University of California
Santa Barbara

The Long and Short of Labor Supply Changes

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Economics

by

Edwin Charles Nusbaum IV

Committee in charge:

Professor Peter Rupert, Chair
Professor Douglas Steigerwald
Professor Javier Birchenall

March 2021

The Dissertation of Edwin Charles Nusbaum IV is approved.

Professor Douglas Steigerwald

Professor Javier Birchenall

Professor Peter Rupert, Committee Chair

March 2021

The Long and Short of Labor Supply Changes

To my stepfather, Jeffrey Sutherland, who has long served as the quiet, organizing force behind my family's success.

.

Acknowledgements

While the completion of my dissertation required hard work and dedication on my part, it would not have been possible without the help of many others. First, I would like to thank Peter Rupert, Doug Steigerwald, Javier Birchenall, and Thomas Cooley as each has taken a keen interest in my success. Their invaluable advice and support has led directly to this document. Beyond my research, we have developed personal relationships. Spending time with them in and out of the office has afforded me more joy and fun than I could have hoped for during my education. I am happy to walk away from this process calling each a friend.

I would also like to thank my friends. Going through this process with a tight-nit cohort has made the trials and tribulations of graduate school significantly easier to overcome. Moreover, no discussion of friends would be complete without mentioning those that became family: Chris Ames, Matty Greer, Rachel Paterno, Rob Gazzara, Jenny Critchley, Tyler Main, and Amanda Rhiel. Since becoming a Stockton Osprey, they have been some of my biggest cheerleaders. Whether offering a breath of fresh air when I visit home or calling me via Zoom to offer support, they have been ready for anything.

Furthermore, I am forever grateful to many in my family: Yolanda Verhoeven, Jeffrey Sutherland, Edwin Nusbaum III, Lauren Sutherland, and Eva Nusbaum. For 29 years, they have given me more love and support than anyone could ask for. They have served as role models as I've discovered what I want out of life, including my graduate education. My sisters, Lauren and Eva, taught me how to conquer adversity. My father, Edwin, has inspired a fire in me to work hard for my success. My mother, Yolanda, and stepfather, Jeff, have been guiding lights in all that I've accomplished. I quite literally would not be here without them.

Finally, words cannot express my indebtedness to my partner, Alina Harper. She has significantly compromised her career to accompany me on this journey and has been a pillar of emotional support throughout. My graduate education has been a roller coaster looping from

the bliss of completing my first peer reviewed economics article to the frustration of hitting proverbial roadblocks for months on end. Alina, more than anyone else, has born the brunt of this whiplash. She has given a large part of her life to this endeavor and I cannot thank her enough. Whatever success this dissertation represents, Alina shares equally.

Curriculum Vitæ

Edwin Charles Nusbaum IV

Department of Economics
University of California, Santa Barbara
1119 North Hall, Santa Barbara, CA 93106-9210

enusbaum@ucsb.edu
www.charlienusbaum.com

RESEARCH INTERESTS

Macroeconomics, Demographic Change, Labor Economics, Applied Econometrics

EDUCATION

University of California, Santa Barbara — Santa Barbara, CA
Ph.D., Economics — *Expected March 2021*

University of California, Santa Barbara — Santa Barbara, CA
M.A., Economics — *June 2016*

Richard Stockton College of New Jersey (*Stockton University*) — Galloway, NJ
B.S. in Applied Physics, Minor in Mathematics — *May 2014*

PUBLICATIONS

Labor Market Dynamics and the Migration Behavior of Married Couples **(2020)**
Review of Economic Dynamics, with Christine Braun and Peter Rupert

Calibration of a Quasi-Adiabatic Magneto-Thermal Calorimeter Used to Characterize Magnetic Nanoparticle Heating **(2013)**
Journal of Nanotechnology in Engineering and Medicine, with Anilchandra Attaluri, Michele Wabler, and Robert Ivkov

WORKING PAPERS

Aggregate Hours and Local Projections with Long-Run Restrictions

Demographic Obstacles to European Growth
with Thomas Cooley and Espen Henriksen

WORKS IN PROGRESS

Age Dependent Sentencing Reform and Economic Mobility

TEACHING EXPERIENCE

Invited Lectures — University of California, Irvine
Markov Processes with Applications to Search Theory — *Summer 2017*
Introduction to Bayesian Estimation of DSGE Models — *Summer 2017*

Head Teaching Assistant — University of California, Santa Barbara
Intermediate Micro I – ECON 10A — *Summer 2017*

Teaching Assistant — University of California, Santa Barbara
Macroeconomic Theory (Ph.D.) – ECON 204A-C — *Fall 2017, Winter 2018, Spring 2018*
Intermediate Micro I – ECON 10A — *Fall 2016, Spring 2017, Summer 2017*
Principles of Macro – ECON 2 — *Winter 2017*
Intermediate Micro II – ECON 100B — *Summer 2016*

Teaching Assistant — Richard Stockton College of New Jersey
Physics II Lab — *Fall 2012, Spring 2013*
Physics I Lab — *Fall 2011, Spring 2012*

ACADEMIC PRESENTATIONS

Society for Economic Dynamics, WUSTL	2019
Macroeconomics Brown Bag Seminar, UC Irvine	2019
Midwest Macro Fall Meetings, Vanderbilt	2018
3rd Annual All California Macroeconomic Conference, Claremont McKenna	2018
13th Annual Economics Graduate Student Conference, WUSTL	2018
Macroeconomics Brown Bag Seminar, UC Irvine	2018
Finn Kydland Macroeconomics Workshop, UC Santa Barbara	2018
UC Graduate Macroeconomics Workshop, UC Irvine	2017
Physics Colloquium, Richard Stockton College	2014
NAMS Research Symposium, Richard Stockton College	2013
Nano-Bio Symposium, Johns Hopkins University	2012

OTHER ACADEMIC CONTRIBUTIONS

U.S. Economic Snapshot Blog *www.econsnapshot.com*
with Thomas Cooley and Peter Rupert

European Economic Snapshot Blog *www.europeansnapshot.com*
with Thomas Cooley and Peter Rupert

Laboratory for Aggregate Economics & Finance Contributor *www.laef.ucsb.edu*

RECENT PROFESSIONAL EXPERIENCE

Keller Rohrback L.L.P. Seattle, WA
Consulting Expert *March 2019 – Present*

UCSB Economic Forecast Project Santa Barbara, CA
Senior Research Analyst *July 2018 – Present*

Research Assistant University of California, Santa Barbara
Kelsey Jack *August 2019 – October 2019*

Research Assistant University of California, Santa Barbara
Heather Royer *September 2018 – December 2018*

Research Assistant University of California, Santa Barbara
Javier Birchenall *September 2017 – December 2017*

PAST PROFESSIONAL EXPERIENCE

Research Assistant Richard Stockton College of New Jersey
Neil Aaronson *Spring 2015*

The Activity Leaders of New Students Richard Stockton College of New Jersey
Student Mentor *Fall 2011 – Fall 2014*

Tutoring Center Richard Stockton College of New Jersey
Math & Physics Tutor *Spring 2014*

Board of Trustees Richard Stockton College of New Jersey
Elected Student Trustee *Spring 2011 – Spring 2013*

Research Assistant Johns Hopkins University
Robert Ivkov *Summer 2012*

Student Senate Public Relations Chair Richard Stockton College of New Jersey
Committee Chair *Spring 2012 – Fall 2012*

Student Senate Representative Richard Stockton College of New Jersey
Public Relations Committee Member *Spring 2011 – Spring 2012*

International Field Day Planning Committee Richard Stockton College of New Jersey
Committee Member *Spring 2012*

ADDITIONAL ACHIEVEMENTS & AWARDS

Department Block Travel Grant	Fall 2018, Spring 2019
Graduate Student Association Travel Grant	Fall 2018
Outstanding Teaching Assistant Award	Spring 2017, Spring 2018
Gretler Research Fellowship	Summer 2017
Graduate Student Association Excellence in Teaching Award Nominee	Fall 2016-Spring 2017
Mercatus Center Adam Smith Fellowship	Fall 2016-Spring 2016
Regent's Fellowship	Fall 2015-Spring 2016
Block Grant Fellowship	Fall 2015-Spring 2016
Summa Cum Laude	Spring 2014
Distinction in Honors	Spring 2014
Program Distinction	Spring 2014
Provost Scholarship	Fall 2010-Spring 2014
Kappa Sigma Scholarship-Leadership Award	2013, 2014
Who's Who Among Students in American Colleges and Universities	2012, 2013
Leadership Education Awareness Development (L.E.A.D.) Certification	2013
Harold E. Taylor Physics Award	2013
Rising Young Physics Student Award	2012
Yitzhak Sharon Physics Book Award	2011

OTHER

Technical Expertise: Python, MatLab, R, Stata, HTML (basic)
Professional Membership: Sigma Pi Sigma, Kappa Sigma
Citizenship: United States, The Netherlands

REFERENCES

Professor Peter Rupert
Department of Economics
University of California, Santa Barbara
peter.rupert@ucsb.edu

Professor Emeritus Thomas Cooley
Stern School of Business
New York University
tcooley@stern.nyu.edu

Professor Douglas Steigerwald
Department of Economics
University of California, Santa Barbara
doug@ucsb.edu

Associate Professor Javier Birchenall
Department of Economics
University of California, Santa Barbara
javier.birchenall@ucsb.edu

Raymond Farrow, Esq.
Keller Rohrback L.L.P.
Seattle, WA
By Request Only

Permissions and Attributions

1. The content of Chapter 2, Appendix A, and Appendix B is the result of a collaboration with Christine Braun and Peter Rupert, and has previously appeared in the *Review of Economic Dynamics* (Braun et al., 2020). At the time of this writing, authors of articles published in the *Review of Economic Dynamics* maintain the right to include such articles in student dissertations circulated for non-commercial purposes. Moreover, all coauthors consent to its inclusion in this dissertation. The published article can be found at `https://doi.org/10.1016/j.red.2020.11.001`.

2. The content of Chapter 3 is the result of a collaboration with Thomas Cooley and Espen Henriksen. It is reproduced here with their consent.

3. All three chapters of this dissertation made extensive use of computational facilities purchased with funds from the National Science Foundation (CNS-1725797) and administered by the Center for Scientific Computing (CSC). The CSC is supported by the California NanoSystems Institute and the Materials Research Science and Engineering Center (MRSEC; NSF DMR 1720256) at UC Santa Barbara.

Abstract

The Long and Short of Labor Supply Changes

by

Edwin Charles Nusbaum IV

The study of the dynamics, causes, and consequences of changes in labor supply is central to understanding modern economies and identifying candidate policies to improve welfare. Each chapter of my dissertation contributes to one or more components of this broad theme by combining applied econometrics techniques with insights from quantitative theoretical models. Using these tools, I aim to address three questions: How do aggregate hours worked change over the business cycle? Can the rise in female labor force participation account for household migration trends? What role do individual labor supply choices play in driving aggregate growth as populations age?

In Chapter 1 of my dissertation, I study the finite sample properties of a novel approach to identifying macroeconomic shocks with long-run restrictions. In contrast to past studies, this approach constructs and constrains the long-run impact of shocks directly using local projections rather than inferring them from vector autoregressions. Through a series of Monte Carlo simulations, I show that the local projections approach can yield substantial reductions in both bias and mean squared error, while also boasting decreased sensitivity to the choice of included lag length and assumed order of integration of the endogenous variables. I then use data from the Bureau of Labor Statistics to revisit a long standing debate on the response of aggregate hours worked to positive productivity shocks. I find that labor hours rise in response to positive productivity shocks and follow a hump-shaped profile thereafter. This result is robust to a number of specification choices and provides new evidence in support of the standard real business cycle model.

My joint work with Christine Braun and Peter Rupert constitutes Chapter 2 of my dissertation, and studies the relationship between the historical rise in female labor force participation and contemporaneous decline in household migration rates. Between 1964 and 2000, the inter-county migration rate of married couples declined by 15%. Concurrently, female labor force participation among married women and the relative wages of women increased by 39 and 14 percentage points, respectively. Using a two location household level search model of the labor market, we show that both the increase in dual earner households and the rise in women's wages contributed significantly to the decline in the migration rate of married households, with each explaining 53% and 20% of the decline, respectively. We further show that this co-location problem has important implications for structural models designed to estimate lifetime earnings inequality.

Finally, I conclude my dissertation in Chapter 3 with joint work with Thomas Cooley and Espen Henriksen wherein we study the growth effects of aging populations in Europe's four largest economies – France, Germany, Italy, and the United Kingdom. Since the early 1990's, GDP per-capita growth in these economies has slowed while at the same time a combination of longer individual life expectancy and declining fertility have led to gradually ageing populations. Using a general equilibrium overlapping generations model, we show that demographic change such as this affects economic growth directly through aggregate savings and labor supply decisions. These decisions are further affected indirectly through additional distortions caused by rising tax rates needed to fund pension systems. We find that the net effect of these forces can account for a significant fraction of the historical growth slowdown and that evolving demographics will continue to drag down growth over the next 20 years. We highlight that the degree to which gains to life expectancy change labor supply decisions is the most important margin through which demographic change affects growth by studying several reforms aimed at increasing late-life labor supply.

Contents

Chapter 1

Aggregate Hours and Local Projections with Long-Run Restrictions

1.1 Introduction

The estimation of impulse response functions is central to understanding the impact of shocks on the macroeconomy. Macroeconomists have largely relied on estimating vector autoregressions (VARs) and imposing the minimum number of additional identifying restrictions to interpret them as structural. One set of identifying restrictions used with VARs are long-run restrictions, wherein practitioners restrict the long-run impact of shocks within the model. For example, the long-standing debate surrounding the response of hours to productivity shocks spurred by Galí (1999) has largely relied on evidence from structural VARs (SVARs) identified by assuming that demand shocks have no long-run effect on productivity growth.[1]

Long-run identification schemes such as these are not robust to two central specification choices: assumed order of integration of the endogenous variables and included lag length. In

[1] Work by Christiano et al. (2003), Christiano et al. (2004), Francis and Ramey (2005), Galí and Rabanal (2005), Fernald (2007), Canova et al. (2010), and Saijo (2019) have all used variants of this approach to show that hours may fall on impact in response to positive productivity shocks.

this paper, I use a novel approach that relies on local projections to identify structural shocks and investigate its finite sample properties. Using Monte Carlo evidence, I show that this approach is significantly more robust in terms of bias to the choice of lag length and order of integration of the endogenous variables than SVARs identified with long-run restrictions. I then provide new evidence that, consistent with Real Business Cycle (RBC) models and in contrast to much of the SVAR evidence, labor hours rise in response to positive productivity shocks.

Identification with long-run restrictions follow the work of Shapiro and Watson (1988), Blanchard and Quah (1989), and King et al. (1991). In this framework, a finite lag VAR is first estimated and the sums of the implied moving average coefficients are subsequently constrained to recover the structural parameters of the model. Cooley and Dwyer (1998) , Ravenna (2007), and Chari et al. (2008) have raised doubts regarding the validity of this approach. Together, they show that an inappropriate choice of lag length can severely bias impulse response functions estimated in this way. An inappropriate lag length not only biases the estimated autoregression (AR) coefficients due to omitted variables, but more importantly ignores terms in the moving average coefficients of the true data generating process.

The identification procedure herein considered, hereafter termed long-run local projections (LRLPs) extends the local projections studied in Jordà (2005) to achieve structural identification. Local projections regress an endogenous variable on lags of itself and other endogenous variables independently for each forecast horizon. In this way, local projections estimate the moving average coefficients of the underlying data generating process directly rather than relying on recursive substitution as in VARs. As a result, the long-run impact of each shock may be constructed and constrained without excluding terms in the moving average representation of the underlying data generating process. Moreover, omitted variable bias is not passed between moving average coefficients through recursive substitution. Because of these features, structural identification using LRLPs may outperform SVARs of similar lag length regardless of the order of integration of the endogenous variables. To the best of my knowledge, I am the first to implement and study the finite sample properties of long-run local projections.

I rely on a two-shock RBC model developed in Chari et al. (2007) and Chari et al. (2008) (hereafter CKM) as the data generating process. This model has several advantages. Its linearized form not only has a VAR(∞) representation for a wide range of parameter values, but its structural parameters can also be identified with long-run restrictions. Moreover, the CKM model has been used to evaluate many empirical models, including SVARs, in McGrattan (2010), Kascha and Mertens (2009), and CKM. For each simulated data series, I estimate the impulse response function of labor hours to a productivity shock using the standard SVAR estimator and LRLPs. I find that the LRLPs yield significant bias reductions relative to the SVAR both in estimating the full impulse response function and the contemporaneous response. For example, LRLPs reduce the bias of the contemporaneous response by 73% and 47% when labor hours are included in first differences and levels, respectively. LRLPs correctly estimates the direction of the contemporaneous response and the shape of the impulse response function in all cases, and can eliminate all of the biases for some specification choices.

I then isolate the small sample bias of the LRLPs from other specification choices by increasing the length of each simulated data series. Similar to the findings of Erceg et al. (2005) in the context of SVARs, I find that the effect of small sample sizes on the structural parameter estimates depends on the empirical specification used.[2] Because the estimated AR suffer from the well known small sample bias first described in Hurwicz (1950), their sum and therefore the estimated structural parameters are also biased in small samples. The bias in the structural parameters, however, results from a non-linear transformation of that in the reduced form coefficients. As a result, the structural estimates may be biased upward even when the reduced form estimates are biased downward. This issue is particularly important at long forecast horizons, but quantitatively small at short forecast horizons. Despite these sensitivities,

[2] Faust and Leeper (1997) make a related point when constructing confidence bands and hypothesis testing after imposing long-run restrictions on estimated VARs.

the LRLPs outperform the SVAR in terms of impact error and integrated bias across all forecast horizons considered.

Having established the advantages of the LRLP method, I show that it has first order implications for existing empirical discussions that rely on SVARs identified with long-run restrictions by revisiting the debate on the response of hours to productivity shocks.[3] The overwhelming conclusion of this literature has been that labor hours fall in response to productivity shocks, a result that is at odds with RBC models à la Kydland and Prescott (1982) and King and Rebelo (1999). Using data from the Bureau of Labor Statistics (BLS) on non-farm private business labor productivity and labor hours spanning 1948Q1 to 2019Q2, LRLPs indicate that labor hours rise in response to a productivity shock and subsequently follow a hump-shaped profile. In contrast to Christiano et al. (2003) and Christiano et al. (2004) who come to a similar conclusion, my results are robust to whether or not labor hours are first differenced and the inclusion of breaks in labor productivity growth.[4] Moreover, my results provide new evidence in support of the standard RBC model.

My findings also contribute to the growing literature assessing the ability of local projections to improve estimates of impulse response functions. Kilian and Kim (2011), for example, use data simulated from a VAR(12) to compare the coverage rates of impulse response functions estimated using local projections and VARs. They argue that the local projections method provides no apparent advantages. Brugnolini (2018) shows that their use of the Akaike (1974) information criterion in choosing the lag length rather than the Schwarz (1978) information criterion drives their results.[5] Both papers apply structural shocks estimated using a VAR with short-run restrictions to the local projections method. Choi and Chudik (2019) instead

[3] Notable exceptions in this literature are Basu et al. (2006) and Sims (2011), who do not rely exclusively on SVARs identified with long-run restrictions. They, however, come to differing conclusions.

[4] Garín et al. (2019) finds that hours decline on impact away from the zero lower bound, but rise when the zero lower bound is binding.

[5] The AIC prefers over-parameterized models. Brugnolini (2018) argues that the exercise presented in Kilian and Kim (2011) effectively asks the local projections method to outperform the true data generating process, which is not possible and therefore provides an inappropriate comparison between local projections and VARs.

tests local projections against several alternative estimation procedures when the sequence of structural shocks is perfectly observed rather than estimated. The closest paper to this study is independent and forthcoming work by Plagborg-Møller and Wolf (forthcoming), who prove that impulse response functions estimated using local projections and VARs are equivalent in population. Because of this result, they suggest that structural identification schemes relying on local projections succeed if and only if the previously described SVAR approach succeed. I compliment the findings of these studies in several important and distinct ways. First, I identify structural shocks using local projections directly rather than relying on externally constructed instruments. Second, I focus squarely on the finite sample properties of structurally identified local projections, of which little is known. Perhaps most importantly, I show that structural identification through LRLPs and the previously described SVAR approach can differ not only quantitatively, but also qualitatively in sample sizes of most interest to empirical researchers. Finally, my empirical application sheds new light on a long-standing debate in macroeconomics.

The rest of the paper is organized as follows. Section 1.2 reviews long-run restrictions in the context of VARs to illustrate the crucial source of bias and describes the local projections alternative. Section 1.3 details both the assumed data generating process and the Monte Carlo approach used to test the long-run local projections estimator. Section 1.4 describes the results of the Monte Carlo exercise. Section 1.5 uses long-run local projections to study the response of aggregate hours to a productivity shock. Section 1.6 summarizes my findings.

1.2 Imposing Long-Run Restrictions

In this section, I begin by presenting a simple structural VAR, highlighting the need to impose additional identifying restrictions. I then review how to recover the structural shocks by imposing long-run restrictions to an estimated VAR and describe why this method is so sensitive to the omission of relevant lags. Finally, I present the long-run local projections approach. I focus on the case with only a single lag throughout to simplify notation and facilitate exposition. The discussion below can be generalized to the case of an arbitrary number of lags.

1.2.1 The Problem with SVARs

The goal of imposing sufficiently many identifying restrictions to an estimated VAR is to orthogonalize the forecast errors, thereby recovering the sequence of structural shocks. Given that the reduced form parameters have been estimated, recovering the contemporaneous correlation matrix of the endogenous variables is sufficient to accomplish this goal. Estimating this matrix through long-run restrictions requires imposing assumptions on the estimated long-run impact of the structural shocks and subsequently inverting the reduced form parameters.

Consider, for example an n-variable structural VAR(1) given by

$$Bx_t = \Gamma_0 + \Gamma_1 x_{t-1} + \varepsilon_t \tag{1.1}$$

where B is the contemporaneous correlation matrix and subsumes the variance of ε_t. B captures both the indirect effect of $\varepsilon_{i,t}$ on $x_{i,t}$ through the other endogenous variables as well as the contemporaneous effect of $\varepsilon_{-i,t}$ on $x_{i,t}$. In this way, the contemporaneous correlation matrix, B, orthogonalizes the shocks that drive the stochastic process, ε_t, so that they may be interpreted as structural (causal) shocks with $\mathbb{E}(\varepsilon_t\varepsilon_t') = I$, where I is the identity matrix. Because each endogenous variable is potentially a function of not only lagged variables, but also the contemporaneous values of each other endogenous variable, rearranging Eq. 1.1 and simply applying OLS methods will yield biased coefficients.

Instead, one must first estimate the reduced form of the structural VAR by inverting B and impose additional identifying restrictions to recover the contemporaneous correlation matrix.

The reduced form VAR implied by Eq. 1.1 is given by

$$x_t = \underbrace{B^{-1}\Gamma_0}_{A_0} + \underbrace{B^{-1}\Gamma_1}_{A_1} x_{t-1} + \underbrace{B^{-1}\varepsilon_t}_{e_t} \tag{1.2}$$

where A_0 and A_1 are transformations of the structural parameters and e_t are forecast errors such that $\mathbb{E}(e_t e_t') = \Omega$ is no longer the identity matrix. Because $\mathbb{E}(x_{t-1}\varepsilon_t) = \mathbb{E}(x_{t-1}e_t) = 0$, the parameters of Eq. 1.2 may be consistently estimated using OLS. Including the covariance matrix of e_t, this yields only $n + n^2 + \frac{n^2+n}{2}$ parameter estimates while the structural VAR is characterized by $n + 2n^2$ parameters.[6] Thus, $\frac{n^2-n}{2}$ restrictions on the long run impact of the structural shocks must be imposed to recover estimates of the structural parameters.

To obtain the long-run impact of a unit impulse of the structural shocks to the endogenous variables, it is instructive to re-write Eq. 1.2 in its moving average form. In particular, recursive substitution and the definition of e_t yields

$$x_t = B^{-1}\varepsilon_t + A_1 B^{-1}\varepsilon_{t-1} + A_1^2 B^{-1}\varepsilon_{t-2} + \ldots \tag{1.3}$$

Thus, the time t impact of a structural shock i periods prior is $A_1^i B^{-1}\varepsilon_{t-i}$. Moreover, the long-run impact of shocks, D, is simply the sum of each of time t impact of those same shocks, i.e. $D = \sum_{i=0}^{\infty} A_1^i B^{-1}$. Further assuming that the eigenvalues of A_1 are less than 1 in modulus, the long-run impact matrix can then be written as

$$D = (I - A_1)^{-1} B^{-1} \tag{1.4}$$

where $(I - A_1)^{-1} = I + A_1 + A_1^2 + \ldots$ are the reduced form moving average coefficients. As only A_1 in Eq. 1.4 can be estimated from Eq. 1.2, an additional equation to pin down D and

[6] Generally, an estimated VAR(p) will provide $n + np + \frac{n^2+n}{2}$ parameter estimates compared to $n + n^2 p + n^2$ parameters in the structural VAR(p).

recover B^{-1} is necessary. Given that $\mathbb{E}(\varepsilon_t \varepsilon_t') = I$, we have $\mathbb{E}(e_t e_t') = \Omega = B^{-1}B^{-1'}$ and so

$$DD' = (I - A_1)^{-1}\Omega(I - A_1')^{-1} \tag{1.5}$$

where A_1 and Ω can both be estimated from the reduced form VAR. Long-run restrictions can now be imposed on D such that Eq. 1.5 holds and A_1 and Ω satisfy Eq. 1.2. Given D, B^{-1} is then obtained by inverting the reduced form moving average coefficients, $(I - A_1)^{-1}$.

Taking a Cholesky decomposition of Eq. 1.4, for example, returns a triangular matrix for D and is equivalent to assuming that structural shocks have no long-run effect on the endogenous variables higher in the Cholesky ordering. By replacing A_1 and Ω with $\widehat{A}_1$ and $\widehat{\Omega}$, respectively, $\widehat{B}$ can be consistently estimated. Given estimates for the contemporaneous correlation matrix and reduced form moving average (lag) coefficients, it is straightforward to obtain the estimated s-step ahead impulse response to a structural shock of size d.

$$\widehat{IR}(s,d) = \widehat{A}_1^s \cdot \widehat{B^{-1}} \cdot d \tag{1.6}$$

As is evident by Eq. 1.3-Eq. 1.5, consistently estimating the moving average coefficients of the underlying data generating process is critical to consistently estimating B and therefore $IR(s,d)$. Lag length mis-specification not only results in inconsistent estimates of the lag coefficients in Eq. 1.2, but more importantly in missing terms in all but the first two moving average coefficients. Suppose, for example, that the true data generating process took the form of a VAR(2) given by $x_t = A_0 + A_1 x_{t-1} + A_2 x_{t-2} + e_t$ rather than the VAR(1) described above. The first two coefficients of the moving average representation of both the true and assumed data generating process are equal to I and A_1. The third coefficient of the moving average representation on the other hand is A_1^2 for a VAR(1) and $A_1^2 + A_2$ for a VAR(2). A similar discrepancy exists between the two for each coefficient in their respective moving average representation after the third. This discrepancy poses first order consequences when imposing long-run restrictions on a VAR with incorrect lag length, which is of particular concern in empirically relevant sample sizes wherein the included lag length is typically constrained. CKM, for example, show that estimated contemporaneous responses may be more than double that of the true response, and may be of the incorrect sign even after controlling for small sample bias.[7]

1.2.2 Identification with Local Projections

The structural VAR methodology developed in the previous section is indeed useful for capturing the dynamics in a high-dimensional system. Estimated VARs, however, are a linear global approximation for the underlying system. That is to say that the dynamics of the system are determined recursively from one-step ahead forecasts. The local projections method provides an alternative.

The local projections method first presented in Jordà (2005) allows for a more direct and flexible estimation of the impulse response function. Rather than relying on recursive forecasts, this method estimates the dynamic relationships of a system at each forecast horizon independently with a collection of regressions. The local projections form of Eq. 1.3 is given by

$$x_{t+s-1} = A_0^{(s)} + A_1^{(s)} x_{t-1} + u_{t+s-1} \qquad s = 1, 2, \ldots, s_{max} \tag{1.7}$$

where $A_0^{(s)}$ is an $n \times 1$ vector of constants, $A_1^{(s)}$ are matrices of coefficients for the lag dependent variable, s denotes the s-step ahead forecast, s_{max} is the maximum forecast horizon, and u_{t+s} is a mean zero forecast error. The forecast error contains information of all shocks from time t to

[7]This discussion puts aside omitted variable bias in the OLS step to better illustrate the central problem with the standard SVAR estimator. The reduced form coefficients of an estimated VAR(1) will of course also be biased due to the omitted second lag.

time $t+s-1$.[8] Jordá denotes the collection of equations given by Eq. 1.7 for $s = 1, 2, \ldots, s_{max}$ as *local projections* due to the fact that each coeffecient is estimated equation by equation for each forecast horizon. As a result, these local projections are a set of local approximations to the true data generating process rather than a single global approximation as is the case with VARs. Jordà and Kozicki (2011) prove that, for a given forecast horizon, the coefficients in Eq. 1.7 can be consistently estimated by simple OLS.[9]

Moreover, the local projections technique allows for a simple method to compute the impulse response functions. In particular, the set of reduced form impulse responses may be orthogonalized as in Eq. 1.6. The key distinction between the two is that rather than using recursive substitution to obtain $\widehat{A}_1^s$, A_1^s is estimated directly with $\widehat{A}_1^{(s)}$. A key question remains, however: how does one recover B? Current approaches to estimating structural impulse response functions using local projections simply apply a contemporaneous correlation matrix obtained elsewhere to the reduced form coefficients. For example, Jordà (2005) accomplishes this by recovering B using an estimated SVAR.[10] This approach, however, causes the shortcomings of SVARs to be passed on to the local projections through the inconsistent estimation of B described previously.[11]

Rather than obtaining B by imposing long-run restrictions on VARs, the local projections can instead be used directly. In particular, recall that a crucial source of bias in estimating B using VARs is missing terms in each moving average coefficient as a result of lag misspecification. Because the local projections method does not rely s-steps of recursive substitution, each projection attempts to estimate its respective moving average coefficient directly. To see this, again consider the set of local projections given by Eq. 1.7 and substitute out x_{t-1} using the

[8]It can be shown that the forecast errors given by u_{t+s-1} are a moving average of the reduced form residuals, $\{e_{t+i}\}_{i=0}^{s-1}$, when the true data generating process is a VAR.

[9]Their result effectively requires that the conditions described in Lewis and Reinsel (1985) for the consistency of VARs hold for each forecast horizon, s.

[10]This is similarly done in the critiques by Kilian and Kim (2011) and Brugnolini (2018).

[11]If the sequence of structural shocks is known (e.g. Hamilton, 1985; Romer and Romer, 2004; Ramey, 2011), then they may instead be included as regressors themselves.

$s = 1$ local projection *only once*.

$$\begin{cases} x_t &= A_0^1 + A_1^{(1)} A_1^{(1)} x_{t-2} + A_1^{(1)} u_{t-1} + u_t \\ x_{t+1} &= A_0^2 + A_1^{(2)} A_1^{(1)} x_{t-2} + A_1^{(2)} u_{t-1} + u_{t+1} \\ x_{t+2} &= A_0^3 + A_1^{(3)} A_1^{(1)} x_{t-2} + A_1^{(3)} u_{t-1} + u_{t+2} \\ \vdots \end{cases} \tag{1.8}$$

Noting several features of Eq. 1.8. First, note that the structural shocks, ε_t, drive the stochastic process and that x_{t-2} summarizes all past shocks by assumption. Further note that u_{t+s-1} summarizes information on $\{e_{t+i}\}_{i=0}^{s-1}$. Then, the only information contained in u_{t-1} is that of the time $t-1$ structural shocks, ε_{t-1}. Thus, $A_1^{(1)}$ is the $t-1$ moving average coefficient, $A_1^{(2)}$ is the $t-2$ moving average coefficient, and so on. The time t moving average coefficient is trivially given by the identity matrix. Given this collection of local projections, the moving average coefficients used construct Eq. 1.5, A_1^i, can simply be replaced by the local projections coefficients, $A_1^{(s)}$. The term $(I - A_1)^{-1}$ in Eq. 1.5 is then given by

$$(I - A_1)^{-1} = \begin{bmatrix} 1 + \sum_s a_{11}^{(s)} & \sum_s a_{12}^{(s)} & \dots & \sum_s a_{1n}^{(s)} \\ \sum_s a_{21}^{(s)} & 1 + \sum_s a_{22}^{(s)} & \dots & \sum_s a_{2n}^{(s)} \\ \vdots & \vdots & \ddots & \vdots \\ \sum_s a_{n1}^{(s)} & \sum_s a_{n2}^{(s)} & \dots & 1 + \sum_s a_{nn}^{(s)} \end{bmatrix} \tag{1.9}$$

The estimation of B then proceeds as before after substituting $\widehat{a}_{ij}$ for a_{ij} and $\widehat{\Omega}$ for Ω, where $\widehat{\Omega}$ obtained from an estimated VAR and the $s = 1$ local projection are equivalent.[12,13]

[12] Christiano et al. (2006) develop a non-parametric method to estimate Ω that may further improve the long-run local projections approach. As shown in Section 1.4, however, this estimator yields significant bias reductions even without relying on the methods developed in Christiano et al. (2006).

[13] As in my discussion of VARs, I have put aside the issue of omitted variable bias. The reduced form coefficients of the local projections will of course be biased due to omitted lags. This issue is present in both long-run local projections and the standard SVAR approach.

This alternative estimator has several distinct advantages. First, any omitted variable bias in the estimated lag coefficients from including too few lagged dependent variables is not compounded across forecast horizons as a result of s-steps of recursive substitution. Second, and most importantly, because the local projections method provides a collection of local approximations rather than a single global approximation as in VARs, summing across the estimated lag coefficients as described above does not result in missing terms in the implied moving average coefficients as is the case with improperly specified VARs. Instead, $A_1^{(s)}$ estimates each moving average coefficient directly with the obvious normalization that the first, i.e. $s = 0$, moving average coefficient is not estimated and is instead given by the identity matrix. This fact may not be obvious at first glance as the long-run local projections estimator still relies on recursive substitution. The key distinction is that the local projections estimator relies only on one-step of recursive substitution. That is, the recursive substitution used in the long-run local projections relies only on the fact that the first moving average coefficient for a VAR of arbitrary lag length is the identity matrix. Thus, relying on one step of recursive substitution does not reintroduce the same issues present in the standard SVAR estimator.

It is immediately evident, that the long-run local projections relies on an appropriate choice of maximum forecast horizon, s_{max}, to construct the sums Eq. 1.9. It is also immediately evident that long-run local projections do not simply trade one problem for another. Increasing the accuracy of impulse response functions estimated using an n-variable VAR with long-run restrictions requires including more lags. Each additional lag results in a degrees of freedom reduction of n. Instead, improvements to local projections with long-run restrictions can be made by increasing the maximum forecast horizon rather than increasing the lag length of each regression. This results in a degrees of freedom reduction of only one. Thus, the long-run local projections method for imposing long-run restrictions is less limited by the length of data series available.[14] I return to this issue in Section 1.4 where I show that empirically relevant

[14]It is well known that properly specified VARs are more efficient than their local projection counterparts and

sample sizes allow for a sufficiently long forecast horizon to improve the estimation of both contemporaneous responses and the subsequent path of the impulse response function using long-run local projections.

1.3 Testing Local Projections with Long-Run Restrictions

To test the properties of the LRLPs against that of the standard SVAR approach, I rely on a two-shock version of a Real Business Cycle model developed in Chari et al. (2007) and CKM as the data generating processes.[15] This model has the distinct advantage that its linearized form satisfies the invertibility conditions described in Fernández-Villaverde et al. (2007) and Lippi and Reichlin (1994) that permit a VAR(∞) representation for a wide range of parameter values. Moreover, Chari et al. (2007) show that including stochastic wedges into a canonical RBC model such as this describes post-war aggregates well. This allows me to test the claim that local projections with long-run restrictions (LRLPs) are more robust than SVARs to specification choices. I begin by simulating the model for 285 quarters 1,000 times and estimating the impulse response of labor to a productivity shock using both an SVAR and LRLPs for each simulated series.[16] I adjust the lag length of both the estimated LRLPs and SVAR, and the maximum forecast horizon of the local projections method to investigate its advantages in empirically relevant samples. I then repeat this exercise using a sample length of 10,000 quarters to investigate the effects of small sample bias for the LRLPs. Throughout, I focus on the response of hours to a productivity shock as this is the statistic most used in the literature to choose between business cycle models. Moreover, this is the statistic I estimate in Section 1.5 where I revisit the hours debate.

so long-run restrictions imposed with VARs dominate asymptotically.

[15] I present only the key equations. Refer to Chari et al. (2007) and CKM for more details and proofs.

[16] The data series used in Section 1.5 is approximately 285 quarters in length.

1.3.1 Two-Shock RBC Model

A unit mass of infinitely lived households maximize expected utility by making a consumption-savings decision and a labor supply decision in frictionless markets, and discount the future at rate β. Households are subject to stochastic labor taxes that are rebated lump-sum in each period. Furthermore, I assume that household preferences are additively separable within and across periods, and are of the CES form. The problem of the representative household is then given by,

$$\max \ \mathbb{E}_{t_0} \sum_{t=t_0}^{\infty} [\beta(1+\gamma)]^t \left[\log c_t + \phi \log(1-l_t)\right] \tag{1.10}$$

subject to

$$c_t + (1+\gamma)k_{t+1} - (1-\delta)k_t = (1-\tau_{lt})w_t l_t + r_t k_t + T_t \tag{1.11}$$

where all variables are in per-capita terms. Here, δ is depreciation, γ is population growth, k_t is per-capita capital stock, w_t is the wage rate, r_t is the return on capital, T_t are lump sum rebates of the stochastic labor wedge, τ_{lt}, and c_t and l_t are per-capita consumption and labor, respectively.

Firms produce a numeraire consumption good using a Cobb-Douglas production function with labor augmenting technology, Z_t. Firm's maximize their profit function, $(k_t)^{\alpha}\,(Z_t l_t)^{1-\alpha} - r_t k_t - w_t l_t$, by choosing labor and capital input each period.

The stochastic variables, $\log Z_t$ and τ_{lt}, are subject to independent mean zero shocks each period given by $\log z_t$ and ε_{lt}, respectively. The covariance matrix of these shocks is therefore given by $\begin{bmatrix} \sigma_z^2 & 0 \\ 0 & \sigma_l^2 \end{bmatrix}$. I further assume that Z_t is log-normally distributed and follows a random walk with drift, μ_z, in logs. τ_{ls} is instead normally distributed and is defined by a stationary

AR(1) process in levels. The evolution of the stochastic variables is then characterized by

$$\log Z_t = \mu_z + \log Z_{t-1} + \log z_t \tag{1.12}$$

$$\tau_{lt} = (1-\rho_l)\mu_l + \rho_l \tau_{lt-1} + \varepsilon_{lt} \tag{1.13}$$

where ρ_l is the persistence of the stochastic labor wedge and $\log z_t$ and τ_{lt} are independent normally distributed shocks. With the assumption that $0 < \rho_l < 1$, innovations in the labor wedge are temporary. To close the model, the aggregate resource constraint is given by $c_t + (1+\gamma)k_{t+1} = y_t + (1-\delta)k_t$.

Log-linearizing about the steady state using standard methods, one may define a state-space system of the model. I assume that labor productivity growth, $\log\left(\frac{y_t}{l_t}\right)$, and the log of labor hours, $\log(l_t)$ are observable. The system is defined this way for several reasons. If more variables are included in the observer equation than shocks driving the process–in this case two– then a subset of the included variables will be a linear combination of the others, precluding a valid VAR(∞) representation of the model (Ingram et al., 1994; Ireland, 2004; Fernández-Villaverde et al., 2007). Moreover, defining the observer equation in this way allows for a VAR(∞) representation of the model that is consistent with the large literature estimating the aggregate response of hours worked mentioned previously–shocks to the stochastic labor wedge have no long-run effect on labor-productivity growth.

To make the model quantitative, I set $\phi = 1.6$, $\alpha = 0.33$, depreciation to be 6%, the rate of time preference to be 2%, population growth to be 1%, and the technology growth rate to be 2%. All parameters are set to match these annualized rates such that a period in the model is equivalent to one quarter. The parameters governing the stochastic variables are set so that $(\mu_z, \mu_l) = (0.005, 0.4)$, the persistent of the stochastic labor wedge is $\rho_l = 0.95$, and the standard deviation of each shock as $(\sigma_z, \sigma_l) = (0.0114, 0.00725)$.[17]

[17] CKM set the standard deviation of productivity shocks by dividing the standard deviation of TFP estimated

1.3.2 Econometric Specification

To investigate the properties of the proposed method, I estimate the model analogue of Eq. 1.7 given by

$$\begin{cases} \Delta\log\left(\frac{y_{t+s}}{l_{t+s}}\right) & = a^{(s)}_{0,\frac{y}{l}} + a^{(s)}_{1,\frac{y}{l}}\Delta\log\left(\frac{y_{t-1}}{l_{t-1}}\right) + a^{(s)}_{1,l}\log l_{t-1} + \dots \\ & \qquad + a^{(s)}_{p,\frac{y}{l}}\Delta\log\left(\frac{y_{t-p}}{l_{t-p}}\right) + a^{(s)}_{p,l}\log l_{t-p} + u^{(s)}_{\frac{y}{l},t+s} \\ \log l_{t+s} & = b^{(s)}_{0,l} + b^{(s)}_{1,\frac{y}{l}}\Delta\log\left(\frac{y_{t-1}}{l_{t-1}}\right) + b^{(s)}_{1,l}\log l_{t-1} + \dots \\ & \qquad + b^{(s)}_{p,\frac{y}{l}}\Delta\log\left(\frac{y_{t-p}}{l_{t-p}}\right) + b^{(s)}_{p,l}\log l_{t-p} + u^{(s)}_{l,t+s} \end{cases} \tag{1.14}$$

for $s = 0, 1, 2, \dots, s_{max}$. As described in Section 1.2, the two main advantages of imposing long-run restrictions using local projections are that each regression is a local approximation and so bias at any one forecast horizon is not compounded due to recursive substitution, the reduced form coefficients estimate the MA coefficients of the underlying data generating process directly, and that additional terms may be included in Eq. 1.9 by simply increasing the s_{max}. While I use a lag length of 4 and a maximum forecast horizon of 25 quarters as a benchmark, I adjust these two specification choices to illustrate the robustness of this method to the choice of lag length and investigate the sensitivity of LRLPs to the choice of s_{max}.

Moreover, I estimate a VAR(p) with labor-productivity growth ordered first to illustrate the advantages of local projections with long-run restrictions. The estimated model is given by

$$\begin{bmatrix} \Delta\log\left(\frac{y_t}{l_t}\right) \\ \log l_t \end{bmatrix} = \begin{bmatrix} A_{11}(L) & A_{12}(L) \\ A_{21}(L) & A_{22}(L) \end{bmatrix} \begin{bmatrix} \Delta\log\frac{y_{t-1}}{l_{t-1}} \\ \log l_{t-1} \end{bmatrix} + \begin{bmatrix} e_{1t} \\ e_{2t} \end{bmatrix} \tag{1.15}$$

where $A_{ij}(L)$ are lag polynomials of order p, and e_{1t} and e_{2t} are reduced form forecast errors. The length of the lag polynomial included in Eq. 1.15 is always consistent with that of Eq. 1.14.

While labor hours in the model are stationary by construction, Dickey-Fuller tests performed

in Prescott (1986) by labor share.

on the data used in Section 1.5 do not reject the null hypothesis of a unit root in hours. Moreover, Francis and Ramey (2005), Galí and Rabanal (2005), and CKM all argue that including labor in levels is inferior to an econometric specification including first differenced labor hours.[18] In light of these considerations, I also estimate the local projections and VAR using a differenced specification by replacing $\log l_t$ with $\Delta \log l_t = \log l_t - 0.99 \cdot \log l_{t-1}$ in Eq. 1.14 and Eq. 1.15.[19]

1.4 Monte Carlo Results

1.4.1 Small Sample Results

The benchmark results wherein 1,000 time series of 285 quarters are simulated and each estimator applied to all series are displayed below.[20] Figure 1.1 compares the mean estimated response of hours to a one standard deviation technology shock to the true response of the model. The results of the econometric specifications incorporating log-labor in differences are shown in Figure 1.1a and the results when including log-labor in levels are shown in Figure 1.1b. Throughout this section, the results from the benchmark local projections method and SVAR are shown in dark blue and red, respectively. The model implied response is always shown in black. Moreover, I present the estimated impulse response function to a forecast horizon of 50 quarters in order to illustrate both the full shape of the estimated impulse response function and how it compares across specification choices. In cases where $s_{max} < 50$, I estimate the local projections to a forecast horizon of 50 but use only the first s_{max} coefficients in the construction of the contemporaneous correlation matrix. I use the remaining $50 - s_{max}$ coefficients only to extend the estimated impulse response function to a forecast horizon of 50 for graphical purposes.

[18] Recent work by Saijo (2019) also uses a first differenced specification, but does not provide a clear justification of this choice.

[19] This avoids over differencing but is quantitatively equivalent to a first differenced specification asymptotically.

[20] The length of these time series are chosen to match those in Section 1.5.

Figure 1.1: Comparison to True Response in Small Samples

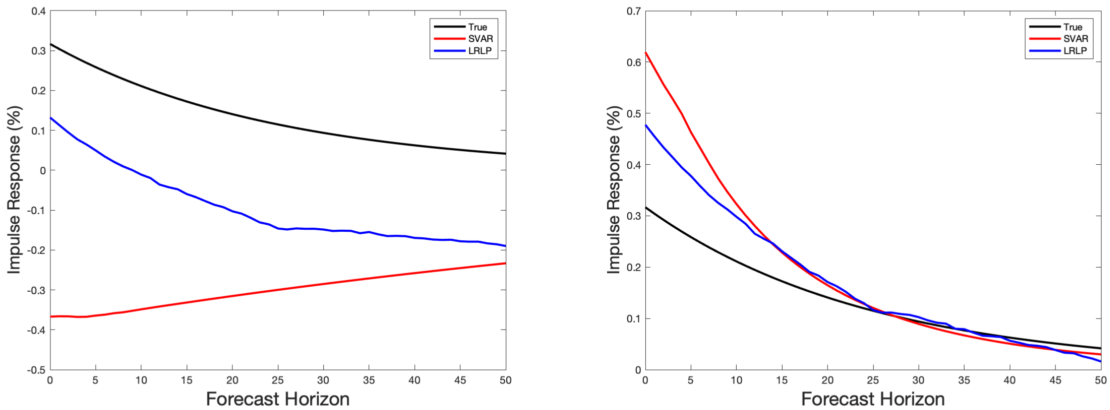

(a) Differences Specification

(b) Levels Specification

On impact, the model implies a 32 basis point rise in hours worked. Relative to the mean response estimated from the SVAR, the bias of the LRLPs is substantially reduced. In fact, the impact response of hours estimated using SVARs is 69 basis points lower and 30 basis points higher than the true response for the differences and levels specification, respectively. The mean impact response of the LRLPs with labor included in differences and in levels is instead 19 basis points lower and 16 basis points higher than the true response, respectively. This translates into bias reductions in the response on impact for differences and levels specification of 72% and 47%, respectively. Moreover, the LRLPs accurately represents the shape of the true response in both specifications whereas the shape of the impulse response function estimated using the SVAR changes drastically with the choice of hours used.

To quantitatively characterize the performance of each method in estimating the full impulse response function, I rely on a normalized integrated root mean squared error defined by

$$\rho\left(\widehat{IR}, IR, \overline{h}\right) = \left(\frac{\int_0^{\overline{h}} \left(\widehat{IR}(h) - IR(h)\right)^2 dh}{\int_0^{\overline{h}} IR(h)^2 dh} \right)^{\frac{1}{2}} \tag{1.16}$$

where $\widehat{IR}(h)$ is the estimated impulse response at horizon h, $IR(h)$ is the true impulse response at time h, and $\overline{h}$ is the maximum forecast horizon of interest. This metric is both robust to scale and punishes deviations symmetrically. Figure 1.2 and Figure 1.3 shows this metric as a function of $\overline{h}$ for the benchmark estimates of the differences and levels specification. The former shows $\rho\left(\mathbb{E}\widehat{IR}, IR, \overline{h}\right)$ and so identifies squarely the integrated bias of each estimator. The latter instead displays $\mathbb{E}\rho\left(\widehat{IR}, IR, \overline{h}\right)$ and so captures the bias-variance tradeoff of each estimator.

Figure 1.2: Integrated Bias

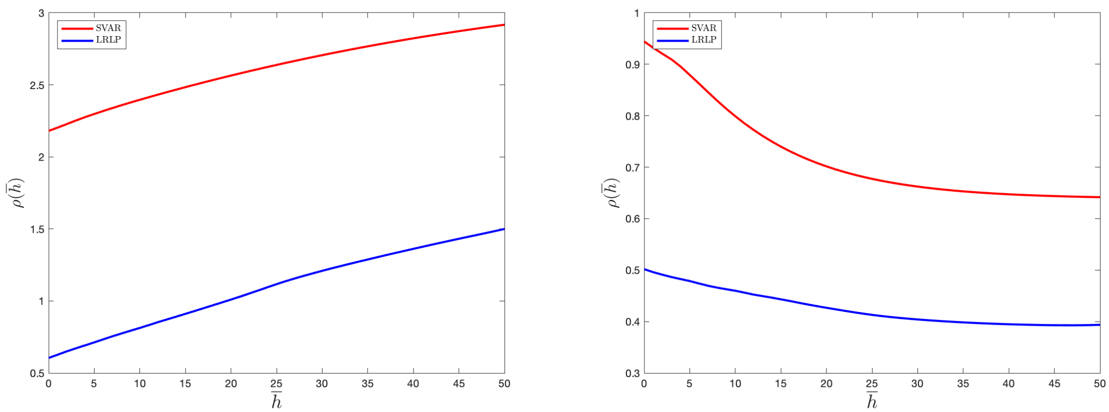

(a) Labor in Differences: $\rho\left(\mathbb{E}\widehat{IR}, IR, \bar{h}\right)$ (b) Labor in Levels: $\rho\left(\mathbb{E}\widehat{IR}, IR, \bar{h}\right)$

In both the levels and differences specification, the LRLPs outperform the SVAR method in terms of bias. Indeed, Figure 1.2a and Figure 1.2b show that integrated bias of the LRLPs is 50% lower than the SVAR method in the differences specification and 39% lower in the levels specification even after 48 quarters (12 years). Reductions in the integrated bias are even more substantial at shorter forecast horizons. Figure 1.3a further shows that the LRLPs outperforms the SVAR method in terms of mean squared error at all horizons considered in the differences specification. In the levels specification, the LRLPs outperform the SVAR in terms of mean squared error at forecast horizons less than 3 years, after which the SVAR performs better. This is a result of the reduced efficiency of local projections at long forecast horizons rather than the

SVAR overtaking LRLPs in terms of bias. Moreover, the reductions in efficiency are relatively small as $\mathbb{E}\rho$ is only 8% lower for the SVAR than the LRLPs even after 24 quarters (i.e. 6 years). Also notice that the mean squared error of the SVAR estimator asymptotes whereas the same is not necessarily true for the LRLPs. This is a direct result of the fact that the SVAR estimator relies on recursive substitution and the true response decays to 0 in the model. At horizons greater than the included lag length, the effect of the shock decays due to multiplication of the estimated AR coefficients in the reduced form VAR. Thus, the estimated effect on hours eventually returns to 0 regardless of bias in $\widehat{B}$ as in the model. Because this is not a feature of the LRLP estimator, the estimated impulse response of a shock does not necessarily decay to 0 at long forecast horizons.

Figure 1.3: Integrated Mean Squared Error

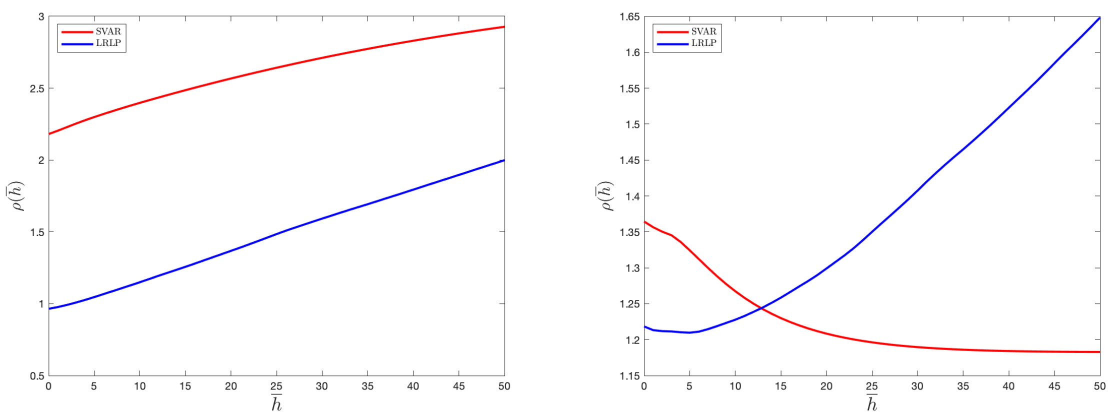

(a) Labor in Differences: $\mathbb{E}\rho\left(\widehat{IR}, IR, \bar{h}\right)$ (b) Labor in Levels: $\mathbb{E}\rho\left(\widehat{IR}, IR, \bar{h}\right)$

Finally, to further illustrate the relative performance of the LRLPs and the standard SVAR estimator, Table 1.1 and Table 1.2 presents the bias and mean squared error, rather than their integrated counterparts, at each forecast horizon. I normalize each by the value of the true impulse response function and the square of this value, respectively, to induce scale invariance. In addition, I highlight in grey rows in which the LRLP estimator yields the correct qualitative result but the SVAR does not. That is, I highlight cases when the LRLP estimator has the correct sign but the SVAR estimator does not. Cases in which the LRLPs do not yield reductions in the given statistic are indicated by “0.00%" in the fourth and seventh columns of each table. Evidently, the LRLP yields large bias reductions at virtually every forecast horizons regardless of the specification used. Moreover, the LRLP estimator yields mean squared error reductions at all forecast horizons considered for the first differences specification. For the levels specification, the LRLPs yield mean-squared error improvements throughout the first year and a half and subsequently do not outperform the SVAR estimator. As explained above, however, this is a result of increasing variance at long forecast horizons rather than SVARs outperforming LRLPs in terms of bias. Still, the LRLPs yield reductions in mean squared error at forecast horizons typically of most interest for both specification choices.

Table 1.1: Bias and MSE in First Differences

	Noramlized Bias			Normalized Mean-Squared Error		
Horizon	*LRLP*	*SVAR*	*% Reduction*	*LRLP*	*SVAR*	*% Reduction*
0	-0.583	-2.16	73.0	1.53	5.02	69.6
1	-0.627	-2.20	71.5	1.61	5.28	69.5
2	-0.676	-2.25	70.0	1.66	5.58	70.3
3	-0.724	-2.31	68.7	1.81	5.92	69.4
4	-0.761	-2.36	67.8	1.91	6.22	69.2
5	-0.808	-2.41	66.5	2.11	6.46	67.3
6	-0.861	-2.46	65.0	2.27	6.73	66.2
7	-0.911	-2.50	63.6	2.48	7.00	64.6
8	-0.957	-2.56	62.5	2.65	7.31	63.8
9	-0.997	-2.60	61.7	2.80	7.61	63.1
10	-1.05	-2.65	60.3	2.98	7.92	62.4
15	-1.34	-2.92	54.0	2.92	9.76	53.9
20	-1.73	-3.24	46.6	6.94	12.1	42.9
25	-2.27	-3.61	37.1	11.7	15.3	23.1

The improvements of the local projections method over the SVAR, of course, depend on the lag length used and maximum forecast horizon, s_{max}, included in the estimation of the contemporaneous response. I therefore re-estimate the impulse response function for each time series after varying the maximum forecast horizon, s_{max}, and lag length, p, included. Figure 1.4a and Figure 1.4b show the results of adjusting s_{max} in the differences and levels

Table 1.2: Bias and MSE in Levels

	Normalized Bias			Normalized Mean-Squared Error		
Horizon	*LRLP*	*SVAR*	*% Reduction*	*LRLP*	*SVAR*	*% Reduction*
0	0.506	0.955	46.8	2.24	2.61	14.3
1	0.496	0.932	46.8	2.23	2.59	13.7
2	0.483	0.904	46.5	2.20	2.53	13.1
3	0.476	0.883	46.1	2.23	2.55	12.4
4	0.464	0.853	45.6	2.19	2.48	11.9
5	0.461	0.791	41.7	2.22	2.32	4.23
6	0.444	0.743	40.3	2.21	2.21	0.00
7	0.427	0.687	37.8	2.26	2.08	0.00
8	0.420	0.628	33.0	2.34	1.96	0.00
9	0.422	0.577	26.9	2.44	1.86	0.00
10	0.412	0.530	22.2	2.51	1.77	0.00
15	0.340	0.326	0.00	3.11	1.48	0.00
20	0.214	0.170	0.00	4.07	1.34	0.00
25	0.013	0.048	73.7	6.09	1.29	0.00

specifications, respectively. The results from adjusting the included lag length are shown in Figure 1.5.

Figure 1.4: Sensitivity to Choice of s_{max}

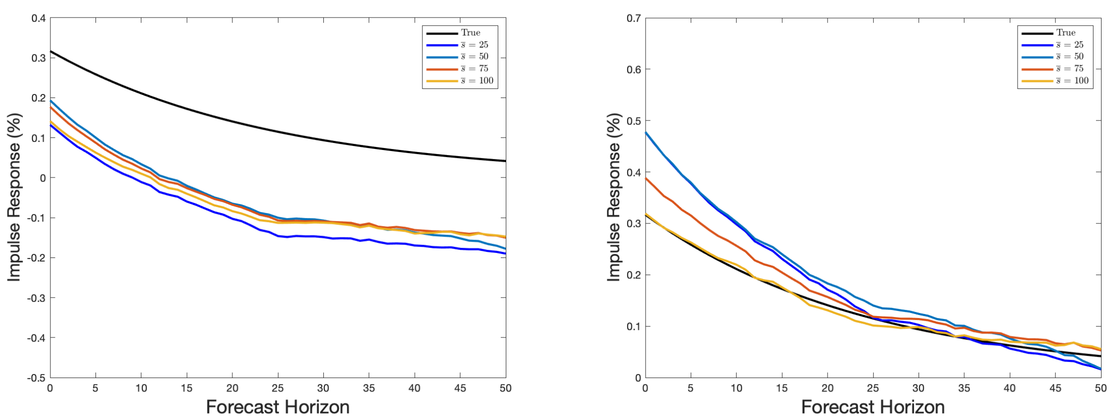

(a) Labor in First Differences (b) Labor in Levels

Several properties of the LRLPs are immediately evident. First, the LRLP estimator is sensitive to the maximum forecast horizon, s_{max}. This results from the fact that s_{max} determines where the summations in Eq. 1.9 are truncated and therefore the number of estimated moving average coefficients included in the construction of $\widehat{B}$. This sensitivity is greater in the levels specification than in the differences specification. In the latter, increasing s_{max} results in a mean estimated impulse response function that is quantitatively more similar to the true response. The bias in the estimated LRLPs with $s_{max} = 100$ is negligible at all but the longest forecast horizons. In the differences specification, the estimated response changes little with s_{max}.

Figure 1.5: Sensitivity to Choice of Lag Length

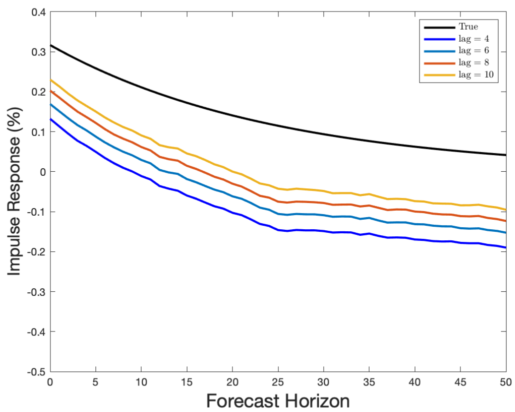

(a) Labor in Differences: LRLP Method

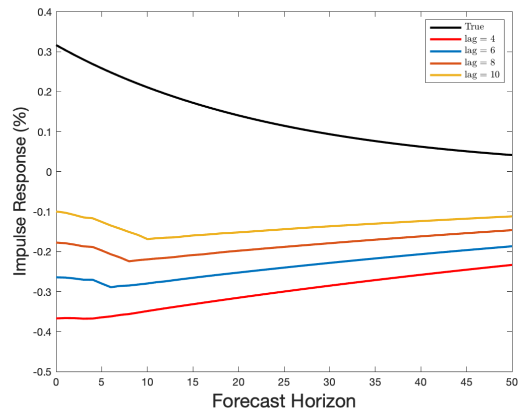

(b) Labor in Differences: SVAR Method

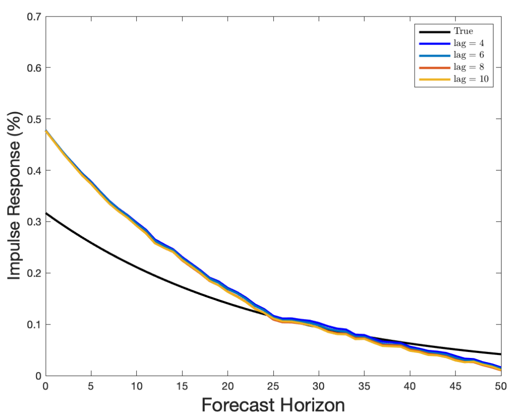

(c) Labor in Levels: LRLP Method

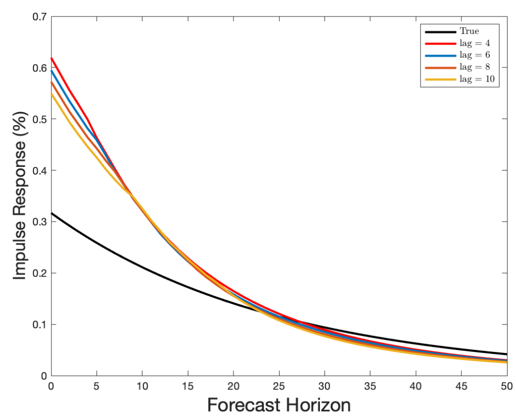

(d) Labor in Levels: SVAR Method

Mirroring the sensitivity of the LRLP estimator to the choice of s_{max}, the choice of included lag length appears to be less of a concern for the levels specification than for the differences specification. In fact, Figure 1.5c shows that the levels specification shows virtually no adjustment in the mean estimated response. Figure 1.5a instead shows that the differences specification does show improvements with the included lag length. Despite these sensitivities, however, it is notable that the shape of the LRLPs remains reflective of the true response regardless of the choice of s_{max} and included lag length. The shape of estimates from the SVAR on the other hand differ not only with the measure of hours used, but also with the choice of included lag length in the case of the differences specification.[21]

1.4.2 Removing Small Sample Bias

To investigate the relative performance of each method without contaminating the results with small sample bias, I repeat the above exercise with 1,000 simulated time series of 10,000 quarters in length. In these exercises, the LRLPs and SVARs suffer from downward bias AR coefficients to a quantitatively insignificant degree. Figure 1.6a to Figure 1.6b mirror Figure 1.1a and Figure 1.1b and show the results of the large sample exercise. I focus on the benchmark specifications as the alternative specifications follow the same qualitative trends as in the previous section.

The contemporaneous response of the LRLP method in differences and levels is 16 basis points and 45 basis points, respectively. Relative to the estimated contemporaneous responses in the small sample case, the former is shifted up by 3 basis points and the latter is shifted downwards by 3 basis points. The non-uniformity of these shifts is consistent with Erceg et al. (2005), who show that the way in which small sample bias of estimated AR coefficients translates to structural parameters estimated using long-run restrictions is dependent on the estimated econometric model. Said differently, downward biased AR coefficients do not necessarily translate into downward bias structural parameters when applying long-run restrictions. The contemporaneous response of the SVAR method is instead -36 basis points and 74 basis points,

[21] Bias reductions from increasing s_{max} and the included lag length are of course accompanied by reductions in efficiency for both estimators.

Figure 1.6: Comparison to True Response without Small Sample Bias

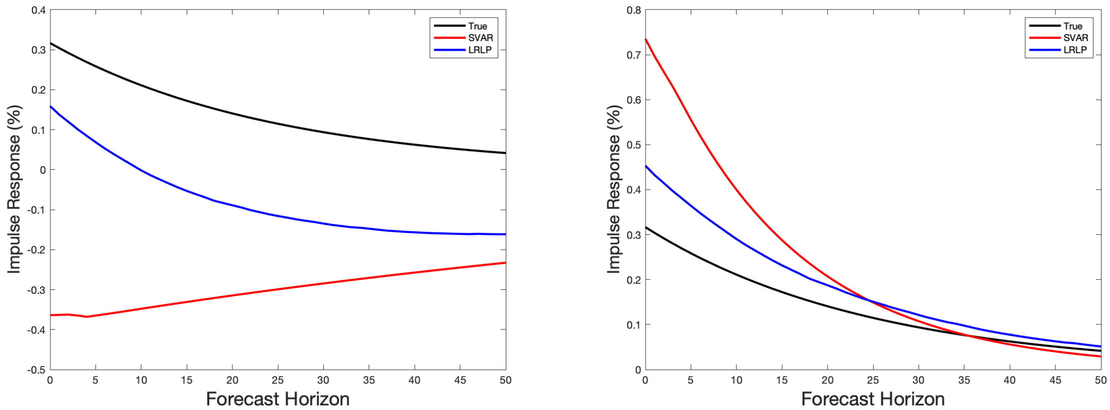

(a) Differences Specification (b) Levels Specification

respectively.

Figure 1.7: Integrated Mean Squared Error without Small Sample Bias

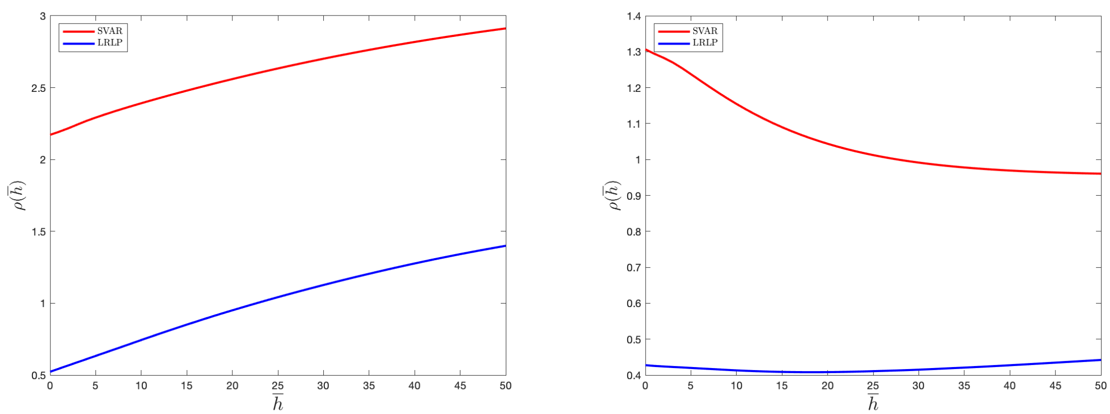

(a) Labor in Differences: $\mathbb{E}\rho\left(\widehat{IR}, IR, \bar{h}\right)$ (b) Labor in Levels: $\mathbb{E}\rho\left(\widehat{IR}, IR, \bar{h}\right)$

Figure 1.7 displays only $\mathbb{E}\rho\left(\widehat{IR}, IR, \bar{h}\right)$ as $\rho\left(\mathbb{E}\widehat{IR}, IR, \bar{h}\right)$ is quantitatively similar. Comparison of Figure 1.7b and Figure 1.3b highlights one apparent advantage of the SVAR method. Because the bias of the SVAR with labor in levels at long forecast horizons is small, its efficiency gains outweigh the bias reductions of the LRLPs at long horizons in this specification. In practice, however, the true bias of each estimator is unknown and the response of hours need not decay to 0 at long horizons. Furthermore, LRLPs outperform the SVAR in terms of mean squared error at forecast horizons of most interest to business cycle researchers (i.e. several years following a shock), and in terms of bias at all forecast horizons herein considered. Given these considerations, the superior bias properties of the LRLP method outwiegh the efficiency gains of SVARs.

1.4.3 Choosing s_{max}

Up until now, I have not discussed how to choose the free parameter, s_{max}, used in the construction of the contemporaneous correlation matrix. While it is possible to use a variety of information criteria to choose the lag length for autoregressive econometric models, the same is not true for the case of s_{max}.[22] Traditional information criteria choose a model parameter value, e.g. the lag length, to minimize some function of the model likelihood function with a penalty term for over-parameterization. Because these criteria depend on the likelihood function, they require the ability to compare observation values to those predicted by the estimated econometric model. In the case of choosing s_{max}, this would require comparing the contemporaneous correlation matrix implied by the estimated LRLP model to the "observed" contemporaneous correlation matrix. This object, however, is unobservable and is itself the subject of the estimator herein proposed. Thus, no comparison is feasible and standard information criteria are not applicable in this context.

Instead, recognize that increasing s_{max} implies a bias variance tradeoff. Figure 1.6 shows that the bias in the estimated contemporaneous correlation matrix decreases as s_{max} rises. Increasing the number of estimated parameter values used in Eq. 1.9, however, increases the uncertainty in their sum and therefore the implied contemporaneous correlation matrix. To balance this tradeoff, practitioners can therefore estimate the contemporaneous correlation

[22] Hacker and Hatemi-J (2008), Ozcicek and Mcmillin (1999), and Lütkepohl (1985) provide a review and comparisons of the selection criteria typically employed.

matrix for a range of s_{max} and choose the specification with the minimum s_{max} such that the estimated contemporaneous correlation matrix becomes both qualitatively and quantitatively similar to its behavior as s_{max} becomes large. Note that in large samples, an arbitrarily large s_{max} should be chosen as the bias-variance tradeoff is quantitatively negligible.

Figure 1.8 shows this procedure for the benchmark specification in the empirically relevant sample size discussed above and used in Section 1.5. The first panel shows the mean behavior across simulations for the contemporaneous response as a function of s_{max} when labor is included in differences. The second panel shows the same when labor is included in levels. When s_{max} is low, both specifications incorrectly suggest that hours decline on impact, a finding echoed by a large strand of the literature. As s_{max} increases, both specifications quickly estimate a contemporaneous response that is qualitatively similar to the model implied response, i.e. that hours initially rise in response to a positive productivity shock. In fact, the estimated response becomes positive after increasing the maximum forecast horizon to only 10-15 quarters depending on whether hours are included in levels or differences. In the differences specification, the contemporaneous response plateaus at a maximum forecast horizon of roughly 25 quarters. In the levels specification, the contemporaneous response also briefly plateaus at around 25 quarters and subsequently slowly declines to become quantitatively indistinguishable from the model implied response. Figure 1.8 suggests that restricting s_{max} to be between 20 and 30 quarters yields a good approximation to the model implied contemporaneous response both qualitatively and quantitatively. Moreover, a comparison of both specifications suggests that the true response–32 basis points–is between roughly +20 and +40 basis points. In contrast, comparison of the SVAR specifications yields a range of roughly -40 basis points to +60 basis points. Both of these facts serve to further highlight the improvements of the LRLP estimator over the traditional SVAR approach discussed in the previous two subsections.

Figure 1.8: Contemporaneous Response vs. s_{max}

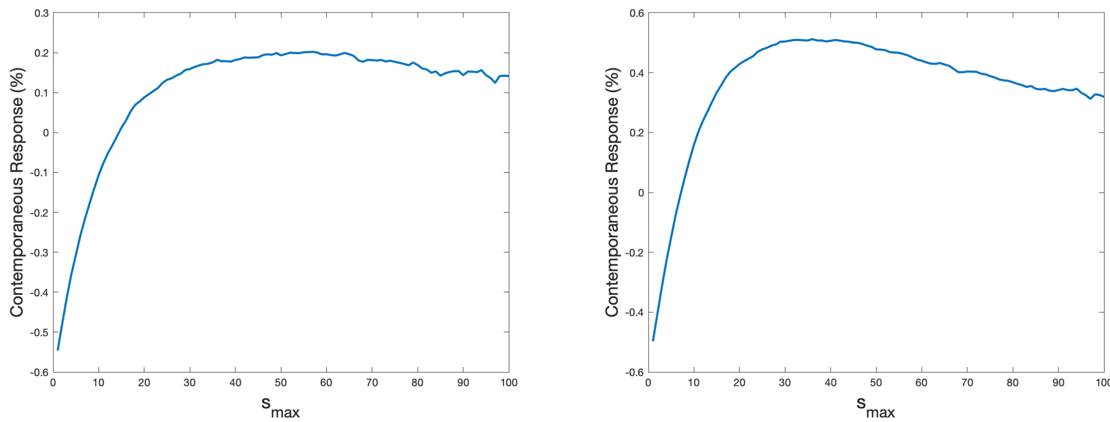

(a) Labor in Differences (b) Labor in Levels

1.5 Estimating the Response of Aggregate Hours

Having established the advantages of the LRLPs, I now revisit the large literature estimating the response of aggregate hours to a productivity shock. As previously discussed, Galí (1999), Christiano et al. (2003), Fernald (2007), Francis and Ramey (2005), Galí and Rabanal (2005), Canova et al. (2010), and Saijo (2019) all present estimates from a structural VAR.[23] Christiano et al. (2004), Basu et al. (2006), and Sims (2011) embed a constructed TFP series in a VAR directly rather than identifying TFP shocks using the estimated VAR itself and come to conflicting conclusions. With the exception of Christiano et al. (2003) and Christiano et al. (2004), the overwhelming conclusion of these papers is that aggregate hours decline–at least on impact–in response to an unanticipated rise in productivity.[24] This contrasts with RBC models à la Kydland and Prescott (1982), King and Rebelo (1999), and the extensions thereof, which predict a rise in hours in response to a positive technology shock. The direct, and often explicit,

[23] Saijo (2019) also presents estimates using the local projections method using the TFP series constructed by Fernald (2014).

[24] Sims (2011) finds that hours rise in response to temporary TFP shocks but fall in response to permanent TFP shocks.

implication of this literature is therefore that this class of models is unreliable.[25]

Because much of the evidence to date has been obtained by imposing long-run restrictions on estimated VARs, it is subject to the critiques previously discussed. I therefore, estimate the response of aggregate hours to a productivity shock using the LRLP estimator. I obtain output per hour (labor productivity) and labor hours data for the non-farm private business sector spanning the 1948Q1-2019Q2 period from the BLS. Labor hours are normalized by the non-institutionalized civilian population aged 16 and over in any given quarter.[26] Labor productivity is first differenced in all empirical specifications. There has been significant debate in the literature as to whether or not labor hours should enter in levels or in first differences with no clear consensus to date. While Dickey-Fuller tests fail to reject the null hypothesis that log-labor hours have a unit root, I do not take a stand on which approach is correct and instead present results when logged hours are included in both levels and first differences. Time series of the logarithm of the raw data are shown in Figure 1.9.

Figure 1.9: BLS Data

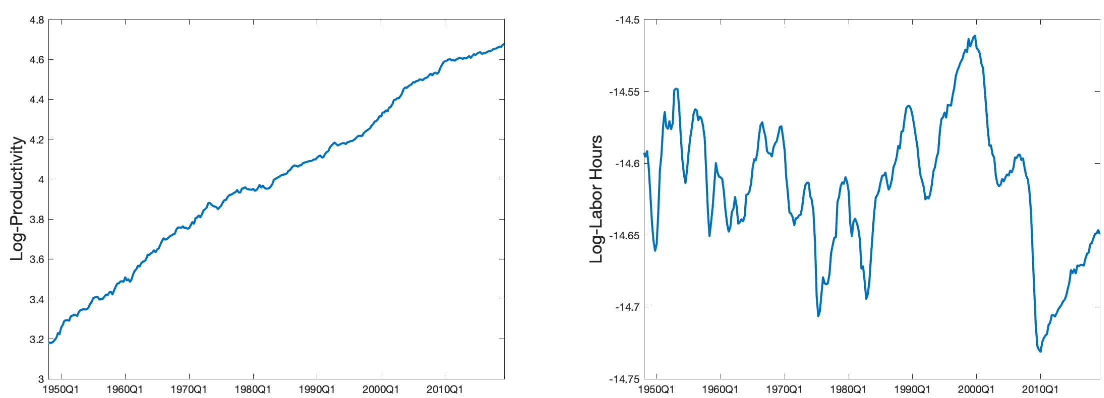

(a) Log-Labor Productivity (b) Log-Labor Hours per Capita

Furthermore, Fernald (2007) and Canova et al. (2010) show that controlling for low-

[25] See Kilian and Ltkepohl (2017) for a more detailed summary of the methodological components of this debate.
[26] The series ID of the BLS data used are PRS85006093, PRS85006033, and LNU00000000, respectively.

frequency changes in labor productivity growth has first order effects on estimated impulse response functions. To control for such considerations, I test for potential differences in subsample means of labor productivity growth in the following way. First, a Chow (1960) test is conducted on the original data for each potential break date until the first break date is found. A second Chow test is then performed on every possible break date after including the previous break dates as regressors. This process is continued until no new break dates are found. Once all break dates have been found, I de-mean labor productivity growth for each subsample (i.e. between break dates). Two break dates were found using this approach: 1973Q1 and 1974Q3.[27]

For comparison, I estimate the response of hours using both LRLPs and an SVAR. The lag length of the local projections is chosen to match that in the SVAR using the Akaike (1974) information criteria with a maximum possible lag length of 10. While some have argued that the the Schwarz (1978) information criteria outperforms alternatives (See e.g. Hacker and Hatemi-J, 2008; Lütkepohl, 1985), long-run restrictions are sensitive to lag truncation bias. As a result, the over parameterized specifications typically chosen by the Akaike (1974) information criteria may be preferable when relying on long-run restrictions for identification. Both information criteria yield the same result in this case. The imposed maximum possible lag length is not binding for any specification.

I construct 90% bootstrapped confidence for the SVARs using the bias-corrected bootstrap algorithm described in Kilian (1998). I use 1,000 replications to estimate the bias in the first step and 2,000 replications to estimate the confidence bands in the second step. An appropriate method for bootstrapping local projections is less clear. While Jordà (2009) discusses the importance of developing bootstrap methods for the local projections method, a blocks-of-blocks bootstrap approach for each forecast horizon, s, described in Kilian and Kim (2011) is the only

[27] Using the partial sample Wald statistic test suggested in Andrews (1993) and Andrews (2003) does not yield a statistically significant break date using the critical values therein presented, though 1972Q4 is close to significant for some presented critical values and virtually identical to first break date implied by the iterative Chow (1960) procedure. Imposing a break only in 1972Q4 does not qualitatively change my results.

proposal to date. Unlike this paper, their estimates of the contemporaneous correlation matrix are obtained from a VAR. In addition, this blocks-of-blocks bootstrap does not adequately consider the correlation between local projections coefficients across forecast horizons. Moreover, Kilian and Kim (2011) show through Monte Carlo simulation that their proposed bootstrap may yield less than nominal coverage. This issue is more or less severe depending on the data generating process used to generate the Monte Carlo samples. In light of these shortcomings, I instead bootstrap the long-run local projections by drawing from the asymptotic joint distribution of the reduced form coefficients and applying the same bias correction method as for the bootstrapped SVARs. I again use 1,000 replications to estimate the bias in a first step and 2,000 replications to estimated the confidence bands in a second step.

Finally, as described in Section 1.2, I must specify a choice of s_{max} to recover the structural parameters and therefore the impulse response function. As discussed previously, the choice of s_{max} directly affects the number of moving average coefficients included in the construction of the long-run impact of each shock. If s_{max} is too small, I may be missing important dynamic relationships of the data generating process. Thus, I follow the heuristic approach discussed in Section 1.4.3 and estimate the contemporaneous correlation matrix for a wide range of s_{max}. Figure 1.10 shows the estimated contemporaneous response for each specification.

The first column of Figure 1.10 shows the contemporaneous response when log-labor is included in levels. The second column shows the estimated contemporaneous response when labor is included in first differences. The overall relationship between s_{max} and the estimated contemporaneous response depicted above suggests that the contemporaneous response rises with s_{max}. The non-smoothness of this relationship results from the fact that the AR coefficients of the local projections method can be choppy relative to their VAR counterparts, a fact highlighted in Ramey (2016). Despite these fluctuations, the tendency of the contemporaneous response to rise with s_{max} is clear. In the levels specification, a maximum forecast horizon of 30 is chosen. A larger s_{max} is required for the differences specification. In this case, I

Figure 1.10: Sensitivity of Results to s_{max}

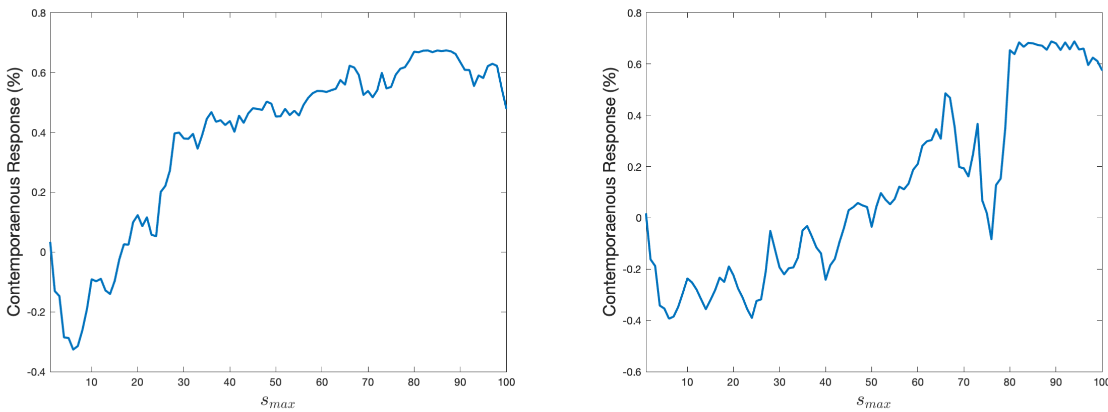

(a) Labor in Levels

(b) Labor in Differences

choose a maximum forecast horizon of 85, though a less conservative choice of 65 provides quantitatively similar results.

Figure 1.11 shows the estimated impulse response of hours to a technology shock. The first column shows results for the long-run local projections and SVAR in levels and the second column shows the results for the these two estimators when log-labor hours is included in first differences. The results from the LRLPs in Figure 1.11a and Figure 1.11b show that labor hours rise on impact in response to a positive technology shock and follow a hump-shaped response thereafter regardless of whether labor hours are included in levels or first differenced. The striking feature of the LRLPs is that, regardless of the chosen specification, the estimated response of hours is qualitatively the same. The SVAR method on the other hand predicts two qualitatively different responses depending on the specification used. The SVAR in levels, however, tightly matches both the level LRLPs and the differenced LRLPs. These results echo the findings of Christiano et al. (2003) that conclusions drawn from the SVAR with labor hours in levels should be preferred.

In all cases except for the SVAR in first differences, the bootstrapped confidence bands for the contemporaneous response contain 0. Hence, I cannot reject the hypothesis that hours do

Figure 1.11: Estimated Impulse Response of Hours to a Productivity Shock

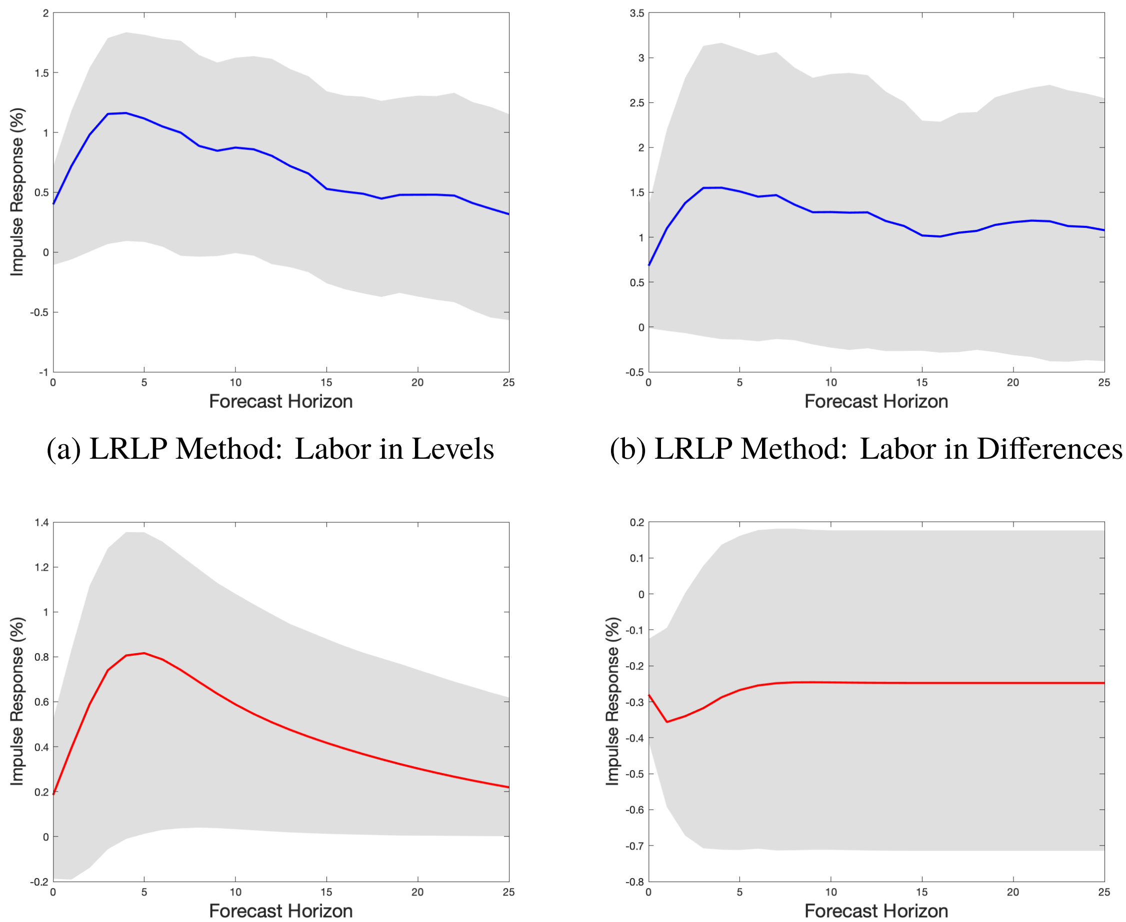

(a) LRLP Method: Labor in Levels

(b) LRLP Method: Labor in Differences

(c) SVAR Method: Labor in Levels

(d) SVAR Method: Labor in Differences

not respond, or even feature a very small negative response, to a positive productivity shock. Still there is substantially more probability mass to the right of zero than to the left of zero. Figure 1.12 shows the bootstrapped distribution of the contemporaneous response for the long-run local projections. For the long-run local projections in levels, approximately 91.8% of the probability mass lies above 0. For the long-run local projections with hours included in first differences, approximately 90.6% of the probability mass lies above 0. Taken together, the results of this section provide new evidence that hours in fact rise in response to a technology shock and that the standard RBC model may in fact be consistent with the data.

Figure 1.12: Bootstrapped Distribution of Contemporaneous Response

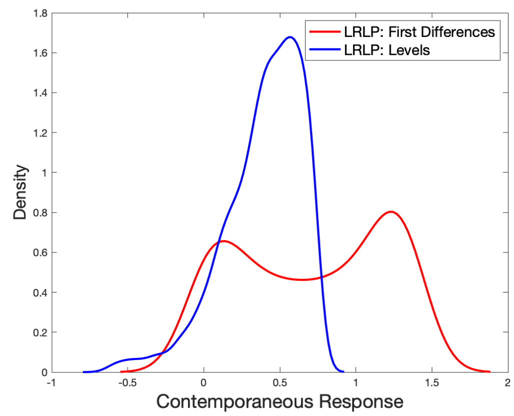

1.6 Conclusion

In this paper, I extend the local projections method to identify structural shocks through long-run restrictions. I show that identifying structural shocks using long-run local projections yields significant reductions in bias relative to SVARs both on impact and for most forecast horizons. Using Monte Carlo evidence, I show that the long-run local projections estimator is less sensitive than standard SVARs identified with long-run restrictions to the choice of included lag length and order of integration of the endogenous variables. Moreover, using this approach, I provide evidence that, in contrast to much of the evidence based on SVARs, aggregate labor hours rise in response to positive productivity shocks and follow a hump-shaped profile thereafter. In fact, I show that over 90% of the bootstrapped probability mass indicates that hours rise on impact in response to a positive productivity shock. This result provides new empirical support for the standard real business cycle model.

I also highlight several previously unexplored issues of importance for future research. First, the bias reductions of structural identification with local projections are illustrated only in the context of long-run restrictions. Structural identification in empirically relevant sample sizes using, for example, sign restrictions may also be improved upon by relying on local projections rather than estimated VARs. Additionally, current methods of bootstrapping with time series data either perform poorly for local projections or do not appropriately accommodate structural identification using local projections. Both of these issues require additional research.

Chapter 2

Labor Market Dynamics and the Migration Behavior of Married Couples

2.1 Introduction

Between 1964 and 2000 the increase in female labor force participation of married women has led to a more than doubling in the fraction of families with both spouses in the labor force, from 36% to 75%. Over the same period there has been a 30% increase in the female to male median wage ratio among married couples. How might such changes affect the willingness of married households to migrate across counties for new work opportunities?

In this paper we document that the intercounty migration rate between 1964 and 2000 of single households increased from 5.4% to 9.1% while the migration rate of married couples declined from 5.7% to 5.0%. After controlling for changes in other demographic characteristics such as age, education, and race, for example, we show that the downward trend for married couples persists whereas the upward migration trend for single households becomes flat. This suggests that the labor market experience of married women may be key in explaining the observed decline in couples' migration.

We estimate how much of the discrepancy in these migration patterns can be accounted for by the above mentioned forces. We find that, in total, the rise of dual labor force households can account for 55% of the decline in migration, whereas rising relative wages of women can account for 16% of the decline. These two mechanisms together can account for 65% of the total decline, implying negative complementarities between the two effects. Moreover, we find that approximately 10% of the rise in dual labor force households is induced by the rising job prospects of women. Accounting for this indirect effect of rising wages results in a wage effect that can account for 20% of the decline in migration. These results are both qualitatively and quantitatively similar in two extensions wherein we include an endogenous participation margin and exogenous, non-job related moves. Consistent with Cooke (2011) and Kaplan and Schulhofer-Wohl (2017) who suggest alternative mechanisms as the primary source of declines in mgiration–namely the Great Recession and increases in technology–for the post 2000 period, we show that female labor force participation among couples instead declined after 2000. Hence, our choice of focusing on the 1964-2000 time period for the effect on migration decisions.

In order to disentangle the composition effect, i.e. more dual searching households, from the wage effect, we construct a two location model with labor market frictions and allow both individuals to receive local and foreign offers while unemployed and employed, interpreting the acceptance of any foreign offer as a move to a new location. Once a move has taken place, only the spouse receiving the foreign offer remains employed. The interaction between mobility and on-the-job search has several implications on the reservation wages of individuals. First, if both spouses are employed, the foreign wage offers for which a household is willing to move is increasing in both spouses current wages, highlighting the fact that increased wages strengthens location ties. Second, if one spouse is employed, the local reservation wage of the unemployed spouse is everywhere decreasing in the employed spouse's wage. This results from increasing location ties as one spouse climbs the job ladder. That is, as the employed spouse's wage increases, both will be less likely to receive acceptable foreign offers, thereby making local offers more attractive and decreasing the reservation wage for the second spouse. The changing location ties of dual searching households, and in particular the strong location ties present when both spouses are employed, is the driving force of our mechanism.

We provide evidence of our mechanisms using household level data from the Current Population Survey (CPS) to test the effect of the joint decision process on migration. We find that households with both members in the labor force are 10% less likely to move. Furthermore, we show that among all households that moved, the relative probability that a dual searching household moves for job related reasons than for other reasons is 26% lower than for couples with only a single searcher. These results are consistent with work by Mincer (1978) and Costa and Kahn (2000) who show that the co-location problem faced by couples has a significant impact on migration decisions.

We also show that ignoring the co-location problem and these changing location ties has implications for estimates of lifetime earnings inequality. Ignoring this additional consideration of spouses within households can bias these estimates by as much as 20% for men, 36% for women, and 23% across all married individuals. This bias has increased for men and decreased for women as increasing wages for women cause fewer male-driven moves for high foreign wage offers and, therefore, more men to optimally enter the state of unemployment as a result of similar female-driven moves throughout their career. This results in dual searching men making choices that differ from their single searching counterparts to a greater degree, implying that the standard search model is no longer a good approximation for these men. The reverse is true for dual searching women.

Our model is similar to that developed in Guler et al. (2012). We add on-the-job search and gender specific offer distributions. Flabbi and Mabli (2018), Rendon and García-Pérez (2018), and Marcassa (2014) also extend the Guler et al. (2012) model to investigate the implications of dual searching households for estimates of lifetime earnings inequality, labor market policy reforms, and spousal unemployment duration, respectively. Karahan and Rhee (2017) argue that aging populations can explain almost half of the historical decline in aggregate migration. Kennan and Walker (2011) studies the effect of expected income on migration decisions. However, none of these papers address the dual-searcher household problem.

Molloy et al. (2017) argue that there is insufficient variation in the fraction of dual searching households to explain this long run decline, however, there is disagreement on this point. Cooke (2013), for example, instead shows that the fraction of dual worker households increased by roughly 15 percentage points from 1980-2010. We show empirically that the more than doubling of such households plays a major role in the migration decline in our sample. Taskin (2016) and Foged (2016) study how migration decisions of couples may change as relative wages change and find that migration is U-shaped in the wife's share of total family income. Neither, however, are able to decompose historical migration trends. Finally, a recent working paper by Guler and Taskin (2018) uses a similar model with marriage and divorce to investigate migration trends. Our paper differs in that we focus on married couples, as other demographic characteristics cannot explain the trend for this group. We allow search on and off the job to differ, and include an empirically representative mix of single and dual searching married couples.

The remainder of the paper proceeds as follows. Section 2.2 presents the CPS data used in our econometric analysis and highlights the key demographic trends underlying our mechanism. Section 2.3 outlines our model and derives the migration rates of dual and single searching households. Section 2.4 presents the results of our calibration and Section 2.5 conducts our quantitative experiment. Section 2.6 concludes.

2.2 Data

The goal of this paper is to understand how the increase in the number of households with both spouses in the labor force and the relative earnings of women have affected the mobility of married couples. We begin this section by documenting the decline in the migration rate of married couples and adjust the series for other demographic changes. Next we turn to the two channels we are interested in: the increase in dual searching households and the rise in the relative earnings of women. We show that both the fraction of dual searching households as well as the relative earnings of women increased rapidly until 2000.

2.2.1 Mobility

In this section we document migration rates for married couples, the group of interest in this paper, discussed below. Our migration data come from the March sample of the Current Population Survey (CPS) through IPUMS (Flood et al., 2018). The variable of interest is the one year mobility question in which respondents were asked if they had changed residence since March of the previous year. Movers are divided into five categories: those who had moved within the same county (intracounty); those who had crossed county lines but stayed in the same state (intercounty-intrastate); those who had resided in a different state (interstate); and those who had migrated from abroad. Throughout the remainder of the paper, we refer to total intercounty migration (the sum of intercounty - intrastate and interstate migration) simply as intercounty migration.

We restrict our sample to civilian households in which the head of the household is at least 16 years old. We define married households as households in which the head of household state they are married with or without their spouse present, the remaining households are labeled as not-married. Further we define single searcher households as married households in which one individual is in the labor force and the other is out of the labor force and dual searcher households as married households in which both individuals are in the labor force.

Figure 2.1: Intercounty Migration Rate by Marital Status

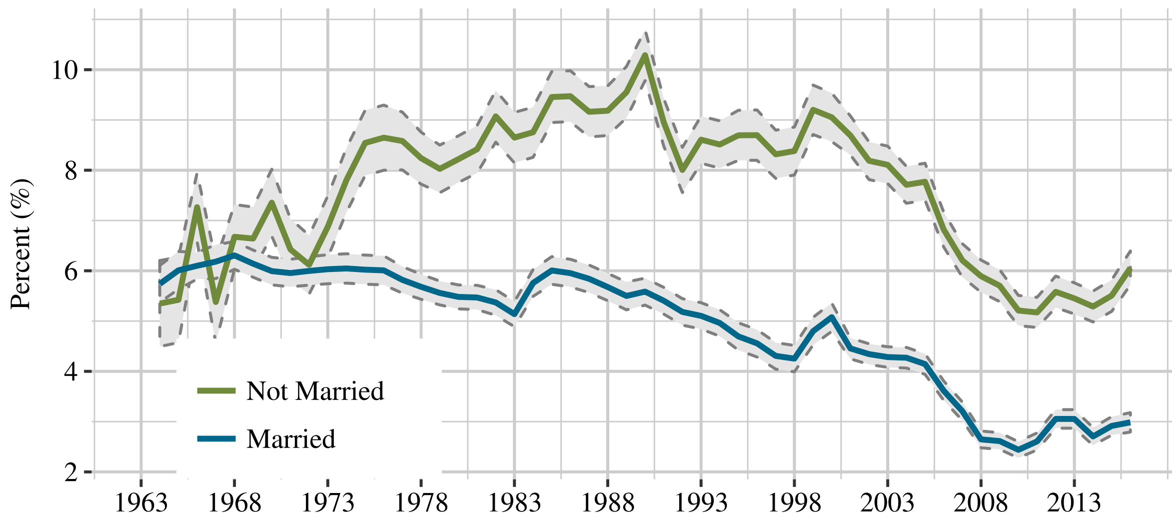

Notes: The 1-year geographic mobility question was not asked between 1972 to 1975, 1977 to 1980, 1985 and 1995 and a cubic spline was used. The figure shows the percent of civilian households that moved across county lines within the previous year by marital status. 95% confidence bands are plotted.

Figure 2.1 shows the intercounty migration rate of civilian households between 1964 and 2015 by marital status of the head of household. The figure reveals that the trends in intercounty migration are different substantially by marital status. While the percent of not-married movers increased from 5.4% in 1964 to 9.1% in 2000, the percent of married movers decreased from 5.7% to 5.0% over the same time period. After 2000, both groups experienced a rapid decrease in their migration rates. As we will show in the following subsections, the trends in dual searcher households began to level out after the late 1990's. Moreover, Kaplan and Schulhofer-Wohl (2012) shows that a change in the CPS imputation method in 2006 can explain almost half of the post 2000 decline and just over 60% of the sharp decline during the 2005-2006 period. As a result, our primary period of interest is 1964-2000 as our mechanism will be able to explain very little of the migration patterns following 2000.[1]

[1] CPS imputation flags are not available prior to 1995. Thus, we cannot drop imputed observations over the

Others have identified that the composition of households had changed over this time period (e.g. Iyigun and Lafortune (2016)). As we discuss in Section 2.3, our model does not consider age, race, and education levels of household members. To remove these potential confounds from our migration data, we adjust the data to control for such characteristics. Figure 2.2 shows the adjusted percent of households that moved across county lines within the year. The adjusted series controls for changes in the age, sex, race, education of the head of household, total real family income, and the number of family members living in the household by estimating Eq. 2.1 in each year and subsequently adding the estimated time dummies to the 1964 migration rates.

$$move_{i,t} = \beta_0 + \beta_1 X_i + \eta_t + \varepsilon_{i,t} \tag{2.1}$$

where X_i contains the vector of controls discussed above. The figure shows that holding constant the composition of these factors at their 1964 levels removes any discernible trend for single searching households until 2000. However, holding the demographic characteristics of married households fixed at their 1964 values flattens the migration trend only slightly. These results indicate that demographic characteristics go a long way in explaining the migration trends of not-married households, but not married couples. Moreover, they provide a middle ground between past studies linking demographics and migration and Molloy et al. (2011), who suggest that demographic characteristics are likely to be unimportant for migration choices.

Indeed, Table 2.1 shows the results of the following linear regression for the 1964-2000 adjusted sample:

$$mrate_{t,g} = \beta_0 + \beta_1 \mathbb{1}(g = Married) + \beta_2 t + \beta_3 \mathbb{1}(g = Married)t + \varepsilon_{t,g} \tag{2.2}$$

vast majority of our period of interest. As discussed in Kaplan and Schulhofer-Wohl (2012), our trends are likely robust to the imputation method used, so long as we use the same imputation method throughout our sample. We therefore also do not drop imputed values for 2000 in our main quantitative exericse.

where $mrate$ is the migration rate at time t for group g and g may take the values $Married$ or $Not\ Married$. The fact that the estimated coefficient on time, $\widehat{\beta}_2$, is not statistically significant suggests that demographic changes account for all of the trend in the migration rate of non married households. Moreover, the negative and significant coefficient for the married-time interaction term, $\widehat{\beta}_3$, indicates that a sizeable residual trend remains after controlling for changes in the demographic characteristics of married households.[2] Together, Figure 2.2 and Table 2.1 show that a residual trend only exists for married households. As a result, only a characteristic unique to married couples can explain this trend. One such aspect is the presence of a working spouse.

Figure 2.2: Adjusted and Unadjusted Intercounty Migration Rates

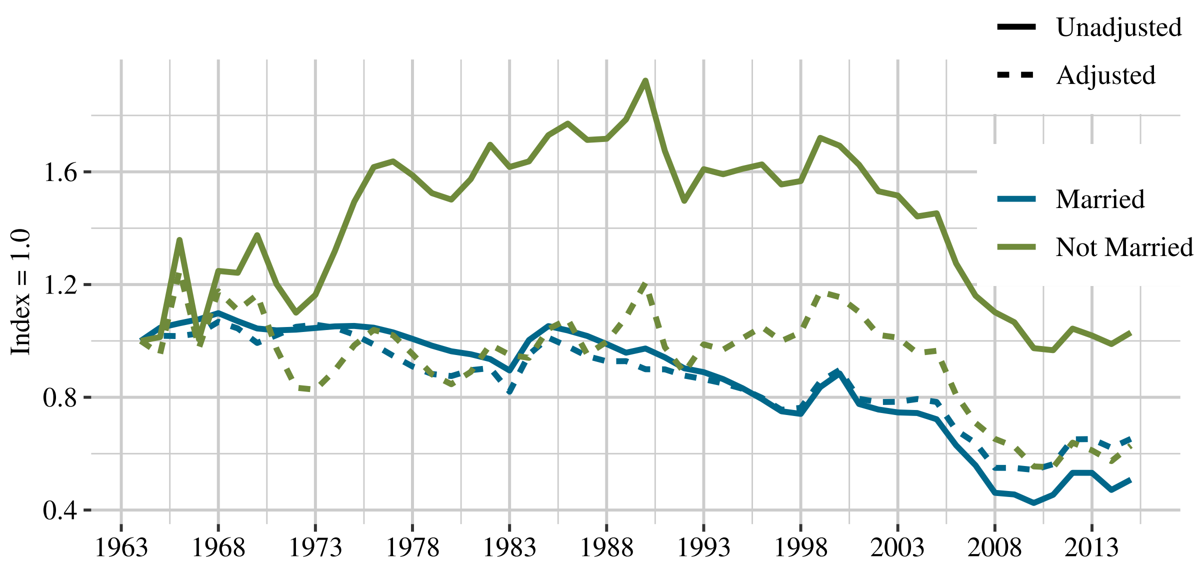

Notes: The figure shows the unadjusted and adjusted percent of households that moved within the previous year. The adjusted series are the sum of the percent of households that moved in 1964 for the group in question and the coefficients on the time dummies of a linear regression of the one year mobility status for that same group on time dummies, age, sex, race and education of the head of household , family size, and total real family income.

[2]An F-test also rejects the null hypothesis of $\beta_2 = \beta_3 = 0$ at the 1% level.

Table 2.1: Time Trend Estimation Results

	Not-Married Time Trend	Married-Time Interaction
Estimate	0.00093	-0.00741^{***}
Robust Standard Error	0.00127	0.00180
N	74	74

$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

2.2.2 Female Labor Force Participation and the Relative Earnings of Women

It is well known that female labor force participation increased rapidly after the end of World War II, from 39% in 1964 to just over 60% in 2000. This trend is even more prevalent among married women. Figure 2.3 shows that both the number of married and not-married women entering the labor force over this time increased substantially and that the labor force participation rate of married women increased at a faster pace, rising by almost 30 percentage points from 1964 to 2000.[3] Importantly, Figure 2.4 shows that the percent of dual searcher households increased from 36% in 1964 to 75% in 2000. The trend in the percent of dual searcher households began to flatten in the late 1990's and has remained almost unchanged at roughly 75%.

The rise in the number of dual searching married households is our first channel of interest, we call this channel the composition effect. As shown in Appendix A, the migration rate of dual searcher households is less than that of single searcher households. This suggests that the increase in female labor force participation which gave rise to a compositional change of married households (the rise in dual searcher households) may have some power in accounting for the decrease in the migration rate of married couples.

The second channel we focus on that may drive down the mobility rate for married couples

[3]Not shown is the labor force participation rate of married men that decreased from 85% in 1964 to 77% in 2000, implying that there may have been some crowding out of married men.

Figure 2.3: Female Labor Force Participation by Marital Status

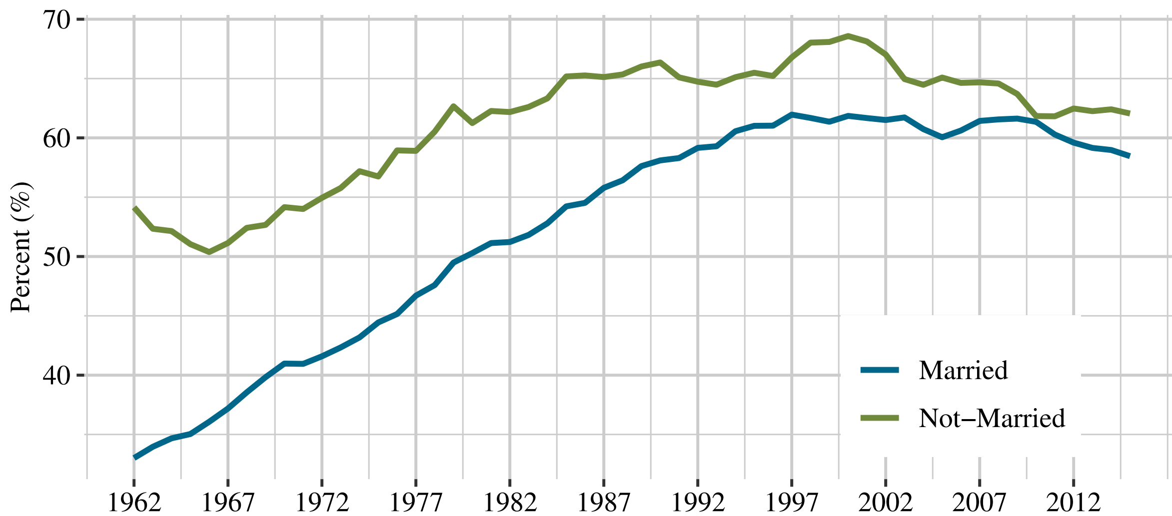

Notes: The figure shows the percent of women, age 16 or older, that are in the labor force by marital status. The data comes from the basic monthly files of the Current Population Survey.

is the ratio of women's to men's earnings. An increase in the ratio of women's to men's earnings may drive down mobility rates of dual searcher households since outside offers must be larger when both individuals are earning more. Figure 2.5 shows the ratio of median yearly earnings of married women to married men from 1962 to 2017. The figure shows that women's earnings relative to men began to increase in the 1980's, rising from about 49% to 64% in 2000. This time period corresponds to the second half of our period of interest as well as the period that saw the largest decrease in the mobility of married couples. We call this channel the wage effect.

In Appendix A we test the effect of the dual searching households and relative wage differences on the migration decisions of married couples using household level data from the March CPS. We show that dual searcher households are 0.4 percentage points less likely to move across county lines and among all households that move, dual searcher households are 26% less likely to move for job related reasons. We also show that dual searcher households

Figure 2.4: Dual Searcher Households

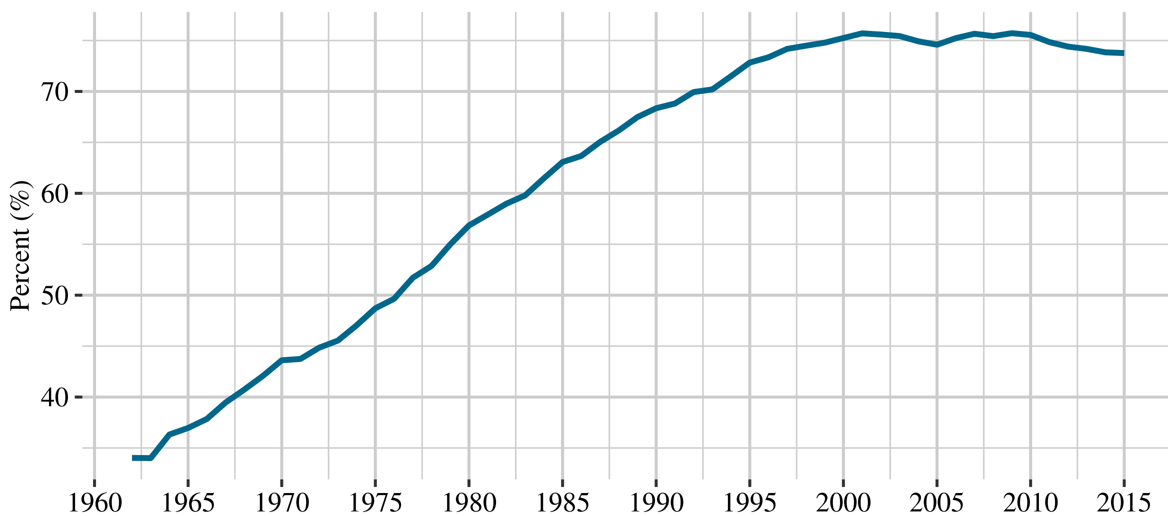

Notes: Plotted is the percent of households in which the head of household is in the labor force that have a spouse who is also in the labor force. Households in which the head of household is not married are included in the sample and the percent of dual earner households is calculated as: (# of married households with both spouses in the labor force)/(Total # of married households). The data comes from the basic monthly files of the Current Population Survey.

that live in states with higher relative earnings of women are less likely to move and show suggestive evidence that among those households that do move, those that face higher relative earnings of women are less likely to move for job related reasons.

We have documented the fact that mobility rates have declined for married couples from 1960 to 2000 while female labor force participation and the fraction of dual searcher households concurrently increased. This evidence suggests that both spouses in the labor force has played a role in decreasing the fraction of married couples that choose to move. In the following section, we develop a model of dual labor search to quantify the extent to which the increase in the fraction of dual searcher households and the rise in the female-to-male median wage ratio have contributed to the observed decrease in mobility.

Figure 2.5: Ratio of Women's to Men's Median Earnings

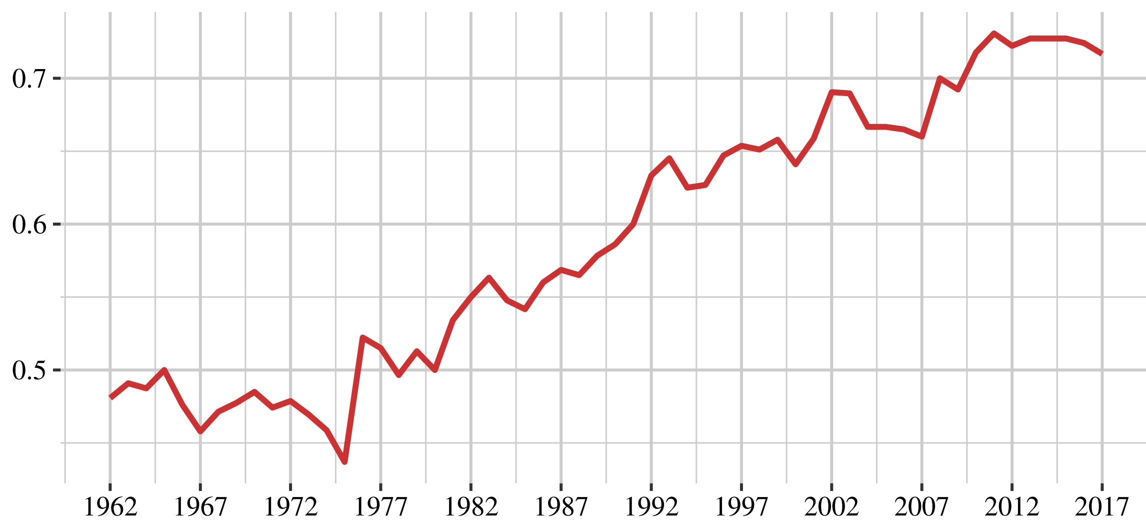

Notes: The figure shows the ratio of women's median yearly earnings to men's median yearly earnings.

2.3 Model

We are interested in modeling the job search problem for married couples under two different scenarios. First, the single searcher household, in which only one spouse is actively searching and receiving job offers. Second, the dual searcher household, in which both spouses are actively searching and receiving job offers. The key features of our model are that searchers can receive either local job offers or foreign job offers and that men and women receive offers from gender specific wage offer distributions.

2.3.1 Environment

Risk neutral households search for jobs and enjoy utility over pooled income similar to Guler et al. (2012).[4] There are two types of households, single searcher households and dual

[4]Examples of alternatives to the unitary model of the household include Dey and Flinn (2005), Gemici (2016), and Lundberg and Pollak (1993).

searcher households. Within dual searcher households, both individuals receive job offers and are ex ante heterogeneous as they have gender specific job prospects and receive gender specific flow utility of unemployment. In the single searcher household, individuals differ as only one is searching for jobs. For now we take the increase in female labor force participation as given and do not model the household's decision of becoming either a single searching household or dual searching household as a benchmark. We relax this assumption and allow for an endogenous participation margin in Section 2.5.

Individuals who are out of the labor force receive flow utility b_O and do not receive job offers. Individuals who are in the labor force but unemployed receive flow utility b_U^i, where $i \in \{M, F\}$ indexes the gender of each household member, and can receive either local or foreign job offers. They receive local offers at wage w drawn from the c.d.f. $F^i(w)$ at exogenous poisson arrival rate α_l^u, and foreign offers at wage w drawn from the same c.d.f. at exogenous poisson rate α_f^u. Individuals also search for jobs while employed and can again receive local or foreign offers at wage w drawn from the same c.d.f. at rate α_l^e and α_f^e, respectively. All jobs separate at exogenous rate δ and households discount utility at rate r. If an individual accepts a foreign offer, the household must quit any locally held jobs and move locations.[5] Here, we abstract from moving costs since we are ultimately interested in the difference between moving rates for single searcher and dual searcher households. While Kennan and Walker (2011) discuss the importance of such costs in migration decisions, the moving costs will simply create a wedge between the reservation wage of local and foreign offers that is similar for both single and dual searching households so long as these costs do not differ across households.[6] As a result, and to simplify our calibration, we exclude these one time fixed costs. As we discuss below, dual searching households will instead incur a cost to migration in the form of lost spousal income.

[5]An extension to allow for non-job related moves is discussed in Appendix B.

[6]Incorporating moving costs may differentially impact single and dual searching households when they are risk averse.

2.3.2 Single Searcher Household

A single searcher household is composed of two individuals: one out of the labor force and the other of gender i in the labor force searching for jobs. Such a household can be in one of two states: employed-out of the labor force with value function $EO^i(w)$ or unemployed-out of the labor force with value function UO^i. The value functions are given by:

$$\begin{aligned} rUO^i = b_O + b_U^i + \alpha_l^u \int \max\{EO^i(w) - UO^i, 0\}\ dF^i(w) \\ + \alpha_f^u \int \max\{EO^i(w) - UO^i, 0\}\ dF^i(w) \end{aligned} \tag{2.3}$$

$$\begin{aligned} rEO^i(w) = b_O + w + \alpha_l^e \int_w EO^i(w') - EO^i(w)\ dF^i(w') \\ + \alpha_f^e \int_w EO^i(w') - EO^i(w)\ dF^i(w') \\ + \delta\left[UO^i - EO^i(w)\right] \end{aligned} \tag{2.4}$$

where b_U^i is the flow value of unemployment for the unemployed spouse.

In either state the household receives flow utility b_O from the spouse that is out of the labor force. Since there is no cost to moving for the single searching household as only one spouse participates in the labor force, the reservation wage for both local and foreign offers is the same. Let R_s^i be the reservation wage such that $EO^i(R_s^i) = UO^i$. This reservation wage is given by the implicit equation,

$$R_s^i - b_U^i = \left(\alpha_l^u + \alpha_f^u - \alpha_l^e - \alpha_f^e\right) \int_{R_s^i}^{\infty} \frac{1 - F^i(w)}{r + \delta + (\alpha_l^e + \alpha_f^e)[1 - F^i(w)]}\ dw. \tag{2.5}$$

The steady state unemployment rate, u_s^i, and steady state distribution of households employed at wage less than or equal to w, $G^i(w)$, are given by Eq. 2.6 and Eq. 2.7.

$$u_s^i = \frac{\delta}{\delta + (\alpha_l^u + \alpha_f^u)[1 - F^i(R_s^i)]} \tag{2.6}$$

$$G^i(w) = \begin{cases} \frac{\delta[F^i(w) - F^i(R_s^i)]}{\{\delta + (\alpha_f^e + \alpha_l^e)[1 - F^i(w)]\}[1 - F^i(R_s^i)]} & w \geq R_s^i \\ 0 & else \end{cases} \tag{2.7}$$

Each are derived as in Burdett and Mortensen (1998). The migration rate for single searcher households of type i is the weighted sum of the migration rate of the unemployed plus the migration rate of the employed, where the weights are given by the mass of households in each employment state. The migration rate of the unemployed is: $\alpha_f^u[1 - F^i(R_s^i)]$. The rate at which workers employed at wage w migrate is $\alpha_f^e[1 - F^i(w)]$. Therefore, the aggregate migration rate for single searcher households of type i, M_s^i, is:

$$M_s^i = u_s^i \cdot \alpha_f^u \left[1 - F^i(R_s^i)\right] + \left(1 - u_s^i\right) \cdot \alpha_f^e \int_{R_s^i}^{\infty} \left[1 - F^i(w)\right] dG^i(w). \tag{2.8}$$

The aggregate migration rate for single searching households is then given by a simple weighted average over all household types as follows:

$$M_s = \xi_M \cdot M_s^M + (1 - \xi_M) \cdot M_s^F \tag{2.9}$$

where ξ_M denotes the fraction of single searching households with the husband in the labor force.

2.3.3 Dual Searcher Household

A dual searcher household is composed of two individuals both of whom are searching for jobs. Such a household can be in one of four states: employed-employed with value function $EE(w, w')$, husband employed-wife unemployed with value function $EU^M(w)$, wife employed-husband unemployed with value function $EU^F(w)$, and unemployed-unemployed with value function UU.

Just as in the single searcher household, the reservation wage for accepting jobs while in the unemployed-unemployed state is the same for both local and foreign offers as neither spouse must quit an existing job.[7] Let R_1^i be the reservation wage for spouse i when both members of the household are unemployed. Because both the offer distribution and flow utility of unemployment for each spouse is different, $EU^i(w)$ will differ by i. As a result, the reservation wage, R_1^i, is also indexed by i. The corresponding value function is:

$$\begin{aligned} rUU = b_U^M + b_U^F + \left(\alpha_l^u + \alpha_f^u\right) \int_{R_1^M}^{\infty} EU^M(w') - UU \, dF^M(w') \\ + \left(\alpha_l^u + \alpha_f^u\right) \int_{R_1^F}^{\infty} EU^F(w') - UU \, dF^F(w'). \end{aligned} \tag{2.10}$$

If one member of the household is employed, several decisions about accepting job offers need to be made. First, if the unemployed spouse receives a local job offer, they may take that offer if either the value of joint employment or the value of switching roles exceeds the current value of single employment. In the former case for example if the wife is employed at wage w' the husband will accept any local offer w such that $EE(w, w') \geq EU^F(w')$ and the household will enter a state of joint employment. Let $R_2^M(w)$ and $R_2^F(w)$ be the reservation wage for men and women to make this transition defined by $EE(w, R_2^F(w)) = EU^M(w)$ and

[7] Ex post inspection of the steady state value functions reveals that a cutoff strategy is optimal for the dual searching household.

$EE(R_2^M(w'), w') = EU^F(w')$, respectively. In the latter case, if the husband receives a wage offer sufficiently high to accept, but not high enough to enter joint employment, each spouse will switch roles and the household will remain in a state of single employment.

Related to this second transition is the fact that both the employed and unemployed spouse may receive a foreign offer. If the foreign offer is received by the spouse that is currently employed, the household will be willing to move for any wage greater than the one it is currently receiving. We do not allow for the possibility that members of the household can live in separate locations or that the household can split up. Therefore, if the unemployed spouse receives an acceptable foreign offer, the employed spouse must quit their job and transition into the unemployed state.

Since individuals are heterogeneous within the household, spouse i will accept the foreign wage offer, w', if and only if $EU^i(w') \geq EU^{-i}(w)$. Thus, the reservation wage to transition from employed-unemployed to unemployed-employed denoted R_3^i is given by $EU^i(w) = EU^{-i}(R_3^{-i}(w))$. Note that this reservation wage is identical for the local switching case discussed above and is not generally the 45° line. The value function for the employed-unemployed state is then given by:

$$
\begin{aligned}
rEU^i(w) = b_U^{-i} + w + &\left(\alpha_l^e + \alpha_f^e\right) \int_w^\infty \left[EU^i(w') - EU^i(w)\right] \, dF^i(w') \\
&+\alpha_l^u \int_{\phi^{-i}(w)}^\infty \max\left\{EE(w, w') - EU^i(w), EU^{-i}(w') - EU^i(w)\right\} dF^{-i}(w') \\
&+\alpha_f^u \int_{R_3^{-i}(w)}^\infty EU^{-i}(w') - EU^i(w) \, dF^{-i}(w') \\
&+\delta\left[UU - EU^i(w)\right]. \qquad (2.11)
\end{aligned}
$$

where $\phi^{-i}(w) = \min\left\{R_2^{-i}(w), R_3^{-i}(w)\right\}$.

If both members of the household are employed, each will accept local job offers above their current wage. If one receives a foreign offer, the household must decide whether or not to accept it and move. If the household chooses to move, the spouse who did not receive the offer transitions into the unemployed state and begins receiving flow utility b_U^{-i}. Moreover, if spouse i loses their job, their partner must decide whether or not to remain employed at their current wage or voluntarily quit and transition to the UU state rather than remain in the $EU^{-i}(w)$ state. Clearly iff $w \geq R_1^{-i}$, spouse $-i$ will remain employed rather than quit. Let $M^i(w, w')$ be the moving reservation wage for spouse i defined as $EE(w, w') = EU^i(M^i(w, w'))$ such that the household decides to move for all foreign offers above $M^i(w, w')$. Then the value function for the employed-employed state is:

$$
\begin{aligned}
rEE(w, w') = w + w' &+ \alpha_l^e \int_w^\infty [EE(w'', w') - EE(w, w')]\, dF^M(w'') \\
&+ \alpha_l^e \int_{w'}^\infty [EE(w, w'') - EE(w, w')]\, dF^F(w'') \\
&+ \alpha_f^e \int_{M^M(w,w')}^\infty \left[EU^M(w'') - EE(w, w')\right] dF^M(w'') \\
&+ \alpha_f^e \int_{M^F(w,w')}^\infty \left[EU^F(w'') - EE(w, w')\right] dF^F(w'') \\
&+ \delta \left[\max\left\{EU^M(w), UU\right\} - EE(w, w')\right] \\
&+ \delta \left[\max\left\{EU^F(w'), UU\right\} - EE(w, w')\right]. \qquad (2.12)
\end{aligned}
$$

Again, notice that because the value of being in the employed-unemployed state differs by the gender of the employed spouse, $M^i(w, w')$ is indexed by gender.

The spillovers between spouses detailed above can be seen as one particular way to micro-found the reduced form location match shocks detailed in Coen-Pirani (2010) needed to match population flows across locations. For example, suppose the male in a $EU^M(w)$ household separates from their job and so the household transitions into the UU state. For the case when the on the job arrival rate is less than that when unemployed, this household will now be more likely to migrate as the previously employed spouse now receives foreign offers at a higher rate.

Additionally, the spouse that was previously unemployed is more likely to cause a move as the household is not tied to any particular location by two employed spouses. On the other hand, suppose the unemployed spouse in the same $EU^M(w)$ household instead accepted a local offer and so the household transitions into the $EE(w,w')$ state. The household is now less likely to move as both spouses receive foreign offers at a lower rate and are tied to their current location by their spouse. These two situations can be interpreted as a bad and good location match shock, respectively. Moreover, this example further illustrates that dual searching households who have recently moved to a particular location are also more likely to out migrate as recent movers are least likely to be in the $EE(w,w')$ state. This is again consistent with Coen-Pirani (2010).

A steady state among dual searcher households consists of a set of four value functions, eight reservations wages, four measures of households, and three steady state distributions of households across jobs. As in the case of single searcher households, the reservation wages, measure of households in each state, and steady state distributions are sufficient to derive the aggregate migration rate. Let uu_d, eu_d^M, eu_d^F, and ee_d be the measure of households in the the unemployed-unemployed state, husband employed-wife unemployed state, wife employed-husband unemployed, and employed-employed state. Moreover, let $T^i(w)$ be the measure of households in the respective employed-unemployed state that are employed at wage less than or equal to w; and, let $H(w,w')$ be the measure of households in the employed-employed state in which one member is employed at wage less than or equal to w and the other is employed at wage less than or equal to w'.

The migration rate for dual searcher households is the weighted sum of the migration rates of all four states. uu_d households can move if either spouse receives an acceptable foreign offer while unemployed whereas ee_d households employed at wages (w,w') move if either spouse receives a foreign offer while employed in excess of $M^M(w,w')$ and $M^F(w,w')$, respectively. eu_d^i households employed at wage w can move for two reasons: if spouse i receives a foreign offer above w while employed or if spouse $-i$ receives a foreign offer while unemployed above $R_3^{-i}(w)$. Thus, the aggregate migration rate for dual searchers, M_d, is given by

$$\begin{aligned}
M_d &= \alpha_f^u \left[2 - F_m(R_1^m) - F_f(R_1^f)\right] \cdot uu \\
&+ \alpha_f^e \left(eu_f \cdot \int_{R_1^F}^{\infty} \left[1 - F_f(w)\right] dT_f(w) + eu_m \cdot \int_{R_1^M}^{\infty} \left[1 - F_m(w)\right] dT_m(w) \right) \\
&+ \alpha_f^u \left(eu_f \cdot \int_{R_1^F}^{\infty} \left[1 - F_m(R_3^m(w))\right] dT_f(w) + eu_m \cdot \int_{R_1^M}^{\infty} \left[1 - F_f(R_3^f(w))\right] dT_m(w) \right) \\
&+ ee \cdot \alpha_f^e \int_{R_1^M}^{\infty} \int_{R_2^f(w)}^{\infty} \left[1 - F_m(M_m(w, w'))\right] d^2 H(w, w') \\
&+ ee \cdot \alpha_f^e \int_{R_1^F}^{\infty} \int_{R_2^m(w')}^{\infty} \left[1 - F_f(M_f(w, w'))\right] d^2 H(w, w')
\end{aligned} \tag{2.13}$$

Finally, the aggregate migration rate of married couples is then a weighted average of that for single and dual searching households:

$$M_{agg} = \zeta_d \cdot M_d + (1 - \zeta_d) \cdot M_s \tag{2.14}$$

where ζ_d is the fraction of dual searching households among married couples.

2.4 Calibration

To carry out our quantitative experiment and decompose the contribution of the increase in female labor force participation and the increase in relative wages of women to the decline in the aggregate inter-county migration rate of married households, we first calibrate the model economy at an annual frequency. We fix a number of parameters and functional forms and use simulated method of moments to calibrate the remaining parameters. We use a time period of one year to calculate our moments. Note that our model is a continuous time model and so calibrating to annual targets still allows for multiple transitions to take place within a year. The discount rate is set to 0.04 to match an annual discount factor of 0.96. We normalize the flow value of being out of the labor force, b_O, to 0. The wage offer distribution, $F^i(\cdot)$, is assumed log-normal. The separation rate, δ, is set to 0.15, which matches closely estimates of involuntary separations detailed in Hall (2005).[8]

This leaves the local and foreign arrival rates of job offers both on and off the job, $\{\alpha_l^e, \alpha_f^e, \alpha_l^u, \alpha_f^u\}$ the flow values of unemployment for both men and women, $\{b_U^M, b_U^F\}$, and the location and shape parameters of the male and female offer distributions, $\{\mu^M, \sigma^M, \mu^F, \sigma^F\}$. These parameters are calibrated by matching key moments in the data. To maintain consistency with our econometric sample in Section 2.2 and Appendix A, we calibrate to the year 2000. In our quantitative exercise in Section 2.5, we then adjust $\{\mu^F, b_U^F\}$ to match the median wage gap in 1964 while holding the ratio of b_U^F to the mean offer fixed.

We use the March Annual Social and Economic Supplement (ASEC) supplement to the CPS to calculate our targeted moments. The median and 90^{th} percentile to median wage ratio of the observed wage of married men and women measured in 1999 dollars pin down the parameters of the offer distribution while the inter-county migration rates for single and dual searchers pin down the foreign arrival rates. Finally, the mass of female and male single searchers, and the mass of dual searching households within each state jointly pin down the local arrival rates and the flow value of unemployment for men and women. The results of the calibration are shown in Table 2.2 and Table 2.3.[9] The model moments closely match the data in all respects except

[8] The separation rate is the weighted annual separation rate for each type of household within our model, namely households of type $EE \to EU^i$, $EU^i \to UU$, and $EO^i \to UO^i$. Each separation rate is calculated as described by Shimer (2012) and weighted by the relative fraction of each type of household in our sample.

[9] We classify individuals as employed if they report usual weekly time spent working of at least 20 hours.

Table 2.2: Calibrated Moments

Moment	Model	Data	Targeted
Single Searcher Mig. Rate	0.057	0.057	✓
Dual Searcher Mig. Rate	0.047	0.047	✓
Mass in EE	0.82	0.79	✓
Mass in EU^M	0.12	0.13	✓
Mass in EU^F	0.056	0.047	✓
Mass in EO^M	0.95	0.89	✓
Mass in EO^F	0.90	0.82	✓
Male Median Wage ($)	$37,342$	$38,000$	✓
Female Median Wage ($)	$23,636$	$23,000$	✓
Male 90-50 Wage Ratio	1.38	2.15	✓
Female $90-50$ Wage Ratio	1.43	2.17	✓
Flow from UU to EU_m	0.98	0.99	
Flow from UU to EU_f	0.84	0.94	
Flow from EU_m to EE	0.80	0.99	
Flow from EU_f to EE	0.98	0.99	
Flow from UO_m to EO_m	0.93	0.99	
Flow from UO_f to EO_f	0.74	0.99	

for the 90^{th} percentile to median wage ratio, which we underestimate. Although the model is only calibrated to match the mass of households in each labor market state, the model also does relatively well in matching the flows between states.

The flow value of unemployment relative to the mean offer for men matches closely that of Shimer (2005). The flow value of unemployment for women is substantially closer to the mean wage offer, more similar to the estimates of Hagedorn and Manovskii (2008). This may reflect a number of factors including the division of homework within the household (Coltrane, 2000), gender stigmas such as those discussed in Evertsson and Nermo (2004), and child care considerations (e.g. Lundberg and Rose, 2000) that we do not model in this paper. Moreover, the fact that our estimated local arrival rates are substantially above their foreign counterparts suggest that households send the majority of job applications to local labor markets. This is consistent with the findings of Marinescu and Rathelot (2018).

Table 2.3: Calibrated Parameters

Parameter	Value	Description
α_l^u	28.591	Local unemp. arrival rate
α_l^e	14.878	Local emp. arrival rate
α_f^u	13.850	Foreign unemp. arrival rate
α_f^e	1.025	Foreign emp. arrival rate
μ_M	9.351	Male location parameter
σ_M	0.490	Male shape parameter
μ_F	8.730	Female location parameter
σ_F	0.546	Female shape parameter
b_U^M	$4,600$	Male flow utility of unemp.
b_U^F	$6,996$	Female flow utility of unemp.

The resulting reservation wages implied by the model are shown in Figure 2.6 and Table 2.4. Panel (a) shows the reservation wages for a spouse in the uu, eu_M, and eu_F states. Noticeably, $R_2^F(w)$ and $R_2^M(w)$ begin above R_1^F and R_1^M, respectively, and decrease below these two reservation wages as the wage of the employed spouse increases. When both spouses are unemployed, and one spouse receives an employment opportunity, they have the opportunity to provide higher consumption for the entire household thereby increasing the reservation wage of the still unemployed spouse; the same reason the reservation wage in a standard job search model increases with an increase in unemployment benefits. On the other hand, however, the now employed spouse also ties the household to their current labor market, decreasing the option value to search of the unemployed spouse and therefore their reservation wage. This is similar in spirit to how the reservation wage of an unemployed individual responds to a decrease in the arrival rate of job offers in the standard search model. When the employed spouse garners a low wage, the former effect dominates whereas the latter effect dominates as the wage of the employed spouse increases.

Panel (b) displays $R_3^M(w)$, $R_3^F(w)$, and the 45-degree line. Note that both $R_3^M(w)$ and $R_3^F(w)$ are bounded below by R_1^M and R_1^F, respectively. Here, the reservation wage for an unemployed male to begin working and their employed spouse to quit either voluntarily or due to a move is everywhere above the 45-degree line, whereas the opposite is true for an unemployed female and employed male. Because men have better job prospects than their wives, the option value for a searching male is higher than that for a searching female. As a result, couples are willing to sacrifice a small amount of consumption today with the wife working so that the unemployed male spouse can search for an even better job. On the other hand, it is costly for couples to allow the female spouse to search while the husband works because he is less able to take advantage of his superior offer distribution. Also of note is that both $R_3^M(w)$ and $R_3^F(w)$ are increasing in the wage of the employed spouse. This indicates that there are some foreign offers that dual searching households reject that their single searching counterparts would accept.

Panels (c) and (d) show the moving reservation wage for dual searching households in the *ee* state. The reservation wage is increasing in both spouses wages, again illustrating the increasing location effect as your spouse's wage increases. Note that this again implies that there are foreign jobs that a dual searcher in the employed state will reject that their single searching counterparts would accept and move. Furthermore, the moving reservation wage is everywhere below the sum of both spouses' wages. Upon moving, spouses enter the *eu* state, increasing the value of search for both spouses. This move frees the household to climb the job ladder quicker by reducing a couple's location ties.

Table 2.4 shows the reservation wages of single searching households in the *uo* state and dual searching households in the *uu* state. The reservation wage of dual searchers in the *uu* state is lower than that of their single searcher counterparts in the *uo* state. This arises due to the fact that accepting a mediocre job offer as a dual searching couple is less harmful than for a single searcher. That is, unemployed spouses are willing to accept comparatively worse offers today to boost household consumption and later quit either because their unemployed counterpart accepts a foreign offer or because they enter the breadwinner cycle.[10] Neither of

[10]This result is analogous to a single searchers problem with a higher job separation rate.

Figure 2.6: Model Reservation Wages

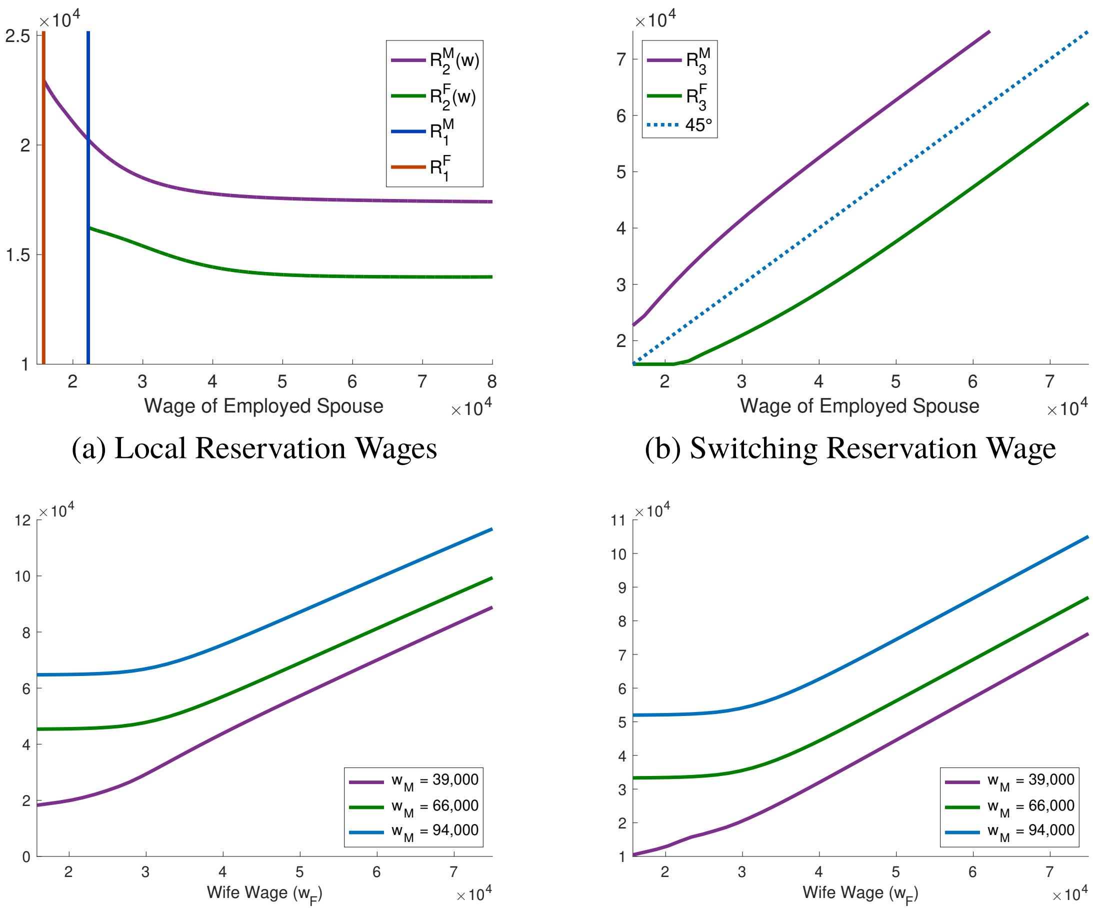

(a) Local Reservation Wages

(b) Switching Reservation Wage

(c) $M^M(w_M, w_F)$

(d) $M^F(w_M, w_F)$

these types of quits are available to single searching couples.

Table 2.4: Reservation Wages in the *uu* and *uo* States

	Men (\$)	Women (\$)
Dual Searcher	22,195	15,815
Single Searcher	24,475	17,020

2.5 Quantitative Experiment

2.5.1 Composition and Wage Effects

We now turn to our main quantitative experiment. In particular, we are interested in the effect of the increase in female labor force participation among married households and the increase in the relative wage of wives on the aggregate migration rate. To find the composition effect, we hold all calibrated parameters constant and adjust the share of dual searching households to match that in 1964. To estimate the wage effect, we fix the wage offer distribution of men and adjust the μ_F to match the 1964 female to male median wage ratio and adjust b_U^F such that the ratio of b_U^F to the mean wage offer remains constant. Moreover, we fix all other parameters including the fraction of dual searchers during this counterfactual.

The findings are given in Table 2.5 where we have normalized the migration rates in the calibration year to 100. We first describe our estimate of the composition effect. The model implies that the aggregate migration rate increased from 4.92% in 2000 to 5.33% in 1964. The composition effect therefore implies an overall change of 0.41 percentage points (8.3%). In the data we see that the migration rate instead increased from 5.0% in 2000 to 5.74% in 1964, a total change of 0.74 percentage points (14.8%). Thus, we conclude that the composition effect accounts for approximately 55.4% of the total decline in intercounty migration of married households seen in the data. If, in addition to changing the fraction fo dual searchers, we also change the fraction of male single searchers among all single searchers to match its 1964 value of 95.8%–up from 78.3% in 2000–then the composition effect accounts for 64.9% of the observed decline.

Next, we investigate the contribution of the wage structure to the decline in the migration rate. In our sample, the median real female wage increased from $15,243 in 1964 to $23,000 in 2000. The median real wage for men increased from $31,273 to $38,000 over this same period. These changes correspond to an increase in the female to male median wage ratio from 0.48

to 0.64. Decreasing the relative wage of married women to target a median wage ratio of 0.48 results in an increase in the migration rate of dual searching households from 4.66% to 4.82%, and in a negligible change in the migration rate of single searching households–it remained at 5.71%. Overall, the model implies an increase in the aggregate migration from 4.92% in 2000 to 5.04% in 1964. The wage effect thus implies a change of 0.12 percentage points (2.4%). Thus, we conclude that the wage effect accounts for 16.2% of the change in intercounty migration of married households. Moreover, the model implies that the composition effect is stronger overall than the wage effect. This indicates that the fact that women are entering the labor force is more important than the fact that they are becoming more similar to men in their roles within the household in explaining these long term migration trends within the United States.

Table 2.5: Quantitative Results

	1964	2000	Change
Composition Effect			
Model	108.2	100.0	8.2
Data	114.8	100.0	14.8
Contribution	–	–	55.4%
Wage Effect			
Model	102.4	100.0	2.4
Data	114.8	100.0	14.8
Contribution	–	–	16.2%
Combined Effect			
Model	109.6	100.0	9.6
Data	114.8	100.0	14.8
Contribution	–	–	64.8%

To measure the degree to which complementarities exist between the two effects, we also conduct a joint counterfactual. In particular, we adjust the female offer distribution to match the female to male median wage ratio in 1964 after fixing the fraction of dual searching households to equal that in 1964. This joint counterfactual implies an aggregate migration rate of 5.39%, a rise of 0.47 percentage points (9.6%). The combination of the forces can therefore account for approximately 65% of the decline in aggregate migration over this time period. Notice that the rise in dual searcher's migration rate resulting from the wage effect is mitigated by the fact that there are simply fewer dual searchers in the counterfactual. Hence the total effect of decreasing women's relative wages and decreasing the fraction of dual searching households is less than the sum of these two individual effects. If we again also change the fraction of male single searchers among all single searchers, then the total effect increases to $\sim 72.9\%$.

Our results indicate that our mechanism was a key force reducing the number of moves by married couples before the 21st century. Consistent with Kaplan and Schulhofer-Wohl (2017) and Cooke (2011) the demographic characteristics herein studied, namely the rise in dual searcher households and relative wages of women, can explain very little of the decline in intercounty migration during the 21st century.

2.5.2 Endogenous Labor Force Participation

To estimate the indirect effect of wages on migration through labor force participation decisions, we include a participation margin. In the model, we allow each member of the household to be endowed with a flow value of non-participation, b^i_O, that has c.d.f. $G^i(b^i_O)$, where i indexes men and women. For simplicity, we assume that $b^M_O \perp b^F_O$. Conditional on a draw of b^i_O, the household decides whether to be a single searching household with the male in the labor force, a single searching household with the female in the labor force, or a dual searching household. Once this decision is made, the participating spouse(s) enter the labor force in the unemployed state and begin searching.[11] We maintain our assumption that only employed and unemployed individuals receive job offers. Several characteristics of our model with a participation decision are worth noting. First, because we have assumed that

[11] Pissarides (2000) and Flinn (2006) model the participation margin analogously in standard search models.

household preferences are linear, the option value to search of the participating member of single searching households is unaffected by the level of b^i_O received by the nonparticipating household member. Second, since the value of non-participation is independent of labor market parameters, conditional on their participation decision, households solve the same problem described in section 3 of the paper.

To calibrate the parameters of G^i and make the participation decision quantitative, we choose a parametric form. We follow Flinn (2006) and set G^i to be the c.d.f. of the exponential distribution with parameter λ^i, which adds two parameters that need to be calibrated: one scale parameter for each distribution. We use the fraction of dual searching households and the fraction of male single searchers among all single searching households as our calibration targets for λ^M and λ^F.

Table 2.6: Calibrated Moments: Participation Decision

Moment	Model	Data
Mass of Dual Searchers	0.752	0.752
Mass of Male Single Searchers	0.783	0.783

Table 2.7: Calibrated Parameters: Participation Decision

Parameter	Value	Description
λ^M	20,224	Mean Value of Male Non-participation
λ^F	34,640	Mean Value of Female Non-participation

Table 2.6 reports the data and model moments and Table 2.7 reports the calibrated parameters. We begin by estimating the total effect by recalibrating μ_F and b^F_U to match the median wage ratio of women and men in 1964 and λ^M and λ^F to match the fraction of dual searching households and the fraction of male single searchers among all single searching households in 1964. We estimate the wage effect by solving for the model implied migration rates using

the participation parameters, λ^M and λ^F, calibrated to 2000 and the wage parameters, μ_F and b_U^F, calibrated to 1964. Similarly we estimate the composition effect by solving for the model implied migration rates using the wage parameters calibrated to 2000 and the participation parameters calibrated to 1964.

Table 2.8: Counterfactual Parameters: Participation Decision

Parameter	Value	Description
λ^M	17,912	Mean Value of Male Non-participation
λ^F	105,384	Mean Value of Female Non-participation

Table 2.9 reports the results of the three counterfactual exercises. The composition effect results in a rise of the aggregate migration rate from 4.92% to 5.37%, or 62% of the total change. This results from the fact that there are both fewer dual searchers and more male single searchers as a fraction of all single searchers. The former decreases from 75.2% to 39.4% whereas the latter rises from 78.3% to 96.5%. Moreover, the migration rate of single searchers rises from 5.71% to 5.83% due to the increase in the fraction of single searching couples that have men in the labor force. Of note is Table 2.8, which shows the counterfactual scale parameters for the distribution of non-participation values. Relative to 2000, the mean value of non-participation for men declined by 11% whereas the mean value of non-participation for women rose by factor of ~ 3. This captures a number of features, including increased stigma, home production expectations, and child rearing, for women in the year 1964 as compared to 2000. In contrast, the outside option for men changes only slightly, implying that stigma and expectations in the home have increased to a far lesser degree than the associated changes for women.

The wage effect instead results in a rise of the migration rate of dual searchers from 4.66% to 4.82% and a negligible change in the migration rate of single searchers. In addition, the fraction of dual searching households among all households decreases from 75.2% to 71.3% as a result of fewer women participating in the labor force. This change accounts for roughly

10% of the rise in dual searching households since 1964, implying the other factors such as stigma were a more important driver than rising female job prospects for the rise of female labor force participation among married women. The total model implied aggregate migration rate is 4.92% in 2000 and 5.07% in 1964; thus the total wage effect accounts for 23.7% of the decline in aggregate migration. The combination of these forces, i.e. the wage and composition effects, results in a rise of the aggregate migration rate from 4.92% to 5.46%. Thus, the total effect is approximately 73%. As in our benchmark model, we see that there are negative complementarities between the wage and composition effects.

Table 2.9: Quantitative Results: Endogenous Participation

	1964	2000	Change
Composition Effect			
Model	109.2	100.0	9.2
Data	114.8	100.0	14.8
Contribution	–	–	62.2%
Wage Effect			
Model	103.5	100.0	3.5
Data	114.8	100.0	14.8
Contribution	–	–	23.7%
Combined Effect			
Model	110.8	100.0	10.8
Data	114.8	100.0	14.8
Contribution	–	–	73.0%

2.5.3 Exogenous Moves

We introduce exogenous moving shocks into the benchmark model to account for the fact that not all moves are a result of job changes. Let η_d and η_s be an exogenous poisson arrival rate of non-job related moves for dual searching households and single searching households, respectively. Further, we assume that the arrival rate of such moves is not optimal, i.e. all employed members of the household become unemployed after an exogenous move. We continue to allow for endogenous participation. The value functions, migration rates, and calibrated parameters for the model with exogenous moves can be found in Appendix B. We calibrate $\eta_d = 0.0287$ and $\eta_s = 0.0314$ to match the fraction of non-job related moves shown in Table A.4.[12] Table B.1 reports the parameter values and Table B.2 reports the data and model moments.

We conduct the same counterfactual exercises as before. The results are both qualitatively and quantitatively similar to those in Section 2.5.2 and reported in Table 2.10. In particular, the combination of changing female labor force participation and rising female wages results in a decline in the migration rate of 7.8% (0.39 percentage points), from 5.40% to 5.01%. Relative to the decline of 14.8% (0.71 percentage points) observed in the data, these two forces can explain roughly 52.7% of the decline. The wage effect, i.e. the change in the migration rate from only adjusting female wage offers, can explain roughly 21.6% of the observed decline. This results from the fact that the migration rate of dual searchers declines from 5.05% to 4.84% and the fraction of dual searching couples rises from 73.3% to 75.2%. Thus, the wage effect causes more households to become location constrained and existing location constraints to become stronger. The migration rate of single searchers is unchanged from the wage effect.

The composition effect, i.e. the change in the migration rate resulting from the change in the mix of dual versus single searching households, can explain approximately 41.9% of the observed decline. The composition effect results from the fact that there are both more dual searching couples and fewer male single searchers, who tend to have higher migration rates than their female counterparts. In this case, the migration rate of single searchers declines from 5.60% to 5.51% and the fraction of dual searchers rises from 37.5% to 75.2%. The dual

[12] The fraction of non related moves is the sum of Family and Other in column Dual-TOT for dual searchers and Single-TOT for single searchers.

Table 2.10: Quantitative Results: Exogenous Moves

	1964	2000	Change
Composition Effect			
Model	106.2	100.0	6.2
Data	114.8	100.0	14.8
Contribution	–	–	41.9%
Wage Effect			
Model	103.2	100.0	3.2
Data	114.8	100.0	14.8
Contribution	–	–	21.6%
Combined Effect			
Model	107.8	100.0	7.8
Data	114.8	100.0	14.8
Contribution	–	–	52.7%

searcher's migration rate is unchanged. Evidently, changes in the outside option of women have been more important for rise in dual searchers than the increase in their wage offers when participating in the labor force. In fact, the counterfactual outside option of women shown in Table B.3 is 3.2 times higher than that in 2000. Finally, there are negative complementarities between the composition and wage effects. This results from the fact that, the large increase in the migration rate of dual searchers in our counterfactual year is mitigated by the fact that there are fewer such households in 1964.

2.5.4 Lifetime Wage Inequality

Flinn (2002) argues that lifetime inequality measures should be estimated using a structural approach. In particular, many large data sets available to researchers are repeated cross-sections rather than long panels of individual earnings. Those panels in which the entire career history is observable are only available for some populations, thereby making generalization difficult.

Moreover, estimates obtained in the absence of a structural model do not allow researchers to investigate the effects of potential labor market reforms or of different existing institutional frameworks across labor markets. Our model has important implications for these types of structural exercises. Flabbi and Mabli (2018) and Bowlus and Robin (2004) have since estimated lifetime earnings inequality in exercises similar to Flinn (2002), however, all three have ignored the co-location problem. Here, we add to their results by estimating the bias of lifetime earnings inequality resulting from ignoring the co-location problem. We further show how the importance of explicitly modeling it has evolved over time. We use our previous calibrations and simulate the career paths of 100,000 households of each type. Each household begins in the unemployment state and ends their career with the first labor market spell, either employment or unemployment, that ends after a total work history of 40 years. We then calculate lifetime wage earnings for spouse j as

$$\omega(j) = \sum_{i=1}^{N} e^{-r\tau_i} \int_{0}^{t_i} w_{i,j} e^{-rv} dv \tag{2.15}$$

where t_i is the duration of labor market spell i for the household, $\tau_{i+1} = \tau_i + t_i$ is the starting time of the $i+1$ labor market spell for the household, and N denotes the number of labor market spells that begin prior to the 40^{th} year. Moreover, we set $w_{i,j} = 0$ when spouse j is unemployed.

Table 2.11 displays the coefficient of variation for men and women when assuming they are single searchers as opposed to dual searchers in both our calibration and counterfactual year. The bottom panel further shows that the coefficient of variation of lifetime earnings for all married individuals when simulating the appropriate mix of dual and single searchers in 1964 and 2000. Our model shows that ignoring the dual search problem can substantially bias estimates of lifetime earnings inequality. For men, ignoring the dual search problem biases estimates of lifetime earnings inequality upwards. This results from the fact that dual searching men optimally choose to enter the state of unemployment as a result of their wives accepting foreign offers and fewer dual searching men accepting higher, but foreign wage offers that would satisfy their single searching counterparts. In 1964, however, this fact was relatively unimportant as female job prospects within the household were relatively unimportant. As female job prospects have improved over time, the co-location problem and therefore the bias in these estimates becomes more severe for men.

Analogous mechanisms are at play for estimates of female lifetime earnings inequality. When women's job prospects were relatively unimportant within the household, more dual searching women optimally choose to enter the state of unemployment as a result of their husbands accepting foreign offers and fewer dual searching women accepting higher, but foreign wage offers that would satisfy their single searching counterparts. This results in very few high wage women relative to low wage women causing the single searcher assumptions to overestimate earnings inequality. Conversely, as female job prospect become more important within the household, a larger fraction–but still not all–of women decline to quit their jobs so that their husbands may accept foreign offers and instead accept high paying foreign offers themselves. As a result, bias in estimates of lifetime earnings inequality decrease, and in this case, become negative.

Our model suggests that lifetime earnings inequality across all married people has declined, with the coefficient of variation falling from 0.46 to 0.38. Much of this decline is due to the fact that men and women are becoming more equal in terms of labor market outcomes. Even so, the model implies that the bias in this measure has increased in magnitude from 3.3% to 23.1%. This results directly from the fact that dual searching households look less and less like their single searching counterparts and that there is a larger fraction of these types of households in 2000 than in 1964. This again illustrates the importance of explicitly modeling the co-location problem when studying inequality.

Despite abstracting from a number of important features detailed in Huggett et al. (2011) that account for earnings inequality, our results suggest that explicitly modeling the co-location

Table 2.11: Lifetime Earnings Inequality

	1964	2000
Men		
Single Searchers	0.1316	0.1715
Dual Searchers	0.1307	0.1426
Bias (%)	0.7	20.3
Women		
Single Searchers	0.1979	0.1473
Dual Searchers	0.1454	0.1548
Bias (%)	36.1	−4.8
All		
Without Dual Searchers	0.4479	0.2919
With Dual Searchers	0.4634	0.3796
Bias (%)	−3.3	−23.1

problem is also important for these estimates. Not doing so has the potential to produce severely biased results, particularly among men during the post 2000 period and women prior to the rise in their relative wages. A potential extension of the mechanisms presented here would include heterogeneity in initial human capital and ability, asset accumulation, and life-cycle skill accumulation in the face of borrowing constraints as in Huggett et al. (2011), Rendon and García-Pérez (2018), and Griffy (2017), respectively. Adding these features would allow us to compare the relative importance of the co-location problem to these other channels in a more detailed way and allow us to assess the impact of potential policy reforms on lifetime inequality for all family types. We leave such extensions to future work.

2.6 Conclusion

Between 1964 and 2000, the fraction of couples in which both spouses are in the labor force nearly doubled while the female to male wage ratio increased by 30%. Contemporaneously, the intercounty migration rate of married couples decreased from 5.74% to 5.0%, whereas the migration rate of single people increased from 5.4% to 9.1%. These differential trends suggest important differences in the decision making process of singles and married couples. Using the March CPS supplement from 1999-2015, we show that dual searching couples are 10% less likely to move than their single searching counterparts. Moreover, among those couples who did move, dual searching couples are 26% less likely to move for job related reasons.

Using a two location job search model with both single and dual searching households we then decompose the contribution of increasing female labor force participation rates and increasing female wages relative to men to the historical decline in migration. We find that the increase in the fraction of dual searching households can account for roughly 55.4% of the decline whereas the rise in relative female wages can account for approximately 16.2% of this decline. Moreover, we show that ignoring the co-location problem biases estimates of lifetime earnings inequality for married individuals downward by up to 23%, and that explicitly modeling this decision within the household has become more important as female job prospects within the household have become more equal.

Chapter 3

Demographic Obstacles to European Growth

3.1 Introduction

One of the most striking characteristics of advanced economies has been the secular rise in life expectancy. During the last 50 years, life expectancy at birth in advanced economies has increased by over ten years and, according to U.N. projections, it is expected to continue increasing.[1] This impressive increase in longevity combined with a decrease in fertility has resulted in ageing populations in most of the developed economies. Ageing populations have powerful implications. In this paper, we estimate the impact of ageing on long-term aggregate economic growth. We find that ageing populations have depressed growth rates in recent decades and will further depress economic growth in the future.

We study the impact of changing demographics for aggregate growth in Europe's four largest economies: France, Germany, Italy, and the United Kingdom. We quantify the impact

[1] Case and Deaton (2017) have documented a recent slight decrease in life expectancy among U.S. males in certain socioeconomic groups due to "Deaths of Despair" – deaths due to suicide, drug overdose and obesity. This is largely a U.S. phenomenon not in evidence in other countries.

of demographic change on aggregate factor supply and demand, and on the growth outcomes of these economies. Since the early 1990's these four economies have experienced a slowdown in long-run growth that is persistent but not uniform. Compared to the prior two decades, annualized long-run growth over the last 20 years fell by between 0.8 percentage points in France and 2.1 percentage points in Italy. At the same time, these countries have experienced persistent increases in longevity and declines in fertility rates. The combination of these two factors has resulted in populations ageing to different degrees within each country. Additionally, we estimate the indirect growth effects of the frictions and distortions that result from the higher marginal taxes that are needed to finance the pension benefits of an ageing society.[2]

Growth accounting allocates growth outcomes to total factor productivity growth, population growth, and changes in factor supplies – specifically capital accumulation and labor supply on both the intensive and extensive margin. Changes in life expectancy and the age-cohort distribution of countries affect all of these channels. An increase in longevity affects individual factor supply decisions whereas changes in the age composition of populations affects the aggregation of individual assets and labor supply. Changes in the aggregation of labor supply also affects measured TFP as a greater or smaller fraction of those choosing to work may be in the most productive years of their lives as the relative cohort distribution changes. The combination of these forces induces general equilibrium effects, with changes in factor prices further affecting individual decisions.

Demographic change may also affect growth indirectly through pension systems. As populations age, and individuals either choose or are forced into retirement, governments' pension liabilities increase. If an increase in the number of retirees is coupled with a shrinking tax base as the relative number of individuals choosing to work decreases, tax rates will have to adjust to balance government budgets. Increases in tax rates introduce additional frictions and

[2] Cooley and Henriksen (2018) show the impact of changing demographics on growth in the U.S. and Japan but do not explicitly model the impact of retirement systems on individual decisions and economic growth.

distortions to individual saving and labor supply decisions.

We use a parsimonious general equilibrium overlapping generations model with a rich demographic structure, endogenous retirement, age heterogeneity in productivity, and a pay-as-you-go pension system to better understand the behavioral responses to increases in life expectancy and changing factor prices, and estimate the effect of changing demographics on economic growth. A structural approach is needed to (i) make projections for future growth, (ii) analyze policy responses to stagnating economic growth, and (iii) have a laboratory to evaluate the welfare consequences of policy alternatives. Often proposals to reform pension systems, for example extending the eligibility age, are evaluated simply on the basis of their impact on cost and output. But, a structural model is useful for evaluating the welfare consequences of reforms.

To account for observed labor supply choices on both the extensive and intensive margin, a key ingredient in our model economy is the dis-utility of labor supply at different ages. The sharp decline in labor force participation at older ages implies that the dis-utility of labor increases at old ages. The convex nature of the dis-utility of labor may be due to both psychological factors– individuals may simply be tired of working after 30+ years in the labor force– or to physiological factors, e.g. declining health and fitness over the life-cycle. This paper is agnostic on the exact causes of the dis-utility of labor over the life cycle, but we calibrate our model to observed age specific labor force participation rates in each of the countries taking as given the pension system in those countries. This enables us to analyze the incentive effects of alternative pension systems. More importantly, it provides a laboratory to evaluate the welfare effects of possible reforms.

We find that the contribution of demographic change to growth is substantial and can account for as much as 70% of the secular growth slowdown in the case of France and Germany. For the United Kingdom and Italy, demographic change can account for 50% and 25% of their respective growth slowdowns. Moreover, our model predicts that demographic change will cause growth to decline further over the next 20 years. In terms of the wider question of "secular stagnation", this paper complements the conclusions of Gordon (2016), Summers (2014, 2016), and others who have written about secular stagnation.

The primary channel through which demographic change operates in our model economy is via changes in labor supply. Decreases in the aggregate employment-to-population ratio is offset to some extent by capital deepening induced by increased savings. At the individual level, increases in life expectancy affect consumption, labor supply and savings decisions as households must adapt to a longer life span and changes in factor prices. A striking fact that the model needs to match is that, despite longer, and presumably healthier lives, life-cycle labor-supply choices, particularly retirement behavior, have changed less. This deepens our understanding of demographics and economic changes as previous studies focusing on the impact of demographic change on interest rates and capital flows have focused on savings choices and generally not considered how demographic change also affects labor-supply choices. For example, Krueger and Ludwig (2007), Backus et al. (2014), and Ferrero (2010) investigate the effect of demographic change on real interest rates and international capital flows, assuming that individuals supply labor inelastically between fixed ages irrespectively of factor prices or changes to life expectancy. In Section 3.3.6, we separately show that both the quantitative and qualitative implications of ageing populations hinge crucially on how individual labor supply responds to increases in life expectancy.[3]

We find that the need to finance pension systems did not lower long-run growth very much over the 1975-1995 period. Over the last 20 years, however, pension systems have decreased growth, sometimes substantially. In Italy, our model indicates that the need to finance pension outlays decreased annual growth by an additional 0.18 percentage points. Moreover, these distortions will decrease annual growth even further over the next 20 years, with labor supply

[3] See also Henriksen and Lambert (2018), Feroli (2003), Sposi (2019), and Bárány et al. (2019) Similarly, also using models solely focusing on savings decisions, Gagnon et al. (2016), Carvalho et al. (2016), and Ikeda and Saito (2014) show the impact of demographic factors on the real interest rate in the United States and Japan.

accounting for most of the decline in projected growth. These results add to the extensive literature on the implications of ageing for the sustainability of social security systems, e.g. Fuster et al. (2007) and İmrohoroğlu et al. (2016), by emphasizing that labor-supply choices are critical and that the big outstanding question is why increases in longevity have not resulted in larger changes in the effective retirement age.[4]

This paper is also related to the literature on late-life labor supply. Erosa et al. (2016) show that a fixed cost to participation is key for matching aggregate Frisch elasticities associated with life-cycle labor supply. Ndiaye (2020) studies endogenous retirement decisions when individuals are exposed to idiosyncratic shocks. Others, such as French (2005), van der Klaauw and Wolpin (2008), and Erosa et al. (2012), argue that social security rules have a sizable impact on retirement behavior. Capatina (2015), Pashchenko and Porapakkarm (2017), and French and Jones (2011) instead study the role of health risk and show that it may be just as important as social security rules in accounting for labor supply choices. While all of these papers account for labor supply over the life cycle and retirement choices, our focus is to estimate how labor and capital supply change over time as factor prices evolve and how these changes affect long-term growth.

Several of the aforementioned papers assume either a constant or linear cost to labor force participation, or allow individuals' time endowment to change only as a result of changing health status. As a result, they either have difficulty matching the labor force participation profiles of both healthy and unhealthy individuals after the age of 60+, or do not attempt to match labor supply profiles that late in life. This, however, is a key part of the life cycle for understanding the effects of ageing populations on economic growth and potential policy reforms to mitigate any adverse welfare effects.

Our paper is also related to Auclert et al. (2019) and Börsch-Supan et al. (2019), who study

[4]A more exhaustive review of such papers includes, but is not limited to, Auerbach and Kotlikoff (1987), De Nardi et al. (1999), Garibaldi et al., eds (2010), Kotlikoff et al. (2007), Kitao (2014), Conesa and Garriga (2016), McGrattan and Prescott (2017), and McGrattan and Prescott (2018).

how changing age-cohort distributions may affect future growth taking decisions as given. A key difference between those papers and this paper is that we model decisions, which allows us, among other things to provide a deeper understanding of the interaction between ageing and factor supply decisions. It also provides a laboratory for making projections, and a welfare measure by which to evaluate reforms.[5]

The rest of the paper is organized as follows. Section 3.2 describes the long term growth and demographic trends in France, Germany, Italy, and the United Kingdom. Section 3.3 describes our model along with two methods of financing our pension systems. Section 3.4 describes the calibration and Section 3.5 our numerical approach. Section 3.6 presents the benchmark results, our historical decompositions, and growth projections. Section 3.7 presents analysis of some pension reforms that have been suggested to increase long-run growth, and some robustness tests. Section 3.8 concludes.

3.2 Growth and Demographic Change

3.2.1 Historical Growth

The world's largest economies have experienced a growth slowdown over the last five decades. Figure 3.1 shows GDP-per-capita trends for the four largest European economies and the United States. In the two decades immediately following World War II, Germany, France, and Italy experienced significant catch up, in large part due to the build-up after wartime destruction. Our goal is to estimate the role that demographic changes may have played since the 1970s.

We can decompose historical growth into its constituent components using growth accounting to determine the contributions of factor inputs and productivity using a standard

[5] Also related are Kopecky (2018), Hopenhayn et al. (2018), and others who study how changing demographics may account for productivity growth through firm dynamics and the decline in entrepreneurial activity.

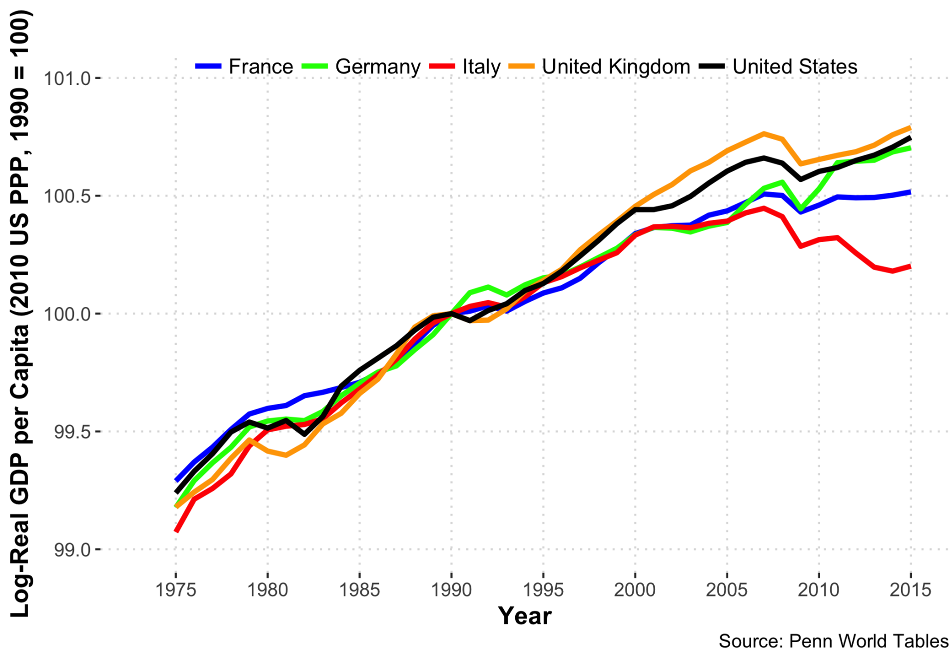

Figure 3.1: Real GDP per Capita in G7 Economies Excluding Canada and Japan

Cobb-Douglas production function:

$$Y = A \cdot K^{\alpha} (L \cdot h)^{1-\alpha} \tag{3.1}$$

where A is TFP, K is the aggregate capital stock, L is the number of workers, and h is the average hours worked by those in the labor force.[6] This implies an expression for growth, which includes both an intensive and extensive labor-supply margin given by h and $\frac{L}{pop}$, respectively.

$$\gamma_Y = \gamma_A + \alpha\gamma_{K/L} + \gamma_{L/pop} + \gamma_{pop} + (1-\alpha)\gamma_h \tag{3.2}$$

In Eq. 3.2, γ_i is the growth rate of component i. Population growth can trivially account for GDP growth and so we exclude it from the rest of our discussion, instead focusing on GDP-per-capita

[6]Our assumed capital share is consistent with that of our calibration discussed in Section 3.4

Table 3.1: Historical Growth Accounting Annualized Growth Rates

	$\gamma_{Y/pop}$	γ_A	$\alpha \cdot \gamma_{K/L}$	$\gamma_{L/pop}$	$(1-\alpha) \cdot \gamma_h$
1975-1995					
France	1.88%	1.61%	0.82%	-0.09%	-0.47%
Germany	2.33%	2.24%	0.76%	-0.09%	-0.57%
Italy	2.43%	1.43%	0.89%	0.20%	-0.09%
United Kingdom	2.31%	1.80%	0.68%	-0.01%	-0.17%
United States	2.20%	1.10%	0.32%	0.77%	0.01%
1995-2014					
France	1.07%	0.73%	0.36%	0.16%	-0.17%
Germany	1.32%	0.94%	0.20%	0.55%	-0.37%
Italy	0.31%	0.01%	0.33%	0.23%	-0.26%
United Kingdom	1.54%	1.05%	0.21%	0.38%	-0.10%
United States	1.53%	1.31%	0.43%	-0.12%	-0.10%

growth. The per-capita growth accounting expression then becomes

$$\gamma_{Y/pop} = \gamma_A + \alpha\gamma_{K/L} + \gamma_{L/pop} + (1-\alpha)\gamma_h \tag{3.3}$$

The annualized results of this exercise are shown in Table 3.1.

Growth accounting highlights both the persistence and heterogeneity across countries in the growth slowdown shown in Figure 3.1. The largest single contributing factor is the decline in TFP growth. Our model economy may shed some, but limited light on whether the decline in TFP growth has demographic roots. Our model economy may, however, account for changes in the supply of capital and labor and make projections for these into the future given demographic projections. In particular, Table 3.1 further shows that understanding the determinants of labor-supply at both the intensive and extensive margins is crucial for understanding the growth experience of these economies. In order to account for these facts, a low-frequency structural change that has first order implications for labor supply and differs across countries is necessary.

3.2.2 Demographic Trends

Life expectancy at birth among the advanced economies has increased steadily as shown in Figure 3.2. Life expectancy among these countries increased from an average of 72.5 to 77.3 and 77.3 to 81.8 between 1975-1995 and 1995-2015, respectively. In other words, every year life expectancy has increased by almost a quarter of a year. U.N. projections suggest that this trend will continue, although at a slightly slower rate, with a predicted rise from 82.7 to 85.5 between 2020 and 2040.

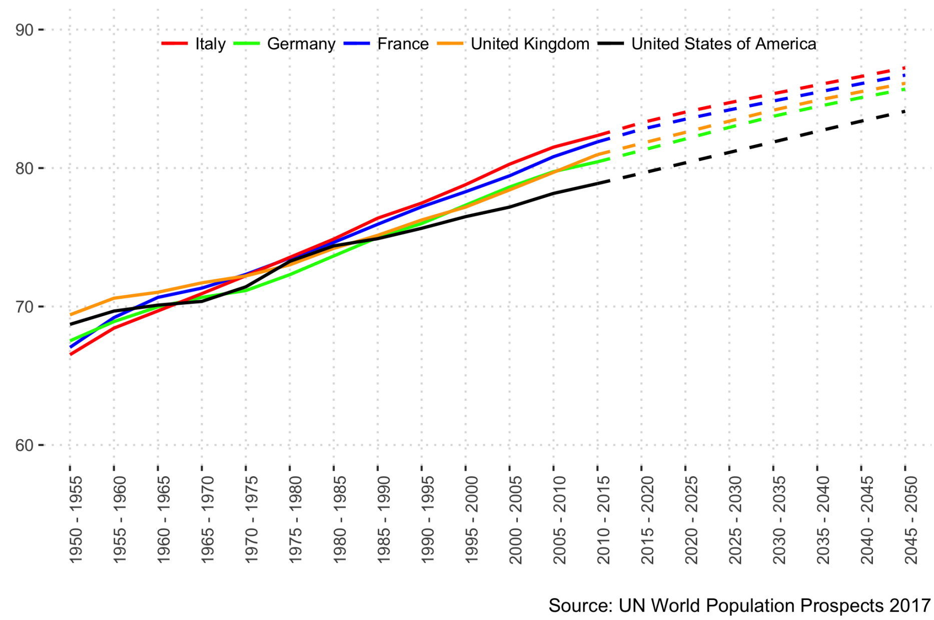

Figure 3.2: Life Expectancy in G7 Economies Excluding Canada and Japan

These dramatic changes in longevity combined with lower fertility have caused populations to age, some significantly. Figure 3.3 illustrates the historical shift in the age-cohort distribution for the United Kingdom, France, Germany and Italy, respectively, and how they differ. Two characteristics stand out. In all economies, population ageing has persisted for several decades and is projected to continue. Second, the historical and projected rightward shifts in the age-cohort distribution differs across these four countries. The low frequency nature of these trends and their differences between countries implies that demographic factors may contribute to both the decline in long-run growth in each country and the different growth histories across countries.

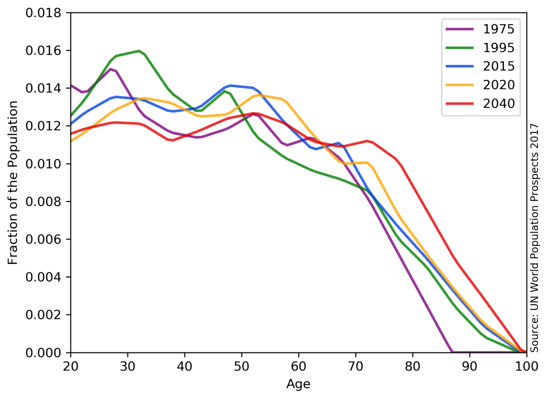

(a) United Kingdom

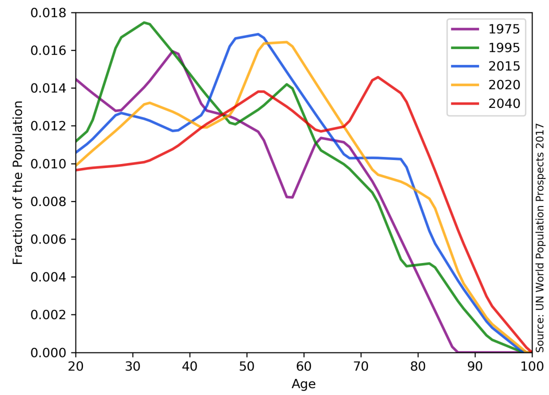

(b) Germany

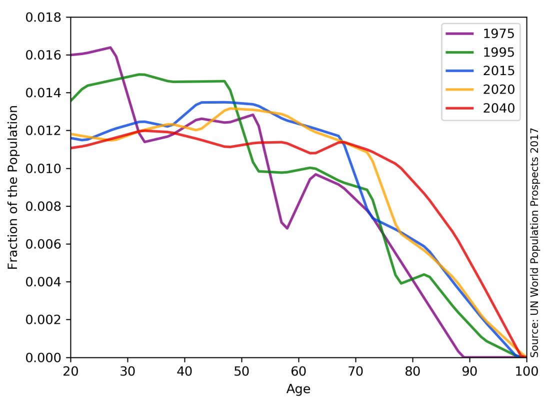

(c) France

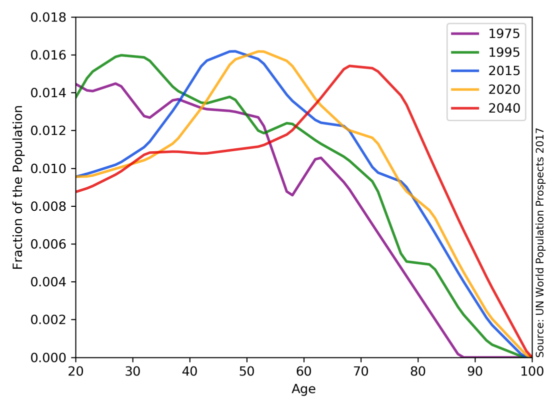

(d) Italy

Figure 3.3: Age-Cohort Distributions

Related to the rise in longevity, and to the rightward movements of the age-cohort distribution, is morbidity. That is, increased lifespans may be accompanied by improved health at old age or may simply result in more years in poor health at the end of life. While limited, both the OECD and Eurostat have recently published old age health indicators. Figure 3.4 displays the fraction of people aged 65+ and older who self report their health as "good" in the left panel and the number of years a person aged 65 can expect to live in good health. With the exception of Italy, the trends in health are not nearly as striking as those in longevity. Perceived old aged health in Britain has remained essentially constant since 1991 and that in Germany has not yet rebounded since a large decline in 2005.[7] Instead this measure increased by approximately 7 percentage points. Elderly Italians, however, have seen large gains to old age health in recent years. After remaining constant from 1994-2006, the fraction reporting good health rose by 23.7 percentage points.

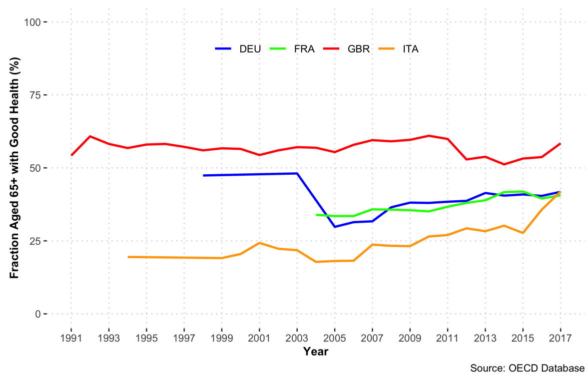

(a) Perceived Health

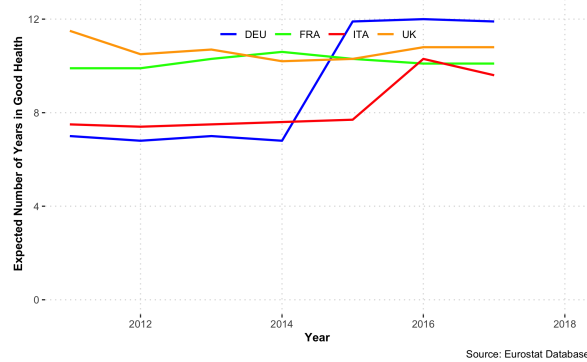

(b) Healthy Life Years at 65

Figure 3.4: Historical Old Age Morbidity

The Healthy Life Years index does not rely on self perceived health and instead defines good health as the absence of disabilities and limitations to functioning. Over the past nine years, there have been no noticeable gains to healthy life years in the United Kingdom and France. That in Germany remained constant through 2014, rose in 2015, and remained constant over the following two years. Italy is similar: healthy life years remained constant through 2015, rose in 2016, and subsequently fell in 2017.

An additional, often overlooked feature of the European experience is that, even as life

[7] If we instead consider only the post 2005 period, the fraction of Germans aged 65+ reporting good health increased from 31.4% to 41.8%.

expectancy has increased substantially, the average retirement age has not.[8] On average, individuals have predominantly allocated additional years of life to retirement.

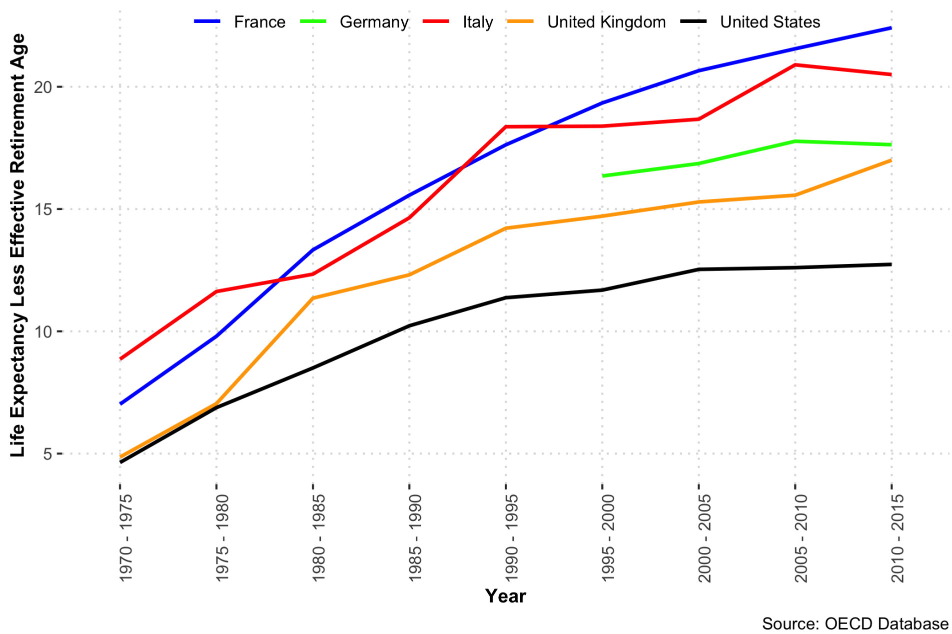

Figure 3.5: Years in Retirement in G7 Economies Excluding Canada and Japan

3.2.3 Growth and demographics

Increases in life expectancy and changes in the age-cohort distribution can affect economic growth through all the five channels identified by the growth accounting exercise in Eq. 3.2. Figure 3.6 shows that labor supply on the intensive margin shows a clear hump-shaped pattern over the life cycle across these European countries. If this hump shape remains unchanged as the cohort distribution shifts, aggregate hours will also change. In addition, labor supply choices on the intensive margin may change as life expectancy increases and factor prices change due to demographic factors.

[8]While this has been historically true for the United States as well, Figure 3.5 shows that the increase in the gap between effective retirement age and life expectancy has been considerably smaller and has remained roughly constant since the mid to late 1990’s.

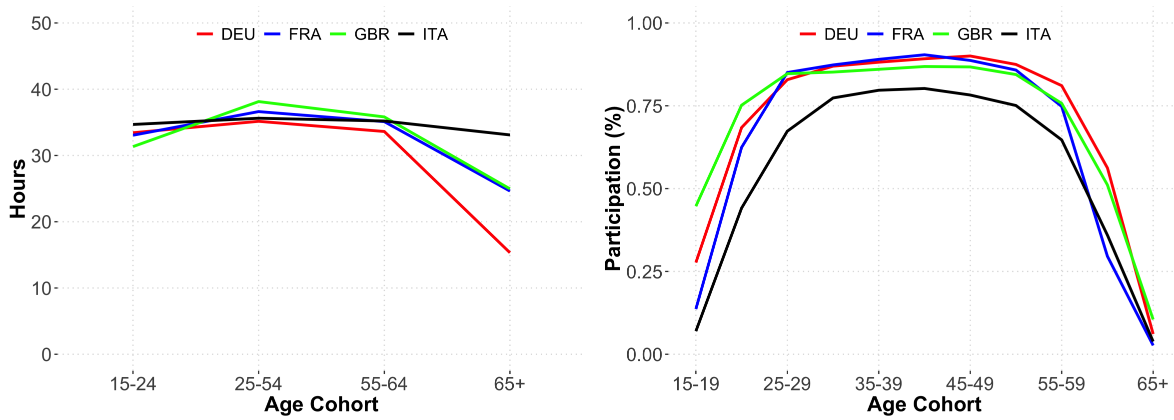

(a) Average Hours Worked (b) Labor Force Participation

Figure 3.6: Life-Cycle Labor Supply in 2015

Similar mechanisms hold for labor supply on the extensive margin. Figure 3.6 also shows a clear hump-shaped pattern of labor force participation over the life cycle across these European countries. The number of individuals in the labor force will therefore be directly affected by shifts in the cohort distributions. In addition, labor supply choices on the extensive margin may also change as life expectancy increases and factor prices change due to demographic factors.

As the average age increases, more individuals will be in their wealthiest years. In order to smooth consumption over a longer expected lifetime, individual savings rates should increase. Both demographic factors shaping individual choices and the aggregation of these choices contribute to capital deepening, ie. an increase in the capital-to-labor ratio.

Ageing populations also affect measured TFP. As populations age, the fraction of the workforce in the most productive years of their lives also changes. Since productivity is measured conditional on number of hours worked, this affects measured TFP.

In addition, the more individuals who have chosen to retire relative to the number of individuals who have chosen to work, the higher are the taxes necessary to finance pensions and other programs supporting retirees. These taxes will distort labor-supply choices both on the extensive and intensive margin. All of these factors, both direct and indirect, affect equilibrium prices through the aggregate capital stock and labor supply.

One can see that labor supply at the extensive margin for older workers has increased slightly in recent years but it is far overshadowed by the rise in life expectancy. Since this decline in labor force participation late in life does not appear to be driven by declines in health, that leaves preferences and retirement policies as possible drivers. Preferences in our model include a cost of participation function that is convex in age and is calibrated for each country to match participation rates at several ages. In a later section we study how policies can affect both participation and welfare at various ages.

3.3 Model

Our model economy is as parsimonious as possible while addressing all the five growth channels identified by the growth accounting exercise. In particular, individuals make labor-supply choices on both extensive and intensive margin and savings choices over the life cycle. In order for the model to match observed retirement behavior, we assume that the disutility of working is increasing with age. In order to distinguish between labor-supply choices on the intensive and the extensive margin, the model is calibrated to idiosyncratic shocks to labor productivity over the life cycle.

The benchmark economy abstracts from pensions. In this economy, individuals fund their own retirement consumption by savings. Subsequently, we introduce a pension system where old age benefits are financed by workers with either lump sum or distorting taxes. We decompose the growth effects of demographic change into a direct effect and an indirect effect operating through the increasing wedges necessary to finance increasing pension outlays. These environments allow us to discuss the extent to which pension systems impose additional obstacles to growth and identifies the margins most affected by them. The benchmark economy is also closed with no capital flows between countries. This constraint is later relaxed.

3.3.1 Households

At each age, i, households maximize their expected discounted utility by choosing consumption and labor supply conditional on their life expectancy

$$\max_{\{c_j,h_j\}} \mathbb{E}_t \sum_{j=i}^{I} s_j \beta^{j-i} u(c_{j,t+j}, h_{j,t+j}) \tag{3.4}$$

where β is the household's discount factor, s_i is the probability that a household lives from age i to $i+1$, and c_i and h_i are consumption and hours worked at age i, respectively. We assume that households participating in the labor force cannot work less than $h = 0.2$ hours. Household preferences are assumed to be additively separable both within and across periods and take the iso-elastic form given by

$$u(c,h) = \frac{c^{1-\sigma}}{1-\sigma} + \chi \frac{(1-h-\theta_i \cdot I_p)^{1-\gamma}}{1-\gamma} \tag{3.5}$$

Here, σ denotes the intertemporal elasticity of substitution, $l = 1-h-\theta_i \cdot I_p$ is effective leisure, and γ defines the curvature over effective leisure. I_p is an indicator function that takes a value of 1 if $h > 0$ and 0 otherwise. Households' cost to participation, θ_i, is allowed to differ by age and is given by the following functional form.

$$\theta_i = \kappa_1 + \kappa_2 \cdot i^{\kappa_3} \tag{3.6}$$

This cost function may capture a number of life-cycle features discussed in the previous section, such as deteriorating health, changes in tolerance to fatigue and stress, and other life-cycle incentives, such as retirement systems. The labor supply literature has largely emphasized the first and last of these three considerations. French (2005) and Capatina (2015), for example, evaluate the role of deteriorating health in late life labor supply. They assume that the cost to participation differs between sick and healthy individuals, but each is fixed over the life cycle. Instead, they allow the probability of negative health to increase with age. Related are Rust and Phelan (1997), Blau and Gilleskie (2006, 2008), and French and Jones (2011) who estimate the retirement incentives induced by medicare and employer provided health insurance. While their findings are mixed, they all find that health insurance can at least partially account for observed retirement behavior. Our specification captures all of these features while preserving the parsimony of our model. Most importantly, however, it can be easily matched to old age labor force participation rates.

Households maximize expected discounted utility subject to their budget constraint, which is given by

$$c_{i,t} + a_{i+1,t+1} = (1 + r_t)a_{i,t} + w_t \cdot h_{i,t} \cdot \psi_i \cdot \eta_{i,t} + b_t \tag{3.7}$$

where c is consumption, a is asset holdings, r is net rate of return on capital, w is the hourly wage rate, h is the number of hours worked, ψ is age-dependent productivity, η is the household's idiosyncratic productivity, and b is accidental bequests. To close the model, we assume that accidental bequests, b_t, from households that exit the model as a result of mortality risk are evenly distributed among all surviving households.

We further assume that households begin their economic lives with no assets and enforce a no-ponzi condition, producing an initial condition and boundary condition given by

$$a_{i_0,t} = 0 \qquad \text{and} \qquad a_{i_{max},t} \geq 0 \tag{3.8}$$

where i_{max} is the maximum allowed age. We maintain Eq. 3.8 throughout the paper.

Households differ both between and within cohorts in several ways. First, household productivity differs between cohorts due to an age specific productivity profile, ψ_i, and within cohorts as a result of idiosyncratic labor productivity shocks. We assume that the idiosyncratic component of individual productivity follows an AR(1) process in logs for each individual given by

$$\ln \eta_{i+1} = \rho \ln \eta_i + \epsilon_{i+1} \tag{3.9}$$

where $\epsilon_i \sim \mathcal{N}\left(0, \sigma^2\right)$ is Gaussian white noise. These two sources of heterogeneity create differences in average hours worked between cohorts and hours dispersion within cohorts, respectively. In particular, these productivity differences incentivize households who are in the most productive years of their lives or who have received a series of high productivity draws to work more hours.

Households also face an endogenous and irreversible retirement decision each period. Consistent with Erosa et al. (2016) and Rogerson and Wallenius (2013), who show the importance of a fixed cost to work in accounting for labor supply elasticities and retirement, the interaction between our cost to participation function, θ_i, and life cycle productivity profile, ψ_i, generates heterogeneity in labor force participation rates between cohorts.[9] As discussed in Section 3.4, our calibrated cost to participation function is increasing in age while our life-cycle productivity profile is hump shaped. Labor force participation rates, therefore decrease over the life cycle since it becomes more costly to remain in the labor force and households are less productive on average at old ages.

The interaction between Eq. 3.9 and Eq. 3.6, further provides a mechanism within our model to generate endogenous workforce composition. As θ_i increases with age, households who have received a series of poor idiosyncratic productivity shocks become less likely to remain in the workforce. As a result, only the most productive households continue working at old ages. This both creates within cohort differences in labor force participation and makes our model consistent with papers showing that life-cycle earnings is flatter than life-cycle productivity, e.g. Rupert and Zanella (2015).

[9] More precisely, Rogerson and Wallenius (2013) emphasize the need for non-convexities in either individual budget constraints or choice sets to generate reasonable intertemporal elasticities of labor.

3.3.2 Technology

We assume that a representative firm with constant returns to scale Cobb-Douglas production technology demands capital and labor, and produces a numeraire good for consumption in perfectly competitive markets. Thus, the firm's problem is given by

$$\max_{K_{d,t},L_{d,t}} \left\{K_{d,t}^{\alpha} L_{d,t}^{1-\alpha} - (r_t + \delta)K_{d,t} - w_t L_{d,t}\right\} \tag{3.10}$$

where $0 < \alpha < 1$ is capital share, $K_{d,t}$ is aggregate capital demand, $L_{d,t}$ is aggregate labor demand measured in efficiency units, r_t is the net real interest rate, and w_t is the real wage rate. Moreover, the aggregate capital stock evolves according to the usual law of motion,

$$K_{t+1} = (1 - \delta)K_t + I_t \tag{3.11}$$

where δ is the depreciation rate and I_t is net investment.

3.3.3 General Equilibrium

An equilibrium in this environment is defined as follows:

1. Households choose savings, consumption, and labor supply on the extensive and intensive margins taking prices and conditional survival probabilities, s_i, as given such that

   - Household's solve the following recursive problem each period:

$$v_{LF}(i, a, \eta) = \max_{c,a',h} \left\{u(c, h) + \beta \cdot s_i \cdot \mathbb{E}_{\eta'|\eta} \max \left\{v_{LF}(i + 1, a', \eta'), v_R(i + 1, a')\right\}\right\}$$

$$v_R(i, a) = \max_{a',c} \left\{u(c, 0) + \beta \cdot s_i \cdot v_R(i + 1, a')\right\}$$

- Decisions are aggregated to get the aggregate supply of capital and labor measured in efficiency units:

$$K_{s,t} = \sum_i x_i \cdot \int_{a\times\eta} a \cdot \mathrm{d}\mu(a, \eta \mid i, t)$$

$$L_{s,t} = \sum_i x_i \cdot \int_{a\times\eta} h \cdot \psi_i \cdot \eta \cdot \mathrm{d}\mu(a, \eta \mid i, t)$$

where x_i is the fraction of the population constituted by cohort i and $\mu(a, \eta \mid i, t)$ is the stationary joint distribution of a and η in time t for cohort i.

2. Firms maximize profits taking prices as given:

$$\max_{K_{d,t}, L_{d,t}} \left\{ K_{d,t}^{\alpha} L_{d,t}^{1-\alpha} - r_t K_{d,t} - w_t L_{d,t} \right\}$$

3. Markets clear:

$$\{r_t, w_t\} \mid K_{s,t} = K_{d,t} \;\&\; L_{s,t} = L_{d,t}$$

3.3.4 Demographics

Our definition of general equilibrium shows that the conditional survival probability at each age, s_i, and the age-cohort distribution, x_i, in each period are sufficient statistics to capture demographics within our model. Several factors influence the evolution of a country's age-cohort distribution. First, changes in mortality rates reduce the number of deaths per year. Second, declines in a country's fertility rate reduces the degree to which aging cohorts are replaced by new, younger individuals. Both of these effects serve to shift the age-cohort distribution right. Lastly, a country's cohort distribution is affected by net migration flows, thereby shifting the cohort distribution either left or right depending on the mix of migrants.

Let $x_t \in \mathbb{R}^I$ denote the vector of length I where each element contains the fraction of the population of age i at time t. Each cohort is endowed with an age specific fertility rate, $f_{i,t}$, and a conditional survival probability, $s_{i,t}$, in each period. Moreover, denote $m_t \in \mathbb{R}^I$ as the vector of net migration of each age group in period t. Then the evolution of the cohort distribution within a given country is given by

$$x_{t+1} = \Gamma_t x_t + m_t \tag{3.12}$$

where,

$$\Gamma_t = \begin{bmatrix} f_{1,t} & f_{2,t} & f_{3,t} & \dots & f_{i,t} \\ s_{1,t} & 0 & 0 & \dots & 0 \\ 0 & s_{2,t} & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & s_{I-1,t} & 0 \end{bmatrix} \tag{3.13}$$

Note that the cohort distributions used in our quantitative exercise do not necessarily equal the stationary distribution implied by Γ_t. Instead, we assume that individuals believe that the current demographic structure and therefore prices will persist in perpetuity. Results should be interpreted as calculating steady states implied by each point along the demographic transition path.

Our model consists of I overlapping generations. We assume that households enter their independent economic lives at age 20 and that no individual can live past age 100. Prior to age 20, households do not work, accumulate assets, or consume. Moreover, all households are born with $a_{i_0} = 0$ net assets.

At each age, i, households face mortality risk. Denoting s_i as the probability of surviving to age $i + 1$ conditional on reaching age i, the unconditional probability of reaching age j is given by $s^j = \prod_{i=1}^{j-1} s_i$. These survival probabilities capture changes to life-expectancy within our model and distort the discount rate at each age. Upon death, any assets saved from age i to $i+1$ are transferred equally across the remaining population in the form of lump sum accidental bequests, b_t.

3.3.5 Pension Systems

To quantify the implications of public defined-benefit pensions and their financing, we introduce a parsimonious pension system with guaranteed old age benefits. Here, we assume that households believe that current pension systems will persist indefinitely and that taxes are adjusted to balance government budgets period-by-period. For computational simplicity, we take the level of real pension benefits to be constant and equal for each eligible household across cohorts above some eligibility age, I_R. In particular, define $\tau_{L,t}$ and t_t to be the labor tax rate and lump sum taxes levied on households, respectively, at time t. The household budget constraint then becomes

$$\begin{aligned} c_{i,t} + a_{i+1,t+1} =& (1 + r_t) \cdot a_{i,t} + (1 - \tau_{L,t}) \cdot w_t \cdot h_{i,t} \cdot \psi_i \cdot \eta_{i,t} \\ &+ b_t - t_{t,i} + p_t \cdot \mathbf{1}(i \geq I_R) \end{aligned} \tag{3.14}$$

where p_t is the level of real pension benefits. Moreover, we assume that lump sum taxes are age dependent. In particular, lump sum taxes used to finance pension benefits are levied only on those who are not eligible for old age benefits. This assumption prevents net transfers received in old age from being distorted by the financing of pension benefits. To close the model, define the budget constraint of the government to be

$$\sum_{i \geq I_R} \mathbf{x}_i \cdot p_t = \tau_{L,t} \cdot w_t \cdot L_t + T_t \tag{3.15}$$

where T_t denotes total lump sum tax revenues.

Pension systems have important implications for individual decisions, particularly household labor supply. First, incentives to accumulate assets and work later in life as life expectancy increases is mitigated relative to a world without pension systems as individuals may rely on social security in addition to individual savings to smooth consumption. Second, as a larger fraction of the population enters retirement, these pension systems can create disincentives to work especially if stopping work is a requirement for collecting the pension.. As populations age, a greater number of households become eligible for benefits causing total pension outlays to increase. If increases in life expectancy do not translate into large increases in labor force participation at old ages, the tax base will decrease and force tax rates to rise. The changes in these tax rates provide greater disincentives to work as populations age and result in even smaller increases in retirement ages as life expectancy rises.

We note that we have opted to include only public pension systems and exclude other government purchases such as public goods expenditures. The inclusion of such expenditures would not qualitative affect our results and would only increase the magnitude of the quantitative effects discussed below. As described in Section 3.4, we allow for only a single tax rate to adjust throughout our quantitative exercises. Including other government expenditures leads to projected tax rates that are, in our view, unreasonably large. To resolve this issue, we would need to include a more complicated tax and transfer system that incorporates adjustable capital taxation. Such an extension would simply complicate our discussion without changing our central conclusions.

3.3.6 Importance of Labor-Supply Decisions

In our model, labor supply decisions late in life have important consequences. We calibrate the model's preferences to match the observed labor supply at various ages given the pension system in place. But, in principle, individuals can supply any amount of labor including working longer as life expectancy increases. The behavior of prices and interest rates will depend on these decisions.

To illustrate the importance of labor supply decisions, we solve a special case of our model with inelastic labor supply. In contrast to our benchmark model, the vast majority of the literature investigating the consequences of aging populations assumes that labor is supplied inelastically until a given retirement age. Key results, such as secular decline in interest rates, rest on this particular assumption.

The focus of literature studying demographic change has largely been on its effects on the supply and demand for capital. The general assumption in the literature has been that labor is supplied inelastically until some exogenous retirement age. The implicit assumption is that retirement age does not change as life expectancy rises. That is to say that gains to longevity, both in the past and in the future, will translate one-for-one into more years spent in retirement with no adjustment to individual work behavior. Understanding how labor supply decisions change as expected longevity and factor prices change are not only of crucial importance for making projections for future real interest rates, but also for making growth projections and estimating the welfare effects of potential reforms.

To illustrate the sensitivity of previous results to the assumptions made about labor supply, we solve a simplified version of our model that reflects those used in the literature. We assume that households make only a consumption-savings choice, are identical within cohort, and supply labor inelastically until an exogenous retirement age. At retirement they exit the labor force and collect social security. In effect, this model is identical to our model after fixing $\{\eta_i, \psi_i\} = \{0, 1\}\ \forall\, i$ and $\chi = 0$. The rest of our model environment is left unchanged. At the benchmark, life expectancy and the exogenous retirement age are first set to 70 and 65, respectively. We then increase life expectancy to 80 years, solve the model for retirement ages between 65 and 75, and compare the model-implied equilibria with the benchmark.[10] The range of retirement ages represent the entire range between two extremes: i) Retirement age does not adjust at all in response to gains to life expectancy and remains at 65 years. All gains to longevity translate one-to-one to more years in retirement. And ii), retirement age adjusts one-for-one with life expectancy, while expected time in retirement remains the same.

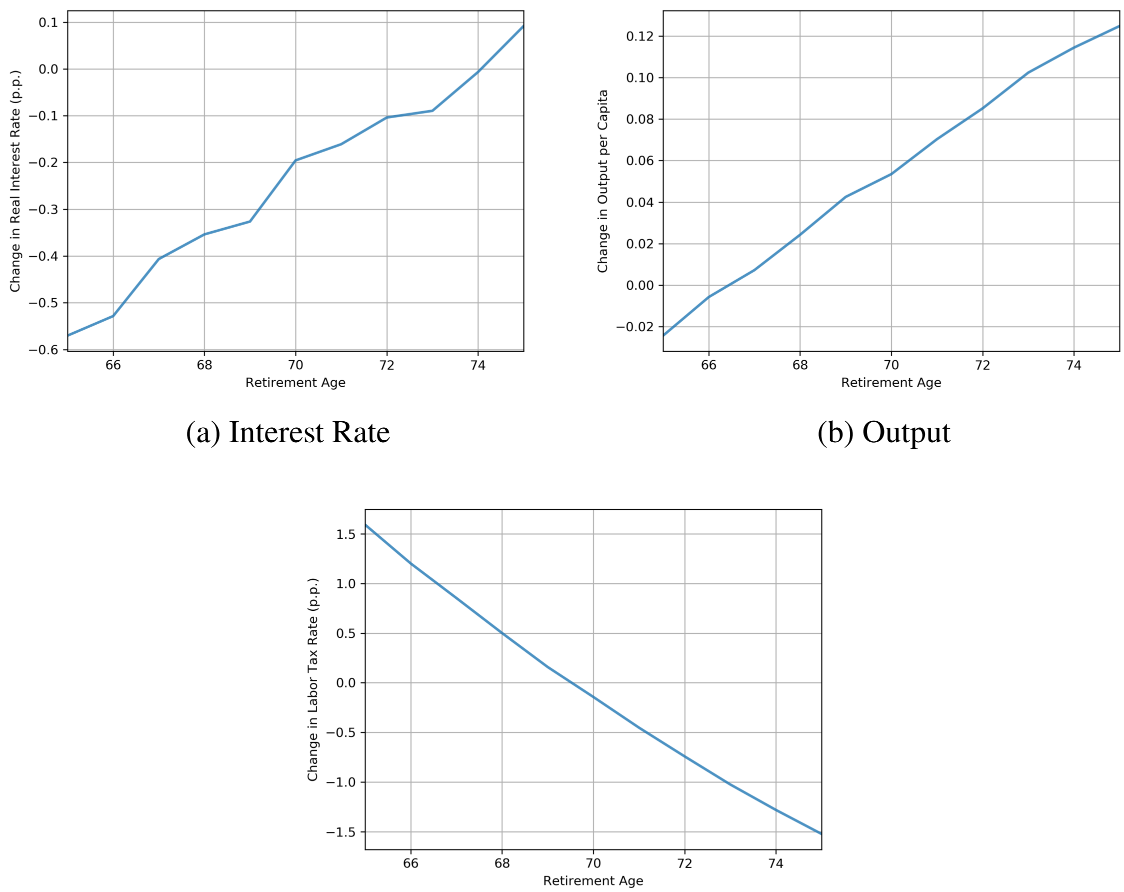

(a) Interest Rate

(b) Output

(c) Tax Rate

Figure 3.7: Sensitivity of the Retirement-Age Assumption on Key Variables and Predictions

Figure 3.7 shows the growth in output, change in equilibrium interest rate, and change in the equilibrium labor tax rate necessary to finance pension outlays relative to the first extreme.

[10] Cohort distributions are given by the steady state distributions implied by each life expectancy at birth.

Clearly, the conclusions drawn are both quantitatively and qualitatively sensitive to assumptions regarding labor supply. In fact, output growth, the change in real interest rates, and the change in budget balancing tax rates may all be positive or negative depending on what assumptions are made. All results would change sign if time in retirement, instead of length of working lives, was held constant. In particular, interest rates would increase if retirement age increases by just six years when life expectancy increases by ten years.

This exercise highlights both the sensitivity of previous results to the assumption that effective retirement age remains constant while life expectancy increases, and it shows that a key margin necessary to understand the effects of demographic change is individual labor supply choice.

3.4 Calibration

We calibrate our benchmark model and pension systems separately for Italy, Germany, France, and the United Kingdom by fixing several preference and production parameters, estimating the age-cohort distributions, survival probabilities, and productivity parameters outside the model, and using simulated method of moments for the remaining parameters. Our calibration target year is 1995 due to the availability of data. The key margins of our model are the labor supply elasticities on both the intensive and extensive margin, particularly at old ages, and tax rate elasticities.

The simulated method of moments procedure minimizes the squared distance between the model implied moments and the associated moments in the data. To solve for the former, candidate parameter values are drawn. Given these parameter values, we draw candidate market clearing prices and solve for individual policy functions. We then simulate a panel of individual decisions and update accidental bequests within the model. We iterate between the solution of individual decision rules and the computation of accidental bequests until we obtain a fixed point. Using the simulated panel from our fixed point solution, we solve for excess capital and labor demand, and update the candidate market clearing prices accordingly. Once we have solved for market clearing prices, we use the associated simulated panel to compute the desired model moments. We continue drawing candidate parameter values until the squared difference between the model and data moments is minimized. We use the described simulated method of moments procedure to calibrate the discount factor, the utility weight on leisure, and the cost to participation function. Our included moments are each country's capital-output ratio, average hours worked, effective retirement age, labor force participation rate of those aged 60-64, and labor force participation rate of those aged 65-69. The rest of our parameters are fixed outside of this approach, i.e. externally calibrated.

3.4.1 Preference Parameters

We first set $\sigma = 1$ in order to obtain balanced growth preferences as in King et al. (1988) and fix the curvature on leisure to be $\gamma = 4$. The remaining preference parameters are the discount factor, β, household's weight on leisure, χ, and the cost to participation parameters, κ_1, κ_2, and κ_3. We set β to match the measured capital-output ratio for each country, which is calculated from the Penn World Tables 9.0 release through the FRED database. χ is targeted to a weighted average of hours worked per year by working households aged 20-64 from the OECD statistical database. The weights are given by the relative size of the workforce at each age. Because of the importance of labor-supply decisions on the extensive margin in driving our results, our goal is to tightly link our cost to participation parameters, $\{\kappa_1, \kappa_2, \kappa_3\}$, to retirement decisions at the end of life. To do so, we calibrate these parameters to match labor force participation rates of those aged 60-64 and 65-69, and the effective retirement age calculated as in Keese (2003).[11] Data for each is obtained from Eurostat and the OECD statistical database, respectively. Finally,

[11] The effective retirement age we use is a weighted labor market exit age starting at age 40, where the weights are the change in labor force participation from age i to $i + 1$.

we restrict participating households to work no less than 20% of their available time.

3.4.2 Technology and Productivity

Our production technology is Cobb-Douglas with a capital share of $\alpha = 0.33$. While standard measures of capital share show significant heterogeneity across economies, Gollin (2002) shows that capital shares are in fact quite stable in the cross section after controlling for self-employed income. Our chosen value reflects his mean estimate. We set the depreciation rate, δ, to match the 1995 real interest rate in each country, where we calculate the real interest rate as the return on 10 year long-term government bonds less inflation. Each is taken from the OECD and World Bank, respectively, through FRED. As in Hansen (1993), we estimate the life-cycle productivity profile and idiosyncratic productivity process from the PSID. [12] Figure 3.8 shows our estimate for ψ_i, and we find the persistence and variance of the idiosyncratic productivity process to be $\rho = 0.97$ and $\sigma^2 = 0.02$, respectively.

3.4.3 The Pension System

There are important differences in the public pension systems of the four European Economies we study. These are discussed extensively in Erosa et al. (2012) and in SHARE, the Survey of European Health and Retirement systems. For our purposes the important features of pension systems that largely drive changes in tax rates and retirement incentives are the eligibility age, I_R, and the level of real old age benefits, p. The former is taken from the Blondal and Scarpetta (1997) and Gruber and Wise (1999). In order to calibrate p_t, we match the level of pension expenditures in the form of non-means tested old age benefits by country as a percentage of GDP and assume that these old age benefits are evenly distributed among the eligible population. This data is obtained from Eurostat.

[12] A similar data set is not readily available to us for the four countries herein considered.

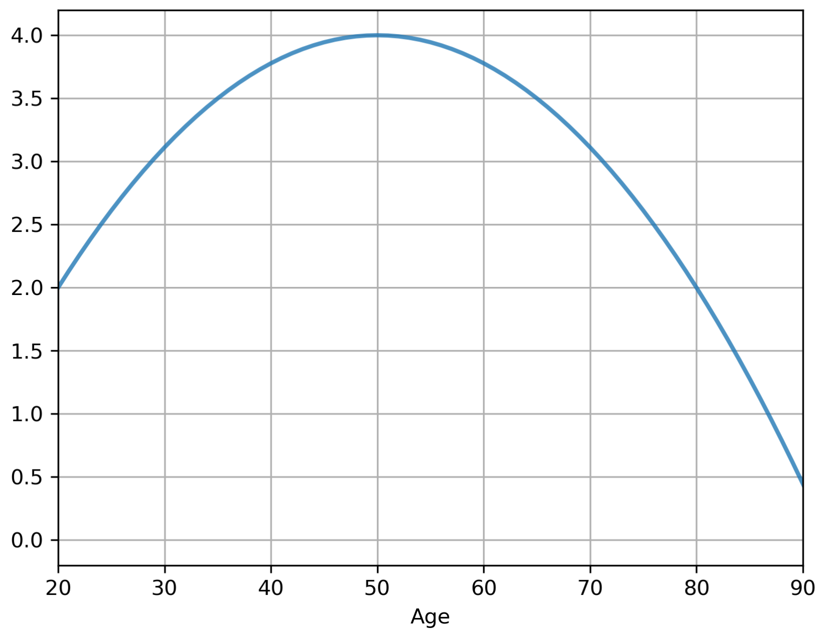

Figure 3.8: Life-Cycle Productivity Profile

We allow only one form of taxes to balance the government budget constraint at any given time. Consequently, we calibrate two models with pension systems: a model with only lump sum taxes and a model with labor income taxes. In our quantitative experiments, we fix p_t and allow the relevant tax rate to adjust to close the model. Finally, note that we fix $p_t = t_t = \tau_{L,t} = 0$ in our benchmark model without pension systems.

3.4.4 Demographics

What remains is to calibrate the survival probabilities and relative size of each age-cohort in each country. The cohort distributions are taken from the United Nations (2017), which gives the cohort distribution in 5 year age bins. We linearly interpolate between the center of each age bin and re-normalize the interpolated distribution to obtain 1 year cohort bins. The interpolated cohort distributions are shown in Figure 3.3.

The one-year survival probabilities are calculated as in Henriksen (2015) using life expectancy data obtained from the United Nations. These estimates are shown in Figure 3.9 and incorporate the fact that mortality rates are a function of both age and life expectancy at birth.

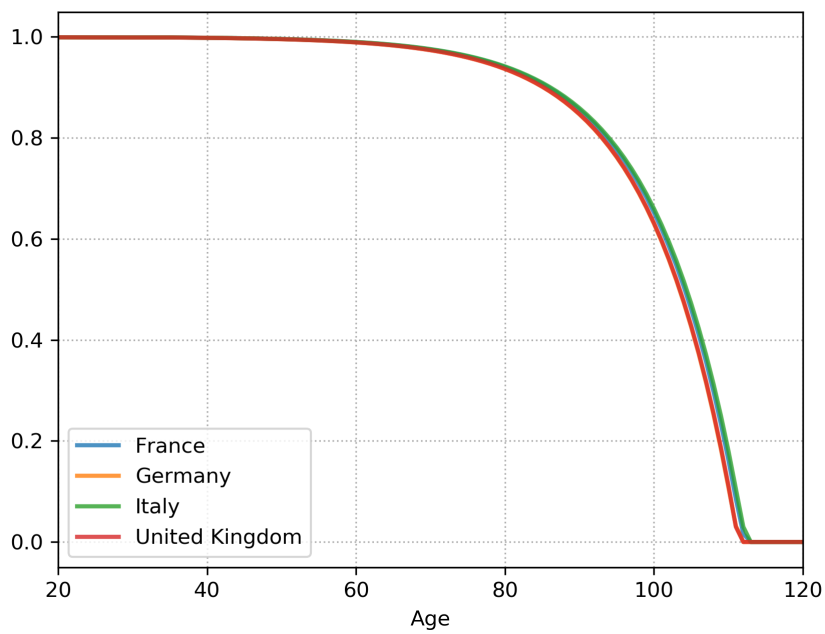

Figure 3.9: 1995 Conditional Survival Probability

The set of moments matched for each country are summarized in Table 3.2. Table 3.3 gives the corresponding model moments and Table 3.4 displays the calibrated parameter values from our simulated method of moments procedure. Of particular note are our calibrated cost to participation functions. For each model and each Country, $\{\kappa_1, \kappa_2, \kappa_3\}$ govern the dis-utility of working. They are calibrated to match observed labor force participation and are broadly consistent with French (2005) and Capatina (2015). Both estimate that the life-cycle probability of poor health is increasing and approximately convex in age with the latter further showing that individual time endowment decreases in expectation over the life-cycle.

Table 3.2: Summary of 1995 Moments by Country

Moment	France	Germany	Italy	U.K.
Labor-force participation rate for ages 60-64	11%	19%	19%	37.2%
Labor-force participation rate for ages 65-69	3.0%	4.5%	6.3%	11%
Avg. hours worked for ages 20-64	0.33	0.33	0.34	0.34
Effective retirement age	59.6	60.3	59.1	62
Real interest rate	5.75%	5.14%	6.96%	5.54%
Capital-to-output ratio	3.23	3.55	4.02	3.02
Pension outlays as a fraction of GDP	9.9%	7.4%	10.2%	7.8%

Table 3.3: Calibration Results

	$LFPR60-64$	$LFPR65-69$	$Retire.\ Age$	K/Y	$Avg.\ Hours\ Worked$
U.K.					
Data	0.37	0.110	62.0	3.02	0.34
Benchmark	0.43	0.046	61.8	3.02	0.32
Lump Sum	0.50	0.110	62.3	3.15	0.32
Labor Tax	0.50	0.110	62.3	3.02	0.32
France					
Data	0.11	0.030	59.6	3.23	0.33
Benchmark	0.28	0.033	60.0	3.27	0.29
Lump Sum	0.34	0.036	60.3	3.21	0.30
Labor Tax	0.29	0.030	59.8	3.27	0.29
Germany					
Data	0.19	0.045	60.3	3.55	0.33
Benchmark	0.38	0.065	61.1	3.65	0.28
Lump Sum	0.38	0.053	60.7	3.57	0.30
Labor Tax	0.36	0.047	60.6	3.43	0.29
Italy					
Data	0.19	0.063	59.1	4.02	0.34
Benchmark	0.28	0.050	59.7	4.16	0.29
Lump Sum	0.27	0.022	59.7	4.37	0.32
Labor Tax	0.27	0.033	59.6	4.16	0.30

Table 3.4: Calibrated Parameters

	β	δ	χ	κ_1	κ_2	κ_3	p
U.K.							
Benchmark	0.944	0.055	0.206	0.0505	0.00181	1.414	–
Lump Sum	0.958	0.055	0.274	0.0503	0.00192	1.331	0.592
Labor Tax	0.955	0.055	0.339	0.0028	0.00225	1.285	0.583
France							
Benchmark	0.942	0.046	0.355	0.0440	0.00203	1.355	–
Lump Sum	0.953	0.046	0.392	0.0370	0.00197	1.306	0.557
Labor Tax	0.957	0.046	0.392	0.0390	0.00266	1.207	0.540
Germany							
Benchmark	0.950	0.043	0.353	0.0495	0.00224	1.310	–
Lump Sum	0.958	0.043	0.332	0.0470	0.00196	1.323	0.602
Labor Tax	0.957	0.043	0.347	0.0420	0.00294	1.196	0.578
Italy							
Benchmark	0.937	0.013	0.350	0.0495	0.00215	1.316	–
Lump Sum	0.952	0.013	0.266	0.0412	0.00232	1.304	0.723
Labor Tax	0.951	0.013	0.415	0.0014	0.00243	1.243	0.673

3.5 Numerical Solution Method

The numerical solution to the model involves solving individuals' consumption–savings and labor–leisure choices, and aggregating those decisions. At any point in time, each individual knows her/his age and their three state variables: savings, and individual productivity. Individuals' consumption-savings and labor-leisure choices can be solved given the realization of the state variables and the current and expected future factor prices. Two crucial questions for any solution algorithm are (i) the source of the distribution of the assets at any point of time and (ii) the individuals' expectations of the path of future factor prices.

Broadly, there are two numerical approaches to modelling the distribution of state variables and individuals' expectations of the path of future factor prices: (1) compute a transition path between two steady states and (2) compute the steady states associated with the individual mortality rates and the aggregate cohort distribution at each point of time.

3.5.1 Transition between two steady states

One approach is to simulate the demographic transition between two steady states.[13] With this approach, the initial steady state distribution of assets is the distribution associated with the initial cohort distribution and the life expectancy at that point in time, while the final steady state is the asset distribution associated with the stationary long-run population distribution. The stationary long-run population distribution is the stationary distribution associated with projected fertility, mortality and immigration rates at the furthest end of where demographers make projections.

When an individual of a given age at a given time makes her or his consumption–savings and labor–leisure decisions, their state variables are their own assets holdings and own idiosyncratic productivity level. They know their conditional survival probabilities so they know with what probability they are going to survive to any given age. They can compute the probability distribution over future idiosyncratic productivity levels, and they can perfectly foresee the entire future path of endogenous factor prices generated by the model.

In other words, both the sequence of factor-supply distributions and the factor-price paths are functions of the initial and terminal distributions. The factor prices clear the markets for labor and capital so they are functions of the asset and labor-supply distributions at each period of time. Since individual asset holdings is an individual state variable, the choices made by anyone alive at the initial state will directly be a function of the initial asset distribution. Through the direct and indirect impact on the factor price paths, all decisions made will be a functions of the initial and terminal steady states.

[13] See e.g. Ríos-Rull (1999) or Backus et al. (2014).

3.5.2 Succession of steady states

An alternative is to compute the steady states associated with the individual mortality rates and the aggregate cohort distribution at each period of time instead of to solving the model as a transition between two steady states. The distribution of assets at each period of time will be the distribution associated with the cohort distribution and life expectancy at that point in time, and the factor prices and the distribution will be consistent with the choices individuals are making in that particular period.

Individuals will still solve their problem given the same set of state variables. Two key differences are that they use the current one-year-younger cohort's asset-holding choice as an approximation for the asset-holding choice that they made the previous year and that they implicitly assume that the current factor prices are the best forecast for future factor prices.

There are three main benefits of the latter numerical approach. First, using the distribution of state variables associated with the current period is no less arbitrary than using a distribution, which is a function of the distribution associated with the population distribution at some initial date chosen by the researcher, eg. 1950 or 1970. Second, it is more reasonable to assume that individuals within the model base their decisions on the conjecture that current factor prices are the best forecasts for future factor prices, rather than assuming they can perfectly foresee the entire deterministic future path. In particular, there remains uncertainty associated both with the future demographic path and the mechanisms for the formation of future factor prices. Thirdly, this latter approach is more parsimonious and transparent.

3.6 Quantitative Results

We first estimate the contribution of demographic change to the historical growth experience. We then look at the implications of projected demographic change for future growth. In the benchmark economy, households save for their own retirement and the economy is closed with no capital flows between countries. We gradually relax these assumptions and look at the role pension systems and capital flows play in the growth outcomes of these economies. For the questions in this paper, pension systems are important and capital flows are not. Finally, we examine some consequences of two types of pension reforms – lower benefits and later retirement ages.

In these quantitative exercises, we fix the calibrated parameters of the model and adjust the conditional survival probabilities and cohort distributions to match those in 1975, in 1995, and in 2015. We then perform growth decompositions from 1975-1995 and 1995-2015 by applying the growth accounting methodology in Eq. 3.2 to the model steady states. Finally, we make growth projections by repeating the above exercise for 2020 and 2040.

3.6.1 Direct Effect of Demographics

We first decompose the per-capita growth effects of demographic change using the benchmark model. Table 3.5 displays a summary these effects in the benchmark model. Table 3.6 and Table 3.7 show the contribution of demographic change to growth over the 1975-1995 and 1995-2015 periods in greater detail. It highlights the changing contribution of ageing populations to growth relative to the growth slowdown discussed in Section 3.2.

Table 3.5: Benchmark Model Historical Annualized Growth Summary

	1975-1995	1995-2014	Change	Percent of Slowdown
United Kingdom	0.39%	0.06%	-0.33	51%
Germany	0.60%	-0.08%	-0.68	67%
France	0.47%	-0.16%	-0.63	77%
Italy	0.57%	-0.02%	-0.59	28%

Demographic change boosted growth during the period from 1975-1995 as the post-war

Table 3.6: 1975-1995 Annualized Model Growth Rates

	$\gamma_{Y/pop}$	γ_A	$\alpha \cdot \gamma_{K/L}$	$\gamma_{L/pop}$	$(1-\alpha) \cdot \gamma_h$
United Kingdom					
Benchmark	0.42	0.02	0.05	0.28	0.05
Lump Sum	0.42	0.02	0.05	0.28	0.07
Labor Tax	0.40	0.02	0.05	0.29	0.05
France					
Benchmark	0.47	0.06	0.09	0.33	-0.01
Lump Sum	0.45	0.05	0.07	0.33	0.01
Labor Tax	0.46	0.06	0.08	0.33	-0.01
Germany					
Benchmark	0.60	0.02	0.07	0.52	-0.02
Lump Sum	0.57	0.03	0.06	0.50	-0.01
Labor Tax	0.56	0.04	0.06	0.48	-0.01
Italy					
Benchmark	0.57	0.01	0.10	0.45	0.01
Lump Sum	0.57	-0.01	0.08	0.46	0.05
Labor Tax	0.51	0.01	0.06	0.43	0.02

generations entered their most productive years. After 1995 their contributions to growth declined throughout the following two decades in all four countries Our model indicates that the changing historical contribution of demographic change to growth is responsible for a decline in annual per-capita growth of 0.33-0.68 percentage points between these two periods. Relative to our growth accounting exercise in Section 3.2, our benchmark results suggest that ageing populations account for roughly 77% of the secular growth slowdown in France, 67% in Germany, 51% in the United Kingdom, and 28% in Italy.

To explore how aging populations effect growth, we look at the past 20 years. Table 3.7 shows that demographic change was a drag on per-capita growth for France, Germany, and Italy while it contributed positively in the United Kingdom. The primary drag on growth comes through changes in labor supply on the extensive and intensive margins and, to smaller degree, through capital accumulation. Indeed, the combination of increases in life expectancy and rightward shifts in the age-cohort distribution leads to capital deepening. These increases in the

Table 3.7: 1995-2015 Annualized Model Growth Rates

	$\gamma_{Y/pop}$	γ_A	$\alpha \cdot \gamma_{K/L}$	$\gamma_{L/pop}$	$(1-\alpha) \cdot \gamma_h$
United Kingdom					
Benchmark	0.06	0.03	0.08	0.03	-0.09
Lump Sum	0.06	0.03	0.07	0.03	-0.07
Labor Tax	0.01	0.05	0.04	-0.01	-0.08
France					
Benchmark	-0.16	0.04	0.09	-0.18	-0.11
Lump Sum	-0.15	0.04	0.07	-0.19	-0.07
Labor Tax	-0.24	0.06	0.05	-0.26	-0.10
Germany					
Benchmark	-0.08	0.05	0.10	-0.11	-0.12
Lump Sum	-0.11	0.04	0.07	-0.13	-0.09
Labor Tax	-0.19	0.06	0.05	-0.18	-0.12
Italy					
Benchmark	-0.02	0.12	0.15	-0.15	-0.15
Lump Sum	-0.11	0.12	0.11	-0.20	-0.14
Labor Tax	-0.18	0.14	0.09	-0.25	-0.17

aggregate capital stock result in significant declines in equilibrium interest rates ranging from 81 basis points in the case of Italy to 43 basis points in United Kingdom. r^* changes between 1975 and 2015: UK: 43 basis points, IT: 81 basis points, GE: 75 basis points, and FR: 75 basis points.

Labor supply is affected along both the intensive and extensive margins. As the work force ages, the disutility of working also rises. The overall effect is that hours worked decreases.

Changes in the employment-population ratio have the largest effect on growth. Over the past 20 years, aging has caused substantial declines in the labor force participation rate. As in the case of labor hours, rightward shifts in the cohort distribution have resulted in a larger fraction of the economically active population being in the right tail of the age-cohort distribution. Older households face strong incentives to retire due to both a higher cost to participating in the labor force and declines in life-cycle productivity at old ages. Conversely, gains to life expectancy incentivize households to increase labor supply at old ages. This endogenous response works to mitigate declines in labor supply as age cohort distributions shift. For example, between 1995 and 2015, the effective retirement age in each country increased by (yrs): UK: 0.93, IT: 1.58, GE: 1.23, and FR: 1.21. In comparison, life expectancy at birth over this same time period increased by an average of 4.5 years across these economies. As discussed above, shifts in the age cohort distribution outweighed increases in effective retirement ages. The net effect is declining labor force participation at older years.

Demographic change also affected measured total factor productivity. Due to our hump shaped life-cycle productivity profile, young populations that age begin to have a larger share of their workforce in the most productive periods of their lives. Thus, aging populations affect measured TFP mechanically through shifts in the age-cohort distribution.

Measured TFP is also affected by an endogenous mechanism briefly discussed in Section 3.3. Because of persistent idiosyncratic productivity shocks and the rising cost of participation, there are strong self selection effects with respect to labor force participation. In particular, those who have received a series of high productivity draws early in life are more likely to remain in the labor force at older ages.

Strong general equilibrium effects further affect all of these margins. As relative prices change, so too do the incentives to accumulate capital and supply labor. Changes to the latter are of particular importance given our results. For example, changes in the wage rate directly affect the level of individual productivity required to remain in the workforce at old ages. Labor force participation rates and the composition of workers are in turn affected. In this case, the negative labor supply effects described above are partially offset by rises in the wage rate induced by capital deepening.

Over the next two decades, our model predicts that the ongoing demographic change will depress per-capita growth further. Indeed, Table 3.8 shows that decreases in the employment-population ratio are projected to become more prominent while capital deepening will decelerate. The benchmark model indicates that Germany and Italy will face the most significant

Table 3.8: 2020-2040 Annualized Model Growth Rates

	$\gamma_{Y/pop}$	γ_A	$\alpha \cdot \gamma_{K/L}$	$\gamma_{L/pop}$	$(1-\alpha) \cdot \gamma_h$
United Kingdom					
Benchmark	-0.24	-0.01	0.05	-0.28	-0.01
Lump Sum	-0.26	-0.03	0.02	-0.28	0.04
Labor Tax	-0.35	-0.01	-0.01	-0.34	0.01
France					
Benchmark	-0.25	-0.01	0.04	-0.32	0.03
Lump Sum	-0.22	-0.03	0.04	-0.31	0.07
Labor Tax	-0.37	-0.02	-0.01	-0.38	0.04
Germany					
Benchmark	-0.52	0.02	0.08	-0.61	-0.01
Lump Sum	-0.48	-0.01	0.06	-0.59	0.06
Labor Tax	-0.71	0.01	-0.01	-0.72	0.01
Italy					
Benchmark	-0.67	-0.05	0.08	-0.76	0.05
Lump Sum	-0.64	-0.09	0.06	-0.79	0.16
Labor Tax	-1.13	-0.05	-0.07	-1.08	0.08

declines to future growth.

3.6.2 Decomposing the Effect of Life Expectancy and Cohort Distributions

Our numerical approach further allows us to decompose the effect of demographics on growth into the effect of the behavioral changes due to life-expectancy gains and the effect of changes in the aggregation of decisions due to shifts in the cohort distribution. To illustrate, in the period 1975-1995 demographics made a positive contribution to growth for France. The benchmark results indicated that demographics gave an annual average contribution to per-capita GDP growth of 0.47 percentage points. Just changing the cohort distribution, while holding individual life expectancy fixed, yielded an average annual contribution to per-capita GDP growth of 0.38 percentage points. The estimated contribution to per-capita annual growth when keeping the cohort distribution fixed and only changing life expectancy is 0.12

percentage points. Implicitly, the general-equilibrium factor-price effect was slightly negative when changing both life expectancy and cohort distribution.

For the following two decades, 1995-2015, the benchmark model estimated that the demographic contribution to per-capita GDP growth in France was -0.16 percentage points per year. Decomposing this into the effect coming from changes in life expectancy and shifts in the cohort distribution, we find that changes in life expectancy was providing a positive contribution to growth of 0.09 percentage points per year, while the contribution from the shift in cohort distribution was -0.30 percentage points per year. Implicitly, the general-equilibrium factor-price effect was slightly negative for this period as well.

This is a general pattern across countries and time periods. While the contribution to economic per-capita growth from the shifts in cohort distribution turns negative as the average age increases, the contribution from realized and projected gains to longevity stays positive. Longer life expectancy provides incentives for more savings and increased labor supply both on the intensive and the extensive margin.

3.6.3 The Effect of Pension Systems

Next, we compare the growth accounting exercise discussed in the previous section to a model with pensions that are funded with lump-sum taxes. Because lump-sum taxes are non-distortionary in nature, this comparison allows us to estimate how old-age transfers alone impact growth. We subsequently describe our results when pensions are funded by a labor income tax, allowing us to estimate the growth effects of increasing distortions that result from pension systems.

Quantitatively, the provision of old-age transfers themselves create few declines in growth. While small, these effects are most strong in Italy. Old-age benefits provide elderly households with additional resources from which to consume. Thus, retired households may rely on social

Table 3.9: Outlays as a Fraction of GDP

	U.K.	France	Germany	Italy
1975				
Lump Sum	7.10%	9.48%	7.58%	8.43%
Labor Tax	7.04%	9.44%	7.63%	8.37%
1995				
Lump Sum	7.85%	9.99%	7.39%	10.18%
Labor Tax	7.81%	9.95%	7.45%	10.24%
2015				
Lump Sum	8.81%	12.79%	10.26%	13.51%
Labor Tax	8.86%	12.95%	10.52%	13.76%
2020				
Lump Sum	9.41%	13.88%	11.05%	14.86%
Labor Tax	9.48%	14.24%	11.38%	15.31%
2040				
Lump Sum	12.51%	16.99%	16.15%	22.35%
Labor Tax	12.84%	17.93%	17.35%	24.49%

security in addition to individual savings to smooth consumption throughout retirement. This mitigates incentives to increase savings and labor supply that would result from increases in life expectancy.

The distortions from these pension outlays have important implications for growth. As populations age, fiscal authorities face increasing liabilities in the form of social security payments. Table 3.9 show just how sharply pension outlays increase in response to aging in our model. Given the declines in labor supply identified in Section 3.6.1, tax rates must increase to balance budgets. The rise in labor tax rates in turn provides an additional disincentive to work, further decreasing labor supply and compounding the problem faced by fiscal authorities.[14]

The quantitative effect of these distortions are shown in Table 3.7 and Table 3.8 for the 1995-2015 and 2020-2040 periods, respectively, and in Table 3.6 for the 1975-2015 period. These distortions were quantitatively unimportant during the 1975-1995 period. Because a large number of previously inactive households began entering the labor force during this period (see

[14] Conesa et al. (2019) suggest that trends in college attainment may mitigate the need to increase tax rates.

Figure 3.3), increases in labor supply provided sufficient tax revenue to offset any need to increase tax rates. Thus, pension systems did not greatly impact growth during 1975-1995.

Over the past two decades, changing demographics imply that there were fewer new workers and more workers in retirement. As a result, tax rates in our model must rise to cope with changing pension obligations. These additional distortions amplify the rise in retirement rates relative to the benchmark and lump sum models. Between 1995 and 2015, effective retirement ages increased by fewer years than in the benchmark model: UK: 0.48, IT: 0.79, GE: 0.71, and FR: 0.64. This endogenous response of old-age labor supply is roughly half of that in the benchmark model. The end result is that annualized growth in the employment/population ratio decreased by an additional 10 basis points in France, Germany, and Italy as a direct result of changing tax rates. The effect in the United Kingdom is smaller, at around 4 basis points.

Higher tax rates also affect capital deepening. Not only does the duration of households' working lives decrease relative to a world without such distortions, but so too does after tax labor income conditional on working. Both features directly effect the resources from which households save. The end result is a decline in capital accumulation. Moreover, pension systems dampen the decline in interest rates previously discussed through this channel. In this case, declines in the equilibrium interest rate instead range from 31 basis points to 50 basis points.

While the effects are smaller, these changing distortions also affect measured TFP growth. As labor taxes rise, the individual productivity level needed to remain in the workforce during old age does as well. Thus, the workforce that remains following a rise in labor taxes is more productive. This amplifies the self-selection mechanism in the model.

All of these effects contributed to a growth decline resulting from pension systems. In total, changing tax rates decreased annual growth by between -0.04 percentage points in the case of the United Kingdom and -0.18 percentage points in Italy throughout the 1995-2015 period. Our model further predicts that the growth effects of these distortions will become bigger over the next 20 years.

3.6.4 Consumption and labor-supply dispersion

Idiosyncratic labor productivity shocks were included in the model to give separate predictions for labor supply on the intensive and extensive margin. In order to check whether the model-generated labor supply behavior is empirically plausible we compute the evolution of the cross-sectional variance of consumption and hours worked over the life cycle. Qualitatively, our results are similar to those reported by Kaplan (2012); the cross-sectional variance of consumption increases gradually over the life-cycle while the cross-sectional variance of hours worked is almost flat until age 50 when it increases slightly.

3.6.5 Passing away with positive debt

With a positive probability of dying and a borrowing constraint that permits positive debt, in equilibrium some households will die with negative net worth, ie. debt. This happens in the data too, and is not considered a "default". Quantitatively, the magnitude of negative bequests in our model is 1.02% of total accidental bequests.

3.6.6 International Capital Flows

In the previous sections, we have assumed that each country a closed economy. Economists have long understood, however, that capital tends to flow between countries to equalize return differences. Heterogeneous demographic change of the kind herein considered has also been studied as an explanation for capital flows between countries (see eg. Backus et al., 2014). To check the robustness of our results, we extend our benchmark model to allow for the movement of capital across borders. Throughout, we maintain our calibration from Section 3.4

Table 3.10: Annualized Model Growth Rates with Capital Flows

	$\gamma_{Y/pop}$	γ_A	$\alpha \cdot \gamma_{K/L}$	$\gamma_{L/pop}$	$(1-\alpha) \cdot \gamma_h$
United Kingdom					
1975-1995	0.49	0.04	0.09	0.32	0.04
1995-2015	0.07	0.03	0.09	0.04	-0.09
2020-2040	-0.28	-0.01	0.04	-0.31	-0.01
France					
1975-1995	0.45	0.06	0.09	0.32	-0.01
1995-2015	-0.14	0.04	0.09	-0.17	-0.10
2020-2040	-0.17	-0.01	0.07	-0.27	0.02
Germany					
1975-1995	0.54	0.02	0.07	0.46	0.01
1995-2015	-0.08	0.06	0.09	-0.10	-0.13
2020-2040	-0.49	0.03	0.07	-0.59	-0.01
Italy					
1975*-1995*	0.54	0.01	0.09	0.45	0.01
1995*-2015*	0.03	0.12	0.17	-0.14	-0.13
2020*-2040*	-0.66	-0.06	0.08	-0.71	0.03

The results of this exercise are presented below in Table 3.10. We place an asterisks next to country-year pairs in which that particular country was a capital importer. All other country-year pairs are capital exporters. A striking feature of Table 3.10 is that, while the introduction of international capital flows induces some quantitative differences with our central findings, it's impact is very small. Notably, there are no qualitative differences with our main results. In particular, we still find that the contribution of ageing populations to growth has declined historically, and will decline further over the next two decades. Here again, the margins most affected by aging populations are those relating to labor supply.

3.7 Policy Reforms

Given the sizable labor supply effects of aging populations, we now consider why individuals exit the labor force when they do, and discuss some policy reforms that could alter those choices.

We stress the welfare effects of policies as measured by the median consumption equivalence of a newborn agent as a measuring stick. Our definition of consumption equivalence is given by

$$\mathbb{E}_{i_0} \sum_{t=i_0}^{I} \beta^t s_t u(\lambda_j c^*_{t,j}, h^*_{t,j}) = \mathbb{E}_{i_0} \sum_{t=i_0}^{I} \beta^t s_t u(c^{**}_{t,j}, h^{**}_{t,j}) \tag{3.16}$$

where λ_j is the consumption equivalence for household j, $\{c^*_{t,j}, h^*_{t,j}\}$ are their pre-reform optimal choices, and $\{c^{**}_{t,j}, h^{**}_{t,j}\}$ are their post reform optimal decisions. Straightforward algebra allows for a simple expression of the consumption equivalence for individual j defined by

$$\lambda_j = e^{\left[V_{j,post} - V_{j,pre}\right]/\phi} \tag{3.17}$$

where $\phi = \sum_{t=i_0}^{I} \beta^t s_t$, $V_{j,post}$ is the post-reform value function for newborn j, and $V_{j,pre}$ is the pre-reform value function for newborn j.

We first assess the role of pension systems in driving decreases in the employment-population ratio. Decreasing the level of old age benefits forces households to rely more heavily on individual savings during retirement. Households may therefore increase labor supply to finance savings increases. Raising the eligibility age incentivizes a longer working life directly. Because households must wait longer to receive any social security payments, households may increase their labor market exit age to avoid a period of low consumption between retirement and social security eligibility. Furthermore, both reforms decrease the fiscal authority's total pension obligations, reducing the degree to which tax rates rise to balance budgets. We consider these reforms by separately reducing per-retiree pension outlays by 5%-20%, and increasing the pension eligibility age by 5 years by the year 2040.[15]

Table 3.11 shows a summary of the effects of reducing the level of pension benefits received when in retirement. Evidently, such a reform mitigates the projected adverse growth effects

[15]We discuss these reforms for the labor tax model only. There are no distortions in the lump-sum taxation model.

of aging populations. Table 3.12 shows the post-reform contribution to 2020-2040 annualized growth of demographic change in more detail. We display only the 20% reduction in pension benefits for conciseness. The need to rely more heavily on personal savings to finance old age consumption causes savings rates to become more sensitive to changes in life expectancy. Second, the indirect effect of pension systems on individual decisions through tax rates becomes weaker. The reduction in benefits reduces the need to raise tax rates, which in turn provides fewer disincentives to work. Relative to the pre-reform economy, the effective retirement age of households can increase substantially depending on the extent of the reform. Table 3.12 further shows that the labor supply channel is again much more important than individual savings.

More importantly, these reforms are welfare improving. The welfare effects of reductions in pension benefits may be significant as distortions to individual incentives begin disappearing depending on the extent of the reform. In France, Germany, and Italy, the median per-period consumption equivalent for newborns quickly rises over 1% as the extent of the reform increases. For significant benefits reductions, median per-period consumption equivalents range from 2.38%-4.16%. The United Kingdom sees more modest welfare gains with consumption equivalence, reaching a peak of 1.59%.

Except for conservative reductions in pension benefits, Table 3.13 shows that increasing the pension eligibility age yields similar increases in annual growth. Increasing the eligibility age operates through similar channels as reductions in pension benefits. Table 3.12 displays the effects of eligibility reform in detail. First, as the eligibility age increases, social security does not become available to households until a later age. This implies that agents must rely more heavily on labor income and private savings for these additional years before they become eligible for pension benefits. Second, as less of the population is eligible for any given level of benefits, government outlays and thus tax rates again increase by less as populations age. These two effects serve again to increase capital accumulation as life expectancy increases and mitigate the reduction in the employment to population ratio as populations age. In the case of

Table 3.11: Projected Effect of x% Decrease in Benefits in Labor Tax Model

	$\Delta\gamma_{Y/pop}$ (pp)	Δ $Retirement\ \ Age$ (yrs.)	CE (%)
UK			
5%	0.03	0.17	0.40
10%	0.07	0.43	0.80
15%	0.11	0.64	1.17
20%	0.15	0.88	1.55
France			
5%	0.09	0.46	0.63
10%	0.14	0.72	1.21
15%	0.18	1.03	1.78
20%	0.27	1.49	2.33
Germany			
5%	0.04	0.24	0.63
10%	0.13	0.59	1.32
15%	0.17	0.80	1.91
20%	0.22	1.01	2.48
Italy			
5%	0.12	0.42	1.11
10%	0.25	0.84	2.13
15%	0.37	1.29	3.11
20%	0.43	1.55	3.93

Table 3.12: Post-Reform Contribution to 2020-2040 Annualized Growth Rates in Labor Tax Model

	$\gamma_{Y/pop}$	γ_A	$\alpha \cdot \gamma_{K/L}$	$\gamma_{L/pop}$	$(1-\alpha) \cdot \gamma_h$
20% Decrease in p					
France	−0.10	−0.01	0.08	−0.20	0.01
Germany	−0.49	0.02	0.07	−0.58	−0.01
Italy	−0.70	−0.04	0.07	−0.80	0.06
United Kingdom	−0.20	−0.01	0.05	−0.23	−0.01
5 Year Increase in I_R					
France	−0.15	−0.02	0.05	−0.22	0.04
Germany	−0.45	0.01	0.07	−0.53	0.01
Italy	−0.77	−0.06	0.03	−0.82	0.08
United Kingdom	−0.18	−0.02	0.05	−0.21	−0.01
3 Year Shift in θ					
France	0.05	−0.01	0.08	−0.14	0.11
Germany	−0.35	0.03	0.04	−0.51	0.08
Italy	−0.67	−0.02	−0.01	−0.78	0.13
United Kingdom	−0.06	0.01	0.02	−0.17	0.08

increasing eligibility ages, increases in both growth and the effective retirement age are larger than for most reductions in pension benefits. Despite this fact, the welfare gains from reductions in pension benefits are larger than increases in the pension eligibility ages, a direct result of sharp increases in the dis-utility of labor experienced by households at old ages.

Our results suggest that pension reforms can induce non-trivial changes in retirement behavior, and therefore aggregate growth, but other important incentives for late life labor supply decisions remain. Two components emphasized previously are health and fitness, which are captured by our model's cost of participation function. Gains to morbidity, for example, may impact the dis-utility of work and thus shift our cost of participation function. While our morbidity data in Section 3.2 do not suggest large, widespread gains to old age health historically, the available data is quite limited. Moreover, it may be the case that the countries considered here see a “take-off" in terms of old age health in the coming years. We estimate

Table 3.13: Contributions to 2020-2040 Annualized Growth Rates for a 5-Year Increase Eligibility Age

	$\Delta\gamma_{Y/pop}$ (pp)	Δ *Retirement Age* (yrs.)	*CE* (%)
Labor Tax			
France	0.22	1.29	1.46
Germany	0.26	1.31	2.11
Italy	0.36	1.43	2.99
United Kingdom	0.17	1.08	1.39

the role of health and fitness in driving our labor supply results by assuming that future gains to life expectancy are accompanied by equivalent rightward shifts in our cost of participation function. Implementing this experiment requires shifting our cost of participation function by three years between 2020 and 2040, approximately the same projected gain in life expectancy.

Table 3.14: 2020-2040 Annualized Growth Rates for a 3 Year Shift in θ

	$\Delta\gamma_{Y/pop}$ (pp)	Δ *Retirement Age* (yrs.)	*CE* (%)
Labor Tax			
France	0.31	1.47	2.00
Germany	0.26	1.53	2.18
Italy	0.46	1.71	2.39
United Kingdom	0.29	1.45	1.73

Table 3.14 shows that in all four countries, the gains to projected growth and increases in the effective retirement age outpace that of any pension reform considered. Table 3.12 shows that the additional gains to growth are indeed a result of increased labor force participation rates and longer average hours worked. These results further highlight the importance of understanding labor supply decisions and life-cycle (dis-)incentives to work for estimating the effects of demographic change. The gains to growth are of course accompanied by welfare

gains as we have decreased our cost of participation function at each age. Still, these welfare gains are due in part to decreased labor tax distortions resulting from higher old age labor force participation. The per-period consumption equivalent welfare gains from this experiment are greater in each country than every pension reform considered except the largest cuts to pension benefits.

3.8 Conclusion

In this paper we have estimated the extent to which demographic trends affect economic growth in the four largest European Economies: France, Germany, Italy, and the United Kingdom. We use a general equilibrium model with a rich demographic structure and funded pension systems to estimate and project the impact of ageing on economic growth through changes in factor supplies. We find that the demographic transition can account for a significant fraction of the historical growth slowdown in these economies, and that evolving demographics will continue to drag down growth over the next 20 years.

Our model framework shows that the large gains to life expectancy are economically more important than declining fertility. Increases to longevity change individuals' savings and labor supply decisions over the life cycle. Changes in longevity are also quantitatively more important than fertility for both the number and the proportion of the population in advanced ages where they have accumulated the most assets, where labor market decisions change, and where retirement is spent.

There has been a lot of research focused on how demographics may change the equilibrium in the market for capital, assuming that labor is supplied inelastically and that gains to longevity do not change the length of working lives. In contrast, our results show that the most important margin through which demographics will affect future growth depends on how gains to individual life expectancy changes labor-supply decisions on both the intensive and extensive margin, and, how those decisions shape savings decisions.

With our baseline calibration, increases in life expectancy translate mostly into longer individual time in retirement. Increases in life expectancy will also, barring offsetting increases in fertility and immigration, imply that the average age will increase. An indirect consequence of individuals choosing to spend more time in retirement and an increase in the number of individuals in the age groups that choose to reduce their labor supply or retire, is that equilibrium tax rates must rise sharply to balance budgets. These higher taxes further distort the participation decision of households.

An important unresolved challenge is to better understand late-life labor supply and retirement decisions as life expectancy changes. If retirement decisions are mainly due to incentives from government programs, there may be both large output and welfare gains from reforming those systems. If instead households have strong intrinsic preferences against working at old ages, policy reforms such as increases to the threshold retirement age, may turn out to be welfare decreasing.

Appendix A

Household-Level Regression Analysis of Migration

A.1 Intercounty Moves

We use the household level data from the March CPS to show the effect that both channels have on the migration decisions of households. First, we show that dual searcher households are less likely to move across county lines and among all households that move, dual searcher households are less likely to move for job related reasons. Second, we show that dual searcher households that live in states with higher relative earnings of women are less likely to move and show suggestive evidence that among those households that do move, those that face higher relative earnings of women are less likely to move for job related reasons.

Since the primary reason for moving was only asked post 1999, we restrict the data from 1999 through 2015. Although this is not our primary time period of interest we use this time period to test cross-sectionally the hypothesis that an increase in dual searching households and in the relative earnings of women is correlated with a decrease in intercounty mobility.

We restrict the sample of households to civilian households in which the head of household is between the ages of 16 and 65. We create 3 samples of households: (1) all households in which the head of household is married (TOT), (2) all households in which the head of household is married and the spouse is present in the home (LT), and (3) all households in the TOT sample plus households that include unmarried partners (COH). Our variables of interest are the labor force status of both spouses. Thus, we divide the households into 2 subgroups: (1) both spouses are in the labor force (Dual) and (2) one spouse is in the labor force and the other is not (Single).

Table A.1: Summary Statistics: Married Households

	Dual LT	Single LT	Dual TOT	Single TOT	Dual COH	Single COH
County Move	0.03	0.04	0.03	0.04	0.04	0.04
Real Family Income	83348.05	63294.73	83256.44	62064.38	80620.84	61534.13
Own Home	0.84	0.75	0.84	0.73	0.82	0.73
Head of Household Characteristics						
Age	43.09	44.05	43.09	43.87	42.69	43.76
White	0.85	0.84	0.85	0.83	0.85	0.84
Black	0.07	0.07	0.07	0.07	0.07	0.07
One race - Other	0.06	0.08	0.06	0.08	0.06	0.08
Multiple races	0.01	0.01	0.01	0.01	0.01	0.01
Less than High School	0.03	0.07	0.03	0.07	0.03	0.07
High School	0.32	0.39	0.32	0.39	0.33	0.40
Some College	0.29	0.24	0.29	0.24	0.29	0.25
College	0.24	0.19	0.24	0.19	0.23	0.19
Advanced Degree	0.13	0.10	0.13	0.10	0.13	0.10
Spouse Characteristics						
Age	43.20	44.46	43.19	44.45	42.76	44.12
White	0.85	0.84	0.85	0.79	0.85	0.84
Black	0.07	0.07	0.07	0.06	0.07	0.07
One race - Other	0.06	0.08	0.06	0.07	0.06	0.08
Multiple races	0.01	0.01	0.01	0.01	0.01	0.01
Less than High School	0.03	0.07	0.03	0.07	0.03	0.07
High School	0.32	0.39	0.32	0.37	0.33	0.40
Some College	0.28	0.24	0.28	0.23	0.29	0.24
College	0.24	0.19	0.24	0.18	0.23	0.18
Advanced Degree	0.13	0.10	0.13	0.15	0.13	0.10
Observations	786,805	362,946	788,243	385,329	832,232	378,948

Table A.1 gives summary statistics for the characteristics of the head of household and their spouse, real family income, and home ownership rates for each subsample. Households across the samples differ in several respects. First, dual searcher households tend to have higher homeownership rates than single searcher households. Second, single searcher households tend to have a lower real family income. Households also differ slightly in educational attainment. In particular, dual searcher households tend to be slightly more educated than single searcher households. Both dual and single searcher households are roughly the same age and demographic makeup. Here we use all intercounty moves, below, section A.2 shows all estimates using only interstate moves.

To test the hypothesis that dual searcher households are less likely to move across county lines we estimate a probit model with the following specification:

$$\mathrm{P}(\mathrm{move}_i = 1) = \Phi(\beta_0 + \beta_1 dual_i + X_i\gamma + \eta_t + \varepsilon_i) \tag{A.1}$$

where dual_i is an indicator for both spouses in household i being in the labor force, X_i is a set of household covariates that include: age, age squared, race, education for both the head of household and the spouse, an indicator for homeownership, real total family income, and an indicator that takes on the value 1 if a child is present in the home. η_t is a year fixed effect and $\Phi(\cdot)$ is the c.d.f. of the normal distribution. We define single searching households to be our base case so that our coefficient of interest is β_1, which we expect to be negative.

Table A.2 gives the estimated coefficients on the labor market indicators. The sign of β_1 indicates that the probability of moving when both spouses are in the labor force is in fact less than that when only one spouse is in the labor force. Table A.3 gives the marginal effects of the labor market indicator for a household in which both spouses are white, 40 years old, have a college degree, own a home, and have a child present in the home in the year 2000. Focusing on the "Total" column of Table A.3 shows that the probability of moving when both spouses are in the labor force is 0.399 percentage points lower than when only one spouse is in the labor force. Given that the average probability of moving across county lines across the entire sample of married households is 3.1%, this effect is quite large.

Table A.2: Probit Estimation Results: Dual Searching

	Total	Living Together	Cohab
dual	-0.0521***	-0.0510***	-0.0255**
	(0.00865)	(0.00867)	(0.00799)
N	508,371	507,457	561,633

Robust Standard errors in parentheses
$^{*}p < 0.05$, $^{**}p < 0.01$, *** $p < 0.001$

Table A.3: Probit Marginal Effects: Dual Searching

	Total	Living Together	Cohab
dual	-0.00399***	-0.00389***	-0.00202**
	(0.000682)	(0.000680)	(0.000639)
N	508,371	507,457	561,633

Marginal effects evaluated for a household in which both spouse are white, 40 years old, with a college degree, own a home and have a child present in the home in the year 2000. Robust standard errors in parentheses.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Next we use the reason for moving response to test if households with both spouses in the labor force are less likely to move for job related reasons. There are 19 categories for the reason for moving variable which we regroup into 4 broader categories. Our main category of interest is “New job or transfer,” which includes only households that indicated that a new job or job transfer was their primary reason for moving. Our second category is “Other job reasons”, which includes all households that indicated they moved to look for work or lost job, for an easier commute, retired, or other job-related reason as their primary reason for moving. Our third category is “Family,” which includes all households that indicated a change in marital status, to establish own household, or other family reason as their primary reason for moving. Our fourth category is "Other," which includes all remaining reasons for moving.[1] Table A.4 gives a summary of the reasons for moving for all of our subgroups. For all subgroups, "Other" is the largest reason for moving and "New job or transfer" is the second largest for all subgroups. For the full sample of single searcher households, the percent of households that moved for "New job or transfer" reasons, 32.3%, is almost the same as the percent of households that moved for other reasons, 34.5%.

Table A.4: Reasons for Moving: Married Households

	Dual-TOT	Single-TOT	Dual-LT	Single-LT	Dual-COH	Single-COH
New job or transfer	27.0	32.3	27.1	32.4	24.5	29.6
Other job reasons	12.1	12.8	12.1	12.9	12.3	12.7
Family	22.4	20.4	22.3	20.3	25.0	22.3
Other	38.5	34.5	38.5	34.4	38.2	35.4
Total	100.0	100.0	100.0	100.0	100.0	100.0
Observations	11,071	6,775	11,028	6,757	15,032	8,035

We use the four broader categories to estimate the probability that a household with both spouses in the labor force will move for job related reasons. Specifically we model the probability that a move occurred for reason $j \in \{\text{New job or transfer, Other job reasons, Family, Other}\} = K$ as,

$$\text{P}(\text{whymove}_i = j) = \frac{e^{(\beta_{0j}+\beta_{1j} dual_i + X_i\gamma_j + \eta_{tj} + \varepsilon_i)}}{1 + \sum_{k\in K} e^{(\beta_{0k}+\beta_{1k} dual_i + X_i\gamma_j + \eta_{tk} + \varepsilon_i)}} \tag{A.2}$$

where the variables are defined as in the probit estimation. Here again, our base case is single searching households so our coefficient of interest is β_{11}. If dual searching households are in fact less likely to move for new jobs than single searching households, then we should expect β_{11} to be negative.

[1] The remaining reasons for moving are: wanted own home - not rent, wanted new or better housing, wanted better neighborhood, for cheaper housing, other housing reason, attend/leave college, change of climate, health reasons, other reasons, natural disaster, and foreclosure or eviction.

Table A.5: Multinomial Logit Results: Dual Searching

	Total	Living Together	Cohabitating
New job or transfer			
dual	-0.300***	-0.303***	-0.273***
	(0.0479)	(0.0479)	(0.0448)
Other job reasons			
dual	-0.0734	-0.0786	-0.0290
	(0.0621)	(0.0622)	(0.0570)
Family			
dual	0.0900	0.0929	0.0948*
	(0.0524)	(0.0526)	(0.0469)
N	14,793	14,745	19,152

Robust standard errors in parentheses. Base case is other reasons for moving.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Table A.5 gives the estimated coefficients on the labor market indicators where all probabilities are relative to moving for other reasons. Again focusing on the "Total" column, we find that the coefficient on *dual* is negative and statistically significant for the new job or transfer reasons for moving. The estimated coefficient implies that the relative probability a household with both spouses in the labor force moves for a new job or transfer is 25.9% lower than for a household with one spouse in the labor force[2]. The coefficients for all other reasons for moving across all subsamples are not statistically different from zero. Thus, we conclude that the strongest evidence of the difference between single and dual searcher household migration patterns points towards differences in the job search process.

Next we turn to the second channel of interest: the relative earnings of women. We construct the dependent variable, earnings ratio (ER), as the ratio of the total wage and salary income of married women to men in 1999 dollars at the state level. Since we do not know the exact location of individuals a year before the interview (we only know whether the individuals have moved and whether the move occurred across state or county lines), we now restrict our sample to households that only moved within a state in order to identify the relative earnings ratio the

[2]The relative probabilities of the multinomial logit are $1 - exp(\beta_j)$.

household faced last year. Since we are interested in the effect of the relative earnings of women on dual searcher households' migration decisions, we further restrict the sample to only dual searcher households.

To test the effect of the relative earnings of women on dual searcher household's migration decisions, we estimate a probit model with the following specification:

$$\mathrm{P}(\text{move}_i = 1) = \Phi(\beta_0 + \beta_1 ER_s + X_i\gamma + \eta_t + \varepsilon_i) \tag{A.3}$$

where ER_s is the relative earnings of women in the state in which the household lives. The rest of the covariates are the same as above. The coefficient of interest is β_1, which we expect to be negative.

Table A.6 gives the estimated coefficients on the relative earnings of women. The sign indicates that the probability of moving is decreasing in the earning ratio. Table A.7 gives the marginal effect of the earnings ratio for a household in which both spouses are white, 40 years old, have a college degree, own a home, and have a child present in the home in the year 2000. Again focusing on the "Total" column, Table A.7 shows that a 10 percentage point increase in the relative earnings of women decreases the probability that a dual searcher household moves by 0.07 percentage points.

Table A.6: Probit Estimation Results: Relative Earnings of Women

	Total	Living Together	Cohab
Earnings Ratio	-0.0160	-0.0156	-0.0328***
	(0.0102)	(0.0102)	(0.00885)
N	371,497	370,816	414,451

Robust Standard errors in parentheses
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Finally, we test whether those dual searching households that moved were less likely to

Table A.7: Probit Marginal Effects: Relative Earnings of Women

	Total	Living Together	Cohab
Earnings Ratio	-0.000669	-0.000648	-0.00143***
	(0.000423)	(0.000420)	(0.000381)
N	371,497	370,816	414,451

Marginal effects evaluated for a household in which both spouse are white, 40 years old, with a college degree, own a home and have a child present in the home in the year 2000. Robust standard errors in parentheses.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

move for job related reasons if they faced a higher earnings ratio. We use a multinomial logit with the same categories as in Eq. A.2 and the same covariates used in Eq. A.3. Again our coefficient of interest is that corresponding to the earnings ratio. Table A.8 gives the estimated coefficients for the multinomial logit. Although the estimated coefficients are negative, they are no longer statistically significant. This is most likely due to the fact that we have a very small sample with all the restrictions in place. However, we take the fact that the estimated coefficients have a negative sign as suggestive evidence that the earnings ratio decreases their probability of moving for job related reasons relative to other reasons for dual searcher households.

Table A.8: Multinomial Logit Results: Relative Earnings of Women

	Total	Living Together	Cohabitating
New job or transfer			
Earnings Ratio	-0.0502	-0.0518	-0.00514
	(0.0844)	(0.0845)	(0.0729)
Other job reasons			
Earnings Ratio	0.0558	0.0519	0.0450
	(0.0906)	(0.0908)	(0.0758)
Family			
Earnings Ratio	-0.0994	-0.107	-0.108
	(0.0700)	(0.0701)	(0.0577)
N	5,675	5,651	7,927

Robust standard errors in parentheses. Base case is other reasons for moving.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

The evidence presented here runs counter with that given in Molloy et al. (2017) and Kaplan and Schulhofer-Wohl (2017) who argue that the co-location problem did not contribute significantly to the historical decline in migration rates. The key differences arise because we condition on married households and do not restrict attention to couples employed in technical occupations as they do. Moreover, we control for self-selection into marriage by conditioning on married couples throughout our analysis and add to our empirics a comparison of moves for job related reasons between dual and single searchers.

A.2 Interstate Moves

Table A.9 gives the summary statistics for the different sample groups for interstate moves. Interstate moves are less, however, single searching household are observed to move more than dual searching household for the living together and total married samples. All other summary states are identical to those presented in Table A.1. Table A.10 gives the probit results for interstate moves. The results confirm those presented in Table A.2 and are of larger magnitude. The probability of moving across state lines when both individuals in the labor force of moving across states is 0.45 less than when only one spouse is in the labor force. Table A.2 shows the reasons for moving across state lines. Across all samples "New job or transfer" is the most common reason for moving. Consistent with intercounty moves, single searching households are more likely to move across state lines for job related reason than dual searching households across all samples. Table A.13 gives the results for the multinomial logit on interstate moves. Again the results are consistent with than those presented in Table A.5 and are of larger magnitude.

Table A.9: Summary Statistics: Interstate Moves

	(1) Dual LT	(2) Single LT	(3) Dual TOT	(4) Single TOT	(5) Dual COH	(6) Single COH
State Move	0.01	0.02	0.01	0.02	0.02	0.02
Total Real Family Income	83348.05	63294.73	83256.44	62064.38	80620.84	61534.13
Own Home	0.84	0.75	0.84	0.73	0.82	0.73
Head of Household Characteristics						
Age	43.09	44.05	43.09	43.87	42.69	43.76
White	0.85	0.84	0.85	0.83	0.85	0.84
Black	0.07	0.07	0.07	0.07	0.07	0.07
One race - Other	0.06	0.08	0.06	0.08	0.06	0.08
Multiple races	0.01	0.01	0.01	0.01	0.01	0.01
Less than High School	0.03	0.07	0.03	0.07	0.03	0.07
High School	0.32	0.39	0.32	0.39	0.33	0.40
Some College	0.29	0.24	0.29	0.24	0.29	0.25
College	0.24	0.19	0.24	0.19	0.23	0.19
Advanced Degree	0.13	0.10	0.13	0.10	0.13	0.10
Spouse Characteristics						
Age	43.20	44.46	43.19	44.45	42.76	44.12
White	0.85	0.84	0.85	0.79	0.85	0.84
Black sp	0.07	0.07	0.07	0.06	0.07	0.07
One race - Other	0.06	0.08	0.06	0.07	0.06	0.08
Multiple races	0.01	0.01	0.01	0.01	0.01	0.01
Less than High School	0.03	0.07	0.03	0.07	0.03	0.07
High School	0.32	0.39	0.32	0.37	0.33	0.40
Some College	0.28	0.24	0.28	0.23	0.29	0.24
College	0.24	0.19	0.24	0.18	0.23	0.18
Advanced Degree	0.13	0.10	0.13	0.15	0.13	0.10
Observations	786,805	362,946	788,243	385,329	832,232	378,948

Table A.10: Probit Estimation Results: Interstate Moves

	Total	Living Together	Cohab
dual	-0.107***	-0.106***	-0.0879***
	(0.0111)	(0.0111)	(0.0102)
N	508,371	507,457	561,633

Robust Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.11: Probit Marginal Effects

	Total	Living Together	Cohab
dual	-0.00454***	-0.00449***	-0.00386***
	(0.000532)	(0.000532)	(0.000490)
N	508,371	507,457	561,633

Marginal effects evaluated for a household in which both spouse are white, 40 years old, with a college degree, own a home and have a child present in the home in the year 2000. Robust standard errors in parentheses.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

Table A.12: Reasons for Insterstate Moves: Married Households

	Dual-TOT	Single-TOT	Dual-LT	Single-LT	Dual-COH	Single-COH
New job or transfer	38.6	43.7	38.6	43.9	34.8	40.4
Other job reasons	11.3	11.6	11.3	11.6	11.6	11.5
Family	21.1	19.8	21.1	19.7	23.5	21.6
Other	29.0	24.9	29.0	24.9	30.1	26.5
Total	100.0	100.0	100.0	100.0	100.0	100.0
Observations	5,154	3,644	5,154	3,633	6,896	4,232

Table A.13: Multinomial Logit Results: Interstate Moves

	Total	Living Together	Cohabitating
New job or transfer			
dual	-0.385***	-0.392***	-0.356***
	(0.0681)	(0.0683)	(0.0630)
Other job reasons			
dual	-0.120	-0.129	-0.0758
	(0.0952)	(0.0954)	(0.0865)
Family			
dual	-0.0222	-0.0184	0.00108
	(0.0784)	(0.0786)	(0.0703)
N	5,208	5,189	6,477

Robust standard errors in parentheses. Base case is other reasons for moving.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix B

A Model of Dual Search with Exogenous Moves

B.1 Model

The value function for single searchers in a model with exogenous moves that arrive at Poisson rate η are:

$$\begin{aligned} rUO^i &= b_O + b_U^i + \alpha_l^u \int \max\{EO^i(w) - UO^i, 0\}\, dF^i(w) \\ &\quad + \alpha_f^u \int \max\{EO^i(w) - UO^i, 0\}\, dF^i(w), \end{aligned} \tag{B.1}$$

$$\begin{aligned} rEO^i(w) &= b_O + w + \alpha_l^e \int_w \left[EO^i(w') - EO^i(w)\right] dF^i(w') \\ &\quad + \alpha_f^e \int_w \left[EO^i(w') - EO^i(w)\right] dF^i(w') \\ &\quad + \delta \left[UO^i - EO^i(w)\right] \\ &\quad + \eta \left[UO^i - EO^i(w)\right]. \end{aligned} \tag{B.2}$$

The migration rate for a type i singe searcher is:

$$M_s^i = u_s^i \cdot \alpha_f^u \left[1 - F^i(R_s^i)\right] + \left(1 - u_s^i\right) \cdot \alpha_f^e \int_{R_s^i}^{\infty} \left[1 - F^i(w)\right] dG^i(w) + \eta. \tag{B.3}$$

The value functions for dual searchers in a model with exogenous moves are:

$$\begin{aligned} rUU = b_U^M + b_U^F &+ \left(\alpha_l^u + \alpha_f^u\right) \int_{R_1^M}^{\infty} \left[EU^M(w') - UU\right] dF^M(w') \\ &+ \left(\alpha_l^u + \alpha_f^u\right) \int_{R_1^F}^{\infty} \left[EU^F(w') - UU\right] dF^F(w'), \end{aligned} \tag{B.4}$$

$$\begin{aligned} rEU^i(w) = b_U^{-i} + w + & \left(\alpha_l^e + \alpha_f^e\right) \int_{w}^{\infty} \left[EU^i(w') - EU^i(w)\right] dF^i(w') \\ & + \alpha_l^u \int_{\phi^{-i}(w)}^{\infty} \max\left\{EE(w, w') - EU^i(w), EU^{-i}(w') - EU^i(w)\right\} \, dF^{-i}(w') \\ & + \alpha_f^u \int_{R_3^{-i}(w)}^{\infty} \left[EU^{-i}(w') - EU^i(w)\right] dF^{-i}(w') \\ & + \delta \left[UU - EU^i(w)\right] \\ & + \eta \left[UU - EU^i(w)\right], \end{aligned} \tag{B.5}$$

$$\begin{aligned} rEE(w,w') = w + w' &+ \alpha_l^e \int_w^\infty \left[EE(w'',w') - EE(w,w')\right] dF^M(w'') \\ &+ \alpha_l^e \int_{w'}^\infty \left[EE(w,w'') - EE(w,w')\right] dF^F(w'') \\ &+ \alpha_f^e \int_{M^M(w,w')}^\infty \left[EU^M(w'') - EE(w,w')\right] dF^M(w'') \\ &+ \alpha_f^e \int_{M^F(w,w')}^\infty \left[EU^F(w'') - EE(w,w')\right] dF^F(w'') \\ &+ \delta \left[\max\left\{EU^M(w), UU\right\} - EE(w,w')\right] \\ &+ \delta \left[\max\left\{EU^F(w'), UU\right\} - EE(w,w')\right] \\ &+ \eta \left[UU - EE(w,w')\right]. \end{aligned} \tag{B.6}$$

The migration rate for a dual searching household is:

$$\begin{aligned} M_d = {} & \alpha_f^u \left[2 - F_m(R_1^m) - F_f(R_1^f)\right] \cdot uu \\ & + \alpha_f^e \left(eu_f \cdot \int_{R_1^F}^\infty \left[1 - F_f(w)\right] dT_f(w) + eu_m \cdot \int_{R_1^M}^\infty \left[1 - F_m(w)\right] dT_m(w) \right) \\ & + \alpha_f^u \left(eu_f \cdot \int_{R_1^F}^\infty \left[1 - F_m(R_3^m(w))\right] dT_f(w) + eu_m \cdot \int_{R_1^M}^\infty \left[1 - F_f(R_3^f(w))\right] dT_m(w) \right) \\ & + ee \cdot \alpha_f^e \int_{R_1^M}^\infty \int_{R_2^f(w)}^\infty \left[1 - F_m(M_m(w,w'))\right] d^2H(w,w') \\ & + ee \cdot \alpha_f^e \int_{R_1^F}^\infty \int_{R_2^m(w')}^\infty \left[1 - F_f(M_f(w,w'))\right] d^2H(w,w') \\ & + \eta. \end{aligned} \tag{B.7}$$

B.2 Calibration

Table B.1: Calibrated Parameters: Model with Exogenous Moves

Parameter	Value	Description
α_l^u	27.770	Local unemp. arrival rate
α_l^e	15.508	Local emp. arrival rate
α_f^u	1.076	Foreign unemp. arrival rate
α_f^e	0.866	Foreign emp. arrival rate
μ_M	9.385	Male location parameter
σ_M	0.518	Male shape parameter
μ_F	8.687	Female location parameter
σ_F	0.550	Female shape parameter
b_U^M	$5,005$	Male flow utility of unemp.
b_U^F	$6,004$	Female flow utility of unemp.
η_s	0.0314	Single Searcher Exogenous move arrival rate
η_d	0.0287	Dual Searcher Exogenous move arrival rate
λ_M	$20,400$	Mean Value of Male Non-participation
λ_F	$34,608$	Mean Value of Female Non-participation

Table B.2: Calibrated Moments: Model with Exogenous Moves

Moment	Model	Data
Single Searcher Mig. Rate	0.0551	0.057
Dual Searcher Mig. Rate	0.0484	0.047
Mass in EE	0.80	0.79
Mass in EU^M	0.13	0.13
Mass in EU^F	0.068	0.047
Mass in EO^M	0.96	0.89
Mass in EO^F	0.93	0.82
Male Median Wage ($)	$39,591$	$38,000$
Female Median Wage ($)	$21,507$	$23,000$
Male 90-50 Wage Ratio	1.43	2.15
Female $90-50$ Wage Ratio	1.45	2.17
Fraction of Dual Searchers	0.752	0.752
Fraction of Male Single Searchers	0.784	0.783

Table B.3: Counterfactual Parameters: Participation Decision with Exogenous Moves

Parameter	Value	Description
λ^M	19,088	Mean Value of Male Non-participation
λ^F	110,672	Mean Value of Female Non-participation

Bibliography

Akaike, Hirotugu, "A new look at the statistical model identification," in "Selected Papers of Hirotugu Akaike," Springer, 1974, pp. 215–222.

Andrews, Donald W. K., "Tests for Parameter Instability and Structural Change With Unknown Change Point," *Econometrica*, 1993, *61* (4), 821–856.

_ , "Tests for Parameter Instability and Structural Change with Unknown Change Point: A Corrigendum," *Econometrica*, 2003, *71* (1), 395–397.

Auclert, Adrien, Hannes Malmberg, Frédéric Martenet, and Matthew Rognlie, "Demographics, Wealth and Interest Rates in the 21st Century," Mimeo, Stanford 2019. Presented at SITE Summer Workshops 2018.

Auerbach, Alan J. and Laurence J. Kotlikoff, *Dynamic Fiscal Policy*, Cambridge University Press, 1987.

Backus, David, Thomas F. Cooley, and Espen Henriksen, "Demography and Low-Frequency Capital Flows," *Journal of International Economics*, 2014, *92* (1), 94–102.

Bárány, Zsófia, Nicolas Coeurdacier, and Stéphane Guibaud, "Capital Flows in an Aging World," Mimeo, SciencePo 2019.

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

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

Blau, David M. and Donna B. Gilleskie, "Health insurance and retirement of married couples," *Journal of Applied Econometrics*, 2006, *21* (7), 935–953.

_ **and** _ , "The Role of Retiree Health Insurance in the Employment Behavior of Older Men," *International Economic Review*, 2008, *49* (2), 475–514.

Blondal, S and Stefano Scarpetta, "Early Retirement in OECD Countries: The Role of Social Security Systems," *OECD Economic Studies*, 1997, *29*, 7–54.

Börsch-Supan, Axel, Duarte Nuno Leite, and Johannes Rausch, "Demographic changes, migration and economic growth in the Euro area," Prepared for the ECB Forum on Central Banking 2019, Sintra, Portugal, June 17–19, 2019 2019.

Bowlus, Audra J. and Jean-Marc Robin, "Twenty Years of Rising Inequality in U.S. Lifetime Labour Income Values," *The Review of Economic Studies*, 2004, *71* (3), 709–742.

Braun, Christine, Charlie Nusbaum, and Peter Rupert, "Labor market dynamics and the migration behavior of married couples," *Review of Economic Dynamics*, 2020.

Brugnolini, Luca, "About Local Projection Impulse Response Function Reliability," Working Paper 440, CEIS 2018.

Burdett, Kenneth and Dale T Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, May 1998, *39* (2), 257–273.

Canova, Fabio, David Lopez-Salido, and Claudio Michelacci, "The effects of technology shocks on hours and output: a robustness analysis," *Journal of Applied Econometrics*, 2010, *25* (5), 755–773.

Capatina, Elena, "Life-cycle effects of health risk," *Journal of Monetary Economics*, 2015, *74*, 67 – 88.

Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio, "Demographics and Real Interest Rates: Inspecting the Mechanism," *European Economic Review*, 2016, *88*, 208–226.

Case, Anne and Angus Deaton, "Mortality and Morbidity in the 21$^{\text{st}}$ Century," *Brookings Papers on Economic Activity*, 2017, *48* (1), 397–476.

Chari, V. V., Patrick J. Kehoe, and Ellen R. McGrattan, "Business Cycle Accounting," *Econometrica*, 2007, *75* (3), 781–836.

Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan, "Are structural VARs with long-run restrictions useful in developing business cycle theory?," *Journal of Monetary Economics*, 2008, *55* (8), 1337 – 1352.

Choi, Chi-Young and Alexander Chudik, "Estimating impulse response functions when the shock series is observed," *Economics Letters*, 2019, *180*, 71 – 75.

Chow, Gregory C., "Tests of Equality Between Sets of Coefficients in Two Linear Regressions," *Econometrica*, 1960, *28* (3), 591–605.

Christiano, Lawrence J, Martin Eichenbaum, and Robert Vigfusson, "What Happens After a Technology Shock?," Working Paper 9819, National Bureau of Economic Research July 2003.

Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson, "The Response of Hours to a Technology Shock: Evidence based on Direct Measures of Technology," *Journal of the European Economic Association*, 2004, *2* (2-3), 381–395.

_ , _ , and _ , "Alternative Procedures for Estimating Vector Autoregressions Identified with Long-run Restrictions," *Journal of the European Economic Association*, 2006, *4* (2/3), 475–483.

Coen-Pirani, Daniele, "Understanding gross worker flows across U.S. states," *Journal of Monetary Economics*, October 2010, *57* (7), 769 – 784.

Coltrane, Scott, "Research on Household Labor: Modeling and Measuring the Social Embeddedness of Routine Family Work," *Journal of Marriage and Family*, 2000, *62* (4), 1208–1233.

Conesa, Juan Carlos and Carlos Garriga, "Intergenerational policy and the measurement of tax incidence," *European Economic Review*, 2016, *83* (C), 1–18.

_ , Timothy J. Kehoe, Vegard M. Nygaard, and Gajendran Raveendranathan, "Implications of Increasing College Attainment for Aging in General Equilibrium," Working Paper 26000, NBER 2019.

Cooke, Thomas J., "It is not Just the Economy: Declining Migration and the Rise of Secular Rootedness," *Population, Space and Place*, 2011, *17* (3), 193–203.

_ , "Internal Migration in Decline," *The Professional Geographer*, 2013, *65* (4), 664–675.

Cooley, Thomas F. and Espen Henriksen, "The Demographic Deficit," *Journal of Monetary Economics*, 2018, *93*, 45–62.

Cooley, Thomas F. and Mark Dwyer, "Business cycle analysis without much theory A look at structural VARs," *Journal of Econometrics*, 1998, *83* (1), 57 – 88.

Costa, Dora L. and Matthew E. Kahn, "Power Couples: Changes in the Locational Choice of the College Educated, 1940–1990*," *The Quarterly Journal of Economics*, 2000, *115* (4), 1287–1315.

Dey, Matthew S. and Christopher J. Flinn, "An Equilibrium Model of Health Insurance Provision and Wage Determination," *Econometrica*, 2005, *73* (2), 571–627.

Erceg, Christopher J., Luca Guerrieri, and Christopher Gust, "Can Long-Run Restrictions Identify Technology Shocks?," *Journal of the European Economic Association*, 2005, *3* (6), 1237–1278.

Erosa, Andrés, Luisa Fuster, and Gueorgui Kambourov, "Labor supply and government programs: A cross-country analysis," *Journal of Monetary Economics*, 2012, *59* (1), 84 – 107.

—, —, and —, "Towards a Micro-Founded Theory of Aggregate Labour Supply," *Review of Economic Studies*, 2016, *83* (3), 1001–1039.

Evertsson, Marie and Magnus Nermo, "Dependence within families and the division of labor: Comparing Sweden and the United States," *Journal of Marriage and Family*, November 2004, *66* (5), 1272–1286.

Faust, Jon and Eric M. Leeper, "When Do Long-Run Identifying Restrictions Give Reliable Results?," *Journal of Business and Economic Statistics*, 1997, *15* (3), 345–353.

Fernald, John, "A quarterly, utilization-adjusted series on total factor productivity," Working Paper 2012-19, Federal Reserve Bank of San Francisco 2014.

Fernald, John G., "Trend breaks, long-run restrictions, and contractionary technology improvements," *Journal of Monetary Economics*, 2007, *54* (8), 2467 – 2485.

Fernández-Villaverde, Jesús, Juan F. Rubio-Ramírez, Thomas J. Sargent, and Mark W. Watson, "ABCs (and Ds) of Understanding VARs," *American Economic Review*, June 2007, *97* (3), 1021–1026.

Feroli, Michael, "Capital Flows Among the G7 Nations: A Demographic Perspective," FEDS Working Paper 2003-54, Board of Governors of the Federal Reserve System 2003.

Ferrero, Andrea, "A structural decomposition of the U.S. trade balance: Productivity, demographics and fiscal policy," *Journal of Monetary Economics*, 2010, *57* (4), 478–490.

Flabbi, Luca and James Mabli, "Household Search or Individual Search: Does It Matter?," *Journal of Labor Economics*, 2018, *36* (1), 1–46.

Flinn, Christopher J., "Labour Market Structure and Inequality: A Comparison of Italy and the U.S.," *The Review of Economic Studies*, 2002, *69* (3), 611–645.

Flinn, Christopher J, "Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates," *Econometrica*, 2006, *74* (4), 1013–1062.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren, "Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [dataset]," *Minneapolis, MN: IPUMS*, 2018.

Foged, Mette, "Family migration and relative earnings potentials," *Labour Economics*, 2016, *42*, 87 – 100.

Francis, Neville and Valerie A. Ramey, "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," *Journal of Monetary Economics*, 2005, *52* (8), 1379 – 1399.

French, Eric, "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour," *Review of Economic Studies*, 2005, *72* (2), 395–427.

_ **and John Bailey Jones**, "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica*, 2011, *79* (3), 693–732.

Fuster, Luisa, Ayşe İmrohoroğlu, and Selahattin İmrohoroğlu, "Elimination of Social Security in a Dynastic Framework," *Review of Economic Studies*, 2007, *74* (1), 113–145.

Gagnon, Etienne, Benjamin K. Johannsen, and David López-Salido, "Understanding the New Normal: The Role of Demographics," Finance and Economics Discussion Series 2016-080, Board of Governors of the Federal Reserve System 2016.

Galí, Jordi and Pau Rabanal, "Technology shocks and aggregate fluctuations: How well does the real business cycle model fit postwar US data?," *NBER macroeconomics annual*, 2005, *19*, 225–288.

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

Garibaldi, Pietro, Joaquim Oliveira Martins, and Jan van Ours, eds, *Ageing, Health, and Productivity: The Economics of Increased Life Expectancy*, Oxford University Press, 2010.

Garín, Julio, Robert Lester, and Eric Sims, "Are Supply Shocks Contractionary at the ZLB? Evidence from Utilization-Adjusted TFP Data," *The Review of Economics and Statistics*, March 2019, *101* (1), 160–175.

Gemici, Ahu, "Family Migration and Labor Market Outcomes," *Working Paper*, 2016.

Gollin, Douglas, "Getting Income Shares Right," *Journal of Political Economy*, 2002, *110* (2), 458–474.

Gordon, Robert J., *The Rise and Fall of American Growth*, Princeton University Press, 2016.

Griffy, Benjamin, "Borrowing Constraints, Search, and Life-Cycle Inequality," *Working Paper*, November 2017.

Gruber, Jonathan and David Wise, *Social Security and Retirement around the World*, National Bureau of Economic Research, Inc, 1999.

Guler, Bulent and Ahmet Ali Taskin, "Dual Income Couples and Interstate Migration," *Working Paper*, February 2018, pp. 1–40.

_ **, Fatih Guvenen, and Giovanni L. Violante**, "Joint-search theory: New opportunities and new frictions," *Journal of Monetary Economics*, 2012, *59* (4), 352–369.

Hacker, Scott R and Abdulnasser Hatemi-J, "Optimal lag-length choice in stable and unstable VAR models under situations of homoscedasticity and ARCH," *Journal of Applied Statistics*, 2008, *35* (6), 601–615.

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

Hall, Robert E, "Job Loss, Job Finding, and Unemployment in the U.S. Economy Over the Past Fifty Years," Working Paper 11678, National Bureau of Economic Research October 2005.

Hamilton, James D., "Historical Causes of Postwar Oil Shocks and Recessions," *The Energy Journal*, 1985, *0* (Number 1), 97–116.

Hansen, G.D, "The Cyclcical and Secular Behaviour of the Labor Input: Comparing Efficiency Units and Hours Worked," *Journal of Applied Econometrics*, 1993, *8* (1), 71–80.

Henriksen, Espen, "Representative Survival Probabilities for Heterogeneous-Agent Economies," Mimeo, UC Davis 2015.

_ **and Frederic Lambert**, ""Imbalances" for the Long Run," Working paper, NYU Stern School 2018.

Hopenhayn, Hugo, Julian Neira, and Rish Singhania, "From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share," Working Paper 25382, NBER 2018.

Huggett, Mark, Gustavo Ventura, and Amir Yaron, "Sources of Lifetime Inequality," *American Economic Review*, December 2011, *101* (7), 2923–54.

Hurwicz, Leonid, "Least squares bias in time series," *Statistical inference in dynamic economic models*, 1950, *10*, 365–383.

Ikeda, Daisuke and Masashi Saito, "The Effects of Demographic Changes on the Real Interest Rate in Japan," *Japan and the World Economy*, 2014, *32*, 37–48.

İmrohoroğlu, Selahattin, Sagiri Kitao, and Tomoaki Yamada, "Achieving fiscal balance in Japan," *International Economic Review*, 2016, *57* (1), 117–154.

Ingram, Beth Fisher, Narayana R. Kocherlakota, and N.E. Savin, "Explaining business cycles: A multiple-shock approach," *Journal of Monetary Economics*, 1994, *34* (3), 415 – 428.

Ireland, Peter N., "A method for taking models to the data," *Journal of Economic Dynamics and Control*, 2004, *28* (6), 1205 – 1226.

Iyigun, Murat and Jeanne Lafortune, "Why Wait? A Century of Education, Marriage Timing and Gender Roles," *IZA Discussion Papers, No. 9671*, Jan 2016.

Jordà, Oscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, March 2005, *95* (1), 161–182.

_ , "Simultaneous Confidence Regions for Impulse Responses," *The Review of Economics and Statistics*, 2009, *91* (3), 629–647.

_ **and Sharon Kozicki**, "Estimation and Inverence by the Method of Projections Minimum Distance: An Application to the New Keynesian Hybrid Phillips Curve," *International Economic Review*, 2011, *52* (2), 461–487.

Kaplan, Greg, "Inequality and the life cycle," *Quantitative Economics*, 2012, *3* (3), 471–525.

_ **and Sam Schulhofer-Wohl**, "Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey," *Demography*, Aug 2012, *49* (3), 1061–1074.

_ **and** _ , "Understanding The Long-Run Decline In Interstate Migration," *International Economic Review*, 2017, *58* (1), 57–94.

Karahan, Fatih and Serena Rhee, "Population aging, migration spillovers, and the decline in interstate migration," Staff Reports 699, Federal Reserve Bank of New York 4 2017.

Kascha, Christian and Karel Mertens, "Business cycle analysis and VARMA models," *Journal of Economic Dynamics and Control*, 2009, *33* (2), 267 – 282.

Keese, Mark, "A Method for Calculating the Average Effective Age of Retirement," Methodology Paper, OECD 2003.

Kennan, John and James R. Walker, "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 2011, *79* (1), 211–251.

Kilian, Lutz, "Small-Sample Confidence Intervals for Impulse Response Functions," *The Review of Economics and Statistics*, 1998, *80* (2), 218–230.

_ **and Helmut Ltkepohl**, "Estimation Subject to Long-Run Restrictions," in "Structural Vector Autoregressive Analysis" Themes in Modern Econometrics, Cambridge University Press, 2017, pp. 297–333.

_ **and Yun Jung Kim**, "How Reliable Are Local Projection Estimators of Impulse Responses?," *The Review of Economics and Statistics*, 2011, *93* (4), 1460–1466.

King, Robert, Charles Plosser, James Stock, and Mark Watson, "Stochastic Trends and Economic Fluctuations," *American Economic Review*, 1991, *81* (4), 819–40.

King, Robert G and Sergio T Rebelo, "Resuscitating real business cycles," *Handbook of macroeconomics*, 1999, *1*, 927–1007.

King, Robert G. Charles I. Plosser, and Sergio T. Rebelo, "Production, Growth, and Business Cycles," *Journal of Monetary Economics*, 1988, *21* (2), 195–232.

Kitao, Sagiri, "Sustainable Social Security: Four Options," *Review of Economic Dynamics*, 2014, *17* (4), 756–779.

Kopecky, Joseph, "An Aging Dynamo: Demographic Change and the Decline of Entrepreneurial Activity in the United States," Mimeo, Trinity College Dublin 2018.

Kotlikoff, Laurence, Kent Smetters, and Jan Walliser, "Mitigating America's demographic dilemma by pre-funding social security," *Journal of Monetary Economics*, 2007, *54* (2), 247–266.

Krueger, Dirk and Alexander Ludwig, "On the consequences of demographic change for rates of returns to capital, and the distribution of wealth and welfare," *Journal of Monetary Economics*, 2007, *54* (1), 49–87.

Kydland, Finn E. and Edward C. Prescott, "Time to Build and Aggregate Fluctuations," *Econometrica*, 1982, *50* (6), 1345–1370.

Lewis, Richard and Gregory C Reinsel, "Prediction of multivariate time series by autoregressive model fitting," *Journal of Multivariate Analysis*, 1985, *16* (3), 393 – 411.

Lippi, Marco and Lucrezia Reichlin, "VAR analysis, nonfundamental representations, blaschke matrices," *Journal of Econometrics*, 1994, *63* (1), 307 – 325.

Lundberg, Shelly and Elaina Rose, "Parenthood and the earnings of married men and women," *Labour Economics*, November 2000, *7* (6), 689–710.

_ **and Robert A. Pollak**, "Separate Spheres Bargaining and the Marriage Market," *Journal of Political Economy*, 1993, *101* (6), 988–1010.

Lütkepohl, Helmut, "Comparison of criteria for estimating the order of a vector autoregressive process," *Journal of time series analysis*, 1985, *6* (1), 35–52.

Marcassa, Stefania, "Unemployment Duration of Spouses: Evidence From France," *LABOUR*, 2014, *28* (4), 399–429.

Marinescu, Ioana and Roland Rathelot, "Mismatch Unemployment and the Geography of Job Search," *American Economic Journal: Macroeconomics*, July 2018, *10* (3), 42–70.

McGrattan, Ellen R., "Measurement with minimal theory," *Federal Reserve Bank of Minneapolis Quarterly Review*, 2010, *33* (1), 2–13.

— **and Edward C. Prescott**, "On financing retirement with an aging population," *Quantitative Economics*, 2017, *8* (1), 75–115.

— **and** — , "An aggregate model for policy analysis with demographic change," *Journal of the Economics of Ageing*, 2018, *11*, 52–61.

Mincer, Jacob, "Family Migration Decisions," *Journal of Political Economy*, 1978, *86* (5), 749–773.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak, "Internal Migration in the United States," *Journal of Economic Perspectives*, September 2011, *25* (3), 173–96.

— **,** — **, and** — , "Job Changing and the Decline in Long-Distance Migration in the United States," *Demography*, Apr 2017, *54* (2), 631–653.

Nardi, Mariacristina De, Selahattin İmrohoroğlu, and Thomas Sargent, "Projected U.S. Demographics and Social Security," *Review of Economic Dynamics*, 1999, *2* (3), 575–615.

Ndiaye, Abdoulaye, "Flexible Retirement and Optimal Taxation," Mimeo, NYU Stern 2020.

Ozcicek, Omer and W.D̃ouglas Mcmillin, "Lag length selection in vector autoregressive models: symmetric and asymmetric lags," *Applied Economics*, 1999, *31* (4), 517–524.

Pashchenko, Svetlana and Ponpoje Porapakkarm, "Work incentives of Medicaid beneficiaries and the role of asset testing," *International Economic Review*, 2017, *58* (4), 1117–1154.

Pissarides, Christopher A, *Equilibrium unemployment theory*, MIT press, 2000.

Plagborg-Møller, Mikkel and Christian K. Wolf, "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, forthcoming.

Prescott, Edward C., "Theory ahead of business-cycle measurement," *Carnegie-Rochester Conference Series on Public Policy*, 1986, *25*, 11 – 44.

Ramey, Valerie A., "Identifying Government Spending Shocks: It's all in the Timing*," *The Quarterly Journal of Economics*, 02 2011, *126* (1), 1–50.

Ramey, Valerie A, "Macroeconomic shocks and their propagation," in "Handbook of macroeconomics," Vol. 2, Elsevier, 2016, pp. 71–162.

Ravenna, Federico, "Vector autoregressions and reduced form representations of DSGE models," *Journal of Monetary Economics*, 2007, *54* (7), 2048 – 2064.

Rendon, Silvio and J. Ignacio García-Pérez, "Family Job Search and Wealth: The Added Worker Effect Revisited," 2018 Meeting Papers 503, Society for Economic Dynamics 2018.

Ríos-Rull, José-Víctor, "Computation of Equilibria in Heterogeneous-Agent Models," in Ramon Marimon and Andrew Scott, eds., *Computational Methods for the Study of Dynamic Economies*, Oxford University Press, 1999, chapter 11, pp. 238–264.

Rogerson, Richard and Johanna Wallenius, "Nonconvexities, Retirement, and the Elasticity of Labor Supply," *American Economic Review*, 2013, *103* (4), 1445–1462.

Romer, Christina D. and David H. Romer, "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, September 2004, *94* (4), 1055–1084.

Rupert, Peter and Giulio Zanella, "Revisiting wage, earnings, and hours profiles," *Journal of Monetary Economics*, 2015, *72* (C), 114–130.

Rust, John and Christopher Phelan, "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets," *Econometrica*, 1997, *65* (4), 781–831.

Saijo, Hikaru, "Technology shocks and hours revisited: Evidence from household data," *Review of Economic Dynamics*, 2019, *31*, 347 – 362.

Schwarz, Gideon, "Estimating the Dimension of a Model," *The Annals of Statistics*, 1978, *6* (2), 461–464.

Shapiro, Matthew D. and Mark W. Watson, "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual*, 1988, *3*, 111–148.

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

_ , "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 2012, *15* (2), 127 – 148.

Sims, Eric, "Permanent and transitory technology shocks and the behavior of hours: A challenge for DSGE models," *Unpublished manuscript. University of Notre Dame, Notre Dame, IN*, 2011.

Sposi, Michael, "Demographics and the Evolution of Global Imbalances," Globalization Institute Working Paper 332, Federal Reserve Bank of Dallas 2019.

Summers, Lawrence H, "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound," *Business Economics*, 2014, *49* (2), 65–73.

_ , "The Age of Secular Stagnation," *Foreign Affairs*, 2016, ( ), 2–9.

Taskin, Ali Ahmet, "Sorted and Settled: Migration Decisions of Dual Income Families," *Working Paper*, February 2016.

United Nations, "World Population Prospects: The 2017 Revision, Key Findings and Advance Tables," Working Paper ESA/P/WP/248, Department of Economics and Social Affairs, Population Division 2017.

van der Klaauw, Wilbert and Kenneth I. Wolpin, "Social security and the retirement and savings behavior of low-income households," *Journal of Econometrics*, 2008, *145* (1), 21–42.