database. Because the latter series underestimates the total number of vacancies, there is a break between the two variables.[6]

2.2.2 Unemployment, transition rates, and compositional changes

Our sample covers is characterised by the recovery from the recession of the early 90s and the business cycle associated with the 2008 recession. Although the rise in unemployment following the Great Recession was relatively limited in the UK compared to some other advanced economies, the unemployment rate did remain persistently elevated for several years after the initial increase. Figure 1 shows that

6 We deal with this break in the state space model estimation, as described below.

Figure 1: Unemployment and vacancies in the UK

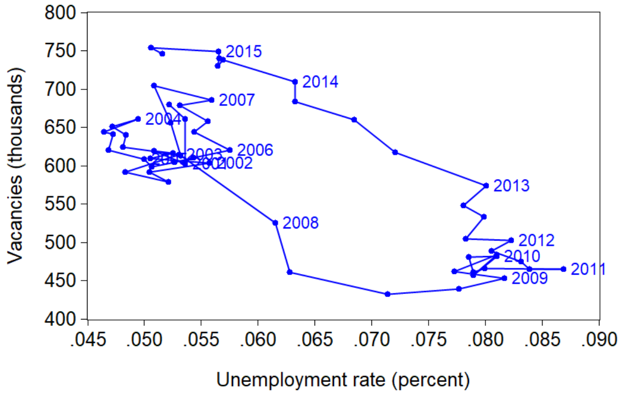

Note. The unemployment rate is measured as the number of unemployed workers over the total number of workers in the labor force within one quarter, using the appropriate LFS weights.

Figure 2: The job finding rate and the job separation rate in the UK

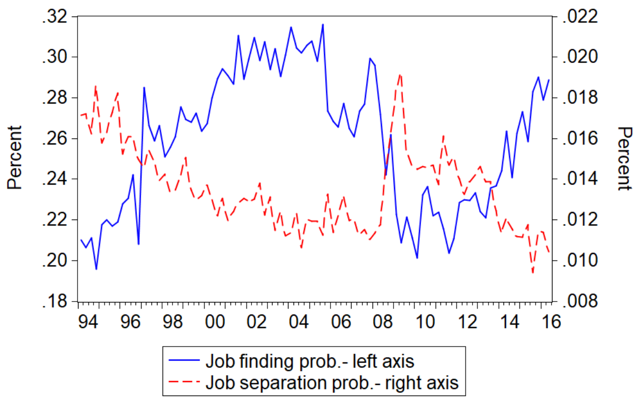

Note. The job finding (separation) rate is measured as the fraction of unemployed (employed) workers in one quarter that moved into employment (unemployment) by the following one. Both variables are computed using the appropriate LFS weights.

the unemployment rate rose from around 5 percent to over 8 percent during the recent recession, while vacancies fell. Only in 2013 has unemployment began to decline, reaching its pre-recession level by 2016. The rise in unemployment was driven by changes in both job creation and destruction. There was a sharp spike in job destruction at the onset of the recession as well as a protracted decline in the job finding rate (Figure 2).[7]

The composition of the pool of unemployed job seekers exhibits substantial changes

7 Note that these transition probabilities are not adjusted for continuous time aggregation. As shown by Shimer (2012), the bias of the discrete time measure can be nontrivial for the separation rate but is limited for the job finding rate.

over the sample period. Figure 3 shows the varying composition along four characteristics: age, education, sex, and duration of unemployment. The first three variables present clear secular trends. Over the past two decades, the proportions of female, young, and more educated workers have increased steadily. Meanwhile, duration of unemployment spells shows a cyclical pattern. As unemployment fell between 1995 and 2005, the share of short-term unemployment increased steadily. The Great Recession of 2008 is then marked by a sharp decline in the share of short-term unemployment and a rise in average duration.

To the extent that observed individual characteristics are associated with variation in one's ability to find employment, cyclical and secular shifts in the composition of job seekers can drive fluctuations in the aggregate job finding rate. In the next section, we outline a two-step estimation to quantify the contribution of compositional changes for the dynamics of job creation over the past two decades.

Figure 3: Compositional changes in the pool of unemployed workers in the UK

Note. All shares are reported as 4-quarter moving averages and were computed using the appropriate LFS weights.

2.3 Generalising the matching function

Our two-step empirical approach fuses insights from the literature on generalised matching function estimation with other work which applies time-varying parameter methods to labor market flows data and the estimation of aggregate matching efficiency. Earlier works tended to model time variation in matching efficiency through deterministic trends (see Petrongolo and Pissarides, 2001, for a review). More recently, Sedláček (2016) applies a more flexible latent variable technique while also specifying a generalised matching function that included a broad pool of job seekers. Hornstein et al. (2015) also make use of a Kalman filter to infer aggregate matching efficiency as the unobserved state in a similar setup.

The DMP model of matching frictions has become the canonical approach to modelling the aggregate job finding probability in the economy. The model assumes the existence of a frictional technology that relates the flow of new job matches (i.e. worker-vacancy pairs) to the stocks of vacancies and unemployment. Given the considerable supporting evidence in the applied literature, we make the conventional assumption of a Cobb-Douglas function with constant returns to scale. Defining m_t as the flow of newly formed job matches within period t, the basic matching function is given by

$$m_t = \mu V_t^{1-\eta} U_t^{\eta} \tag{1}$$

where V_t denotes the supply of available job vacancies, U_t is unemployment, or the pool of job seekers more generally, and $\eta \in (0, 1)$. Matching efficiency is defined as the scale parameter μ in equation (1).

In its basic form, equation (1) is an aggregate relation which abstracts from heterogeneity in the input variables V_t and U_t. However, data on the variation of job finding propensities across observable individual characteristics shows that this assumption is restrictive in practice. We therefore generalise equation (1) to allow for time variation in μ and heterogeneity in the "search intensity" of different groups of job seekers.[8] Denoting the search effectiveness of worker type j as s_j, assumed to be

8 There is the additional question of the underlying occupational structure of the labor market, and whether a single aggregate matching technology which pools workers and jobs from all industries is realistic. Barnichon and Figura (2015), for example, assumed that the labor market is segmented, with individuals searching only for jobs within their occupation of previous employment and geographic location, with each segment characterised by a separate matching function. We do not make this assumption because we lack data on vacancies by occupation. Despite having data on vacancies by industry, we opted for not including industry of previous employment among the first-stage con-

time-invariant, the matching function in generalized form is

$$m_t = \mu_t V_t^{1-\eta}(s_t U_t)^\eta, \tag{2}$$

where $s_t = \sum_j \frac{U_{jt}}{U_t} s_j$. Aggregate search intensity fluctuates over time due to changes in the unemployment shares of worker types, U_{jt}/U_t . For now, attention is restricted to the unemployed pool and we will consider expanding the searcher pool in Section 2.5 below. Under the assumption of random matching, meaning that each job seeker of a given type has the same probability of being matched to a vacant position, the continuous-time job finding rate of the type-j unemployed worker is

$$f_{jt} = \frac{s_j}{s_t}\frac{m_t}{U_t}. \tag{3}$$

Data on individual employment transitions are combined with a parameterization of the function s_j in order to estimate the dependence of job finding rates on individual traits using the relation in (3), following Barnichon and Figura (2015). The results are used to generate a time series for aggregate search intensity, s_t. Subsequently, dividing both sides of (2) by unemployment and taking logs yields a regression equation for the aggregate job finding rate:

$$\log f_t = \log \mu_t + \eta \log s_t + (1-\eta)\log \theta_t, \tag{4}$$

where $\theta_t = V_t/U_t$ is labor market tightness. We now describe in more detail how we take equations (2) and (4) to the data with a two-stage estimation approach.

2.3.1 First stage: micro-estimation

Individual search intensity, s_j, is parametrized using micro-data on individual labor market transitions from the UK Labour Force Survey (LFS). The search function s_j is simply assumed to be an exponential function of observable characteristics:

$$s_j = \exp(\beta X_j)$$

where X_j is a vector of worker characteristics. From our data, we observe whether each job seeker transitioned into employment in a given time period. The log-likelihood function is therefore set up as

trols because the data show that a surprisingly large number of individuals who lose their jobs end up finding work in a different industry.

$$l(\beta) = \sum_t \sum_j \sum_{i=1}^{N_j} \{y_{it} \log(F_{jt}) + (1 - y_{it}) \log(1 - F_{jt})\}, \tag{5}$$

where y_{it} takes a value of 1 if the individual finds a job in period t and 0 otherwise, and F_{jt} is the discrete time-adjusted job finding probability, and N_j is the total number of workers of type j. Given that the data are only observed at discrete intervals, the continuous-time job finding rate, which is assumed to be constant within each quarter, is converted to a discrete-time quarterly probability. Formally, $F_{jt} = 1 - \exp(-\frac{s_j}{s_t} f_{jt})$.

Maximum Likelihood Estimation (MLE) is employed to recover estimates of the parameter vector β. Given the structure of the matching process, what matters for individual transitions is *relative* search effectiveness. That is, search effectiveness is only identified up to a normalizing constant in our model. We therefore impose the normalization that aggregate search intensity equals 1 over the time period used for estimation, which is four periods of quarterly data for the year 1994. The implied assumption is that the estimated impact of individual characteristics on job finding do not change over the rest of the sample (i.e. β is fixed over time). Changes in aggregate search intensity over time therefore occur only through changes in the shares of job seekers across the different categories. We discuss in detail the robustness of this assumption below. Once s_j has been estimated for each worker type, the time-varying aggregate s_t can be computed by multiplying each s_j by the respective share of worker type j for each periods.

2.3.2 Second stage: macro-estimation

The second stage of the process estimates the matching elasticity parameter, η, and the path of μ_t. To this end, we cast the model in state space form, treating μ_t as an unobserved time-varying state variable. The state equation for matching efficiency is assumed to be a random walk,

$$\log \mu_t = \log \mu_{t-1} + \nu_t, \quad \nu_t \sim N(0, \sigma_\nu^2) \tag{6}$$

where ϵ_t represents innovations to matching efficiency. As a robustness check, we also use an alternative specification in which the efficiency state is allowed to be a stationary autoregressive process. The observation equation is obtained by adding an independent identically-distributed error term to (4) to make the job finding rate a "noisy" observation of matching efficiency:

$$\log f_t = \log \mu_t + \eta \log s_t + (1-\eta)\log \theta_t + \epsilon, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2). \tag{7}$$

In practice, we also treat the true stock of vacancies as an additional unobserved state variable. We make this choice as a remedy to the fact that prior to the introduction of the current national survey vacancies data were only available from job centers data. As it is not generally required for firms to post their job openings at job centres, these estimates typically suffer from incomplete coverage. Improvements in coverage over time due to modernisation can then induce a false perceived upward trend in the vacancies stock, which would bias the measured matching efficiency path. We assume that the job centre data on vacancies is only an imperfect signal of the true underlying stock of vacancies for the sample period to which this applies (i.e., prior to 2002). The numerator of tightness, V_t, is therefore also an unobserved state, and is assumed to follow a random walk with an independently distributed error term. The full state space model is specified in Appendix 2.A.

The state space model is estimated by maximising the likelihood function provided by a Kalman filter. The parameters of interest to estimate are the elasticity η and the standard deviations of the shocks. Using the estimates of the parameters, Kalman smoothing is used to obtain an estimate of the μ_t series.

2.4 Baseline results

First stage

Table 1 presents the coefficients from the first-stage MLE. The "reference group" is comprised of male individuals with no General Certificate of Secondary Education (GCSE) aged 16 to 25 who are short-term unemployed. Individuals in this category are roughly 50 percent more likely to find a job than the average searcher in 1994.[9] Educational attainment is positively associated with the chances of finding a job, as the dummy variables associated with higher academic qualifications have higher positive values and are statistically significant. Meanwhile, being female, older age groups, and duration of unemployment are all negatively correlated with job finding probability. Figure 2.C.1 in the appendix shows the implied distribution of relative search intensity over the population based on our results for the year 1994, illustrating large variation across groups.

9 Taking the exponent of the coefficient, $\exp(0.4) = 1.5$.

Table 1: Maximum-likelihood coefficients from the first (micro) stage estimation

	Coef.	**Std. Err.**	**P-val.**
Reference group	0.401	0.051	0.000
Female	-0.047	0.036	0.184
Age 26-35	-0.100	0.044	0.022
Age 36-45	-0.020	0.049	0.681
Age 46-55	-0.230	0.055	0.000
Age 56-65	-0.645	0.083	0.000
Other qual.	0.242	0.055	0.000
GCSE qual.	0.331	0.053	0.000
A-level qual.	0.331	0.052	0.000
Higher educ.	0.601	0.058	0.000
3-6 months	-0.264	0.047	0.000
6 month-1 year	-0.570	0.049	0.000
1-2 years	-0.915	0.055	0.000
2+ years	-1.435	0.061	0.000
Number of obs. = 16,961			
Wald $\chi^2(13)$ = 1131.60			
Prob. > χ^2 = 0.0000			
Log pseudolikelihood = -4885.2914			

Note. The table reports the coefficients for the relative search intensity estimated via MLE, as shown in Equation (5) using the four quarters of year 1994 from the LFS.

Once the vector β has been estimated, we obtain a time series of aggregate search intensity using the appropriate unemployment weights, which is plotted in Figure 4 together with the job finding rate. Comparing the two series, aggregate search intensity follows cyclical fluctuations similar to those of the job finding rate but also has an upward trend that is absent from the latter. Figure 5 decomposes aggregate search intensity into the observable characteristics that we measure.[10] Fluctuations in the unemployment duration composition account for the swings in search effectiveness at the business cycle frequency. As Figure 3 shows, the share of long-term unemployment falls steadily from 1995 to 2003. It then rises, with a sharp increase over the years of the Great Recession, until 2012 and falls henceforth. These cyclical movements are reflected clearly in the yellow area of Figure 5. On the other hand, the secular trend is accounted for mainly by education and, to a much lesser extent,

10 To compute the decomposition, we slightly changed the estimation of the likelihood function, using the linear functional form $s_{ij} = \beta X_i$ rather than the exponential one. Given the fact that all X's are dummy variables, this change allows for aggregate search intensity to be computed simply as $s_t = \sum_j s_j U_{jt}/U_t$, where each j represents a category of sex, age, education, and unemployment duration. The reason why we do not use this approach in the main model is that the linear formulation does not guarantee that the estimated s_j's are positive, although they all turn out to be for the estimation run on our sample of data.

age. Over time, the unemployment pool has become better educated, yielding an overall improvement in search effectiveness of almost 8 percent between 1995 and 2008. Gender has not had a measurable impact on search intensity.

Figure 4: Aggregate search intensity and the job finding rate

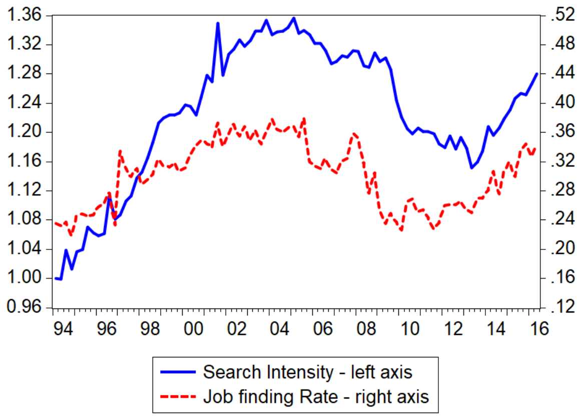

Note. the job finding rate is measured as the fraction of unemployed workers in one quarter that moved into employment by the following one. The search intensity index, normalized to 1 in 1994q1, is computed as the average search intensity of the unemployed population, using the coefficients from Table 1. All variables are computed using the appropriate LFS weights.

Figure 5: Decomposing changes in search intensity (4-quarter moving average)

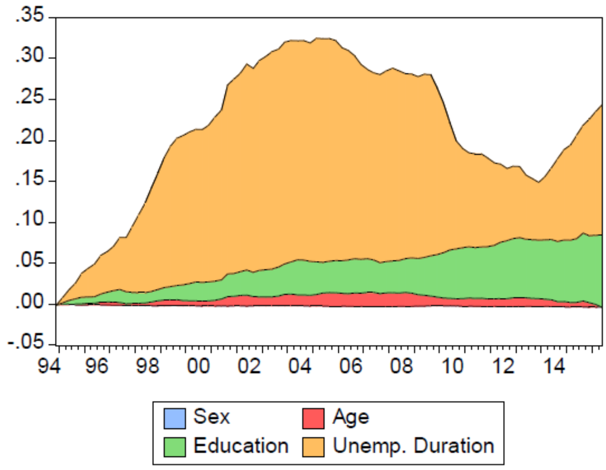

Note. the search intensity index, is computed as the average search intensity of the unemployed population, using the coefficients from Table 1. All variables are computed using the appropriate LFS weights.

Second stage

The search intensity index which was computed in the first stage of the estimation process is now treated as an observable variable in the aggregate matching regression (4). Standard Kalman smoothing techniques are then applied to obtain the expected path of the unobserved matching efficiency over the period. Table 2 shows the results of the estimation. For comparison, results are also reported for a version of the state space model excluding search intensity. The elasticity of the job finding rate with respect to tightness is close to 0.3 in both cases, which is consistent with other results in the empirical literature, but the point estimate is about 20 % higher when search intensity is omitted from the model. This suggests that specifications which fail to control for fluctuations in search intensity will have an upwardly biased matching elasticity estimate deriving from the pro-cyclicality of search effectiveness, which are caused by fluctuations in average unemployment duration as described previously.

Figure 6 plots the path of matching efficiency from the aggregate matching function augmented with search intensity of the unemployed (left panel) and for the standard un-augmented specification (right panel). Controlling for the characteristics of the unemployed, the estimated path of matching efficiency evinces a steady decline from 1997 to 2009, with only a partial recover since then. We fail to find a sustained negative impact of the 2008 recession on matching efficiency, as the decline started several years before. The estimated decline in matching efficiency pre-dating the recession is large; the labor market was about 14% less efficient at matching workers with job openings in 2008 than it was in 1995. However, there has been a partial recovery since then, implying that matching efficiency is about 7 percent lower at the end of the sample compared to the beginning.

Failing to control for time variation in the aggregate composition of job seekers significantly affects the resulting path of matching efficiency. When search intensity is not controlled for, it appears as though matching efficiency was stable, or even slightly rising, in the earlier part of the sample, before beginning a decline prior to the Great Recession which has stalled but not unwound since. By 2016, matching efficiency is about 4% below its 1995 value.

In Appendix 2.C, Figure 2.C.2 reports the smoothed path of vacancies, estimated from the state space model. Prior to 2001, the job center data serves as signal for the true state. From 2001 onwards the state is fully observed and coincides with the actual ONS series for vacancies.

Table 2: Second-stage estimation results for the baseline model

	(1) With search intensity	(2) Without search intensity
1- η	0.277	0.333
	(0.048)	(0.042)
Log likelihood	234.77	234.82
Observations	90	90
Search Intensity	Yes	No

Note. The table reports the MLE coefficients and the log likelihood from the state space model of the baseline aggregate matching function with and without search intensity as an exogenous observed state. Standard errors are reported in parenthesis.

Figure 6: Smooth estimate path of matching efficiency (1995q1=1)

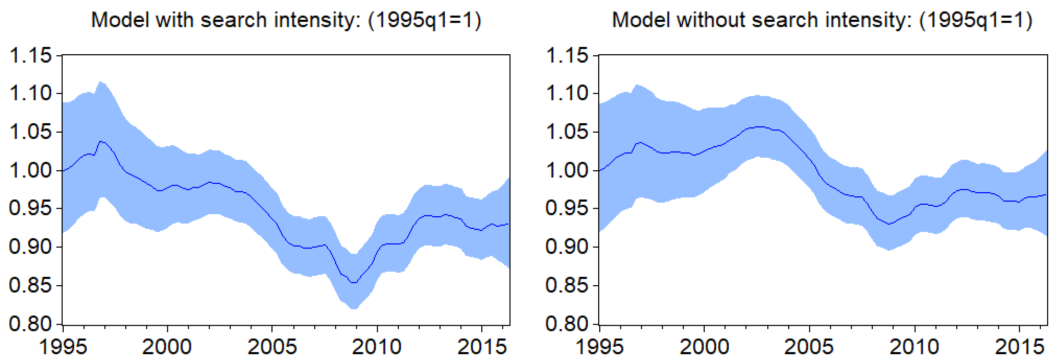

Note. the dark blue line reports the Smooth estimate of the unobserved state variable μ_t, normalized to 1 in the first period of the sample, while the dark blue area represents the 95 percent confidence interval.

Sensitivity analysis

Before discussing extensions to the model, we test the robustness of the baseline results with a few variations in the first stage estimation. In the first case, we exclude Government Training Schemes (GTS) from the unemployment-to-employment transitions. GTS inflows may be reflective of government policies rather than labor demand. Flows into GTSs are arguably not subject to the same frictional matching process as regular job openings. Omitting such flows from the definition of the job finding probability may hence have implications for the results, with our baseline treatment of the data possibly overstating the true vacancy yield, particularly during the recession. Although GTSs constitute only a small portion of all transitions into employment, the share of newly employed workers in GTSs increased significantly after the 2008 recession, as show in Figure 2.C.3. The first column of Table 2.C.1 shows that the coefficients of the first stage are not particularly affected by this restriction on the type of exits from unemployment.

The second robustness check addresses the issue of true duration of a joblessness spell. The LFS includes a question on the time since an individual last had a job. For a non-trivial portion of the sample, the duration of unemployment is shorter than the total time without a job, implying at least one transitional period out of the labor force. Being jobless, whether inactive or seeking, is detrimental for human capital and hence employment prospects. Only considering unemployment duration can therefore neglect long periods of non-employment for some individuals.

We assess the results' sensitivity to this issue with two different specifications. In the first one, we generate a binary variable for having spent time in the inactive state. In the second one, we replace unemployment duration with joblessness duration, regardless of whether the individual was searching or not for the entire time. Columns 2 and 3 of Table 2.C.1 show that the coefficients of the first-stage MLE for these two specifications are very similar to those of the baseline model. In Column 2, the dummy variable for having had a spell of inactivity is negative and statistically significant. This result may either indicate an adverse impact of having spent time outside the labor force or simply account for the extra length of the joblessness spell for this group of workers. Interestingly, in this specification the coefficients of the duration categories are similar in value to the baseline model. However, when using the effective time out of a job (Column 3), the magnitude of the coefficients is attenuated, implying a less adverse effect of duration. The possible explanation for the attenuation is that a portion of workers reporting a short duration of unemployment in fact had spent a long time out of the labor force, which compresses the true disparities in job finding prospects.

As a final robustness check of the first stage, we introduce controls for different reasons for unemployment. Many works, and most recently Ahn and Hamilton (2016), show that those workers who separated from their last job voluntarily subsequently have a higher job finding probability compared to those who suffered an involuntary layoff. We hence include two dummy variables in the fourth specification: one for quitting the last job and one for being involuntarily separated. In this case, the reference group is composed of those who have never had a job (mostly young workers who recently came into the labor force), those who separated for health, family, or other reasons, those offered early retirement, and those whose temporary job contract came to an end. Because the classification of reasons for unemployment in the LFS has changed over time and does not cover all observations, utilizing this variable entails a loss of observations. Since the coverage of the variable is particularly limited in the years 1994 and 1995, we carry out the estimation for 1996. Both the quit and layoff dummies are positive and statistically significant, with the former being larger than the latter. This result is in line with intuition: the reference group includes job seekers without work experience as well as individuals with family or health problems. These groups are less likely to move back into employment compared to those who terminated their last job with the intention to find another one. Furthermore, the result that voluntary quits are less disruptive to future employment opportunities (i.e. $\beta_{quit} > \beta_{layoff}$) is also in line with the findings of the literature (see for instance Ahn and Hamilton, 2016).

Moving on to the second stage, Table 2.C.2 reports the estimated elasticity parameter of the matching function for the four first-stage robustness exercises. The case without search intensity is only reported for the GTS exclusion because that is the only robustness check in which we change not just the search intensity measure but also the aggregate job finding probability. In all cases, the estimates are in line with the baseline results from Table 2: the elasticity with respect to vacancies estimated with search intensity is lower than in the canonical model. Finally, Figure 2.C.4 shows the Smooth estimate of the path of matching efficiency. Consistent with the baseline result, for all four checks except the GTS exclusion, matching efficiency in the intensity-augmented models begins falling in the late 90s and reaches its trough by 2009. In the GTS exclusion, the intensity-adjusted specification does not show a decline in efficiency until 2002. However, the GTS exclusion without controlling for intensity (right panel of the top subfigure) yields an upward path for matching efficiency in the period 1995-2003. Thus, the result that matching efficiency in the augmented model shows a more declining path prior to the Great Recession still holds.

Another robustness check we carry out involves relaxing the unit root assumption of the efficiency process μ_t in the second stage. In Table 2.C.4 we report the results of the second stage where μ_t is an AR(1) process $\log \mu_t = \alpha + \rho \log \mu_{t-1} + \nu_t$ with an autoregressive coefficient ρ and intercept α, which are also estimated. The estimated ρ coefficient is 0.94 for the model with search intensity and 0.89 for the un-augmented specification. The two coefficients are close to 1, implying that the unit root assumption is not a very restrictive one. Figure 2.C.5 reports the resulting estimated path. As visible from the figure, the resulting path is not particularly different from the baseline.

The choice of the reference year for the first stage is also a potentially important factor in the results. The key assumption for the first stage is that the β's, estimated on a baseline year only, are constant over the years, and hence the contribution of different individual characteristics to search intensity does not vary with time. To assess the robustness of this assumption, we repeat the MLE on each year separately and plot the estimated coefficients for each year (with 95% confidence intervals) in Figure 2.C.6. The plot shows that most coefficients exhibit only minor variations over the years. Only three coefficients show substantial changes or clear trends: "age 56-65", "unemployed for longer than 2 years", and the reference group.

In itself, variation in the values is not necessarily a problem as it may simply result from changes in the composition of job seekers. For instance, as the unemployed pool becomes more educated, the advantage of high-education workers in finding a job relative to the "average" worker falls and hence the respective coefficient would also be smaller. However, changes may also be driven by true shifts in relative search intensities. The latter case would result in changes in the path of the aggregate s_t over time compared to our baseline results. To assess the impact of estimating the coefficients in different years, Figure 2.C.7 plots aggregate search intensity, normalized to 1 in 1994q1, using the estimated β's from 1994, 2000, 2005, 2010, and 2015. It is clear from the graphs that qualitatively the results do not change based on the year used for the first stage. In all cases, search intensity presents an upward path until the early 2000s, followed by a dip and partial recovery after the Great Recession. Quantitatively, the main difference arises from the magnitude of the fluctuations. For the estimation conducted with the 2005 data, intensity peaks at 1.4, while for 2010 and 2015 the maximum value is slightly below 1.3. The 1994 series is somewhat in the middle of these extremes, meaning that it can be interpreted as a more balanced candidate to serve as reference year.

2.5 Extensions

So far, we have associated job seekers with the definition of unemployment provided by the International Labor Organization. In this section, our first extension involves extending the sample beyond the standard definition of unemployment by incorporating "marginally attached" individuals and workers undertaking on-the-job search. The second extension proposes a preliminary attempt to account for heterogeneity on the labor demand side, which remains relatively unexplored in the wider literature.

2.5.1 Expanding the set of job seekers

In the baseline model, we only consider transitions from unemployment to employment. However, this flow accounts for only about half of all newly employed workers (Gomes, 2012). Therefore, we now expand the definition of job seekers to include job-to-job moves and employment inflows from inactivity. We thus re-interpret U_t in the matching function 2 as a measure of all job seekers: not just the unemployed (U), but also on-the-job searchers (OJS) and "marginally attached" (MA) inactive workers, each with an associated level of search intensity. What distinguishes these groups of seekers from the unemployed is the intensity of job search in the matching function, and possibly the degree to which observable characteristics affect their job finding probability.[11]

Following Gomes (2012), OJS individuals are defined as those employed workers who state that they are looking for another job. For these workers, a job-to-job transition is identified in the LFS micro-data as an OJS worker who is employed in both quarters but whose employment tenure in the second quarter is below three months. One caveat is that this measure possibly overstates the total number of job-to-job flows because it does not account for possible spells of involuntary unemployment between the two jobs. The MA category is defined as those individuals not in the labor force who are not actively searching for work but would be willing and able to start a job in the next two weeks.

Some descriptive analysis of the main differences between unemployed searchers, MA, and OJS individuals, provided by Figure 7, offers several insights. OJS activity is mildly procyclical and evinces a slightly increasing trend. Meanwhile the level of MA searchers remains fairly stable over the years and is subject to very mild counter-cyclical fluctuations.[12] Overall, both series show much smaller volatility than

11 For this reason, we also allow the β coefficients of education, age, and sex to differ for each group.

12 A caveat: the relative acyclicality of the raw figure of MA workers may mask cyclical fluctuations in the transitions from inactivity to unemployment. For instance, if MA individuals were to be more likely to actively search for work (and hence be classified as unemployed) during recoveries, the constant MA level would imply a commensurate increase in the inflows into MA status from

unemployment at business cycle frequency. As a result, the composition of the entire pool of job seekers varies over time, with the unemployed and OJS workers accounting for the main changes. The share of OJS workers rose from 30 percent in 1994 to more than 40 percent by 2007. It then dropped sharply during the Great Recession and eventually recovered beyond its 2007 point by 2016. The flow rate into employment varies dramatically across the three groups, both in its level and in the magnitude of its fluctuations. The job finding rate of the unemployed is the highest and the most pro-cyclical. OJS workers have lower chances of finding new employment, and MA searchers lower still. These two transition rates are also less volatile than the unemployment-employment transition probability. Furthermore, while the latter has almost returned to its pre-2007 peak, the job OJS finding rate remains below it by 2016.

The bottom right panel of Figure 7 shows the estimated path of aggregate search intensity for each of the different job seeker definitions (the estimated coefficients are reported in Table 2.C.3). The path of search intensity for the unemployed (dashed green line) is merely a rescaled version of the baseline s_t series. The paths of intensity for MA and OJS workers do not show any cyclical movement because they do not include any variation in duration. When including MA and OJS types, aggregate search intensity (black line) only rises moderately in the first decade of the sample and by 2016 it is almost back to the original level. Importantly, part of the movements in aggregate search intensity are due to variation within each group of seekers, while others are caused by the changing composition across groups.

Table 3 presents the second-stage state space estimation results for the extended sample with and without search intensity. Unlike the baseline case, the addition of search intensity leads to an estimate of the matching elasticity with respect to vacancies that is higher (by about 0.05) compared to the un-augmented matching function. This suggests that not controlling for fluctuations in the number of OJS job seekers can bias downwards the elasticity estimate since their share in the aggregate searcher pool is procyclical and they have a low likelihood of matching.

The resulting smoothed path of matching efficiency for the full U+MA+OJS sample is shown in Figure 8. Comparing the results of this model with the baseline one from Figure 6 (reproduced as the red lines in Figure 8), two differences are vis-

other labor force statuses. Singleton (2017) provides a comprehensive analysis of flows in and out of inactivity.

Figure 7: Comparison of unemployed, MA, and OJS: level, share, job finding probability and estimated search intensity.

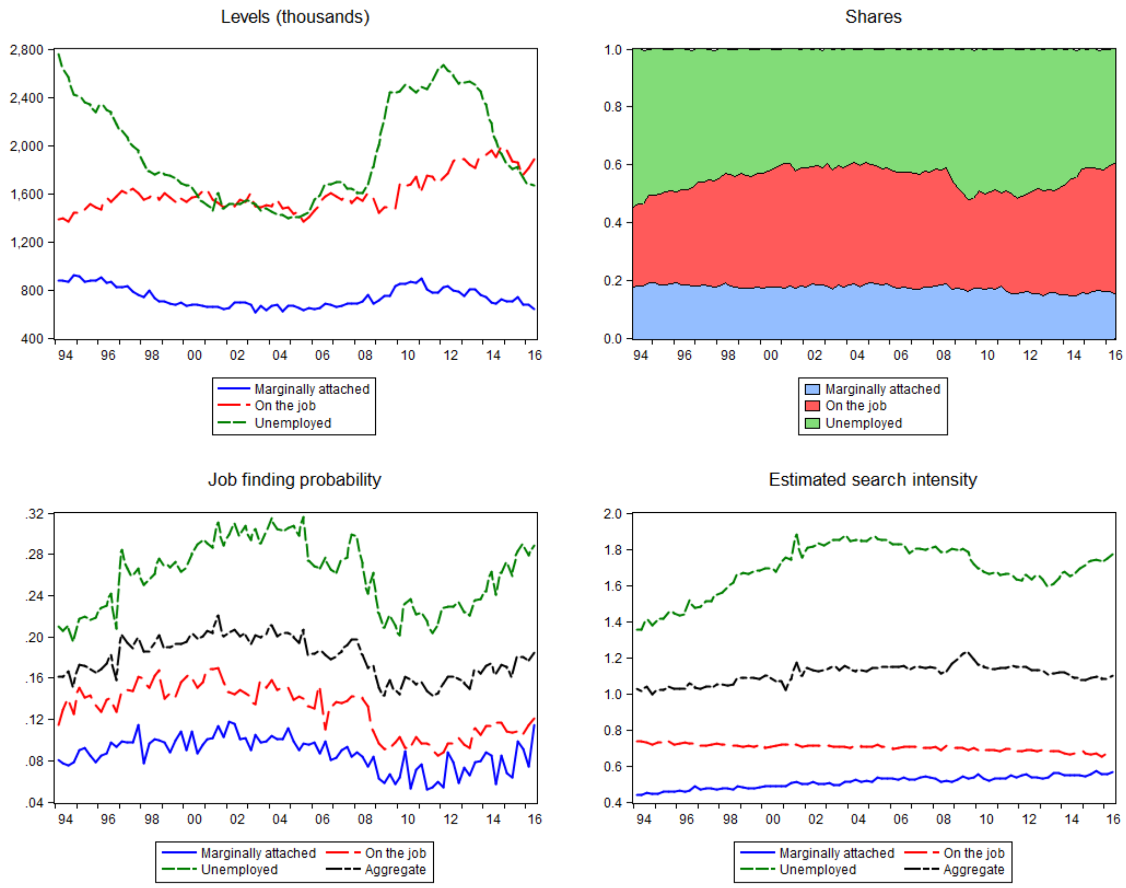

Note. the job finding rate is measured as the fraction of unemployed workers in one quarter that moved into employment by the following one. The search intensity index, normalized to 1 in 1994q1, is computed as the average search intensity of the unemployed population, using the coefficients from Table 2.C.3. All variables are computed using the appropriate LFS weights.

ible. First, the path of matching efficiency prior to the Great Recession has a sharper downward slope, reaching a deeper trough. In the intensity-augmented specification, the minimum search intensity is 77 percent of the initial value. This trough is almost 10 percentage points lower than in the baseline sample that is restricted to unemployed searchers. Second, the path of the model without search intensity remains almost flat until 2005 and does not present the slight increase visible in Figure 6 for the same period.

The explanation for the second observation originates from the differences between unemployed workers and the OJS pool. As outlined above, the dynamics of the job

Table 3: Second-stage estimation results for the model with the expanded set of job seekers.

	(1) With search intensity	(2) Without search intensity
1- η	0.341 (0.050)	0.292 (0.048)
Log likelihood	246.3781	246.3871
Observations	90	90
Search Intensity	Yes	No
Job Seeker Definition	U+MA+OJS	U+MA+OJS

Note. The table reports the MLE coefficients and the log likelihood from the state space model of the aggregate matching function using the expanded pool of job seekers with and without search intensity as an exogenous observed state. Standard errors are reported in parenthesis.

Figure 8: Smooth estimate of matching efficiency assuming an expanded job seeker definition (1995=1).

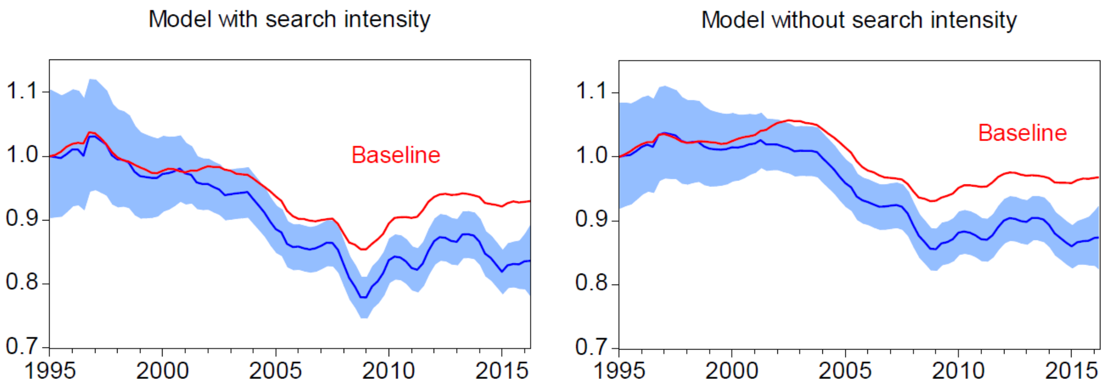

Note. the dark blue line reports the Smooth estimate of the unobserved state variable μ_t, normalized to 1 in the first period of the sample, while the dark blue area represents the 95 percent confidence interval. The red line reports the baseline case from Figure 6 for comparison.

finding rate are affected by the compositional changes occurring within the expanded pool of job seekers. Because of the increase in the share of OJS workers in the period 1995-2005, the aggregate job finding rate does not rise as steeply as when only the unemployed are included in the searcher pool.

The results of this section indicate that the halt in the decline of matching efficiency in 2008 appears to be robust to considering a wider measure job seekers, accounting for both changes in the shares of the different types of seekers as well as in the composition of the pool of individuals within each type. However, the resulting dynamics from adding MA and, more importantly, OJS workers show that focusing narrowly on unemployment is not representative of job search behavior in the broader labor market.

So far, the analysis has only controlled for variable search intensity along the supply side of the labor market. In particular, we have assumed that all vacancies are supplied by firms which recruit with homogeneous effectiveness. In the next subsection, we describe how we attempt to relax this assumption and how it influences our results.

2.5.2 Recruiting intensity

In a two-sided labor market, controlling for observable characteristics which influence the search effectiveness of job seekers, as done so far, only captures half of the story. Recruiting intensity of firms can play an equally important role. Recent research for the US by Gavazza et al. (2016) shows that firms adjust recruiting effort as the number of job seekers per vacancy changes, leading to fluctuations in vacancy yields as recruitment effort drops during recessions.

We can take a tentative step in this direction by exploiting cross-sectional variation in recruiting effectiveness across industries. To the extent that vacancy yields differ across sectors, then either business cycle fluctuations, trend shifts in industrial composition, or changes in the recruiting intensity within each industry may all affect measured aggregate matching efficiency.

Figure 9 shows the shares of vacancies and of flows from unemployment to employment for each of ten industries.[13] Two observations can be made. First, for most industries there is a wedge between the vacancy share and the hiring share, indicating differences in recruiting intensity across sectors. Second, some industries display large changes over time in the way the two shares track each other, particularly around the Great Recession. For instance, Financial Services experienced a permanent drop in the share of matches but not in the share of vacancies, indicating a potential fall in recruiting intensity. The opposite occurred for Construction and the Public Sector.

Under the assumption that job seekers search in all industries (i.e. search is not segmented), the aggregate matching function can be expanded to include recruiting intensity through a vacancy-augmenting term as follows:

13 Vacancies by industry are produced using the Jobs Centre Data from NOMIS until 2001 and the ONS series from 2001 onwards. To compute recruiting intensity over the whole period we chain-link the NOMIS series to the ONS ones. The industry classification, which changed, is harmonized to create a consistent set of industries throughout the sample. The LFS contains a question on current industry of employment, which is used to measure flows from unemployment into different industries using the same classification as for vacancies.

Figure 9: Match and vacancy shares by industry.

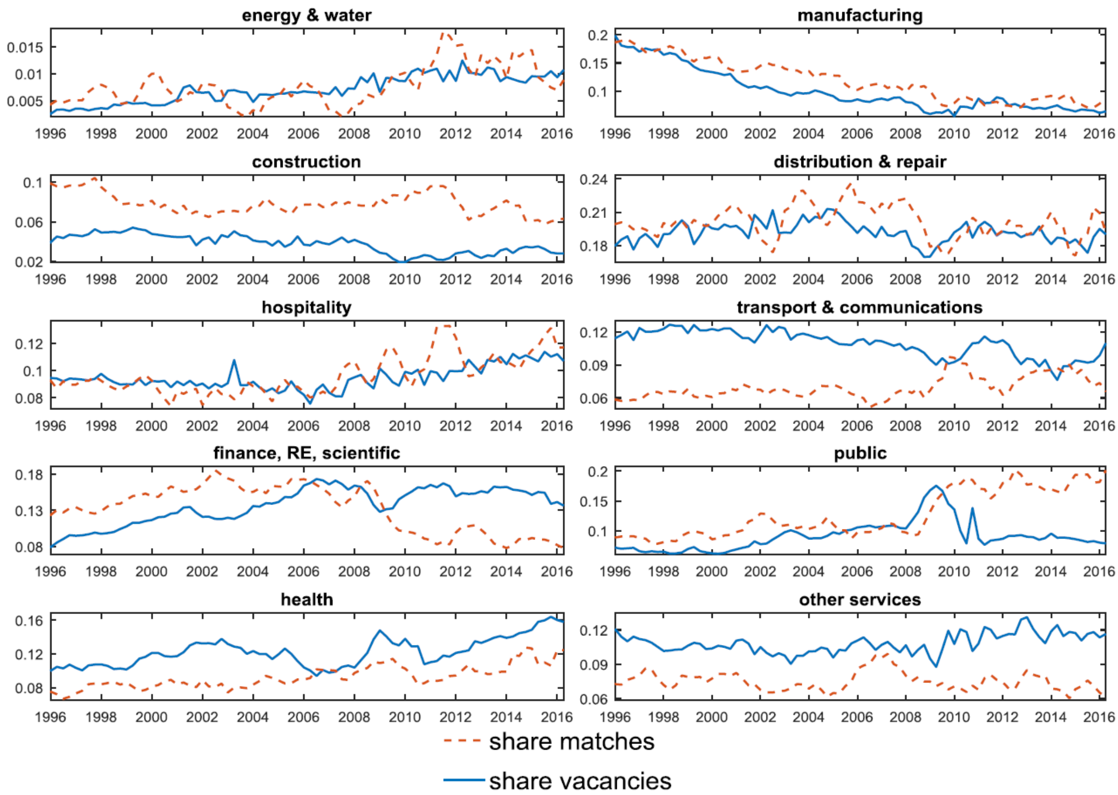

Note. the share of vacancies for each industry is provided by the ONS, while the share of matches is computed using the reported industry of employment for newly employed workers in a given quarter in the LFS. Both shares are reported as 4-quarter moving averages.

$$m_t = \mu_t (r_t V_t)^{1-\eta} (s_t U_t)^{\eta}, \tag{8}$$

where $r_t = \sum_i \frac{V_{it}}{V_t} r_{it}$. The term r_{it} represents the industry-i specific recruiting intensity, which, unlike search intensity, we assume to be time-varying. The first-stage estimation for search intensity allows for a wide set of covariates in age, education, duration, and sex. For recruiting intensity, industry is the only dimension of heterogeneity. In practice, we choose to allow for the r_{it}'s to vary over time in relation to a reference industry. Consequently, the discrete-time probability for an unemployed individual of type j to find a job in industry i is

$$F_{ijt} = \frac{r_{jt} V_{jt}}{r_t V_t} F_{jt} = \frac{r_{jt} V_{jt}}{r_t V_t} (1 - \exp(-f_{jt})) \tag{9}$$

Hence the probability of entering a job in industry i conditional on finding a job is independent of the worker's individual characteristics, which implies that the relative recruiting intensities of different industries can be estimated separately from the search intensity of individuals. In each period t, the set of relative recruiting intensities r_{it}'s can be estimated as the set of coefficients satisfying the equations

$$\frac{r_{it}V_{it}}{r_{\hat{i}t}V_{\hat{i}t}} = \frac{m_{it}}{m_{\hat{i}t}}, \tag{10}$$

where one industry $\hat{i}$ is taken as the reference industry in all periods, such that $r_{\hat{i}t} = 1, \ \forall t$. The r_{it}'s recovered through this method thus represent recruiting efforts relative to the reference industry. However, under the assumption that the intensity in the reference industry has not changed over time, the aggregate r_t can also be interpreted as absolute changes in recruiting intensity. Based on the graphic evidence of Figure 9, manufacturing is the industry that seems closest to satisfying this assumption. The shares of vacancies and of hiring have a correlation of 0.85, signalling a stable relationship between the two. Taking manufacturing as the reference group, whose recruiting effort is constant and normalized to 1, changes in all the other r_{it}'s and in the vacancies V_{it} determine the progression of aggregate recruiting intensity.

Figure 10 reports the estimated changes in r's and the derived aggregate intensity. Aggregate recruiting intensity was on a slow downward trend for the period 1995-2008, initiating a recovery in the aftermath of the Great Recession. Clearly, the relative ability of firms to hire in different sectors seems to vary over time, such that compositional shifts cannot be considered the only drivers of recruiting intensity.

Including our measure of recruiting intensity in the state-space estimation of the matching function is straightforward, as the signal equation becomes:

$$\log f_t = \log \mu_t + \eta \log s_t + (1-\eta) \log r_t + (1-\eta) \log \theta_t + \epsilon_t \tag{11}$$

We carry out the MLE on the sample comprised only of the unemployed. Table 4 presents the estimates of the parameters adding recruiting intensity to the aggregate matching function, with and without including search intensity. Interestingly, in both cases the estimated value of η, representing elasticity of the matching function with respect to vacancies, is lower than in the baseline model. The result that including search intensity yields a lower estimated value, however, still holds. As shown in the left panel of Figure 11, over the period 1995-2002, matching efficiency remains fairly constant, before turning downwards until the outbreak of the Great Recession. The inclusion of recruitment intensity tempers, but does not eliminate, the pre-crisis decline in matching efficiency compared to the baseline model. This is because over the period 1995-2002 the rise in search intensity and the fall in recruiting intensity effectively offset each other, leaving matching efficiency relatively stable. While this extension can explain some of the decline in matching efficiency at the beginning of the sample, we still observe deteriorating matching efficiency before the crisis, which has only partially unwound since then.

Figure 10: Recruiting intensity by industry.

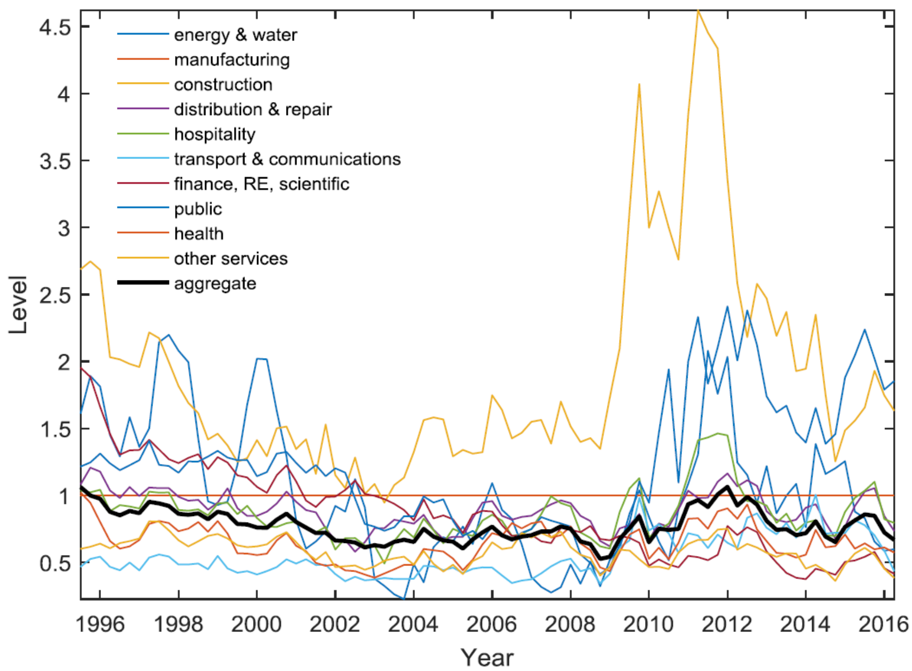

Note. recruiting intensity for each industry is computed with respect to manufacturing, as indicated in equation (10). The aggregate intensity is computed as the mean intensity based on the share of vacancies by each industry.

Table 4: Second-stage estimation results with recruiting intensity

	(1) With search intensity	(2) Without search intensity
1- η	0.23	0.248
	(0.041)	(0.041)
Log likelihood	232.54	228.992
Observations	84	84
Search Intensity	Yes	No
Recruiting Intensity	Yes	Yes
Job Seeker Definition	U	U

Note. The table reports the MLE coefficients and the log likelihood from the state space model of the aggregate matching function including recruiting intensity as an exogenous observed state and with or without search intensity as an exogenous observed state. Standard errors are reported in parenthesis.

Figure 11: Smooth estimate of matching efficiency controlling for recruiting intensity (1995=1).

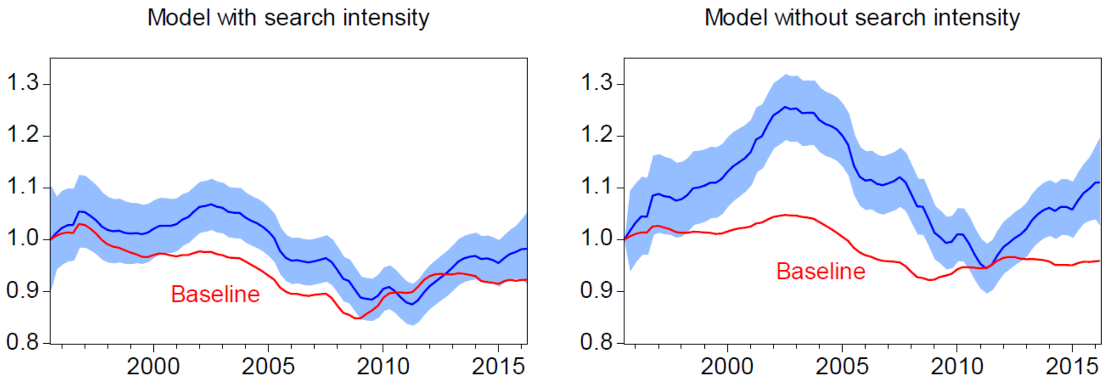

Note. the dark blue line reports the Smooth estimate of the unobserved state variable μ_t, normalized to 1 in the first period of the sample, while the dark blue area represents the 95 percent confidence interval. The red line reports the baseline case from Figure 6 for comparison.

2.6 The Beveridge Curve with heterogeneous workers

In this section, we map changes in matching efficiency and search intensity into shifts of the Beveridge Curve. In the standard DMP model with homogeneous labor, the implicit steady-state relationship between vacancies and unemployment is given by

$$\overline{U} = \frac{\delta}{\delta + F\left(\frac{\overline{V}}{\overline{U}}\right)}, \tag{12}$$

where δ is the exogenous separation rate, $F\left(\frac{\overline{V}}{\overline{U}}\right)$ is the job finding rate, and the bar over unemployment and vacancies indicates steady-state values. In our model with heterogeneous job searchers, the steady state composition of unemployment is itself a function of labor market tightness. Additionally, the inclusion of worker heterogeneity entails a much larger set of steady-state conditions. For a given value of vacancies, the aggregate level of unemployment is computed by solving a set of equations similar to (12), one for each worker type. Since the specific job finding rate of each group depends on the composition of the unemployment pool, aggregate restrictions need to be contemporaneously satisfied. Furthermore, while most worker characteristics are assumed to be fixed, duration dependence of employment prospects implies that transitions from short-term to long-term unemployment (more specifically, across the different duration categories available in the LFS) must also satisfy steady state

Figure 12: The Beveridge Curve with heterogeneity.

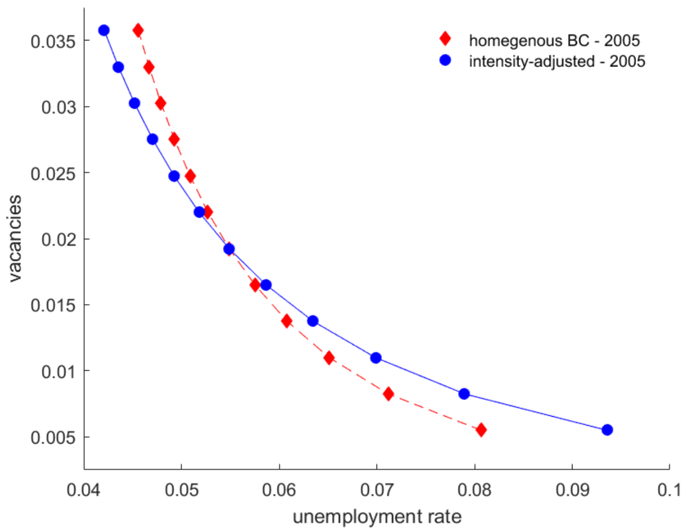

Note. The red diamonds plot the canonical Beveridge Curve computed using Equation (12). The blue circles plot the Beveridge Curve with heterogeneous job seekers computed through the set of equations described in Section 2.B.

restrictions.[14] Appendix 2.B contains the details of the solution.

To assess the impact of labor force heterogeneity, we compute the Beveridge Curve with the set of conditions outlined above using the matching efficiency value for 2005q1 and the labor force composition from the LFS for the same quarter.[15] For comparison, we then compute the homogeneous Beveridge Curve using equation (12), calibrated to have the same value of V when U is equal to 5.5 percent.

As shown in Figure 12, heterogeneity pivots the Beveridge Curve relative to a model with homogeneous unemployment, effectively flattening the curve.[16] When vacancies are high (low), unemployment is lower (higher) than in homogeneous case. Heterogeneity therefore increases the implicit elasticity of steady-state unemployment with respect to vacancies.

Figure 13 shows how the composition of steady state unemployment changes with

14 For simplicity, 10-year age groups are assumed to be fixed in the exercise.

15 Although we have experimented with group-specific separation rates, for the analysis presented below we opted for applying the average quarterly separation rate to all groups. Despite the important role of heterogeneity in separations, our choice was dictated by the overall noisiness of the estimates of the separation rates for some small groups of workers in different years. Because such noisiness could affect the analysis of shifts in the Beveridge Curve over time, we preferred to impose the more restrictive assumption of homogeneous separations to maintain the focus of the analysis on hiring prospects.

16 In Figure 12 as well as the following figures we normalize vacancies by total labor force (employed and unemployed) to make the empirical data comparable with the computed Beveridge Curve. Hence, the values of vacancies on the y-axes can be interpreted as a "vacancy rate".

the level of labor demand, illustrating that variation in the duration distribution of unemployment accounts for the pivot in Figure 12. The intuition is as follows. Higher vacancies raise job finding prospects for all groups, but in the steady state this shift implies that fewer workers reach long-term unemployment and so the long-term share declines, raising the average level of search intensity. Since this change in average duration affects all groups of workers, the overall composition of unemployment along the dimensions of sex, education, and age evince only extremely small changes. Hence, along the Beveridge Curve unemployment duration is almost the sole source of changes in the search intensity.

This framework can shed light on possible shifts of the Beveridge Curve over time. Improvements in matching efficiency or in the search intensity of the labor force shift the Beveridge curve inwards, so that a given level of job openings will be associated with a lower steady state unemployment rate. Our results from the previous section indicate that over the initial part of the sample, before the 2008 recession, the improvement in job seeker quality was offset by a decline in general matching efficiency. These counterbalancing forces tended to keep the Beveridge Curve relatively stable. In more recent years, matching efficiency stopped declining but the search intensity of the labor force continued to rise. The implication is that the Beveridge Curve may have shifted inwards due to the ongoing improvement in search intensity.

To carry out a simple exercise, Figure 14 plots the Beveridge Curve for 2005 and 2016. These years are fairly comparable since 2005 can be interpreted as a moment of stability in the UK labor market and by 2016 the economy reached similar unemployment rate and was 8 years past the Great Recession. As the figure shows, the 2016 curve lies slightly to the left of the 2005, implying a lower steady-state unemployment for a given level of vacancies. For instance, for the vacancy rate of 2.2 percent, close to the 2001-2006 average, the steady state unemployment rate drops from 5.2 to 4.55 percent. In Figure 15 we plot two alternative Beveridge Curves where all parameter values are set to their 2005 level except either matching efficiency or the labor force share of different worker groups, which are set to their 2016 levels. The two alternative curves hence partial out the individual effect of each channel. As the figure shows, both a slight recovery of the matching efficiency and a change in labor force composition contributed to the shift. While the former mechanism derives

Figure 13: Changing composition of the unemployment pool along the Beveridge curve.

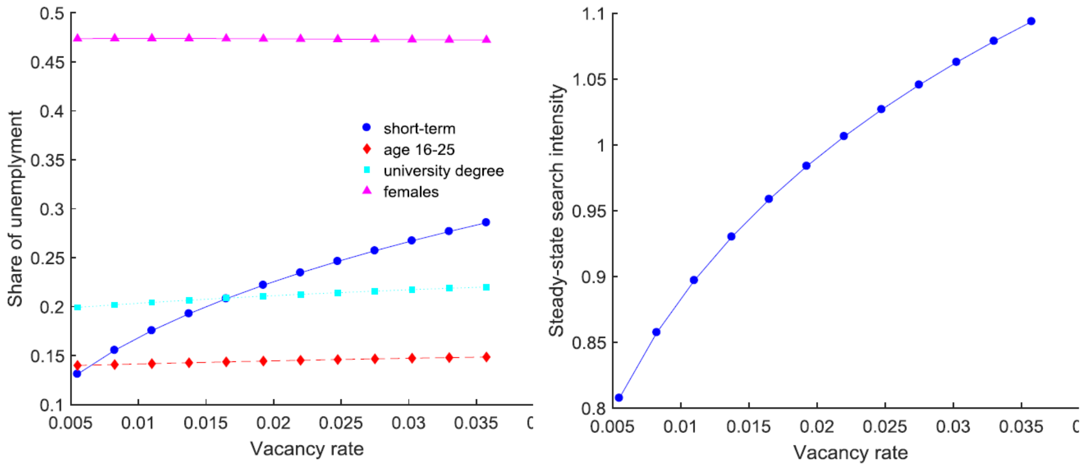

Note. The left plot shows how the share of different worker characteristics change over steady-state values of vacancies along the heterogeneous Beveridge Curve. The right plot shows the level of aggregate search intensity over steady-state values of vacancies along the heterogeneous Beveridge Curve.

from the estimated path of matching efficiency shown in the previous sections, the latter reflects demographic trends. In the Appendix, Figure 2.C.8 plots the change in the labor force share of different groups between 2016 and 2005 against their relative search intensity (based on the 2005 average within the whole labor force). The main trend emerging from the graph is the fall in the share of groups with lower intensity and the rise in the share of numerous groups with greater propensity to job matching.

2.7 Conclusion

In this work, we investigated the role of labor force heterogeneity and matching efficiency in determining unemployment-to-employment flows in the UK. Following the approach of Barnichon and Figura (2015), we expand a canonical aggregate matching function with a term representing the average search intensity of the unemployment pool. Search intensity is estimated with micro-data on individual transitions into employment using the UK labor Force Survey. In the second stage of our two-step procedure, the aggregate matching function is estimated as a state space model in which efficiency is an unobservable time-varying process.

We find that search intensity has progressively increased since the mid-90s along

Figure 14: Shift in the Beveridge curve from 2005 to 2016.

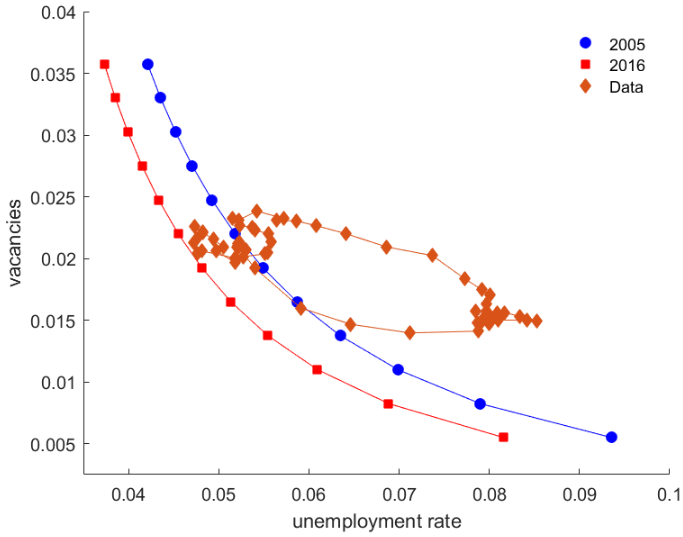

Note. The blue circles plot the heterogeneous Beverdige Curve using the 2005q1 value of matching efficiency and the 2005q1 labor force composition. The red squares plot a similar curve for 2016q1. The maroon diamonds report the empirical quarterly series for the vacancy rate (vacancies over total number of filled and vacant jobs) and the unemployment rate.

Figure 15: Decomposing the shift in the Beveridge Curve.

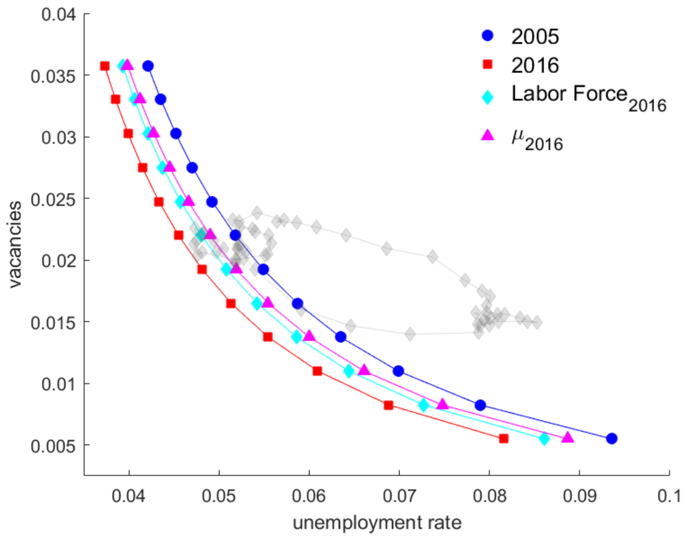

Note. The blue circles plot the heterogeneous Beverdige Curve using the 2005q1 value of matching efficiency and the 2005q1 labor force composition. The red squares plot a similar curve for 2016q1. The light blue diamonds plot the partial Beveridge Curve with the 2005 value of matching efficiency and the 2016 labor force composition. The purple triangles plots the partial Beveridge Curve with the 2005 labor force composition and the 2016 value of matching efficiency.

with rising educational attainment and an increasingly younger labor force. Additionally, the duration dependence of job finding probabilities entails that aggregate search intensity also contains a quantitatively important cyclical component associated with the share of short-term unemployment. The estimated path of matching efficiency shows a downward trend that begins prior to the 2008 recession, partially recovering since then.

We also consider two extensions to the baseline model. First, we expand the sample to on-the-job-search workers and marginally attached individuals, two groups of workers with very different characteristics from the unemployed. The results on the expanded pool show an even larger drop in matching efficiency in the first part of the sample. The key takeaway from this exercise is that the outcome of the analysis depends on the boundaries of the definition of a job seeker. The second extension models heterogeneity in recruiting capacity across industries with a similar approach. We find that a fall in aggregate recruiting intensity may provide a partial explanation for the decrease in matching efficiency up to 2002 from in the baseline model.

In the final section, we quantify how our estimates of search intensity and matching efficiency may have shifted the Beveridge Curve over time. Labor heterogeneity is shown to increase the elasticity of steady state unemployment with respect to vacancies due to endogenous changes in the steady-state duration distribution of unemployment along the curve, which acts to pivot the Beveridge Curve relative to a canonical model of homogeneous job searchers. In the last decade, improvements in the composition of the labor force and a partial recovery in matching efficiency are likely to have shifted the Beveridge Curve inward.

The methodology and results of this work leave scope for future research in multiple directions. First, a natural additional step would be to strengthen the evidence on the joint effect of labor market heterogeneity on both the matching process and wage formation. In this realm, an objective of particular interest would be to understand the importance of labor force composition for the recent "productivity puzzle" of the UK. A second important direction of work would entail disentangling the effect of observed and unobserved heterogeneity, as done by Ahn and Hamilton (2016) for the US, while also estimating a structural matching function of the aggregate economy.

2.A State space form, Kalman filtering and smoothing

This appendix presents the state space formulation of the aggregate matching function. We also briefly outline the use of the Kalman filter for the estimation of the parameters via maximum likelihood and the Kalman smoother to recover an estimate of the unobserved state. We leave the details of these techniques to more extensive references, such as Koopman et al. (1999).

For a given variable x_t, let $\hat{x}_t$ be its natural logarithm. For the period 1995q1-2000q4 we use the log job finding rate $\hat{f}_t$ and the change in the log of vacancies measured from the job center data $\Delta\hat{v^a}_t$ as the signal vector. The unobserved state vector includes the log efficiency $\hat{\mu}_t$, the log of the true vacancy value $\hat{v}_t$, and the one-quarter lag of the latter. Two further exogenous observed states are the log search intensity $\hat{s}_t$ and the log level of job seekers $\hat{u}_t$. Hence the specification is as follows:

$$\begin{bmatrix} \hat{f}_t \\ \Delta\hat{v^a}_t \end{bmatrix} = \begin{bmatrix} 1 & (1-\eta) & 0 \\ 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} \hat{\mu}_t \\ \hat{v}_t \\ \hat{v}_{t-1} \end{bmatrix} + \begin{bmatrix} \eta & -(1-\eta) \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{s}_t \\ \hat{u}_t \end{bmatrix} + \begin{bmatrix} \epsilon_t^1 \\ \epsilon_t^2 \end{bmatrix} \tag{13}$$

Starting in 2001, we use the vacancies series from the ONS $\hat{v^b}_t$ and assume that it equals the true vacancy level. Hence the specification is as follows:

$$\begin{bmatrix} \hat{f}_t \\ \hat{v^b}_t \end{bmatrix} = \begin{bmatrix} 1 & (1-\eta) \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \hat{\mu}_t \\ \hat{v}_t \end{bmatrix} + \begin{bmatrix} \eta & -(1-\eta) \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \hat{s}_t \\ \hat{u}_t \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} \epsilon_t^1 \tag{14}$$

Both states are assumed to be unit roots:

$$\begin{bmatrix} \hat{\mu}_t \\ \hat{v}_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \hat{\mu}_{t-1} \\ \hat{v}_{t-1} \end{bmatrix} + \begin{bmatrix} \nu_t^1 \\ \nu_t^2 \end{bmatrix} \tag{15}$$

The set of parameters to be estimated is $\Gamma = \{\eta, \sigma_{\epsilon 1}, \sigma_{\epsilon 2}, \sigma_{\nu 1}, \sigma_{\nu 2}\}$, where the last four items are the standard deviations of the error terms, which are assumed to be independent from each other and over time and normally distributed with mean zero.

The estimation of the parameters via maximum likelihood applies the Kalman filter interpretation to the two equations above. Let the signal equation be

$$y_t = Az_t + B\alpha_t + \epsilon_t,$$

the unobserved state equation be

$$\alpha_t = C\alpha_{t-1} + \nu_t,$$

and the variance of all disturbances be

$$\Omega = \mathrm{Var} \begin{bmatrix} \epsilon_t \\ \nu_t \end{bmatrix} = \begin{bmatrix} H & 0 \\ 0 & Q \end{bmatrix}$$

Given a set of parameters $\gamma \in \Gamma$, the Kalman filter, whose details are omitted, consists in obtaining $\alpha_{t|t}$ and $P_{t|t}$, which are the expected value and the variance of the unobserved state α_t based on the observed signal y_t and using all information on the state from period $t-1$. In particular, let

$$\alpha_{t|t-1} = \mathbb{E}_{t-1}\{\alpha_t\} = C\alpha_{t-1|t-1}$$

$$P_{t|t-1} = \mathbb{E}_{t-1}\{(\alpha_t - \alpha_{t|t-1})(\alpha_t - \alpha_{t|t-1})'\} = CP_{t-1|t-1}C' + Q$$

Based on the above equations, the Kalman filter can be used to compute a one-step ahead expectation of the state y_t:

$$y_{t|t-1} = \mathbb{E}\{y_t|\alpha_{t|t-1}, z_t\} = Az_t + B\alpha_{t|t-1}.$$

The above implies that the one-step ahead prediction error and its variance are

$$\tilde{\epsilon_t} = \epsilon_{t|t-1} = y_t - y_{t|t-1},$$

$$\tilde{F_t} = F_{t|t-1} = \text{Var}\{\tilde{\epsilon_t}\} = BP_{t|t-1}B' + H.$$

Since both the forecast error and its variance depend on the parameters of the state space model, these are estimated by maximizing the likelihood function of the sequence of forecast errors $\{\tilde{\epsilon_t}\}_1^T$. For a given set of parameters γ, the log-likelihood function is:

$$L(\gamma) = -\frac{T}{2}\log(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\log\left(|\tilde{F_t}|\right) - \frac{1}{2}\sum_{t=1}^{T}\tilde{\epsilon}_t'\tilde{F}_t^{-1}\tilde{\epsilon}_t$$

.

The maximum likelihood estimation hence consists in finding

$$\hat{\gamma} = \text{argmax}L(\gamma).$$

The Kalman smoothing procedure is then carried out based on the estimated parameters $\hat{\gamma}$. While filtering aims at finding the expected value of α_t only using all information up to time t, Kalman smoothing takes advantage of the entire information regarding the signal available to the econometrician: i.e. the full series $y^T = \{y_t\}_1^T$. In other words, it allows to find

$$\alpha_{t|T} = \mathbb{E}\{\alpha_t|y^T\}.$$

2.B Beveridge Curve computation

This section describes the method used to compute the Beveridge Curve accounting for worker heterogeneity with respect to demographic groups and unemployment duration.

The solution hinges on finding a measure of unemployment U_j of each group j defined by categories of age, sex, education, and duration, which is consistent with the a steady-state value of market tightness $\theta = V/U$, where $U = \sum_j U_j$ is aggregate unemployment.

For simplicity, age is assumed to be fixed, so that workers do not transition across 10-year age groups. Given the broad age brackets, this simplification is not crucial for the results. Sex and education are also assumed to be fixed. The only characteristic for which transitions have to be computed is duration of unemployment. The LFS contains five duration categories: less than one quarter, between 1 and 2 quarters, between 2 and 4 quarters, between 4 and 8 quarters, and more than 8 quarters. The main challenge is that job finding probabilities depend on duration of unemployment and that some LFS duration groups encompass multiple quarters. This requires appropriately computing a set of transition equations from one duration group to the next. Below is a general sketch of the procedure.

Let e_j be the measure of workers of a given type $j \in \mathbb{J}$ defined over sex, age, and education. The total labor force is normalized to 1, such that $\sum_j e_j = 1$. Let U_j be the total unemployment of that type and U_{jt} be the unemployment level at duration group $t \in \mathbb{T}$, such that $\sum_t U_{jt} = U_j$. Based on the β coefficients estimated in Section 2.4, let s_{jt} be the search intensity of each type-duration group. For a given θ and aggregate search intensity s, the discrete-time job finding probability is $F_{jt} = \big(1 - \exp(-\frac{s_{jt}}{s}\mu s^{\sigma}\theta^{1-\sigma})\big)$.

For a given level of vacancies V, we guess a set of unemployment measures $U_{jt}^{(0)}$ over types $\mathbb{J}$ and duration groups $\mathbb{T}$. Based on these measures we compute total unemployment $U^{(0)} = \sum_j \sum_t U_{jt}^{(0)}$ and aggregate search intensity $\overline{s} = \sum_j \sum_t \frac{U_{jt}^{(0)}}{U^{(0)}} s_{jt}$. Through these, each job finding probability F_{jt} can be computed. Based on the above, a new iteration $U_{jt}^{(1)}$ of the unemployment measures can be computed. To be consistent with the duration bins available in the LFS, the following transition equations are necessary:

$$
\begin{aligned}
U_{j1}^{(1)} &= (e_j - U_j^{(0)})\delta && \text{group 1: less than 1 quarter}\\
U_{j2}^{(1)} &= (1 - F_{j1})U_{j1}^{(1)} && \text{group 2: between 1 and 2 quarters}\\
U_{j3.1}^{(1)} &= (1 - F_{j2})U_{j2}^{(1)} && \text{group 3.1: between 2 and 3 quarters}\\
U_{j3.2}^{(1)} &= (1 - F_{j3})(1 - F_{j2})U_{j2}^{(1)} && \text{group 3.2: between 3 and 4 quarters}\\
U_{j4.1}^{(1)} &= (1 - F_{j3})U_{j3.2}^{(1)} && \text{group 4.1: between 4 and 5 quarters}\\
U_{j4.2}^{(1)} &= (1 - F_{j4})(1 - F_{j3})U_{j3.2}^{(1)} && \text{group 4.2: between 5 and 6 quarters}\\
U_{j4.3}^{(1)} &= (1 - F_{j4})^2(1 - F_{j3})U_{j3.2}^{(1)} && \text{group 4.3: between 6 and 7 quarters}\\
U_{j4.4}^{(1)} &= (1 - F_{j4})^3(1 - F_{j3})U_{j3.2}^{(1)} && \text{group 4.4: between 7 and 8 quarters}\\
U_{j5}^{(1)} &= \frac{(1 - F_{j4})U_{j4.4}^{(1)}}{F_{j5}} && \text{group 5: 8 quarters or more}
\end{aligned}
$$

where δ is the separation rate. Some of the groups are then aggregated into the LFS categories:

$$
\begin{aligned}
U_{j3}^{(1)} &= U_{j3.1}^{(1)} + U_{j3.2}^{(1)}\\
U_{j4}^{(1)} &= U_{j4.1}^{(1)} + U_{j4.2}^{(1)} + U_{j4.3}^{(1)} + U_{j4.4}^{(1)}
\end{aligned}
$$

After all the U_{jt}'s have been computed in this way, convergence is checked as $\sum_j \sum_t \|U_{jt}^{(0)} - U_{jt}^{(1)}\| < \epsilon$ for a tolerance value $\epsilon > 0$. If the condition is not satisfied, we restart the process using $U_{jt}^{(1)}$ as the initial guess. If the condition is satisfied, then the guess $U_{jt}^{(0)}$ is the set of steady-state measures of unemployment for each type j and duration group t consistent with the value V of vacancies.

By repeating this process for multiple values of V, the Beveridge Curve can be recovered.

2.C Additional tables and figures

Table 2.C.1: Sensitivity analysis of first-stage MLE estimation

	(1)		(2)		(3)		(4)	
	No GTS UE flow		Previously Inactive Control		Duration since last job		Quit and Layoff Controls	
	Coef.	P-Val.	Coef.	P-Val.	Coef.	P-Val.	Coef.	P-Val.
Reference group	0.245	0.000	0.465	0.000	0.257	0.000	0.527	0.000
Female	0.005	0.900	0.033	0.363	-0.010	0.770	-0.058	0.166
Age 26-35	-0.054	0.250	-0.037	0.405	-0.122	0.010	-0.160	0.002
Age 36-45	0.014	0.780	0.038	0.440	-0.078	0.110	-0.139	0.013
Age 46-55	-0.183	0.000	-0.190	0.001	-0.287	0.000	-0.200	0.002
Age 56-65	-0.545	0.000	-0.611	0.000	-0.697	0.000	-0.639	0.000
Other qual.	0.268	0.000	0.235	0.000	0.246	0.000	0.124	0.053
GCSE qual.	0.370	0.000	0.324	0.000	0.371	0.000	0.237	0.000
A-level qual.	0.415	0.000	0.321	0.000	0.380	0.000	0.258	0.000
Higher educ.	0.650	0.000	0.608	0.000	0.664	0.000	0.438	0.000
3-6 months	-0.306	0.000	-0.272	0.000	-0.015	0.770	-0.422	0.000
6 month-1 year	-0.663	0.000	-0.589	0.000	-0.257	0.000	-0.667	0.000
1-2 years	-1.025	0.000	-0.962	0.000	-0.681	0.000	-0.891	0.000
2+ years	-1.724	0.000	-1.493	0.000	-1.192	0.000	-1.559	0.000
Previously Inactive			-0.384	0.000				
Quit							0.141	0.026
Layoff							0.102	0.029
Sample	16,961		16,961		16,961		11,631	
Log-pseudolik.	-4,509		-4,857		-4,978		-4,080	

The reference group is comprised of male individuals with no GCSE qualification aged 16 to 25 who are short-term unemployed. The first column excludes Government Training Schemes from the unemployment exit flows. The second column includes a dummy variable for having spent time out of the labor force prior to unemployment. The third column uses effective duration of the joblessness spell rather than effective duration of unemployment. The fourth columns adds controls for whether the separation from the last job was a voluntary quit or a layoff (the reference group includes separations for family and health reasons, other reasons, and those who never had a job). Because the variables on reasons for separation changed over the course of the LFS and due to the large fractions of missing values in the early years, this specification is estimated on the year 1996. The reported p-values are for a two-tailed test.

Table 2.C.2: Second-stage state space estimation results for the robustness checks.

	No GST UE flows		Inactivity Spell	Effective Duration	Quits & Layoffs
1-η	0.293	0.358	0.291	0.314	0.267
	(0.045)	(0.041)	(0.050)	(0.047)	(0.052)
Log likelihood	248.01	243.46	245.42	246.82	234.82
Observations	90	90	90	90	90
Search Intensity	Yes	No	Yes	Yes	Yes
Job seeker definition	U	U	U	U	U

Note. The specification in the first column excludes Government Training Schemes from the unemployment exit flows. The second column includes a dummy variable for having spent time out of the labor force prior to unemployment in the first-stage estimation. The third column uses effective duration of the joblessness spell rather than effective duration of unemployment in the first-stage estimation. The fourth column includes dummy variables for voluntary quits and layoffs in the first-stage estimation. Standard errors are reported in parenthesis.

Table 2.C.3: Coefficients from the first-stage Maximum Likelihood Estimation in the expanded sample.

	Unemployed		Marginally Attached		On-the-job Search	
	Coef.	P-val.	Coef.	P-val.	Coef.	P-val.
Reference group	0.673	0.000				
MA or OJS dummy			-1.298	0.000	-0.409	0.000
Female	-0.047	0.188	0.081	0.475	0.127	0.057
Age 26-35	-0.097	0.027	-0.316	0.013	-0.407	0.000
Age 36-45	-0.016	0.740	-0.570	0.000	-0.520	0.000
Age 46-55	-0.227	0.000	-0.559	0.000	-0.525	0.000
Age 56-65	-0.640	0.000	-0.692	0.001	-0.521	0.055
Other qual.	0.244	0.000	-0.087	0.564	-0.456	0.000
GCSE qual.	0.333	0.000	0.186	0.162	-0.483	0.000
A-level qual.	0.331	0.000	0.205	0.161	-0.505	0.000
Higher educ.	0.599	0.000	0.344	0.037	-0.914	0.000
3-6 months	-0.259	0.000				
6 month-1 year	-0.568	0.000				
1-2 years	-0.912	0.000				
2+ years	-1.433	0.000				

Number of obs. = 33,065
Wald $\chi^2(33) = 2182.9$
Prob. $> \chi^2 = 0.0000$
Log pseudolikelihood = -7910.47

The expanded pool of job seekers includes marginally attached (MA) individuals and on-the-job-search (OJS) workers. The reference group is comprised of male individuals with no GCSE qualification aged 16 to 25 who are short-term unemployed. All coefficients on sex, education, and age are estimated separately for each type of job seeker. The reported p-values are for a two-tailed test.

Table 2.C.4: Second-stage state space estimation results using an AR(1) process for matching efficiency.

Parameter	**(1) With search intensity**	**(2) Without search intensity**
1-η	0.293	0.346
	(0.045)	(0.042)
α	-0.061	-0.086
	(0.046)	(0.053)
ρ	0.940	0.891
	(0.045)	(0.069)
Log likelihood	236.946	236.648
Observations	90	90
Search Intensity	Yes	No

Figure 2.C.1: Distribution of estimated search intensity among job seekers in 1994.

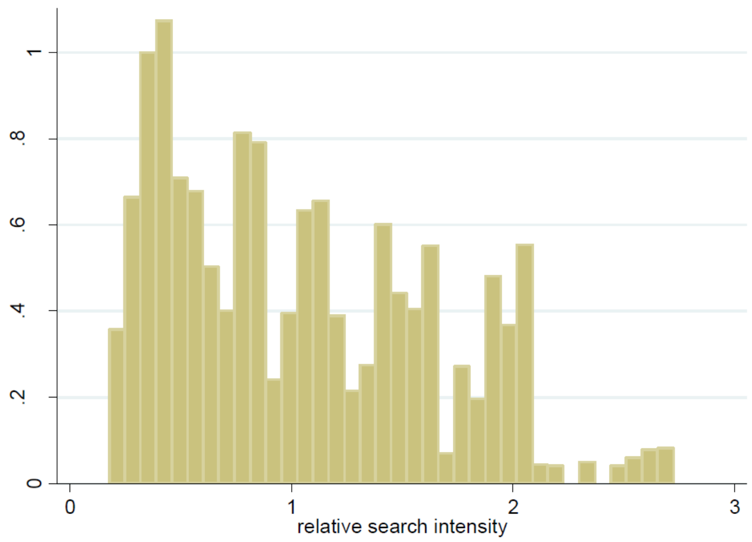

Note. The histogram displays the distribution of search intensities (s_j) in the 1994q1 sample obtained from the first-stage micro-level estimation.

Figure 2.C.2: Smooth estimate of vacancies from the state space estimation.

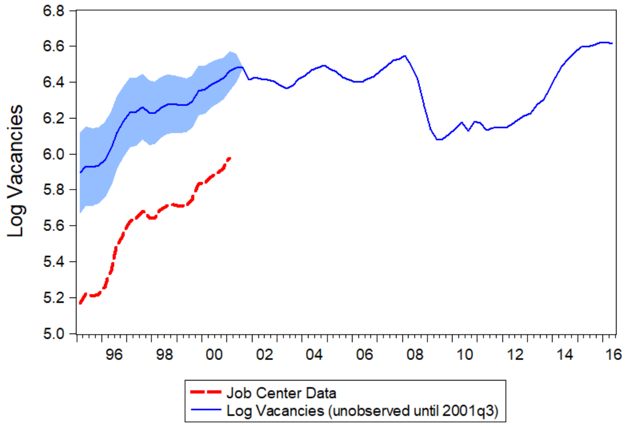

Note. The red line represents the vacancies series taken from job center data. The blue line represents the exogenous smooth path of vacancies estimated via the state space model, which coincides with the ONS series starting from 2001. Both series are reported in hundreds of thousands.

Figure 2.C.3: Fraction of Government Training Schemes among all new job matches from unemployment.

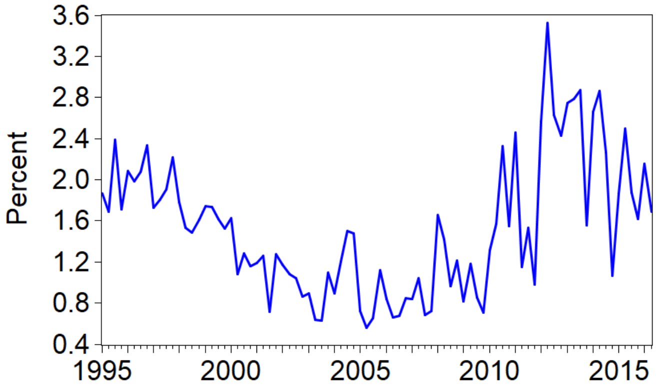

Note. The solid blue line reports the share of Government Training Schemes among all transitions from unemployment to employment in a given quarter.

Figure 2.C.4: Smooth estimate path of matching efficiency for four first-stage robustness checks.

(a) Excluding government training schemes from unemployment-to-employment transitions

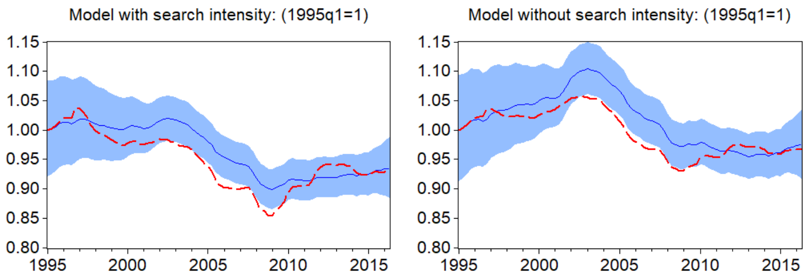

(b) Adding dummy variable for spell of inactivity.

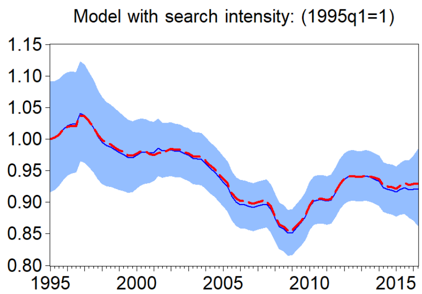

(c) Using effective duration of joblessness instead of duration of unemployment.

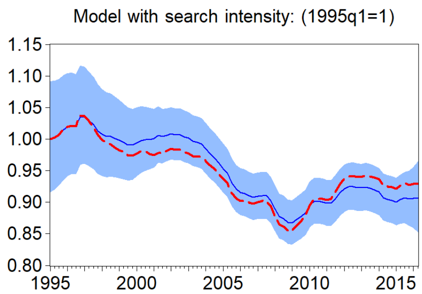

(d) Adding dummies for voluntary quits and involuntary layoffs.

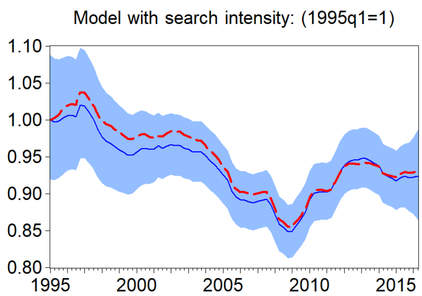

Note. the dark blue line reports the Smooth estimate of the unobserved state variable μ_t, normalized to 1 in the first period of the sample, while the dark blue area represents the 95 percent confidence interval. The red dashed line reports the forecast for the baseline model.

Figure 2.C.5: Smooth estimate path of matching efficiency for four first-stage robustness checks.

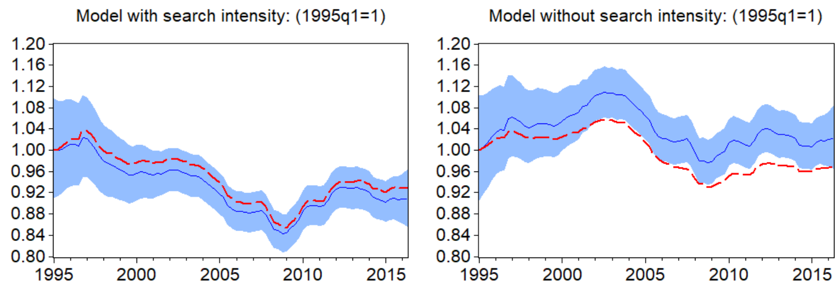

Note. the dark blue line reports the Smooth estimate of the unobserved state variable μ_t, normalized to 1 in the first period of the sample, while the dark blue area represents the 95 percent confidence interval. The red dashed line reports the forecast for the baseline model.

Figure 2.C.6: Estimated coefficients from the first-stage MLE in different years.

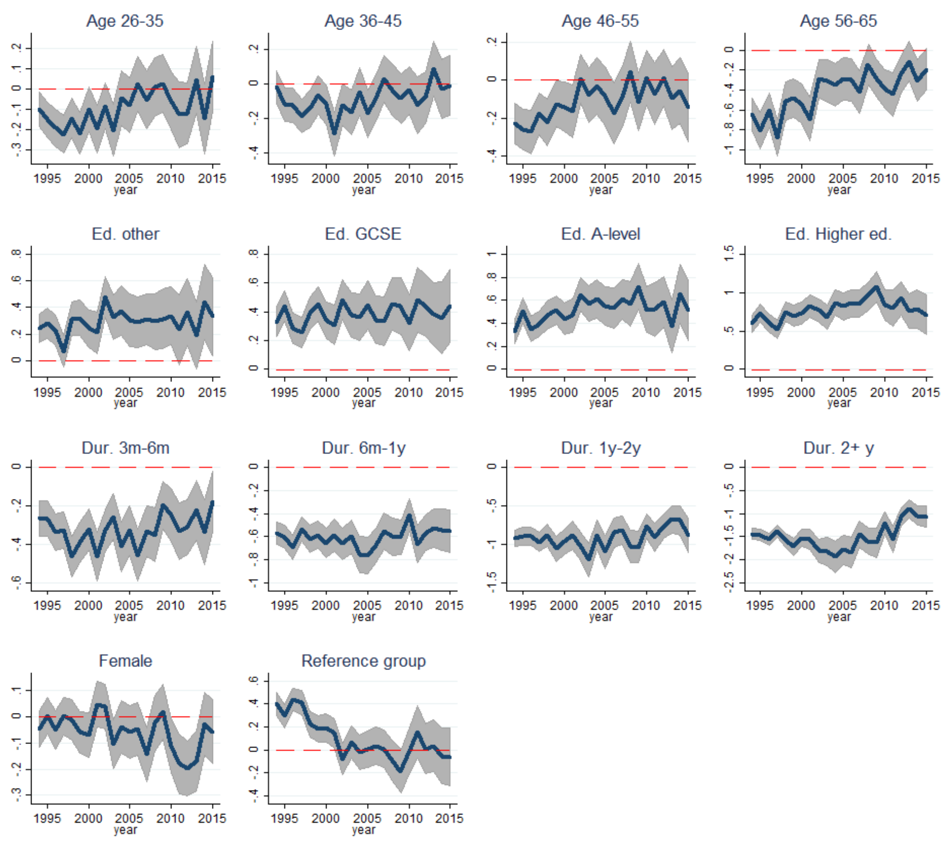

Note. Each plot reports the estimated value for the coefficient of the respective variable in the first-stage MLE carried out over different years. The 95 percent confidence intervals are reported in grey.

Figure 2.C.7: Aggregate search intensity with parameters estimated in different years.

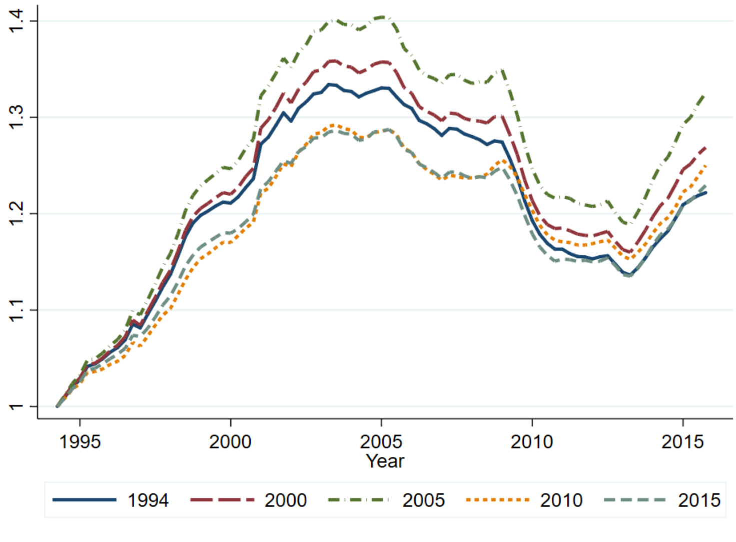

Note. Each line represents the path of aggregate search intensity s_t computed using the coefficients estimated on different baseline years. All series are normalized to be equal to one in 1994q1 and are reported in 4-quarter moving averages.

Figure 2.C.8: Changes in labor force share between 2005 and 2016 against relative search intensity in 2005.

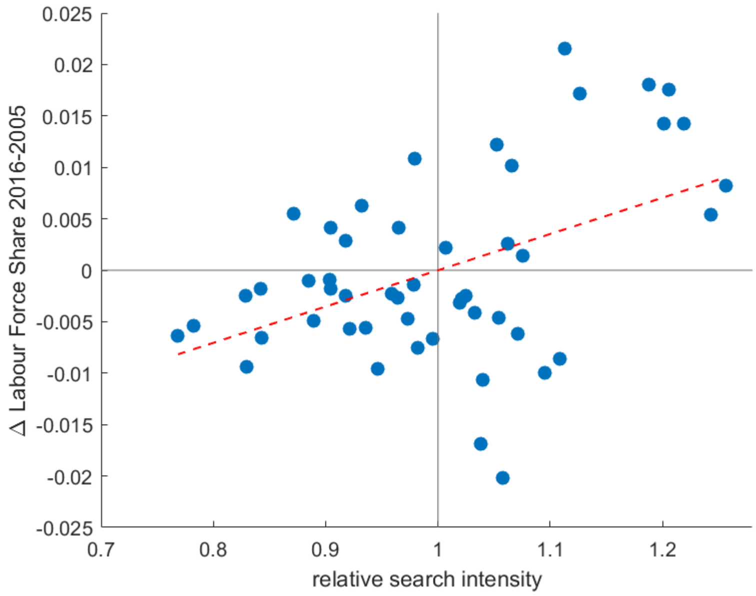

Note. The x-axis reports search intensity relative to 2005 for each worker type defined by age, sex, and education categories. The y-axis reports the change in the share of each group from 2005 to 2016. The dashed red line reports the least squares best fit of the relationship.

3 State dependence in labor market fluctuations: evidence, theory, and policy implications [1]

3.1 Introduction

Numerous studies show that fluctuations in macroeconomic variables differ across phases of the business cycle.[2] This paper builds on this strand of research and identifies systematic changes in the cyclical properties of labor market variables that are linked to the state of aggregate productivity: movements in the unemployment rate and the job separation rate are considerably larger in periods of low productivity.

To explain this novel finding and assess the policy implications, we develop a Diamond-Mortensen-Pissarides (DMP) search model with endogenous job separation and on-the-job search (OJS) that entails state dependence in labor market fluctuations across different states of aggregate productivity (i.e., low versus high) and in reaction to contractionary and expansionary shocks. In a state with low aggregate productivity, the firm retains profits by setting a high threshold of individual productivity that yields a surplus in the match, dismissing jobs with individual productivity below the threshold. Under standard assumptions, the high threshold is associated with a high density of jobs. Therefore an exogenous movement in aggregate productivity that changes the threshold generates large shifts in the job separation rate and the unemployment rate in periods of low productivity. Conversely, in a state with high aggregate productivity, the firm sets a low threshold for efficient matches that is associated with a low density of jobs and therefore an equivalent change in productivity produces limited movements in the job separation and unemployment rates.

The model shows that labor market responses are different across contractionary

1 This chapter is co-authored with Francesco Zanetti (University of Oxford). I would like to thank seminar participants the University of Oxford and the European Economic Association (Geneva) for extremely valuable comments.

2 Seminal studies are those by Neftci (1984), Sichel (1989) and Beaudry and Koop (1993), followed by more recent studies by McKay and Reis (2008), Bachmann et al. (2013) and Mumtaz and Surico (2015).

and expansionary shocks. In the aftermath of a contractionary productivity shock, the firm raises the threshold of individual productivity for efficient matches and terminates jobs with productivity lower than the threshold, instantly increasing the job separation rate. In response to a positive shock, the firm decreases the threshold but the fall in job terminations is more moderate since the productivity of most workers is higher than the threshold, leaving them unaffected by the more relaxed reservation productivity. This mechanism makes the job separation rate and the unemployment rate more responsive to contractionary productivity shocks than to expansionary ones.

To study the impact of a policy which directly interacts with the separation margin, we enrich the model with layoff taxes. The tax is levied on the firm for the termination of existing jobs but averted on matches that fail to result in job relations. [3] The layoff tax increases the surplus of job relations that continue into the next period since they forego the tax payment but it reduces the surplus of new matches in the prospect of paying the layoff tax in the future. The tax commands a high threshold of individual productivity to make new matches profitable and therefore discourages OJS. The pool of job seekers diminishes, and the firm's recruiting costs for establishing a profitable match rise, leading to a decrease in hiring. Overall, the layoff tax considerably reduces the job finding rate and increases the pool of workers subject to job separation by discouraging OJS. These complementary forces generate a rise in the unemployment rate. In particular, the effect of the tax crucially hinges on the endogenous response of OJS.[4]

To illustrate the importance of state dependence coming from job destruction, we use the model with layoff taxes to assess whether an unexpected and permanent removal of the tax generates distinct transitional dynamics and welfare effects in states with low and high aggregate productivity. In the long run, the elimination of the tax generates a fall in the unemployment rate and a rise in output that is welfare-enhancing regardless of the initial states of aggregate productivity. In the short run, however, the reform generates sharp differences in the transitional dynamics of labor

3 To the best of our knowledge, this is the first study to introduce firing taxes in a general equilibrium model with endogenous job separation and OJS. Layoff taxes are modeled as a deadweight loss, similar to Pissarides (2000). We focus on layoff taxes since Cacciatore and Fiori (2016) show that they are effective policies in reducing inefficiencies of unemployment fluctuations. In addition, an array of studies show that they are powerful in affecting labor market outcomes (see, for example, Campolmi and Faia (2011), Zanetti (2011) and references therein).

4 As we discuss in section 3.6.2, there is no established consensus on the effect of layoff taxes on the level of the unemployment rate. We contribute to this realm of research by showing that OJS plays a critical role for the effect of layoff taxes on unemployment.

market variables across initial states. The unemployment rate gradually declines in the state with high aggregate productivity whereas it suddenly contracts in the state with low aggregate productivity. These temporary differences disappear after four quarters, but they produce significant welfare differences. The tax removal raises the surplus of establishing a job relation and induces firms to post vacancies and workers to search on the job to extract the enhanced benefits of forming an employment relation. Search efforts are stronger in the state with high aggregate productivity since the joint surplus of forming a job is larger while aggregate productivity is high. The considerable rise in search efforts in the state with high aggregate productivity generates large temporary welfare losses caused by deadweight search costs of matching frictions. The total discounted welfare gain of a tax removal enacted in the state with low aggregate productivity is 3.5 percent larger than the same reform in the state with high aggregate productivity.

Our analysis relates to empirical and theoretical studies on the asymmetry of fluctuations over the business cycle. On the empirical side, the studies by Neftci (1984), Altissimo and Violante (2001), Panagiotidis and Pelloni (2007), Barattieri et al. (2014), Benigno et al. (2015) and Caggiano et al. (2014) show that unemployment and wages move differently across phases of the business cycles. Compared to these studies, we establish state dependence in labor market fluctuations linked to the level of aggregate productivity, and we extend the analysis to job transition rates. On the theoretical side, the analysis is related to studies that use structural models to investigate state dependence and nonlinearities of macroeconomic variables over the business cycle. In particular, several works examine asymmetric dynamics through the lens of search models. Petrosky-Nadeau and Zhang (2017), Kohlbrecher and Merkl (2016), and Ferraro (2016) develop search models that replicate labor market asymmetries over the business cycle.[5] Different from the these studies, our analysis considers state dependence rather than skewness in unconditional distributions. Furthermore, we include OJS in the model and show that it plays an important role for the asymmetries because of its interaction with firms' job destruction decisions.

Finally, this work relates to the growing literature exploring the state-dependent

5 Outside of the search and matching literature, Kim and Ruge-Murcia (2009), Aruoba et al. (2013) and Gortz and Tsoukalas (2013) show that standard models replicate important nonlinearities in the data if enriched with asymmetric adjustment costs. Abbritti and Fahr (2013), Benigno and Ricci (2011) and Benigno et al. (2015) show that asymmetric wage rigidities change the propagation of exogenous disturbances and have important welfare effects for stabilization policies.

effect of labor market policies on aggregate fluctuations. In the context of search and matching, Michaillat (2014) develops a model of the labor market that partitions jobs into public and private sectors and shows that the effect of government spending is nonlinear and less effective in periods of economic expansions since public employment crowds out private employment. Focusing on the timing of structural refoms, Cacciatore and Fiori (2016), Cacciatore et al. (2015), and Cacciatore et al. (2016) show that the effect of labor market deregulation changes nonlinearly with the level of unemployment and that the responses of labor market variables are different across phases of the business cycles. Our analysis complements these findings by focusing on a search model with OJS and on the transitional dynamics and welfare effects of the removal of layoff taxes implemented at distinct states of aggregate productivity. We show that the interaction of the firing tax with endogenous OJS decisions is crucial for the long-run impact on aggregate unemployment. With regards to transitional dynamics, OJS is quantitatively important because it affects the aggregate increase in market tightness once the tax is removed.

The remainder of the paper is structured as follows. Section 3.2 presents the empirical findings. Section 3.3 lays out the model. Section 3.4 discusses the mechanisms that generate state dependence in labor market fluctuations across distinct states of aggregate productivity and in response to contractionary and expansionary shocks. Section 3.5 presents model simulations and quantitative results. Section 3.6 assesses the implementations of labor market reforms. Section 3.7 concludes.

3.2 Empirical evidence

This section isolates systematic differences in the fluctuations of labor market variables linked to the state of productivity and the resulting skewness in the distribution of unemployment.

We use quarterly series for the log (un)employment rate, the job finding rate, the job separation rate, log output and log labor productivity over the period 1950:I-2014:IV. To extract the cyclical component of variables, we use an HP filter with a smoothing parameter equal to 1,600.[6] Panel (a) in Figure 1 plots cyclical movements in the unemployment rate (solid line) against periods of labor productivity below the median value (shaded area). Especially in the early part of the sample,

6 Appendix 3.A.1 provides details on the data sources.

the largest movements in the unemployment rate occur during period of below-trend productivity. This dynamic is less pronounced since the early 1990s, where longer periods of slow decreases in unemployment have been named "jobless recoveries." Panel (b) in the figure shows histograms of the unemployment rate in periods with labor productivity below (left entry) and above (right entry) the historical median. A comparison across the two histograms reveals that changes in the unemployment rate are significantly larger in periods of aggregate productivity below the median value. This first pass to the data outlines systematic differences in the variability of the unemployment rate that are linked to the state of productivity. In the subsequent analysis, we use statistical methods to isolate significant changes in the cyclical properties of a wider range of labor market variables across distinct states of productivity.

Figure 1: The unemployment rate and states of average labor productivity

(a) Cyclical component

(b) Growth rates

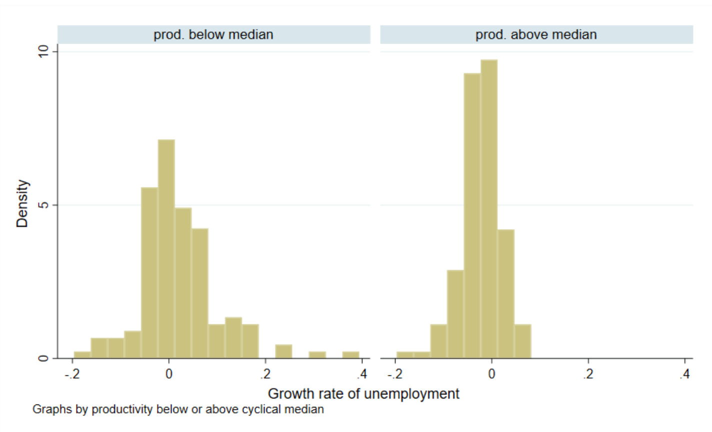

Note. Panel (a) shows the cyclical movements in the unemployment rate (solid line) together with periods of labor productivity below the median value (shaded area). Panel (b) shows histograms of the unemployment rate when productivity is below (left entry) and above (right entry) the median value.

Table 1: Standard deviation for different states of productivity

	Level			Growth Rate		
	p< 50^{th} percentile (1)	p> 50^{th} percentile (2)	$\frac{\sigma_{p<50}}{\sigma_{p>50}}$ (3)	p< 50^{th} percentile (4)	p> 50^{th} percentile (5)	$\frac{\sigma_{p<50}}{\sigma_{p>50}}$ (6)
Unemployment	0.146	0.120	1.23	0.0829	0.0401	2.06
Job Finding Rate	0.037	0.034	1.07	0.0240	0.0192	1.25
Job Separation Rate	0.0021	0.0013	1.57	0.0019	0.00157	1.25
Employment	0.009	0.007	1.32	0.0051	0.0023	2.15
Output	0.022	0.016	1.34	0.0153	0.0090	1.71
Productivity	0.012	0.009	1.25	0.0099	0.0074	1.34

Note. Appendix 3.A.1 provides data sources. The data is quarterly over the period 1950:I-2014:IV. The series of the (un)employment rate, output and productivity are in logs. Series are HP-filtered with a smoothing parameter equal to 1,600. Growth rates are log differences of quarterly averages. The third and sixth columns report the ratios of the standard deviations below and above the historical median of productivity.

Table 1 shows the standard deviation of the HP-filtered series (Columns 1-3) and growth rates of the series (Columns 4-6) over the sample period. The table reports the standard deviation of the variables when labor productivity is below (Columns 1 and 4) and above (Columns 2 and 5) the median value of its cyclical component, and the ratio of the two (Columns 3 and 6). The entries consistently show that the standard deviation of most of the variables in levels is 20 to 30 percent larger in periods of productivity below its median value, and even twice larger in growth rate. Interestingly, the difference in the volatility of the job finding rate (in levels) across states of productivity is more moderate.

We undertake a series of robustness checks, which are included in Appendix 3.I.1. First, we check that the state dependence remains when using yearly growth rates of the labor market variables as a measure of volatility (Table 3.I.1). Additionally, in Table 3.I.2 we show that results continue to hold if we use a smoothing parameter for the HP filter equal to 10^5, as suggested by Shimer (2005), or set the regimes based on the productivity series by Fernald (2014). We also inspect that the state dependence is robust to two specific sub-periods of the data: the Great Moderation (1980-2007) and the full pre-Great Recession period (1950-2007). Third, we consider only the more marked cases of low and high productivity, using as thresholds the 25^{th} and 75^{th} percentiles of average labor productivity, respectively, thus excluding all observations in the second and third quartiles.

We further study the sensitivity to defining the low and high productivity states using alternative variables: yearly growth rates of productivity, NBER recession dates, and quarterly growth rates of productivity (both as 4-quarter moving average and in its raw series).[7] As Table 3.I.3 shows, the result is broadly robust, especially for the labor market variables in levels. The only threshold definition for which the results are not particularly robust is the raw series of quarterly growth rates of productivity. This result is explained by the fact that the quarterly growth rate is not indicative of whether the aggregate level of productivity is high or low. As Table 3.I.4, the correlation between our baseline definition of states and the quarterly growth rate is only 0.2, while the regimes based on the other definitions have higher correlations. This point is reiterated by a cross-tabulation of the identification of regimes across the alternative definitions (Table 3.I.5).

These findings further point to large and systematic differences in the variability of labor market variables across states of aggregate productivity. To ensure results are not driven by the larger volatility of labor productivity itself in periods with low aggregate productivity, we use regression analysis that estimates the elasticity of labor market variables with respect to productivity controlling for the initial level of productivity. Table 2 regresses the log (un)employment rate, the job finding rate and the separation rate on log productivity ($\log p_t$), a dummy variable equal to one when productivity is above its historical median value (High-p_t) and an interaction term between the two variables (High-p_t*$\log p_t$) that captures the differential effect of productivity in periods with high economic activity. Since the response of labor market variables to changes in aggregate conditions may be delayed, we include the explanatory variables with a lag, and to capture persistence, we also include the dependent variable with a lag.[8] Column 1 shows that the unemployment rate is negatively correlated to current-period productivity. The interaction term is positive, implying that in times of productivity above the median value the negative correlation between productivity and the unemployment rate is reduced. Column 2 shows that the job finding rate is positively correlated with current-period productivity. The coefficient in the interaction term is negative, therefore indicating that the positive correlation is smaller in states with high productivity. The weak statistical significance is consistent

7 For all series but the NBER recessions dates, the threshold is based on the median value of the variable. For the recession dates, we base the low state as the quarters of economic recession.

8 The results continue to hold if we have no lags or if we include additional (two and three) lags of the explanatory variables.

Table 2: Regression analysis, specification in levels.

Variables	$\log U_t$ (1)	JFR_t (2)	SR_t (3)	$\log E_t$ (4)
$\log p_t$	-4.834***	0.842***	-0.120***	0.294***
	(0.582)	(0.226)	(0.0166)	(0.032)
High-p_t	0.0155*	-0.001	0.000	-0.001**
	(0.009)	(0.003)	(0.000)	(0.000)
log p_t*High-p_t	2.734***	-0.679*	0.104***	-0.172***
	(0.890)	(0.353)	(0.026)	(0.050)
log p_{t-1}	-0.171	0.374	0.025	0.070*
	(0.630)	(0.244)	(0.018)	(0.036)
High-p_{t-1}	-0.003	0.000	0.000	-0.000
	(0.009)	(0.003)	(0.000)	(0.000)
log p_{t-1}*High-p_{t-1}	0.977	-0.182	-0.029	-0.111**
	(0.906)	(0.357)	(0.027)	(0.051)
Lagged dependent	0.884***	0.775***	0.345***	0.878***
	(0.024)	(0.034)	(0.0593)	(0.022)
Constant	-0.026***	0.005*	-0.000**	0.002***
	(0.007)	(0.003)	(0.000)	(0.000)
Observations	259	259	259	259
R-squared	0.890	0.762	0.479	0.911

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note. The dependent variables are the log unemployment rate, the job-finding rate, the separation rate, and the log employment rate, respectively. All time series are HP-filtered with a smoothing 1,600 parameter. The explanatory variables are the log productivity, a dummy variable equal to one when productivity at time t is larger than the historical median, and an interaction term between the log productivity and the dummy variable.

with the fact that the ratio of volatilities from Table 1 is only marginally greater than one. Column 3 shows that the separation rate responds negatively to current-period productivity. The coefficient of the interaction term is large and positive, implying a much reduced negative correlation between these variables in states with high productivity. An equivalent result holds for the employment rate in Column 4. Overall, the results from regression analysis corroborate the evidence above, pointing to larger labor market fluctuations in the state with productivity below the median value. Table 3 performs the same exercise using the growth rate of the variables and shows that results continue to point to a larger response of labor market variables in period of productivity growth below the median value.[9]

A recent study by Ferraro (2016) establishes that the employment rate is negatively

9 Appendix 3.I.2 shows that results continue to hold if we use a smoothing parameter in the HP filter equal to 10^5, the productivity series by Fernald (2014) and different sample periods.

Table 3: Regression analysis, specification in growth rates.

Variables	ΔU_t (1)	ΔJFR_t (2)	ΔSR_t (3)	ΔE_t (4)
Δp_t	-3.500***	0.777***	-0.0854***	0.193***
	(0.507)	(0.184)	(0.015)	(0.030)
High-p_{t-1}	-0.042***	0.011***	-0.001**	0.002***
	(0.011)	(0.0038)	(0.000)	(0.001)
$\Delta p_t *$ High-p_{t-1}	1.694**	-0.596**	0.067***	-0.096*
	(0.826)	(0.299)	(0.025)	(0.049)
Δp_{t-1}	-2.529***	0.762***	0.014	0.194***
	(0.555)	(0.201)	(0.017)	(0.033)
High-p_{t-2}	-0.035***	0.007*	0.000	0.00288***
	(0.010)	(0.004)	(0.000)	(0.001)
$\Delta p_{t-1} *$ High-p_{t-2}	1.555*	-0.041	-0.027	-0.142***
	(0.830)	(0.301)	(0.025)	(0.049)
Constant	0.064***	-0.016***	0.000*	-0.004***
	(0.007)	(0.002)	(0.000)	(0.000)
Observations	258	258	258	258
R-squared	0.360	0.231	0.121	0.383

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note. Variables are expressed in growth rates. Columns (1)-(4) report the growth rates of the unemployment rate, the job-finding rate, the separation rate, and the employment rate, respectively. The explanatory variables are the growth rate of productivity, a dummy variable equal to one when productivity at time t is larger than the historical median, and an interaction term between the growth rate of productivity and the dummy variable.

skewed in levels and growth rate and that the job finding and job separation rates are jointly responsible for the asymmetric distribution of the employment rate. To ensure that our analysis appraises this dimension in the data, we test for skewness in the unemployment rate, the job separation rate and the job finding rate. Table 4 shows the excess skewness test by Bai and Ng (2005) and reports the skewness coefficient for the distribution of each variable together with the associated one-tailed p-value in parentheses. The entries show statistically significant positive (negative) skewness in levels and growth rates of the unemployment (employment) rate. The job finding rate is not significantly skewed in levels and growth rates. Meanwhile, the job separation rate shows significant positive skewness both in levels and growth rates.[10] Overall, these findings suggest that fluctuations in the job separation rate display larger upward movements than downward ones, implying sharp rises during recessionary periods. The same feature is subsequently acquired by the unemployment

10 In Appendix 3.I.3, Table 3.I.8 provides extensive robustness checks on the skewness of the series.

rate.[11]

Table 4: Skewness in US labor market variables.

	HP-filtered	Growth rate
Unemployment	0.403	1.385
	(0.060)	(0.045)
Job Finding Rate	-0.292	-0.385
	(0.103)	(0.160)
Separation Rate	0.772	0.478
	(0.014)	(0.042)
Employment	-0.666	-1.214
	(0.005)	(0.010)
Output	-0.509	-0.098
	(0.023)	(0.356)
Productivity	-0.363	0.070
	(0.052)	(0.270)

Note. The table reports the skewness coefficients of US labor market variables for the period 1950:I-2014:IV. The skewness test is based on Bai and Ng (2005). The P-values for one-tailed tests are reported in parenthesis.

3.3 The model

This section lays out a DMP search and matching model with endogenous job separation and OJS. The model is based on den Haan et al. (2000), Krause and Lubik (2007), Thomas and Zanetti (2009), and Fujita and Ramey (2012). However, it differs from these studies in allowing for the separation of newly-established jobs and, in Section 3.6, by introducing layoff taxes on the termination of existing jobs.

Economic environment and timing. A continuum of households of mass 1 and a continuum of firms operate in a discrete time environment. Households supply labor to firms inelastically. Matching frictions in the labor market prevent full employment. Firms pay a fixed cost for each vacancy posted to recruit new workers. However, neither are all vacancies filled nor are all job-seekers hired in a given period. Employed workers produce a single consumption good, whose price is normalized to one, and they may search for a new job while employed. In every period t, production by a single worker depends on aggregate labor productivity, a_t, and individual, idiosyncratic productivity, x. For each a_t, there is a reservation level of individual productivity $x^r(a_t)$, below which jobs are not mutually efficient and are dismissed.

11 Appendix 3.I.3 shows that results continue to hold if we use a smoothing parameter in the HP filter equal to 10^5, the productivity series by Fernald (2014), different sample periods and the specification of the variables in growth rates.

Similarly, there is a level of individual productivity $x^s(a_t)$, below which employed workers find it efficient to pay a fixed cost to search for other jobs.[12]

Within each period t, the timing of events is as follows. At the start of period, firms post vacancies that are matched with job seekers by the end of the period. Employed workers produce and may search for a new job within the period. At the end of the period, a fraction of employed workers is exogenously separated, and another fraction of employed workers obtains a new draw of individual productivity. At the beginning of the next period $t+1$, aggregate productivity a_{t+1} and the individual productivity x of each worker are observed. Each firm converts profitable matches into jobs, and each worker that searches on the job decides whether to move to a new firm or remain in the current job.

The matching function. A matching function encapsulates search frictions in the labor market. In each period t, the constant-returns-to-scale matching function establishes the number of matches between job seekers and vacancies:

$$m_t = m(u_t + \psi_t, v_t) = \gamma(u_t + \psi_t)^{1-\eta} v_t^{\eta}, \tag{1}$$

where u_t is unemployment, ψ_t is the mass of OJS workers, v_t are vacancies, and $0 < \eta < 1$. The sum of employed OJS workers and unemployed workers forms the number of effective job searchers. The probability for a job seeker to fill a vacancy and for a vacancy to be filled can be expressed in terms of the "labor market tightness," defined as the ratio of vacancies to job seekers, $\theta_t = v_t/(u_t + \psi_t)$. The probability of a job seeker to find a suitable vacancy is $p(\theta_t) = m(u_t + \psi_t, v_t)/(u_t + \psi_t) = m\big(1, v_t/(u_t + \psi_t)\big)$ and the probability for the firm to find a suitable worker is $q(\theta_t) = m((u_t + \psi_t), v_t)/v_t = m\big((u_t + \psi_t)/v_t, 1\big)$.[13]

Production and matched workers. Each firm manufactures a unique final good by hiring labor. Each hired worker produces $a_t x$ units of output. Aggregate pro-

12 Below, we use the following notation: when x has a time subscript, it refers to an aggregate variable (e.g. x_t^r is the individual productivity threshold that applies to all firms given an aggregate productivity level a_t). Without a subscript, it refers to any individual productivity level for a match independent of aggregate states.

13 As we discuss below, due to individual productivity shocks, $p(\theta_t)$ and $q(\theta_t)$ cannot be interpreted as the job finding and job filling probabilities, respectively. We therefore refer to them as the "contact" probabilities for workers and firms, respectively. Also note that labor market tightness includes also OJS workers and hence differs from the empirically observable vacancy/unemployment ratio.

ductivity a_t follows the auto regressive process:

$$\ln a_{t+1} = \rho \ln a_t + \epsilon_{t+1}, \tag{2}$$

where $\epsilon \sim N(0, \sigma^2)$ and $\|\rho\| < 1$. During each period $t+1$, an existing worker maintains the previous individual productivity level with probability $(1-\lambda)$, and with probability λ, the worker receives a new productivity drawn from the constant distribution $F(x)$ over the domain $[x_L, x_H]$. Job seekers matched in period t also receive a productivity value from the same distribution in the beginning of period $t+1$.

Job separation and job creation. During each period t, total job separations comprise exogenous and endogenous terminations. Existing workers are separated from their jobs with the exogenous probability of $s < 1$. Given aggregate productivity a_t, the firm establishes a threshold of individual productivity $x^r(a_t)$, below which existing matches are mutually inefficient. All workers whose individual productivity satisfies $x \leq x^r(a_t)$ are dismissed whereas if $x > x^r(a_t)$, the job relation continues in the next period.

On-the-job search. A worker may search for a new job at the cost k^s. An employed job searcher is matched to a firm from the same pool as the unemployed job seekers and therefore is subject to the same matching frictions. Once matched, the worker receives an idiosyncratic x from the distribution $F(x)$ as any other newly-matched job seeker. If the draw of individual productivity is below the reservation threshold, the match is discontinued and the employed job searcher stays in the original job. Also, as any existing worker, the job searcher who remains with her current firm draws a new individual productivity with probability λ and faces exogenous job separation. Each firm applies the same separation threshold to employed and unemployed job seekers.[14]

14 This simplifying assumption abstracts from the fact that the actual outside option for employed job seekers is their current employment contract rather than unemployment. This simplification avoids the issue of heterogeneity in wage bargaining and hence the fact that new wages depend on the value of x for the current contract and the value of x from the previous employer. These dynamics would substantially complicate the aggregation for the solution of the model because the entire distribution of x over employed workers, which is history-dependent, would become a relevant state variable for firms' decisions. Within the microeconomic literature, the details of wage bargaining from on-the-job search have been considered by Postel-Vinay and Robin (2004) and Shimer (2006), among others. See also Gottfries (2018) for a recent contribution.

Recursive formulation. Four value functions solve the model: the value of unemployment (U), the value of a vacancy (V), the *joint* value of a match (M) and the *joint* surplus of a match (S). The joint surplus of a match is split in constant proportions through Nash bargaining for wages, assigning the fraction ϕ of the joint match surplus to the worker and the fraction $1-\phi$ to the firm. The value of unemployment is:

$$U(a_t) = b + \beta\mathbb{E}_t\Big[U(a_{t+1}) + p(\theta_t)\phi\int_{x_L}^{x_H} S(a_{t+1}, x')dF(x')\Big]. \tag{3}$$

Equation (3) shows that the value of unemployment is equal to the opportunity cost of working (i.e., the flow value of unemployment b) and the expected benefits that finding a job brings in the next period. In period $t+1$, the prospective worker encounters a suitable vacancy with probability $p(\theta_t)$, and, if the match is mutually profitable, the worker gains a fraction (ϕ) of the total surplus on top of value of staying unemployed. Otherwise the job seeker remains unemployed, gaining the continuation value $U(a_{t+1})$.

The value of an open vacancy is:

$$V(a_t) = -k + \beta\mathbb{E}_t\Big[V(a_{t+1}) + q(\theta_t)(1-\phi)\int_{x_L}^{x_H} S(a_{t+1}, x')dF(x')\Big]. \tag{4}$$

Equation (4) shows that the present value of an open vacancy is equal to the fixed cost of posting the vacancy (k) and the expected benefits that the vacancy brings in the next period. In period $t+1$, the firm finds a prospective worker with probability $q(\theta_t)$, and if the match is profitable, the firm gains a fraction $(1-\phi)$ of the total surplus. Otherwise, the vacancy remains open, giving the firm a continuation value $V(a_{t+1})$. In equilibrium, the free-entry condition leads firms to post vacancies until their expected value is equal to zero in each period (i.e. $V(a_t) = 0$, for all t). This equilibrium condition applied to equation (4) yields the job-creation condition:

$$\frac{k}{q(\theta_t)} = (1-\phi)\beta\mathbb{E}_t\Big[\int_{x_L}^{x_H} S(a_{t+1}, x')dF(x')\Big]. \tag{5}$$

Equation (5) shows that the expected cost of a match (left-hand side of the equation) is equal to the expected benefit that the match brings into the firm if the job is established (right-hand side of the equation). With this formulation, the problem can be recast in terms of choosing a given market tightness $\theta(a_t)$ for a level of aggregate productivity.

For each given vector of (a_t, x), an employment relationship is established if the match is mutually efficient, and therefore the joint value of establishing a job relation is greater than the value of the outside options (i.e., the individual values from separation). Thus, the joint value of a firm-worker match is:

$$M(a_t, x) = \max\Big\{M^{n,c}(a_t, x), M^{s,c}(a_t, x), U(a_t) + V(a_t)\Big\}, \tag{6}$$

where $M^{n,c}(a_t, x)$ is the joint value of a continued match without OJS, $M^{s,c}(a_t, x)$ is the joint value of the continued match with OJS, and $U(a_t) + V(a_t)$ is the joint value of the outside option.

The joint value of a continued match without OJS is:

$$\begin{aligned} M^{n,c}(a_t, x) = a_t x + \beta \mathbb{E}_t \Big\{ & U(a_{t+1}) + V(a_{t+1}) \\ & + (1-s)\Big[(1-\lambda)S(a_t, x) + \lambda \int_{x_L}^{x_H} S(a_{t+1}, x')dF(x')\Big]\Big\}. \end{aligned} \tag{7}$$

Equation (7) shows that the value of a continued match is equal to production plus the expected continuation value of the work relationship. Meanwhile the value of a continued match while searching on the job is

$$\begin{aligned} M^{s,c}(a_t, x) = a_t x - k^s + \beta \mathbb{E}_t \Big\{ & U(a_{t+1}) + V(a_{t+1}) \\ & + \Big[1 - p(\theta_t)\overline{F(x^r_{t+1})}\Big](1-s)\Big[(1-\lambda)S(a_{t+1}, x) \\ & + \lambda \int_{x_L}^{x_H} S(a_{t+1}, x')dF(x')\Big] + p(\theta_t)\phi \int_{x_L}^{x_H} S(a_{t+1}, x')dF(x')\Big\}. \end{aligned} \tag{8}$$

where $\overline{F(x^r_{t+1})} = (1 - F[x^r(a_{t+1})])$.

The last term uses the fact that $\int_{x^r(a_{t+1})}^{x_H} S(a_{t+1}, x')dF(x') = \int_{x_L}^{x_H} S(a_{t+1}, x')dF(x')$, which represents the expected surplus that may accrue to the worker if she is matched with another firm and the match is continued. This event materializes with probability $p(\theta_t)\overline{F(x^r_{t+1})}$, and encompasses all the values of x above the reservation threshold x^r_{t+1}.

The joint surplus of a match equals the value of a match, M, net of the outside option for the worker, U, and the firm, V (i.e. S=$M - U - V$). Thus, the value function for the joint surplus of a continuing match is:

$$S(a_t, x) = \max[S^{n,c}(a_t, x), S^{s,c}(a_t, x), 0], \tag{9}$$

where $S^{n,c}(a, x)$ is surplus of the match when the job relation continues without OJS and $S^{n,c}(a, x)$ is the surplus of a continued match with OJS. The surpluses are defined as follows:

$$S^{n,c}(a_t,x) = a_t x - b + \beta \mathbb{E}_t \Big\{ (1-s) \Big[(1-\lambda) S(a_{t+1},x) + \lambda \int_{x_L}^{x_H} S(a_{t+1},x') dF(x') \Big] - p(\theta_t) \phi \int_{x_L}^{x_H} S(a_{t+1},x') dF(x') \Big\}, \quad (10)$$

$$S^{s,c}(a_t,x) = a_t x - k^s - b + \beta \mathbb{E}_t \Big\{ \Big[1 - p(\theta_t) \overline{F(x^r_{t+1})} \Big] (1-s) \Big[(1-\lambda) S(a_{t+1},x) + \lambda \int_{x_L}^{x_H} S(a_{t+1},x') dF(x') \Big] \Big\}. \quad (11)$$

A worker searches while on the job if $S^{s,c}(a_t,x) \geq S^{n,c}(a_t,x)$. The presence of OJS introduces a threshold $x^S(a_t)$, below which it is efficient to search on the job. For values of the threshold $x^r(a_t) < x^S(a_t) < x_H$, it is efficient to incur in the search costs for all $x \in \big(x^r(a_t), x^S(a_t)\big]$. Substituting equations (10) and (11) into the condition $S^{s,c}(a_t,x) \geq S^{n,c}(a_t,x)$ yields

$$k^s \leq \beta \mathbb{E}_t \Big\{ - p(\theta_t) \overline{F(x^r_{t+1})} (1-s) \Big[(1-\lambda) S(a_{t+1},x) + \lambda \int_{x_L}^{x_H} S(a_{t+1},x') dF(x') \Big] + p(\theta_t) \phi \int_{x_L}^{x_H} S(a_{t+1},x') dF(x') \Big\}. \quad (12)$$

with equality, equation (12) determines the efficient threshold under which workers engage in OJS. Intuitively, the cost of searching has to be smaller than the increase in the continuation value coming from possibly finding a new match.

Finally, for a given aggregate state a_t, the individual productivity threshold for exogenous separations is the value of x which makes the joint surplus of continuing a match equal to zero, such that[15]

$$S^c(a_t, x^r(a_t)) = 0. \quad (13)$$

Labor flows and transition rates. Labor flows depend on the distribution of x across employed matches. The distribution of individual productivity among employed workers is history dependent: $G_t(x) = Pr(X < x \big| a^t)$, where a^t represents the history of aggregate productivity shocks $\{a_0, a_1, ..., a_t\}$ realized up to time t. The conditional distribution is determined by the measure of employed workers over individual productivity, $e_t(x)$, which follows a law of motion determined by the flows

15 As $S(a_t,x)$ is monotonically increasing in x, the individual productivity threshold $x^r(a_t)$ is unique, and $S(a_t,x) > 0 \;\forall\; x > x^r(a_t)$. Appendix 3.B provides a detailed discussion.

between unemployment and employment and within employment. For those workers whose individual productivity is in the OJS interval $(x_t^r, x_t^s]$:

$$\begin{aligned} e_{t+1}(x) = {} & p(\theta_t)[1 - e_t(x_H)][F(x) - F(x_{t+1}^r)] + p(\theta_t)[F(x) - F(x_{t+1}^r)]e_t(x_t^s) \\ & + (1-s)\Big\{\lambda[F(x) - F(x_{t+1}^r)]\Big[e_t(x_H) - p(\theta_t)\overline{F(x_{t+1}^r)}e_t(x_t^s)\Big] \\ & + (1-\lambda)\Big[e_t(x) - e_t(x_{t+1}^r)\Big]\Big[1 - p(\theta_t)\overline{F(x_{t+1}^r)}\Big]\Big\}. \end{aligned} \tag{14}$$

For the non-searching workers with $x > x_t^s$:

$$\begin{aligned} e_{t+1}(x) = {} & p(\theta_t)[1 - e_t(x_H)][F(x) - F(x_{t+1}^r)] + p(\theta_t)[F(x) - F(x_{t+1}^r)]e_t(x_t^s) \\ & + (1-s)\Big\{\lambda[F(x) - F(x_{t+1}^r)]\Big[e_t(x_H) - p(\theta_t)\overline{F(x_{t+1}^r)}e_t(x_t^s)\Big] \\ & + (1-\lambda)\Big[e_t(x) - e_t(x_t^s) + (1 - p(\theta_t)\overline{F(x_{t+1}^r)})[e_t(x_t^s) - e_t(x_{t+1}^r)]\Big]\Big\}. \end{aligned} \tag{15}$$

Gross flows from employment to unemployment represent the total mass of workers separated from a job between two periods:

$$\begin{aligned} EU_{t+1} = {} & s\Big[e_t(x_H) - p(\theta_t)\overline{F(x_{t+1}^r)}e_t(x_t^s)\Big] \\ & + (1-s)\Big\{\lambda F(x_{t+1}^r)\Big[e_t(x_H) - p(\theta_t)\overline{F(x_{t+1}^r)}e_t(x_t^s)\Big] \\ & + (1-\lambda)e_t(x_{t+1}^r)\Big[1 - p(\theta_t)\overline{F(x_{t+1}^r)}\Big]\Big\} \end{aligned} \tag{16}$$

The job separation rate is then defined as the probability that an employed worker in period t is not employed in period $t+1$: $SR_t = EU_{t+1}/[e_t(x_H)]$. Similarly, the gross unemployment to employment (UE) flow is the total mass of workers who start a new job from unemployment:

$$UE_{t+1} = u_t p(\theta_t)\overline{F(x_{t+1}^r)},$$

and the job finding rate (JFR) is defined as the probability that an unemployed worker in period t is not unemployed in period $t+1$: $JFR_t = UE_{t+1}/u_t$. The job-to-job rate (JJR) is measured as the ratio of gross employment to new employment (EE) flows over total employment:

$$JJR_{t+1} = \frac{EE_{t+1}}{e_t(x_H)} = \frac{e_t(x_t^s)p(\theta_t)\overline{F(x_{t+1}^r)}}{e_t(x_H)}$$

3.4 Mechanisms for state-dependent fluctuations

The model embeds state-dependent dynamics in labor market variables across two dimensions. First, the response of labor market variables to the productivity shock is stronger in a state with low aggregate productivity. Second, the job separation rate and the unemployment rate are more responsive to a contractionary productivity shock than to an expansionary shock whereas the job finding rate responds symmetrically across shocks.

Figure 2: Distribution for $F'(x)$ and $G'(x)$

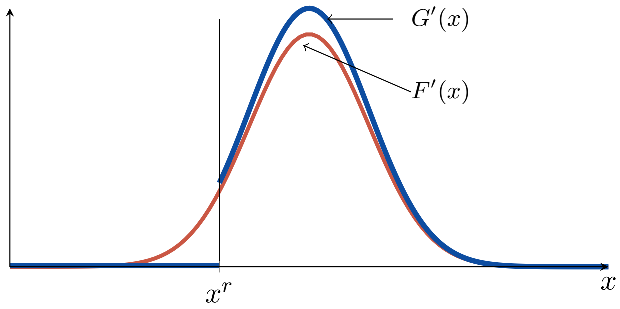

Note. The figure shows the p.d.f. for $F(x)$ (labelled $F'(x)$, red line) and $G(x)$ (labelled $G'(x)$, blue line).

The distinct responses of labor market variables over states of aggregate productivity is generated by the effect of changes in the individual productivity threshold $x^r(a)$ on the distributions of individual productivity for newly-established matches $F(x)$ and for continuing jobs, $G(x)$.[16] Figure 2 shows an illustrative probability density function for for the x of new matches (i.e. $F'(x)$) and incumbent ones (i.e. $G'(x)$) in red and blue, respectively. The difference between the two distributions is that $G(x)$ has zero mass below the individual productivity threshold $x^r(a)$ since jobs with productivity lower than the threshold are terminated, whereas $F(x)$ is continuous and twice differentiable since the productivity of new jobs is positively defined across the whole domain of individual productivity.[17] For both distributions, workers whose productivity is below the searching threshold x^s and above $x^r(a)$ search on the job.

Movements in the individual productivity threshold generate distinct responses in the job separation rate in relation to the state of aggregate productivity. Panel (a) in Figure 3 shows the initial productivity threshold x_0^r on the distribution for continuing jobs $G'(x)$ that is associated with a high level of aggregate productivity.[18] In response

16 To simplify notation, we drop the time index from the distributions. Given the timing assumption in the model $G(x)$ refers to $G_{t-1}(x)$ whereas $F(x)$ does not vary with time.

17 We make the standard assumption that $F(x)$ is continuous, twice differentiable and unimodal since it proxies the wage distribution in the data, as examined in Moscarini (2005).

18 Given that the surplus is increasing in both x and a, equation (13) implies that a high level of aggregate productivity is associated with a low individual productivity threshold. Assuming that the threshold always lies below the distribution mode, a lower threshold is located in a region of the

Figure 3: States of aggregate productivity and the job separation rate

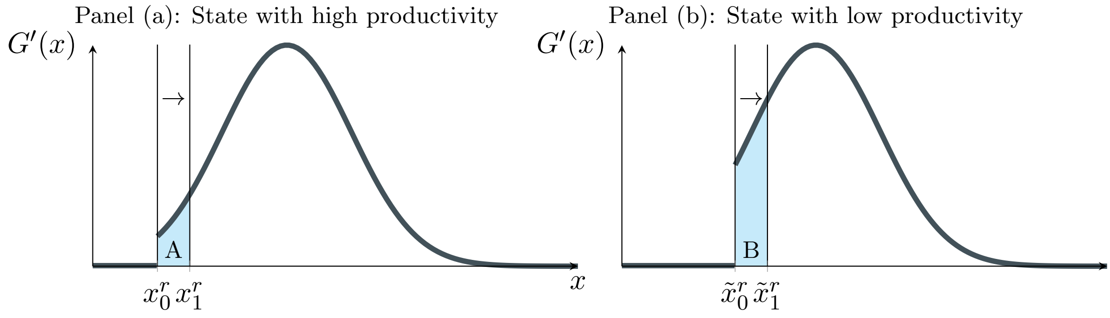

Note. An increase of the threshold of individual productivity from $\tilde{x}_0^r$ to $\tilde{x}_1^r$ generates a larger response in the job separation rate in states with low aggregate productivity than an equivalent increase of the threshold of individual productivity from x_0^r to x_1^r in states with high aggregate productivity. The shaded area shows the mass of jobs sensitive to job separation in response to the change in the threshold.

to a fall in aggregate productivity, the individual productivity threshold increases from x_0^r to x_1^r, leading to a rise in the job separation rate equal to shaded area A. Panel (b) shows the effect of an equivalent fall in aggregate productivity from an initially low level of aggregate productivity. In this instance, the individual productivity threshold is high and located in the domain of the distribution with high density. The same fall in aggregate productivity increases the individual productivity threshold from $\tilde{x}_0^r$ to $\tilde{x}_1^r$, leading to a rise in the job separation rate equal to shaded area B, which is larger than area A. Thus, the effect of a shock on the mass of jobs exposed to movements in the individual productivity threshold differs across levels of aggregate productivity, and the response of the job separation rate to the aggregate productivity shock is larger when aggregate productivity is low. By the same principle, the job finding rate may exhibit stronger responses in states with low aggregate productivity. However, these effects turn out to quantitatively weak, as we discuss in the next section.[19]

The second important feature of the model is the larger response of the job separation rate and the unemployment rate to a contractionary productivity shock compared to an expansionary one. To gain intuition, Figure 4 shows the effect of opposite movements in the threshold on the distribution for continuing jobs $G'(x)$. An increase in the threshold from x_0^r to x_1^r in response to a contraction in aggregate productivity generates the separation of jobs whose productivity is below x_1^r, represented by the shaded area. On the contrary, a decrease in the threshold from x_0^r to x_2^r in response to

distribution associated with lower density.

19 Appendix 3.C reports a graphical representation of the mechanism underpinning the distinct response of the job finding rate at different levels of productivity.

Figure 4: Contractionary and expansionary shocks and the job separation rate

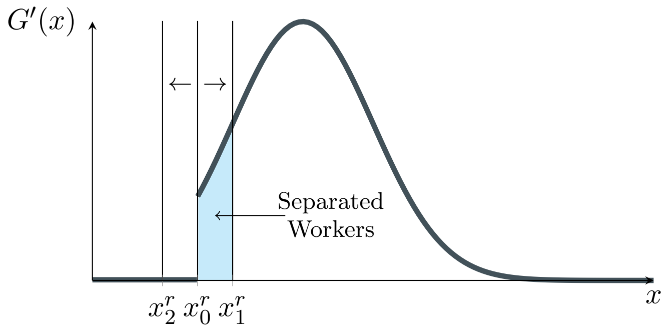

Note. The threshold of individual productivity increases (decreases) from x_0^r to x_2^r (x_1^r) in response to contractionary (expansionary) productivity shocks. The shaded area shows the mass of jobs sensitive to job separation in response to a contractionary technology shock.

an expansion in aggregate productivity involves no job separation since the threshold moves in a region of the distribution with zero mass (i.e. $G(x_2) = 0$). Hence, while a contractionary productivity shock that increases the threshold sharply raises the job separation rate and the unemployment rate, an expansionary productivity shock that decreases the threshold does not generate an equivalent fall in these variables. The fall in separation only comes through the effect of the change in the reservation threshold on the proportion λ of workers who receive a new shock and on the newly matched job-seekers who now face lower efficiency standards to be hired. This distinct reaction to expansionary and contractionary shocks is absent in the job finding rate since the distribution for new hires is positively defined across the domain of $F(x)$.[20]

3.5 Model simulation and quantitative results

This section presents the calibration of the model and the quantitative results. It compares simulated moments in the model against those in the data. It investigates the extent to which the model replicates the observed changes in the magnitude of fluctuations at distinct states of aggregate productivity and the skewness of labor market variables. Finally, it presents generalized impulse response functions to isolate the dynamic responses of labor market variables in different states of aggregate productivity (i.e. high versus low) and to contractionary and expansionary shocks.

20 In theory, distinct responses in the job finding rate with respect to the sign of the shock are possible. The fall in job creation from an increase in the reservation threshold need not be as large as a fall due to an equal decrease in the reservation threshold. In our model, we find that this effect is extremely small and the size of responses in the job finding rate is almost identical across contractionary and expansionary shocks.

3.5.1 Calibration

To allow the theoretical framework to embed state-dependent dynamics, we solve the model non-linearly iterating over the policy function on a discretized state space, following the approach in Tauchen (1986). Appendix 3.D reports the solution procedure.

We calibrate the model at monthly frequency. The discount factor β is set equal to $0.953^{(1/12)}$, as in Shimer (2005). The cost of posting a vacancy κ is set equal to 0.17 to match the derived calculations on costs of a job opening based on survey results cited in Barron and Bishop (1985) and Barron et al. (1997).[21] The flow value of unemployment b is set to equal 0.71, as in Hall and Milgrom (2008), which is between the value of 0.4 in Shimer (2005) and the value of 0.95 in Hagedorn and Manovskii (2008). The elasticity of the matching function with respect to vacancies η is set equal to 0.5 in the range of empirical estimates in Petrongolo and Pissarides (2001). To satisfy the Hosios (1990) condition, which ensures that the equilibrium of the decentralized economy is Pareto efficient, we assume that the elasticity of labor market tightness with respect to vacancies is equal to the firm's bargaining power $(1-\phi)$ i.e., $\eta = (1-\phi) = 0.5$. The parameter of match efficiency γ is set equal to 0.47 to match the empirical average job finding rate of 0.45, as in Hagedorn and Manovskii (2008). The exogenous separation probability s is set equal to 0.022 to match the average job separation rate of 0.03. The probability of receiving a new individual productivity shock λ and the variance of individual productivity σ_x are set to equal 0.05 and 0.13, respectively, to match the quarterly autocorrelation and standard deviation of the HP-filtered log separation rate, equal to 0.54 and 0.055, respectively. The distribution of individual productivity shocks is a truncated log-normal density function with the lower bound equal to zero (i.e. $x_L = 0$) and the upper bound set to have less than 1 percent of the mass of the distribution above it (i.e. $x_H = 1.55$). The mean of the log distribution of individual productivity μ_x is set to -0.087 to normalize the long-run average productivity in the economy to 1. The cost of searching on the job k^s is set to equal 0.128 to match the mean monthly job-to-job transition rate of 3.2, calculated from the CPS data, as in Fujita and Ramey (2012). The autoregressive parameter ρ and the standard deviation σ of the aggregate productivity process are

21 As in Fujita and Ramey (2012), the value is derived from a calculation of the costs based on survey results cited in the above papers.

Table 5: Parameter values

Parameter	Description	Value
β	Discount factor	$0.953^{(1/12)}$
κ	Vacancy cost	0.17
κ_s	OJS cost	0.128
b	Flow value of unemployment	0.71
η	Elasticity of matching with respect to vacancies	0.5
γ	Matching function efficiency parameter	0.47
ϕ	Worker's bargaining power	0.5
s	Exogenous job separation rate	0.022
λ	Arrival rate of individual productivity shocks	0.05
x_L	Lower bound of individual productivity shocks	0
x_H	Upper bound of individual productivity shocks	1.55
μ_x	Mean of log individual productivity shocks	-0.087
σ_x	Standard deviation of individual productivity shocks	0.13
ρ	Persistence parameter of aggregate productivity	0.973
σ	Standard deviation of aggregate productivity shocks	0.0068

set equal to 0.973 and 0.0068, respectively, to match the autocorrelation and standard deviation of HP-filtered log labor productivity at quarterly frequency, as in Hagedorn and Manovskii (2008). Table 5 summarizes the calibration of parameters. Table 3.D.1 in Appendix 3.D shows that the simulated targeted moments are close to the empirical counterparts in the data.

3.5.2 Business cycle statistics

Table 6 compares the standard deviation, the correlation coefficient with productivity, and the autocorrelation coefficient for selected variables in the data (top panel) against the corresponding statistics in the simulated model (bottom panel). The moments are based on a set of 1,000 simulations of the same length as the empirical data. The model accurately reproduces the standard deviation of output 0.021 in the data. The simulated standard deviation of unemployment and the job finding rate of 0.101 and 0.063, respectively, are very close to those in the data (0.137 and 0.089, respectively). The simulated standard deviation of vacancies, equalling 0.040, is approximately three times smaller than the value of 0.138 in the data. Likewise, the simulated standard deviation of vacancy-to-unemployment ratio (v/u) equal to 0.137 is approximately half the value in the data.[22]

22 The relatively larger fluctuations of the job finding rate compared to vacancies indicate that in the model the productivity threshold x_t^r rather than posting vacancies is the main channel through which firms adjust their recruiting decisions. Firms adjust x_t^r to determine endogenous separation, which

The model replicates accurately the sign of the correlation coefficient of the variables with productivity. The unemployment rate and the job separation rate are negatively correlated with productivity whereas the rest of the variables are positively correlated. The correlations in the model are larger than those in the data since productivity shocks are the only exogenous source of aggregate fluctuations. It is worth noting that in the model productivity is positively correlated with vacancies and negatively correlated with the unemployment rate, which replicates important stylized facts in the data (i.e. the Beveridge Curve).[23] Finally, the autocorrelation coefficients for most variables in the model are similar to those in the data.

Table 6: Labor market statistics in the data and the model

Data	p	U	JFR	SR	E	V	V/U	Y
σ_X	0.013	0.137	0.089	0.055	0.009	0.138	0.262	0.021
Corr(p_t,X_t)	1.000	-0.229	0.212	-0.556	0.232	0.394	0.316	0.661
Corr(X_t,X_{t-1})	0.763	0.891	0.840	0.535	0.899	0.907	0.905	0.843
Model	**p**	**U**	**JFR**	**SR**	**E**	**V**	**V/U**	**Y**
σ_X	0.013	0.101	0.063	0.055	0.008	0.040	0.137	0.021
Corr(p_t,X_t)	1.000	-0.970	0.978	-0.951	0.929	0.927	0.991	0.991
Corr(X_t,X_{t-1})	0.769	0.822	0.781	0.647	0.821	0.586	0.778	0.794

Note. The table reports cyclical statistics for average labor productivity (p), the unemployment rate (U), the job finding rate (JFR), the job separation rate (SR), the employment rate (E), vacancies (V), the V/U ratio (V/U), and output (Y). The simulated moments are computed as means of 1,000 simulations of 1,380 monthly periods. After discarding the first 600 observations in each simulation, the remaining series are aggregated at quarterly frequency with the same length of the observed series for the period 1950:I-2014:IV.

3.5.3 State-dependent fluctuation and skewness

This section assesses whether the model is able to replicate the distinct responses of labor market variables related with the different states of aggregate productivity and evaluates the degree of skewness in the model.

also affects the expected value of new matches. The adjustment in reservation productivity mitigates the fluctuations in the expected surplus of new matches and hence dampens the response of vacancies.

23 As discussed by Fujita and Ramey (2012), OJS is critical for the negative correlation between vacancies and the unemployment rate when the model includes endogenous job destruction. In response to a negative productivity shock, unemployment rises sharply but the total mass of job seekers does not change as much because the endogenously separated workers were all in the OJS pool. Was it not for the fall in OJS, firms could have adjusted the v/u ratio downwards while simultaneously increasing vacancies due to the high separation rate.

Table 7 shows the standard deviations of simulated labor market variables in levels (Columns 1-3) and growth rates (Columns 4-6) associated with productivity below and above its median value. Columns 1 and 2 shows that the standard deviations of the simulated variables are larger in periods of productivity below the median value. Column 3 shows the relative standard deviation between the variance of the variables when productivity is below and above the median value. The entries show that all the variables, except the job finding rate, have larger fluctuations in periods with low aggregate productivity. Once again, the result is substantially smaller for the job finding rate. The relative standard deviation of most variables in levels is within the range of 1.16-1.58, while it is only 1.08 for the job finding rate. A similar result holds for the statistics derived using the growth rates of the variables (Columns 4-6). These patterns are similar to those in the data reported in Table 1. Overall, the model fares better in levels rather than growth rates because its original frequency is monthly but first differences are computed over quarterly averages, which mitigates some of the inherent nonlinearities in the monthly changes.

Table 7: Standard deviation of simulated variables for different states of productivity

	Levels			Growth Rates		
	p< 50^{th} percentile (1)	p> 50^{th} percentile (2)	$\frac{\sigma_{p<50}}{\sigma_{p>50}}$ (3)	p< 50^{th} percentile (4)	p> 50^{th} percentile (5)	$\frac{\sigma_{p<50}}{\sigma_{p>50}}$ (6)
Unemployment	0.0921	0.0792	1.17	0.0657	0.0578	1.14
Job Finding Rate	0.0237	0.0222	1.08	0.0188	0.0174	1.08
Separation Rate	0.0017	0.0012	1.41	0.0017	0.0012	1.47
Employment Rate	0.0079	0.0051	1.58	0.0056	0.0035	1.62
Output	0.0188	0.0164	1.16	0.0144	0.0126	1.14
Productivity	0.0112	0.0116	0.978	0.0091	0.0093	0.980

Note. Entries are averages of 1,000 simulations over 1,380 monthly periods. After discarding the first 600 observations in each simulation, the remaining series are aggregated at quarterly frequency and have the same length as the period 1950:I-2014:IV.

Table 8 reports the excess skewness of selected variables in the model. The entries show that the model replicates the sign of skewness in the data from Table 4. In the model as in the data, the unemployment rate and the job separation rate are positively skewed both in levels and growth rate. The employment and output are negatively skewed. The model generates a low degree of skewness in the job finding rate, which corroborates the empirical findings on the insignificant degree of skewness in the job finding rate in the data. Overall, the coefficients of skewness in the model are lower than those in the data. This difference results from the quarterly aggregation, which averages out differences in the original simulated monthly series.[24]

Table 8: Skewness in simulated data

	Levels	Growth rate
Unemployment	0.180	0.159
Job Finding Rate	-0.082	-0.023
Separation Rate	0.519	0.400
Employment	-0.579	-0.195
Output	-0.175	-0.048
Productivity	0.035	0.017

Note. Entries are averages of 1,000 simulations over 656 monthly periods. After discarding the first 400 observations in each simulation, the remaining series are aggregated at quarterly frequency and have the same length as the period 1950:I-2014:IV.

To summarize, the model replicates reasonably well the different volatility of labor market variables at distinct states of aggregate productivity, and it reproduces accurately the signs of skewness in the data.

3.5.4 Dynamic response of labor market variables

In this section we investigate the extent to which the dynamic responses of labor market variables are different across states with high and low aggregate productivity and in response to an expansionary or a contractionary shock. To this end, we compute generalized Impulse Response Functions (IRFs) to both positive and negative productivity shocks equal to one quarterly standard deviation at two points of the business cycle: one in which labor productivity is very low (i.e. below the 10^{th} percentile of its distribution) and one in which it is very high (i.e. above the 90^{th} percentile).

Figures 5-7 plot generalized impulse response functions (IRFs) for the job separation rate, the job finding rate, and the unemployment rate to a productivity shock (solid line) together with the 5^{th}-95^{th} percentiles (shaded area).[25] Each figure reports the response of the variables in states with high and low aggregate productivity

24 Table 3.I.9 in the Appendix shows skewness coefficients for the series at monthly frequencies, which are of similar magnitude to those in the data.

25 Appendix 3.E describes the computational method to derive generalized IRFs.

(left and right panels, respectively) and to contractionary and expansionary shocks (bottom and top panels, respectively).[26]

Figure 5: Responses of the job separation rate across productivity states and shocks

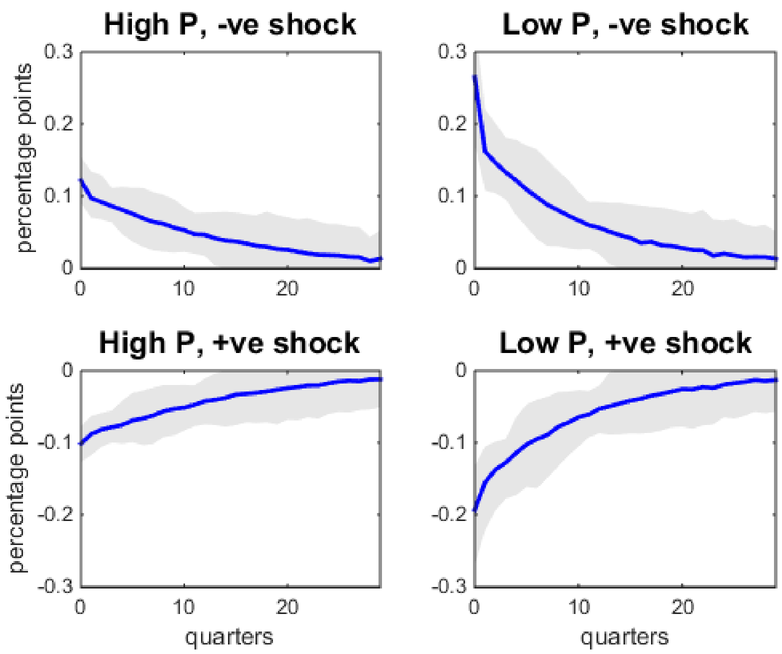

Note. The solid line represents the mean IRF value in each period. The shaded area represents the 5^{th} and 95^{th} percentiles of the IRF values. Responses of the variables in periods with high and low aggregate productivity are in top and bottom panels, respectively. Responses to positive and negative shocks are in left and right panels, respectively. Units on the y-axis are percentage points.

Figure 5 shows that the response of the job separation rate is more than twice as large in the state with low productivity compared to the state with high productivity and that responses to a contractionary shock are stronger than those to an expansionary shock, especially in the state with low aggregate productivity. As discussed in Section 3.4, the mechanism that generates these distinct dynamics is straightforward. The different responses with respect to the state with aggregate productivity (left versus right panels) originate from the effect of shifts in the reservation threshold for individual productivity on the job separation rate. In the state with high aggregate productivity, the threshold is low and located in a region of individual productivity distribution with low density. Aggregate shocks that move the threshold displace a limited number of workers and therefore have a limited effect on the job separation rate. By contrast, in the state with low aggregate productivity, the threshold of in-

26 In the simulation, a state with high aggregate productivity is considered to be a quarter in which aggregate productivity is above the 90^{th} percentile. A state with low productivity is a quarter in which aggregate productivity is below the 10^{th} percentile.

dividual productivity is high and located in a region of the distribution of individual productivity with high density. Thus, an identical aggregate productivity shock that moves the threshold upward displaces a larger fraction of workers, thereby generating a larger jump in the job separation rate.

The distinct responses in the job separation rate with respect to a contractionary and an expansionary shock (top versus bottom panels) originate from the truncation of the distribution of individual productivity, which has zero mass for values below the reservation threshold since those jobs are inefficient. As discussed in Section 3.4, a contractionary productivity shock increases the threshold of individual productivity, leading to the termination of jobs with an individual productivity below the threshold. By contrast, an expansionary productivity shock lowers the productivity threshold by relaxing the firms' threshold of individual productivity for efficient matches. The effect on the job separation rate is limited because the individual productivity levels of existing jobs already exceed the individual productivity threshold and workers avoid job separation via OJS. The separation rate falls to the extent that a fraction $\lambda \mathrm{F}(x^r)$ of existing jobs receive a new individual productivity below the separation threshold.

Figure 6: Responses of the job finding rate across productivity states and shocks

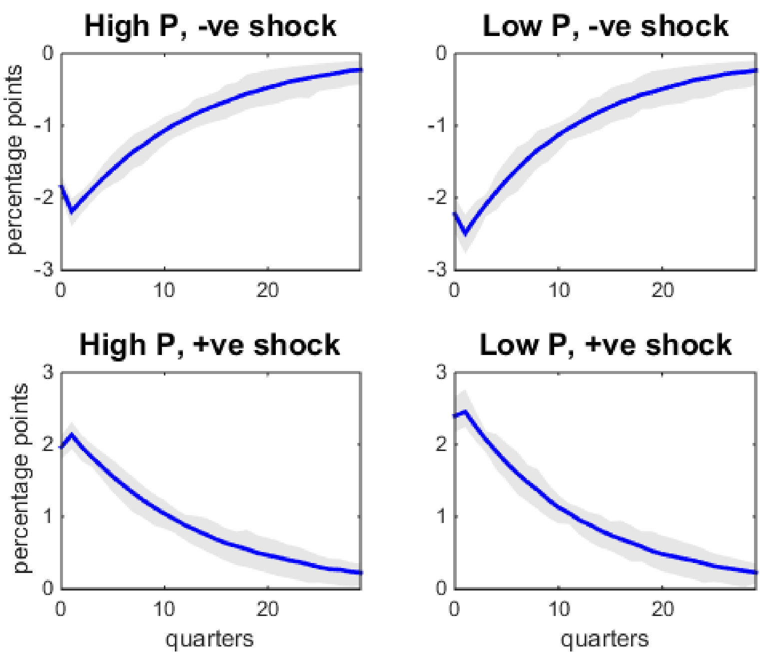

Note. The solid line represents the mean IRF value in each period. The shaded area represents the 5^{th} and 95^{th} percentiles of the IRF values. Responses of the variables in periods with high and low aggregate productivity are in top and bottom panels, respectively. Responses to positive and negative shocks are in left and right panels, respectively. Units on the y-axis are percentage points.

Figure 7: Responses of the unemployment rate across productivity states and shocks

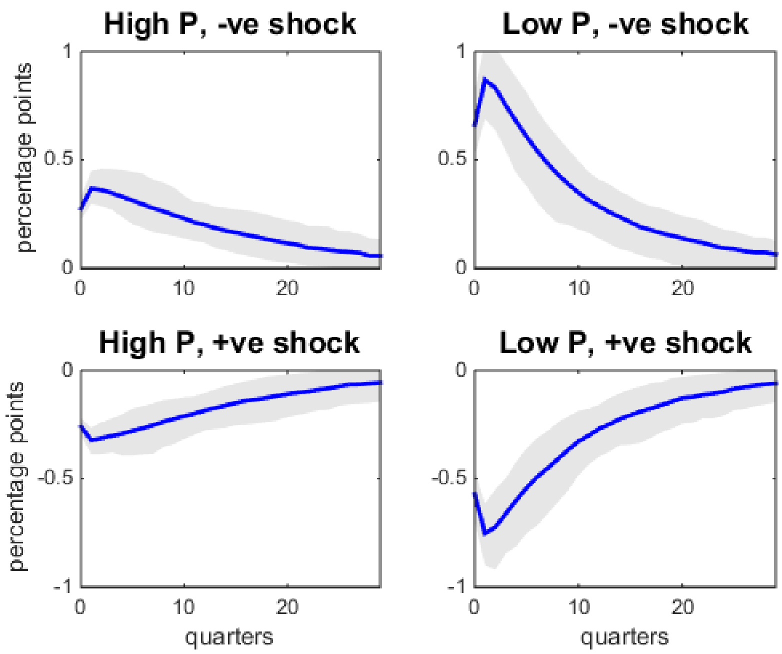

Note. The solid line represents the mean IRF value in each period. The shaded area represents the 5^{th} and 95^{th} percentiles of the IRF values. Responses of the variables in periods with high and low aggregate productivity are in top and bottom panels, respectively. Responses to positive and negative shocks are in left and right panels, respectively. Units on the y-axis are percentage points.

Figure 6 shows the responses of the job finding rate. The entries reveal that the job finding rate reacts differently across states with low and high aggregate productivity (left versus right panels) whereas responses are similar across contractionary and expansionary shocks (top versus bottom panels). Consistent with the finding in Section 3.2, the quantitative differences across states of aggregate productivity are substantially smaller than those for the job separation rate. The more moderate difference in state-dependent fluctuations results from OJS. Movements in OJS counteract fluctuations in the unemployment rate and dampen changes in labor market tightness, limiting the impact of shocks on a firms' hiring and on the resulting job finding rate.[27]

In Figure 7, the response of the unemployment rate is almost twice as large in the state with low productivity (left versus right panels), with moderate differences in the magnitude of responses across contractionary and expansionary shocks (top versus

27 Petrosky-Nadeau and Zhang (2013) show that the DMP model with a constant job separation generates a larger elasticity in the job finding rate during periods of deep recession, driven by the higher sensitivity of the joint surplus of forming a match in periods with low productivity. Our analysis shows that the interaction between endogenous job separation and OJS reduces the elasticity of the job finding rate, because the mass of all job seekers fluctuates less than unemployment and contractionary and expansionary shocks have a similar effect on the job finding rate.

bottom panels). In the model, the unemployment rate results from changes in the job separation rate and the job finding rate. The distinct responses in unemployment are thus primarily inherited from those in the job separation rate.

3.6 State dependence and labor market reforms

As the previous sections identified job destruction as the main source of asymmetries over the business cycle, in this section we consider the impact of a policy that directly interferes with firms' firing decisions. We enrich the model with labor market protection in the form of a layoff tax levied on the dismissal of established jobs. The analysis investigates the effect of the layoff tax on the long-run equilibrium of the model and assesses whether the permanent tax removal in states with low or high aggregate productivity generates critical differences in the transitional dynamics and welfare.

3.6.1 Introducing layoff taxes

The model described in Section 3.3 is enriched with a "wasteful" layoff tax τ that the firm must pay to cover administrative costs and layoff procedures whenever a worker is (endogenously or exogenously) separated.[28] Firms whose workers move to another job with OJS do not incur into the layoff tax. Layoff taxes are not levied on the separation of newly-established matches, and therefore the joint value of an employment relationship for new matches (indexed by N) and continuing matches (indexed by O) are distinct and defined as

$$M^N(a_t, x) = \max[M^{c,s}(a_t, x), M^{c,n}(a_t, x), U(a_t) + V(a_t)], \tag{17}$$

$$M^O(a_t, x) = \max[M^{c,s}(a_t, x), M^{c,n}(a_t, x), U(a_t) + V(a_t) - \tau], \tag{18}$$

where $M^N(a_t, x)$ and $M^O(a_t, x)$ denote the joint value for new and existing workers, respectively, which account for on-the-job searchers. $M^{c,s}(a_t, x)$ and $M^{c,n}(a_t, x)$ are the joint values for continuing the job relationship with or without OJS, respectively,

28 See Fella (2007), Ljungqvist (2002), Postel-Vinay and Turon (2014), and Cozzi and Fella (2016) for a discussion on the role of employment protection measures in matching models.

and are defined as:

$$M^{c,n}(a_t,x) = a_t x + \beta \mathbb{E}_t \Big\{ U(a_{t+1}) + V(a_{t+1}) - \tau \\ + (1-s)\Big[(1-\lambda)S^O(a_t,x) + \lambda \int_{x_L}^{x_H} S^O(a_{t+1},x')dF(x')\Big]\Big\}, \quad (19)$$

$$M^{c,s}(a_t,x) = a_t x - k^s + \beta \mathbb{E}_t \Big\{ U(a_{t+1}) + V(a_{t+1}) - \tau \\ + \Big(1 - p(\theta)\overline{F(x^r_{t+1})}\Big)(1-s)\Big[(1-\lambda)S(a_{t+1},x) \\ + \lambda \int_{x_L}^{x_H} S^O(a_{t+1},x')dF(x')\Big] + p(\theta)\phi \int_{x(a_{t+1})}^{x_H} S^O(a_{t+1},x')dF(x')\Big\}. \quad (20)$$

The total surplus equals the value of establishing a match net of the outside options to the worker and the firm. Thus, the value functions for the joint surpluses for new and continuing jobs are:

$$S^N(a_t,x) = \max[S^{N,c}(a_t,x), 0], \quad (21)$$

$$S^O(a_t,x) = \max[S^{O,c}(a_t,x), 0], \quad (22)$$

where $S^{N,c}(a,x)$ and $S^{O,c}(a,x)$ represent the total surpluses in case the worker and the firm establish a new match or continue an existing job relationship, respectively, accounting for the optimal choice of OJS. The total surpluses for new and old matches, without and with OJS are:

$$S^{N,n}(a_t,x) = a_t x - b + \beta \mathbb{E}_t \Big\{ (1-s)\Big[(1-\lambda)S^O(a_{t+1},x) + \lambda \int_{x_L}^{x_H} S^O(a_{t+1},x')dF(x')\Big] \\ - \tau - p(\theta_t)\phi \int_{x_L}^{x_H} S^N(a_{t+1},x')dF(x')\Big\}, \quad (23)$$

$$S^{N,s}(a_t,x) = a_t x - k^s - b + \beta \mathbb{E}_t \Big\{ \Big[1 - p(\theta)\overline{F(x^{N,r}_{t+1})}\Big](1-s)\Big[(1-\lambda)S^O(a_{t+1},x) \\ + \lambda \int_{x_L}^{x_H} S^O(a_{t+1},x')dF(x')\Big] - \tau\Big\}, \quad (24)$$

$$S^{O,n}(a_t,x) = S^{N,n}(a_t,x) + \tau, \quad (25)$$

$$S^{O,s}(a_t,x) = S^{N,c}(a_t,x) + \tau. \quad (26)$$

Equations (23) and (24) show that the surpluses of newly-established job relations (with or without OJS) are reduced by the expected layoff tax if the job is dismissed in the future. Equations (25) and (26) show that the surpluses for existing job relations (with or without OJS) entail an intertemporal tradeoff between the benefit of foregoing tax payment if the job is not severed in the present period t and the cost of having to pay the layoff tax if the worker is dismissed in the future. Important for our analysis, the difference in the surpluses for newly-hired and existing workers generates distinct thresholds of individual productivity. The reservation productivity at which new matches become inefficient is higher than the reservation threshold for existing workers because firms are not discouraged from discontinuing the newly-formed match at time t as they are for incumbent workers. Consequently, the firm retains existing workers with individual productivity in the range $x \in \big(x_t^{r,O}(a), x_t^{r,N}(a)\big]$ but does not hire new matches with individual productivity in the same interval. Within this range of individual productivity, it is inefficient to pay layoff taxes to dismiss existing workers but it is efficient to refuse new matches to which layoff taxes do not apply.

A worker's decision to search on the job is influenced by the prospects of obtaining a successful match. The productivity threshold that applies to this expectation, and hence to the decision to search on the job, is $x_t^{r,N}(a)$ since employed and unemployed job seekers are identical to the hiring firm. Additionally, the cutoff level for OJS is the same across new and incumbent workers.[29] Appendix 3.F outlines the laws of motion of employment in the presence of layoff taxes.

We set the value for the layoff tax equivalent to 15 percent of average monthly productivity ($\tau = 0.15$), which makes the tax approximately equivalent to 5 percent of average quarterly wages, as in Llosa et al. (2014).

Figure 8 shows how the tax changes the relevant thresholds for workers' and firms' decisions. The reservation productivity for firing incumbent workers is only slightly lower than in the no-tax case. That is because, for an incumbent match, the tax creates a small wedge between the surplus in the baseline and the tax economies. For new matches, however, the tax shifts the reservation threshold $x_\tau^{r,N}$ upward.[30] At the same time, the threshold for OJS shifts downward. For a given level of market tightness, firms stricter productivity requirements induce the formation of fewer matches. As a result, the highest level of match-specific productivity at which workers decide to search for another job also falls. For low levels of aggregate productivity, the OJS threshold is sufficiently low to be located below the threshold for new matches, im-

29 Given that the tax (and future taxes) enter linearly into both the surplus from searching and not searching, as seen in equations (23)-(26), the threshold level for OJS x^s that satisfies $S^{N,n}(a_t, x) = S^{N,s}(a_t, x)$ is the same as that which satisfies $S^{O,n}(a_t, x) = S^{O,s}(a_t, x)$.

30 An intuition for why the threshold for incumbent workers only decreases by a small amount, while that of new workers rises by almost the entire value of the tax can be given by a steady-state version of the model. There, the wedge between the surplus functions for new workers in the no-tax and the tax economies is close to $-\bar{s}\beta\tau/\big(1-(1-\bar{s})\beta\big) = -0.128$, using a steady state separation rate $\bar{s} = 0.03$. Meanwhile for incumbent workers the wedge is $\tau - \bar{s}\beta\tau/\big(1-(1-\bar{s})\beta\big) = 0.0215$.

Figure 8: Thresholds for separation, job creation, and OJS, over match-specific productivity and aggregate productivity in the baseline model and the model with layoff taxes.

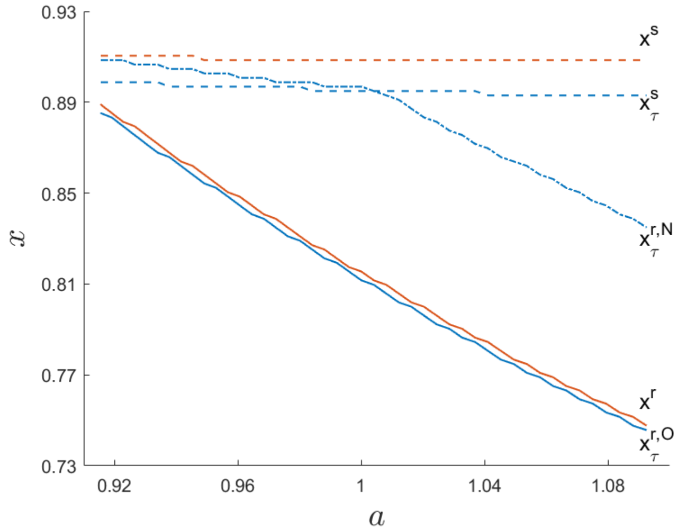

Note. The x-axis reports the level of aggregate productivity a. The y-axis reports the level of individual productivity x. The red lines report the thresholds for the baseline case, where x^s is the OJS threshold and x^r is the reservation productivity level for both job separations and formation of new matches. The blue lines report the thresholds for the model with employment protection with $\tau = 0.15$, where x^s_τ is the OJS threshold, $x^{r,O}_\tau$ is the reservation productivity level for incumbent workers, and $x^{r,N}_\tau$ is the threshold for the formation of new matches.

plying that no new match chooses OJS. Overall, the levels and slope of the thresholds imply that the tax reduces the mass of OJS workers, and especially so during low productivity times.

3.6.2 The long-run effect of layoff taxes

Table 9 compares the long-run values of key labor market variables in the version of the model without (Column 1) and with (Column 2) layoff taxes. The tax increases the long-run unemployment rate by two percentage points from 6.8 percent to 8.7 percent, as a result of the large fall in the job finding rate and the broadly constant rate of job separation. The effect of the tax on the long-run value of the job separation rate is limited. The reason is straightforward. While introducing the tax slightly decreases the efficiency threshold for continuing jobs and therefore decreases endogenous separation, it also increases the efficiency threshold for new jobs and therefore discourages OJS. Overall, OJS decreases from 6.6 percent to 3.4 percent, leaving a greater fraction of workers subject to job separation. These two opposing forces offset each other, and the job separation rate effectively remains unchanged.

The tax lowers the job finding rate from 44.5 percent to 32.4 percent. The fall in the job finding rate originates from the increase in the threshold of efficient matches for new hires $x_t^{r,N}(a)$ and the fall in vacancies. The higher productivity threshold for new matches discourages on-the-job search and consequently reduces the total number of job seekers, increasing the search costs per vacancy filled accrued to firms, which react by decreasing vacancy posting. Thus, the reduction in OJS amplifies the contractionary effect of the layoff tax on the job finding rate and generates a large rise in the unemployment rate.

The model involves a positive relation between layoff taxes and the unemployment rate in the long run. There is no established consensus on the effect of layoff taxes on the unemployment rate in the literature.[31] Mortensen and Pissarides (1999) show that layoff taxes reduce incentives both to create and destroy jobs and the net effect of these forces on labor market tightness is ambiguous. Ljungqvist (2002) shows that the effect of layoff taxes on unemployment depends on theoretical setups rather then alternative calibrations. We contribute to this realm of research by showing that OJS plays a critical role for the effect of layoff taxes on the unemployment rate.[32]

3.6.3 Short-run effects and welfare of layoff tax removal

This section investigates whether the timing of an unexpected, permanent removal of the layoff tax enacted in alternative states of aggregate productivity is critical for transitional dynamics and welfare. We compare the effect of the reform enacted at a level of labor productivity below the 10^{th} percentile and above the 90^{th}.

Figure 9 shows transitional paths for the unemployment rate (solid line) that result from the removal of the layoff tax in the low productivity state (right panel) and the high productivity state (left panel), together with the 10^{th}-90^{th} interval of

31 See Nickell et al. (2005) and references therein for a review on the literature. A recent study by Bentolila et al. (2012) documents that changes in firing costs have a different impact on unemployment in different economies.

32 In this respect, our findings contrast those in Postel-Vinay and Turon (2014) who develop a search and matching model with on-the-job search and renegotiation over severance packages that generates a negative relation between layoff taxes and the unemployment rate. Our analysis builds on the standard DMP model enriched with OJS. The aforementioned study enriches the standard DMP model across several dimensions, assuming wage renegotiation based on mutual consent, endogenous severance packages, and a minimum wage. The opposite result on the effect of layoff taxes on the unemployment rate results from the interaction of the layoff tax with on-the-job searches and minimum wages.

Table 9: The long-run effects of a layoff tax

	Baseline $\tau = 0$ (1)	Layoff tax $\tau = 0.15$ (2)
Unemployment rate	0.068	0.087
Job finding rate	0.445	0.324
Separation Rate	0.030	0.030
Job-to-job rate	0.032	0.012
On-the-job search	0.066	0.034
Employment Rate	0.932	0.913
Vacancies	0.179	0.158
V/U	2.832	1.919
Productivity	1.000	1.004

Note. The table shows the long-run averages of labor market variables in the baseline model ($\tau = 0$) and in the alternative model with layoff taxes ($\tau = 0.15$).

the differences across states (shaded area).[33] The dashed line shows the long-run difference between the unemployment rate in the economy with the tax eliminated in the first period and the economy with the tax always in place.[34] In the long run, the tax removal leads to a similar fall in the average unemployment rate of approximately 2.2 percentage points across states with initially low and high aggregate productivity since both economies converge to the equilibrium over the long run. In the first period of the reform, the unemployment rate falls by approximately 1 percentage point in both initial states of aggregate productivity. However, in the subsequent periods, the transitional dynamics of the unemployment rate differ significantly across states of aggregate productivity. In the state with high productivity, the reform causes the unemployment rate to gradually decrease towards its long-run equilibrium whereas in the state with low productivity, the decline of unemployment is immediate, and the change in the unemployment rate remains below its long-run equilibrium for a protracted number of periods.

To interpret these sharp differences in the transitional paths of the unemployment rate across states of aggregate productivity, we consider the transitional dynamics

33 The transition path of unemployment from the outset of the reform at time t is computed as $\mathrm{TP}(u_t|e_0(x), \{a_t\}_0^t) = u_t^R - u_t^\tau$, where the superscripts R and τ indicate the reform and tax scenario, respectively, assuming an initial state $e_0(x)$ and a productivity sequence $\{a_t\}_0^t$. Appendix 3.G discusses the simulation procedure.

34 The reported long-run difference of -2.3 is slightly different from the value reported in Table 9 since the value in the table is computed using a set of short simulations equivalent to the 1950-2014 period, as explained in Appendix 3.D. Meanwhile, the dotted line in the figures is the long-run mean-difference between the two economies taken as the average difference after 25 quarters.

Figure 9: Transition path of the unemployment rate to a permanent elimination of the layoff tax

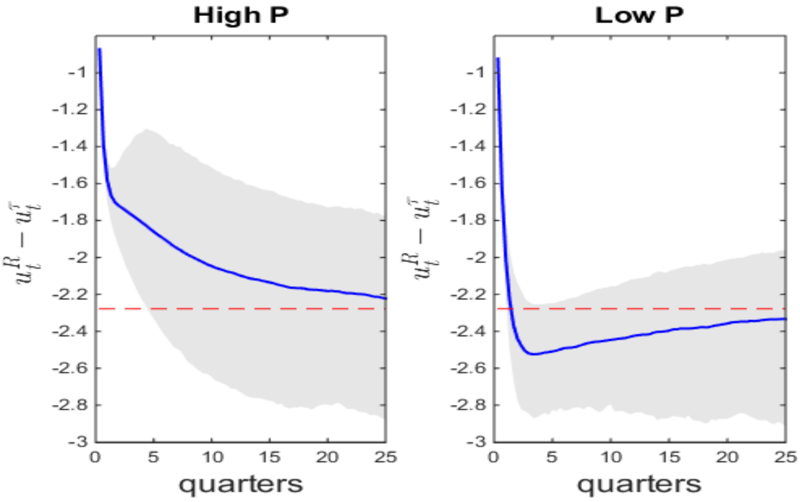

Note. The solid blue lines represent the average difference in the unemployment rate (in percentage points) of the economy where at time 0 the layoff tax was abolished from that of the same economy where the tax persists from 2,000 simulations of the model. The shaded grey area represents the 10^{th}-90^{th} percentile interval of the differences. The dashed red line represents the long-run average difference between the two economies.

of the job finding rate, the job separation rate, the job-to-job transition rate, and vacancies after the tax removal across states of aggregate productivity, which are reported by Figure 3.H.1 in the appendix. The tax removal has a limited effect on the job separation rate across states of aggregate productivity. The reason is straightforward. While removing the tax raises the individual productivity threshold and therefore increases endogenous job separation, it also increases OJS, which enables workers to avoid endogenous separation by moving to a new job. These two opposing forces offset each other and leave the overall response of the job separation rate broadly unchanged across states of aggregate productivity.[35] Meanwhile, the fall of the individual productivity threshold for new matches stimulates OJS and increases the pool of job seekers. This force reduces the firm's search cost of filling a vacancy, which leads to a sharp increase in the number of vacancies and the job finding rate on impact.[36] In the aftermath of tax reform, the job finding rate immediately rises in both states of aggregate productivity. In the high-productivity state, the job-finding

35 The short-run dynamics include an immediate rise in separation due to the increase in x^r, followed by a fall due to the rise in OJS. Quantitatively, both effects are small and similar across initial states.

36 The job-to-job rate shows a small increase on impact since in period t, the mass of OJS workers is predetermined in period $t-1$, but the probability of having a successful match is determined by the current fall in the threshold. The job-to-job rate sharply increases in period $t+1$ when a large number of newly employed workers search on the job. The rise in the mass of OJS workers rapidly stabilizes to its long-run value.

rate overshoots the long-run equilibrium while in the low-productivity state the rate gradually rises towards the long-run equilibrium. This critical difference originates from the larger surplus of forming a job relation in the high-productivity state, which leads firms to post a large amount of vacancies in the period of the tax removal to reap the benefits of forming a job while productivity is high. However, since the initial unemployment is higher in the low-productivity state as a result of the high job separation rate, even a small rise in the job finding rate yields a large fall in unemployment in the state with low productivity, as depicted in the right entry in Figure 9.

How do these sharp differences in the dynamic responses of labor market variables influence welfare across states of aggregate productivity? The tax removal unambiguously increases welfare in the long run since firms stop paying the wasteful tax. Resource allocation becomes efficient because firms terminate the low-productive jobs they had retained to avoid the payment of the layoff tax and recruit high-productivity workers whose hiring was prevented in anticipation of payment of the tax in future periods.

In the short run, however, the timing of the reform is critical for welfare since the transitional dynamics of labor market variables is notably different across distinct states of aggregate productivity.To investigate the relevance in the timing of structural reforms, we proxy welfare with the flow value of the economy that comprises output and the flow value of unemployment net of hiring costs, OJS costs, and layoff taxes.[37] It is straightforward to derive the welfare gain of the tax removal by subtracting the flow value of the economy with the layoff tax from the flow value of the economy without the tax in each period:

$$\Delta FV_t = \Big(a_t \int_{x_L}^{x_H} x \;\; de_t^R(x) + bu_t^R - kv_t^R - k^s \phi_t^R\Big) \\ - \Big(a_t \int_{x_L}^{x_H} x \;\; de_t^\tau(x) + bu_t^\tau - kv_t^\tau - k^s \phi_t^\tau - EU_t^\tau \tau\Big), \quad (27)$$

where the superscript R indicates variables in the economy with tax reform and the superscript τ indicates variables in the economy with the layoff tax still in place. Equation (27) tracks the welfare change of the tax removal during each period t.

Figure 10 plots equation (27) for the initial 25 quarters across 2,000 simulations and shows the net welfare gains from the removal of the layoff tax for the state with high

37 See Ljungqvist and Sargent (2012) for a similar approach to approximate welfare.

Figure 10: Transition path of the aggregate flow value to a permanent elimination of the layoff tax

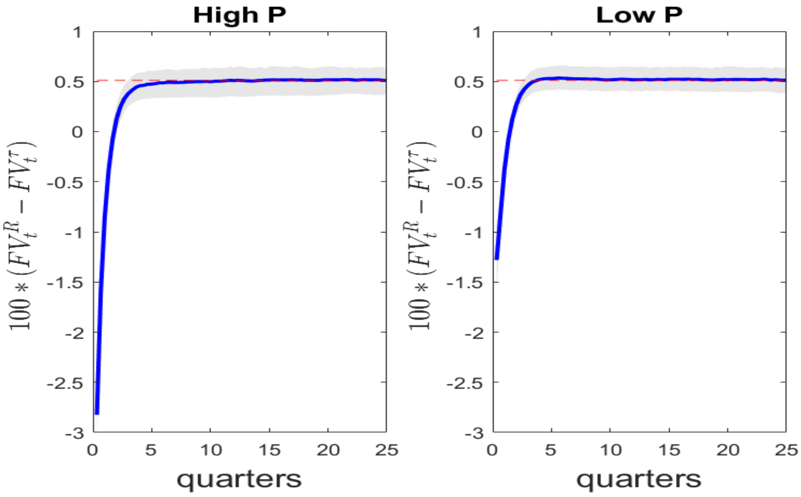

Note. The solid blue lines represent the average difference in the aggregate flow value of the economy where at time 0 the layoff tax was abolished from that of the same economy where the tax persists from 2,000 simulations of the model. The shaded grey area represents the 10^{th}-90^{th} percentile interval of the differences. The dashed red line represents the long-run average difference between the two economies.

and low aggregate productivity (left and right panels, respectively). A tax removal generates a contemporaneous welfare loss approximately twice as large in the state with high productivity compared to the state with low aggregate productivity. When the reform is enacted in the state with high aggregate productivity, the welfare loss is equal to a reduction of approximately 2.8 units in the flow value of the economy, compared to the reduction of approximately 1.2 units in the state with low aggregate productivity.[38] Welfare losses are short lived across different states of aggregate productivity as they disappear after four quarters once the economy reaches its long-run equilibrium. Over the long run, the welfare gain from the reform is equivalent to approximately 0.5 welfare units, as indicated by the dashed red line. We use equation (27) to compute an overall measure of welfare gain from the tax removal by deriving the present discounted gain from the tax removal (i.e. the weighed discounted sum of future flow values, $\sum_0^\infty \beta^t \mathbb{E}_0\{\Delta FV_t\}$). We find that the flow value associated with the implementation of the tax removal during periods with low productivity is 3.5 percent larger than the value associated with the same reform enacted in states with high aggregate productivity.

The analysis reveals that state dependence in labor market fluctuations is critical

38 The flow value of welfare is implicitly normalized by marginal utility of consumption, which is equal to 1 since preferences are linear in consumption.

Figure 11: Vacancies, vacancy-matching probability, on-the-job search, gross job destruction (EU flows), unemployment, and output over the distribution of labor productivity

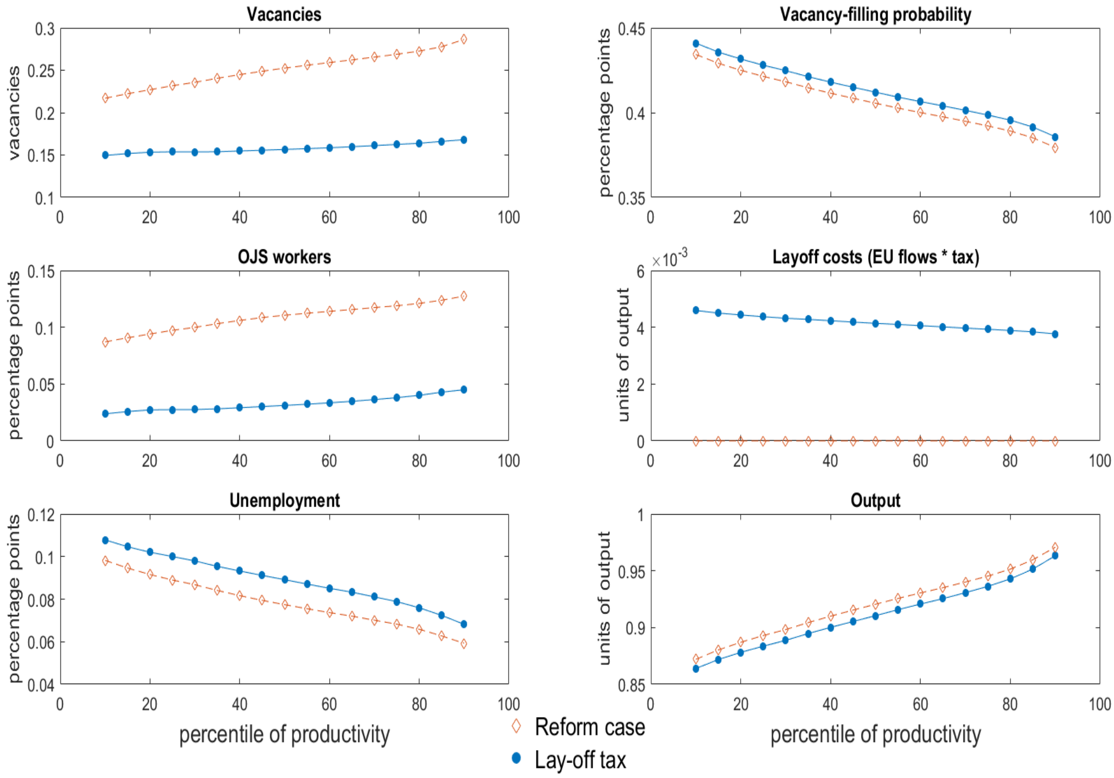

Notes. Blue solid circles represent average levels of the given variable at the respective percentile of labor productivity. Red empty diamonds represent the average instantaneous level of the same variable when the unannounced tax reform occurs (i.e. $\tau = 0$) at the respective productivity percentile. The values are computed from a simulation of 3,000 monthly periods (1,000 quarters).

for welfare and the tax removal involves sharp welfare losses in the first period of the reform. Thereafter the benefits quickly outweigh the costs. To identify the sources of welfare losses in the first period of the tax removal, Figure 11 compares the mean of selected variables in the economy with the layoff tax (circle line) against the mean of the same variables in the first period of the tax removal (diamond line) over the percentiles of the distribution of aggregate productivity. The difference between the lines represents the immediate effect of the tax removal on each variable at each percentile of aggregate productivity. The reform generates an immediate rise in output and fall in the unemployment rate that are proportional across percentiles of aggregate productivity. Instead, the magnitudes of the rises in vacancy posting and the mass of workers searching on the job, both of which are costly, increase with the level of aggregate productivity. The immediate increase in vacancies and OJS are 70 and 25 percent larger, respectively, when the reform is initiated at the 90^{th} percentile of aggregate productivity compared to the 10^{th}. The reason is intuitive. The incentive for firms and workers to establish a new job relation depends on the joint surplus of forming a match, which increases with the level of aggregate productivity. Therefore, search efforts also increase proportionally to the level of aggregate productivity. These differences are directly related to the short-run costs of the reform. In good times, the rise in vacancies is associated with a low vacancy-filling probability, which lead to a higher recruiting cost per matched firm. Furthermore, in good times, unemployment is low and OJS is high, such that the share of OJS workers in the whole pool of job seekers is high. Because of the cost k^s, the total cost of searching is higher in good times. The benefits directly related to the elimination of firing costs only vary moderately across percentiles of aggregate productivity, as depicted in the center-right diagram in the figure.

Overall, the analysis shows that tax removal involves important short-run tradeoffs mainly related to the deadweight losses of search costs that are considerably larger in states with high aggregate productivity. The timing of the labor market reforms is critical. The tax elimination in states with low productivity involves higher short-run welfare gains than during states with high aggregate productivity. Over the long run, the labor market reform is welfare enhancing across states of aggregate productivity.

3.7 Conclusion

This paper isolates important state dependence in labor market fluctuations in the US over the business cycle. The volatilities of the unemployment rate and the job separation rate are larger in periods of low aggregate labor productivity. A DMP model enriched with endogenous job separation and on-the-job search captures this state dependence through the interaction between the distribution of match-specific productivity and firms' reservation threshold for efficient matches, replicating these empirical regularities well. Our application establishes critical differences of labor market reforms enacted in distinct states of the economy for the transitional dynamics of labor market variables and welfare.

The analysis may be extended in both empirical and theoretical directions. On the empirical front, non-linear Vector Autoregression (VAR) techniques could be employed to further investigate the difference across macroeconomic regimes in the elasticity of transition rates to technology shocks. For instance, a Threshold VAR or a Regime-Switching VAR, where the autoregressive model is linear conditional on the economy's regime, would provide a first-order approximation of a nonlinear data-generating process. On the theoretical side, the mechanism of state dependence based on endogenous job separation may be recast in a comprehensive model that accounts for a broader range of real and nominal rigidities that are needed to replicate several business cycle properties in the data. The more general framework may unveil important interactions between state dependence of labor market dynamics and a broad set of macroeconomic variables. This will prove challenging, however, because it requires a non-linear solution to a complex model. It would also be interesting to use the framework to study the design of optimal labor market reforms. Future work could extend the analysis to determine the optimal provision of labor market reforms at different states of productivity and in response to contractionary and expansionary shocks. The analysis can be further extended to assess a wide range of labor market institutions (e.g. unemployment benefits, and hiring subsidies among others) to provide a comprehensive appraisal of the welfare implications of alternative labor market reforms. These investigations remain outstanding tasks for future research.

3.A Data appendix

3.A.1 Data sources

The analysis uses the following time series: real Gross Domestic Product (GDP), average labor productivity, the unemployment rate, vacancies, the job finding rate, and the separation rate. Real GDP is the non-farm business output as provided by the Bureau of Labor Statistics (BLS), while labor productivity is output per worker in the non-farm business sector. Both series were downloaded from Federal Reserve Bank of St. Louis Database (FRED). Unemployment is also provided by the BLS via FRED. The monthly job separation and job finding probabilities are computed following the continuous-time adjustment proposed by Shimer (2012). While we leave the details to the original paper, the essence of continuous time adjustment is to estimate the transition probabilities between unemployment and employment as discrete time probabilities derived from continuous-time hazard rates that are assumed to be constant within each month. This method controls for the bias of simultaneously estimating two related discrete probabilities. We used the original series provided on Rober Shimer's web page from 1950 to 2007, and extend them using the BLS data until 2014. The monthly series are then averaged over the respective quarters.

3.A.2 Empirical distribution of labor market variables

Figure 3.A.1a shows the distribution of quarterly averages in the levels (top panel) and growth rates (bottom panel) of the unemployment rate, the employment rate, productivity, output, the job finding rate and the job separation rate over the sample period. The unemployment rate is positively skewed both in levels and growth rates and similarly the employment rate is negatively-skewed. The separation rate is positively skewed both in levels and growth rates. The distribution of the job finding rate shows a slightly negative skewness in levels, which disappears in growth rates. A similar result holds for average labor productivity.

3.B Uniqueness of productivity threshold $x^r(a_t)$

Assuimng the threshold for OJS $x^s(a)$ lies above the threshold $x^r(a)$, and setting $S^s(a, z) = 0$, we can rearrange (11) into

$$
\begin{aligned}
b + k^s = a_t x_t^r + \beta \mathbb{E}_t \Big\{ (1-s)\Big(1 - p(\theta_t)\overline{F(x_{t+1}^r)}\Big)\Big((1-\lambda)S(a_{t+1}, x) \\
+ \lambda \int_{x_L}^{x_H} S(a_{t+1}, x')dF(x')\Big)\Big\}
\end{aligned}
$$

The first term on the RHS is continuous, strictly increasing, and bounded below in x over the support $[x_L,\ x_H]$. For a given θ_t, the term $\Big(1 - p(\theta_t)(1 - F(x_{t+1}^r))\Big)$ is continuous and strictly increasing in x. By the properties of $S(a, x)$, the second term on the RHS is bounded below by 0 and is continuous and weakly increasing in x. Therefore, the RHS is bounded below by 0, strictly increasing and continuous in x. These conditions are sufficient for the uniqueness of the value $x^r(a_t)$. Furthermore, as the LHS of the equation is constant, and $S(a, x)$ is increasing in a, then $x^r(a)$ must be decreasing in a.

The proof easily extends to the two cutoffs in the model with layoff taxes, as the layoff tax merely imply adding the tax terms to the LHS.

Figure 3.A.1: Empirical skewness of labor market variables

(a) Histograms of quarterly averages for the period 1950:I-2014:IV.

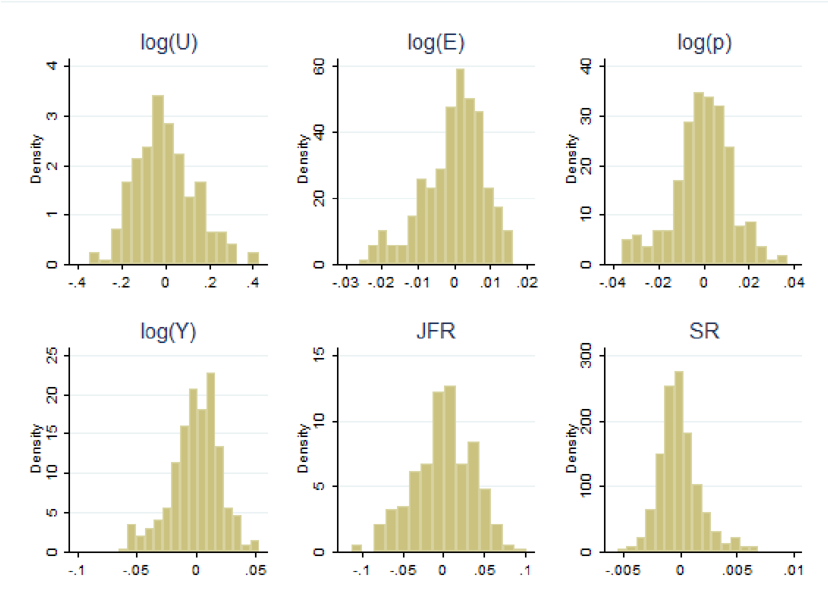

(b) Histograms of quarterly growth rates averages for the period 1950:I-2014:IV.

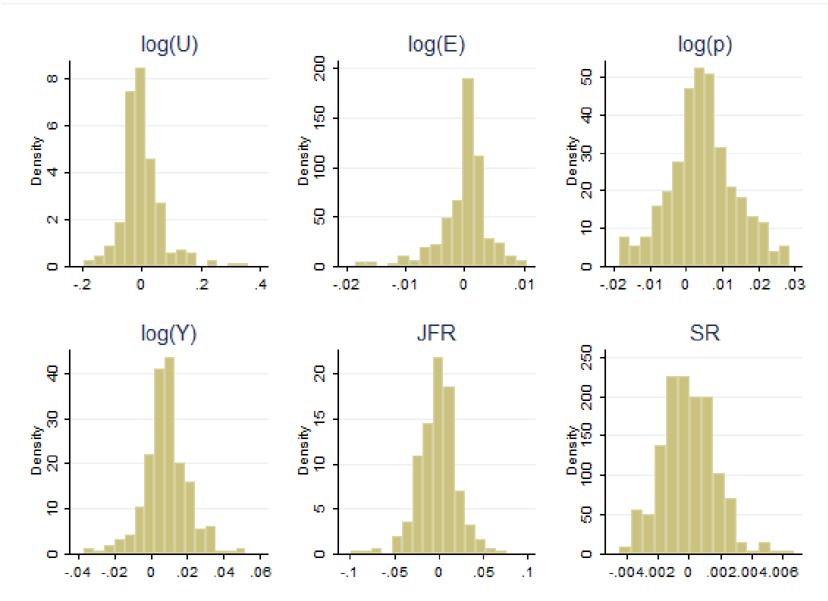

Note. U is the log unemployment rate, E is the log employment rate, p is log labor productivity, Y is log output, JFR is the job finding rate, SR is the job separation rate. Level variables are HP filtered with a smoothing parameter equal to 1,600. Growth rates are computed as log differences of quarterly averages Appendix 3.A.1 provides data sources.

3.C Job finding rate and asymmetry with respect to the state of the economy: a graphic example

Figure 3.C.1: Illustrative diagram of the mechanism driving asymmetries with respect to state of the economy in the job finding rate.

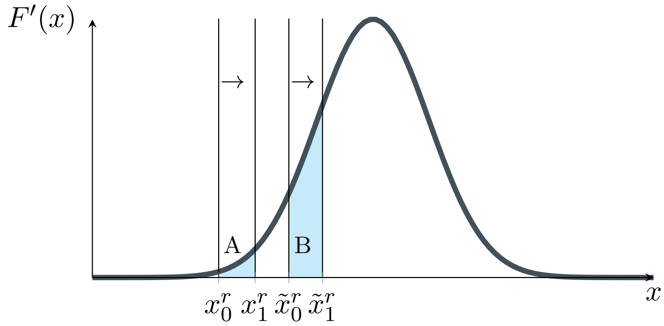

Figure 3.C.1 shows the effect of a movement in the individual productivity threshold on the p.d.f. for new workers $F'(x)$ at different levels of productivity. The figure shows that the mass of jobs sensitive to movements in the individual productivity threshold depends on the location in the support of the distribution of individual productivity shocks. The closer the productivity threshold is to the mode of the density function, the larger the increase in the mass of non-formed matches generated by a rise in the reservation threshold. For this reason, the mass of new matches affected by movements in the reservation threshold are larger when the threshold is already high (i.e. when productivity is low). The figure shows that for the same increase in the reservation threshold from x_0^r to x_1^r and from $\tilde{x}_0^r$ to $\tilde{x}_1^r$, respectively, the increase in the cumulative density function is lower when the threshold is low (shaded area A) compared to when it is high (shaded area B). Thus, the effect of a shock on the mass of new matches exposed to movements in the threshold is different across different levels of productivity. Note that this mechanism holds to the extent that $x^r(a)$ lies in the region of the domain of x where $F''(x) > 0$. In such cases, changes in the individual productivity threshold affect a larger proportion of workers when the equilibrium individual productivity threshold is higher.

3.D Solution of the model and targeted moments

We solve the model nonlinearly using an iterative procedure. To generate an accurate solution, we set the number of grid points of the state space to 45 for the discretized AR(1) process a_t, following the approach described in Tauchen (1986), and 800 for the individual productivity x. The iteration starts with a guess for the policy function of the market tightness $\theta^{(0)}(a)$. Using the guess, we compute the match surplus for all values of x and a and derive the individual productivity threshold $x^r(a)$ and the OJS threshold $x^s(a)$. Using these results, $\theta^{(1)}(a)$ is computed through the free entry condition and used again to compute the surplus. The process is repeated until the norm of $||\theta^{(0)}(x) - \theta^{(1)}(x)||$ is below a chosen critical value.[39] Using the result from Pissarides (2000), under a linear production function, hiring and layoff decisions do not depend on the aggregate level of employment or on the distribution of

39 The value we use is 10^{-8}.

individual productivity. Hence, the only relevant state variable for the policy function is the aggregate productivity factor.

To obtain business cycle statistics, we run 1,000 simulations of the model. Each simulation comprises 1,380 monthly periods. After discarding the first 600 periods, we take quarterly averages of all the simulated series to create a time series of the same length as the data. For each simulation we compute the relevant business cycle moments after taking logs and HP-filtering the series. The simulated business cycle statistics are the average of each moment across the simulations.

Table 3.D.1: Model targets

	Target	Model
Job finding rate - mean	0.45	0.442
Separation rate - mean	0.03	0.0308
Separation rate - standard deviation	0.055	0.0545
Separation rate autocorrelation	0.535	0.647
Job-to-job rate - mean	0.032	0.032
Productivity - mean	1	1.001
Productivity - standard deviation	0.013	0.013
Productivity - autocorrelation	0.757	0.7691

Note. The moments are computed as the means of 1,000 simulations of 1380 monthly periods. After discarding the first 600 observations in each simulation, the remaining series are aggregated at quarterly frequency and have the same length as the period 1950:I-2014:IV.

3.E Computation of Generalized IRFs

We resort to numerical simulations of the model to produce the response of the variables at different points in the state space. To implement the computation, it is critical to establish the starting points for the IRFs, which we describe below together with the procedure used to compute the IRFs.

After obtaining the firm's policy functions for job separation and job finding rates, we simulate the model for 10,000 monthly periods by generating a random sequence of the Markov process for productivity and then computing the relevant policy variables and state variables.[40] We then obtain the stationary distribution of labor productivity and compute the 10^{th} and 90^{th} percentiles. We split the simulated time series into two samples. The first sample includes all of the periods in which productivity is equal to 10^{th} percentile: this is the "bad-times" sample. The second sample includes all period in which productivity is equal to the 90^{th} percentile: this is the "good-times" sample. For each observation in a given sample, we collect the following variables: u_t, a_t, $e_t(x)$ and $e_{t-1}(x)$. We compute four IRFs as combinations of the following conditions: productivity is either "high" or "low", and the economy is hit by either a positive or a negative one-standard deviation productivity shock.

Each IRF is obtained through a series of 1,000 simulations. As a starting point for each simulation, we draw a random observation from the relevant sample (either "bad" or "good ") with replacement: $\{u_0, a_0, e_0(x), e_{-1}(x)\}$. We then simulate a continuous Markov path of productivity from the preset starting value a_0 to a_T, where $T = 90$ months. We compute corresponding values of $\{u_t, e_t(x), sr_t, jfr_t, jjr_t\}_{t=0}^{T}$, where sr_t

40 We start with an initial 200 observations that are then discarded.

is the separations rate, jfr_t the job finding rate, and jjr_t is the job-to-job transition rate. We compute the "alternative" history in which the path of productivity is initially hit by a further positive or negative shock in the initial period: i.e. $\{\tilde{a_t}\}_{t=0}^T$ such that $\tilde{a_0} = a_0 + \sigma$ or $\tilde{a_0} = a_0 - \sigma$. The corresponding variables under the alternative aggregate productivity path are $\{\tilde{u_t}, \widetilde{e_t(x)}, \tilde{sr_t}, \tilde{jfr_t}, \tilde{jjr_t}\}_{t=0}^T$. Because the value of $\tilde{a_t}$ does not fall on one of the notes of the discretized grid, we compute the paths of all variables using a linear interpolation of the polciy functions. After taking quarterly averages, the IRF for a given variable is computed as the difference between the values under the alternative and the baseline history: e.g. $\mathrm{d}u_q = \tilde{u_q} - u_q$ for the unemployment response, where q represents a one-quarter period. The 10^{th}, 50^{th}, and 90^{th} percentile IRFs are calculated as the relevant percentiles of $\mathrm{d}u_t$ (or any other variable) across all simulations at each $t = 0, T$.[41]

3.F Model with layoff taxes: additional details

The measure of employed workers with individual productivity below x has the following law of motion. For those workers whose individual productivity is in the OJS interval $(x_t^{r,O}, x_t^s]$:

$$\begin{aligned} e_{t+1}(x) = {} & p(\theta_t)\big(1 - e_t(x_H)\big)\big(F(x) - F(x_{t+1}^{r,N})\big) + p(\theta_t)\big(F(x) - F(x_{t+1}^{r,N})\big)e_t(x_t^s) \\ & + (1-s)\Big[\lambda\big(F(x) - F(x_{t+1}^{r,O})\big)\Big(e_t(x_H) - p(\theta_t)\overline{F(x_{t+1}^{r,N})}e_t(x_t^s)\Big) \\ & \qquad + (1-\lambda)\Big(e_t(x) - e_t(x_{t+1}^{r,O})\Big)\Big(1 - p(\theta_t)\overline{F(x_{t+1}^{r,N})}\Big)\Big], \end{aligned}$$

where $\overline{F(x)} = 1 - F(x)$. For the non-searching workers, with $x > x_t^s$:

$$\begin{aligned} e_{t+1}(x) = {} & p(\theta_t)(1 - e_t(x_H))\big(F(x) - F(x_{t+1}^{r,N})\big) + p(\theta_t)\big(F(x) - F(x_{t+1}^{r,N})\big)e_t(x_t^s) \\ & + (1-s)\Big[\lambda\big(F(x) - F(x_{t+1}^{r,O})\big)\Big(e_t(x_H) - p(\theta_t)\overline{F(x_{t+1}^{r,N})}e_t(x_t^s)\Big) \\ & + (1-\lambda)\Big(e_t(x) - e_t(x_t^s) + \big(1 - p(\theta_t)\overline{F(x_{t+1}^{r,N})}\big)\big(e_t(x_t^s) - e_t(x_{t+1}^{r,O})\big)\Big)\Big]. \end{aligned}$$

3.G Simulation of the layoff tax removal

We use a long simulation of the model economy with the tax in place to obtain a distribution of state variables $\{e_{t-1}(x), a_t\}$ when labor productivity is at the 10^{th} and 90^{th} percentiles, representing the trough and peak of the productivity cycle. From these distributions we draw (with replacement) a sample of 2,000 initial starting points. For each starting point we simulate the ensuing path of aggregate productivity through a random sequence of exogenous innovations for 25 quarters (75 months). For each sampled technology series we then compute the path of the economy under two scenarios. In the first scenario there are no changes to the economy, and the tax is expected to remain in place forever. In the second scenario there is an unannounced permanent elimination of the tax in period 0. Specifically, in the first case the path of the economy is computed using the firms' policy function for the model with tax, while in the second case the policy function used from period 0 onward is the solution for the baseline model without the tax. This alternative scenario is interpreted as the "structural reform" case.

41 It is important to note that the plotted IRFs therefore do not represent a specific and unique response, but simply the percentiles of the distribution of responses in period t.

3.H Transition path of labor market variables

Figure 3.H.1: Transition path of the separation rate, job finding rate, and job-to-job rate to a permanent elimination of the layoff tax.

(a) Separation rate

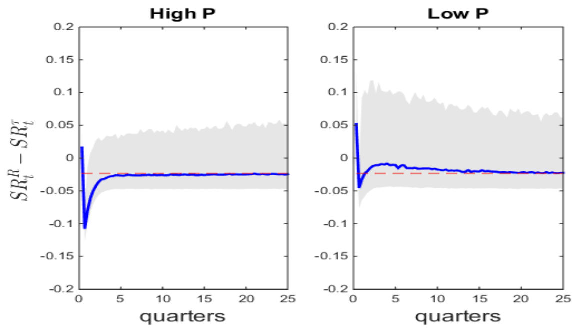

(b) Job finding rate

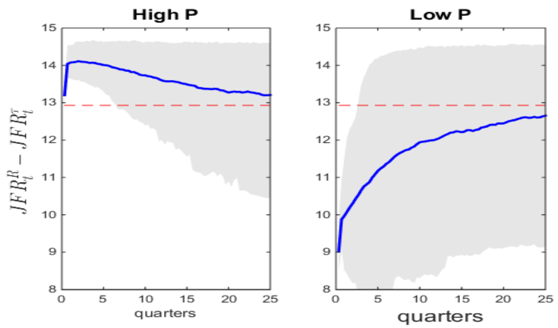

(c) Job-to-job rate

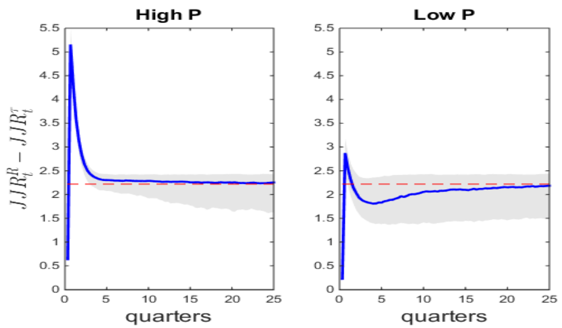

(d) Vacancies

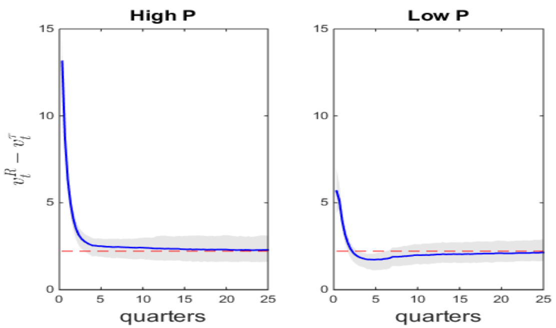

Note. The solid blue lines represent the average difference in the respective variable for the economy where at time 0 the layoff tax was abolished from that of the same economy where the tax persists from 2,000 simulations of the model. The shaded grey area represents the 10th-90th percentile interval of the differences. The dashed red line represents the long-run average difference between the two economies.

3.I Robustness checks

3.I.1 Volatility of labor market variables

Table 3.I.1: Using yearly growth rates as a measure of volatility of labor market variables.

	Volatility of YoY Growth Rate		
	p$<$ 50$^{\text{th}}$ percentile	p$>$ 50$^{\text{th}}$ percentile	$\frac{\sigma_{<50}}{\sigma_{>50}}$
Growth rate			
Unemployment	0.2302	0.1651	**1.39**
Job Finding Rate	0.0578	0.0464	**1.24**
Separation Rate	0.0031	0.0028	**1.09**
Employment Rate	0.0147	0.0091	**1.62**
Output	0.0361	0.0274	**1.32**
Productivity	0.0191	0.0168	**1.14**

Note. The table reports the standard deviation of year-on-year growth rates of labor market variables across quarters in which the cyclical component of log labor productivity is either below or above its historical median. For a given quarter, the growth rate is computed as the (log) difference with the value from 4 quarters later.

Table 3.I.2: Robustness of the variables' standard deviation to different time samples, to an HP-weight of 10^5, and to using the labor productivity series by Fernald (2014).

	Baseline			Pre Great Recession			Great Moderation		
	p< 50^{th} percentile	p> 50^{th} percentile	$\frac{\sigma_{<50}}{\sigma_{>50}}$	p< 50^{th} percentile	p> 50^{th} percentile	$\frac{\sigma_{<50}}{\sigma_{>50}}$	p< 50^{th} percentile	p> 50^{th} percentile	$\frac{\sigma_{<50}}{\sigma_{>50}}$
Levels									
Unemployment	0.1465	0.1194	**1.23**	0.1471	0.1164	**1.26**	0.1180	0.0947	**1.25**
Job Finding Rate	0.0371	0.0345	**1.07**	0.0375	0.0343	**1.09**	0.0315	0.0298	**1.06**
Separation Rate	0.0021	0.0013	**1.57**	0.0021	0.0014	**1.52**	0.0018	0.0011	**1.67**
Employment Rate	0.0095	0.0072	**1.32**	0.0092	0.0066	**1.40**	0.0088	0.0061	**1.44**
Output	0.0221	0.0165	**1.34**	0.0222	0.0165	**1.34**	0.0218	0.0115	**1.90**
Productivity	0.0119	0.0095	**1.26**	0.0122	0.0096	**1.26**	0.0105	0.0071	**1.48**
Growth rates									
Unemployment	0.0829	0.0401	**2.06**	0.0833	0.0415	**2.01**	0.0476	0.0308	**1.54**
Job Finding Rate	0.0240	0.0192	**1.25**	0.0242	0.0196	**1.23**	0.0169	0.0156	**1.08**
Separation Rate	0.0576	0.0497	**1.16**	0.0020	0.0016	**1.25**	0.0017	0.0013	**1.34**
Employment Rate	0.0051	0.0023	**2.15**	0.0049	0.0024	**2.06**	0.0034	0.0024	**1.44**
Output	0.0153	0.0090	**1.71**	0.0153	0.0091	**1.69**	0.0112	0.0068	**1.64**
Productivity	0.0099	0.0074	**1.34**	0.0100	0.0075	**1.34**	0.0083	0.0059	**1.42**

	HP weight 10^5			25^{th} – 75^{th} Cutoffs			ALP series by Fernald (2014)		
	p< 50^{th} percentile	p> 50^{th} percentile	$\frac{\sigma_{<50}}{\sigma_{>50}}$	p< 25^{th} percentile	p> 75^{th} percentile	$\frac{\sigma_{<25}}{\sigma_{>75}}$	p< 50^{th} percentile	p> 50^{th} percentile	$\frac{\sigma_{<50}}{\sigma_{>50}}$
Level									
Unemployment	0.139	0.130	**1.07**	0.148	0.120	**1.23**	0.139	0.13	**1.07**
Job Finding Rate	0.036	0.037	0.97	0.035	0.034	**1.02**	0.036	0.04	0.97
Separation Rate	0.002	0.001	**1.53**	0.002	0.001	**1.58**	0.002	0.00	**1.53**
Employment Rate	0.009	0.008	**1.19**	0.010	0.007	**1.37**	0.009	0.01	**1.19**
Output	0.022	0.018	**1.26**	0.021	0.018	**1.16**	0.022	0.02	**1.26**
Productivity	0.012	0.010	**1.24**	0.012	0.010	**1.23**	0.012	0.01	**1.25**
Growth rate									
Unemployment	0.085	0.040	**2.10**	0.090	0.037	**2.46**	0.085	0.04	**2.10**
Job Finding Rate	0.025	0.019	**1.31**	0.023	0.018	**1.24**	0.025	0.02	**1.31**
Separation Rate	0.002	0.001	**1.32**	0.002	0.002	**1.30**	0.002	0.00	**1.32**
Employment Rate	0.005	0.002	**2.41**	0.006	0.002	**2.35**	0.005	0.00	**2.41**
Output	0.015	0.009	**1.72**	0.017	0.010	**1.74**	0.015	0.01	**1.72**
Productivity	0.010	0.008	**1.30**	0.010	0.008	**1.28**	0.010	0.01	**1.35**

Note. The table reports the standard deviation of labor market variables, both in levels and growth rates, across states of low and high labor productivity for the baseline case and for a battery of robustness checks. The checks include: considering only the pre-Great Recession and the Great Moderation periods, using an HP-filter weight of 10^5 for labor productivity and the other variables in levels, using the 25^{th} and 75^{th} percentiles of productivity as thresholds, and using the factor-intensity adjusted measure of labor productivity by Fernald (2014). The ratios with a value above 1 are reported in bold font.

Table 3.I.3: Standard deviation of labor market variables using alternative definitions of low- and high- productivity regimes.

	Threshold: YoY Growth Rate of Productivity			Threshold: NBER Recessions		
	p< 50^{th} percentile	p> 50^{th} percentile	$\sigma_{<50}/\sigma_{>50}$	Recession	Recovery	$\frac{\sigma_{\text{Recession}}}{\sigma_{\text{Recovery}}}$
Level						
Unemployment	0.1434	0.1303	**1.10**	0.1683	0.1187	**1.42**
Job Finding Rate	0.0371	0.0368	**1.01**	0.0421	0.0328	**1.28**
Separation Rate	0.0020	0.0016	**1.24**	0.0022	0.0013	**1.68**
Employment Rate	0.0095	0.0075	**1.26**	0.0111	0.0073	**1.52**
Output	0.0234	0.0181	**1.29**	0.0233	0.0169	**1.37**
Productivity	0.0126	0.0103	**1.22**	0.0103	0.0095	**1.09**
Growth rate						
Unemployment	0.0718	0.0579	**1.24**	0.0802	0.0428	**1.87**
Job Finding Rate	0.0203	0.0235	0.86	0.0241	0.0189	**1.28**
Separation Rate	0.0016	0.0019	0.89	0.0027	0.0015	**1.81**
Employment Rate	0.0046	0.0031	**1.47**	0.0048	0.0025	**1.87**
Output	0.0130	0.0115	**1.12**	0.0159	0.0101	**1.58**
Productivity	0.0093	0.0090	**1.03**	0.0124	0.0083	**1.49**

	Threshold: Quarterly Growth Rate of Productivity (4Q-MA)			Threshold: Quarterly Growth Rate of Productivity		
	p< 50^{th} percentile	p> 50^{th} percentile	$\sigma_{<50}/\sigma_{>50}$	p< 50^{th} percentile	p> 50^{th} percentile	$\sigma_{<50}/\sigma_{>50}$
Level						
Unemployment	0.1435	0.1299	**1.10**	0.1297	0.1397	0.93
Job Finding Rate	0.0372	0.0366	**1.02**	0.0340	0.0387	0.88
Separation Rate	0.0020	0.0016	**1.22**	0.0020	0.0017	**1.17**
Employment Rate	0.0095	0.0075	**1.28**	0.0081	0.0087	0.93
Output	0.0234	0.0182	**1.29**	0.0218	0.0206	**1.06**
Productivity	0.0126	0.0102	**1.23**	0.0131	0.0117	**1.12**
Growth rate						
Unemployment	0.0713	0.0581	**1.23**	0.0796	0.0511	**1.56**
Job Finding Rate	0.0202	0.0236	0.85	0.0233	0.0212	**1.10**
Separation Rate	0.0016	0.0019	0.87	0.0018	0.0017	**1.05**
Employment Rate	0.0046	0.0031	**1.46**	0.0047	0.0033	**1.42**
Output	0.0130	0.0114	**1.14**	0.0132	0.0115	**1.15**
Productivity	0.0093	0.0090	**1.03**	0.0090	0.0093	0.97

Note. The table reports the standard deviations across low and high productivity states using different definitions of the two regimes. The alternative definitions are: NBER recession dates, yearly growth rates of productivity (computed as 4-quarter log differences), quarterly growth rates (computed as log differences) both using a 4-quarter moving average and in their raw values. The ratios with a value above 1 are reported in bold font.

Table 3.I.4: Correlations across alternative definitions of productivity regimes.

	ALP	**HP-Filter** 10^5	**Fernald Measure**	**NBER Recessions**	**Yearly Growth Rates**	**Quarterly Growth Rates (4Q-MA)**	**Quarterly Growth Rates**
ALP	1						
HP-Filter 10^5	0.60	1					
Fernald Measure	0.52	0.40	1				
NBER Recessions	0.40	0.36	0.34	1			
Yearly Growth Rates	0.40	0.34	0.17	0.37	1		
Quarterly Growth Rates (4Q-MA)	0.38	0.32	0.15	0.38	0.98	1	
Quarterly Growth Rates	0.20	0.11	0.08	0.22	0.32	0.34	1

Note. The table reports the correlations of different definitions of regimes, where in each case the high-productivity state is given a value of 1 and the low-productivity one a value of 0. The alternative definitions are presented in descending order based on their correlation with the baseline definition (average labor productivity). ALP stands for average labor productivity, which is the baseline definition. The alternative definitions are: ALP using an HP-filter weight of 10^5, ALP based on the factor-intensity adjusted measure of Fernald (2014), NBER recession dates, yearly growth rates of productivity (computed as 4-quarter log differences), quarterly growth rates (computed as log differences) both using a 4-quarter moving average and in their raw values.

Table 3.I.5: Cross-tabulation of low and high states across alternative threshold definitions with the baseline threshold.

Threshold	**State**	**Average Labor Productivity**		
		Below Median	**Above Median**	**All**
HP-Filter 10^5	Below Median	80%	20%	100%
	Above Median	20%	80%	100%
Fernald Measure	Below Median	76%	24%	100%
	Above Median	24%	76%	100%
NBER Recessions	Recession	93%	7%	100%
	Recovery	41%	59%	100%
Yearly Growth Rates	Below Median	70%	30%	100%
	Above Median	30%	70%	100%
Quarterly Growth Rates (4Q-MA)	Below Median	69%	31%	100%
	Above Median	31%	69%	100%
Quarterly Growth Rates	Below Median	60%	40%	100%
	Above Median	40%	60%	100%

Note. For each alternative regime definition and for both the low- and high-productivity states, the table reports the percent of quarters in which ALP (i.e. the baseline regime definition) indicates a low state (i.e. productivity below median) or a high one (i.e. productivity above median). The alternative definitions are presented in descending order based on their correlation with ALP, based on Table 3.I.4. The alternative definitions are: ALP using an HP-filter weight of 10^5, ALP based on the factor-intensity adjusted measure of Fernald (2014), NBER recession dates, yearly growth rates of productivity (computed as 4-quarter log differences), quarterly growth rates (computed as log differences) both using a 4-quarter moving average and in their raw values.

3.I.2 Regression analysis

Table 3.I.6: Robustness checks of the regressions in Table 2 for the unemployment rate and the job finding rate.

	Unemployment				JFR			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log p_t	-4.781***	-3.916***	-3.553***	-5.197***	0.838***	0.618**	0.865***	1.172***
	(0.631)	(0.546)	(0.618)	(0.631)	(0.242)	(0.251)	(0.259)	(0.239)
High p_t	0.012	0.017*	-0.000	0.006	-0.001	-0.003	-0.003	0.000
	(0.010)	(0.009)	(0.013)	(0.009)	(0.004)	(0.004)	(0.005)	(0.003)
log p_t * High p_t	2.622***	2.371***	2.302**	3.010***	-0.693*	-0.251	-0.659*	-1.384***
	(0.963)	(0.838)	(0.947)	(0.953)	(0.378)	(0.388)	(0.392)	(0.369)
log p_{t-1}	-0.057	-0.281	1.465**	0.179	0.350	0.396	-0.357	0.199
	(0.684)	(0.584)	(0.647)	(0.664)	(0.262)	(0.266)	(0.270)	(0.252)
High p_{t-1}	-0.005	0.000	0.003	0.008	0.002	0.003	0.003	-0.007*
	(0.010)	(0.009)	(0.013)	(0.009)	(0.004)	(0.004)	(0.005)	(0.003)
log p_{t-1} * High p_{t-1}	0.924	0.435	-1.245	-0.081	-0.118	-0.446	0.640	0.878**
	(0.981)	(0.847)	(0.950)	(0.958)	(0.383)	(0.391)	(0.394)	(0.371)
Lagged Dependent	0.874***	0.908***	0.925***	0.829***	0.761***	0.837***	1.003***	0.716***
	(0.026)	(0.024)	(0.020)	(0.025)	(0.037)	(0.039)	(0.046)	(0.035)
Constant	-0.023***	-0.023***	-0.013	-0.023***	0.004	0.004	0.001	0.006**
	(0.008)	(0.006)	(0.009)	(0.007)	(0.003)	(0.003)	(0.004)	(0.003)
Observations	231	152	259	259	231	152	259	259
R-squared	0.883	0.935	0.918	0.894	0.759	0.831	0.726	0.782
Sample	1950-2007	1980-2007	1950-2014	1950-2014	1950-2007	1980-2007	1950-2014	1950-2014
HP-filter	1600	1600	10^5	1600	1600	1600	10^5	1600
Productivity Series	ALP	ALP	ALP	Fern.	ALP	ALP	ALP	Fern.

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note. The dependent variables are the log unemployment rate and the job-finding rate. The table reports robustness checks for the baseline regressions. In column (1) and (5), we restrict the sample to the pre-Great Recession period. In columns (2) and (6), we restrict the sample to the Great Moderation period. In columns (3) and (7) we use variables HP-filtered with a smoothness weight of 10^5. In columns (4) and (8) we use the factor-intensity adjusted measure of labor productivity from Fernald (2014). The explanatory variables are the log productivity, a dummy variable equal to one when productivity at time t is larger than the historical median, and an interaction term between the log productivity and the dummy variable.

Table 3.I.7: Robustness checks of the regressions in Table 2 for the job separation rate and the employment rate.

	SR				Employment			
	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
$\log \mathrm{p}_t$	-0.119***	-0.139***	-0.081***	-0.131***	0.288***	0.274***	0.226***	0.335***
	(0.018)	(0.020)	(0.016)	(0.018)	(0.035)	(0.039)	(0.037)	(0.036)
High p_t	0.000	0.000	-0.000	0.000	-0.001*	-0.001**	-0.000	-0.001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
$\log \mathrm{p}_t *$ High p_t	0.105***	0.149***	0.080***	0.072**	-0.161***	-0.176***	-0.162***	-0.223***
	(0.029)	(0.031)	(0.024)	(0.029)	(0.053)	(0.059)	(0.057)	(0.054)
$\log \mathrm{p}_{t-1}$	0.028	0.021	0.047***	0.046**	0.060	0.042	-0.061	0.028
	(0.020)	(0.021)	(0.016)	(0.019)	(0.038)	(0.042)	(0.039)	(0.038)
High p_{t-1}	-0.000	0.000	-0.000	-0.000	-0.000	-0.000	-0.000	-0.001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
$\log \mathrm{p}_{t-1} *$ High p_{t-1}	-0.033	-0.080**	-0.060**	-0.022	-0.104*	-0.049	0.042	-0.028
	(0.030)	(0.032)	(0.025)	(0.029)	(0.054)	(0.060)	(0.057)	(0.054)
Lagged Dependent	0.348***	0.233***	0.505***	0.313***	0.862***	0.893***	0.932***	0.828***
	(0.063)	(0.075)	(0.057)	(0.061)	(0.025)	(0.025)	(0.020)	(0.023)
Constant	-0.000*	-0.001***	-0.000	-0.000	0.002***	0.002***	0.001***	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)
Observations	231	152	259	259	231	152	259	259
R-squared	0.459	0.536	0.451	0.484	0.904	0.933	0.926	0.914
Sample	1950-2007	1980-2007	1950-2014	1950-2014	1950-2007	1980-2007	1950-2014	1950-2014
HP-filter	1600	1600	$10\hat{5}$	1600	1600	1600	$10\hat{5}$	1600
Productivity Series	ALP	ALP	ALP	Fern.	ALP	ALP	ALP	Fern.

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note. The dependent variables are the separation rate and the log employment rate. The table reports robustness checks for the baseline regressions. In column (1) and (5), we restrict the sample to the pre-Great Recession period. In columns (2) and (6), we restrict the sample to the Great Moderation period. In columns (3) and (7) we use variables HP-filtered with a smoothness weight of 10^5. In columns (4) and (8) we use the factor-intensity adjusted measure of labor productivity from Fernald (2014). The explanatory variables are the log productivity, a dummy variable equal to one when productivity at time t is larger than the historical median, and an interaction term between the log productivity and the dummy variable.

3.I.3 Skewness

Table 3.I.8: Robustness of the variables' skewness to different time samples and to using an HP-weight of 10^5.

	Baseline		Pre- Great Recession		Great Moderation		HP weight 10^5
	Level	Growth Rate	Level	Growth Rate	Level	Growth Rate	Level
Unemployment	0.402	1.386	0.452	1.401	0.099	0.918	0.113
Job Finding Rate	-0.293	-0.385	-0.337	-0.391	-0.199	-0.042	-0.202
Separation Rate	0.772	0.478	0.747	0.475	0.780	0.289	0.727
Employment Rate	-0.666	-1.214	-0.724	-1.143	-0.762	-0.660	-0.668
Output	-0.509	-0.098	-0.531	-0.037	-1.074	-0.859	-0.326
Productivity	-0.363	0.070	-0.366	0.076	-0.872	-0.261	-0.291

Note. The table reports sensitivity analysis on the skewness of labor market variables (in both levels and growth rates). The three checks restrict the sample to the pre-Great Recession period and the Great Moderation period, and using an HP-filter weight of 10^5.

Table 3.I.9: Simulated skewness of monthly growth rates in the model.

	Average Skewness
Unemployment	0.855
Job Finding Rate	-0.063
Separation Rate	0.573
Employment Rate	-0.998
Output	-0.193
Productivity	0.067

Note. The table reports the average skewness of the simulated data. Entries are averages of 1,000 simulations over 1,380 monthly periods, discarding the first 600 observations in each simulation, the remaining series have the same length as the period 1950:I-2014:IV.

Bibliography

Abbritti, M. and S. Fahr (2013): "Downward Wage Rigidity and Business Cycle Asymmetries," *Journal of Monetary Economics*, 60, 871–886.

Ahn, H. J. and J. D. Hamilton (2016): "Heterogeneity and Unemployment Dynamics," Working Paper.

Altissimo, F. and G. L. Violante (2001): "The Non-linear Dynamics of Output and Unemployment in the U.S," *Journal of Applied Econometrics*, 16, 461–486.

Aruoba, S. B., L. Bocola, and F. Schorfheide (2013): "Assessing DSGE Model Nonlinearities," *National Bureau of Economic Research Working Paper Series*, No. 19693.

Attanasio, O., P. Levell, H. Low, and V. Sanchez-Marcos (2015): "Aggregating Elasticities: Intensive and Extensive Margins of Female Labor Supply," 1–45.

Attanasio, O., H. Low, and V. Sanchez-Marcos (2005): "Female Labor Supply as Insurance against Idiosyncratic Risk," *Journal of the European Economic Association*, 3, 755–764.

Attanasio, O., H. Low, and V. Sánchez-Marcos (2008): "Explaining Changes in Female Labor Supply in a Life-Cycle Model in Female Labor Supply Explaining Changes Life-Cycle Model," *The American Economic Review*, 98, 1517–1552.

Attanasio, O. P., R. Bottazzi, H. W. Low, L. Nesheim, and M. Wakefield (2012): "Modelling the Demand for Housing over the Life Cycle," *Review of Economic Dynamics*, 15, 1–18.

Bachmann, R., R. J. Caballero, and E. M. R. A. Engel (2013): "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model," *American Economic Journal: Macroeconomics*, 5, 29–67.

Bachmann, R. and M. Sinning (2016): "Decomposing the Ins and Outs of Cyclical Unemployment," *Oxford Bulletin of Economics and Statistics*, 78, 853–876.

Bai, J. and S. Ng (2005): "Tests for Skewness, Kurtosis, and Normality for Time Series Data," *Journal of Business & Economic Statistics*, 23, 49–60.

Bajari, P., P. Chan, D. Krueger, and D. Miller (2013): "A Dynamic Model of Housing Demand: Estimation and Policy Implications," *International Economic Review*, 54, 409–442.

Banks, J., Z. Smith, and M. Wakefield (2002): "The Distribution of Financial Wealth in the UK: Evidence from 2000 BHPS Data," *IFS working Paper*.

Barattieri, A., S. Basu, and P. Gottschalk (2014): "Some Evidence on the Importance of Sticky Wages," *American Economic Journal: Macroeconomics*, 6, 70–101.

Barnichon, R. and A. Figura (2015): "Labor Market Heterogeneities and the Aggregate Matching Function," *American Economic Journal: Macroeconomics*, 7, 222–249.

Barron, J. M., M. C. Berger, and D. A. Black (1997): "Employer Search, Training, and Vacancy Duration," *Economic Inquiry*, 35, 167–192.

Barron, J. M. and J. Bishop (1985): "Extensive Search, Intensive Search, and Hiring Costs: New Evidence on Employer Hiring Activity," *Economic Inquiry*, 23, 363–382.

Beaudry, P. and G. Koop (1993): "Do Recessions Permanently Change Output?" *Journal of Monetary Economics*, 31, 149–163.

Benigno, P. and L. A. Ricci (2011): "The Inflation-Output Trade-Off with Downward Wage Rigidities," *American Economic Review*, 101, 1436–66.

Benigno, P., L. A. Ricci, and P. Surico (2015): "Unemployment and Productivity in the Long Run: The Role of Macroeconomic Volatility," *The Review of Economics and Statistics*, 97, 698–709.

Bentolila, S., P. Cahuc, J. J. Dolado, and T. Le Barbanchon (2012): "Two-Tier Labour Markets in the Great Recession: France Versus Spain," *The Economic Journal*, 122, F155–F187.

Blau, F. D. and L. M. Kahn (2007): "Changes in the Labor Supply Behavior of Married Women : 1980 2000," *Journal of L*, 25, 393–438.

Blundell, B. R., L. Pistaferri, and I. Saporta-eksten (2016): "Consumption Inequality and Family Labor Supply," *American Economic Review*, 106, 387–435.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98, 1887–1891.

Borella, M., M. De Nardi, and F. Yang (forthcoming): "The Aggregate Implications of Gender and Marriage," *Journal of the Economics of Ageing*.

Bottazzi, R., H. Low, and M. Wakefield (2007): "Why Do Home Owners Work Longer Hours?" *IFS Working Paper*, WP10/07, 1–33.

Burbidge, J. B., L. Magee, and a. L. Robb (1988): "Alternative Transformations to Handle Extreme Values of the Dependent Variable," *Journal of the American Statistical Association*, 83, 123–127.

Cacciatore, M., R. Duval, G. Fiori, and F. Ghironi (2016): "Market Reforms in the Time of Imbalance," *Journal of Economic Dynamics and Control*, 72, 69–93.

Cacciatore, M. and G. Fiori (2016): "The Macroeconomic Effects of Goods and Labor Marlet Deregulation," *Review of Economic Dynamics*, 20, 1–24.

Cacciatore, M., G. Fiori, and F. Ghironi (2015): "The Domestic and International Effects of Euro Area Market Reforms," Mimeo, HEC Montreal, North Carolina State University and University of Washington.

Caggiano, G., E. Castelnuovo, and N. Groshenny (2014): "Uncertainty Shocks and Unemployment Dynamics in U.S. Recessions," *Journal of Monetary Economics*, 67, 78 – 92.

Campolmi, A. and E. Faia (2011): "Labor Market Institutions and Inflation Volatility in the Euro Area," *Journal of Economic Dynamics and Control*, 35, 793–812.

Cozzi, M. and G. Fella (2016): "Job Displacement Risk and Severance Pay," *Journal of Monetary Economics*, 84, 166 – 181.

Darby, M., J. Haltiwanger, and M. Plant (1986): "The Ins and Outs of Unemployment: The Ins Win," *NBER Working Paper*.

Darby, M. R., J. Haltiwanger, and M. Plant (1985): "Unemployment Rate Dynamics and Persistent Unemployment under Rational Expectations," *American Economic Review*, 75, 614–637.

Davis, S. J. and J. Haltiwanger (2014): "Labor Market Fluidity and Economic Performance," *National Bureau of Economic Research Working Paper Series*, No. 20479, 98.

Deaton, A. S. and C. H. Paxson (1993): "Saving , Growth, and Aging in Taiwan," *NBER Working Paper*, 4330.

Del Boca, D. and A. Lusardi (2003): "Credit Market Constraints and Labor Market Decisions," *Labour Economics*, 10, 681–703.

Del Río, A. and G. Young (2008): "The Impact of Unsecured Debt on Financial Pressure among British Households," *Applied Financial Economics*, 18, 1209–1220.

den Haan, W. J., G. Ramey, and J. Watson (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90, 482–498.

Disney, R. and J. Gathergood (2013): "House Prices , Wealth Effects and Labour Supply Richard Disney and John Gathergood by," Working Paper.

Druedahl, J. (2015): "The Demand for Housing over the Life-Cycle under Long-Term Gross Debt Contracts," Working Paper.

Elsby, M. W. L., B. Hobijn, and A. Şahin (2015): "On the Importance of the Participation Margin for Labor Market Fluctuations," *Journal of Monetary Economics*, 72, 64–82.

Elsby, M. W. L., J. C. Smith, and J. Wadsworth (2011): "The Role of Worker Flows in the Dynamics and Distribution of UK Unemployment," *Oxford Review of Economic Policy*, 27, 338–363.

FCA (2014): "Consultation Guidance on the Financial Policy Committee's Recommendation on Loan to Income Ratios in Mortgage Lending," Tech. Rep. August, Financial Conduct Authority.

Fella, G. (2007): "When Do Firing Taxes Matter?" *Economics Letters*, 97, 24–31.

Fernald, J. G. (2014): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," *Federal Reserve Bank of San Francisco Working Paper Series*, 2012-19.

Fernández-Villaverde, J. and D. Krueger (2011): "Consumption and Saving Over the Life Cycle: How Important Are Consumer Durables?" *Macroeconomic Dynamics*, 15, 725–770.

Ferraro, D. (2016): "The Asymmetric Cyclical Behavior of the U.S. Labor Market," Working paper series, Arizona State University.

Fortin, N. M. (1995): "Allocation Inflexibilities, Female Labor Supply, and Housing Assets Accumulation: Are Women Working to Pay the Mortgage?" *Journal of Labor Economics*, 13, 524–557.

FSA (2004): "FSA Guide to Mortgages," Tech. rep., Financial Services Authority.

——— (2009): "Mortgage Market Review," Tech. Rep. October, Financial Services Authority.

Fujita, S. and G. Ramey (2012): "Exogenous versus Endogenous Separation," *American Economic Journal: Macroeconomics*, 4, 68–93.

Gavazza, A., S. Mongey, and G. L. Violante (2016): "Aggregate Recruiting Intensity," Working Paper.

Gomes, P. (2012): "Labour Market Flows: Facts from the United Kingdom," *Labour Economics*, 19, 165–175.

Gortz, C. and J. D. Tsoukalas (2013): "Learning, Capital-Embodied Technology and Aggregate Fluctuations," *Review of Economic Dynamics*, 16, 708 – 723.

Gottfries, A. (2018): "Bargaining with Renegotiation in Models with on-the-Job Search," Working Paper.

Hagedorn, M. and I. Manovskii (2008): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, 98, 1692–1706.

Hall, R. and P. Milgrom (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98, 1653–1674.

Hall, R. E. and S. Schulhofer-Wohl (2018): "Measuring Matching Efficiency with Heterogeneous Jobseekers," *American Economic Journal: Macroeconomics*, 10, 1–32.

Hobijn, B. and A. Şahin (2012): "Beveridge Curve Shifts across Countries since the Great Recession," Working Paper.

Honoré, B. E. (1992): "Trimmed Lad and Least Squares Estimation of Truncated and Censored Regression Models with Fixed Effects," *Econometrica*, 60, 533–565.

Hornstein, A., M. Kudlyak, and F. R. B. Richmond (2015): "Estimating Matching Efficiency with Variable Search Effort," Working Paper.

Hosios, A. J. (1990): "On The Efficiency of Matching and Related Models of Search and Unemployment," *Review of Economic Studies*, 57, 279–298.

Hsu, J. W., D. A. Matsa, and B. T. Melzer (2017): "Unemployment Insurance as a Housing Market Stabilizer," *American Economic Review*, 108, 49–81.

Hyslop, D. R. (2001): "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings," *American Economic Review*, 91, 755–777.

Iacoviello, M. (2008): "Household Debt and Income Inequality, 1963 2003," *Journal of Money, Credit and Banking*, 40, 929–965.

Iacoviello, M. and M. Pavan (2013): "Housing and Debt over the Life Cycle and over the Business Cycle," *Journal of Monetary Economics*, 60, 221–238.

Kaplan, G. and G. L. Violante (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82, 1199–1239.

Kim, J. and F. J. Ruge-Murcia (2009): "How Much Inflation is Necessary to Grease the Wheels?" *Journal of Monetary Economics*, 56, 365–377.

Klinger, S. and E. Weber (2016): "Decomposing Beveridge Curve Dynamics By Correlated Unobserved Components," *Oxford Bulletin of Economics and Statistics*, 78, 877–894.

Kohlbrecher, B. and C. Merkl (2016): "Business Cycle Asymmetries and the Labor Market," Working Paper.

Koopman, S. J., N. Shephard, and J. A. Doornik (1999): "Statistical Algorithms for Models in State Space using SsfPack 2.2," *Econometrics Journal*, 2, 107–160.

Krause, M. U. and T. A. Lubik (2007): "The (Ir)Relevance of Real Wage Rigidity in the New Keynesian Model with Search Frictions," *Journal of Monetary Economics*, 54, 706–727.

Kroft, K., F. Lange, M. J. Notowidigdo, and L. F. Katz (2016): "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Nonparticipation," *Journal of Labor Economics*, 34, S7–S54.

Krueger, A. B., J. Cramer, and D. Cho (2014): "Are the Long-Term Unemployed on the Margins of the Labor Market?" *Brookings Papers on Economic Activity*, 229–300.

Lise, J. and S. Seitz (2007): "Consumption Inequality and Intra-Household Allocations," *IFS Working Paper.*

Ljungqvist, L. (2002): "How Do Lay-off Costs Affect Employment?" *The Economic Journal*, 112, 829–853.

Ljungqvist, L. and T. J. Sargent (2012): *Recursive Macroeconomic Theory, Third Edition*, vol. 1 of *MIT Press Books*, The MIT Press.

Llosa, G., L. Ohanian, A. Raffo, and R. Rogerson (2014): "Firing Costs and Labor Market Fluctuations: A Cross-Country Analysis," Working Paper.

Low, H. W. (2005): "Self-insurance in a Life-Cycle Model of Labour Supply and Savings," *Review of Economic Dynamics*, 8, 945–975.

Lundberg, S. (1985): "The Added Worker Effect," *Journal of Labor Economics*, 3, 11–37.

Mankart, J. and R. Oikonomou (2017): "Household Search and the Aggregate Labour Market," *The Review of Economic Studies*, 1735–1788.

McKay, A. and R. Reis (2008): "The Brevity and Violence of Contractions and Expansions," *Journal of Monetary Economics*, 55, 738–751.

Mian, A., K. Rao, and A. Sufi (2013): "Household Balance Sheets, Consumption, and the Ecoomic Slump," *Quarterly Journal of Economics*, 128, 1687–1726.

Michaillat, P. (2014): "A Theory of Countercyclical Government Multiplier," *American Economic Journal: macroeconomics*, 6, 190–217.

Mitman, K. (2016): "Macroeconomic Effects of Bankruptcy and Foreclosure Policies," *American Economic Review*, 106, 2219–2255.

Morchio, I. (2016): "Work Histories and Lifetime Unemployment," Working Paper.

Mortensen, D. T. and C. A. Pissarides (1999): "Unemployment Responses to 'Skill-Biased' Technology Shocks: The Role of Labour Market Policy," *The Economic Journal*, 109, 242–265.

Moscarini, G. (2005): "Job Matching and the Wage Distribution," *Econometrica*, 73, 481–516.

Mumtaz, H. and P. Surico (2015): "The Transmission Mechanism In Good And Bad Times," *International Economic Review*, 56, 1237–1260.

Neftci, S. N. (1984): "Are Economic Time Series Asymmetric over the Business Cycle?" *Journal of Political Economy*, 92, 307–28.

Nickell, S., L. Nunziata, and W. Ochel (2005): "Unemployment in the OECD Since the 1960s. What Do We Know?*," *The Economic Journal*, 115, 1–27.

ONS (2016): "Annual Survey of Hours and Earnings: 2015 Provisional Results." Tech. rep., Office of National Statistics.

Panagiotidis, T. and G. Pelloni (2007): "Nonlinearity in the Canadation and U.S Labor Markets: Univariate and Multivariate Evidence from a Battery of Tests," *Macroeconomic Dynamics*, 11, 613–637.

Patterson, C., A. ahin, G. Topa, and G. L. Violante (2016): "Working Hard in the Wrong Place: A Mismatch-Based Explanation to the UK Productivity Puzzle," *European Economic Review*, 84, 42–56.

Petrongolo, B. and C. Pissarides (2001): "Looking into the black box: A survey of the matching function," *Journal of Economic Literature*, 39, 390–431.

Petrosky-Nadeau, N. and L. Zhang (2013): "Unemployment Crises," Working Paper 19207, National Bureau of Economic Research.

——— (2017): "Solving the Diamond-Mortensen-Pissarides Model Accurately," *Quantitative Economics*, 8, 611–650.

Pissarides, C. A. (2000): *Equilibrium Unemployment Theory, 2nd Edition*, vol. 1 of *MIT Press Books*, The MIT Press.

Postel-Vinay, F. and J.-M. Robin (2004): "To Match or Not to Match?: Optimal Wage Policy with Endogenous Worker Search Intensity," *Review of Economic Dynamics*, 7, 297–330.

Postel-Vinay, F. and H. Turon (2014): "The Impact of Firing Restrictions on Labour Market Equilibrium in the Presence of on-the-Job Search," *Economic Journal*, 124, 31–61.

Şahin, A., J. Song, G. Topa, and G. L. Violante (2014): "Mismatch Unemployment," *American Economic Review*, 104, 3529–3564.

Sedláček, P. (2016): "The Aggregate Matching Function and Job Search from Employment and out of the Labor Force," *Review of Economic Dynamics*, 21, 16–28.

Sedláček, P. (2014): "Match Efficiency and Firms' Hiring Standards," *Journal of Monetary Economics*, 62, 123–133.

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

——— (2006): "On-the-Job Search and Strategic Bargaining," *European Economic Review*, 50, 811–830.

——— (2012): "Wage rigidities and jobless recoveries," *Journal of Monetary Economics*, 59, S65–S77.

Sichel, D. E. (1989): "Are Business Cycles Asymmetric? A Correction," *Journal of Political Economy*, 97, 1255–1260.

Singleton, C. (2017): "Long-Term Unemployment and the Great Recession: Evidence from UK Stocks and Flows," Working Paper.

Smith, J. C. (2011): "The Ins and Outs of UK Unemployment," *Economic Journal*, 121, 402–444.

Stephens, M. J. (2002): "Worker Displacement and the Added Worker Effect," *Journal of Labor Economics*, 20, 504–537.

Sutton, A. (2013): "On the Determinants of UK Unemployment and the Great Recession: Analysing the Gross Flows Data," *Applied Economics*, 45, 3599–3616.

Tauchen, G. (1986): "Finite State Markov Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181.

Thomas, C. and F. Zanetti (2009): "Labor Market Reform and Price Stability: An Application to the Euro Area," *Journal of Monetary Economics*, 56, 885–899.

Wooldridge, J. M. (2010): *Econometric Analysis of Cross Section and Panel Data, Second Edition*, MIT Press.

Wu, C. and D. Krueger (2016): "How Much Consumption Insurance in Bewley Models with Endogenous Family Labor Supply ?" Working Paper.

Yang, F. (2009): "Consumption over the Life Cycle: How Different is Housing?" *Review of Economic Dynamics*, 12, 423–443.

Zanetti, F. (2011): "Labor Market Institutions and Aggregate Fluctuations in a Search and Matching Model," *European Economic Review*, 55, 644–658.

Zhao, B. (2018): "Too Poor to Retire? Housing Prices and Retirement," *Review of Economic Dynamics*, 27, 27–47.