From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2019

ESSAYS ON MACROECONOMIC POLICY WITH HETEROGENEOUS AGENTS

Bernstein, Joshua

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01bc386n08p

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

ESSAYS ON MACROECONOMIC POLICY WITH HETEROGENEOUS AGENTS

JOSHUA MARK SAMUEL BERNSTEIN

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE BY

THE DEPARTMENT OF

ECONOMICS

ADVISER: MIKHAIL GOLOSOV

JUNE 2019

© Copyright by Joshua Mark Samuel Bernstein, 2019. All rights reserved.

Abstract:

What are the consequences of microeconomic heterogeneity for the transmission and design of macroeconomic policies? This dissertation collects together 3 essays that contribute to our understanding of this important question.

In chapter one, I study the transmission of aggregate shocks in a New Keynesian model in which households' incomes are heterogeneously exposed to changes in aggregate income, and borrowing frictions limit opportunities for aggregate risk sharing. I analytically show that shock transmission is asymmetric: output responds more to contractionary shocks than to expansionary shocks of equal magnitude. Estimating key model parameters using the micro evidence on heterogeneous consumption exposures to changes in output generates asymmetric responses of output to monetary policy shocks that can explain at least 60% of the empirical asymmetry.

In chapter two, co-authored with David Arnold, we study the macroeconomic effects of Employment Protection Legislation (EPL) in Brazil. The fact that Brazilian EPL only affects jobs with tenures greater than three months causes a spike in the job termination hazard at three months, the size of which identifies the behavioral response of firms to the imposition of EPL. We estimate a structural model that maps this response into macroeconomic outcomes of interest. The imposition of EPL causes a 1.3% drop in GDP, which is driven by a 0.9 percentage point increase in the unemployment rate. Intuitively, the reduction in vacancy creation by firms caused by EPL widens the gap between the efficient level of vacancy creation and the observed equilibrium level.

Finally, in chapter three, I study optimal taxation in an economy in which households make consumption, labor supply, and fertility choices. I derive sufficient statistics for the sign and shape of optimal wedges on child quantity, goods investment and time investment, and provide intuition for the main economic forces at play. Distorting

fertility choices relaxes incentive constraints, which facilitates redistribution, but also may discourage households to earn income, thus hampering the redistributive strength of the income tax. A quantitative exercise demonstrates that the bulk of welfare gains available from subsidizing investments in children can be obtained using feasible linear subsidies on child investment goods.

Acknowledgments:

Completing my doctoral studies at Princeton would not have been possible without the generous support of many people.

Firstly, my advisor Mikhail Golosov. His willingness to listen to my ideas, even when they were far from fully formed, and to give me sharp insights and constructive criticism undoubtedly helped me to improve my understanding of macroeconomic theory, and to better envision its use in attempting to answer important policy questions. I will always be incredibly grateful that he continued to advise me after he left Princeton at the end of my second year.

As the job market approached, Mark Aguiar and Gianluca Violante were an invaluable source of encouragement and support that truly helped me navigate the most stressful three months of my life. Their congratulations when I secured my future employment meant an enormous amount to me.

I must also thank the wider community of Princeton Economics for creating such a highly intellectual and stimulating environment in which I could develop my skills as an economist. Being able to learn from such great minds has been an experience that I will never forget. Similarly, finding myself surrounded by like-minded students during my studies has been hugely satisfying. Special thanks go to Christian Wolf, David Arnold, and Charis Katsiardis for their support and inspiration throughout my time at Princeton.

It goes without saying that family has played an important role as I have navigated the PhD process. I especially thank my parents for encouraging me to attend Princeton, and for their continued support and guidance as I embark on the early stages of my academic career in the USA. I also thank Stan and Michelle DeGroote for welcoming me into their home.

Finally, I thank Jessa DeGroote for her unending support and belief that I could and would complete my doctoral studies. Her unwavering confidence in me has been a steady beacon of light that has helped me develop, both as an economist, and also as a person.

For my family

Contents

Abstract iii

Acknowledgments v

1 A Model of the Asymmetric Transmission of Aggregate Shocks 1

1.1 Introduction 1

1.2 Environment 6

1.3 Asymmetric Transmission of Aggregate Shocks 12

1.3.1 Asymmetric Output Responses to Monetary Policy Shocks . . 12

1.3.2 Other Aggregate Shocks 21

1.3.3 Inflation Responses 22

1.3.4 Idiosyncratic Shocks 23

1.3.5 Heterogeneous Preferences 24

1.4 Quantitative Exercise 25

1.4.1 Generalizing the Mechanism 26

1.4.2 A Semi-Structural Model 29

1.4.3 Parameter Estimation 32

1.4.4 Data 36

1.4.5 Results . . . 40
1.4.6 Quantitative Assessment . . . 42
1.5 Empirical Evidence of Monetary Policy Asymmetry . . . 43
1.5.1 Empirical Specification . . . 44
1.5.2 Results . . . 44
1.5.3 Robustness Checks . . . 46
1.6 Conclusion . . . 47
References . . . 49
1.A Additional Figures . . . 59
1.B Robustness for CEX Regressions . . . 61
1.C Additional Results for Section 1.3 . . . 62
1.C.1 Idiosyncratic Risk . . . 66
1.C.2 Heterogeneous Preferences . . . 71
1.D Proofs of Analytical Results . . . 73

2 Employment Protection Legislation and Efficiency in a Frictional Labor Market: Evidence from Brazil 104

2.1 Introduction . . . 104
2.2 Data and Institutional Setting . . . 108
2.2.1 Data . . . 108
2.2.2 Institutional Details . . . 109
2.2.3 Sample Selection . . . 110
2.3 The Impact of EPL on the Job Termination Hazard . . . 112

2.3.1 Estimating Bunching in Job Terminations 112
2.3.2 Temporary Contracts 115
2.3.3 Summary 121
2.4 Model 121
2.4.1 Belief Dynamics 123
2.4.2 Value Functions 124
2.4.3 Wage Determination via Nash Bargaining 129
2.4.4 Discontinuities at T_1 131
2.4.5 Hazard Rate Spike at Tenure $t = T_1$ 132
2.4.6 Ergodic Distribution of Beliefs 134
2.4.7 The Matching Function 136
2.4.8 Vacancy Posting 136
2.4.9 Equilibrium 137
2.4.10 Aggregate Output and Welfare 138
2.4.11 Efficiency 138
2.5 Estimation 140
2.5.1 Calibration 140
2.5.2 Estimation 142
2.5.3 Identification 143
2.6 Results 144
2.6.1 Estimation 144
2.6.2 The Equilibrium Effects of EPL 146

2.7 Robustness Exercises . . . 149
2.7.1 The Value of Leisure, b . . . 149
2.7.2 The Elasticity of Matches to Unemployment, η . . . 151
2.8 Conclusion . . . 152
References . . . 153
2.A Additional Figures . . . 155
2.B Proofs . . . 159
2.B.1 Proof of Proposition 11 . . . 159
2.C Numerical Implementation . . . 159
2.C.1 Algorithm to solve the model . . . 159
2.D Data Appendix . . . 160
2.D.1 Overview . . . 160
2.D.2 Sample Selection . . . 161
2.D.3 Variable Definitions . . . 161

3 Optimal Taxation and Fertility Policies **163**

3.1 Introduction . . . 163
3.2 A Model of Fertility Choice . . . 168
3.3 The Planning Problem . . . 170
3.3.1 Elasticity Concepts . . . 172
3.3.2 Wedges in the Optimal Allocation . . . 173
3.4 Optimal Fertility Policies Under Quasi-linearity . . . 174
3.4.1 The Optimal Income Wedge, $\tau_y^*(\theta)$ . . . 175

3.4.2 The Optimal Child Quantity Wedge, $\tau_n^*(\theta)$ 175
3.4.3 The Optimal Child Goods Investment Wedge, $\tau_K^*(\theta)$ 177
3.4.4 The Optimal Child Time Investment Wedge, $\tau_H^*(\theta)$ 179
3.5 Optimal Fertility Policies Without Quasi-linearity 180
3.5.1 Elasticity Concepts 180
3.5.2 Optimal Wedges 181
3.6 Connection to Optimal Commodity Taxation 185
3.7 Implementation 187
3.8 Quantitative Exercise 189
3.8.1 Baseline Calibration 189
3.8.2 Results 192
3.9 Implementation and Policy Comparisons 194
3.9.1 Implementation: computing a version of $T^*(y, K, n)$ 194
3.9.2 A Review of Tax Policy Towards Families in the USA 196
3.9.3 Results 197
3.10 Welfare Gains From Simple Policies 199
3.10.1 Augmented Planning Problem 199
3.10.2 Quantitative Results 200
3.11 Conclusion 201
References 203
3.A The Planning Problem 207
3.B Proof of Proposition 9 211
3.C Policy Comparisons 212

Chapter 1

A Model of the Asymmetric Transmission of Aggregate Shocks

1.1 Introduction

A growing body of papers studies how departing from the representative household paradigm affects the transmission of aggregate shocks in New Keynesian models. When households are no longer identical, some may be more exposed to the effects of aggregate shocks than others. For example, Guvenen and co-authors (2014, 2017) use high quality administrative income data for the US to document significant heterogeneity in how households' incomes co-move with business cycle movements in GDP. However, this heterogeneity in the incidence of aggregate shocks is somewhat understudied by the existing literature, which focuses on the effects of idiosyncratic and uninsurable income risk on the transmission of aggregate shocks. In this paper, I contribute towards filling this gap.

I study the transmission of aggregate shocks when households' incomes are heterogeneously exposed to changes in aggregate income (output), and borrowing frictions

limit the opportunities for risk sharing. I couple this model of the household sector with a standard New Keynesian supply-side: firms are monopolistically competitive and are subject to costly price adjustments, and nominal interest rates are set according to a Taylor rule. In this setting, I explore the transmission mechanism of aggregate shocks both theoretically and numerically.

Using a simple version of my model, I analytically establish my main result: output responds more to contractionary monetary policy shocks than to expansionary shocks of equal magnitude. This result follows from the fact that households' incomes are heterogeneously exposed to changes in aggregate income, and that borrowing frictions prevent households from fully sharing this aggregate income risk.

In equilibrium, borrowing frictions have two effects. First, they tie adjustments in the real interest rate to the consumption responses of households who are unconstrained in equilibrium. Intuitively, market clearing prices must reflect the choices of household who are able to adjust to price changes on the margin.

Second, borrowing frictions limit the amount of asset trading that occurs in equilibrium. Therefore, households are unable to fully insulate their consumption streams from their heterogeneous exposures to aggregate shocks, causing their consumption paths to partially inherit the heterogeneous exposures of their incomes. Together, these two effects create an asymmetric relationship between output and the real interest rate that lies at the center of my results.

As a simple example, suppose that the elasticity of intertemporal substitution (EIS) is one, and consider a transitory shock to output of $+1\%$. In equilibrium, households' consumptions increase, but by different amounts due to the limited asset trading that occurs. Furthermore, unconstrained households must experience the largest consumption increase among all households. Intuitively, since all other households are borrowing-constrained, their consumption increase must be smaller by definition.

Hence, unconstrained households experience a consumption increase larger than 1%. The Euler equation of these households then implies that the real interest rate must fall by more than 1% in equilibrium.

In response to a transitory shock to output of -1%, households' consumptions decrease by different amounts in equilibrium. In this case, unconstrained households experience the smallest consumption drop among all households since the consumption drop of borrowing-constrained households must be greater by definition. Therefore, unconstrained households experience a drop in consumption of less than 1%. Their Euler equation then implies that the real interest rate must rise by less than 1% in equilibrium.

Crucially, this asymmetry mechanism applies in both flexible and sticky price economies. When prices are sticky, the monetary authority controls the real rate via her choice of nominal rate. In this scenario, the mechanism implies that a 1% interest rate cut will cause output to increase by less than 1%, while a 1% interest rate hike will cause output to fall by more than 1%. Therefore, output responds asymmetrically to monetary policy shocks.

I establish analytically that this mechanism applies to the transmission of two other aggregate shocks commonly used in the New Keynesian literature: TFP shocks and cost-push shocks (direct shocks to inflation). In each case, the transmission of the shock that increases output in equilibrium is weaker than the transmission of an equal and opposite shock that decreases output. I also study the responses of inflation to these shocks and show that the direction of the asymmetry depends the type of shock hitting the economy: inflation inherits the output asymmetry in response to monetary policy shocks, but exhibits the opposite asymmetry pattern for TFP and cost-push shocks.

I confirm that output response asymmetry is robust to the introduction of idiosyncratic risk. I show that, in the empirically relevant case of very persistent idiosyncratic shocks, the asymmetry of the output responses to monetary policy shocks is unaffected by idiosyncratic risk. Intuitively, when idiosyncratic shocks are very persistent, a household does not expect her consumption to change for idiosyncratic reasons, so that her incentive to borrow or save is driven mainly by her income sensitivity to changes in output, as in the case without idiosyncratic risk.

Similarly, I show that output response asymmetry occurs when households have heterogeneous EISs. In this case, binding borrowing constraints restrict the increase in consumption of high EIS households in response to expansionary shocks, and also prevent low EIS households from achieving small consumption declines in response to contractionary shocks.

The asymmetry mechanism relies on the fact that unconstrained households' consumptions must be the most exposed to increases in output, but the least exposed to falls in output. Furthermore, the size of the asymmetry is explicitly linked to the range of these consumption exposure coefficients in the population of households. Using micro data on household consumption from the Consumer Expenditure Survey, I estimate a lower bound for the ratio of these coefficients of 2.9. Inserting this ratio into the model then implies that the output response to a contractionary monetary policy shock is approximately three times larger than the response to an expansionary monetary policy shocks of equal size.

I finish by showing that an asymmetry of this magnitude is consistent with the macro-econometric evidence for asymmetric output responses to monetary policy shocks. Using local projection methods (Jorda, 2005), I estimate that the maximal response of output to a 1% contractionary monetary policy shock is approximately five times larger than the maximal response to a 1% expansionary shock. Therefore,

the quantitative mechanism is capable of explaining at least 60% of the empirical asymmetry.

Related Literature I contribute to a growing literature that studies how departing from the representative household paradigm affects the transmission of monetary policy and other aggregate shocks in New Keynesian models. A large body of work has replaced the representative household with the assumption that households face idiosyncratic and uninsurable income risk that causes ex-ante identical households to experience ex-post heterogeneous time paths of income and consumption. This class of Heterogeneous Agent New Keynesian (HANK) models has been used to study the decomposition of monetary transmission (Auclert, 2017; Kaplan et al., 2018), the power of forward guidance (McKay et al., 2016; Werning, 2015), and the determinacy of interest rate rules (Acharya and Dogra, 2018), among other issues. Relative to these papers, I consider a novel dimension of household heterogeneity, and study its implications for the transmission of a variety of aggregate shocks, not limited to monetary policy.

An important exception in the extant literature, and the closest forebear to my paper, is Bilbiie (2018), who analyzes the transmission of monetary policy in a tractable class of Two Agent New Keynesian (TANK) models. In this class of models, the first set of households do not face any frictions in the asset market, while the second set face severe frictions that prevent them from both borrowing and saving. These households thus live "hand-to-mouth", and consume their entire income in each period with a marginal propensity to consume of one. Using this set up, Bilbiie shows how the response of output to interest rate changes is amplified when changes in aggregate income fall mainly on the hand-to-mouth households, and dampened otherwise.

While Bilbiie's model also features heterogeneous income exposures, the responses of output remain symmetric in his framework. The lack of asymmetry follows from

the extreme way in which asset markets are modeled in TANK frameworks: the first group of households has complete access to asset markets, while the second is completely barred from borrowing or saving any amount. This assumption implies that only the first group of households adjust their consumption in response to interest rate changes. Hence, the equilibrium output response simply coincides with the consumption response of this group, and is therefore symmetric in the sign of the interest rate change. I show that relaxing this assumption so that all households only face constraints to borrowing generates asymmetric transmission of aggregate shocks that aligns well with the macro-econometric evidence for such asymmetry.

Finally, in contemporaneous work, Patterson (2019) argues that contractionary shocks are amplified when households who are highly exposed to the fall in aggregate income also have high marginal propensities to consume (MPC), a fact that she documents in the data. My results complement this empirical finding by providing a structural theory of the amplification of contractionary shocks, and by exploring its consequences for expansionary shocks, thus establishing my key result on asymmetric transmission.

The paper proceeds as follows: section 1.2 describes the economic environment. Section 1.3 establishes my main theoretical result on asymmetric output responses, and discusses extensions and robustness. I estimate key model parameters in section 1.4, and compare the implied output response asymmetry to the empirical evidence in section 1.5. Section 1.6 concludes.

1.2 Environment

The economy is populated by a unit mass of households indexed by $i \in [0, 1]$. Each household has preferences over her infinite sequence of final good consumption $\{c_{i,t}\}$

and labor supply hours $\{n_{i,t}\}$ given by

$$\mathbb{E}_1 \left[\sum_{t=1}^{\infty} \delta^{t-1} u\left(c_{i,t}, n_{i,t}\right) \right]$$

where $\delta \in (0,1)$ is a time discount factor, $\mathbb{E}_t$ is an expectations operator conditioned on time t information, and u is strictly increasing and concave in c and strictly decreasing and concave in n.

If household i works for $n_{i,t}$ hours, she supplies $\theta_{i,t} n_{i,t}$ units of effective labor, where $\theta_{i,t}$ is her labor productivity and is subject to idiosyncratic shocks, as described below. Effective labor earns the nominal wage $P_t w_t$ where P_t is the nominal price of the final consumption good, and w_t is the real wage. In addition to wage income, household i receives a fixed share s_i of dividends from intermediate goods firms d_t measured in consumption units.[1]

When households have heterogeneous dividend shares and labor productivities, their incomes are heterogeneously sensitive to changes in aggregate income. For example, if total wage income increases more than total dividend income, households with higher labor productivities will benefit more than households with higher dividend shares. In order to smooth these heterogeneous exposures to aggregate income fluctuations and to insure against idiosyncratic income shocks, I assume that households can trade a nominal, risk-less one period bond $b_{i,t}$, that earns the real rate r_t given by

$$1 + r_t = \frac{1 + \iota_{t-1}}{1 + \pi_t}$$

where ι_t is the nominal interest rate, and $\pi_t = \frac{P_t - P_{t-1}}{P_{t-1}}$ is inflation. All households begin with zero assets, $b_{i,0} = 0$ for all i.

[1] Interpreting s_i as a household's equity holdings, I assume that trading equity is sufficiently costly so that households do not trade at business cycle frequencies in response to aggregate shocks.

In this economy, households face income risk due to idiosyncratic and aggregate shocks. Without additional restrictions on the structure of financial markets, households could achieve large amounts of self-insurance against these shocks using only the 1 period bond (Krusell and Smith, 1998). Therefore, I follow the literature on incomplete markets and heterogeneous households and assume that households are subject to ad hoc borrowing constraints,

$$b_{i,t} \geq -\underline{b}_{i,t}$$

where $\underline{b}_{i,t} \geq 0$ for all i, t. The fact that the constraint may depend on time and the identity of each household captures, in reduced form, the fact that different households may face borrowing frictions of varying severity at different points in time. For example, higher income households are likely to face less stringent restrictions on their borrowing capacity than lower income households.

In sum, and taking all prices and dividends as given, household i solves

$$\max_{\{c_{i,t}, n_{i,t}, b_{i,t}\}_t} \mathbb{E}_1 \left[\sum_{t=1}^{\infty} \delta^{t-1} u\left(c_{i,t}, n_{i,t}\right) \right]$$

subject to

$$c_{i,t} + b_{i,t} = w_t \theta_{i,t} n_{i,t} + s_i d_t + \left(1 + r_t\right) b_{i,t-1}$$

$$b_{i,t} \geq -\underline{b}_{i,t}$$

$$b_{i,0} = 0$$

$$\theta_{i,0} = \theta_i$$

A representative competitive final good firm packages the unit mass of intermediate goods indexed by $j \in [0,1]$, using the CES production function

$$Y_t = \left(\int_0^1 y_t(j)^{\frac{\Phi_t - 1}{\Phi_t}} dj \right)^{\frac{\Phi t}{\Phi t - 1}}$$

where $\Phi_t > 1$ is the elasticity of substitution across intermediate inputs, and is subject to aggregate shocks. Taking the price of each input and the price of the final good as given, the firm solves

$$\max_{\{y(j)\}_j} P_t \left(\int_0^1 y_t(j)^{\frac{\Phi_t - 1}{\Phi_t}} dj \right)^{\frac{\Phi t}{\Phi_t - 1}} - \int_0^1 p_t(j) y_t(j) dj$$

Optimization yields a demand function for each intermediate good

$$y_t(j) = \left(\frac{p_t(j)}{P_t} \right)^{-\Phi_t} Y_t$$

and a nominal price index

$$P_t = \left(\int_0^1 p_t(j)^{1-\Phi_t} dj \right)^{\frac{1}{1-\Phi_t}}$$

Each intermediate good j is produced by a monopolistically competitive firm employing effective labor $E_t(j)$ in the production function

$$y_t(j) = A_t E_t(j)$$

where A_t is aggregate TFP and is also subject to aggregate shocks. Each firm faces its own demand curve, and chooses its path of prices to maximize profits subject to quadratic price adjustment costs (Rotemberg, 1982):

$$\max_{\{p_t(j)\}_t} \mathbb{E}_1 \left[\sum_{t=1}^{\infty} \delta^{t-1} \left(p_t(j) y_t(j) - P_t w_t E_t(j) - \frac{\xi^p}{2} \left(\frac{p_t(j)}{p_{t-1}(j)} - 1 \right)^2 P_t Y_t \right) \right]$$

subject to

$$y_t(j) = \left(\frac{p_t(j)}{P_t} \right)^{-\Phi_t} Y_t$$

$$y_t(j) = A_t E_t(j)$$

$$p_0(j) = P_0$$

where, for simplicity, I assume that firms discount future profits using the discount factor of households.[2] I focus on the symmetric equilibrium in which $p_t(j) = P_t$, $y_t(j) = Y_t$, and $E_t(j) = E_t$ for all $j \in [0,1]$. In this case, the aggregate dividend in period t is given by

$$d_t = Y_t \left(1 - \frac{w_t}{A_t} - \frac{\xi^p}{2} \pi_t^2 \right)$$

A monetary authority sets the nominal interest rate on the asset according to the Taylor rule

$$1 + \iota_t = \frac{1}{\delta} (1 + \pi_t)^{\phi_\pi} \left(\frac{Y_t}{Y^T} \right)^{\phi_y} e^{v_t}$$

where $\phi_\pi > 1$, $\phi_y \geq 0$, v_t is a monetary policy shock, and Y^T is some fixed target level of aggregate income (output).[3]

In each period, the labor, final good, and bond market must clear:

$$E_t = \int_0^1 \theta_{i,t} n_{i,t} di$$

$$\int_0^1 c_{i,t} di = Y_t \left(1 - \frac{\xi^p}{2} \pi_t^2 \right)$$

[2]This assumption is innocuous for my theoretical results since I approximate around equilibria in which the true stochastic discount factor is constant.

[3]Throughout the paper, I refer to output and aggregate income interchangeably.

$$\int_0^1 b_{i,t} di = 0$$

Households are subject to idiosyncratic shocks to their labor productivity. Specifically, labor productivity for household i follows the AR(1) process

$$\log \theta_{i,t} = \rho_\theta \log \theta_{i,t-1} + (1 - \rho_\theta) \log \theta_i + \epsilon_{i,t}$$

where $\rho_\theta \in (0, 1)$ and $\epsilon_{i,t}$ is an i.i.d. random variable with mean zero and variance σ_ϵ^2. Aggregate shocks affect aggregate TFP, the elasticity of substitution among intermediate inputs (the source of so called "cost-push" shocks), and the innovations to monetary policy, all of which evolve as AR(1) processes,

$$\log A_t = \rho_a \log A_{t-1} + (1 - \rho_a) \log \bar{A} + \epsilon_t^a$$

$$\log \Phi_t = \rho_\Phi \log \Phi_{t-1} + (1 - \rho_\Phi) \log \bar{\Phi} + \epsilon_t^\Phi$$

$$v_t = \rho_v v_{t-1} + \epsilon_t^v$$

where $\rho_a, \rho_\Phi, \rho_v \in (0, 1)$, and $\left\{\epsilon_t^a, \epsilon_t^\Phi, \epsilon_t^v\right\}_t$ are i.i.d. random variables, each with mean zero and variance Σ^2. Due to their effects on inflation, and following the New Keynesian convention, I refer to $\epsilon_t^\Phi < 0$ as a positive cost-push shock and $\epsilon_t^\Phi > 0$ as a negative cost-push shock.[4]

[4]Intuitively, $\epsilon_t^\Phi > 0$ increases the elasticity of demand faced by each monopolist producer, and hence causes firms to lower their prices.

Equilibrium

Definition 1. Given initial conditions $\left\{\{b_{i,0}, \theta_{i,0}\}_i, P_0\right\}$, an equilibrium is a sequence $\left\{\{c_{i,t}, n_{i,t}, b_{i,t}\}_{i,t}, \{y_t(j)\}_{j,t}, P_t, d_t, w_t, \iota_t\right\}_t$ such that

1. $\{c_{i,t}, n_{i,t}, b_{i,t}\}_{i,t}$ solve the household problem for each i.
2. $\{y_t(j)\}_{j,t}$ solve the final good firms' problem.
3. $\{P_t\}_t$ solve the intermediate goods firms' problem.
4. $\{d_t\}_t$ satisfies the dividend equation.
5. $\{\iota_t\}_t$ satisfies the Taylor rule.
6. Markets clear at every time $t \geq 1$.

1.3 Asymmetric Transmission of Aggregate Shocks

In this section, I obtain analytical results regarding the responses of output to aggregate shocks. I begin with the case of monetary policy shocks, and then discuss the extension to cost-push and TFP shocks. I also discuss inflation responses, and generalizations of the result to the inclusion of idiosyncratic risk, and heterogeneous preferences.

1.3.1 Asymmetric Output Responses to Monetary Policy Shocks

In order to achieve tractability, I simplify the model economy in various dimensions.

Assumption 1. *Let*

$$\rho_a, \rho_\Phi, \rho_v = 0$$

$$\Sigma \to 0$$

$$u(c,n) = \frac{\left(c - \frac{n^{1+\varphi}}{1+\varphi}\right)^{1-\sigma}}{1-\sigma}$$

$$\sigma_\epsilon = 0$$

$$\underline{b}_{i,t} \to 0 \ \forall i,t$$

The first two conditions restrict aggregate stochastic parameters to be i.i.d. over time, and to have "small" variances so that local approximation techniques are valid.

The next two conditions restrict aspects of the household problem. The Greenwood et al. (1988) (GHH) specification of utility is common in the business cycle literature, and is tractable since it sets income effects on labor supply to zero. The forth condition sets idiosyncratic risk to zero so that $\theta_{i,t} = \theta_i$ for all i and t. I relax this restriction in section 1.3.4.

The final condition restricts trading in the asset market, and should be interpreted as follows: as the $\left\{\underline{b}_{i,t}\right\}$ parameters approach zero, the sizes of the asset positions taken by households who would borrow in response to a shock get closer to zero due to the binding borrowing constraint. Since the asset is in zero net supply, in general equilibrium, the asset positions of saving households must also get closer to zero. Therefore, the household budget constraint implies that, as the $\left\{\underline{b}_{i,t}\right\}$ parameters approach zero, the equilibrium consumption choices of a household become well approximated by her income choices. I study the equilibrium under this approximation.

Importantly, this condition is not the same as imposing autarky. Instead, the limit condition implies that households can take arbitrarily small positions in the asset in equilibrium. This then requires that prices and quantities adjust in equilibrium so that the asset market clears. In particular, the adjustment must be such that households who save in equilibrium optimally choose a vanishingly small asset position that offsets the vanishingly small positions taken by borrowing-constrained households. Therefore, this assumption buys tractability without losing the key transmission mechanism from interest rates to savings choices.[5]

Deterministic Equilibrium Under assumption 1, I can define a deterministic equilibrium as an equilibrium when all aggregate shocks are set to zero in all periods. In this equilibrium, inflation is always zero, aggregate prices and quantities are constant, and all household choices of consumption, labor supply, and asset positions are fixed over time since they do not face any risk.

Definition 2. Suppose that assumption 1 holds, and assume that there are no aggregate shocks. Then, given initial conditions $\left\{\{b_{i,0}, \theta_{i,0}\}_i, P_0\right\}$, a deterministic equilibrium is a sequence $\left\{\{c_i, n_i, b_i\}_i, \{y(j)\}_j, P, d, w, \iota\right\}$ such that

1. $\{c_i, n_i, b_i\}_i$ solve the household problem for each i.
2. $\{y(j)\}_j$ solve the final good firms' problem.
3. $P = P_0$ solves the intermediate goods firms' problem.
4. d satisfies the dividend equation.
5. Markets clear at every time $t \geq 1$.

[5]Werning (2015) uses a similar assumption to analyze how the cyclicality of idiosyncratic risk affects the power of forward guidance.

I consider the dynamics of the economy in response to aggregate shocks around this deterministic equilibrium.[6] Formally, the following lemma condenses the economy's equilibrium dynamics to a set of necessary and sufficient conditions expressed as log deviations around the deterministic equilibrium, where I use $\hat{x}_t = \log x_t - \log x$ to denote such a deviation. The proofs of this and all other results are contained in the appendix.

Lemma 1. *Under assumption 1, the economy's first order equilibrium dynamics in response to monetary policy shocks satisfy the system*

$$\iota_t = \rho + \phi_\pi \pi_t + \phi_y \hat{y}_t + \epsilon_t^v$$

$$\pi_t = \frac{\bar{\Phi} - 1}{\xi^p} \varphi \hat{y}_t + \delta \mathbb{E}_t \left[\pi_{t+1}\right]$$

$$\min_i \left\{ \mathbb{E}_t \left[\hat{\tilde{c}}_{i,t+1}\right] - \hat{\tilde{c}}_{i,t} \right\} = \frac{1}{\sigma} \left(\iota_t - \mathbb{E}_t \left[\pi_{t+1}\right] - \rho\right)$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t \ \forall i$$

where $\rho = -\log \delta$, $\{\beta_i^y\}_i$ *depend only on model primitives, and* $\tilde{c}$ *is consumption net of the disutility of labor supply,* $\tilde{c} = c - \frac{n^{1+\varphi}}{1+\varphi}$.

$\phi_\pi > 1$ *and* $\phi_y \geq 0$ *are sufficient to ensure that the system has a unique steady state,* $\hat{y}_t = 0$, $\pi_t = 0$, $\hat{\tilde{c}}_{i,t} = 0 \ \forall i$.

The first two equations of lemma 1 are standard features of New Keynesian models. The first equation is the Taylor rule for the nominal interest rate where the target level of output Y^T is set to the level in the deterministic equilibrium. The second equation is the New Keynesian Phillips Curve (NKPC) linking inflation, and output.

[6] In deterministic economies with heterogeneous households, the wealth distribution may be indeterminate (Sorger, 2000). However, my restriction that $\underline{b}_{i,t} \to 0$ imposes that all households hold zero assets in all periods, thus breaking the indeterminacy, and ensuring uniqueness of the deterministic equilibrium.

Intuitively, if output is higher today holding, firms face higher marginal costs of labor ceteris paribus, and so will optimally choose to raise their prices, leading to inflation.

The final two equations characterize the two effects that borrowing frictions have in equilibrium. The third equation says that equilibrium adjustments in the real interest rate are tied to the consumption responses of unconstrained households. To see this, first note that the right-hand side of the equation is the deviation of the real rate from its steady state value, multiplied by the relevant preference parameter, the elasticity of intertemporal substitution (EIS), $\frac{1}{\sigma}$. This deviation is then equated to the lowest expected consumption growth among all households, which must refer to the expected consumption growth of households who are unconstrained in equilibrium. By definition, borrowing-constrained households must obtain a current consumption level that is excessively low, and hence have an excessively high expected consumption growth rate. Therefore, unconstrained households must have a lower expected consumption growth rate than these households, which then must be the lowest among the population of households in equilibrium.

The forth equation characterizes how the heterogeneous exposures of household income to changes in output transmit into heterogeneous consumption exposures due to the limited asset trading that occurs in equilibrium. The sensitivity coefficients $\{\beta_i^y\}$ measure the per cent change in household i's consumption for a 1% change in output in equilibrium. Intuitively, the pattern of $\{\beta_i^y\}$ depends on the underlying heterogeneity in income sensitivities, and the tightness of the borrowing constraints that restrict asset trading in equilibrium. For example, in the absence of any asset trading frictions, households would fully insulate their consumption from their heterogeneous income sensitivities, and β_i^y would be fixed across i.

Under assumption 1 however, the no-borrowing limit restriction implies the other extreme: a household's consumption inherits the sensitivity of her income to changes

in output. Heterogeneity in consumption sensitivities therefore reflects heterogeneity in the underlying income sensitivities.

There are two sources of heterogeneity in the household income sensitivities, which are most clearly seen using the explicit expression for β_i^y given by

$$\beta_i^y = \frac{\varphi \frac{\bar{\Phi}-1}{\bar{\Phi}} \frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta} + s_i \left(\frac{1}{\bar{\Phi}} - \varphi \frac{\bar{\Phi}-1}{\bar{\Phi}} \right)}{\frac{\varphi}{1+\varphi} \frac{\bar{\Phi}-1}{\bar{\Phi}} \frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta} + s_i \frac{1}{\bar{\Phi}}}$$

where $\Theta = \int_0^1 \theta_i^{\frac{1+\varphi}{\varphi}} di$. First, when output increases, households with higher labor productivities receive more of the corresponding increase in wage income, so that β_i^y is increasing in θ_i. Second, since TFP is fixed in the case of monetary policy shocks, higher wages cause the dividend share of aggregate income to decline, so that β_i^y is decreasing in s_i.

It is simple to show that $\beta_i^y > 0$ when the labor productivity effect dominates the dividend effect on household income and hence consumption,

$$\beta_i^y > 0 \iff \frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta} > s_i \left(1 - \frac{1}{\varphi \left(\bar{\Phi} - 1 \right)} \right)$$

Anticipating the empirical results, which find positive consumption sensitivities for all i, I assume that this condition holds for all i from now on.

Given the consumption response of saver households $i = S$, the equilibrium output response must mechanically satisfy

$$\hat{y}_t = \frac{\hat{\bar{c}}_{S,t}}{\beta_S^y}$$

Hence, if β_S^y is different for contractionary and expansionary shocks, then the equilibrium path of output will be shock-dependent, and asymmetric, as demonstrated by the following example.

A Numerical Example Let there be two household types, $i \in \{1, 2\}$, and assume that prices are fixed, $\xi^p \to +\infty$, so that the central bank directly controls the real interest rate. Suppose that $\phi_y = 0$ and$\frac{1}{\sigma} = 1$, and assume that group 1 households' consumption is more sensitive to changes in output than group 2 households' consumption: $\beta_1^y = 2$, and $\beta_2^y = 0.5$.[7]

Given real interest rate shocks of $\epsilon_t^v = \pm 1\%$, the responses of output are found by solving the system

$$\min_i \left\{ -\hat{\tilde{c}}_{i,t} \right\} = \frac{1}{\sigma} \epsilon_t^v$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t \ \forall i$$

where I have used $\mathbb{E}_t \left[\hat{\tilde{c}}_{i,t+1} \right] = 0$ because the interest rate shock is i.i.d. over time.

When the real interest rate increases by 1%, unconstrained households reduce their consumption by 1% because $\frac{1}{\sigma} = 1$. Since all other households are borrowing-constrained, their equilibrium reductions in consumption are greater than 1%. Hence, the equilibrium reduction in output must be greater than 1%.

Specifically, given output falls in equilibrium, group 2 households must be unconstrained since they experience the smallest drop in current consumption and therefore have the lowest expected consumption growth. Hence, $\hat{c}_{2,t} = -1\%$. Inverting the sensitivity equation for group 2 households, $\hat{\tilde{c}}_{2,t} = 0.5\hat{y}_t$, implies that $\hat{y}_t = \frac{-1\%}{0.5} = -2\%$, so that output falls by 2% in equilibrium.

[7] Technically, the steady state of the system is no longer unique when $\xi^p \to +\infty$ and $\phi_y = 0$. However, I abstract from this complication for the purposes of this example.

When the real interest rate decreases by 1%, unconstrained households increase their consumption by 1%. Since all other households are borrowing-constrained, their equilibrium increase in consumption is less than 1%. Hence, the equilibrium increase in output must be less than 1%.

Specifically, given output rises in equilibrium, group 1 households must be unconstrained since they experience the largest increase in current consumption and therefore have the lowest expected consumption growth. Hence, $\hat{c}_{1,t} = 1\%$. Inverting the sensitivity equation for group 1 households, $\hat{c}_{1,t} = 2\hat{y}_t$, implies that $\hat{y}_t = 0.5\%$ so that output increases by 0.5% in equilibrium.

Therefore, output responds more to contractionary monetary policy shocks than to expansionary monetary policy shocks. Furthermore, because unconstrained households always adjust their consumption growth by 1% in equilibrium, the ratio of the contractionary response to the expansionary response coincides with the ratio of consumption sensitivity coefficients. I come back to this connection in section 1.4.

In general, we have the following closed-form representation of the asymmetric output responses to monetary policy shocks.

Proposition 1. *Under assumption 1, the first order equilibrium dynamics of output in response to monetary policy shocks are given by*

$$\hat{y}_t = \begin{cases} -\frac{1}{\underline{\beta}+\frac{1}{\sigma}\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y}\frac{1}{\sigma}\epsilon_t^v & if \quad \epsilon_t^v > 0 \\ \\ -\frac{1}{\bar{\beta}+\frac{1}{\sigma}\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y}\frac{1}{\sigma}\epsilon_t^v & if \quad \epsilon_t^v < 0 \end{cases}$$

where

$$\bar{\beta} = \max_i \{\beta_i^y\}$$

$$\underline{\beta} = \min_i \{\beta_i^y\}$$

$\bar{\beta} > \underline{\beta}$ implies that output responds more to positive (contractionary) monetary policy shocks than to negative (expansionary) monetary policy shocks of equal magnitude. In other words, output responds asymmetrically to monetary policy shocks.[8]

In response to an expansionary monetary policy shock, interest rates fall, and unconstrained households increase their current consumption. By virtue of being borrowing-constrained in equilibrium, all other households must increase their consumption by a smaller amount. Therefore, the equilibrium increase in output is smaller than the consumption response of unconstrained households alone. Equivalently, unconstrained households' consumption increase is the most sensitive to the increase in output in equilibrium, as captured by the forth equation of the lemma, evaluated for

[8]The equilibrium dynamics exist as long as $\underline{\beta} + \frac{1}{\sigma}\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y > 0$, which is implied by $\beta_i^y > 0$ for all i. If this condition fails, then there does not exist an equilibrium output response to a contractionary monetary policy shock, $\epsilon_t^v > 0$, of the piece-wise linear form presented in proposition 1. However, given $\underline{\beta}$, parameters ϕ_π, ϕ_y, and ξ^p can always be chosen to ensure that the condition holds.

the unconstrained household $i = U$,

$$\hat{\hat{c}}_{U,t} = \bar{\beta}\hat{y}_t, \ \bar{\beta} = \max_i \{\beta_i^y\}$$

In response to a contractionary monetary policy shock, unconstrained households decrease their current consumption. By virtue of being borrowing-constrained in equilibrium, all other households must experience a larger decrease in their consumption. Therefore, the equilibrium decrease in output is larger than the consumption response of unconstrained households alone. Equivalently, unconstrained households' consumption is the least sensitive to changes in output in equilibrium,

$$\hat{\hat{c}}_{U,t} = \underline{\beta}\hat{y}_t, \ \underline{\beta} = \min_i \{\beta_i^y\}$$

Since the size of the consumption response of unconstrained households is the same for both positive and negative equilibrium real interest rate changes, output must respond more to contractionary monetary policy shocks than to expansionary monetary policy shocks.

For completeness, I note that in the knife-edge case of $\bar{\beta} = \underline{\beta}$, there is no asymmetry. In this case, all households' incomes are equally sensitive to changes in aggregate income so that there is no incentive for households to trade the asset in response to an aggregate shock. Therefore, borrowing constraints do not play a role in determining the equilibrium response of output to monetary policy shocks.

1.3.2 Other Aggregate Shocks

It is simple to derive the equivalent of proposition 1 for the cases of cost-push shocks and TFP shocks. In both cases, the same asymmetry emerges: output responds more to shocks that increase interest rates and lower output, and less to shocks that

decrease interest rates and increase output. I briefly sketch the intuition below. Full details of the analysis are in the appendix.

The asymmetry of cost-push shocks follows from the fact that a positive cost-push shock creates inflation which causes interest rates to rise via the Taylor rule, while a negative cost-push shock causes interest rates to fall. These interest rate movements then initiate the same mechanism as above, causing output to respond more to the contractionary movement than to the expansionary movement.

Similarly, a positive TFP shock causes deflation and hence lower interest rates, while a negative TFP shock creates inflation and higher interest rates. Therefore, the same mechanism implies that output will respond more to negative (contractionary) TFP shocks than to positive (expansionary) TFP shocks.

1.3.3 Inflation Responses

Solving the system of equations in lemma 1 yields equilibrium responses of both output and inflation. In contrast to output, the direction of the asymmetry of inflation is shock-dependent. I outline the key economic mechanisms below, and relegate the derivations to the appendix.

In the case of monetary policy shocks, inflation inherits the asymmetry of output. This occurs because the response of inflation is entirely determined by the response of output via the logic of the NKPC: higher output implies higher marginal costs which causes firms to increase their prices, thus raising inflation. Therefore, inflation moves in the same direction as output in response to monetary shocks. Since output responds more to contractionary monetary shocks than to expansionary shocks, inflation inherits this asymmetry.

In the cases of cost-push and TFP shocks, the asymmetry of the inflation response is the opposite to that of output: inflation responds more to shocks that increase output.

This occurs because the responses of inflation are determined by two forces. First, the movement of output affects inflation via the NKPC as in the monetary shock case. Second, both cost-push and TFP shocks directly affect firms' marginal costs and so directly affect inflation (positive cost-push shocks increase marginal costs and inflation, while positive TFP shocks lower marginal costs and inflation). Crucially, this effect pushes inflation in the opposite direction to the movement of output such that the first effect offsets the second. The strength of this offsetting force inherits the asymmetry of output's responses to both types of shock so that inflation responds more overall when output responds less and the offsetting force is weaker. Therefore inflation responds with the opposite asymmetry to output.

1.3.4 Idiosyncratic Shocks

It is straightforward to extend the analysis of this section to the case in which idiosyncratic risk is "small", so that $\sigma_\epsilon \to 0$. In this case, the asymmetry of output responses to monetary policy shocks depends on the persistence of idiosyncratic shocks. In particular, when uninsurable idiosyncratic shocks are very persistence ($\rho_\theta \to 1$), the output response asymmetry is the same as in the case without idiosyncratic risk, so that proposition 1 still applies. I outline the key economic mechanisms below, and relegate the derivations to the appendix.

In the presence of idiosyncratic risk, there is an additional channel through which households can have low consumption growth, and hence be unconstrained in equilibrium. When a household expects to experience a drop in her idiosyncratic labor productivity, her consumption growth will be low to the extent that the drop in labor productivity is uninsurable and hence transmits to her consumption. Importantly, this novel channel is independent of aggregate shocks hitting the economy, and so cannot be a source of asymmetry.

However, when the process for uninsurable labor productivity shocks is very persistent ($\rho_\theta \to 1$), households expect their labor productivity to remain approximately constant across consecutive periods. This renders the novel savings channel inactive. As a result, savings choices are entirely determined by household income sensitivities to changes in output. Therefore, the responses of output to monetary policy shocks are identical to the economy without idiosyncratic risk.

The empirical evidence suggests that the process for idiosyncratic shocks to labor income has a very persistent component, so that $\rho_\theta \to 1$ is a good approximation to reality (see, for example, Storesletten et al., 2004; Guvenen et al., 2016). Furthermore, related evidence suggests that households are very well insured against shocks to the transitory component (Blundell et al., 2008; Heathcote et al., 2014), so that these shocks do not affect consumption growth computations.

In order to obtain analytical insights, I assume that the idiosyncratic risk is "small" so that first order approximations are valid. This approach rules out second order phenomena, such as the cyclicality of idiosyncratic income risk, that may also effect the responses of output to monetary policy shocks (and other aggregate shocks). These effects are the focus of Werning (2015) and Acharya and Dogra (2018), who show how the cyclicality of idiosyncratic income risk affects the size of the output responses to both positive and negative monetary policy shocks. Importantly, this effect is symmetric in the sign of the shock, and so does not affect the asymmetry that I am interested in.

1.3.5 Heterogeneous Preferences

In my benchmark model, I follow the New Keynesian literature and assume that households have homogeneous EISs given by $\frac{1}{\sigma}$.[9] However, the asset pricing literature

[9]Technically, $\frac{1}{\sigma}$ is the EIS for "net consumption" when preferences are of the GHH form. However, abstracting from this complication does not affect the intuition.

has suggested that heterogeneous EISs may help to reconcile macroeconomic models and asset pricing facts (Guvenen, 2009), and have empirical grounding (Vissing-Jorgensen, 2002; Guvenen, 2006). In the case of heterogeneous EISs, the asymmetry still emerges, but for different reasons. I outline the key mechanism below, and relegate the derivations to the appendix.

When households have heterogeneous EISs, their consumption responses to changes in the real interest rate are heterogeneous. In equilibrium, these different consumption responses occur via asset trading. For example, in response to a contractionary monetary policy shock, households with larger EISs will save, lending to households with smaller EISs in equilibrium.

However, binding borrowing constraints limit the asset trading that occurs in equilibrium and causes output to respond asymmetrically. Consider a contractionary shock. When the real interest rate falls, households try to decrease their consumption by an amount dictated by their EIS. However, households with small EISs become borrowing-constrained in equilibrium and so experience larger consumption drops than they would in the absence of the constraint. Therefore, the overall equilibrium decrease in output is amplified.

When interest rates rise after an expansionary shock, households increase their consumption. In this case, households with large EISs become borrowing-constrained in equilibrium, and so experience smaller consumption increases then they would were they not constrained. Therefore, the overall equilibrium increase in output is dampened, thus creating asymmetric responses of output to monetary policy shocks.

1.4 Quantitative Exercise

In this section, I quantitatively assess the asymmetry of output responses to monetary policy shocks. I first generalize my theoretical insights to highlight the key parameters

that I need to measure to quantify the asymmetry, and then turn to parameter estimation.

1.4.1 Generalizing the Mechanism

The theoretical analysis has shown that the mechanism generating the output response asymmetry has two building blocks: the Euler equation that holds when some households are borrowing-constrained, and the presence of heterogeneous consumption sensitivities to changes in output in equilibrium. While proposition 1 was derived under assumption 1, the intuition suggests that such stringent restrictions are not necessary to generate output response asymmetry.

In order to generalize the asymmetry mechanism while maintaining tractability, I now adopt a semi-structural approach.[10] I first show that the Euler equation in the simple model continues to hold in a wide class of economic models. I then combine this structural equation with a reduced form description of household consumption that generalizes the forth equation from the simple model to allow for idiosyncratic shocks to consumption. While this equation does not have explicit structural foundations, it maintains a close connection to the empirical consumption literature, and highlights the key moments that I ultimately need to estimate in order to quantitatively assess the asymmetry mechanism.

The Euler Equation The Euler equation stems from household consumption and savings decisions. To this end, I now consider a household who chooses consumption, labor supply, and asset positions in a nominal risk-free bond and $J \geq 0$ other arbitrary assets to maximize her infinite horizon utility subject to a budget constraint and a

[10]Semi-structural approaches are somewhat common in the inflation dynamics literature, for example Kichian et al. (2010). A structural New Keynesian Phillips Curve is combined with a reduced form equation for firms' marginal costs in order to improve estimation efficiency and forecasting properties of the model.

general asset position constraint, as in Berger et al. (2019). Let $j \in \mathcal{J}$ number the arbitrary additional assets, where $|\mathcal{J}| = J$. A household i solves

$$\max_{\left\{c_{i,t}, n_{i,t}, b_{i,t}, \left\{a_t^j\right\}_j\right\}_t} \mathbb{E}_1 \sum_{t=1}^{\infty} \delta^{t-1} \left(\frac{c_{i,t}^{1-\sigma}}{1-\sigma} - \psi \frac{n_{i,t}^{1+\varphi}}{1+\varphi} \right)$$

subject to

$$c_{i,t} + b_{i,t} + \sum_{j \in \mathcal{J}} q_t^j a_{i,t}^j = w_t \theta_{i,t} n_{i,t} + d_{i,t} + (1 + r_t) b_{i,t-1} + \sum_{j \in \mathcal{J}} a_{i,t-1}^j$$

$$\mathcal{B}_{i,t} \left(b_{i,t}, \left\{ a_{i,t}^j \right\} \right) \geq 0$$

where the $\mathcal{B}_{i,t}$ function has the property that

$$\frac{\partial \mathcal{B}_{i,t}}{\partial b} \left(b_{i,t}, \left\{ a_{i,t}^j \right\} \right) \geq 0$$

This problem nests a large class of incomplete market models: households choose paths for consumption and labor, together with positions in a risk-free asset and an arbitrary set $\mathcal{J}$ of other assets to maximize their utility subject to a budget constraint and asset market frictions, summarized by the second and third constraints. The second constraint is a general constraint on asset positions, while the third constraint implies that taking a more positive position in the risk-free asset loosens the constraint, i.e. that the risk-free asset is subject to borrowing frictions of some kind.

The FOCs for the risk-free asset imply that

$$\mathbb{E}_t \left[\delta \left(\frac{c_{i,t+1}}{c_{i,t}} \right)^{-\sigma} (1 + r_{t+1}) \right] \leq 1$$

for all households, with strict inequality when the borrowing constraint binds for the risk-free asset choice $b_{i,t}$. In general equilibrium, at least one household must be unconstrained, so that

$$\max_i \mathbb{E}_t \left[\delta \left(\frac{c_{i,t+1}}{c_{i,t}} \right)^{-\sigma} (1 + r_{t+1}) \right] = 1$$

Taking a first order log-linear expansion of this condition yields the Euler equation from the simple model (written in growth rate form),

$$\min_i \left\{ \mathbb{E}_t \left[\Delta \log c_{i,t+1} \right] \right\} = \frac{1}{\sigma} \left(\iota_t - \mathbb{E}_t \left[\pi_{t+1} \right] - \rho \right)$$

which therefore holds in much more general economic environments.

Heterogeneous Consumption Sensitivities Given this Euler equation, the asymmetry in the simple model came from heterogeneity in the consumption sensitivity parameters, β_i^c, in the equation

$$\Delta \log c_{i,t} = \beta_i^c \Delta \hat{y}_t$$

where I have written the equation in growth rate form. The derivation of this relationship relied on the lack of idiosyncratic risk, and the borrowing constraints going to zero. Unfortunately, departing from these assumptions breaks the tractability of the model and prevents a similarly simple equation from being derived in more general models.

In order to make progress, I extend the simple expression for household consumption growth to allow for idiosyncratic shocks,

$$\Delta \log c_{i,t} = \beta_i^{c,v} \Delta \hat{y}_t + u_{i,t}$$

where the evidence for the transmission of idiosyncratic shocks from income to consumption (Blundell et al., 2008; Heathcote et al., 2014) suggests that $u_{i,t}$ is mean independent across time, such that $\mathbb{E}_t [u_{i,t+1}] = 0$. As described in section 1.3.4, this condition implies that idiosyncratic shocks do not affect the asymmetry mechanism.[11]

Although this equation is reduced form in nature, it captures the two dimensions in which household consumption deviates from aggregate consumption (output) in equilibrium: systematic heterogeneity in exposure to aggregate shocks, and idiosyncratic shocks to consumption. Furthermore, this specification is simple to estimate given data on household consumption and output. I exploit this feature in the estimation section below.

1.4.2 A Semi-Structural Model

In order to derive the asymmetry result in a more general setting using the two building blocks above, I now combine these two equations with a Taylor rule for the nominal interest rate, and a New Keynesian Phillips Curve that links inflation and output. While the Taylor rule is a standard description of monetary policy, the Phillips curve specification is reduced form in nature. This reflects the fact that models without GHH preferences do not admit closed form expressions for the Phillips curve.[12]

[11] I verify that the idiosyncratic shocks satisfy this condition as part of the estimation process.

[12] In recent work, Berger et al. (2019) show that household heterogeneity manifests in the slope parameter of the NKPC, κ_y. Given my focus is on household heterogeneity at the Euler equation, I abstract from this channel, and take κ_y as a fixed parameter that I calibrate using the relevant data.

Therefore, let $\mathcal{S}$ be the system

$$\iota_t = \rho + \phi_\pi \pi_t + \phi_y \hat{y}_t + v_t$$

$$\pi_t = \kappa_y \hat{y}_t + \delta \mathbb{E}_t \left[\pi_{t+1}\right]$$

$$\min_i \left\{\mathbb{E}_t \left[\Delta \log c_{i,t+1}\right]\right\} = \frac{1}{\sigma}\left(\iota_t - \mathbb{E}_t \left[\pi_{t+1}\right] - \rho\right)$$

$$\Delta \log c_{i,t} = \beta_i^{c,v} \Delta \hat{y}_t + u_{i,t} \ \forall i$$

that describes the equilibrium dynamics of output $\hat{y}_t$, inflation π_t, nominal interest rates ι_t, and household consumption $\hat{c}_{i,t}$, in response to monetary policy shocks v_t.

In the empirically relevant case of persistent monetary policy shocks (Gertler and Karadi, 2015; Christiano et al., 2005), the inherent non-linearity of system $\mathcal{S}$ prevents me from studying the full equilibrium dynamics. Therefore, I instead derive the responses of output to one-time, zero probability ("MIT") monetary policy shocks of the form $v_t = \rho_v^{t-1} v_1$, where v_1 is the zero-probability monetary policy shock, and $\rho_v \in (0, 1)$ is a persistence parameter. This approach is tractable since once the shock hits, the economy transitions deterministically back to the steady state of the system.

Proposition 2. *In system $\mathcal{S}$, the response of output to a one time, zero probability monetary policy shock v_1 with persistence ρ_v, is given by*

$$\hat{y}_t = \begin{cases} -\frac{1}{(1-\rho_v)\underline{\beta}^{c,v} + \frac{1}{\sigma}(\phi_\pi - \rho_v)\frac{\kappa_y}{1-\delta\rho_v} + \frac{1}{\sigma}\phi_y} \frac{1}{\sigma}\rho_v^{t-1} v_1 & \text{if} \quad v_1 > 0 \\ \\ -\frac{1}{(1-\rho_v)\bar{\beta}^{c,v} + \frac{1}{\sigma}(\phi_\pi - \rho_v)\frac{\kappa_y}{1-\delta\rho_v} + \frac{1}{\sigma}\phi_y} \frac{1}{\sigma}\rho_v^{t-1} v_1 & \text{if} \quad v_1 < 0 \end{cases}$$

where

$$\underline{\beta}^{c,v} = \min_i \{\beta_i^{c,v}\}$$

$$\bar{\beta}^{c,v} = \max_i \{\beta_i^{c,v}\}$$

As in the case of i.i.d shocks, output responds more to the contractionary monetary policy shock than to the expansionary shock. Appealing to the structural interpretation of the system yields the same intuition as before: when $v_1 > 0$, the decrease in savers' consumption is the smallest among all households, so that output falls a lot. When $v_1 < 0$, the increase in savers' consumption is the largest among all households so that output increases a little.

In order to assess the quantitative magnitude of the output response asymmetry, I define the ratio of the contractionary response to the expansionary response,

$$\mathcal{R} = \frac{(1-\rho_v)\,\bar{\beta}^{c,v} + \frac{1}{\sigma}\,(\phi_\pi - \rho_v)\,\frac{\kappa_y}{1-\delta\rho_v} + \frac{1}{\sigma}\phi_y}{(1-\rho_v)\,\underline{\beta}^{c,v} + \frac{1}{\sigma}\,(\phi_\pi - \rho_v)\,\frac{\kappa_y}{1-\delta\rho_v} + \frac{1}{\sigma}\phi_y}$$

The key parameters for quantifying the asymmetry are $\bar{\beta}^{c,v}$ and $\underline{\beta}^{c,v}$, which measure the highest and lowest equilibrium sensitivities of household consumption to changes in output. In particular, when $\frac{1}{\sigma}\,(\phi_\pi - \rho_v)\,\frac{\kappa_y}{1-\delta\rho_v} + \frac{1}{\sigma}\phi_y = 0$, the response ratio is simply the ratio of the consumption sensitivities, $\mathcal{R} = \frac{\bar{\beta}^{c,v}}{\underline{\beta}^{c,v}}$.

The parameters $\bar{\beta}^{c,v}$ and $\underline{\beta}^{c,v}$ are "sufficient statistics" for computing the output response asymmetry (Chetty, 2009).[13] In other words, to compute $\mathcal{R}$, I only need to know the values of $\bar{\beta}^{c,v}$ and $\underline{\beta}^{c,v}$, and do not need quantitative information on the underlying structural mechanism that generates them. In my setting, this means that I do not need to know quantitative details concerning borrowing constraints or the heterogeneity of income sensitivities. This is convenient because I can estimate $\bar{\beta}^{c,v}$ and $\underline{\beta}^{c,v}$ directly using micro data on household consumption, and then plug these estimates into $\mathcal{R}$ to immediately quantify the asymmetry.

1.4.3 Parameter Estimation

Estimates for $\bar{\beta}^{c,v}$ and $\underline{\beta}^{c,v}$ are obtained by first estimating the sensitivities of household consumption to changes in output driven by monetary policy shocks, and then taking the maximum and minimum of these estimates.

Let $\{c_{i,t}\}_{i,t}$ and $\{Y_t\}_t$ be data on household consumption and output respectively. The forth equation of the system $\mathcal{S}$ suggests estimating linear regressions of the form

$$\Delta \log c_{i,t} = \alpha_i + \beta_i \Delta \log Y_t + u_{i,t}$$

where β_i measures the sensitivity of household i's consumption to changes in output. In this specification, $\{u_{i,t}\}$ capture idiosyncratic consumption shocks and measurement errors in consumption.

In order to measure the sensitivities of household consumption to changes in output driven by monetary policy shocks, the variation in $\Delta \log Y_t$ must be due to monetary policy shocks only. However, the variation in raw output data is driven by multiple aggregate shocks hitting the economy simultaneously in each period. Running the

[13] Sufficient statistics approaches have recently become popular in macroeconomics. See, for example, Auclert and Rognlie (2017).

above regression would therefore result in estimates of $\{\beta_i\}$ that measure the sensitivity of household consumption to changes in output driven by multiple shocks, and would not correspond to the theoretical parameters $\{\beta_i^{c,v}\}$.

In order to alleviate this issue, I first project the output data onto a set of identified, lagged monetary policy shocks (described in more detail below), $Z_t = \left(\epsilon_{t-1}^v, ..., \epsilon_{t-L}^v\right)$,

$$\Delta \log Y_t = \alpha_y + Z_t'\gamma + e_t$$

so that the fitted values $\left\{\Delta \hat{\log} Y_t\right\}$ capture the variation in $\Delta \log Y_t$ driven by monetary policy shocks only. This specification is consistent with the structural vector-autoregression paradigm, in which aggregate variables are expressible as a moving average of the (infinite) history of structural shocks (see Barnichon and Matthes (2016) for a review and extension of this approach to the non-linear case).[14]

Using these fitted values, I then consider the second stage regression

$$\Delta \log c_{i,t} = \alpha_i + \beta_i \Delta \hat{\log} Y_t + u_{i,t}$$

which correctly identifies β_i as the sensitivity of household consumption to changes in output driven by monetary policy shocks only.

Intuitively, this process amounts to Two-Stage-Least-Squares (2SLS) estimation, where the first stage extracts the variation in $\Delta \log Y_t$ due to monetary policy shocks only, and the second stage estimates the household sensitivity parameters using this variation alone.

[14]My theoretical results suggest that $\Delta \log Y_t$ should depend non-linearly on the history of monetary policy shocks. However, for the purposes of extracting the variation in $\Delta \log Y_t$ driven by monetary policy shocks, I abstract from this complication. I investigate non-linear responses in section 2.3.

Household Consumption Data Estimating the above regression requires a sufficiently long panel dataset on household consumption at business cycle frequency. The Consumer Expenditure Survey (CEX) is the most suitable dataset, but is known to have measurement error problems (Aguiar and Bils, 2015), and only features the same household for four consecutive quarters.

In order to alleviate these issues, I group households together within the CEX data and estimate pooled OLS regressions instead. This approach helps to mitigate the effects of measurement error in the cross-section, and creates longer synthetic panels in the time series dimension. Similar methods are common among analyses that use CEX data to analyze trends and fluctuations in household consumption (see, for example, Parker and Vissing-Jorgensen, 2009; Primiceri and van Rens, 2009; De Giorgi and Gambetti, 2017).

As mentioned, $u_{i,t}$ captures both idiosyncratic shocks to consumption and measurement error at the household level. While the idiosyncratic shock component is uncorrelated with $\Delta \hat{\log} Y_t$ by definition, the measurement error component may not be, and so could induce bias into the estimates of consumption sensitivities. However, as long as measurement error is independent across households, and each group consists of a sufficiently large number of households in each period, applying a cross-sectional Law of Large Numbers implies that the composite measurement error term is approximately zero in every period for a given group, and is therefore uncorrelated with $\Delta \hat{\log} Y_t$. I assume that this condition holds in my analysis.

The choice of grouping naturally affects the estimates obtained from running regressions at the group level. Formally, let $\mathcal{G}$ be a surjective function that maps household i in period t, i.e. the household-period tuple (i, t), into a finite set of groups $\{1, 2, ..., G\}$. $\mathcal{G}$ represents an arbitrary group formation process, and nests fixed group assignment as a special case, $\mathcal{G}(i, t)$ fixed for all t.

Given a choice of $\mathcal{G}$ function, consider the pooled OLS regression for a group $g \in \{1, 2, ..., G\}$,

$$\Delta \log c_{i,t} = \alpha_g + \beta_g \Delta \hat{\log} Y_t + e_{i,t}$$

where the pooling occurs over the set $\{(i,t) : \mathcal{G}(i,t) = g\}$ of household-periods assigned to group g. Estimating this regression for each group implies that the key parameters for quantifying the asymmetry can be estimated as $\bar{\beta}^{c,v} = \max_g \left\{\hat{\beta}_g\right\}$ and $\underline{\beta}^{c,v} = \min_g \left\{\hat{\beta}_g\right\}$.

When the $\mathcal{G}$ function assigns each household i to a fixed group over time, the implied asymmetry parameters will always be weakly bounded by the true asymmetry parameters $\max_i \{\beta_i\}$ and $\min_i \{\beta_i\}$. Therefore, pooled OLS using fixed group assignments will always weakly underestimate the true asymmetry.

Proposition 3. *Suppose the model for household consumption growth is given by*

$$\Delta \log c_{i,t} = \alpha_i + \beta_i \Delta \hat{\log} Y_t + u_{i,t}$$

If $\mathcal{G}$ does not depend on t for all i, then the asymmetry parameters implied by the pooled OLS regressions

$$\Delta \log c_{i,t} = \alpha_g + \beta_g \Delta \hat{\log} Y_t + e_{i,t}$$

are weakly bounded by $\max_i \{\beta_i\}$ *and* $\min \{\beta_i\}$*, i.e.*

$$\max_g \left\{plim_{T\to\infty} \hat{\beta}_g\right\} \leq \max_i \{\beta_i\}$$

$$\min_g \left\{plim_{T\to\infty} \hat{\beta}_g\right\} \geq \min_i \{\beta_i\}$$

Intuitively, when group assignments are fixed over time, the estimated consumption exposure of a group g is a convex combination of the consumption exposures of each household in that group. Therefore, each group's consumption exposure is weakly smaller than the largest household exposure, and weakly larger than the smallest household exposure. This immediately says that the asymmetry implied by the estimates must be bounded by the true asymmetry at the household level.

When $\mathcal{G}$ assigns households to different groups over time, it is difficult to say whether the implied asymmetry from pooled OLS over- or underestimates the true asymmetry. As an extreme example, suppose that $\beta_i = 1$ for all i (so that the true asymmetry is nil) and consider the following assignment process for a fixed group g. When $\Delta \hat{\log} Y_t > 0$, assign households with the highest consumption growths to group g. When $\Delta \hat{\log} Y_t < 0$, assign households with the lowest consumption growths to group g. Such a process will result in an estimate of $\hat{\beta}_g$ much larger than 1, due to the selection bias created by the assignment mechanism's dependence on idiosyncratic shocks, and will therefore overestimate the true asymmetry. Furthermore, the opposite assignment process will clearly result in an underestimate of the true asymmetry.

In light of this discussion, I choose as a benchmark, an assignment mechanism that is fixed over time, so that the estimated asymmetry is known to be a lower bound on the true asymmetry (in the limit $T \to \infty$). In practice, this amount to defining groups based on household characteristics that are fixed in the sample of households that I observe.

1.4.4 Data

Monetary Policy Shocks In order to extract the variation in $\Delta \log Y_t$ driven by monetary policy shocks, I follow Coibon et al. (2017), who use the methods introduced by Romer and Romer (2004) to identify innovations to monetary policy that are

orthogonal to economic conditions. Formally, the authors run the regression

$$\Delta FFR_t = x_t'\Gamma + \epsilon_t^v$$

where ΔFFR_t is the change in the federal funds rate from period $t-1$ to t, and x_t is a vector of controls that contains forecasts of GDP growth, inflation, and the unemployment rate taken from the Greenbooks at each Federal Open Market Committee meeting. The residuals from this regression, $\{\hat{\epsilon}_t^v\}$, are then taken as the series of monetary policy shocks, with the interpretation that $\hat{\epsilon}_t^v > 0$ is a contractionary shock, and $\hat{\epsilon}_t^v < 0$ is an expansionary shock.

Using this method, Coibon et al. (2017) generate a series of monetary policy shocks at a monthly frequency from 1969 to 2008, which I plot in figure 1.1. The shocks are evenly spread over positive and negative values, and are very volatile during the Volcker disinflation period in the early 1980s.

Output Data As my measure of output, I use quarterly growth rates of per-capita personal consumption expenditures of non-durable goods and services (at a monthly frequency), taken from the NIPA, deflated using the personal consumption expenditure price deflator.

My choice to use growth in per-capita personal consumption expenditures as the right-hand side variable reflects two considerations. First, the theoretical models I have studied in this paper have all abstracted from capital investment and government spending, so that aggregate consumption is the theoretically consistent measure of total output. Second, unlike measures of GDP, personal consumption expenditures are available at a monthly frequency, which enables me to exploit all of the variation in the micro-data and to maintain a reasonable sample size.

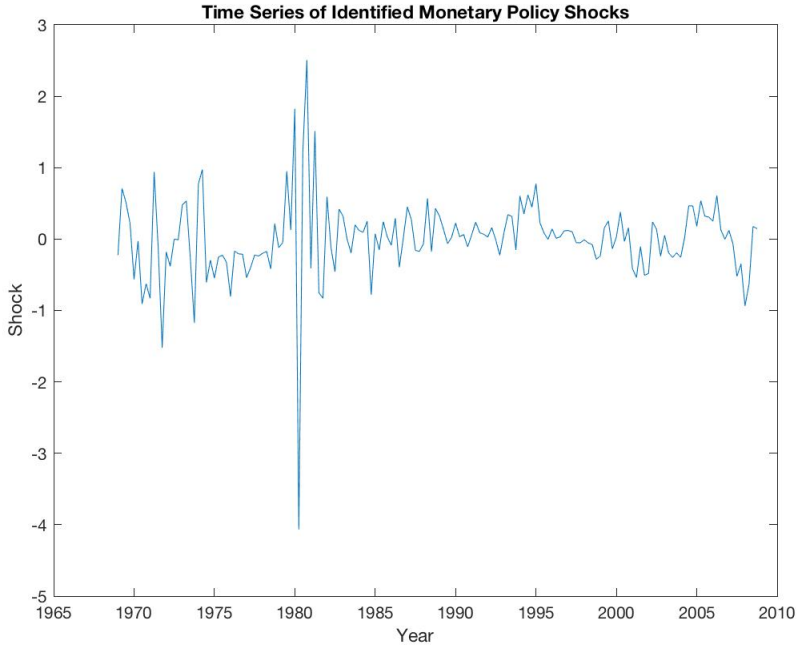


Figure 1.1: Identified Monetary Policy Shocks from Coibon et al. (2017). The authors run the regression $\Delta FFR = x_t'\gamma + \epsilon_t$ where ΔFFR_t is the change in the federal funds rate from period $t-1$ to t, and x_t is a vector of controls that contains forecasts of GDP growth, inflation, and the unemployment rate taken from the Greenbooks at each Federal Open Market Committee meeting. The residuals from this regression, $\{\hat{\epsilon}_t\}$, are then taken as the series of monetary policy shocks, with the t interpretation that $\hat{\epsilon}_t > 0$ is a contractionary shock, and $\hat{\epsilon}_t < 0$ is an expansionary shock.

Consumption Data I use the CEX surveys from 1996 to 2009 to measure consumption of non-durables and services at the household level. In order to ensure the consistency of consumption measurements between the CEX and NIPA, I sum across the relevant categories of expenditure in the CEX, and define non-durable and services consumption as total expenditures on food, services, heating fuel, public and private transport, personal care, and clothing and footwear.[15] I deflate nominal expenditures using the personal consumption expenditure price deflator.

I restrict the sample to urban households, not in student status, where the household head is of working age (25-64), and only consider households who respond to all four interview waves. In order to remove consumption variation caused by factors outside of my model, I first regress log real consumption on a polynomial in age of the household head, family size, and number of children under the age of eighteen, and use the residuals from this regression as my measures of household consumption.

Each household reports their consumption four times at three month intervals. From these reports, I compute three quarterly growth rates of log consumption for each household. Since different households are interviewed each month, I have quarterly growth rates of household consumption, available at a monthly frequency.

Proposition 3 suggests that grouping households together based on a fixed attribute is a useful benchmark to estimate a lower bound on the asymmetry coefficients. In the CEX data, the best candidate for this is the level of education of the household head.[16] Over the year long cycle during which the household reports consumption, the education level of the household head is fixed and is certainly exogenous to changes in output over the same period. Therefore, I sort households into five groups based on the education level of the household head: less than high school, high school, some

[15] My results are robust to variations in this definition.

[16] The very short panel nature of the CEX data implies that other potential fixed attributes such as permanent income are difficult to plausibly compute. Education is of course likely to be correlated with this and other fixed attributes.

college, full college, and beyond college (advanced degree). Note that each group consists of thousands of households in each period so that the composite measurement error term is plausibly zero.

Empirical Specification For the first stage regression

$$\Delta \log Y_t = Z_t' \gamma + e_t$$

I project $\Delta \log Y_t$ onto a vector of the ninety six most recent identified monetary policy shocks $Z_t = \left(\hat{\epsilon}^v_{t-1}, ..., \hat{\epsilon}^v_{t-96}\right)$. This allows the effects of monetary policy shocks to persist for up to eight years. Since the empirically relevant range of monetary policy shock persistence is two to three years (Gertler and Karadi (2015) and Christiano et al. (2005)), the choice of $L = 96$ is a reasonable approximation of the history of shocks that matter for variation in output growth. In appendix 1.B, I show that my results are robust to variations in the lag length L.

All regressions are weighted using the CEX survey weights provided in the data sets.

1.4.5 Results

Table 1.1 shows the estimated coefficient $\hat{\beta}_g$ for each education group, together with its standard error, which I cluster at the household level, and total sample size. The estimated coefficients are strong increasing with respect to education. A 1% increase in the growth of aggregate consumption caused by monetary policy shocks is associated with a 3.58% increase in the consumption growth of households with an advanced degree, but a 1.23% increase in the consumption growth of households with only a high-school diploma.

	Less than High School	High School	Some College	Full College	Advanced Degree
$\hat{\beta}_g$	1.28	1.23	1.71	2.86	3.58
s.e.	1.23	0.79	0.71	0.77	1.12
n	9,621	21,396	27,025	19,206	11,022

Table 1.1: Estimated $\left\{\hat{\beta}_g\right\}$ exposure coefficients across household groups with different education levels using monthly data over the period 1996-2008. Standard errors are clustered at the household level.

These results imply an estimate for the sensitivity ratio of $\frac{\bar{\beta}^{c,v}}{\underline{\beta}^{c,v}} \approx 2.9$. Therefore, the most sensitive households are approximately three times as sensitive to changes in aggregate consumption than the least sensitive households.

This finding is in line with previous studies of heterogeneous consumption sensitivities. For example, Parker and Vissing-Jorgensen (2009) group households in period t by their consumption level in period $t-1$, and find a sensitivity ratio of 5. While this estimate is larger than the lower bound of 2.9, the grouping strategy fails the conditions in proposition 3 so that it likely yields a biased estimate the true sensitivity ratio.

The slight "U-shaped" pattern of sensitivities is also consistent with the evidence on heterogeneous income sensitivities. For example, Guvenen et al. (2017) run a similar regression using worker level income data and unconditional variation in GDP growth across percentiles of the permanent income distribution, and find a "U-shaped" pattern of sensitivities such that the highest and lowest permanent income workers are the most sensitive to unconditional changes in GDP growth. This finding supports the theory that borrowing constraints cause household consumption to inherit the sensitivity of household income to changes in output.

1.4.6 Quantitative Assessment

Given estimates for $\bar{\beta}^{c,v}$ and $\underline{\beta}^{c,v}$, the other key parameters in $\mathcal{R}$ are the slope of the NKPC, κ_y, and the coefficient in output in the Taylor rule, ϕ_y. I set $\phi_y = 0$, which is in line with the existing literature that uses calibrated Taylor rules. In order to set κ_y, I appeal to the empirical evidence from the literatures on inflation forecasting and estimation of the NKPC.

Both of these literatures suggest that κ_y is very small. The forecasting literature suggests that $\kappa_y = 0$ is very plausible (Atkeson and Ohanian, 2001), while the estimation literature tends to find κ_y around 0.05, but with a decent dose of uncertainty (Schorfheide, 2008). Therefore, as a convenient benchmark, I set $\kappa_y = 0$.

When $\kappa_y = 0$, the asymmetry ratio is $\mathcal{R} = 2.9$. Therefore, the output response to a contractionary monetary policy shock is three times as large as the output response to an expansionary monetary policy shock of equal magnitude. I compare this asymmetry to the macro evidence for asymmetry in the next section.

For completeness, figure 1.2 plots $\mathcal{R}$ as a function of κ_y using a standard calibration of the other parameters.[17] The ratio declines as κ_y increases, but remains above 2.3 throughout the range, which covers the most plausible values of κ_y away from zero.

Intuitively, when output increases after an expansionary shock, $\kappa_y > 0$ implies that inflation also increases. Higher inflation causes high nominal rates via the Taylor rule, which offsets some of the initial expansionary shock. The same logic implies that $\kappa_y > 0$ causes deflation to offset the contractionary shock. Since the initial output response is larger for a contractionary shock, the offsetting force is larger too, which shrinks the overall asymmetry.

[17] I set $\rho_v = 0.6$ to reflect the quarterly persistence of monetary policy shocks estimated in the data (Christiano et al., 2005; Gertler and Karadi, 2015). I set $\phi_\pi = 1.25$ and $\phi_y = 0$, which is a commonly used specification for the Taylor rule, and set $\frac{1}{\sigma} = 0.67$ in line with estimates for the EIS (Vissing-Jorgensen, 2002). Finally, I set $\delta = 0.995$, which is consistent with an annual real interest rate of 2%.

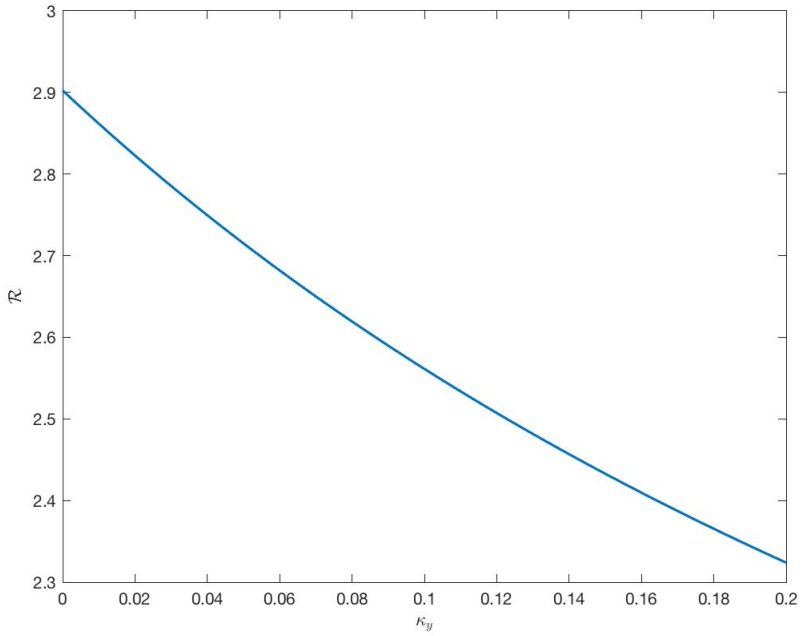


Figure 1.2: Asymmetry ratio $\mathcal{R}$ as a function of κ_y when $\rho_v = 0.6$, $\phi_\pi = 1.25$, $\phi_y = 0$, $\sigma = 1.5$, and $\delta = 0.995$.

1.5 Empirical Evidence of Monetary Policy Asymmetry

The micro evidence on heterogeneous consumption sensitivities implies that contractionary monetary policy shocks are three times more powerful than expansionary monetary policy shocks. In this section, I show that this result is in line with the macro-econometric evidence for asymmetric monetary policy transmission. Specifically, I use local projection methods (Jorda, 2005) to demonstrate that contractionary monetary policy shocks are approximately four times more powerful than expansionary shocks.[18]

[18]The literature on asymmetric monetary policy goes back to at least Cover (1992) and DeLong and Summers (1988), who both find contractionary shocks are more powerful than expansionary shocks. More recently, Angrist et al. (2013), and Barnichon and Matthes (2016), introduce novel methodologies to measure asymmetric effects, and also find that contractionary monetary policy shocks are more powerful than expansionary shocks.

1.5.1 Empirical Specification

I follow Jorda (2005), and estimate the impulse response of output to monetary policy shocks using local projection methods. Formally, I estimate the specification

$$y_{t+h} = \alpha^h + \beta^{h,+} \max\{\hat{\epsilon}_t^v, 0\} + \beta^{h,-} \min\{\hat{\epsilon}_t^v, 0\} + \sum_{l=0}^{L} \gamma_{y,l}^h y_{t-l} + \sum_{l=1}^{L} \gamma_{FFR,l}^h FFR_{t-l} + u_{t+h}^h$$

for horizons $h = 1, ..., H$. Here, $\{y_t\}$ is linearly de-trended output (in logs), $\{\hat{\epsilon}_t^v\}$ is the series of identified monetary policy shocks, and $\{FFR_t\}$ is the federal funds rate. The estimated coefficients $\left\{\hat{\beta}^{h,+}\right\}_1^H$ and $\left\{\hat{\beta}^{h,-}\right\}_1^H$ are the impulse responses of y to positive and negative shocks of unit size respectively.

I use quarterly frequency data over the period 1969 - 2008. In order to be consistent with the micro-data evidence, I use per-capita aggregate consumption of non-durables and services as my measure of output. I set $L = 1$, and note that the inclusion of contemporaneous aggregate consumption as a regressor is consistent with the convention that monetary policy shocks only affect measures of aggregate demand with a 1 period delay (Christiano et al., 1999). Finally, I estimate the system of equations over $h = 1, ..., H$ jointly, and compute Driscoll-Kraay (1998) standard errors that are robust to arbitrary serial and cross-sectional correlation across time and horizons.

1.5.2 Results

Figure 1.3 plots the estimated impulse responses of output to contractionary (positive) and expansionary (negative) monetary policy shocks of 1% size over fifteen quarters. The dashed lines are 90% confidence intervals. For ease of comparison, I have multiplied the expansionary response by -1. Both impulse responses exhibit the

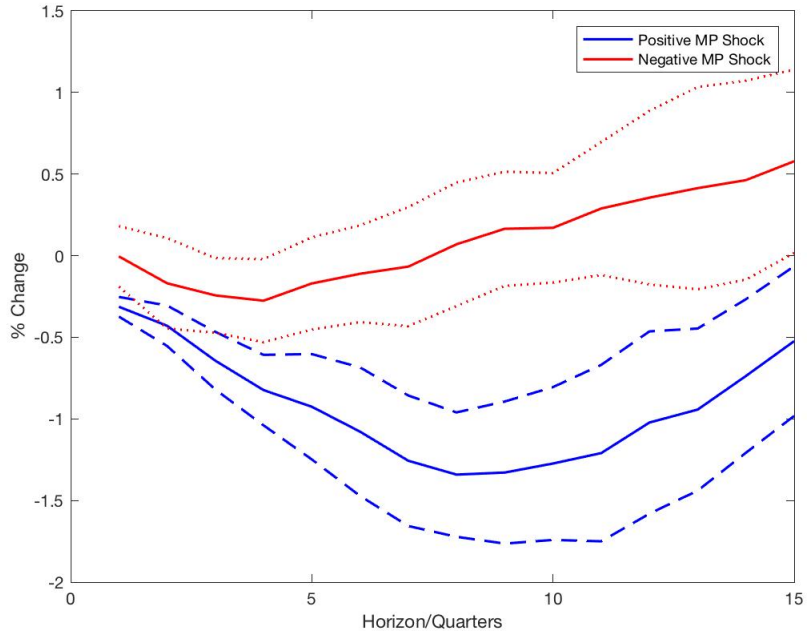


Figure 1.3: Impulse responses of aggregate consumption (from NIPA) estimated using local projection methods. The dashed lines are 90% confidence intervals computed using Driscoll-Kraay standard errors.

"U-shape" that is a common feature of output responses to monetary policy shocks (Christiano et al., 1999).[19]

The contractionary shock generates a maximum response that is approximately four times as large as the maximum response to an expansionary shock. The asymmetry is statistically significant after about one year, by which time the effect of the expansionary shock has started to died out, but the contractionary shock is still causing further declines in output.

As a simple metric of comparison, I compare the ratio of the maximum responses in the data to the ratio of responses in the model, $\mathcal{R}$. According to this metric, the asymmetry estimated in the macro data is reasonably consistent with the asymmetry implied by the micro-data. The fact that the sensitivity ratio implied by the micro

[19]Since my simple model does not contain ingredients such as consumption habits or investment frictions that are typically found in medium-scale DSGE models, it cannot generate the "hump-shaped" impulse responses found in the data.

data is a lower bound implies that a quantitative version of model can explain at least 60% of the asymmetry found in the macro data, and could plausibly explain much more if we can estimate the true exposure ratio at a more granular level of household heterogeneity than education.

1.5.3 Robustness Checks

Here, I show that the asymmetric responses of output to monetary policy shocks are robust to regression specifications with different lag and control variable structures, sample restrictions that exclude the Volcker disinflation period, and when I change the dependent variable to GDP. All figures are in the appendix.

My baseline choice of $L = 1$ is optimal according to the Bayesian Information Criterion (BIC) given by

$$T \log (RSS/T) + k \log T$$

where RSS is the residual sum of squares from the regressions and T is the sample length. I also consider the Akaike Information Criterion (AIC), which is given by

$$T \log (RSS/T) + 2k$$

and also suggests an optimal choice of $L = 1$. Furthermore, figure 1.B plots the impulse responses for $L \in \{2, 3, 4, 5\}$, and shows that the asymmetry is similar to the baseline specification in all cases.

The baseline regression includes aggregate demand and the federal funds rate as control variables. However, most New Keynesian models imply that inflation is also determined as part of the equilibrium system, and so affects the path of aggregate demand. To this end, figure 1.5 plots the impulse responses with inflation (measured by the Personal Consumption Expenditure deflator) as an additional control variable

that follows the same lag structure as aggregate demand. The asymmetry is essentially unchanged.

It is well known that the Volcker disinflation period in the early 1980s resulted in volatile monetary policy, as exhibited by the large shocks in figure 1.1. While these shocks provide useful variation in the explanatory variable, it is useful to check that they are not the driving force behind the result. Therefore, in figure 1.6 I plot the impulse responses from the baseline regression having restricted the sample to 1985Q1 onwards, thus dropping the entire Volcker episode. While the smaller sample results in much wider confidence intervals, the asymmetry is still clear to see, with contractionary shocks having twice the effect of expansionary shocks. Note that in this case, the micro evidence can explain all of the asymmetry.

Finally, I run the baseline regression with real GDP as the dependent variable instead of aggregate consumption. Figure 1.7 plots the impulse responses, which exhibit similar levels of asymmetry, although they are slightly more noisily estimated.

1.6 Conclusion

When output falls in response to a contractionary monetary policy shock, the decrease in consumption of saver households is necessarily the smallest among all households. Therefore, the fall in output is greater than the response of saver households alone. In contrast, when output increases in response to an expansionary monetary policy shock, the increase in consumption of saver households is necessarily the largest among all households. Therefore, the increase in output is smaller than the response of saver households alone. Hence, output responds more to contractionary monetary shocks than to expansionary shocks of equal magnitude.

The micro-data suggests that the largest sensitivity of household consumption to changes in output is approximately three times the size of the smallest sensitivity.

When inflation is unresponsive to changes in output, output should respond three times as much to contractionary shocks than to expansionary shocks. This quantitative result can therefore explain at least 60% of the asymmetry found in the macro data.

The mechanism in this paper applies to any aggregate shock. It would therefore be interesting to investigate the asymmetric transmission of other aggregate shocks, and to see how well the model does at explaining the asymmetry.

References

Abbritti, M. and Fahr, S.: 2013, Downward wage rigidity and business cycle asymmetries, *Journal of Monetary Economics* **60**(7), 871–886.

Acemoglu, D. and Scott, A.: 1997, Asymmetric business cycles: Theory and time-series evidence, *Journal of Monetary Economics* **40**(3), 501–533.

Acharya, S. and Dogra, K.: 2018, Understanding HANK: Insights from a PRANK.

Aguiar, M. and Bils, M.: 2015, Has consumption inequality mirrored income inequality?, *American Economic Review* **105**(9), 2725–56.

Ahn, S., Kaplan, G., Moll, B., Winberry, T. and Wolf, C.: 2018, When inequality matters for macro and macro matters for inequality, *NBER Macroeconomics Annual* **32**(1), 1–75.

Ai, H. and Bhandari, A.: 2018, Asset Pricing with Endogenously Uninsurable Tail Risk, *Technical report*, National Bureau of Economic Research.

Aiyagari, S. R.: 1994, Uninsured idiosyncratic risk and aggregate saving, *The Quarterly Journal of Economics* **109**(3), 659–684.

Alvarez, F. and Jermann, U. J.: 2000, Efficiency, equilibrium, and asset pricing with risk of default, *Econometrica* **68**(4), 775–797.
URL: *https://doi.org/10.1111/1468-0262.00137*

Angrist, J. D., Jordà, Ò. and Kuersteiner, G.: 2013, Semiparametric estimates of monetary policy effects: string theory revisited, *Technical report*, National Bureau of Economic Research.

Atkeson, A. E. O. et al.: 2001, Are Phillips Curves Useful for Forecasting Inflation?, *Federal Reserve bank of Minneapolis quarterly review* **25**(1), 2–2.

Attanasio, O. P. and Paiella, M.: 2011, Intertemporal consumption choices, transaction costs and limited participation in financial markets: reconciling data and theory, *Journal of Applied Econometrics* **26**(2), 322–343.

Attanasio, O. P. and Weber, G.: 1995, Is consumption growth consistent with intertemporal optimization? Evidence from the consumer expenditure survey, *Journal of political Economy* **103**(6), 1121–1157.

Auclert, A.: 2017, Monetary policy and the redistribution channel, *Technical report*, National Bureau of Economic Research.

Auclert, A. and Rognlie, M.: 2017, A note on multipliers in NK models with GHH preferences, *Report, Stanford.[269]* .

Auclert, A. and Rognlie, M.: 2018, Inequality and aggregate demand, *Technical report*, National Bureau of Economic Research.

Baqaee, D.: 2016, Asymmetric inflation expectations, downward rigidity of wages, and asymmetric business cycles, *Downward Rigidity of Wages, and Asymmetric Business Cycles (June 20, 2016)* .

Barnichon, R. and Brownlees, C.: 2018, Impulse response estimation by smooth local projections, *Review of Economics and Statistics* (0).

Barnichon, R. and Matthes, C.: 2017a, Functional approximations of impulse responses: New insights into the asymmetric effects of monetary policy, *Technical report*, Working paper.

Barnichon, R. and Matthes, C.: 2017b, Understanding the Size of the Government Spending Multiplier: it's in the Sign.

Barnichon, R., Matthes, C., Sablik, T. et al.: 2017, Are the Effects of Monetary Policy Asymmetric?, *Richmond Fed Economic Brief* (March), 1–4.

Barro, R. J. and Redlick, C. J.: 2011, Macroeconomic effects from government purchases and taxes, *The Quarterly Journal of Economics* **126**(1), 51–102.

Basu, S. and House, C. L.: 2016, Allocative and remitted wages: new facts and challenges for keynesian models, *Handbook of macroeconomics*, Vol. 2, Elsevier, pp. 297–354.

Bayer, C., Lütticke, R., Pham-Dao, L. and Tjaden, V.: 2019, Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk, *Econometrica* **87**(1), 255–290.

Berger, D., Dovis, A. and Bocola, L.: 2018, Accounting for Heterogeneity, *Technical report*.

Bhandari, A., Evans, D., Golosov, M. and Sargent, T. J.: 2018, Inequality, business cycles, and monetary-fiscal policy, *Technical report*, National Bureau of Economic Research.

Bilbiie, F. O.: 2018, Monetary policy and heterogeneity: An analytical framework.

Blundell, R., Pistaferri, L. and Preston, I.: 2008, Consumption inequality and partial insurance, *American Economic Review* **98**(5), 1887–1921.

Boppart, T., Krusell, P. and Mitman, K.: 2018, Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative, *J. Econ. Dyn. Control* **89**, 68–92.

Brinca, P., Chari, V. V., Kehoe, P. J. and McGrattan, E.: 2016, Accounting for business cycles, *Handbook of Macroeconomics*, Vol. 2, Elsevier, pp. 1013–1063.

Broer, T., Hansen, N.-J. H., Krusell, P. and Öberg, E.: 2016, The new keynesian transmission mechanism: A heterogenous-agent perspective, *Technical report*, National Bureau of Economic Research.

Bunn, P., Le Roux, J., Reinold, K. and Surico, P.: 2018, The consumption response to positive and negative income shocks, *Journal of Monetary Economics* **96**, 1–15.

Carrillo, J., Fève, P., Matheron, J. et al.: 2012, *Monetary policy inertia or persistent shocks: a DSGE analysis*, BiblioGov.

Challe, E. and Giannitsarou, C.: 2014, Stock prices and monetary policy shocks: A general equilibrium approach, *Journal of Economic Dynamics and Control* **40**, 46–66.

Chari, V. V., Kehoe, P. J. and McGrattan, E. R.: 2007, Business cycle accounting, *Econometrica* **75**(3), 781–836.

Chari, V. V., Kehoe, P. J. and McGrattan, E. R.: 2009, New keynesian models: not yet useful for policy analysis, *American Economic Journal: Macroeconomics* **1**(1), 242–66.

Chetty, R.: 2009, Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods, *Annu. Rev. Econ.* **1**(1), 451–488.

Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L. and Van Rooij, M.: 2017, Asymmetric consumption effects of transitory income shocks.

Christiano, L. J., Eichenbaum, M. and Evans, C. L.: 1999, Monetary policy shocks: What have we learned and to what end?, *Handbook of macroeconomics* **1**, 65–148.

Cochrane, J. H.: 2017, The new-keynesian liquidity trap, *Journal of Monetary Economics* **92**, 47–63.

Coibion, O., Gorodnichenko, Y., Kueng, L. and Silvia, J.: 2017, Innocent bystanders? monetary policy and inequality, *Journal of Monetary Economics* **88**, 70–89.

Cover, J. P.: 1992, Asymmetric effects of positive and negative money-supply shocks, *The Quarterly Journal of Economics* **107**(4), 1261–1282.

Davis, S. J., Kubler, F. and Willen, P.: 2006, Borrowing costs and the demand for equity over the life cycle, *The Review of Economics and Statistics* **88**(2), 348–362.

De Giorgi, G. and Gambetti, L.: 2017, Business cycle fluctuations and the distribution of consumption, *Review of Economic Dynamics* **23**, 19–41.

De Long, J. B., Summers, L. H., Mankiw, N. G. and Romer, C. D.: 1988, How does macroeconomic policy affect output?, *Brookings Papers on Economic Activity* (2), 433.

Driscoll, J. C. and Kraay, A. C.: 1998, Consistent covariance matrix estimation with spatially dependent panel data, *Review of economics and statistics* **80**(4), 549–560.

Farhi, E. and Werning, I.: 2016, A theory of macroprudential policies in the presence of nominal rigidities, *Econometrica* **84**(5), 1645–1704.

Farhi, E. and Werning, I.: 2017, Monetary policy, bounded rationality, and incomplete markets, *Technical report*, National Bureau of Economic Research.

Fernández-Villaverde, J. and Krueger, D.: 2007, Consumption over the life cycle: Facts from consumer expenditure survey data, *The Review of Economics and Statistics* **89**(3), 552–565.

Fuster, A., Kaplan, G. and Zafar, B.: 2018, What would you do with $ 500? spending responses to gains, losses, news and loans, *Technical report*, National Bureau of Economic Research.

Galí, J.: 2015, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press.

Garín, J., Lester, R. and Sims, E.: 2016, Are supply shocks contractionary at the zlb? evidence from utilization-adjusted tfp data, *Review of Economics and Statistics* (00).

Gertler, M. and Karadi, P.: 2015, Monetary policy surprises, credit costs, and economic activity, *American Economic Journal: Macroeconomics* **7**(1), 44–76.

Gervais, M. and Klein, P.: 2010, Measuring consumption smoothing in cex data, *Journal of Monetary Economics* **57**(8), 988–999.

Gornemann, N., Kuester, K. and Nakajima, M.: 2016, Doves for the rich, hawks for the poor? distributional consequences of monetary policy.

Greenwood, J., Hercowitz, Z. and Huffman, G. W.: 1988, Investment, capacity utilization, and the real business cycle, *The American Economic Review* pp. 402–417.

Guvenen, F.: 2006, Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective, *Journal of Monetary Economics* **53**(7), 1451–1472.

Guvenen, F.: 2007, Do stockholders share risk more effectively than nonstockholders?, *The Review of Economics and Statistics* **89**(2), 275–288.

Guvenen, F.: 2009, A parsimonious macroeconomic model for asset pricing, *Econometrica* **77**(6), 1711–1750.

Guvenen, F. and Kaplan, G.: 2017, Top income inequality in the 21st century: Some cautionary notes, *Technical report*, National Bureau of Economic Research.

Guvenen, F., Kaplan, G. and Song, J.: 2014, How risky are recessions for top earners?, *American Economic Review* **104**(5), 148–53.

Guvenen, F., Kaplan, G., Song, J. and Weidner, J.: 2017, Lifetime incomes in the united states over six decades, *Technical report*, National Bureau of Economic Research.

Guvenen, F., Karahan, F., Ozkan, S. and Song, J.: 2016, What do data on millions of us workers reveal about life-cycle earnings dynamics?, *Federal Reserve Bank of New York Staff Report* (710).

Guvenen, F., Ozkan, S. and Song, J.: 2014, The nature of countercyclical income risk, *Journal of Political Economy* **122**(3), 621–660.

Hamilton, J. D.: 2018, Why you should never use the hodrick-prescott filter, *Review of Economics and Statistics* **100**(5), 831–843.

Heathcote, J., Perri, F. and Violante, G. L.: 2010, Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006, *Review of Economic dynamics* **13**(1), 15–51.

Heathcote, J., Storesletten, K. and Violante, G. L.: 2014, Consumption and labor supply with partial insurance: An analytical framework, *American Economic Review* **104**(7), 2075–2126.

Jappelli, T. and Pistaferri, L.: 2010, The consumption response to income changes.

Jordà, Ò.: 2005, Estimation and inference of impulse responses by local projections, *American economic review* **95**(1), 161–182.

Kaplan, G., Moll, B. and Violante, G. L.: 2018, Monetary policy according to hank, *American Economic Review* **108**(3), 697–743.

Kaplan, G. and Violante, G. L.: 2010, How much consumption insurance beyond self-insurance?, *American Economic Journal: Macroeconomics* **2**(4), 53–87.
URL: *http://www.aeaweb.org/articles?id=10.1257/mac.2.4.53*

Kaplan, G. and Violante, G. L.: 2014, A model of the consumption response to fiscal stimulus payments, *Econometrica* **82**(4), 1199–1239.

Karlan, D. S.: 2005, Using experimental economics to measure social capital and predict financial decisions, *American Economic Review* **95**(5), 1688–1699.

Kehoe, P. J., Midrigan, V. and Pastorino, E.: 2018, Evolution of modern business cycle models: Accounting for the great recession, *Journal of Economic Perspectives* **32**(3), 141–66.

Kichian, M., Rumler, F. and Corrigan, P.: 2010, Semi-structural models for inflation forecasting, *Technical report*, Bank of Canada Working Paper.

Korinek, A. and Simsek, A.: 2016, Liquidity trap and excessive leverage, *American Economic Review* **106**(3), 699–738.

Krusell, P. and Smith, Jr, A. A.: 1998, Income and wealth heterogeneity in the macroeconomy, *Journal of political Economy* **106**(5), 867–896.

Lucas Jr, R. E. and Stokey, N. L.: 1983, Optimal fiscal and monetary policy in an economy without capital, *Journal of monetary Economics* **12**(1), 55–93.

Luetticke, R.: 2018, Transmission of monetary policy with heterogeneity in household portfolios.

Malloy, C. J., Moskowitz, T. J. and Vissing-Jørgensen, A.: 2009, Long-run stockholder consumption risk and asset returns, *The Journal of Finance* **64**(6), 2427–2479.

McKay, A., Nakamura, E. and Steinsson, J.: 2016, The power of forward guidance revisited, *American Economic Review* **106**(10), 3133–58.

McKay, A. and Reis, R.: 2016, The role of automatic stabilizers in the us business cycle, *Econometrica* **84**(1), 141–194.

Parker, J. A.: 2001, The consumption risk of the stock market, *Brookings Papers on Economic Activity* **2001**(2), 279–348.

Parker, J. A.: 2017, Why don't households smooth consumption? evidence from a $ 25 million experiment, *American Economic Journal: Macroeconomics* **9**(4), 153–83.

Parker, J. A. and Vissing-Jorgensen, A.: 2009, Who bears aggregate fluctuations and how?, *American Economic Review* **99**(2), 399–405.

Parker, J. A. and Vissing-Jorgensen, A.: 2010, The increase in income cyclicality of high-income households and its relation to the rise in top income shares, *Technical report*, National Bureau of Economic Research.

Patterson, C.: 2019, The Matching Multiplier and the Amplification of Recessions.

Primiceri, G. E., Schaumburg, E. and Tambalotti, A.: 2006, Intertemporal disturbances, *Technical report*, National Bureau of Economic Research.

Ramey, V. A.: 2011, Can government purchases stimulate the economy?, *Journal of Economic Literature* **49**(3), 673–85.

Ravn, M. O., Sola, M. et al.: 2004, Asymmetric effects of monetary policy in the united states, *Review-Federal Reserve Bank of Saint Louis* **86**, 41–58.

Ravn, M. O. and Sterk, V.: 2016, Macroeconomic fluctuations with hank & sam: an analytical approach.

Romer, C. D. and Romer, D. H.: 2004, A new measure of monetary shocks: Derivation and implications, *American Economic Review* **94**(4), 1055–1084.

Schmitt-Grohé, S. and Uribe, M.: 2004, Optimal fiscal and monetary policy under sticky prices, *Journal of economic Theory* **114**(2), 198–230.

Schorfheide, F.: 2008, Dsge model-based estimation of the new keynesian phillips curve, *FRB Richmond Economic Quarterly* **94**(4), 397–433.

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

Smets, F. and Wouters, R.: 2007, Shocks and frictions in us business cycles: A bayesian dsge approach, *American economic review* **97**(3), 586–606.

Smith, M., Yagan, D., Zidar, O. M. and Zwick, E.: 2019, Capitalists in the twenty-first century, *Technical report*, National Bureau of Economic Research.

Sorger, G.: 2000, Income and wealth distribution in a simple model of growth, *Economic Theory* **16**(1), 23–42.

Storesletten, K., Telmer, C. I. and Yaron, A.: 2004, Consumption and risk sharing over the life cycle, *Journal of monetary Economics* **51**(3), 609–633.

Tenreyro, S. and Thwaites, G.: 2016, Pushing on a string: Us monetary policy is less powerful in recessions, *American Economic Journal: Macroeconomics* **8**(4), 43–74.

Vissing-Jørgensen, A.: 2002a, Limited asset market participation and the elasticity of intertemporal substitution, *Journal of political Economy* **110**(4), 825–853.

Vissing-Jorgensen, A.: 2002b, Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures, *Technical report*, National Bureau of Economic Research.

Werning, I.: 2015, Incomplete markets and aggregate demand, *Technical report*, National Bureau of Economic Research.

Appendix

1.A Additional Figures

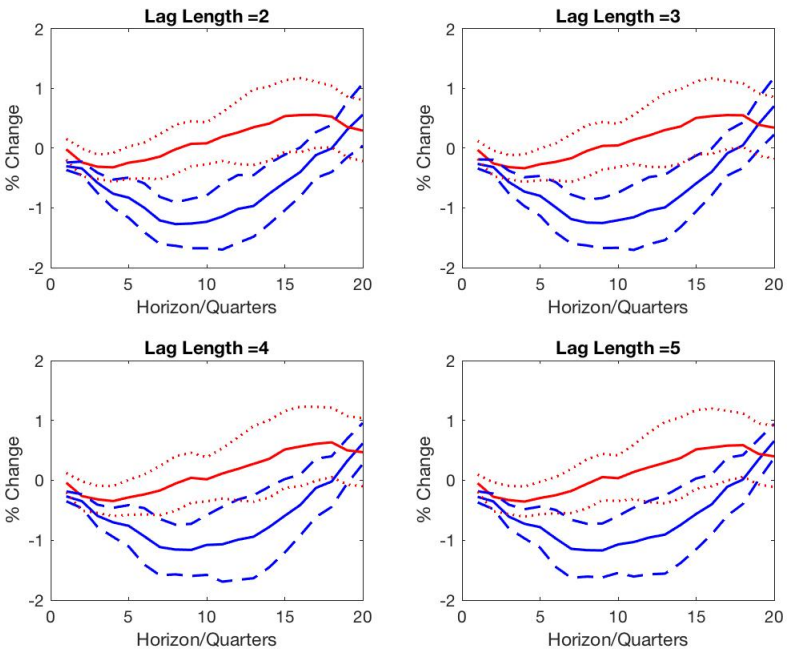


Figure 1.4: Impulse responses of real GDP (from NIPA) estimated using local projection methods and different lag structures.

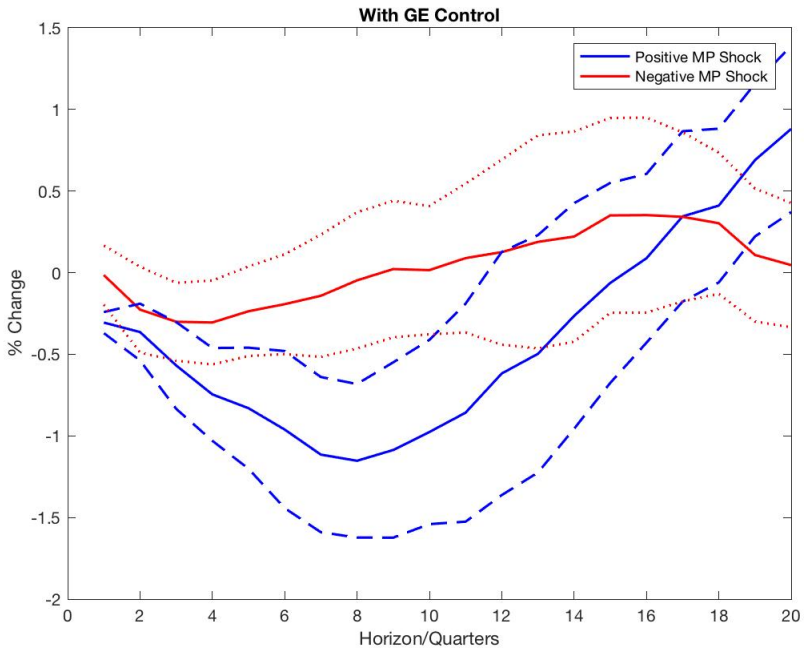


Figure 1.5: Impulse responses of real GDP (from NIPA) estimated using local projection methods with inflation as an additional control variable.

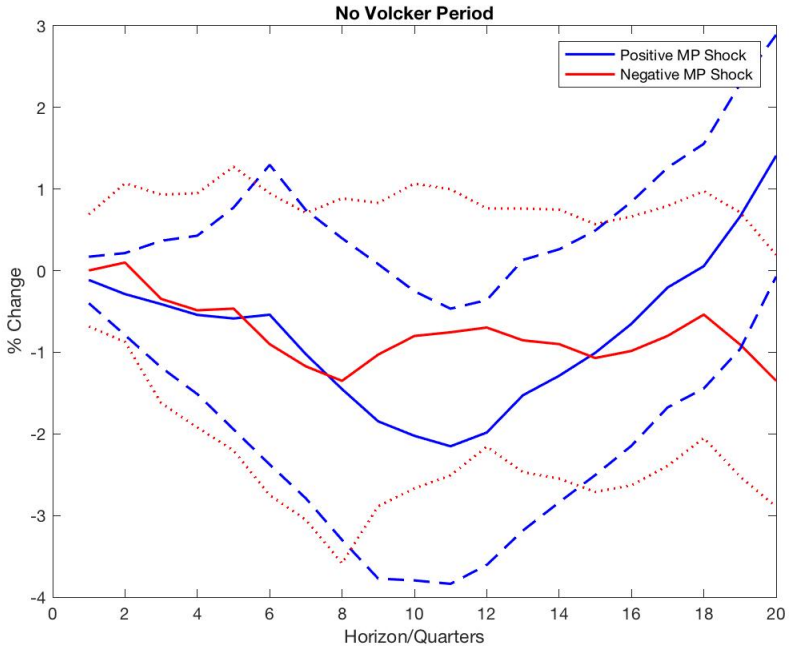


Figure 1.6: Impulse responses of real GDP (from NIPA) estimated using local projection methods, using only the post-Volcker sample.

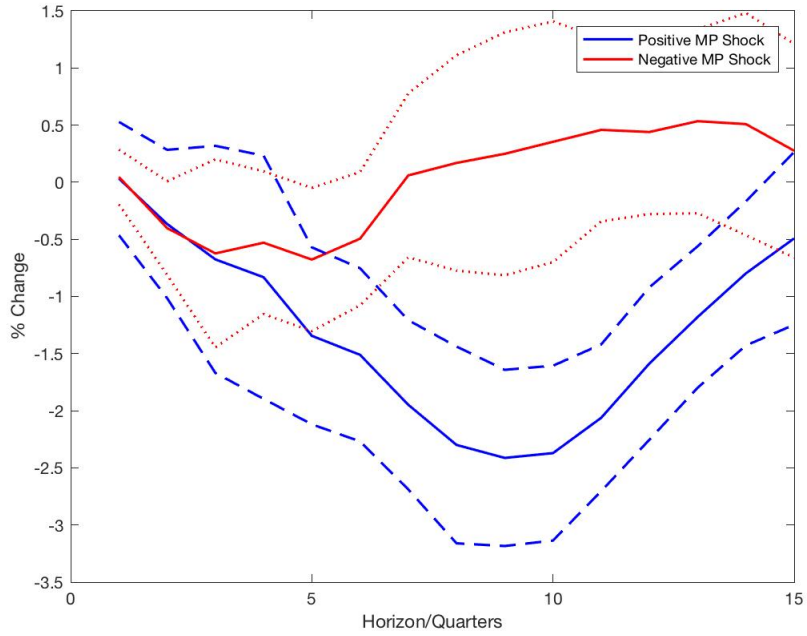


Figure 1.7: Impulse responses of real aggregate consumption of non-durables and services (from NIPA) estimated using local projection methods.

1.B Robustness for CEX Regressions

$\hat{\beta}_g$	$L = 84$	**L = 96**	$L = 108$	$L = 120$	$L = 132$
Less than High School	0.791	1.272	1.230	1.157	1.202
	(1.332)	(1.227)	(1.206)	(1.185)	(1.180)
High School	0.315	1.233	1.544**	1.405*	1.269*
	(0.856)	(0.785)	(0.761)	(0.747)	(0.744)
Some College	1.057	1.711**	1.991***	2.027***	2.070***
	(0.774)	(0.713)	(0.684)	(0.669)	(0.659)
College	2.822***	2.856***	2.803***	2.842***	2.648***
	(0.838)	(0.774)	(0.753)	(0.743)	(0.740)
Advanced Degree	4.176***	3.579***	3.623***	3.523***	3.628***
	(1.240)	(1.123)	(1.078)	(1.066)	(1.066)

Robust standard errors clustered at household level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

1.C Additional Results for Section 1.3

Proposition 4. *Under assumption 1, the first order equilibrium dynamics of output in response to cost-push shocks are given by*

$$\hat{y}_t = \begin{cases} \frac{\phi_\pi \frac{1}{\xi^p}}{\bar{\beta}+\frac{1}{\sigma}\phi_\pi \frac{\Phi-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y} \frac{1}{\sigma}\epsilon_t^\Phi & if \quad \epsilon_t^\Phi > 0 \\ \\ \frac{\phi_\pi \frac{1}{\xi^p}}{\underline{\beta}+\frac{1}{\sigma}\phi_\pi \frac{\Phi-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y} \frac{1}{\sigma}\epsilon_t^\Phi & if \quad \epsilon_t^\Phi < 0 \end{cases}$$

where

$$\bar{\beta} = \max_i \{\beta_i^y\}$$

$$\underline{\beta} = \min_i \{\beta_i^y\}$$

Proposition 4 shows the equilibrium dynamics of output as a function of the contemporaneous cost-push shock, and the heterogeneity in exposures of household income to aggregate income captured by the set of coefficients $\{\beta_i^y\}$. Similar to the monetary shock case, the dynamics of output depend on the sign of the cost-push shock, and the fact that $\bar{\beta} > \underline{\beta}$ implies that output responds more to positive cost-push shocks ($\epsilon_t^\Phi < 0$) than to negative cost-push shocks ($\epsilon_t^\Phi > 0$).

Proposition 5. *Under assumption 1, the first order equilibrium dynamics of output in response to TFP shocks are given by*

$$\hat{y}_t = \begin{cases} c_y^+ \epsilon_t^a & if \quad \epsilon_t^a > 0 \\ \\ c_y^- \epsilon_t^a & if \quad \epsilon_t^a < 0 \end{cases}$$

where

$$c_y^+ = \frac{\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)}{\bar{\beta}^{TFP,+} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y}\frac{1}{\sigma} > 0$$

$$c_y^- = \frac{\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)}{\underline{\beta}^{TFP,-} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y}\frac{1}{\sigma} > c_y^+$$

$$\bar{\beta}^{TFP,+} = \max_i \left\{ \beta_i^y + \frac{\beta_i^a}{c_y^+} \right\}$$

$$\underline{\beta}^{TFP,-} = \min_i \left\{ \beta_i^y + \frac{\beta_i^a}{c_y^-} \right\}$$

Proposition 5 shows the equilibrium dynamics of aggregate output as a function of the contemporaneous TFP shock, and the heterogeneity in exposures of household consumption (and hence income) to aggregate income captured by the set of coefficients $\left\{\beta_i^y + \frac{\beta_i^a}{c_y^{a,+}}\right\}$ and $\left\{\beta_i^y + \frac{\beta_i^a}{c_y^{a,-}}\right\}$. Relative to the previous analyses, an additional complication stems from the fact that TFP shocks directly affect households' incomes since an increase in TFP increases the non-labor share of income and hence raises the exposure of households with large dividend shares. Therefore, the coefficients $c_y^{a,-}$, $c_y^{a,+}$, $\bar{\beta}^{TFP}$, and $\underline{\beta}^{TFP}$ must be determined jointly unlike in the previous analyses where the $\{\beta_i^y\}$ coefficients were direct functions of primitives. Similar to the other shocks, the dynamics of output depend on the sign of the TFP shock, and the fact that $c_y^- > c_y^+$ implies that output responds more to negative TFP shocks than to positive TFP shocks.[20]

[20]Technically, it is possible that $c_y^+, c_y^- < 0$ due to the complementarity between consumption and labor supply created by my assumption of GHH preferences. Specifically, when TFP increases, there is a direct effect on labor supply: holding aggregate demand fixed, labor supply falls since TFP is higher. Given GHH preferences, this fall in labor supply causes net consumption to increase. Holding interest rates fixed, this causes households to lower their demand for consumption in order to maintain a smooth time path of marginal utility. Therefore, a positive TFP shock has a direct contractionary effect on output when households have GHH preferences. Given this effect is purely a result of my special assumption on preferences, I assume that parameters are such that this channel does not dominate the the usual effects that are the focus of my analysis, so that $c_y^+, c_y^- > 0$.

Proposition 6. *Under assumption 1, the first order equilibrium dynamics of inflation in response to monetary policy shocks are given by*

$$\pi_t = \begin{cases} -\frac{\frac{\bar{\Phi}-1}{\xi^p}\varphi}{\underline{\beta}+\frac{1}{\sigma}\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y}\frac{1}{\sigma}\epsilon_t^v & if \quad \epsilon_t^v > 0 \\ \\ -\frac{\frac{\bar{\Phi}-1}{\xi^p}\varphi}{\bar{\beta}+\frac{1}{\sigma}\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y}\frac{1}{\sigma}\epsilon_t^v & if \quad \epsilon_t^v < 0 \end{cases}$$

where

$$\bar{\beta} = \max_i \{\beta_i^y\}$$

$$\underline{\beta} = \min_i \{\beta_i^y\}$$

Proposition 6 shows the equilibrium dynamics of inflation as a function of the contemporaneous monetary shock, and the heterogeneity in exposures of household income to aggregate income captured by the set of coefficients $\{\beta_i^y\}$. Similar to output, the dynamics of inflation depend on the sign of the monetary shock, and the fact that $\bar{\beta}^{MP} > \underline{\beta}^{MP}$ implies that inflation responds more to positive (contractionary) monetary shocks than to negative (expansionary) shocks, thus mimicking the asymmetry of output.

Proposition 7. *Under assumption 1, the first order equilibrium dynamics of inflation in response to cost-push shocks are given by*

$$\pi_t = \begin{cases} -\frac{1}{\xi^p}\left(\frac{\bar{\beta}^{CP}+\frac{1}{\sigma}\phi_y}{\bar{\beta}^{CP}+\frac{1}{\sigma}\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y}\right)\epsilon_t^\Phi & if \quad \epsilon_t^\Phi > 0 \\ \\ -\frac{1}{\xi^p}\left(\frac{\underline{\beta}^{CP}+\frac{1}{\sigma}\phi_y}{\underline{\beta}^{CP}+\frac{1}{\sigma}\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y}\right)\epsilon_t^\Phi & if \quad \epsilon_t^\Phi < 0 \end{cases}$$

where

$$\bar{\beta}^{CP} = \max_i \{\beta_i^y\}$$

$$\underline{\beta}^{CP} = \min_i \{\beta_i^y\}$$

Proposition 7 shows the equilibrium dynamics of inflation as a function of the contemporaneous cost-push shock, and the heterogeneity in exposures of household income to aggregate income captured by the set of coefficients $\{\beta_i^y\}$. Similar to output, the dynamics of inflation depend on the sign of the shock, and the fact that $\bar{\beta}^{CP} > \underline{\beta}^{CP}$ implies that inflation responds more to negative cost-push shocks ($\epsilon_t^{\Phi} > 0$) than to positive cost-push shocks. Hence inflation exhibits the opposite asymmetry to output.

Proposition 8. *Under assumption 1, the first order equilibrium dynamics of inflation in response to TFP shocks are given by*

$$\pi_t = \begin{cases} c_\pi^+ \epsilon_t^a & \text{if} \quad \epsilon_t^a > 0 \\ \\ c_\pi^- \epsilon_t^a & \text{if} \quad \epsilon_t^a < 0 \end{cases}$$

where

$$c_\pi^+ = -\frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)\frac{\bar{\beta}^{TFP,+} + \frac{1}{\sigma}\phi_y}{\bar{\beta}^{TFP,+} + \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)} < 0$$

$$c_\pi^- = -\frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)\frac{\underline{\beta}^{TFP,-} + \frac{1}{\sigma}\phi_y}{\underline{\beta}^{TFP,-} + \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)} > c_y^+$$

$$\bar{\beta}^{TFP,+} = \max_i \left\{\beta_i^y + \frac{\beta_i^a}{c_y^+}\right\}$$

$$\underline{\beta}^{TFP,-} = \min_i \left\{\beta_i^y + \frac{\beta_i^a}{c_y^-}\right\}$$

Proposition 8 shows the equilibrium dynamics of inflation as a function of the contemporaneous TFP shock, and the heterogeneity in exposures of household income to aggregate income captured by the set of coefficients $\left\{\beta_i^y + \frac{\beta_i^a}{c_y^+}\right\}$ and $\left\{\beta_i^y + \frac{\beta_i^a}{c_y^-}\right\}$. Similar to output, the dynamics of inflation depend on the sign of the shock, and the fact that $c_\pi^+ < c_\pi^-$ implies that inflation responds more to positive TFP shocks than to negative TFP shocks. Hence inflation exhibits the opposite asymmetry to output.

1.C.1 Idiosyncratic Risk

In this section, I relax the restriction that $\sigma_\epsilon = 0$, so that households experience idiosyncratic shocks to their labor productivity $\theta_{i,t}$. In order to maintain tractability, I assume instead that the idiosyncratic risk is "small", and study the economy in the limit $\sigma_\epsilon \to 0$. In addition, I assume that all households have the same average productivity level, $\theta_i = \bar{\theta}$ for all i. Since the response of output to interest rate changes lies at the heart of all of the asymmetry results, I focus on the case of monetary policy shocks.

For clarity, I state the assumption I require for tractability, which replaces assumption 1. All common restrictions have the same interpretation as before.

Assumption 2. *Let*

$$\rho_a, \rho_\Phi, \rho_v = 0$$

$$\Sigma \to 0$$

$$u(c,n) = \frac{\left(c - \frac{n^{1+\varphi}}{1+\varphi}\right)^{1-\sigma}}{1-\sigma}$$

$$\sigma_\epsilon \to 0 \ \forall i$$

$$\theta_i = \bar{\theta} \ \forall i$$

$$\underline{b}_{i,t} \to 0 \ \forall i, t$$

Stationary Equilibrium When all aggregate shocks are set to zero in all periods, I can define a stationary equilibrium. In this equilibrium, aggregate quantities and prices are constant over time, while households' choices of consumption and labor supply change over time as a function of their labor productivity, which is subject to idiosyncratic shocks.

Under assumption 2, there is a unique ergodic distribution of labor productivities in the economy, Λ_θ. Since the no-borrowing restriction implies that the wealth distribution is degenerate in equilibrium, Λ_θ is sufficient to compute cross-sectional averages of household consumption, labor supply, and income variables. Hence, the stationary equilibrium is unique.

Definition 3. Under assumption 2, and given initial conditions $\left\{\left\{b_{i,0},\theta_{i,0}\right\}_i, P_0\right\}$ where $\theta_{i,0} \sim \Lambda_\theta$ and $b_{i,0} = 0$ for all i, the unique stationary equilibrium is a sequence $\left\{\left\{c_{i,t}, n_{i,t}, b_{i,t}\right\}_{i,t}, \left\{y\left(j\right)\right\}_j, P, d, w, \iota\right\}$ such that

1. $\left\{c_{i,t}, n_{i,t}, b_{i,t}\right\}_{i,t}$ solve the household problem for each i.
2. $\left\{y\left(j\right)\right\}_j$ solve the final good firms' problem.
3. $P = P_0$ solves the intermediate goods firms' problem.
4. d satisfies the dividend equation.
5. Markets clear at every time $t \geq 1$.

In order to derive analytical results, I consider the dynamics of the economy in response to aggregate shocks around this unique stationary equilibrium. Formally, the following lemma condenses the economy to a set of four equations, analogous to lemma 1 for the case without idiosyncratic risk. Similar to that case, aggregate variables are expressed in log deviations around their values in the stationary equilibrium. Variables for household i are expressed as linear functions of these log deviations and also of log deviations of her labor productivity $\theta_{i,t}$ from its mean value θ_i. The approximation in the labor productivity dimension is valid since I work in the neighborhood of zero idiosyncratic risk, $\sigma_\epsilon \to 0$.

Lemma 2. *Under assumption 2, and to first order, the economy admits the following representation:*

$$\iota_t = r + \phi_\pi \pi_t + \phi_y \hat{y}_t + \epsilon_t^v$$

$$\pi_t = \frac{\bar{\Phi} - 1}{\xi^p} \varphi \hat{y}_t - \frac{\bar{\Phi} - 1}{\xi^p} (1 + \varphi) \epsilon_t^a - \frac{1}{\xi^p} \epsilon_t^\Phi + \delta \mathbb{E}_t [\pi_{t+1}]$$

$$\min_i \left\{ \mathbb{E}_t \left[\hat{\tilde{c}}_{i,t+1} \right] - \hat{\tilde{c}}_{i,t} \right\} = \frac{1}{\sigma} (\iota_t - \mathbb{E}_t [\pi_{t+1}] - \rho)$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t + \beta_i^a \epsilon_t^a + \beta_i^\theta \hat{\theta}_{i,t} \ \forall i$$

where r is the real interest rate in the stationary equilibrium, $\left\{ \beta_i^y, \beta_i^a, \beta_i^\theta \right\}_i$ depend only on model primitives, and where $\tilde{c}$ is consumption net of the disutility of labor supply, $\tilde{c} = c - \frac{n^{1+\varphi}}{1+\varphi}$.

$\phi_\pi > 1$ and $\phi_y \geq 0$ are sufficient to ensure that the system has a unique steady state, $\hat{y}_t = 0$, $\pi_t = 0$, $\hat{\tilde{c}}_{i,t} = \beta_i^\theta \hat{\theta}_{i,t} \ \forall i$.

Remark 1. The stationary real interest rate is given by

$$r = \rho - \sigma (1 - \rho_\theta) \max_i \left\{ \beta_i^\theta \hat{\theta}_{i,t} \right\} < \rho$$

In the presence of idiosyncratic labor productivity risk and binding borrowing constraints, households use the asset to build a buffer stock of savings as a self insurance mechanism. This increases the demand for the asset, and drives down the equilibrium real interest rate below the discount rate of households, $r < \rho$.

The empirical evidence suggests that the process for idiosyncratic shocks to labor income has a very persistent component (see, for example, Storesletten et al. (2004) and Guvenen et al. (2016)). In my setting, this is the case in which $\rho_\theta \to 1$.

Conveniently, this limit case also allows me to solve the system explicitly for the responses of output to monetary policy shocks.

Proposition 9. *Under assumption 2, and in the limit as $\rho_\theta \to 1$, the first order equilibrium dynamics of output in response to monetary policy shocks are given by*

$$\hat{y}_t = \begin{cases} -\frac{1}{\underline{\beta}+\frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y}\frac{1}{\sigma}\epsilon_t^v & \text{if} \quad \epsilon_t^v > 0 \\ \\ -\frac{1}{\bar{\beta}+\frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi+\frac{1}{\sigma}\phi_y}\frac{1}{\sigma}\epsilon_t^v & \text{if} \quad \epsilon_t^v < 0 \end{cases}$$

where

$$\bar{\beta} = \max_i \{\beta_i^y\}$$

$$\underline{\beta} = \min_i \{\beta_i^y\}$$

Proposition 9 shows that the responses of output are identical to the case without idiosyncratic risk. Therefore, the asymmetry is robust to the inclusion of idiosyncratic shocks that are very persistent.

In the presence of idiosyncratic risk, there are two distinct channels through which households can have low consumption growth, and hence be unconstrained in equilibrium. The first is the source of asymmetry that I have analyzed in earlier sections. When households' incomes are heterogeneously exposed to changes in aggregate income, binding borrowing constraints prevent the equalization of consumption sensitivities of output changes, and cause contractionary shocks to have larger effects than expansionary shocks.

The novel channel is due to idiosyncratic risk and is independent of heterogeneous income exposures, and so cannot be a source of asymmetry. When a household expects

to experience a drop in her labor productivity, her consumption growth will be low to the extent that the drop in labor productivity is uninsurable and hence transmits to her consumption.

When the process for labor productivity is very persistent ($\rho_\theta \to 1$), households expect their labor productivity to remain approximately constant across consecutive periods. This renders the novel savings channel inactive. As a result, savings choices are entirely determined by exposures to changes in output. In this case, the responses of output to monetary policy shocks are identical to the economy without idiosyncratic risk, as shown by the proposition.

Idiosyncratic shocks are often modeled as the sum of a persistent and transitory component. However, the empirical evidence suggests that households are very well insured against transitory shocks (Blundell et al., 2008; Heathcote et al., 2014), so that they do not affect consumption growth computations. Therefore, adding transitory shocks to labor productivity with a sufficiently rich set of contracts to provide insurance against them would complicate the model greatly without providing additional insights.

1.C.2 Heterogeneous Preferences

Let household i have preferences given by

$$u_i(c, n) = \frac{\left(c - \frac{n^{1+\varphi}}{1+\varphi}\right)^{1-\sigma_i}}{1-\sigma_i}$$

where $\sigma_i > 0$ for all i. The follow lemma summarizes the dynamics of the economy, and is the natural extension of lemma 1.

Lemma 3. *Under assumption 1, the economy's first order equilibrium dynamics in response to monetary policy shocks satisfy the system*

$$\iota_t = \rho + \phi_\pi \pi_t + \phi_y \hat{y}_t + \epsilon_t^v$$

$$\pi_t = \frac{\bar{\Phi} - 1}{\xi^p} \varphi \hat{y}_t + \delta \mathbb{E}_t \left[\pi_{t+1} \right]$$

$$\min_i \left\{ \sigma_i \left(\mathbb{E}_t \left[\hat{\tilde{c}}_{i,t+1} \right] - \hat{\tilde{c}}_{i,t} \right) \right\} = \iota_t - \mathbb{E}_t \left[\pi_{t+1} \right] - \rho$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t \ \forall i$$

where $\rho = -\log \delta$, $\{\beta_i^y\}_i$ *depend only on model primitives, and* $\tilde{c}$ *is consumption net of the disutility of labor supply,* $\tilde{c} = c - \frac{n^{1+\varphi}}{1+\varphi}$.

$\phi_\pi > 1$ *and* $\phi_y \geq 0$ *are sufficient to ensure that the system has a unique steady state,* $\hat{y}_t = 0$, $\pi_t = 0$, $\hat{\tilde{c}}_{i,t} = 0 \ \forall i$.

Given this lemma, we have the following closed-form representation of the asymmetric output responses to monetary policy shocks.

Proposition 10. *Under assumption 1, and when households have heterogeneous σ_i parameters, the first order equilibrium dynamics of output in response to monetary policy shocks are given by*

$$\hat{y}_t = \begin{cases} -\frac{1}{\underline{\beta\sigma}+\phi_\pi \frac{\Phi-1}{\xi^p}\varphi+\phi_y}\epsilon_t^v & if \quad \epsilon_t^v > 0 \\ \\ -\frac{1}{\bar{\beta\sigma}+\phi_\pi \frac{\Phi-1}{\xi^p}\varphi+\phi_y}\epsilon_t^v & if \quad \epsilon_t^v < 0 \end{cases}$$

where

$$\bar{\beta\sigma} = \max_i \{\sigma_i \beta_i^y\}$$

$$\underline{\beta\sigma} = \min_i \{\sigma_i \beta_i^y\}$$

Hence output responds more to contractionary monetary policy shocks than to expansionary shocks of equal magnitude. Clearly in the case $\sigma_i = \sigma$ for all i, the result simplifies to proposition 1.

1.D Proofs of Analytical Results

Proof of Lemma 1 I prove the lemma for the general case of all shocks.

The first equation is simply the Taylor rule for nominal interest rates, and is derived by taking natural logs of

$$1 + \iota_t = \frac{1}{\delta}(1+\pi_t)^{\phi_\pi}\left(\frac{Y_t}{Y^T}\right)^{\phi_y} e^{v_t}$$

and using $\log(1+x) \approx x$ for small x.

In order to derive the remaining equations, I first derive an expression for household income, $y_{i,t}$. Under GHH preferences, labor supply of household i is given by

$$n_{i,t} = \theta_{i,t}^{\frac{1}{\varphi}} w_t^{\frac{1}{\varphi}}$$

so that total household income is given by

$$y_{i,t} = w_t^{\frac{1+\varphi}{\varphi}} \theta_{i,t}^{\frac{1+\varphi}{\varphi}} + s_i d_t$$

Aggregating the labor supply condition over all households, and using the production function yields an expression for aggregate output,

$$Y_t = A_t w_t^{\frac{1}{\varphi}} \int_0^1 \theta_{i,t}^{\frac{1+\varphi}{\varphi}} di$$

so that the real wage is given by

$$w_t = \left(\frac{Y_t}{A_t \int_0^1 \theta_{i,t}^{\frac{1+\varphi}{\varphi}} di} \right)^{\varphi}$$

To first order, resource costs of inflation are zero. Hence dividends are given by

$$d_t = Y_t \left(1 - \frac{w_t}{A_t} \right)$$

Defining

$$\Theta_t = \int_0^1 \theta_{k,t}^{\frac{1+\varphi}{\varphi}} dk$$

and substituting these expressions into the equation for household income yields

$$y_{i,t} = \left(\frac{Y_t}{A_t \Theta_t} \right)^{1+\varphi} \Theta_t \left(\frac{\theta_{i,t}^{\frac{1+\varphi}{\varphi}}}{\Theta_t} - s_i \right) + s_i Y_t$$

Using $\theta_{i,t} = \theta_i$ for all t implies that

$$y_{i,t} = \left(\frac{Y_t}{A_t \Theta}\right)^{1+\varphi} \Theta \left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta} - s_i\right) + s_i Y_t$$

can be expressed as a function only of Y_t and A_t,

$$y_{i,t} = f_i\left(Y_t, A_t\right)$$

where the function index i stems from cross-sectional heterogeneity in θ_i and s_i.

The second equation is the New Keynesian Phillips Curve (NKPC), and follows from two steps. First, log linearizing the FOC of the intermediate goods firms' problem around the zero inflation deterministic equilibrium yields

$$\pi_t = \frac{\Phi_t}{\xi^p}\frac{w_t}{A_t} - \frac{\Phi_t - 1}{\xi^p} + \delta \mathbb{E}_t\left[\pi_{t+1}\right]$$

where the product terms are approximated by

$$\frac{\Phi_t}{\xi^p}\frac{w_t}{A_t} = \frac{\bar{\Phi}}{\xi^p}\frac{\bar{\Phi} - 1}{\bar{\Phi}} + \frac{\bar{\Phi} - 1}{\xi^p}\left(\log \Phi_t - \log \bar{\Phi}\right) + \frac{\bar{\Phi} - 1}{\xi^p}\left(\log \frac{w_t}{A_t} - \log \frac{\bar{\Phi} - 1}{\bar{\Phi}}\right)$$

and

$$\frac{\Phi_t - 1}{\xi^p} = \frac{\bar{\Phi} - 1}{\xi^p} + \frac{\bar{\Phi}}{\xi^p}\left(\log \Phi_t - \log \bar{\Phi}\right)$$

so that

$$\pi_t = \frac{\bar{\Phi} - 1}{\xi^p}\left(\log \frac{w_t}{A_t} - \log \frac{\bar{\Phi} - 1}{\bar{\Phi}}\right) - \frac{1}{\xi^p}\left(\log \Phi_t - \log \bar{\Phi}\right) + \delta \mathbb{E}_t\left[\pi_{t+1}\right]$$

Second, aggregation of the household labor supply condition stemming from GHH preferences yields

$$w_t = \left(\frac{Y_t}{A_t \Theta_t}\right)^{\varphi}$$

Taking logs yields

$$\log w_t - \log A_t = \varphi \log Y_t - (1+\varphi) \log A_t - \varphi \log \Theta_t$$

Using $\Theta_t = \Theta$, a first-order Taylor expansion of this equation around the deterministic equilibrium then yields

$$\log \frac{w_t}{A_t} - \log \frac{\bar{\Phi}-1}{\bar{\Phi}} = \varphi \hat{y}_t - (1+\varphi)\, \hat{a}_t$$

Substituting this into the NKPC yields

$$\pi_t = \frac{\bar{\Phi}-1}{\xi^p} \varphi \hat{y}_t - \frac{\bar{\Phi}-1}{\xi^p} (1+\varphi)\, \hat{a}_t - \frac{1}{\xi^p} \left(\log \Phi_t - \log \bar{\Phi} \right) + \delta \mathbb{E}_t \left[\pi_{t+1} \right]$$

Next, consider the third equation, which is the Euler equation for the economy. To derive this equation, note that the Euler equation for household i is given by

$$\tilde{c}_{i,t}^{-\sigma} \geq \delta \mathbb{E}_t \left[\tilde{c}_{i,t+1}^{-\sigma} (1 + r_{t+1}) \right]$$

where the inequality is strict if the borrowing constraint binds, and $\tilde{c} = c - \frac{n^{1+\varphi}}{1+\varphi}$ is consumption net of the disutility of labor supply (this occurs due to the GHH preference specification).

Therefore, the Euler equation features a "distortion" only if household i would like to borrow in equilibrium,

$$\tilde{c}_{i,t}^{-\sigma} > \delta \mathbb{E}_t \left[\tilde{c}_{i,t+1}^{-\sigma} (1 + r_{t+1}) \right] \iff b_{i,t} < 0$$

Hence, in equilibrium, there exists a household $i^*(t)$ such that

$$1 = \delta \mathbb{E}_t \left[(1 + r_{t+1}) \left(\frac{\tilde{c}_{i^*(t),t+1}}{\tilde{c}_{i^*(t),t}} \right)^{-\sigma} \right] \geq \delta \mathbb{E}_t \left[(1 + r_{t+1}) \left(\frac{\tilde{c}_{i,t+1}}{\tilde{c}_{i,t}} \right)^{-\sigma} \right]$$

for all $i \neq i^*(t)$, where $i^*(t)$ satisfies

$$i^*(t) \in \arg\max_i \mathbb{E}_t \left[(1 + r_{t+1}) \left(\frac{\tilde{c}_{i^*(t),t+1}}{\tilde{c}_{i^*(t),t}} \right)^{-\sigma} \right]$$

The aggregate Euler equation is therefore given by

$$1 = \delta \max_i \mathbb{E}_t \left[(1 + r_{t+1}) \left(\frac{\tilde{c}_{i,t+1}}{\tilde{c}_{i,t}} \right)^{-\sigma} \right]$$

Taking logs and using the first-order approximation $\log \mathbb{E}_t [x_{t+1}] \approx \mathbb{E}_t [\log x_{t+1}]$ yields

$$0 = \log \delta + \max_i \left\{ \mathbb{E}_t [\log (1 + r_{t+1})] - \sigma \mathbb{E}_t \left[\log \left(\frac{\tilde{c}_{i,t+1}}{\tilde{c}_{i,t}} \right) \right] \right\}$$

Writing $\rho = -\log \delta$ and using the approximation $\log(1 + r) \approx r$ together with the definition of the real interest rate simplifies the equation to

$$0 = i_t - \mathbb{E}_t [\pi_{t+1}] - \rho - \sigma \min_i \{ \mathbb{E}_t [\log \tilde{c}_{i,t+1}] - \log \tilde{c}_{i,t} \}$$

where I have also used the fact that

$$\max_i \{-X_i\} = -\min_i \{X_i\}$$

Using the log deviation around the deterministic equilibrium

$$\hat{\tilde{c}}_{i,t} = \log \tilde{c}_{i,t} - \log \tilde{c}_i$$

yields the third equation

$$\min_i \left\{ \mathbb{E}_t \left[\hat{\tilde{c}}_{i,t+1} \right] - \hat{\tilde{c}}_{i,t} \right\} = \frac{1}{\sigma} \left(i_t - \mathbb{E}_t \left[\pi_{t+1} \right] - \rho \right)$$

For the final equation, use the fact that in equilibrium

$$\tilde{c}_{i,t} = y_{i,t} - \frac{\theta_{i,t}^{\frac{1+\varphi}{\varphi}} w_t^{\frac{1+\varphi}{\varphi}}}{1+\varphi}$$

so that

$$\tilde{c}_{i,t} = y_{i,t} - \frac{\theta_{i,t}^{\frac{1+\varphi}{\varphi}} \left(\frac{Y_t}{A_t \Theta_t} \right)^{1+\varphi}}{1+\varphi}$$

Therefore, when $\theta_{i,t} = \theta_i$,

$$\log \tilde{c}_{i,t} = \log \left(y_{i,t} - \frac{\theta_i^{\frac{1+\varphi}{\varphi}} \left(\frac{Y_t}{A_t \Theta} \right)^{1+\varphi}}{1+\varphi} \right)$$

which can be linearly approximated around the deterministic equilibrium as

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t + \beta_i^a \epsilon_t^a$$

for some coefficients β_i^y and β_i^a that depend only model primitives, as required to complete the representation.

Note that

$$\log \tilde{c}_{i,t} = \log \left(y_{i,t} - \frac{\theta_i^{\frac{1+\varphi}{\varphi}} \left(\frac{Y_t}{A_t \Theta} \right)^{1+\varphi}}{1+\varphi} \right)$$

implies

$$\log \tilde{c}_{i,t} = \log \left(\left(\frac{Y_t}{A_t \Theta} \right)^{1+\varphi} \Theta \left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta} \frac{\varphi}{1+\varphi} - s_i \right) + s_i Y_t \right)$$

which has a first order expansion

$$\hat{\tilde{c}}_{i,t} \approx \frac{(1+\varphi)\left(\frac{Y}{A\Theta}\right)^{1+\varphi}\Theta\left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}-s_i\right)+s_iY}{\left(\frac{Y}{A\Theta}\right)^{1+\varphi}\Theta\left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}-s_i\right)+s_iY}\hat{y}_t$$

$$-\frac{(1+\varphi)\left(\frac{Y}{A\Theta}\right)^{1+\varphi}\Theta\left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}-s_i\right)}{\left(\frac{Y}{A\Theta}\right)^{1+\varphi}\Theta\left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}-s_i\right)+s_iY}\epsilon_t^a$$

In the deterministic equilibrium,

$$w=\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}=\left(\frac{Y}{A\Theta}\right)^{\varphi}$$

$$\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1+\varphi}{\varphi}}=\left(\frac{Y}{A\Theta}\right)^{1+\varphi}$$

$$Y=\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1}{\varphi}}\bar{A}\Theta$$

so that

$$\hat{\tilde{c}}_{i,t} \approx \frac{(1+\varphi)\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1+\varphi}{\varphi}}\Theta\left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}-s_i\right)+s_i\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1}{\varphi}}\bar{A}\Theta}{\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1+\varphi}{\varphi}}\Theta\left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}-s_i\right)+s_i\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1}{\varphi}}\bar{A}\Theta}\hat{y}_t$$

$$-\frac{(1+\varphi)\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1+\varphi}{\varphi}}\Theta\left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}-s_i\right)}{\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1+\varphi}{\varphi}}\Theta\left(\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}-s_i\right)+s_i\left(\frac{\bar{\Phi}-1}{\bar{\Phi}}\bar{A}\right)^{\frac{1}{\varphi}}\bar{A}\Theta}\epsilon_t^a$$

Hence

$$\beta_i^y=\frac{\varphi\frac{\bar{\Phi}-1}{\bar{\Phi}}\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}+s_i\left(1-(1+\varphi)\frac{\bar{\Phi}-1}{\bar{\Phi}}\right)}{\frac{\bar{\Phi}-1}{\bar{\Phi}}\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi}+s_i\frac{1}{\bar{\Phi}}}$$

$$\beta_i^a = -\frac{\varphi\frac{\bar{\Phi}-1}{\bar{\Phi}}\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta} - (1+\varphi)\frac{\bar{\Phi}-1}{\bar{\Phi}}s_i}{\frac{\bar{\Phi}-1}{\bar{\Phi}}\frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta}\frac{\varphi}{1+\varphi} + s_i\frac{1}{\bar{\Phi}}}$$

so that

$$\beta_i^y > 0 \iff \frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta} > s_i\left(1 - \frac{1}{\varphi\left(\bar{\Phi}-1\right)}\right)$$

$$\beta_i^a > 0 \iff \frac{\theta_i^{\frac{1+\varphi}{\varphi}}}{\Theta} < s_i\frac{1+\varphi}{\varphi}$$

I assume that both conditions are satisfied. □

Proof of Propositions 1 and 6 Recall the representation

$$i_t = \rho + \phi_\pi\pi_t + \phi_y\hat{y}_t + \epsilon_t^v$$

$$\pi_t = \frac{\bar{\Phi}-1}{\xi^p}\varphi\hat{y}_t + \delta\mathbb{E}_t\left[\pi_{t+1}\right]$$

$$\min_i\left\{\mathbb{E}_t\left[\hat{\tilde{c}}_{i,t+1}\right] - \hat{\tilde{c}}_{i,t}\right\} = \frac{1}{\sigma}\left(i_t - \mathbb{E}_t\left[\pi_{t+1}\right] - \rho\right)$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y\hat{y}_t$$

where I have set $\epsilon_t^a = \epsilon_t^\Phi = 0$ by assumption.

Consider $\epsilon_t^v > 0$, and suppose that the solution takes the form

$$\hat{y}_t = c_y^+\epsilon_t^v$$

$$\pi_t = c_\pi^+\epsilon_t^v$$

where $c_y^+ < 0$. Substituting these guesses into the system and simplifying yields

$$c_\pi^+ = \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^+$$

$$\min_i \left\{-\beta_i^y c_y^+ \epsilon_t^v\right\} = \frac{1}{\sigma}\left(\phi_\pi c_\pi^+ \epsilon_t^v + \phi_y c_y^+ \epsilon_t^v + \epsilon_t^v\right)$$

so that

$$\min_i \left\{-\beta_i^y c_y^+ \epsilon_t^v\right\} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^+ + \phi_y c_y^+ + 1\right)\epsilon_t^v$$

By the supposition, $-c_y^+ \epsilon_t^v > 0$ so that

$$-c_y^+ \epsilon_t^v \underline{\beta} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^+ + \phi_y c_y^+ + 1\right)\epsilon_t^v$$

where

$$\underline{\beta} = \min_i \left\{\beta_i^y\right\}$$

Hence

$$c_y^+ = -\frac{1}{\underline{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y}\frac{1}{\sigma}$$

$$c_\pi^+ = -\frac{1}{\underline{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y}\frac{1}{\sigma}\frac{\bar{\Phi}-1}{\xi^p}\varphi$$

where $c_y^+ < 0$ since

$$\underline{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y > 0$$

For completeness, suppose $c_y^+ > 0$. The same steps yield

$$\min_i \left\{-\beta_i^y c_y^+ \epsilon_t^v\right\} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^+ + \phi_y c_y^+ + 1\right)\epsilon_t^v$$

By supposition, $c_y^+ \epsilon_t^v > 0$ so that

$$-c_y^+ \epsilon_t^v \bar{\beta} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^+ + \phi_y c_y^+ + 1\right)\epsilon_t^v$$

where

$$\bar{\beta} = \max_i \left\{\beta_i^y\right\} > 0$$

Hence

$$c_y^+ = -\frac{1}{\bar{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y}\frac{1}{\sigma} < 0$$

where the inequality follows from

$$\bar{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y > 0$$

Therefore, we have a contradiction.

Now consider $\epsilon_t^v < 0$, and suppose that the solution takes the form

$$\hat{y}_t = c_y^- \epsilon_t^v$$

$$\pi_t = c_\pi^- \epsilon_t^v$$

where $c_y^- < 0$. Substituting these guesses into the system and simplifying yields

$$c_\pi^- = \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- > 0$$

$$\min_i \left\{-\beta_i^y c_y^- \epsilon_t^v\right\} = \frac{1}{\sigma}\left(\phi_\pi c_\pi^- \epsilon_t^v + \phi_y c_y^- \epsilon_t^v + \epsilon_t^v\right)$$

so that

$$\min_i \left\{-\beta_i^y c_y^- \epsilon_t^v\right\} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- + \phi_y c_y^- + 1\right)\epsilon_t^v$$

By the supposition, $c_y^- \epsilon_t^v > 0$ so that

$$-c_y^- \epsilon_t^v \bar{\beta} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- + \phi_y c_y^- + 1\right)\epsilon_t^v$$

where

$$\bar{\beta} = \max_i \{\beta_i^y\}$$

Hence

$$c_y^- = -\frac{1}{\bar{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y} \frac{1}{\sigma} < 0$$

$$c_\pi^+ = -\frac{1}{\bar{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y} \frac{1}{\sigma} \frac{\bar{\Phi}-1}{\xi^p}\varphi_\Theta < 0$$

as required.

For completeness, suppose that $c_y^- > 0$ so that $-c_y^- \epsilon_t^v > 0$ and the same steps as above lead to

$$-c_y^- \epsilon_t^v \underline{\beta} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- + \phi_y c_y^- + 1\right)\epsilon_t^v$$

so that

$$c_y^- = -\frac{1}{\underline{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y} \frac{1}{\sigma} < 0$$

which is a contradiction.[21] $\square$

Proof of Propositions 4 and 7 Recall the representation

$$i_t = \rho + \phi_\pi \pi_t + \phi_y \hat{y}_t$$

$$\pi_t = \frac{\bar{\Phi}-1}{\xi^p}\varphi \hat{y}_t - \frac{1}{\xi^p}\epsilon_t^\Phi + \delta \mathbb{E}_t\left[\pi_{t+1}\right]$$

$$\min_i \left\{\mathbb{E}_t\left[\hat{\tilde{c}}_{i,t+1}\right] - \hat{\tilde{c}}_{i,t}\right\} = \frac{1}{\sigma}\left(i_t - \mathbb{E}_t\left[\pi_{t+1}\right] - \rho\right)$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t$$

where I have set $\epsilon_t^a = \epsilon_t^v = 0$ by assumption.

[21] In order to derive closed form expressions for the impulse responses, I have assumed that $\mathbb{E}_t\left[\hat{y}_{t+1}\right] = \mathbb{E}_t\left[\pi_{t+1}\right] = 0$. Given that aggregate shocks are i.i.d., this assumption amounts to the following restriction on the distributions of the shocks,

$$\mathbb{E}_t\left[\hat{y}_{t+1}\right] = 0 \iff c_y^+ \mathbb{E}\left[\epsilon_{t+1} | \epsilon_{t+1} > 0\right] \Pr\left(\epsilon_{t+1} > 0\right) = -c_y^- \mathbb{E}\left[\epsilon_{t+1} | \epsilon_{t+1} < 0\right] \Pr\left(\epsilon_{t+1} < 0\right)$$

However, since I work in the limiting case $\Sigma \to 0$, both sides of the equality are approximately equal for small enough ϵ_{t+1}.

Consider $\epsilon_t^{\Phi} > 0$, and suppose that the solution takes the form

$$\hat{y}_t = c_y^+ \epsilon_t^{\Phi}$$

$$\pi_t = c_\pi^+ \epsilon_t^{\Phi}$$

where $c_y^+ > 0$. Substituting these guesses into the system and simplifying yields

$$c_\pi^+ = \frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^+ - \frac{1}{\xi^p}$$

$$\min_i \left\{ -\beta_i^y c_y^+ \epsilon_t^{\Phi} \right\} = \frac{1}{\sigma} \left(\phi_\pi c_\pi^+ \epsilon_t^{\Phi} + \phi_y c_y^+ \epsilon_t^{\Phi} \right)$$

so that

$$\min_i \left\{ -\beta_i^y c_y^+ \epsilon_t^{\Phi} \right\} = \frac{1}{\sigma} \left(\phi_\pi \left(\frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^+ - \frac{1}{\xi^p} \right) + \phi_y c_y^+ \right) \epsilon_t^{\Phi}$$

By supposition, $c_y^+ \epsilon_t^{\Phi} > 0$ so that

$$-c_y^+ \epsilon_t^{\Phi} \bar{\beta} = \frac{1}{\sigma} \left(\phi_\pi \left(\frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^+ - \frac{1}{\xi^p} \right) + \phi_y c_y^+ \right) \epsilon_t^{\Phi}$$

where

$$\bar{\beta} = \max_i \{ \beta_i^y \} > 0$$

Hence

$$c_y^+ = \frac{\phi_\pi \frac{1}{\xi^p}}{\bar{\beta} + \frac{1}{\sigma} \phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi + \frac{1}{\sigma} \phi_y} \frac{1}{\sigma} > 0$$

$$c_\pi^+ = -\frac{1}{\xi^p} \left(\frac{\bar{\beta} + \frac{1}{\sigma} \phi_y}{\bar{\beta} + \frac{1}{\sigma} \phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi + \frac{1}{\sigma} \phi_y} \right)$$

as required.

For completeness, suppose that $c_y^+ < 0$ so that $-c_y^+ \epsilon_t^\Phi > 0$ and

$$c_y^+ = \frac{\phi_\pi \frac{1}{\xi^p}}{\underline{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y} \frac{1}{\sigma} > 0$$

where

$$\underline{\beta} = \min_i \{\beta_i^y\}$$

which is a contradiction.

Now consider $\epsilon_t^\Phi < 0$, and suppose that the solution takes the form

$$\hat{y}_t = c_y^- \epsilon_t^\Phi$$

$$\pi_t = c_\pi^- \epsilon_t^\Phi$$

where $c_y^- > 0$. Substituting these guesses into the system and simplifying yields

$$c_\pi^- = \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- - \frac{1}{\xi^p}$$

$$\min_i \left\{-\beta_i^y c_y^- \epsilon_t^\Phi\right\} = \frac{1}{\sigma}\left(\phi_\pi c_\pi^- \epsilon_t^\Phi + \phi_y c_y^- \epsilon_t^\Phi\right)$$

so that

$$\min_i \left\{-\beta_i^y c_y^- \epsilon_t^\Phi\right\} = \frac{1}{\sigma}\left(\phi_\pi\left(\frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- - \frac{1}{\xi^p}\right) + \phi_y c_y^-\right)\epsilon_t^\Phi$$

By supposition, $-c_y^- \epsilon_t^\Phi > 0$ so that

$$-c_y^- \epsilon_t^\Phi \underline{\beta}^{CP} = \frac{1}{\sigma}\left(\phi_\pi\left(\frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- - \frac{1}{\xi^p}\right) + \phi_y c_y^-\right)\epsilon_t^\Phi$$

where

$$\underline{\beta} = \min_i \{\beta_i^y\}$$

Hence

$$c_y^- = \frac{\phi_\pi \frac{1}{\xi^p}}{\underline{\beta}^{CP} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y} \frac{1}{\sigma} > 0$$

$$c_\pi^+ = -\frac{1}{\xi^p}\left(\frac{\underline{\beta}^{CP} + \frac{1}{\sigma}\phi_y}{\underline{\beta}^{CP} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y}\right)$$

as required.

For completeness, suppose that $c_y^- < 0$ so that $c_y^- \epsilon_t^\Phi > 0$ and

$$c_y^- = \frac{\phi_\pi \frac{1}{\xi^p}}{\bar{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y} \frac{1}{\sigma} > 0$$

where

$$\bar{\beta} = \max_i \{\beta_i^y\}$$

which is a contradiction. □

Proof of Propositions 5 and 8 Recall the representation

$$i_t = \rho + \phi_\pi \pi_t + \phi_y \hat{y}_t$$

$$\pi_t = \frac{\bar{\Phi}-1}{\xi^p}\varphi \hat{y}_t - \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)\epsilon_t^a + \delta \mathbb{E}_t[\pi_{t+1}]$$

$$\min_i \left\{\mathbb{E}_t\left[\hat{\tilde{c}}_{i,t+1}\right] - \hat{\tilde{c}}_{i,t}\right\} = \frac{1}{\sigma}(i_t - \mathbb{E}_t[\pi_{t+1}] - \rho)$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t + \beta_i^a \epsilon_t^a$$

where I have set $\epsilon_t^\Phi = \epsilon_t^v = 0$ by assumption.

Consider $\epsilon_t^a > 0$, and suppose that the solution is given by

$$\hat{y}_t = c_y^+ \epsilon_t^a$$

$$\pi_t = c_\pi^+ \epsilon_t^a$$

where $c_y^+ > 0$. Substituting these guesses into the the system and simplifying yields

$$c_\pi^+ = \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^+ - \frac{\bar{\Phi}-1}{\xi^p}\left(1+\varphi\right)$$

$$\min_i \left\{-\left(\beta_i^y c_y^+ + \beta_i^a\right)\epsilon_t^a\right\} = \frac{1}{\sigma}\left(\phi_\pi c_\pi^+ + \phi_y c_y^+\right)\epsilon_t^a$$

so that

$$\min_i \left\{-\left(\beta_i^y c_y^+ + \beta_i^a\right)\epsilon_t^a\right\} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^+ - \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\left(1+\varphi\right) + \phi_y c_y^+\right)\epsilon_t^a$$

Now define $\beta_i^{TFP,+} = \beta_i^y + \frac{\beta_i^a}{c_y^+}$ so that

$$\min_i \left\{-\beta_i^{TFP,+} c_y^+ \epsilon_t^a\right\} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^+ - \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\left(1+\varphi\right) + \phi_y c_y^+\right)\epsilon_t^a$$

where $c_y^+ \epsilon_t^a > 0$ by the supposition, so that

$$c_y^+ = \frac{\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\left(1+\varphi\right)}{\bar{\beta}^{TFP,+} + \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)}\frac{1}{\sigma}$$

$$c_\pi^+ = -\frac{\bar{\Phi}-1}{\xi^p}\left(1+\varphi\right)\frac{\bar{\beta}^{TFP,+} + \frac{1}{\sigma}\phi_y}{\bar{\beta}^{TFP,+} + \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi_\Theta + \phi_y\right)}$$

where

$$\bar{\beta}^{TFP,+} = \max_i \left\{\beta_i^y + \frac{\beta_i^a}{c_y^+}\right\}$$

To ensure that $c_y^+ > 0$, first rewrite

$$c_y^+ = \frac{\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\left(1+\varphi\right)\frac{1}{\sigma} - \beta_{i^*}^a}{\beta_{i^*}^y + \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)}$$

where

$$i^* = \arg\max_i \left\{\beta_i^y + \frac{\beta_i^a}{c_y^+}\right\}$$

Then, I impose the restriction

$$\beta_i^a < \phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} (1 + \varphi) \frac{1}{\sigma}$$

for all i, which guarantees $c_y^+ > 0$ as required. This restriction requirement is purely a result of using GHH preferences: a positive TFP shock lowers labor supply ceteris paribus, which boosts a household's net consumption due to the strong complementarity between labor supply and consumption embodied by GHH preferences. If this complementarity is strong enough, output can fall in response to a positive TFP shock in equilibrium since households are happy to consume less but also work less. Since this feature of GHH preferences is not the focus of my analysis, I rule it out by assumption.

Now consider $\epsilon_t^a < 0$, and suppose that the solution is given by

$$\hat{y}_t = c_y^- \epsilon_t^a$$

$$\pi_t = c_\pi^- \epsilon_t^a$$

where $c_y^- > 0$. Substituting these guesses into the the system and simplifying yields

$$c_\pi^- = \frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^- - \frac{\bar{\Phi} - 1}{\xi^p} (1 + \varphi)$$

$$\min_i \left\{ - \left(\beta_i^y c_y^- + \beta_i^a \right) \epsilon_t^a \right\} = \frac{1}{\sigma} \left(\phi_\pi c_\pi^- + \phi_y c_y^- \right) \epsilon_t^a$$

Now define $\beta_i^{TFP,-} = \beta_i^y + \frac{\beta_i^a}{c_y^-}$ so that

$$\min_i \left\{ -\beta_i^{TFP,-} c_y^- \epsilon_t^a \right\} = \frac{1}{\sigma} \left(\phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^- - \phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} (1 + \varphi) + \phi_y c_y^- \right) \epsilon_t^a$$

where $-c_y^- \epsilon_t^a > 0$ by the supposition, so that

$$c_y^- = \frac{\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)}{\underline{\beta}^{TFP,-} + \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)} \frac{1}{\sigma} > 0$$

$$c_\pi^- = -\frac{\bar{\Phi}-1}{\xi^p}(1+\varphi) \frac{\underline{\beta}^{TFP,-} + \frac{1}{\sigma}\phi_y}{\underline{\beta}^{TFP,-} + \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)}$$

where

$$\underline{\beta}^{TFP,-} = \min_i \left\{ \beta_i^y + \frac{\beta_i^a}{c_y^-} \right\}$$

To ensure that $c_y^- > 0$, rewrite

$$c_y^- = \frac{\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)\frac{1}{\sigma} - \beta_{i^{**}}^a}{\beta_{i^{**}}^y + \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)}$$

where

$$i^{**} = \arg\min_i \left\{ \beta_i^y + \frac{\beta_i^a}{c_y^-} \right\}$$

The restriction

$$\beta_i^a < \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)\frac{1}{\sigma}$$

for all i then ensures that $c_y^- > 0$.

Finally, in order to prove that $c_y^- > c_y^+$, define the functions

$$g^+(c) = \frac{\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)}{\sigma \max_i \left\{ \beta_i^y + \frac{\beta_i^a}{c} \right\} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y} - c$$

$$g^-(c) = \frac{\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)}{\sigma \min_i \left\{ \beta_i^y + \frac{\beta_i^a}{c} \right\} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y} - c$$

where $g^{-}(c) > g^{+}(c)$ for all $c > 0$, and the coefficients c_y^{-} and c_y^{+} satisfy

$$g^{-}\left(c_y^{-}\right) = 0$$

$$g^{+}\left(c_y^{+}\right) = 0$$

Then, the fact that $g^{-}(c) > g^{+}(c)$ implies that $c_y^{+} < c_y^{-}$ as required. □

Proof of Lemma 2 The first equation is simply the Taylor rule for nominal interest rates. In order to derive the remaining equations, I first derive an expression for household income, $y_{i,t}$. Under GHH preferences, labor supply of household i is given by

$$n_{i,t} = \theta_{i,t}^{\frac{1}{\varphi}} w_t^{\frac{1}{\varphi}}$$

so that total household income is given by

$$y_{i,t} = w_t^{\frac{1+\varphi}{\varphi}} \theta_{i,t}^{\frac{1+\varphi}{\varphi}} + s_i d_t$$

Aggregating the labor supply condition over all households, and using the production function yields an expression for aggregate output,

$$Y_t = A_t w_t^{\frac{1}{\varphi}} \int_0^1 \theta_{i,t}^{\frac{1+\varphi}{\varphi}} di$$

so that the real wage is given by

$$w_t = \left(\frac{Y_t}{A_t \int_0^1 \theta_{i,t}^{\frac{1+\varphi}{\varphi}} di} \right)^{\varphi}$$

To first order, resource costs of inflation are zero. Hence dividends are given by

$$d_t = Y_t \left(1 - \frac{w_t}{A_t}\right)$$

Defining

$$\Theta_t = \int_0^1 \theta_{k,t}^{\frac{1+\varphi}{\varphi}} dk$$

and substituting these expressions into the equation for household income yields

$$y_{i,t} = \left(\frac{Y_t}{A_t \Theta_t}\right)^{1+\varphi} \Theta_t \left(\frac{\theta_{i,t}^{\frac{1+\varphi}{\varphi}}}{\Theta_t} - s_i\right) + s_i Y_t$$

In the stationary equilibrium, $\Theta_t = \Theta$ by construction, so that

$$y_{i,t} = \left(\frac{Y_t}{A_t \Theta}\right)^{1+\varphi} \Theta \left(\frac{\theta_{i,t}^{\frac{1+\varphi}{\varphi}}}{\Theta} - s_i\right) + s_i Y_t$$

The second equation is the New Keynesian Phillips Curve (NKPC), and follows from two steps. First, log linearizing the FOC of the intermediate goods firms' problem around the zero inflation stationary equilibrium yields

$$\pi_t = \frac{\Phi_t}{\xi^p} \frac{w_t}{A_t} - \frac{\Phi_t - 1}{\xi^p} + \delta \mathbb{E}_t \left[\pi_{t+1}\right]$$

where the product terms are approximated by

$$\frac{\Phi_t}{\xi^p} \frac{w_t}{A_t} = \frac{\bar{\Phi}}{\xi^p} \frac{\bar{\Phi} - 1}{\bar{\Phi}} + \frac{\bar{\Phi} - 1}{\xi^p} \left(\log \Phi_t - \log \bar{\Phi}\right) + \frac{\bar{\Phi} - 1}{\xi^p} \left(\log \frac{w_t}{A_t} - \log \frac{\bar{\Phi} - 1}{\bar{\Phi}}\right)$$

and

$$\frac{\Phi_t - 1}{\xi^p} = \frac{\bar{\Phi} - 1}{\xi^p} + \frac{\bar{\Phi}}{\xi^p} \left(\log \Phi_t - \log \bar{\Phi}\right)$$

so that

$$\pi_t = \frac{\bar{\Phi}-1}{\xi^p}\left(\log\frac{w_t}{A_t} - \log\frac{\bar{\Phi}-1}{\bar{\Phi}}\right) - \frac{1}{\xi^p}\left(\log\Phi_t - \log\bar{\Phi}\right) + \delta\mathbb{E}_t\left[\pi_{t+1}\right]$$

Second, aggregation of the household labor supply condition stemming from GHH preferences yields

$$w_t = \left(\frac{Y_t}{A_t\Theta_t}\right)^{\varphi}$$

Taking logs yields

$$\log w_t - \log A_t = \varphi\log Y_t - (1+\varphi)\log A_t - \varphi\log\Theta_t$$

Around the stationary equilibrium, $\Theta_t = \Theta$, and a first-order Taylor expansion of this equation yields

$$\log\frac{w_t}{A_t} - \log\frac{\bar{\Phi}-1}{\bar{\Phi}} = \varphi\hat{y}_t - (1+\varphi)\,\hat{a}_t$$

Substituting this into the NKPC yields

$$\pi_t = \frac{\bar{\Phi}-1}{\xi^p}\varphi\hat{y}_t - \frac{\bar{\Phi}-1}{\xi^p}(1+\varphi)\,\hat{a}_t - \frac{1}{\xi^p}\left(\log\Phi_t - \log\bar{\Phi}\right) + \delta\mathbb{E}_t\left[\pi_{t+1}\right]$$

Next, consider the third equation, which is the Euler equation for the economy. To derive this equation, note that the Euler equation for household i is given by

$$\tilde{c}_{i,t}^{-\sigma} \geq \delta\mathbb{E}_t\left[\tilde{c}_{i,t+1}^{-\sigma}\left(1+r_{t+1}\right)\right]$$

where the inequality is strict if the borrowing constraint binds, and $\tilde{c} = c - \frac{n^{1+\varphi}}{1+\varphi}$ is consumption net of the disutility of labor supply (this occurs due to the GHH preference specification).

Therefore, the Euler equation features a "distortion" only if household i would like to borrow in equilibrium,

$$\tilde{c}_{i,t}^{-\sigma} > \delta \mathbb{E}_t \left[\tilde{c}_{i,t+1}^{-\sigma} \left(1 + r_{t+1}\right) \right] \iff b_{i,t} < 0$$

Hence, in equilibrium, there exists a household $i^*(t)$ such that

$$1 = \delta \mathbb{E}_t \left[\left(1 + r_{t+1}\right) \left(\frac{\tilde{c}_{i^*(t),t+1}}{\tilde{c}_{i^*(t),t}} \right)^{-\sigma} \right] \geq \delta \mathbb{E}_t \left[\left(1 + r_{t+1}\right) \left(\frac{\tilde{c}_{i,t+1}}{\tilde{c}_{i,t}} \right)^{-\sigma} \right]$$

for all $i \neq i^*(t)$, where $i^*(t)$ satisfies

$$i^*(t) \in \arg\max_i \mathbb{E}_t \left[\left(1 + r_{t+1}\right) \left(\frac{\tilde{c}_{i^*(t),t+1}}{\tilde{c}_{i^*(t),t}} \right)^{-\sigma} \right]$$

The aggregate Euler equation is therefore given by

$$1 = \delta \max_i \mathbb{E}_t \left[\left(1 + r_{t+1}\right) \left(\frac{\tilde{c}_{i,t+1}}{\tilde{c}_{i,t}} \right)^{-\sigma} \right]$$

Taking logs and using the first-order approximation $\log \mathbb{E}_t \left[x_{t+1}\right] \approx \mathbb{E}_t \left[\log x_{t+1}\right]$ yields

$$0 = \log \delta + \max_i \left\{ \mathbb{E}_t \left[\log \left(1 + r_{t+1}\right) \right] - \sigma \mathbb{E}_t \left[\log \left(\frac{\tilde{c}_{i,t+1}}{\tilde{c}_{i,t}} \right) \right] \right\}$$

Writing $\rho = -\log \delta$ and using the approximation $\log\left(1 + r\right) \approx r$ together with the definition of the real interest rate simplifies the equation to

$$0 = i_t - \mathbb{E}_t \left[\pi_{t+1}\right] - \rho - \sigma \min_i \left\{ \mathbb{E}_t \left[\log \tilde{c}_{i,t+1}\right] - \log \tilde{c}_{i,t} \right\}$$

where I have also used the fact that

$$\max_i \left\{-X_i\right\} = -\min_i \left\{X_i\right\}$$

Using the log deviation around the stationary equilibrium and mean labor productivity level

$$\hat{\tilde{c}}_{i,t} = \log \tilde{c}_{i,t} - \log \tilde{c}_i$$

yields the third equation

$$\min_i \left\{ \mathbb{E}_t \left[\hat{\tilde{c}}_{i,t+1} \right] - \hat{\tilde{c}}_{i,t} \right\} = \frac{1}{\sigma} \left(i_t - \mathbb{E}_t \left[\pi_{t+1} \right] - \rho \right)$$

For the final equation, use the fact that in equilibrium

$$\tilde{c}_{i,t} = y_{i,t} - \frac{\theta_{i,t}^{\frac{1+\varphi}{\varphi}} w_t^{\frac{1+\varphi}{\varphi}}}{1+\varphi}$$

so that

$$\tilde{c}_{i,t} = y_{i,t} - \frac{\theta_{i,t}^{\frac{1+\varphi}{\varphi}} \left(\frac{Y_t}{A_t \Theta} \right)^{1+\varphi}}{1+\varphi}$$

Therefore,

$$\log \tilde{c}_{i,t} = \log \left(y_{i,t} - \frac{\theta_i^{\frac{1+\varphi}{\varphi}} \left(\frac{Y_t}{A_t \Theta} \right)^{1+\varphi}}{1+\varphi} \right)$$

which can be linearly approximated around the stationary equilibrium and mean labor productivity level as

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t + \beta_i^a \epsilon_t^a + \beta_i^\theta \hat{\theta}_{i,t}$$

for some coefficients β_i^y, β_i^a, and β_i^θ that depend only model primitives, as required to complete the representation. □

Proof of Remark 1 In the stationary equilibrium, the economy is summarized by the two equations

$$\min_i \left\{ \mathbb{E}_t \left[\hat{\tilde{c}}_{i,t+1} \right] - \hat{\tilde{c}}_{i,t} \right\} = \frac{1}{\sigma} \left(r - \rho \right)$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^\theta \hat{\theta}_{i,t} \ \forall i$$

so that

$$\min_i \left\{ \beta_i^\theta \left(\mathbb{E}_t \left[\hat{\theta}_{i,t+1} \right] - \hat{\theta}_{i,t} \right) \right\} = \frac{1}{\sigma} (r - \rho)$$

Since

$$\hat{\theta}_{i,t} = \rho_\theta \hat{\theta}_{i,t-1} + \epsilon_{i,t}$$

we have

$$-(1 - \rho_\theta) \max_i \left\{ \beta_i^\theta \hat{\theta}_{i,t} \right\} = \frac{1}{\sigma} (r - \rho)$$

$$r = \rho - \sigma (1 - \rho_\theta) \max_i \left\{ \beta_i^\theta \hat{\theta}_{i,t} \right\} < \rho$$

as required. □

Proof of Proposition 9 Recall the system

$$\iota_t = r + \phi_\pi \pi_t + \phi_y \hat{y}_t + \epsilon_t^v$$

$$\pi_t = \frac{\bar{\Phi} - 1}{\xi^p} \varphi \hat{y}_t + \delta \mathbb{E}_t \left[\pi_{t+1} \right]$$

$$\min_i \left\{ \mathbb{E}_t \left[\hat{\tilde{c}}_{i,t+1} \right] - \hat{\tilde{c}}_{i,t} \right\} = \frac{1}{\sigma} (\iota_t - \mathbb{E}_t \left[\pi_{t+1} \right] - \rho)$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t + \beta_i^\theta \hat{\theta}_{i,t} \ \forall i$$

where I have set $\epsilon_t^a = \epsilon_t^\Phi = 0$ by assumption. Substitution yields

$$\min_i \left\{ \beta_i^y \left(\mathbb{E}_t \left[\hat{y}_{t+1} \right] - \hat{y}_t \right) + \beta_i^\theta \left(\mathbb{E}_t \left[\hat{\theta}_{i,t+1} \right] - \hat{\theta}_{i,t} \right) \right\}$$

$$=$$

$$\frac{1}{\sigma} \left(\phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi \hat{y}_t + \phi_y \hat{y}_t + \epsilon_t^v + r - \rho \right)$$

Suppose $\epsilon_t^v > 0$. Guess a solution of the form $\hat{y}_t = c_y^+ \epsilon_t^v$ with $c_y^+ < 0$. Using this and the fact that

$$\hat{\theta}_{i,t+1} = \rho_\theta \hat{\theta}_{i,t} + \epsilon_{i,t+1}$$

$$-\frac{1}{\sigma}(r-\rho) = (1-\rho_\theta)\max_j\left\{\beta_j^\theta \hat{\theta}_{j,t}\right\}$$

yields

$$\min_i\left\{-\beta_i^y c_y^+ \epsilon_t^v + (1-\rho_\theta)\left(\max_j\left\{\beta_j^\theta\hat{\theta}_{j,t}\right\} - \beta_i^\theta\hat{\theta}_{i,t}\right)\right\}$$
$$=$$
$$\tfrac{1}{\sigma}\left(\phi_\pi\tfrac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)c_y^+\epsilon_t^v + \tfrac{1}{\sigma}\epsilon_t^v$$

Taking the limit $\rho_\theta \to 1$ then implies that $(1-\rho_\theta)\left(\max_j\left\{\beta_j^\theta\hat{\theta}_{j,t}\right\} - \beta_i^\theta\hat{\theta}_{i,t}\right)$ is of second order so that the equation becomes

$$\min_i\left\{-\beta_i^y c_y^+\epsilon_t^v\right\} = \frac{1}{\sigma}\left(\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)c_y^+\epsilon_t^v + \frac{1}{\sigma}\epsilon_t^v$$

which implies

$$c_y^+ = -\frac{1}{\underline{\beta} + \frac{1}{\sigma}\phi_\pi\frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y}\frac{1}{\sigma}$$

where

$$\underline{\beta} = \min_i\{\beta_i^y\}$$

Now suppose $\epsilon_t^v < 0$. Guess a solution of the form $\hat{y}_t = c_y^-\epsilon_t^v$ with $c_y^- < 0$. Using this and the fact that

$$\hat{\theta}_{i,t+1} = \rho_\theta\hat{\theta}_{i,t} + \epsilon_{i,t+1}$$

$$-\frac{1}{\sigma}(r-\rho) = (1-\rho_\theta)\max_j\left\{\beta_j^\theta \hat{\theta}_{j,t}\right\}$$

yields

$$\min_i\left\{-\beta_i^y c_y^- \epsilon_t^v + (1-\rho_\theta)\left(\max_j\left\{\beta_j^\theta\hat{\theta}_{j,t}\right\} - \beta_i^\theta\hat{\theta}_{i,t}\right)\right\}$$
$$=$$
$$\tfrac{1}{\sigma}\left(\phi_\pi\tfrac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right)c_y^-\epsilon_t^v + \tfrac{1}{\sigma}\epsilon_t^v$$

Taking the limit $\rho_\theta \to 1$ then implies that $(1-\rho_\theta)\left(\max_j\left\{\beta_j^\theta \hat{\theta}_{j,t}\right\} - \beta_i^\theta \hat{\theta}_{i,t}\right)$ is of second order so that the equation becomes

$$\min_i \left\{-\beta_i^y c_y^- \epsilon_t^v\right\} = \frac{1}{\sigma}\left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y\right) c_y^- \epsilon_t^v + \frac{1}{\sigma}\epsilon_t^v$$

which implies

$$c_y^- = -\frac{1}{\bar{\beta} + \frac{1}{\sigma}\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \frac{1}{\sigma}\phi_y}\frac{1}{\sigma}$$

where

$$\bar{\beta} = \max_i \left\{\beta_i^y\right\}$$

as required. □

Proof of Proposition 10 Recall the system

$$\iota_t = \rho + \phi_\pi \pi_t + \phi_y \hat{y}_t + \epsilon_t^v$$

$$\pi_t = \frac{\bar{\Phi}-1}{\xi^p}\varphi \hat{y}_t + \delta \mathbb{E}_t\left[\pi_{t+1}\right]$$

$$\min_i \left\{\sigma_i\left(\mathbb{E}_t\left[\hat{\tilde{c}}_{i,t+1}\right] - \hat{\tilde{c}}_{i,t}\right)\right\} = \left(\iota_t - \mathbb{E}_t\left[\pi_{t+1}\right] - \rho\right)$$

$$\hat{\tilde{c}}_{i,t} = \beta_i^y \hat{y}_t \ \forall i$$

Consider $\epsilon_t^v > 0$, and suppose that the solution takes the form

$$\hat{y}_t = c_y^+ \epsilon_t^v$$

$$\pi_t = c_\pi^+ \epsilon_t^v$$

where $c_y^+ < 0$. Substituting these guesses into the system and simplifying yields

$$c_\pi^+ = \frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^+$$

$$\min_i \left\{ -\sigma_i \beta_i^y c_y^+ \epsilon_t^v \right\} = \left(\phi_\pi c_\pi^+ \epsilon_t^v + \phi_y c_y^+ \epsilon_t^v + \epsilon_t^v \right)$$

so that

$$\min_i \left\{ -\sigma_i \beta_i^y c_y^+ \epsilon_t^v \right\} = \left(\phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^+ + \phi_y c_y^+ + 1 \right) \epsilon_t^v$$

By the supposition, $-c_y^+ \epsilon_t^v > 0$ so that

$$-c_y^+ \epsilon_t^v \underline{\beta\sigma} = \left(\phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^+ + \phi_y c_y^+ + 1 \right) \epsilon_t^v$$

where

$$\underline{\beta\sigma} = \min_i \left\{ \sigma_i \beta_i^y \right\}$$

Hence

$$c_y^+ = -\frac{1}{\underline{\beta\sigma} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p} \varphi + \phi_y}$$

$$c_\pi^+ = -\frac{1}{\underline{\beta\sigma} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p} \varphi + \phi_y} \frac{\bar{\Phi} - 1}{\xi^p} \varphi$$

where $c_y^+ < 0$ since

$$\underline{\beta\sigma} + \phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi + \phi_y > 0$$

For completeness, suppose $c_y^+ > 0$. The same steps yield

$$\min_i \left\{ -\sigma_i \beta_i^y c_y^+ \epsilon_t^v \right\} = \left(\phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^+ + \phi_y c_y^+ + 1 \right) \epsilon_t^v$$

By supposition, $c_y^+ \epsilon_t^v > 0$ so that

$$-c_y^+ \epsilon_t^v \bar{\beta\sigma} = \left(\phi_\pi \frac{\bar{\Phi} - 1}{\xi^p} \varphi c_y^+ + \phi_y c_y^+ + 1 \right) \epsilon_t^v$$

where

$$\bar{\beta\sigma} = \max_i \{\sigma_i \beta_i^y\} > 0$$

Hence

$$c_y^+ = -\frac{1}{\bar{\beta\sigma} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y} < 0$$

where the inequality follows from

$$\bar{\beta\sigma} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y > 0$$

Therefore, we have a contradiction.

Now consider $\epsilon_t^v < 0$, and suppose that the solution takes the form

$$\hat{y}_t = c_y^- \epsilon_t^v$$

$$\pi_t = c_\pi^- \epsilon_t^v$$

where $c_y^- < 0$. Substituting these guesses into the system and simplifying yields

$$c_\pi^- = \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- > 0$$

$$\min_i \left\{-\sigma_i \beta_i^y c_y^- \epsilon_t^v\right\} = \left(\phi_\pi c_\pi^- \epsilon_t^v + \phi_y c_y^- \epsilon_t^v + \epsilon_t^v\right)$$

so that

$$\min_i \left\{-\sigma_i \beta_i^y c_y^- \epsilon_t^v\right\} = \left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- + \phi_y c_y^- + 1\right)\epsilon_t^v$$

By the supposition, $c_y^- \epsilon_t^v > 0$ so that

$$-c_y^- \epsilon_t^v \bar{\beta\sigma} = \left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- + \phi_y c_y^- + 1\right)\epsilon_t^v$$

where

$$\bar{\beta\sigma} = \max_i \{\sigma_i \beta_i^y\}$$

Hence

$$c_y^- = -\frac{1}{\bar{\beta\sigma} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y} < 0$$

$$c_\pi^+ = -\frac{1}{\bar{\beta\sigma} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y} \frac{\bar{\Phi}-1}{\xi^p}\varphi_\Theta < 0$$

as required.

For completeness, suppose that $c_y^- > 0$ so that $-c_y^- \epsilon_t^v > 0$ and the same steps as above lead to

$$-c_y^- \epsilon_t^v \underline{\beta\sigma} = \left(\phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi c_y^- + \phi_y c_y^- + 1\right)\epsilon_t^v$$

so that

$$c_y^- = -\frac{1}{\underline{\beta\sigma} + \phi_\pi \frac{\bar{\Phi}-1}{\xi^p}\varphi + \phi_y} < 0$$

which is a contradiction. □

Proof of Proposition 2 Recall the system

$$\iota_t = \rho + \phi_\pi \pi_t + \phi_y \hat{y}_t + v_t$$

$$\pi_t = \kappa_y \hat{y}_t + \delta \mathbb{E}_t \left[\pi_{t+1}\right]$$

$$\min_i \left\{\mathbb{E}_t \left[\hat{c}_{i,t+1}\right] - \hat{c}_{i,t}\right\} = \frac{1}{\sigma}\left(\iota_t - \mathbb{E}_t \left[\pi_{t+1}\right] - \rho\right)$$

$$\hat{c}_{i,t} = \beta_i^{c,v} \hat{y}_t + u_{i,t} \; \forall i$$

Consider $v_1 > 0$, and guess linear solution $\hat{y}_t = c_y^+ v_t$, $\pi_t = c_\pi^+ v_t$ with $c_y^+ < 0$. Substitution yields

$$c_\pi^+ = \frac{\kappa_y}{1 - \delta\rho_v} c_y^+$$

$$\min_i \left\{ \left(\beta_i^{c,v} \rho_v - \beta_{i,t}^{c,v} \right) c_y^+ v_t \right\} = \frac{1}{\sigma} \left(\left(\phi_\pi - \rho_v \right) \frac{\kappa_y}{1 - \delta \rho_v} c_y^+ + \phi_y c_y^+ + 1 \right) v_t$$

$-c_y^+ v_t > 0$ implies

$$c_y^+ = -\frac{1}{\left(1 - \rho_v\right) \underline{\beta}^{c,v} + \frac{1}{\sigma} \left(\phi_\pi - \rho_v \right) \frac{\kappa_y}{1 - \delta \rho_v} + \frac{1}{\sigma} \phi_y} \frac{1}{\sigma}$$

where

$$\underline{\beta}^{c,v} = \min_i \left\{ \beta_i^{c,v} \right\}$$

Suppose $v_1 < 0$, and guess linear solution $\hat{y}_t = c_y^- v_t$, $\pi_t = c_\pi^- v_t$ with $c_y^- < 0$. Substitution yields

$$\min_i \left\{ \left(\rho_v - 1 \right) \beta_i^{c,v} c_y^- v_t \right\} = \frac{1}{\sigma} \left(\left(\phi_\pi - \rho_v \right) \frac{\kappa_y}{1 - \delta \rho_v} c_y^- + \phi_y c_y^- + 1 \right) v_t$$

where $\left(1 - \rho_v\right) c_y^- v_t > 0$ implies

$$c_y^- = -\frac{1}{\left(1 - \rho_v\right) \bar{\beta}^{c,v} + \frac{1}{\sigma} \left(\phi_\pi - \rho_v \right) \frac{\kappa_y}{1 - \delta \rho_v} + \frac{1}{\sigma} \phi_y} \frac{1}{\sigma}$$

where

$$\bar{\beta}^{c,v} = \max_i \left\{ \beta_i^{c,v} \right\}$$

as required. □

Proof of Proposition 3 Pooled OLS estimation for group g yields

$$\hat{\beta}_g = \frac{\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y} \right) \Delta \log c_{i,t}}{\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y} \right)^2}$$

where $\bar{y} = \frac{1}{T}\sum \Delta \hat{\log} Y_t$, and the summation over i is read as "sum over all households i such that $\mathcal{G}(i,t) = g$". Substituting in the true model for household consumption yields

$$\hat{\beta}_g = \frac{\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)\left(\alpha_i + \beta_i \Delta \hat{\log} Y_t + u_{i,t}\right)}{\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)^2}$$

$$\hat{\beta}_g = \frac{\sum_t \sum_i \left(\alpha_i \Delta \hat{\log} Y_t + \beta_i \left(\Delta \hat{\log} Y_t\right)^2 + u_{i,t}\Delta \hat{\log} Y_t - \bar{y}\alpha_i - \beta_i \bar{y}\Delta \hat{\log} Y_t - \bar{y}u_{i,t}\right)}{\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)^2}$$

$$\begin{aligned}\hat{\beta}_g &= \frac{\sum_t \Delta \hat{\log} Y_t \sum_i \alpha_i + \sum_t \left(\Delta \hat{\log} Y_t\right)^2 \sum_i \beta_i + \sum_t \Delta \hat{\log} Y_t \sum_i u_{i,t}}{\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)^2} \\ &\quad - \frac{\bar{y}\sum_t \sum_i \alpha_i + \bar{y}\sum_t \Delta \hat{\log} Y_t \sum_i \beta_i + \bar{y}\sum_t \sum_i u_{i,t}}{\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)^2}\end{aligned}$$

$$\begin{aligned}\hat{\beta}_g &= \frac{\sum_t \Delta \hat{\log} Y_t \alpha_{g,t} + \sum_t \left(\Delta \hat{\log} Y_t\right)^2 \beta_{g,t} + \sum_t \Delta \hat{\log} Y_t u_{g,t}}{\frac{1}{n}\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)^2} \\ &\quad - \frac{\bar{y}\sum_t \alpha_{g,t} + \bar{y}\sum_t \Delta \hat{\log} Y_t \beta_{g,t} + \bar{y}\sum_t u_{g,t}}{\frac{1}{n}\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)^2}\end{aligned}$$

where

$$\alpha_{g,t} = \frac{1}{n_{g,t}} \sum_{i \in g,t} \alpha_i$$

$$\beta_{g,t} = \frac{1}{n_{g,t}} \sum_{i \in g,t} \beta_i$$

$$u_{g,t} = \frac{1}{n_{g,t}} \sum_{i \in g,t} u_{i,t}$$

are parameter means over households in group g in period t. Note that $\beta_{g,t} \in (\min_i \beta_i, \max_i \beta_i)$ by definition.

Continuing,

$$\hat{\beta}_g = \frac{\frac{1}{T}\sum_t \left(\Delta \hat{\log} Y_t\right)\left(\Delta \hat{\log} Y_t\right)\beta_{g,t} - \bar{y}\frac{1}{T}\sum_t \Delta \hat{\log} Y_t \beta_{g,t} + \frac{1}{T}\sum_t \Delta \hat{\log} Y_t \alpha_{g,t}}{\frac{1}{n_{g,t}}\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)^2}$$

$$- \frac{\bar{y}\frac{1}{T}\sum_t \alpha_{g,t} - \frac{1}{T}\sum_t \Delta \hat{\log} Y_t u_{g,t} + \bar{y}\frac{1}{T}\sum_t u_{g,t}}{\frac{1}{n_{g,t}}\sum_t \sum_i \left(\Delta \hat{\log} Y_t - \bar{y}\right)^2}$$

Hence as $T \to \infty$, we can apply a suitable Law of Large Numbers (e.g. Proposition 7.5 in Hamilton (1994)) to obtain

$$\hat{\beta}_g \to^p \frac{Cov\left(\Delta \hat{\log} Y_t, \Delta \hat{\log} Y_t \beta_{g,t}\right) + Cov\left(\Delta \hat{\log} Y_t, \alpha_{g,t} + u_{g,t}\right)}{V\left[\Delta \hat{\log} Y_t\right]}$$

where, if $\mathcal{G}$ does not alter group assignments over time (i.e. is exogenous to changes in $\Delta \hat{\log} Y_t$), then

$$Cov\left(\Delta \hat{\log} Y_t, \alpha_{g,t} + u_{g,t}\right) = 0$$

and

$$Cov\left(\Delta \hat{\log} Y_t, \Delta \hat{\log} Y_t \beta_{g,t}\right) = V\left[\Delta \hat{\log} Y_t\right] \frac{1}{n}\sum_{i \in g} \beta_i$$

so that

$$\hat{\beta}_g \to^p \frac{1}{n}\sum_{i \in g} \beta_i \in \left(\min_i \left\{\beta_i\right\}, \min_i \left\{\beta_i\right\}\right)$$

as required. □

Chapter 2

Employment Protection Legislation and Efficiency in a Frictional Labor Market: Evidence from Brazil

2.1 Introduction

Employment Protection Legislation (EPL) is a pervasive feature of many modern labor markets. However, there exists no consensus on the general equilibrium effects of such legislation, in particular, on the effects of EPL on unemployment and aggregate output. This lack of consensus is attributable to 2 issues. First, the existing reduced-form evidence linking EPL and the unemployment rate finds both positive and negative relationships, and is also silent on the structural mechanisms that mediate the mapping from EPL to the unemployment rate. Second, the broad and often non-monetary nature of EPL makes it difficult to measure holistically, which precludes getting a handle on its total equilibrium impact.[1]

[1]For example, in our country of interest, Brazil, EPL consists of both monetary costs to firms, and non-monetary costs in the form of legal recourse and termination notice periods.

In this paper, we analyze the case of EPL in Brazil, and develop techniques that allow us to tackle both issues. First, we exploit the tenure-dependence built into the EPL in Brazilian labor markets in order to cleanly identify its effect on the hazard rate of job termination. Intuitively, since EPL only takes effect once a job reaches a tenure of three months, we should observe a spike in the hazard rate just before three months, as firms avoid incurring the fixed cost imposed by EPL by firing many workers just before EPL kicks in. Therefore, the size of the hazard rate spike identifies the behavioral effect of EPL on firms' termination decisions. Importantly, this identification strategy does not require us to measure EPL directly, thus overcoming measurement issues associated with the broad nature of EPL in Brazil.

In order to empirically assess firms' responses to the timing of EPL, we use administrative employer-employee matched data from the *Relação Anual de Informações Sociais* (RAIS), 2002-2007. The large sample size allows us to precisely uncover the tenure-dependence in the hazard rate of job termination, as shown in figure 2.1. There is a sharp increase in firing just before the three month probationary periods ends and EPL becomes active. This firing behavior creates a visible spike in the job termination hazard rate, the size of which is informative about the cost of EPL. Importantly, we document that this spike is robust across different industries and skill levels, and is not driven by temporary contracts, short term labor demand fluctuations, or increased worker turnover.

Given this clean identification of the effect of EPL on the hazard rate of job termination, our second contribution is to build and estimate a structural model that maps EPL into macroeconomic outcomes of interest such as aggregate output and the unemployment rate. Specifically, we extend the model of Moscarini (2005), which combines endogenous job destruction through learning about match quality with a frictional labor market à la Diamond-Mortensen-Pissarides, to allow for tenure-dependent EPL, which we model as a fixed cost that firms must pay if they terminate a worker with at

least three months' tenure. This fixed cost captures the equivalent real cost to firms of the wide variety of EPL regulation that firms face in reality.

Our model estimation strategy mirrors our approach to identifying the impact of EPL on the hazard rate of job termination. In particular, we use the spike in the empirical hazard rate schedule to infer the size of the fixed cost parameter that firms must pay. This strategy of inferring hard-to-measure policy parameters from their identified effects in the data builds on Garicano et al. (2016), and ensures that the estimated model parameter captures the equivalent real cost that EPL imposes on firms. In addition, variation in the empirical hazard rate across different tenures allows us to identify and estimate model parameters governing the bargaining share of workers, and the variance of the idiosyncratic shocks to a worker's production. Finally, we calibrate the remaining parameters in line with the rich literature that studies frictional labor markets.

Using our estimated model, we perform a counterfactual exercise and compare macroeconomic outcomes when we remove EPL from the economy. In our baseline setting, we find that EPL lowers aggregate output by 1.3%. The drop in output is driven by a violation of the Hosios (1990) efficiency condition that applies in our economy. In particular, our estimated parameters imply that firms do not create enough vacancies in equilibrium so that the job finding rate is inefficiently low relative to the solution to the planning problem. The presence of EPL then exacerbates this inefficiency since it weakens firms' incentives to create new vacancies even more. Hence EPL lowers aggregate output.

Our results highlight a potential unintended consequence of EPL. While EPL creates a system of benefits that workers receive when they enter unemployment, the transfer of resources from firms to workers upon which EPL is built lowers the value that firms ascribe to employment matches in the first place. The reduction in value that firms receive as a result of EPL makes entering into new employment contracts less

attractive, which leads them to decrease their rates of vacancy creation, and hence causes an increase in steady state unemployment. This behavioral response by firms turns out to be far more quantitatively important for macroeconomic outcomes, than the readily observable effect of EPL on firms' termination decisions.

Related Literature Our paper relates to two distinct literatures. First, our clean identification of the impact of EPL on firms' job termination decisions contributes to an extensive literature that explores the impact of EPL on labor market and macroeconomic outcomes (Daruich et al., 2017; Di Tella and MacCulloch, 2005; Garibaldi and Violante, 2005; Kugler and Pica, 2008; Lazear, 1990). These papers generally rely on two identification strategies. The first uses cross-country variation in EPL to estimate the causal impact of EPL on unemployment, while the other uses within-country variation, for example, by exploiting EPL reforms or size-contingent laws to estimate similar causal effects. Relative to these methods, we also exploit contingencies in the design of EPL in Brazil: the discrete activation of EPL for jobs with at least three months' tenure allows us to cleanly identify its affect on the hazard rate of job termination. Furthermore, by estimating the behavioral response of firms at the microeconomic level, our estimation strategy does not require such strong identification assumptions as is the case when estimating the causal effect on macroeconomic outcomes. Instead, we build and estimate a structural model that makes this mapping explicit.

Our structural estimation exercise relates to the literature that studies the effects of EPL on macroeconomic outcomes through the lens of a search model (Garibaldi and Violante, 2005; Ljungqvist, 2002). In these papers, EPL is modeled as a fixed cost that applies to jobs of all tenures, thus missing the tenure-dependence that is a key feature of many labor markets, including Brazil's (Cahuc et al., 2016). Our contribution, therefore, is to capture the tenure-dependence of EPL within a structural model, and

to use our rich data to estimate key model parameters rather than relying solely on calibration.

The paper proceeds as follows: section 2.2 introduces our data and institutional setting within the Brazilian labor market. In section 2.3, we document the hazard rate spike created by the EPL, and show that it is robust feature of a wide variety of labor submarkets. We develop our theoretical framework in section 2.4, and describe the estimation and calibration procedure in section 2.5. We present our main quantitative exercise together with robustness checks in sections 2.6 and 2.7. Finally, section 2.8 concludes.

2.2 Data and Institutional Setting

2.2.1 Data

Our analysis utilizes administrative data from the *Relação Anual de Informações Sociais* (RAIS), years 2002-2007. The RAIS data contains linked employer-employee records from a mandatory survey administered by the Brazilian Ministry of Labor and Employment (MTE). Fines are levied on firms which provide inaccurate or incomplete information on the survey.

Each entry in the RAIS dataset is a employee-employer match. Each individual, firm, and establishment are assigned unique administrative identifiers which do not change over time. Importantly, the data track each the tenure of each employer-employee match (job). For our analysis, we bin tenure into 15 day intervals due to "heaping" in the distribution of tenure (i.e. it is much more likely to observe a 30 days job spell than a 29 day job spell). The data include additional information about the job, such as occupation, wage, hours, type of labor contract, whether the job has ended, and why the job has ended, and also contain demographic data on individuals, such as

education, gender, ethnicity, and occupation. For more information about the dataset and the definition of variables, see section 2.D in the appendix.

2.2.2 Institutional Details

In Brazil, EPL is composed of many components. For example, formal sector workers in Brazil are guaranteed severance pay if dismissed without cause, yearly bonuses equivalent to one month's salary, and 30 days' notice for any separation.[2] Furthermore, in the event of a separation, the employer firm must pay a "firing penalty", which is equal to roughly one month of the worker's salary for every year the worker has been employed at the firm.

The key feature of EPL in Brazil that facilitates our analysis is its tenure-dependence: all dimensions of EPL only apply to firms and workers that have been in an employer-employee match for at least three months. This sharp discontinuity in the cost of EPL as a function of tenure is the basis of our identification strategy. Intuitively, the jump in EPL costs at three months incentivizes firms to terminate matches just before this tenure is reached to avoid incurring the higher costs of termination should the match deteriorate soon after three months. To show that firms do indeed respond to this discontinuity in the cost of EPL, figure 2.1 plots the job termination hazard rate as a function of tenure. There is a sharp increase in firing just before the three month probationary periods ends and EPL becomes active. This firing behavior creates a visible spike in the job termination hazard rate, which is informative about the cost of EPL. Intuitively, if the cost of EPL is higher, firms will be more cautious about hiring a worker beyond three months, resulting in a larger spike in the job termination hazard at three months.

[2]In the event that a worker is given advance notice of a termination, the worker must be allotted time to search for a new job.

2.2.3 Sample Selection

Our identification strategy hinges on the spike in job terminations at three months' tenure being solely driven by the timing of EPL. A natural confounder is therefore the presence of workers on temporary three month contracts. In Brazil, temporary contracts are subject to approval by the Ministry of Labor (MTE) and about 5 percent of workers at a given time are employed under such contracts. These contracts are approved to meet temporary increases in demand and many of these contracts last for three months. Therefore, a spike in the job termination hazard may naturally arise at three months due to the existence of such contracts.[3] Given the focus of our paper is on the effects of EPL on permanent employment contracts, we therefore eliminate temporary contracts from the majority of our empirical and theoretical analysis.[4]

In addition to eliminating temporary contracts, we restrict attention to workers aged 18-65, and working in full-time jobs (at least 35 hours per week). We exclude individuals with invalid identifiers (less than one percent of the data). Column 1 of table 2.1 presents summary statistics for the population of 18-65 year olds. Column 2 presents summary statistics for jobs which last less than or equal to three months. Short-duration workers are slightly younger (30.4 vs. 31.5), less likely to be a college graduate (3.8 percent vs. 7.2 percent), are paid lower monthly wages (670.90 Real vs 819.21 Real), and are less likely to be in public administration jobs (7.4 percent vs. 2.2 percent) and more likely to be in agricultural jobs (14.0 percent vs. 9.3 percent). In total, there are 92,023,307 jobs corresponding to 29,438,306 unique workers. 24,427,409 jobs last three months or less (i.e. 26.5 percent of all jobs).

[3]While the presence of EPL may also theoretically cause increased substitution towards temporary contracts (Daruich et al., 2017), the regulated usage of temporary contracts likely limits this substitution in Brazil.

[4]Importantly for our study, the RAIS dataset includes information on the type of contract.

	All Jobs	Short Duration Jobs
Panel A: Demographics	(1)	(2)
Age	31.535	30.405
High School Graduate	0.336	0.305
College Graduate	0.072	0.038
Male	0.658	0.693
Panel B: Job Characteristics		
Monthly Wage	819.212	670.909
Tenure	13.332	1.681
Hours	43.108	43.444
Panel C: Firm Characteristics		
Manufacturing	0.199	0.185
Agriculture	0.093	0.140
Public Administration	0.074	0.022
Health and Education	0.039	0.023
All Other Sectors	0.595	0.631
Unique Workers	29,438,306	13,312,346
Number of Jobs	92,023,307	24,427,409

Table 2.1: Descriptive Statistics for Estimation Sample, 2002-2007. Column 1 reports descriptive statistics for jobs between age 18-65 who began a job after after 1986, excluding workers on temporary contracts. Column 2 reports descriptive statistics for jobs which last less than three months. Tenure is measured in months. Wages are denominated in Brazilian Real.

2.3 The Impact of EPL on the Job Termination Hazard

In this section we document a visible spike in the job termination hazard rate at a tenure of three months, and argue that it predominantly reflects the early termination of permanent employment contracts caused by the tenure dependence built into EPL in Brazil.

2.3.1 Estimating Bunching in Job Terminations

To summarize the quantitative magnitude of the spike in the job termination hazard rate at three months tenure, we estimate a "bunching" statistic. This summary statistic is useful as it is comparable across different labor submarkets, and can be used to explore potential driving mechanisms of the hazard rate spike. We follow the public finance literature Kleven (2016) to estimate bunching. Specifically, we fit a flexible polynomial to the empirical job termination hazard, excluding data from around the notch point T, which is the tenure length in which EPL takes affect. Formally, let $B_j = \{15, 30, ...\}$ define the bins and H_j indicate the hazard rate in bin j (i.e. H_{90} denotes the probability a job ends between 75 and 90 days, given the job has lasted for 75 days). To estimate the counterfactual hazard, we estimate the following regression:

$$H_j = \sum_{i=0}^{q} \beta_i \cdot (B_j)^i + \sum_{i=-R}^{R} \gamma_i \cdot \mathbf{1}[B_j = i] + \varepsilon_j^0 \tag{2.1}$$

where q is the order of the polynomial and R denotes the width of the excluded region around the notch in firing costs. In practice, we set $q = 10$ and $R = 15$, and therefore exclude any tenure durations which end between 75 and 105 days in the estimation of the counterfactual hazard rate. We use the results from Equation 2.1 to estimate

the counterfactual hazard as:

$$\hat{H}_i = \sum_{i=0}^{q} \hat{\beta}_i (B_j)^i \tag{2.2}$$

The excess mass is defined as the difference in the true hazard rate and the counterfactual hazard rate at $T = 90$:

$$B = H_T - \hat{H}_T \tag{2.3}$$

The normalized excess mass, which we will refer to as bunching, b, is defined as the excess mass divided by the counterfactual hazard rate at tenure duration T.

$$b = \frac{B}{\hat{H}_T} \tag{2.4}$$

To compute standard errors for the excess mass and bunching, we generate hazards and excess mass by resampling the residuals in equation 2.1. The standard error is then equal to standard deviation of the distribution of excess mass estimates over 500 bootstrap samples.

Figure 2.1 displays the job termination hazard rate. There is a visible spike in the hazard rate at a tenure of three months, which results in significant bunching.[5] We find that the excess mass is equal to 1.4, with a standard error of 0.234, indicating that the true hazard rate is more than double the predicted counterfactual hazard rate at three months' tenure.

The sizable excess mass indicates a non-trivial response by firms, who alter their job termination decisions in the presence of tenure-dependent EPL. An intuitive explanation, and indeed the mechanism we are most interested in, is that the presence

[5]There is another much smaller spike in the hazard rate at around six months' tenure. Van Doornik et al. (2018) shows that this spike is due to "fake" separations. If a worker is fired after six months of tenure, the worker can receive unemployment insurance from the government. This incentivizes firms to fire workers and then split the unemployment insurance. In order to focus on the effects driven by the timing of EPL, and to keep the structural model as parsimonious as possible, we abstract from the effects of EPL beyond the bunching that occurs at three months' tenure.

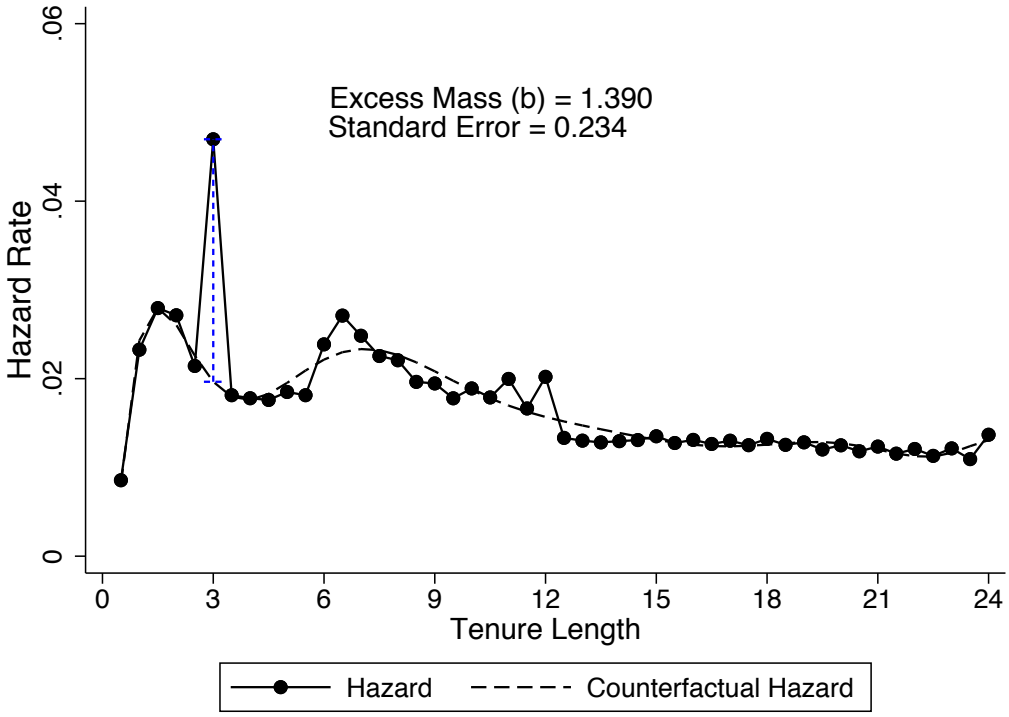


Figure 2.1: This figure plots the layoff hazard rate. Tenure duration is binned into 15 day intervals. The dashed line is a tenth-degree polynomial fitted to the empirical hazard rate, excluding points 15 days away from the notch, as in Equation (2.1). The vertical dotted line displays the excess mass B, while the normalized excess mass b and standard error is reported in the figure. The standard error is computed using a residual bootstrap procedure.

of EPL for jobs with tenures of at least three months incentivizes firms to terminate a significant number of jobs earlier than they otherwise would have. This mechanism is formalized by our structural model in section 2.4, in which the size of EPL exactly pins down the size of the hazard rate spike at three months' tenure.

In order to justify this intuition empirically, we now argue that the observed bunching is indeed predominantly driven by the early termination of permanent employment contracts, as opposed to firms optimally choosing to hire workers on short-term or temporary contracts.

2.3.2 Temporary Contracts

An advantage of our data is that we can drop all matches labeled as temporary contracts to ensure that our bunching analysis is not confounded by mechanical job terminations at three months' tenure.[6]

However, hiring workers on temporary contracts is itself a regulated process in Brazil.[7] Therefore, in order to sidestep these regulatory frictions, firms desiring short term employment arrangements may simply hire workers on permanent contracts with the intention of simply firing them after three months thus mimicking the temporary contract. This behavior of creating artificial temporary contracts is not directly observable in our data, and could confound our bunching estimates. We now examine 2 key reasons that could cause firms to create artificial temporary contracts.

Demand Volatility Firms that face volatile demand for their product will naturally have volatile labor demand that is best served via short term employment

[6]To understand the role of temporary contracts, figure 2.7 plots the job termination hazard rate which includes temporary contracts. As can be seen in the figure, the amount if bunching is now equal to 1.9.

[7]To hire a worker under a temporary contracts, firms must get permission from the Ministry of Labor. They must also establish that the worker needs to be hired on a temporary contract in order to meet seasonal fluctuations in demand.

contracts. Such firms will therefore find it optimal to create artificial temporary contracts when official temporary contracts are unavailable.

One simple way to discern how much of the bunching behavior is driven by this behavior is to compare bunching across different industries, where some industries are naturally more prone to short-term labor hiring than others. Specifically, we estimate bunching separately across industries at the 3-digit level, and then correlate bunching with the month-to-month variation in employment in the given 3-digit industry. If labor demand volatility is a significant driver of bunching, we would expect industries with higher employment volatility to also display greater magnitudes of bunching.

To estimate demand volatility, we compute a normalized measure of monthly employment changes in sector j at time t as:

$$\Delta E_{j,t} = \frac{hires_{jt} - fires_{jt}}{hires_{jt}} \tag{2.5}$$

We then compute volatility of sector j as $V_j = Var(\Delta E_{j,t} - \Delta E_{j,t-1})$. In words, we create time series of month-to-month net employment changes scaled by the total number of hires. We then take the variance of the first difference as our measure of employment volatility. We then correlate this measure with bunching by running the following regression:

$$ln(\hat{V}_j) = \beta_0 + \beta_1 \hat{b}_j \tag{2.6}$$

where β_1 capture the correlation between bunching and employment volatility.

Figure 2.2 shows the results of this regression as well as a scatterplot of bunching across industries. The first thing to notice in the figure is that bunching is positive in every single three-digit industry, indicating that positive bunching is an important feature across many industries. Additionally, it does not appear as if bunching is strongly correlated with demand volatility. Demand volatility is actually *negatively*



Figure 2.2: This figure plots the bunching estimate by sector (defined by three-digit CNAE classification) and correlates the bunching to volatility in employment. The volatility is calculated as the standard deviation of month-to-month employment changes over the course of a year.

correlated with bunching, although the correlation is not significant. This suggests that demand volatility does not play a quantitatively important role in determining the amount of bunching that we estimate.

High Worker Turnover In addition to volatile labor demand creating demand for artificial temporary contracts, the presence of EPL itself could incentivize firms to create such contracts. In particular, if production is unaffected by high worker turnover, then it could be profit-maximizing for firms to cycle through workers on a three month tenure basis and to hence avoid ever having to pay the cost imposed by EPL. This behavior would create bunching that would confound our estimates.

Intuitively, this channel seems most prevalent for lower skill occupations (constant replacement of engineers, for example, seems very unrealistic). Therefore, in order to examine how much it contributes towards the bunching we observe, we divide occupations into different skill levels, where skill level is defined by the International Standard Classification of Occupations (ISCO). Low-skill occupations are characterized by the performance of simple and routine physical tasks, and includes occupations such as cleaners and construction laborers. Medium-skill jobs involve performing more complex tasks, such as operating machinery, and includes occupations such as office clerks and skilled craftsman. High-skill jobs require workers to perform complex tasks and in many cases, some form of advanced education. High-skill occupations include technicians, managers and professionals.[8].

Given these definitions, we expect the channel to be stronger for low-skill workers, as the tasks they perform require little training. As can be seen in figure 2.3, however, bunching occurs across all skill levels. For example, bunching in both the high-skill and low-skill categories is equal to 1.5. While it is true that job termination is in general higher in low-skill occupations, the excess firing at three months is similar across skill levels.

To show this result is consistent, we estimate bunching across all three-digit occupations. To capture a crude measure of skill-level, we use the average wage in the occupation. In figure 2.4, we plot the estimated bunching against the average log monthly wage. As can be seen in the figure, there is a slight but insignificant negative correlation between average log monthly wage and bunching. This suggests that bunching is not driven by firms rotating through low-skill workers that are relatively easy to replace quickly, but is a feature even in high-skill, high-wage occupations, and hence must reflect the early termination of truly permanent employment contracts.

[8]ISCO also breaks down high-skill occupations into medium-high skill and high-skill. For this paper, we have aggregated these two groups and defined them as high-skilled

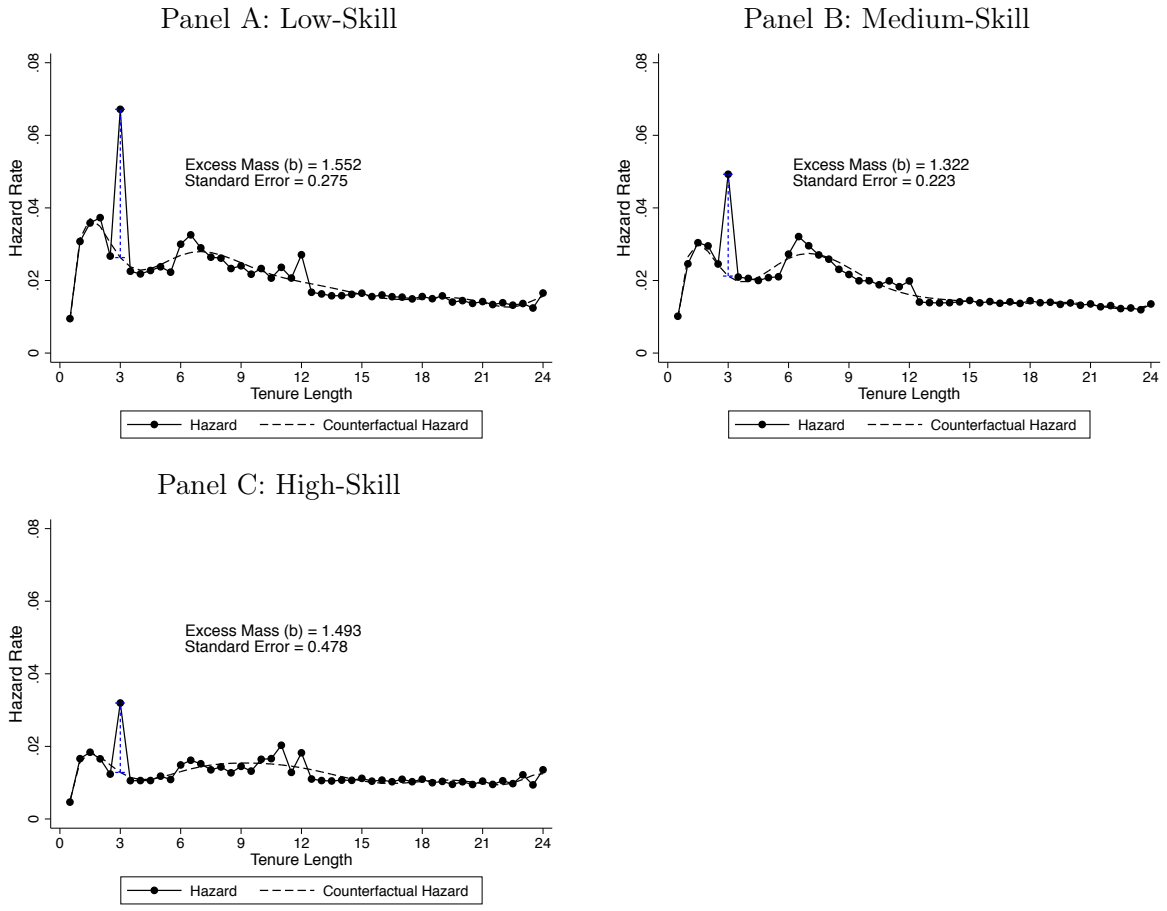


Figure 2.3: This figure plots the layoff hazard rate by different occupation skill levels. Skill level is defined by the International Standard Classification of Occupations (ISCO). Low skill occupations are characterized by the performance of simple and routine physical tasks, and includes occupations such as cleaners and construction laborers. Medium-skill jobs involve performing more complex tasks, such as operating machinery, and includes occupations such as office clerks and skilled craftsman. High-skill jobs require workers to perform complex tasks and requires significant practical knowledge, and is composed of technicians, managers and scientific professionals. Tenure duration is binned into 15 day intervals. The dashed line is a tenth-degree polynomial fitted to the empirical hazard rate, excluding points 15 days away from the notch, as in Equation 2.1. The vertical dotted line displays the excess mass B, while the normalized excess mass b and standard error is reported in the figure. The standard error is computed using a residual bootstrap procedure.

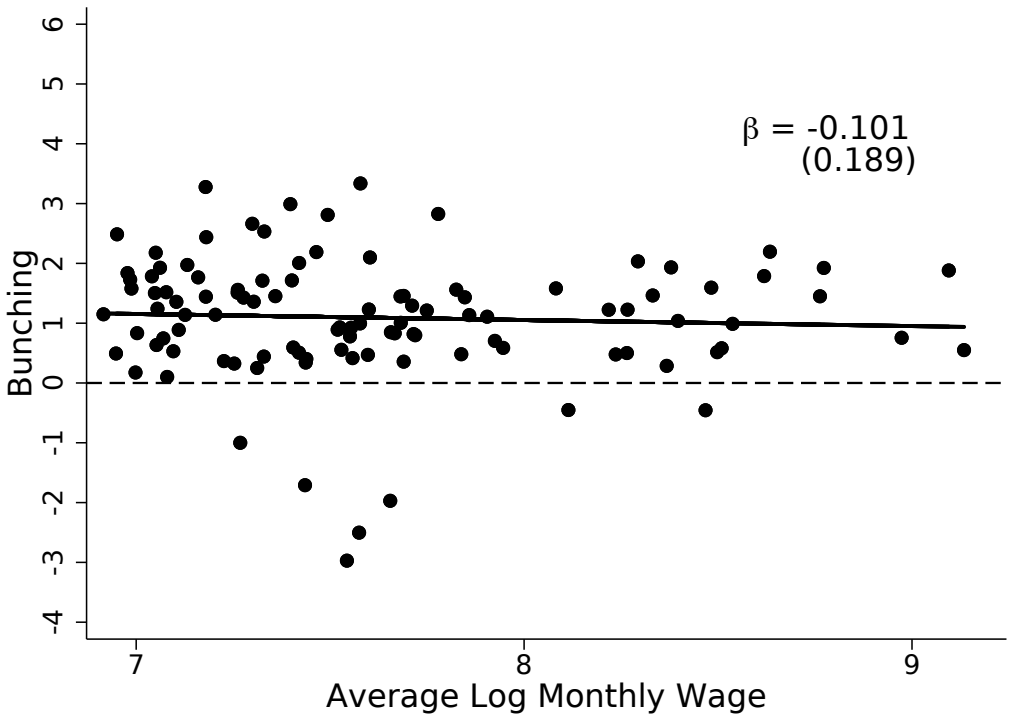


Figure 2.4: This figure plots the bunching estimate by occupation (defined by three-digit ISCO identifier) and correlates the bunching to average wages. The coefficient displayed is the estimate of the coefficient of a regression of the bunching estimate on the variable on the x-axis.

2.3.3 Summary

In this section we have documented a visible spike in the job termination hazard rate at a tenure of three months, and have argued that it predominantly reflects the early termination of permanent employment contracts caused by the tenure dependence built into EPL in Brazil.

While our empirical analysis exploits this tenure dependence allows us to cleanly identify the effect of EPL on the hazard rate of job termination via firms' decisions at the micro level, it cannot say how these decisions aggregate up and affect macroeconomic outcomes such as unemployment and aggregate output. In order to address these questions, we now develop a structural framework that formalizes the mapping between EPL, hazard rates, and general equilibrium macroeconomic outcomes.

2.4 Model

We wish to study the consequences of EPL for macroeconomic outcomes such as unemployment and GDP. In order to accomplish this task, a structural model must include equilibrium unemployment, and endogenous job turnover that responds to the incentives to terminate jobs early created by EPL. To this end, we introduce EPL into the general equilibrium model of Moscarini (2005), which merges the theories of job turnover and unemployment in a tractable manner. As such, our model exposition closely follows his, except where we emphasize the new role played by EPL.

A final good is produced in continuous time by pairwise firm-worker matches. The match-specific productivity of a match, μ, is ex-ante uncertain, and evolves stochastically over the life of the match. Upon forming the match, both firm and worker share a common prior on μ that is independent of their histories. The prior puts mass on

two points $\left\{\mu^L, \mu^H\right\}$ where $\mu^L < \mu^H$, and μ^L denotes a bad match, μ^H a good match. Let $p_0 = \Pr\left(\mu = \mu^H\right)$ be the initial prior that the match is good.

Final good production is linear in match productivity, so that in a small interval dt, production of the good X is given by

$$dX_t = \mu dt + \sigma dZ_t$$

where dZ_t is a standard Brownian motion. In other words, output is a noisy indicator of true match quality, with the noise being scaled by $\sigma > 0$. The presence of noise creates an inference problem that firms and workers solve by using the information provided by the history of output, denoted by the filtration $\mathcal{F}_t^X$, to update their prior belief in a Bayesian manner,

$$p_t = \Pr\left(\mu = \mu^H | \mathcal{F}_t^X\right)$$

When the belief is low enough, firms and workers will optimally choose to terminate the match.[9] However, upon termination, firms must pay a firing cost $\kappa(t)$ that depends on the tenure t of the match in the following way:

$$\kappa(t) = \begin{cases} 0 & \text{if} \quad t < T_1 \\ \\ \kappa & \text{if} \quad t \geq T_1 \end{cases}$$

where $\kappa > 0$ is a constant, and T_1 is a tenure after which the firing cost becomes active. This firing cost represents the effect of EPL on firms in the model economy.

[9] As noted by Moscarini (2005), similar job termination dynamics could be achieved using observable productivity shocks to a worker's output. However, the learning formulation is both tractable, and generates an average wage-tenure relationship that is increasing and concave, as in the data.

Ultimately, we will estimate κ to match the spike in the hazard rate of job termination so that we can interpret it as the effective real cost of the EPL regulation imposed on firms in reality.

In addition to this endogenous separation, matches are also subject to exogenous termination at rate $\delta > 0$.

There is a large mass of ex-ante homogeneous firms, who can post vacancies at flow cost $c > 0$ when unmatched (the large mass ensures free-entry in equilibrium). We normalize the mass of workers to unity, and assume that unemployed workers receive a flow value of leisure b. All agents are risk-neutral, and discount future payoffs at rate $r > 0$. In order to split the surplus generated by a match, we assume that the firm and worker use a generalized Nash bargaining rule, and denote the bargaining weight of workers by $\beta \in (0, 1)$.

We study the steady state equilibrium of this economy. As such, aggregate variables do not have a time subscript, and we use $t \geq 0$ to unambiguously denote tenure for individual level variables from now on.

2.4.1 Belief Dynamics

The solution to the continuous time inference problem for p_t has a well known solution in the form of the stochastic differential equation

$$dp_t = p_t \left(1 - p_t\right) s d\bar{Z}_t$$

where $s = \frac{\mu^H - \mu^L}{\sigma}$ is the signal-to-noise ratio and

$$d\bar{Z}_t = \frac{1}{\sigma}\left(dX_t - \left(p_t \mu^H + \left(1 - p_t\right)\mu^L\right) dt\right)$$

is a standard Brownian motion with respect to the filtration generated by X_t, $\mathcal{F}_t^X$.[10] Intuitively, beliefs move faster the more uncertain is the current belief ($p(1-p)$ has a maximum at $p = \frac{1}{2}$). In addition, it is useful to define

$$\Sigma(p) = \frac{1}{2} s^2 p^2 (1-p)^2$$

which is interpreted as half the variance of the change in beliefs over a small change in tenure.

2.4.2 Value Functions

Define $\bar{\mu}(p) = p\mu^H + (1-p)\mu^L$ as the expected productivity of a match when the belief is p, and let V be the value to a vacant firm of opening a vacancy. In equilibrium, free entry into the vacancy posting market ensures that $V = 0$.

Given the presence of firing costs, it is useful to divide the analysis into two stages of tenure: $t \in [0, T_1)$, and $t \geq T_1$. Throughout, let superscripts on functions and variables $i \in \{1, 2\}$ refer to each respective tenure stage, and let subscripts denote derivatives. E.g. J^i_{xx} is the second derivative of the function J^i with respect to the argument x.

$t \geq T_1$:

We proceed by backward induction, and start in stage two of a job's tenure. Since the firing cost is fixed, firm and worker value functions only depend the current belief, which is therefore the natural state variable.

[10] For a discussion of this result, see Moscarini (2005) and the references therein.

The value to a firm of a match with current belief p must satisfy the Hamilton-Jacobi-Bellman (HJB) equation

$$rJ^2(p) = \bar{\mu}(p) - w^2(p) + \Sigma(p) J^2_{pp}(p) - \delta J^2(p)$$

where w^2 is the wage of a match with belief p.

The flow value of the match to the firm equals the flow profits (production minus the wage) plus capital gains stemming from the change in beliefs and the possibility of exogenous separation.

As discussed, when the belief of the match reaches some threshold $\underline{p}^{2,J}$, the firm will optimally choose to terminate the match. Optimality requires that the firm's choice of threshold satisfy the two boundary conditions

$$J^2\left(\underline{p}^{2,J}\right) = -\kappa$$

$$J^2_p\left(\underline{p}^{2,J}\right) = 0$$

The first condition states that, at the termination threshold $\underline{p}^{2,J}$, the value of the match to the firm equals the value of termination, which is equal to the value of entering the vacancy posting stage (zero in equilibrium) minus the firing cost that the firm must pay to terminate the match in the second stage of tenure.[11]

The second condition states that, at the termination threshold, the slope of the value function must be zero. Intuitively, this condition ensures that the firm is indifferent between terminating the match as soon as the belief hits the threshold, and waiting for a small amount of time to see what happens.[12]

[11]This condition is often referred to as the "value matching" condition in the optimal stopping literature.

[12]This condition is referred to as the "smooth pasting" condition.

The value to an employed worker of a match with current belief p satisfies the HJB equation

$$rW^2(p) = w^2(p) + \Sigma(p) W^2_{pp}(p) - \delta\left(W^2(p) - U\right)$$

The value of the match to the worker equals the flow benefit (the wage) plus capital gains stemming from the change in beliefs and the possibility of exogenous separation, where U is the value to the worker of entering unemployment, and is defined shortly.

Similarly to the firm, when the belief of the match reaches some threshold $\underline{p}^{2,W}$, the worker will optimally choose to terminate the match. Optimality requires that the worker's threshold choice satisfy the two boundary conditions

$$W^2\left(\underline{p}^{2,W}\right) = U$$

$$W^2_p\left(\underline{p}^{2,W}\right) = 0$$

The first condition states that value of the match at the threshold is equal to the value to the worker from termination, which is equal to the value of unemployment. The second condition states that, at the termination threshold, the slope of the value function must be zero so that the worker is indifferent between terminating the match as soon as the belief hits the threshold, and waiting for a small amount of time to see what happens.

$t \in [0, T_1)$**:**

For matches of tenure $t \in [0, T_1)$, the value functions depends on both the belief p and tenure t since the time until firing costs become non-zero changes with tenure.

The value to a firm of a match with current belief p and tenure t must satisfy the HJB equation

$$rJ^1(p,t) = \bar{\mu}(p) - w^1(p,t) + J_t^1(p,t) + \Sigma(p) J_{pp}^1(p,t) - \delta J^1(p,t)$$

where w^1 is the wage of a match with belief p and tenure t. This has the same interpretation as the stage two HJB equation, except that the capital gains now includes the change in value due to the increase in tenure and hence the reduction in time until the firing cost becomes active.

As in stage two, when the belief reaches a tenure threshold, the firm will choose to terminate the match. The key difference is that the thresholds now depend on tenure, and so satisfy the boundary conditions

$$J^1\left(\underline{p}^{1,J}(t), t\right) = 0$$

$$J_p^1\left(\underline{p}^{1,J}(t), t\right) = 0$$

The first condition ensures that at the threshold $\underline{p}^{1,J}$ at tenure t, the value to the firm of the match equals the value of entering the vacancy posting stage, which is now zero since there is no firing cost to pay. The second condition ensures that the threshold is optimal by a similar logic to the stage two case.

In addition to these conditions, there is another boundary condition in the tenure dimension that pins down the function $J^1(p, T_1)$,

$$J^1(p, T_1) = J^{*,2}(p) \;\; \forall p \geq \underline{p}^{1,J}(T_1)$$

where $J^{*,2}$ is determined in equilibrium. Importantly, $J^{*,2} \neq J^2$, since the firm has the option to terminate the match an instant before $t = T_1$ as zero cost, while waiting

until $t = T_1$ would incur a positive firing cost when $\kappa > 0$. We will later show that this jump is firing costs at $t = T_1$ results in a discontinuity in the wage function and hence a discontinuity in the value function for the firm.

The value to an employed worker of a match with current belief p and tenure t must satisfy the HJB equation

$$rW^1(p,t) = w^1(p,t) + W_t^1(p,t) + \Sigma(p) W_{pp}^1(p,t) - \delta\left(W^1(p,t) - U\right)$$

As in the firm case, in stage one, worker belief thresholds now depend on tenure, and satisfy the boundary conditions

$$W^1\left(\underline{p}^{1,W}(t), t\right) = U$$

$$W_p^1\left(\underline{p}^{1,W}(t), t\right) = 0$$

As in the firm case, $W^1(p, T_1)$ is pinned down by the additional boundary condition

$$W^1(p, T_1) = W^{*,2}(p) \;\; \forall p \geq \underline{p}^{1,W}(T_1)$$

where $W^{*,2}$ is determined in equilibrium. Again, $W^{*,2} \neq W^2$ for the same reasons as in the firm case.

Finally, let U be the value of unemployment, which satisfies the HJB equation

$$rU = b + \lambda\left(W^1(p_0, 0) - U\right)$$

The value of being unemployed equals the flow value from leisure plus the gain from becoming employed times the probability of entering a match, given by the job finding rate λ.

2.4.3 Wage Determination via Nash Bargaining

It is again useful to proceed by backward induction.

$t \geq T_1$**:**

Given the firm and worker value functions, we can define the surplus of a match with current belief p and tenure $t \geq T_1$ as the sum of the values of the match to the firm and worker minus their respective outside options,

$$S^2(p,t) = J^2(p) + \kappa + W^2(p) - U$$

The generalized Nash bargaining protocol selects a wage according to

$$w^2(p) = \arg\max_w \left(W^2(p) - U\right)^\beta \left(J^2(p) + \kappa\right)^{1-\beta}$$

which has the FOC

$$\beta\left(J^2(p) + \kappa\right) = (1-\beta)\left(W^2(p) - U\right)$$

Note that this condition together with the first boundary conditions of the firm and worker HJB equations imply that

$$\underline{p}^{2,J} = \underline{p}^{2,W} = \underline{p}^2$$

so that optimal firm and worker belief thresholds coincide at all tenures in stage two. Using the expressions for the value functions to solve for the wage function yields

$$w^2(p) = (1-\beta)\, b + \beta\left(\bar{\mu}(p) + \lambda J^1(p_0, 0) + (r+\delta)\,\kappa\right)$$

which shows how the worker's wage is a weighted average of her outside option (leisure during unemployment) and her inside option (the expected surplus flow plus the continuation value from a match). Importantly, we see that the firing cost enters positively. Intuitively, the worker can exploit the fact that the firm must pay a firing cost today in order to terminate the match, to increase her share of the surplus.

$t \in [0, T_1)$**:**

Similarly to stage two, we can define the surplus of a match with current belief p and tenure t as the value of the match to the firm and worker minus their respective outside options,

$$S^1(p,t) = J^1(p,t) + W^1(p,t) - U$$

where the firing cost is now zero. The Nash bargaining solution is hence given by

$$w^1(p,t) = \arg\max_w \left(W^1(p,t) - U\right)^\beta \left(J^1(p,t)\right)^{1-\beta}$$

which has the FOC

$$\beta\left(J^1(p,t)\right) = (1-\beta)\left(W^1(p,t) - U\right)$$

Note that this condition together with the first boundary conditions of the firm and worker HJB equations imply that

$$\underline{p}^{1,J}(t) = \underline{p}^{1,W}(t) = \underline{p}^1(t)$$

so that optimal firm and worker belief thresholds coincide at all tenures. Using the expressions for the value functions to solve for the wage function yields

$$w^1(p,t) = (1-\beta)\, b + \beta\left(\bar{\mu}(p) + \lambda J^1(p_0, 0)\right)$$

which has a similar interpretation to the stage two tenure case.

2.4.4 Discontinuities at T_1

Recall that we have postulated a discontinuity in the value functions at $t = T_1$, where the firing cost jumps from zero to a positive number. This discontinuity stems from the change in wages as tenure approaches T_1 from below.

For $t < T_1$, the wage is given by

$$w^1(p,t) = (1-\beta)\, b + \beta\left(\bar{\mu}(p) + \lambda J^1(p_0, 0)\right)$$

while for $t \geq T_1$, the wage is given by

$$w^2(p,t) = (1-\beta)\, b + \beta\left(\bar{\mu}(p) + \lambda J^1(p_0, 0) + (r+\delta)\,\kappa\right)$$

Since $\kappa(t)$ jumps at $t = T_1$, the wage will also jump. In particular, holding the belief fixed, as tenure crosses T_1, the wage jumps by an amount

$$\omega^1 = \beta(r+\delta)\,\kappa > 0$$

This jump reflects the fact that for tenures $t \geq T_1$, the worker can increase her share of the surplus by exploiting the fact that the firm must pay a firing cost in order to terminate the match.

Using this logic, we can now derive the functions $J^{*,2}$ and $W^{*,2}$ as the values of a match with belief p at tenure T_1 to firms and workers were the wage not to jump by ω^1 due to the firing cost.

Proposition 11. *The functions $J^{*,2}$ and $W^{*,2}$ are given by*

$$J^{*,2}(p) = J^2(p) + \beta\kappa$$

$$W^{*,2}(p) = W^2(p) - \beta\kappa$$

Intuitively, we see that the jump size is increasing in the bargaining share of workers, since a larger share implies that workers can extract more of the surplus by exploiting the firing cost that firms must pay.

2.4.5 Hazard Rate Spike at Tenure $t = T_1$

The key feature of our model is that firms must pay a firing cost in order to terminate a match that has tenure of at least T_1. The following logic formalizes how this firing cost creates a spike in the hazard rate of match termination at T_1, thus providing an explicit micro-foundation for the link between EPL and the hazard rate spike that we documented in the data.

Define the firm-specific thresholds $\underline{p}^{1,J}(T_1)$ and $\bar{p}^{1,J}(T_1)$, that satisfy the conditions

$$J^1\left(\underline{p}^{1,J}(T_1), T_1\right) = 0$$

$$J^2\left(\bar{p}^{1,J}(T_1)\right) = J^1\left(\bar{p}^{1,J}(T_1), T_1\right) - \beta\kappa = 0$$

where $\bar{p}^{1,J}(T_1) > \underline{p}^{1,J}(T_1)$ since firm value functions are increasing in the current belief. Given these threshold definitions, any match with belief $p \in \left[\underline{p}^{1,J}(T_1), \bar{p}^{1,J}(T_1)\right]$ at tenure $t \to T_1$ will be immediately terminated since the value of such a match will

instantaneously drop to $J^2(p, T_1) < 0$. This extra termination creates a spike in the hazard rate at T_1.

Intuitively, when tenure reaches T_1, there is an interval of beliefs such that a match with a belief in that interval would not be terminated were it not for the presence of a firing cost. While such matches are reasonably productive, they are not productive enough to warrant the firm continuing the match and paying the firing cost if productivity deteriorates later. All matches with beliefs in this interval are therefore terminated at tenure T_1, thus creating a spike in the hazard rate at tenure T_1.

Analogously, we can define the worker-specific thresholds $\bar{p}^{1,W}(T_1)$ and $\underline{p}^{1,W}(T_1)$ that satisfy the conditions

$$W^1\left(\bar{p}^{1,W}(T_1), T_1\right) = U$$

$$W^1\left(\underline{p}^{1,W}(T_1), T_1\right) = U - \beta\kappa$$

where $\bar{p}^{1,W}(T_1) > \underline{p}^{1,W}(T_1)$ since worker value functions are increasing in the current belief. Then, at tenure T_1, any match with belief $p \in \left[\underline{p}^{1,W}(T_1), \bar{p}^{1,W}(T_1)\right]$ will not be terminated since the value of such a match will instantaneously jump to $W^2(p, T_1) \geq U$.

However, applying the FOC from the Nash bargaining protocol to beliefs as $t \to T_1$ yields

$$\bar{p}^{1,W}(T_1) = \underline{p}^{1,J}(T_1)$$

so that the interval of no termination for workers lies strictly in a range of beliefs that would have resulted in match termination at an earlier tenure. Therefore, we can safely ignore the effect of firing costs on worker termination incentives and focus solely on firms.

2.4.6 Ergodic Distribution of Beliefs

The Markovian nature of the process for beliefs ensures that there exists a stationary distribution of beliefs in equilibrium.

We again divide to derivation into 2 stages, according to match tenure, and define 2 belief distribution functions f^1, and f^2, where the function $f^i(p,t)$ is the distribution over the unit interval of beliefs for a cohort of matches that began production at the same moment in calendar time, and have reached tenure t.[13]

$t \leq T_1$**:**

We can characterize the evolution of the distribution of beliefs using the Kolmogorov Forward Equation (KFE), which states that f^1 evolves according to the equation

$$\frac{\partial}{\partial t} f^1(p,t) = \frac{\partial^2}{\partial p^2}\left[\Sigma(p) f^1(p,t)\right] - \delta f^1(p,t)$$

The KFE states that the change in density at a belief p for a small change in tenure is the sum of two components. First, beliefs move around in the distribution according the evolution equation. The change in density caused by these movements is captured by the first term on the right hand side. Second, at any belief, a fraction δ of matches end exogenously, causing a negative change to the density captured by the second term.

The initial distribution satisfies the conditions

$$\int_{\underline{p}^1(0)}^{1} f^1(p,0)\, dp = \lambda u$$

[13]The f functions are not strictly distributions since I do not require them to sum to 1. Instead, their total mass will be total employment, and hence 1 minus total unemployment since we have normalized the mass of workers to unity.

$$f^1(p,0) = \Delta(p - p_0)$$

where we recall that λ is the finding rate and u is the mass of unemployed workers, and we define Δ as the Dirac delta function that places all mass at the initial prior p_0.

In addition, at each tenure $t < T_1$, f^1 satisfies the boundary condition

$$f^1\left(\underline{p}^1(t), t\right) = 0$$

which ensures that there is always zero mass at the termination threshold at tenure t, since when a match belief reaches the threshold, it is immediately terminated.

$t \geq T_1$:

The KFE again states that f^2 evolves according to

$$\frac{\partial}{\partial t} f^2(p,t) = \frac{\partial^2}{\partial p^2}\left[\Sigma(p) f^2(p,t)\right] - \delta f^2(p,t)$$

At tenure $t = T_1$, we initialize f^2 with the condition

$$f^2(p, T_1) = f^1(p, T_1)\, \mathbf{1}\left\{p \in \left[\bar{p}^{1,J}(T_1), 1\right]\right\}$$

where $\mathbf{1}$ is the indicator function. This condition ensures that only matches with belief above the threshold $\bar{p}^{1,J}$ survive the termination process at tenure T_1.

Similarly to stage one, at each tenure $t \geq T_1$, f^2 satisfies the boundary condition

$$f^2\left(\underline{p}^2, t\right) = 0$$

which has the same interpretation as in stage one.

Ergodic Distribution of Beliefs

Combining f^1 and f^2 yields the ergodic distribution of beliefs,

$$g(p) = \int_0^{T_1} f^1(p,t)\,\mathbf{1}\left\{p \geq \underline{p}^1(t)\right\} dt + \int_{T_1}^{\infty} f^2(p,t)\,\mathbf{1}\left\{p \geq \underline{p}^2\right\} dt$$

Hence the mass of unemployed workers is given by

$$u = 1 - \int_0^1 g(p)\,dp$$

2.4.7 The Matching Function

Let v be the mass of open vacancies. Together with the mass of unemployed workers u, v vacancies generate m new job matches according to the matching function

$$m = zu^{\eta}v^{1-\eta}$$

where z is matching efficiency. This implies the job finding rate $\lambda = z\theta^{1-\eta}$, and job filling rate $q = z\theta^{-\eta}$, where $\theta = \frac{v}{u}$ is labor market tightness.

2.4.8 Vacancy Posting

A vacant firm has a value function V that satisfies

$$(r+q)\,V = qJ^1(p_0, 0) - c$$

where q is the job filling rate, and c is the per-period cost of maintaining a vacancy. Free entry into the market for vacancies implies that, in equilibrium, $V = 0$.

2.4.9 Equilibrium

Definition 4. A stationary equilibrium is a set of scalars

$$\left\{\lambda, q, \theta, u, v, \underline{p}^2\right\}$$

and functions

$$\left\{J^1, J^2, W^1, W^2, U, \underline{p}^1, w^1, w^2, f^1, f^2\right\}$$

such that:

1. $\left\{J^1, J^2, W^1, W^2, U, \underline{p}^1, \underline{p}^2\right\}$ satisfy their HJB equations and boundary conditions.
2. $\{w^1, w^2\}$ satisfy the Nash bargaining condition.
3. $\{f^1, f^2\}$ satisfy their KFEs, and boundary and initial conditions.
4. $\{\lambda, q, \theta\}$ satisfy their definitions.
5. v ensures free entry in the vacancy posting market.
6. u is consistent with total unemployment implied by $\{f^1, f^2\}$.

In the case without EPL, $\kappa = 0$, and the economy features a unique equilibrium, as shown by Moscarini (2005). When $\kappa > 0$, an existence and uniqueness proof is unavailable due to the lack of tractability imposed by the firing cost. However, in the neighborhood of $\kappa = 0$, it is natural to conjecture that an equilibrium exists and is unique. Our numerical exercises suggest that this conjecture is indeed true.

2.4.10 Aggregate Output and Welfare

Since production is linear in productivity beliefs, we can define aggregate output as

$$Y = \mu^H \int_0^1 pg\left(p\right) dp + \mu^L \int_0^1 \left(1-p\right) g\left(p\right) dp - cv$$

which has the simple interpretation of the number of good matches times the productivity of a good match plus the number of bad matches times the productivity of a bad match minus the flow costs of vacancy creation.

Exploiting the linearity of preferences, and recalling that $u = 1 - \int_0^1 g\left(p\right) dp$, we can define steady state welfare as

$$\mathcal{W} = b\left(1 - \int_0^1 g\left(p\right) dp\right) + \mu^H \int_0^1 pg\left(p\right) dp + \mu^L \int_0^1 \left(1-p\right) g\left(p\right) dp - cv$$

which is the flow value of leisure times the mass of unemployed workers plus the flow value of production, given by aggregate output. This expression also indicates that firing costs affect welfare by altering the shape of the ergodic distribution of beliefs (and hence productivities), and the level of unemployment, both of which depend on the function g.

2.4.11 Efficiency

Although our model contains both endogenous job creation and job destruction, the Hosios (1990) condition still applies to the economy without EPL (see Rogerson et al. (2005) for a simple example). Formally, when $\kappa = 0$, the equilibrium coincides with the solution to the planning problem when the worker's bargaining share is equal to the elasticity of matches with respect to the unemployment rate, i.e. when $\beta = \eta$.

Application of the Hosios (1990) logic follows from the fact that the planner is subject to the same belief inference problem as firms and workers once a new match is formed. As such, the planner chooses the same belief threshold at which to terminate matches as firms and workers do in the absence of EPL. Given this, the only remaining margin is vacancy creation, which is subject to the usual efficiency considerations.

Vacancy creation by firms imposes 2 externalities on the economy. First, an additional vacancy lowers the job filling rate of all other vacant firms, thus creating a negative externality. Second, an additional vacancy increases the job finding rate for all unemployed workers, thus creating a positive externality.

When $\beta > \eta$, the positive externality dominates, and firms do not create enough vacancies in equilibrium. Intuitively, the relatively high value of the workers' bargaining share lowers the firms' share of the match surplus, and hence their incentive to create vacancies. The lack of vacancy creation translates into an inefficiently high unemployment rate in equilibrium.

Conversely, when $\beta < \eta$, the negative externality dominates, and firms create too many vacancies in equilibrium. Intuitively, the relatively low value of the workers' bargaining share raises the firms' share of the match surplus, and hence their incentive to create vacancies. The excess of vacancy creation translates into an inefficiently low unemployment rate in equilibrium.

In light of these effects, whether EPL improves efficiency and welfare depends crucially on how it affects vacancy creation by firms, which is governed by the free-entry condition

$$c = qJ^1(p_0, 0)$$

Intuitively, the presence of EPL lowers the value to firms of entering into a new match, $J^1(p_0, 0)$, since there is a positive probability that they will have to pay $\kappa > 0$ in order

to terminate the match in the future. The free-entry condition then implies that EPL causes a decline in equilibrium vacancy creation.

Therefore, a necessary condition for EPL to be welfare-improving is $\beta < \eta$, so that vacancy creation is too high in equilibrium without EPL. Whether this condition holds depends on the value of the parameters estimated from the data.

2.5 Estimation

The model features 13 parameters: $T_1, \mu^L, \mu^H, r, \eta, z, b, p_0, c, \delta, \sigma, \beta, \kappa$. Since our key data variation is the hazard rate schedule, we use these data to identify and estimate the exogenous separation rate, δ, the noise parameter σ, the bargaining share β, and the firing cost parameter κ. We normalize and or calibrate the remaining parameters.

2.5.1 Calibration

Table 2.2 summarizes our calibration. We set $T_1 = 3$ months to reflect the tenure at which firing costs and other labor market regulations become active. Since μ^L and μ^H simply set the location and scale of production in the economy, we normalize them to $\mu^L = 0$ and $\mu^H = 1$ respectively. We set the annual discount rate $r = 7.5\%$ which is in line with Brazilian interest rates.

The vacancy posting cost parameter is set at $c = 0.05$ so that the total cost of hiring a new worker is approximately equal to paying average wages for 2 months, which is in line with Mortensen and Pissarides (1999).

Following Shimer (2005), we normalize market tightness to $\theta = 1$, and set the matching efficiency parameter z to target a 15-day job finding rate of $\lambda = 0.02$ in our baseline economy. Our target for λ is motivated by appendix figure 2.8, which plots the hazard rate of finding a new job for workers who become unemployed

in our sample. For workers who spend between 15 days and 1 year out of the formal employment sector, the average 15-day job finding rate is approximately 0.02. Furthermore, this choice implies a steady state unemployment rate of approximately 32%, which is consistent with recent Brazilian data, where we interpret unemployment as including both true unemployment, and employment in the informal sector.

The weight on unemployment in the matching function, η, is set to 0.25, in line with the estimates in Hoek (2007) who estimates a Cobb-Douglas matching function for Brazil. Given the importance of η in establishing the efficiency properties of the economy, we consider a range of other values for η in our robustness exercises.

As a baseline, we set the value of leisure to $b = 0.4$. This choice reflects two considerations. First, b includes unemployment benefits, which are valued at between 15% and 80% of a worker's previous wage in Brazil. Second, as argued by (Hagedorn and Manovskii, 2008), the value of leisure parameter should be such that workers are indifferent between working and not working on the margin, which implies that b is equal to average labor productivity. Our choice of $b = 0.4$ strikes a balance between each of theses forces, and results in a value of leisure equal to 65% of the mean wage. Given the uncertainty of how to calibrate b, we consider a range of alternative values in the robustness exercises.

Finally, the initial prior belief parameter, p_0, is pinned down endogenously by the free entry condition for vacancy posting (given all other parameters),

$$c = z\theta^{-\eta} J^1 (p_0, 0)$$

Table 2.2: Parameter Calibration

Parameter	Value	Target
T_1	3 months	Timing of EPL
μ^L	0	Normalization
μ^H	1	Normalization
r	7.5%	Short-term real interest rate
θ	1	Normalization
z	0.02	$\lambda = 0.02$
c	0.05	2 months' average wages
η	0.25	Elasticity of matches to unemployment
b	0.4	Unemployment benefits, average labor productivity
p_0	0.52	Free-entry condition

2.5.2 Estimation

We estimate δ, σ, β, and κ via the method of simulated moments. Formally, letting $\Xi = (\delta, \sigma, \beta, \kappa)$ be the vector of parameters, define the vector of model-generated hazard rates for the first 4 years of tenure as

$$H^{mod}(\Xi) = \left(h_1^{mod}(\Xi), h_2^{mod}(\Xi), ..., h_{96}^{mod}(\Xi)\right)$$

where the vector length of 96 reflects the 15 day spacing of hazard rates to match the data. We use 4 years of hazard rate data to cleanly measure the hazard rate after the effects of the learning process have died out so that we can identify δ. Similarly, define the vector of hazard rates in the data as

$$H = (h_1, h_2, ..., h_{96})$$

Then, we choose parameters Ξ to solve

$$\min_{\Xi} \left(H^{mod}(\Xi) - H\right) W \left(H^{mod}(\Xi) - H\right)'$$

where W is a weighting matrix. We set W to the identity matrix in our baseline estimation. Finally, we obtain standard errors using a bootstrap procedure with 100 replications.

2.5.3 Identification

Parameters δ, σ, and β determine the shape of the hazard rate schedule in the absence of firing costs.[14] The exogenous separation rate δ determines the level to which the hazard rate converges as tenure increases and the learning process becomes less prevalent. Therefore, δ is identified by the long run level of the empirical hazard rate.

The speed at which the hazard rate converges to its long run value, and hence the slope of the hazard rate schedule at medium to longer term tenures, is governed by σ since this noise parameter determines the "speed" at which beliefs move around the interval $[0, 1]$: the higher is sigma, the noisier is the belief updating process, the slower beliefs move, and the slower is the convergence.

At shorter term tenures, the average level of the hazard rate schedule is governed by the workers' bargaining share parameter β. Intuitively, when β is higher, firms receive a smaller fraction of the match surplus, so that they have a weaker incentive to terminate an existing match since new matches are less valuable to them. Therefore, a higher value for β implies that short term hazard rates are lower.

[14] Although we cannot formally prove identification of these parameters from the hazard rate data, simple numerical simulations confirm our intuition for how each parameter affects the shape of the overall hazard rate schedule.

Finally, the identification of κ follows similar arguments to the empirical identification: the fact that EPL only becomes active at three months of tenure creates a strong incentive for firms to terminate matches just before this tenure is reached. This behavior creates a spike in the hazard rate schedule at three months' tenure. Hence, we use the spike in the empirical hazard to infer the value of κ, that summarizes the real cost of EPL to the firms.

2.6 Results

2.6.1 Estimation

Hazard Rates Figure 2.5 plots the hazard rate as a function of job tenure both in the data and in the estimated model. Although we only use 4 parameters to target the entire hazard rate schedule, our sparse parameterization does a good job of capturing key features of the hazard rate schedule shape. In particular, the presence of a fixed firing cost at tenures beyond three months results in a spike in the model hazard rate at three months, just as in the data.

Parameter Estimates Table 2.3 shows the estimation results in our baseline model. Standard errors are computed using a bootstrap procedure. Note that the large sample size of our data set ensures that the parameters are very precisely estimated.

We estimate $\hat{\delta} = 0.003$, which closely matches the hazard rate at long tenures in the data. At these tenures, there is very little left for the firm and worker to learn about the quality of the match, which by construction must be of which quality. As such, the only force for termination is the exogenous rate δ.

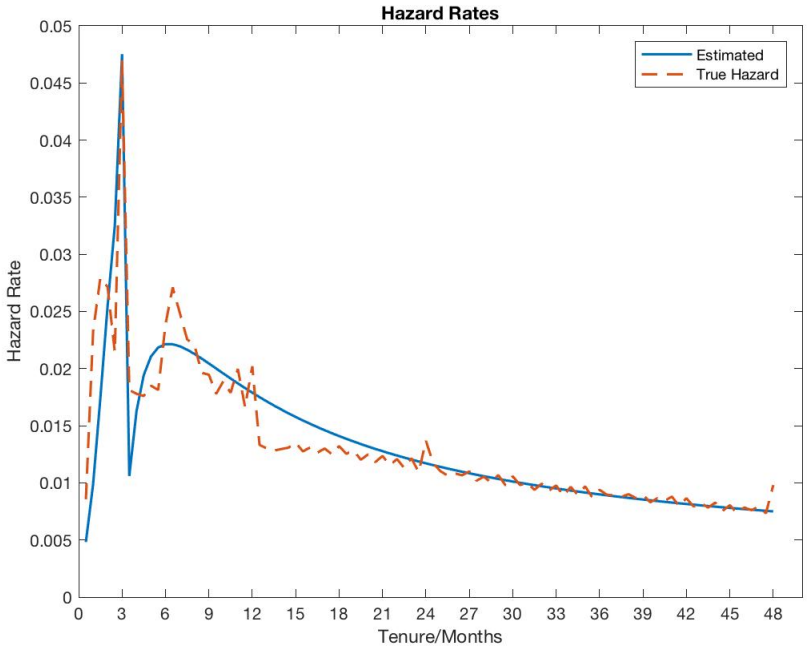


Figure 2.5: This figure plots the empirical layoff hazard rate (dashed) as well as the hazard rate from the estimated model (solid). Tenure duration is binned into 15 day intervals to estimate the empirical layoff hazard.

The estimate of $\hat{\sigma} = 14$ ensures that the estimated hazard rate converges to its long run value at a similar rate to the empirical hazard.[15] The relatively large value of $\hat{\sigma}$, together with the implied hazard rate slope, indicates that it takes approximately 24 months of tenure for the firm and worker to solve the inference problem, and to be confident of the match's true quality.

The value of $\hat{\beta} = 0.63$ implies that workers receive 63% of the surplus generated by a match. Note that this is larger than the elasticity of unemployment in the matching function, $\eta = 0.25$, which violates the Hosios (1990) condition for efficiency.

Finally, the estimate of the EPL parameter $\hat{\kappa} = 0.25$ implies that the cost to a firm of terminating a job at a tenure of more than three months is approximately 1.7%

[15]Recall that the empirical hazard rate schedule is subject to additional frictions that occur at tenures beyond three months. These create small but visible bumps in the hazard rate that our model cannot match by design. Instead, the estimated hazard rate schedule fits the average decline in the hazard rate over tenures beyond three months.

Parameter	Estimate	Standard Error
δ	0.003	0.00001
σ	14	0.14
β	0.63	0.004
κ	0.25	0.00001

Table 2.3: This table reports the estimated model parameters: δ is the rate of exogenous job destruction, σ controls the speed of employer learning, β is the worker's bargaining share and κ is the real cost to firms of EPL. The model is estimated using a method of simulated moments procedure which chooses parameters to match the empirical job termination hazard. Standard errors are computed using a bootstrap procedure which re-samples at the worker level, estimates the job termination hazard, and re-estimates the model parameters. The standard errors are equal to the standard deviation of the estimated model parameters across 100 bootstrap replications.

of the mean annual wage in the economy. Recall that this parameter incorporates the wide variety of labor market regulations in place in Brazil. Its estimated value is therefore a monetary summary of how all of this regulation impacts firms.

2.6.2 The Equilibrium Effects of EPL

We now turn to our main counterfactual exercise, and study how macroeconomic outcomes change when we set $\kappa = 0$, thus removing EPL from the economy.

Hazard Rates Figure 2.6 plots the hazard rate in the data, in the estimated model, and in the counterfactual model with the EPL removed ($\kappa = 0$). In the absence of firing costs, the hazard rate function resembles the standard hump shape (Farber, 1994). In particular, the hazard rate without firing costs does not spike at three months, and is uniformly higher than the hazard with firing costs at tenures greater than three months, but lower than the hazard with firing costs at tenures less than three months. In other words, the presence of firing costs causes firms to terminate

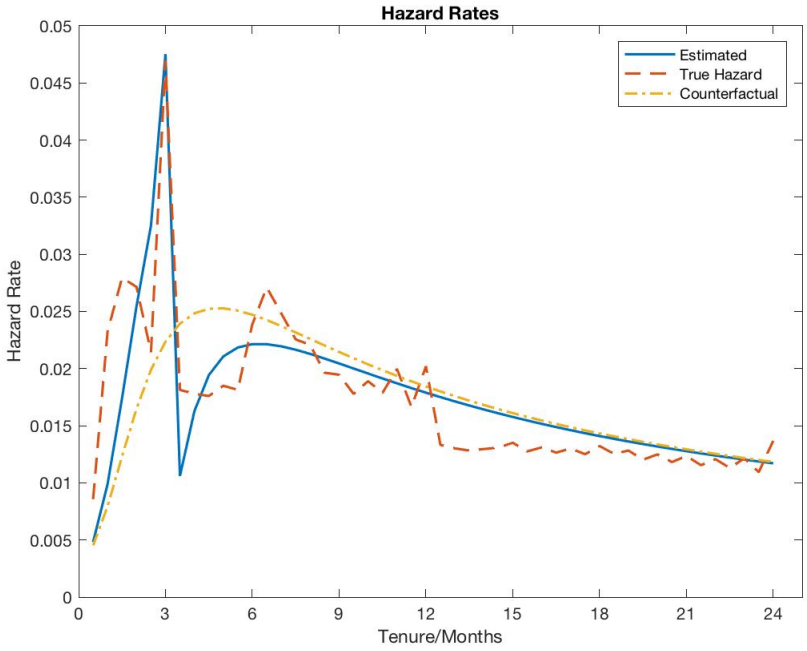


Figure 2.6: This figure plots the empirical layoff hazard rate (dashed), the estimated hazard rate generated by the model (solid), and a counterfactual hazard rate (dotted) when firing costs are set to zero. Tenure duration is binned into 15 day intervals to estimate the empirical layoff hazard.

matches earlier than they otherwise would, just as the theory predicted, and as we found in the data.

Macroeconomic Outcomes Table 2.4 summarizes the key macroeconomic implications of this change in hazard rate pattern, where the change is measured from the economy without EPL to the economy with EPL. Imposing EPL, holding all other structural parameters fixed, leads to an decrease in output of 1.29%, and an increase in the unemployment rate of 0.9 percentage points. As a result, welfare also decreases by 0.39%. Therefore, the imposition of EPL has a negative effect on the aggregate economy.

The negative impact of EPL follows immediately from the fact that $\beta > \eta$ in the estimated structural model, which violates the necessary condition for EPL to be

	Δ
Y	-1.29%
u	0.92 p.p
$\mathcal{W}$	-0.39%

Table 2.4: This table reports the change in macroeconomic outcomes when EPL is imposed on the economy, relative to the counterfactual economy without EPL.

welfare-improving. Intuitively, $\beta > \eta$ implies that vacancy creation is already inefficiently low in the absence of EPL. Imposing EPL then causes a further decline in vacancy creation, which exacerbates the positive externality on workers in the economy, and thus worsens the inefficiency.

Decomposing the Fall in Output A simple decomposition sheds light on the driving force behind the decline in output, Y, where we abstract from the change in vacancy creation costs for simplicity since it is quantitatively much smaller. Letting Y^c, u^c, and g^c be output, unemployment, and the distribution over beliefs in the counterfactual economy without EPL, we can write

$$Y - Y^c = \int_0^1 p\left(g\left(p\right) - g^c\left(p\right)\right) dp$$

which can be decomposed into 2 parts,

$$Y - Y^c = \underbrace{\int_0^1 p\left(g\left(p\right) - \frac{1-u^c}{1-u} g\left(p\right)\right) dp}_{\text{Unemployment Effect}} + \underbrace{\int_0^1 p\left(\frac{1-u^c}{1-u} g\left(p\right) - g^c\left(p\right)\right) dp}_{\text{Productivity Effect}}$$

The “Unemployment Effect" measures how much of the change in output is driven by the change in unemployment, holding the distribution over beliefs fixed at g. Intuitively, if unemployment did not change, then this term would be zero.

Given this term, the "Productivity Effect" measures the residual change, and captures how much of the change in output is driven by the change in the distribution over beliefs, and hence productivity.

Applying our estimation to this decomposition shows that the "Unemployment Effect" accounts for 99% of the change in output.[16] In other words, the dominant effect of EPL is to raise unemployment. EPL has only minor effects on the ergodic distribution of beliefs and hence productivity.

That the drop in output is driven by an increase in unemployment highlights an important unintended consequence of EPL. While EPL creates a system of benefits that workers receive when they enter unemployment, the transfer of resources from firms to workers upon which EPL is built lowers the value that firms ascribe to employment matches in the first place. The reduction in value that firms receive as a result of EPL makes entering into new employment contracts less attractive, which leads them to decrease their rates of vacancy creation, and hence causes an increase in steady state unemployment.

2.7 Robustness Exercises

In this section we consider the robustness of our results to changes in the value of calibrated parameters: the value of leisure b, and the elasticity of matches to unemployment η.

2.7.1 The Value of Leisure, b

We consider 2 additional values, $b \in \{0.2, 0.6\}$, which imply values of leisure equal to 36% and 80% of the mean wage respectively. In each case, we re-estimate the model,

[16] This result is unchanged if we consider the alternative version of this decomposition, holding beliefs fixed at g^c in the "Unemployment Effect" term.

b	0.2	0.4	0.6
δ	0.002	0.003	0.003
σ	11	14	12
β	0.75	0.63	0.51
κ	0.18	0.25	0.27
c	0.05	0.05	0.06

Table 2.5: This table presents results which vary the leisure parameter b from 0.2 to 0.6. In each case, we recalibrate the parameter c and then re-estimate the model to find δ, σ, β and κ. We then perform the counterfactual by setting the firing cost parameter to zero and computing the changes in aggregates.

b	0.2	0.4	0.6
ΔY	-2.03%	-1.29%	-1.58%
$\Delta \mathcal{W}$	-1.24%	-0.38%	-0.32%
Δu	1.5 p.p	0.9 p.p	1.1 p.p

Table 2.6: This table presents results which vary the leisure parameter b from 0.2 to 0.6. In each case, we recalibrate the parameter c and then re-estimate the model to find σ, β and κ. We then perform the counterfactual by setting the firing cost parameter to zero and computing the changes in aggregates.

and re-calibrate the vacancy posting cost parameter c to make each case comparable. We then compute the results of our counterfactual exercise in which we set the firing cost parameter to zero.

Table 2.5 shows the values of the estimated and re-calibrated parameters in each robustness exercise, where we have reprinted the baseline for ease of comparison. Table 2.6 shows the corresponding results of the counterfactual exercise for each value of b.

The results show that EPL has consistently negative effects on macroeconomic outcomes. For example, when $b = 0.2$, the imposition of EPL causes a 2% drop in GDP. The negative impacts are again driven by the fact that $\beta > \eta$ for all calibrated values of b, so that EPL only serves to exacerbate the inefficiently low vacancy creation.

2.7.2 The Elasticity of Matches to Unemployment, η

We consider 2 additional values, $\eta \in \{0.1, 0.4\}$. In each case, we re-estimate the model, and then compute the results of our counterfactual exercise in which we set the firing cost parameter to zero.

Table 2.7 shows the values of the estimated and re-calibrated parameters in each robustness exercise, where we have reprinted the baseline for ease of comparison. Table 2.8 shows the corresponding results of the counterfactual exercise for each value of η.

The results show that, as η increases, the negative effect of EPL is attenuated. For example, when $\eta = 0.1$, the drop in GDP caused by the imposition of EPL is 2.5%, while when $\eta = 0.4$, the drop is only 0.39%. In order to interpret this finding, it is first useful to note that the estimated value of β does not change across the different calibrations of η.

Given this, as η increases, the gap between $\beta > \eta$ is reduced, and the inefficiency of low private vacancy creation is attenuated. Therefore, the further reduction in vacancy creation caused by EPL has a smaller negative effect on the economy than in the case.

η	0.1	0.25	0.4
δ	0.003	0.003	0.003
σ	14	14	14
β	0.64	0.63	0.63
κ	0.25	0.25	0.25

Table 2.7: This table presents results which vary the elasticity η from 0.25 to 0.75. In each case, we re-estimate the model to find δ, σ, β, and κ. We then perform the counterfactual by setting the firing cost parameter to zero and computing the changes in aggregates.

η	0.1	0.25	0.4
ΔY	-2.56%	-1.29%	-0.39%
$\Delta \mathcal{W}$	-0.87%	-0.38%	-0.14%
Δu	1.8 p.p	0.9 p.p	0.23 p.p

Table 2.8: This table presents results which vary the elasticity η from 0.25 to 0.75. In each case, we re-estimate the model to find δ, σ, β, and κ. We then perform the counterfactual by setting the firing cost parameter to zero and computing the changes in aggregates.

2.8 Conclusion

We make two contributions to the literature that studies the effects of EPL on macroeconomic outcomes. First, we exploit tenure-dependence in the design of Brazilian EPL to obtain clean identification of its effect on firms' decisions to terminate jobs. Second, we estimate a structural model of a frictional labor market augmented to include EPL, in order to compute the counterfactual implications of removing EPL on unemployment and aggregate output.

Our baseline estimation results imply that EPL has substantial negative effects on macroeconomic outcomes. These negative effects stem from the fact that the data suggest that private vacancy creation is inefficiently low, and that EPL only serves to exacerbate this inefficiency by reducing vacancy creation even more relative to the efficient level.

Tenure-dependence is a common feature of EPL across many countries. It would be interesting to apply our methods we have developed here to other settings, in order to further improve our understanding of the macroeconomic impacts of EPL.

References

Best, M. C., Cloyne, J., Ilzetzki, E. and Kleven, H.: 2018, Estimating the elasticity of intertemporal substitution using mortgage notches, *Technical report*, National Bureau of Economic Research.

Cahuc, P., Charlot, O. and Malherbet, F.: 2016, Explaining the spread of temporary jobs and its impact on labor turnover, *International Economic Review* **57**(2), 533–572.

Daruich, D., Di Addario, S. and Saggio, R.: 2017, The effects of partial employment protection reforms: Evidence from italy, *Technical report*, Mimeo, Princeton University.

Di Tella, R. and MacCulloch, R.: 2005, The consequences of labor market flexibility: Panel evidence based on survey data, *European Economic Review* **49**(5), 1225–1259.

Farber, H. S.: 1994, The analysis of interfirm worker mobility, *Journal of Labor Economics* **12**(4), 554–593.

Garibaldi, P. and Violante, G. L.: 2005, The employment effects of severance payments with wage rigidities, *The Economic Journal* **115**(506), 799–832.

Garicano, L., Lelarge, C. and Van Reenen, J.: 2016, Firm size distortions and the productivity distribution: Evidence from france, *American Economic Review* **106**(11), 3439–79.

Hagedorn, M. and Manovskii, I.: 2008, The cyclical behavior of equilibrium unemployment and vacancies revisited, *American Economic Review* **98**(4), 1692–1706.

Hoek, J.: 2007, Labor flows in formal and informal labor markets in brazil, *Unpublished manuscript* .

Hosios, A. J.: 1990, On the efficiency of matching and related models of search and unemployment, *The Review of Economic Studies* **57**(2), 279–298.

Kleven, H. J.: 2016, Bunching, *Annual Review of Economics* **8**, 435–464.

Kugler, A. and Pica, G.: 2008, Effects of employment protection on worker and job flows: Evidence from the 1990 italian reform, *Labour Economics* **15**(1), 78–95.

Lazear, E. P.: 1990, Job security provisions and employment, *The Quarterly Journal of Economics* **105**(3), 699–726.

Ljungqvist, L.: 2002, How do lay-off costs affect employment?, *The Economic Journal* **112**(482), 829–853.

Menezes-Filho, N. A. and Muendler, M.-A.: 2011, Labor reallocation in response to trade reform, *Technical report*, National Bureau of Economic Research.

Mortensen, D. T. and Pissarides, C. A.: 1999, New developments in models of search in the labor market, *Handbook of labor economics* **3**, 2567–2627.

Moscarini, G.: 2005, Job matching and the wage distribution, *Econometrica* **73**(2), 481–516.

Rogerson, R., Shimer, R. and Wright, R.: 2005, Search-theoretic models of the labor market: A survey, *Journal of economic literature* **43**(4), 959–988.

Shimer, R.: 2005, The cyclical behavior of equilibrium unemployment and vacancies, *American economic review* **95**(1), 25–49.

Van Doornik, B. N., Schoenherr, D. and Skrastins, J.: 2018, Unemployment insurance, strategic unemployment, and firm-worker collusion.

Appendix

2.A Additional Figures

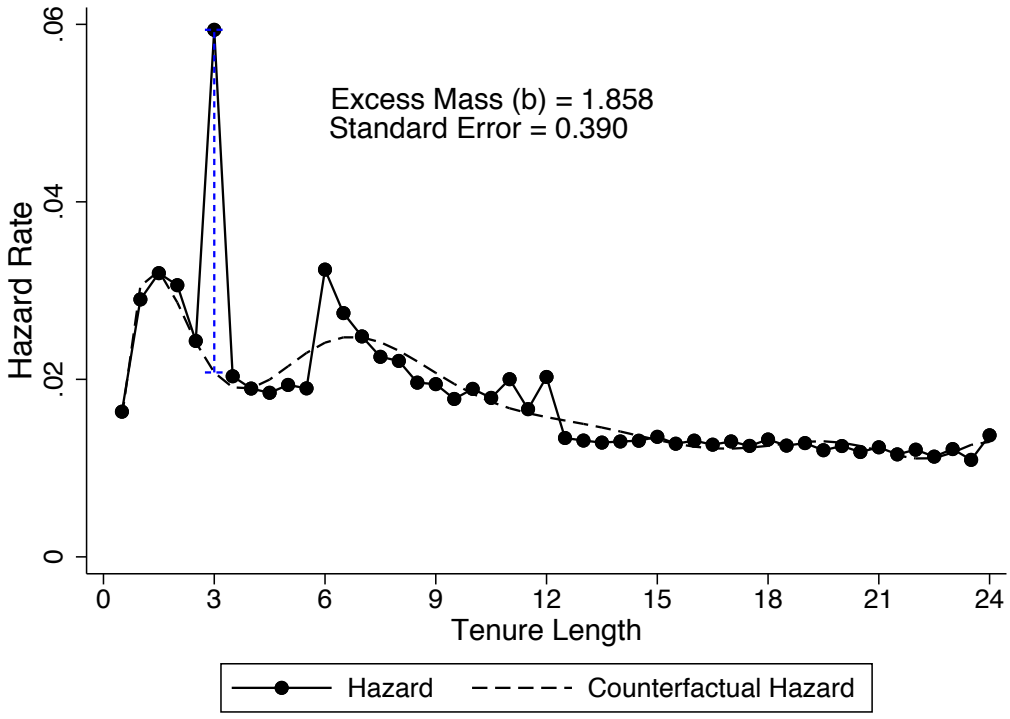


Figure 2.7: This figure plots the job termination hazard rate which includes temporary contracts.

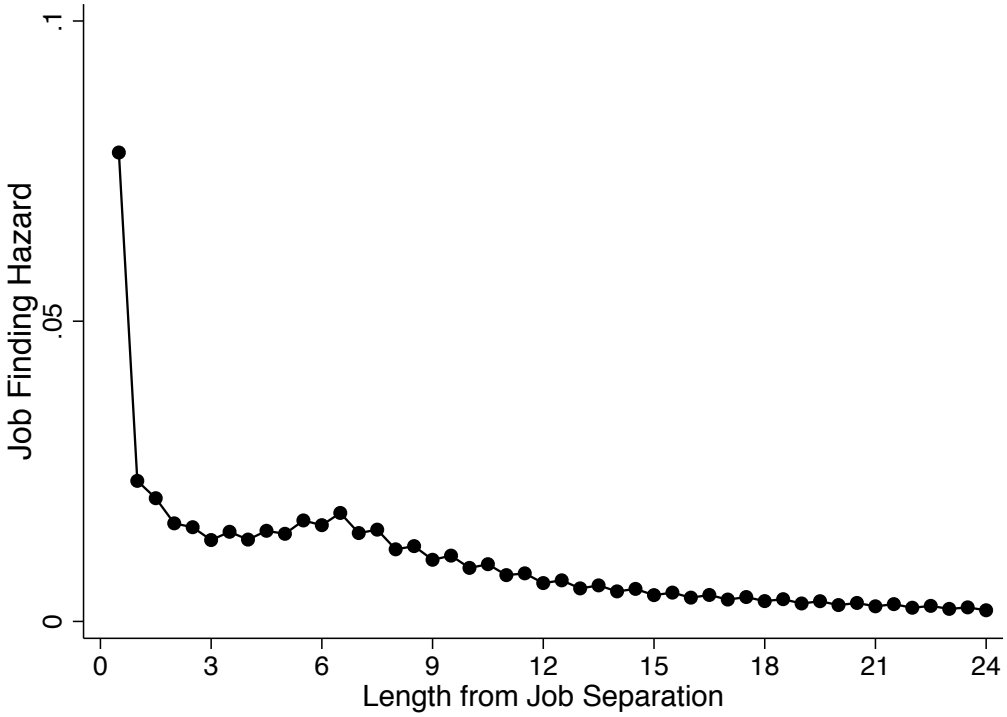


Figure 2.8: This figure plots the job finding hazard rate as a function of duration of unemployment, for workers who become unemployed in our sample.

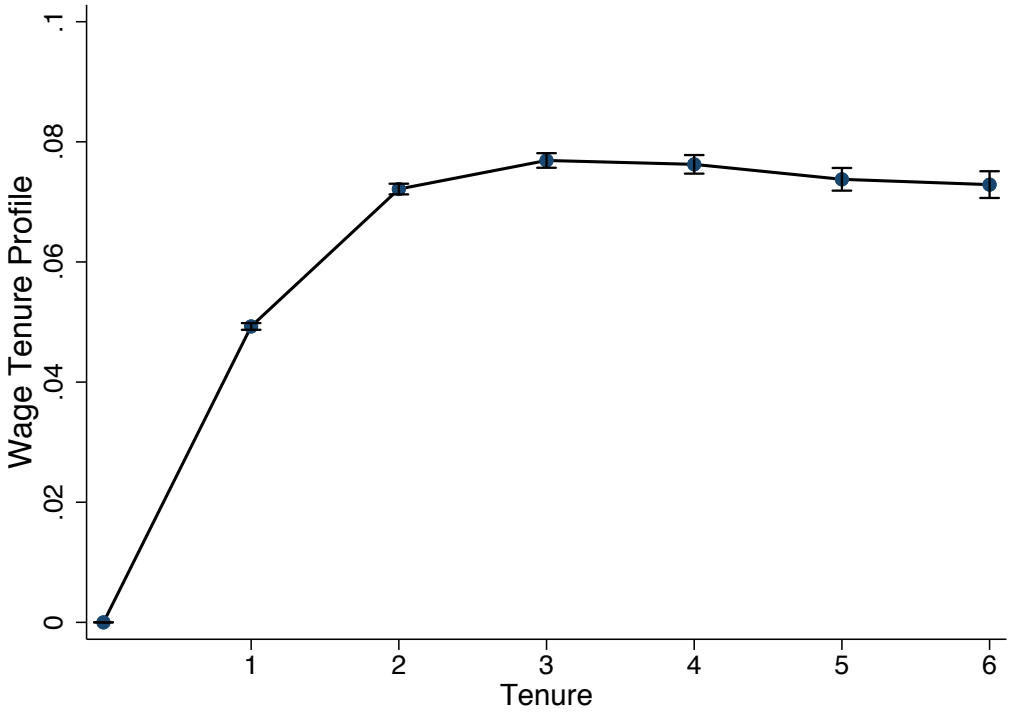


Figure 2.9: This figure plots the return to tenure within a job spell. Standard errors are clustered at the worker level.

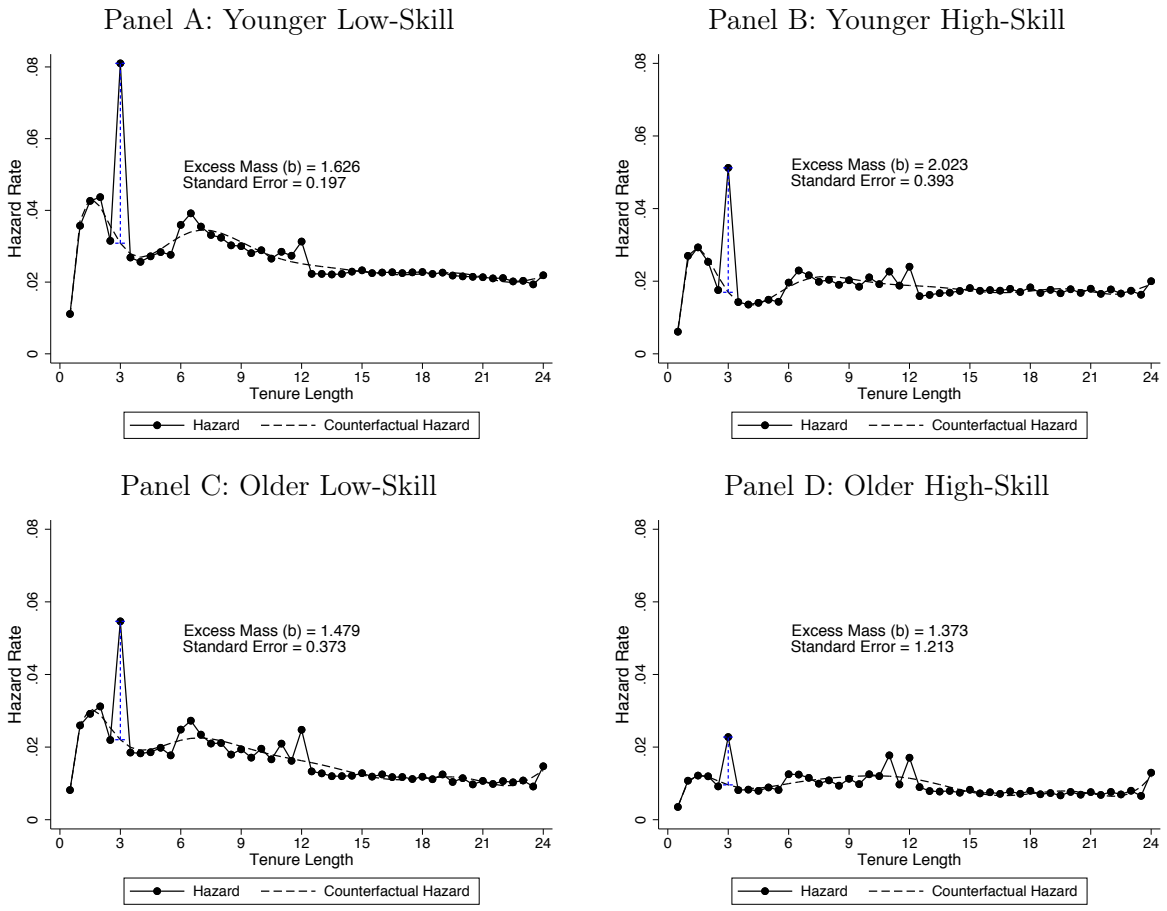


Figure 2.10: This figure plots the layoff hazard rate by different occupation skill levels and age. Skill level is defined by the International Standard Classification of Occupations (ISCO). Younger workers are individuals who are between 20 and 25 years of age at the beginning of the employment spell. Older workers are individuals who are between 40 and 45 years of age at the beginning of the employment spell. Tenure duration is binned into 15 day intervals. The dashed line is a tenth-degree polynomial fitted to the empirical hazard rate, excluding points 15 days away from the notch, as in Equation 2.1. The vertical dotted line displays the excess mass B, while the normalized excess mass b and standard error is reported in the figure. The standard error is computed using a residual bootstrap procedure.

2.B Proofs

2.B.1 Proof of Proposition 11

Following Moscarini (2005), the solution to the firm's HJB in stage 2 is of the form

$$J^2(p) = j(p) + \frac{\bar{\mu}(p) - w^2(p)}{r+\delta}$$

where the j function does not depend on the wage.

At tenure $t = T_1$, the function $J^{*,2}$ is defined as the value to the firm of a given match in the absence of the firing cost κ. Holding the belief of the match fixed, we know that the only effect of the firing cost is to increase the wage by $\omega^1 = \beta(r+\delta)\kappa$. Therefore, $J^{*,2}$ is given by

$$J^{*,2}(p) = j(p) + \frac{\bar{\mu}(p) - w^2(p) + \beta(r+\delta)\kappa}{r+\delta} = J^2(p) + \beta\kappa$$

The argument for $W^{*,2}$ is analogous.

2.C Numerical Implementation

2.C.1 Algorithm to solve the model

Solving the model can be broken into two steps: first solve for the equilibrium firm value functions and belief thresholds, and then solve for the equilibrium unemployment rate.

Solving for J^i and $\underline{p}^i$

After substituting the wage into the firm value functions, we essentially have to solve for three HJB equations together with the optimal belief threshold functions. To do this, we use a finite-difference approximation to the HJB equation, and exploit the fact that the optimal stopping problem characterizing the thresholds can be solved as a linear complementarity problem (see the information and MATLAB codes on Benjamin Moll's website). Specifically, we start by solving for J^2 and $\underline{p}^2$ for $t \geq T_1$ since these objects are stationary. Given these objects, we can then work backwards to solve for J^1, and $\underline{p}^1$.

Solving for f^i

Given the thresholds and a guess of the unemployment rate, we can use a similar finite-difference approximation to solve the KFE forward in time, using the appropriate initial and boundary conditions, to get f^1, and f^2. Using these distributions to compute an implied unemployment rate then yields a simple iterative scheme to find the equilibrium unemployment rate. Given these objects, all other equilibrium objects can be computed using the relationships stated in the theoretical exposition.

2.D Data Appendix

2.D.1 Overview

The *Relação Anual de Informações Sociais* (RAIS) is an employer-employee matched dataset which includes information on all workers and firms in the formal sector of Brazil. The main use of the RAIS is to compute federal wage-supplements (*Abono Salarial*). While not reporting can in theory result in fines, these fines are rarely

issued in practice. However, workers and firms are incentivized to provide accurate wage information given the federal public wage-supplement is based on the wage reported in the RAIS.

2.D.2 Sample Selection

In the RAIS, workers are identified by an individual-specific PIS (Programa de Integração Social), a unique time-invariant worker identifier similar to a social security number. We follow Menezes-Filho and Muendler (2011) and drop workers with PIS identifiers less than 11 digits, as these are not valid identifiers. Errors in worker identifiers may be caused by (1) bad compliance and bookkeeping errors or (2) to allow workers to withdraw from their severance account through fake layoffs and rehires. We eliminate jobs for workers which begin on the same day for the same employer. A single employer may report multiple accounts for one worker so that the workers may access their employer-funded severance payment account, which by law should only be accessed in the case of a firing or for health-related reasons. However, individuals must work at an employer for more than six months in order to access the FGTS account. Therefore, the spike in the job termination hazard cannot be due to employers reporting multiple jobs for the same worker.

2.D.3 Variable Definitions

PIS: A PIS is a worker identifier that is unique to a given worker over time.

Occupation: Occupations are defined by the Classificação Brasileira de Ocupações (CBO) into 2355 distinct groups. We map these occupations to International Standard Classification of Occupations (ISCO) for comparability. Additionally, ISCO classifies occupations by skill level, where occupations that require more training or credentials, and require more specialized work have higher skill levels.

Sector: Sectors are reported under the *CNAE* four-digit classification (*Classificação Nacional de Atividade Econômica*) for 654 industries.

Wage: Wage refers to total payments, including regular salary payments, holiday bonuses, performance-based and commission bonuses, tips, and profit sharing agreements, divided by total months worked during the year for that employer. Payments that are not considered part of the wage include severance payments for layoffs and indemnity pay for maternal leave.

Tenure: The duration the worker has been employed at the establishment. We recode the tenure duration so that it increases in increments of two weeks.

Chapter 3

Optimal Taxation and Fertility Policies

3.1 Introduction

I study a static economy in which households make consumption, labor supply, and fertility choices, which comprise of how many children to have, and how much to invest in them in terms of goods (e.g. food, clothing, and schooling) and time (e.g. early childcare, nighttime reading, and transportation).[1] Using this economy, I consider a planning problem in which a benevolent social planner wishes to redistribute from high income households to lower income households, but does not observe each households' earnings ability, and therefore faces an information asymmetry friction which rules out household specific lump sum taxation. Applying the mechanism design approach, I solve for the optimal allocation, and derive novel expressions for the optimal distortions of fertility choices as functions of sufficient statistics that

[1]The very nature of pregnancy means that many children are of course "unplanned". However, I abstract from such uncertainty in my benchmark specification in order to focus on the trade off between incentive provision and redistribution that lies at the heart of optimal taxation problems.

capture the key economic forces at play. I then use relevant data to pin down these sufficient statistics in order to quantitatively assess the size and sign of the optimal distortions in reality.

I begin by studying the case in which households have preferences that are linear in private consumption. In this case, I show that the elasticity of the fertility choice (child quantity or investment) with respect to earnings ability is a sufficient statistic for the sign of the optimal distortion of that choice, where the distortion is defined as the wedge between the marginal utilities of the fertility choice and consumption.

To grasp the intuition for this result, consider as an example the case of child quantity choice, and suppose that households with higher earnings ability prefer to have fewer children, so that the elasticity of child quantity with respect to earnings ability is negative. In this case, the optimal distortion will be negative, and child quantity will be subsidized on the margin. This subsidy is optimal because is makes it encourages low ability households to have more children, which makes is more costly for high ability households to masquerade as low ability households since child quantity is observable to the planner. This relaxes the incentive constraint that the planner faces so that she can increase the income tax on high ability households which generates more redistribution, and thus increases social welfare.

Similar logic can applied to the choices of child goods investment and child time investment. In each case, the optimality of taxing or subsidizing the choice hinges on the elasticity with respect to the earnings ability parameter. By matching the sign of the elasticity, higher ability households find it more costly to imitate a lower ability household which relaxes the incentive constraint that choices must satisfy. This slack can then be taken up by increasing the income tax burden on higher ability households, which allows more redistribution to take place since total income tax revenue is higher.

I then extend the model to allow households to have preferences over consumption that feature diminishing marginal utility. In this case, the signs of the optimal wedges are jointly determined by two forces. The first is the incentive channel just described. The second stems from the effect that taxing fertility choices now has on a household's optimal choice of income. Intuitively, when fertility choices are taxed (or subsidized), they affect household consumption via the budget constraint. If consumption falls as a result of a tax on some fertility choice, a household may choose to earn more income to offset this loss of consumption. To the extent that this income effect on labor supply occurs, the planner would prefer to tax fertility choices, since increasing household income increases tax revenue from the income tax thus enabling more redistribution to occur. Therefore, it is optimal to subsidize a fertility choice only when the benefits from relaxing incentives outweigh the costs associated with higher income. As in the linear case, I express this trade off in terms of the relevant estimable elasticities.

Having studied properties of optimal wedges, I then provide a result that characterizes the set of tax functions that will actually implement the optimal allocation. This is a non-trivial step since the tax function has to ensure that each household wants to choose the bundle prescribed to them in the optimal allocation designed by the planner, which requires careful treatment of all feasible bundles including those outside of the optimal allocation. The construction builds on Werning (2011) who considers the case of implementing a capital tax given a sequence of savings wedges.

In order to quantify the optimal distortions prescribed by the model, I use data on household fertility choices to pin down the key elasticities that determine the sign and size of the optimal wedges. In the data, lower income households have more children and spend a larger fraction of their income on child-related goods. These patterns imply that it is optimal to subsidize both child quantity and child investments, with larger subsidies for poorer households. Quantitatively, an additional child warrants a subsidy of 3% of median income on average ($2400/year), while expenditures on

children should be subsidized at a rate of 9%. This subsidy is much larger than what is currently observed in the US, and highlights the welfare gains from the optimal tax treatment of child expenditure choices that is arguably missing in the current system. I also show that most of these gains are achievable by introducing a simple linear subsidy on goods related to child investments. This is important, since in reality, observing these choices may be infeasible, thus ruling out non-linear taxation.

Related Literature A core contribution of this paper is to extend the Mirrleesian analysis of taxation to include fertility decisions in order to analyze the properties of optimal taxes (or subsidies) on child quantity and child investments. There are two key advantages to my approach. First, I am able to express optimal distortions in terms of estimable elasticities, thus connecting the optimal tax treatment of fertility choices to the growing literature that estimates the relationships between fertility choices and other variables such as consumption and income. Second, and related to this empirical link, I develop clear intuition for the properties of the optimal wedges, in particular their signs, by exploiting the interpretation of these elasticities. Neither of these have clear precedents in the existing literature on fertility taxation.

There is a group of papers published in the early 2000s that study the optimal taxation of fertility choices. However, relative to my paper, the approaches taken either restrict the set of taxes available (e.g. linear taxes only), or consider non-linear taxation in the presence of only two levels of household ability.

Balestrino et al. (2002) study the optimal non-linear income and child quantity taxes, and the optimal linear child investment tax in an economy with households that can be differentiated along two dimensions: earnings ability, and child rearing ability. The multi-dimensional nature of the type-space makes the analysis complicated even in the two-by-two case that the authors consider. Restricting to only two sets of households also renders the formulas for optimal taxes hard to interpret since they are expressed

in terms of primitives of the model rather than observable quantities. Finally, I find their conclusion about optimal child quantity taxes misleadingly emphasizes the role of child quantity taxes as an offset for the effect of other distortions, a result that relies on their specific assumptions.

Cigno and Pettini (2002) restrict attention to only linear taxes on consumption, child investments and child quantity and find that if higher wage parents invest more in their children, then it is optimal to tax consumption and investments, but to subsidize child quantity. This result has the flavor of results that I derive here. A key difference is that I do not restrict any taxes to be linear, which means that I can link each tax to its relevant elasticity thus making my results more general.

In each of these papers, it is assumed that households care about the quantity and quality of children. This assumption builds on the seminal contribution of Becker (1960) who first conceived of modeling children like other consumption goods, and of the trade off between the quality and quantity of children that a household chooses to have. While this mechanism has been used in the study of growth (e.g. de la Croix and Doepke, 2003), and earnings risk (Sommer, 2015), empirical evidence for the mechanism is somewhat limited. Black et al. (2005) study detailed Norwegian data and establish that family size effects on child quality outcomes are in fact almost entirely due to "birth order effects", where older siblings tend to do better than their younger siblings in later life outcomes such as earnings. Furthermore, while Rosenzweig and Zhang (2009) find evidence of a quantity-quality trade off in Chinese data, they conclude that it is a quantitatively small effect.

In addition to the lack of evidence for quantity-quality theory, there are a host of other potential mechanisms that can explain observed empirical relationships between fertility choices are other variables such as income. Jones et al. (2008) establish the negative correlation between income and quantity of children, and then explore a range of theories to try and explain it, from opportunity costs of time, to het-

erogeneous tastes for children. They conclude that no one theory can explain all the relevant facts simultaneously, and that more research is needed. This lack of consensus on the underlying mechanism motivates my theoretical approach, in which I avoid specifying a particular mechanism, and instead derive results that depend only on reduced form elasticities that are independent of underlying mechanisms.

Finally, to solve the optimal taxation problem, I build on the methodology first developed by Mirrlees (1971), and then further explored by Diamond (1998) and Saez (2001) for income taxation. In addition, I follow the recent trend in optimal taxation since Saez (2001) and derive formulas for optimal taxes that emphasize the roles of parameters that can potentially be estimated in the data, rather than relying on equations that are opaque combinations of primitives of the model.

The paper proceeds as follows. Section 3.2 introduces the household problem laying out the preferences of households that populate the economy. These are vital ingredients to the planning problem which is described in section 3.3. Section 3.4 explores the theoretical properties of the optimal wedges under quasi-linear preferences, while section 3.5 does so without this restriction. Given these results, I connect them to the literature on optimal commodity taxation in section 3.6. In section 3.7, I show how to design a tax function that can implement the optimal allocation. Sections 3.8 and 3.9 solve a quantitative version of the model, and compare the optimal tax policies to those observed in the US currently. In light of this, section 3.10 shows studies the potential welfare gains from implementing simpler policies that those required by the optimal allocation. Section 3.11 concludes.

3.2 A Model of Fertility Choice

There is a unit mass of households who each make choices over consumption, labor effort, child quantity, child goods investment, and child time investment, denoted

by c, l, n, K and H respectively. K and H refer to the total amount of goods and time that are devoted to the household's n children. Each household is further characterized by a one dimensional parameter $\theta \in \Theta = [\underline{\theta}, \infty)$ where $\underline{\theta} \geq 0$. Let the c.d.f. and p.d.f. of the distribution of θ be F and f respectively, where I assume that the p.d.f. exists. As is standard in the optimal taxation literature, I shall refer to θ as a household's ability. However, θ plays two distinct roles in my model, in contrast to the singular role it is usually given.[2]

First, θ captures a household's productivity in producing the single output good in the economy. Specifically, I adopt the usual linear production function and assume that if a household of ability θ supplies l units of labor effort, she earns income $y = \theta l$, where I have implicitly (and without loss of generality) normalized the wage rate to one. Second, θ directly impacts a household's optimal choices of child quantity n and child investments K and H. This interaction is born out in the utility function of households, which I now state (writing $l = y/\theta$):

$$U(c, l, n, K, H) = u(c) - \phi(y/\theta, H) + v(n, K, H, \theta)$$

where u is increasing and at least weakly concave, ϕ is increasing and at least weakly convex in each argument, and v is concave in n and increasing and at least weakly concave in K and H. Note that both the benefits and costs of child quantity are subsumed in the v function. While this is somewhat different to other specifications (for example de la Croix and Doepke, 2003), doing so offers two benefits. First I am able to skip discussion of the non-convexities that arise when considering per-child investments $k = K/n$, $h = H/n$, by working with aggregate variables directly (Alvarez, 1999). Second, deriving taxes that depend on aggregate quantities is simpler

[2]It would of course be more general to let θ be a multi-dimensional parameter. However, as is well understood, solving mechanism design problems with multi-dimensional hidden types is prohibitively difficult. Therefore, I proceed under the assumption that the same parameter enters the utility function twice.

since I do not have to account for how a tax on per-child investment may affect child quantity decisions, and vice versa.

The function v permits a general specification of how θ interacts with the variables n, K, and H. This is advantageous for two reasons. Firstly, more structure is unnecessary since I shall derive optimal taxes whose key features depend only on reduced form elasticities rather than primitives of the model such as v. Second, it is unclear what sort of structural relationship to specify since there is little consensus on what mechanisms govern the way households make fertility choices. As surveyed by Jones et al. (2008), there are a wide range of potential explanations for the observed empirical relationships between income and fertility choices, from standard substitution effects, to models with heterogeneous tastes for children. As such, it would be premature for me to specify an exact mechanism of fertility choice since more research is needed to fully understand which kinds of mechanisms are likely to be the dominant forces at play.

3.3 The Planning Problem

In the baseline model specification, the planner is welfarist and so only cares about the households' utilities. For concreteness, I consider an additive welfare function, with Pareto weights $\{\alpha(\theta)\}_\Theta$ (normalized so that $\mathbb{E}[\alpha(\theta)] = 1$) capturing the planner's preference for redistribution beyond the implicit desire for insurance captured by the risk averse household preferences. I assume that $\partial\alpha/\partial\theta < 0$ so that the planner would like to redistribute resources from higher ability households to lower ability ones. This reflects three potential motivations. First, there could be purely normative reasons for wanting to redistribute. Second, there is the usual idea that the planner would like to insure households against the risk of having low earnings ability in the absence of private insurance markets for this risk. Finally, and novel to this framework, the

planner would like to insure the children born to each household against the risk of having parents with low ability.

If the planner could observe θ, her desired level of redistribution could be achieved by a system of lump sum taxes and transfers conditional on θ. However, as in all taxation models, I assume that the planner cannot observe θ and in fact can only observe a household's choices of income, child quantity and investments $\{y, n, K, H\}$, which prevents the planner from even inferring θ from household choices[3].

Given this asymmetric information environment, I follow the recent taxation literature and use mechanism design to solve the planning problem (Golosov et al., 2003). Specifically, the planner chooses an allocation $A = \{c(\theta), y(\theta), n(\theta), K(\theta), H(\theta)\}_\Theta$ that maximizes social welfare subject to resource feasibility and incentive-compatibility:

$$\max_A \int_\Theta \alpha(\theta) U(\theta) f(\theta) d\theta$$

subject to

$$U(\theta) = u(c(\theta)) - \phi\left(\frac{y(\theta)}{\theta}, H(\theta)\right) + v(n(\theta), K(\theta), H(\theta), \theta)$$

$$\theta = \arg\max_{\theta'} u(c(\theta')) - \phi\left(\frac{y(\theta')}{\theta}, H(\theta')\right) + v(n(\theta'), K(\theta'), H(\theta'), \theta)$$

$$\int_\Theta (c(\theta) + K(\theta)) f(\theta) d\theta \leq \int_\Theta y(\theta) f(\theta) d\theta$$

The first constraint simply defines household utility under truth-telling in the allocation A. The second constraint ensures that each household truthfully reports their type when faced with the allocation A, as is required in a direct mechanism. The last

[3]It is certainly questionable whether some kinds of investment are really observable, especially those involving time. However, from a pedagogical perspective, it makes sense to begin with my assumption of full observability, since relaxing this assumption is a simple extension. Intuitively, when a choice becomes unobservable, taxes on observable goods change to reflect how well targeting them also allows the planner to target the unobserved choice. For an example in the context of human capital, see section 7 of Stantcheva (2015).

constraint ensures that the allocation is resource-feasible, i.e. that total consumption and goods investment is less than or equal to total output.

As is standard, I simplify the set of incentive constraints (ICs) by using the First Order Approach (Farhi and Werning, 2013), which replaces each household's global incentive constraint with a necessary envelope condition from their private optimization. I then verify ex-post that the global ICs are indeed satisfied. Therefore, the planner's problem that I solve can be stated as

$$\max_{A} \int_{\Theta} \alpha(\theta) U(\theta) f(\theta) d\theta$$

subject to

$$U(\theta) = u(c(\theta)) - \phi\left(\frac{y(\theta)}{\theta}, H(\theta)\right) + v(n(\theta), K(\theta), H(\theta), \theta)$$

$$U_{\theta}(\theta) = \frac{y(\theta)}{\theta^2}\phi_l\left(\frac{y(\theta)}{\theta}, H(\theta)\right) + v_{\theta}(n(\theta), K(\theta), H(\theta), \theta)$$

$$\int_{\Theta} (c(\theta) + K(\theta)) f(\theta) d\theta \leq \int_{\Theta} y(\theta) f(\theta) d\theta$$

This problem can be solved using standard optimal control techniques. Details of the first order conditions can be found in the appendix.

3.3.1 Elasticity Concepts

A contribution of this paper is to express optimal taxes in terms of quantities that are estimable, much like Saez (2001) does for income taxation. As such, I now introduce the relevant elasticities that enter the formulas for the optimal taxes.

First, define the compensated elasticity of labor supply with respect to ability by

$$\epsilon^c(\theta) = \frac{\partial \log l^c(\theta)}{\partial \log \theta}$$

where $l^c(\theta)$ is the compensated labor supply function of a household with ability θ. Note that $l^c(\theta)$ is conditional on all other household choices, so that $\epsilon^c(\theta)$ measures the percent change in labor supply for a one percent change ability holding n, K, and H fixed.

Next define the conditional demand elasticity for choice $j \in \{n, K, H\}$ with respect to ability by

$$\epsilon_\theta^j(\theta) = \frac{\partial \log j(\theta)}{\partial \log \theta}$$

where $j(\theta)$ is the optimal choice of j for a given θ, holding other choices fixed. Hence, $\epsilon_\theta^j(\theta)$ measures the percent change in demand for j for a one percent change in ability holding l and all $j' \in \{n, K, H\} \setminus \{j\}$ fixed.

Finally, define the conditional demand semi-elasticity for choice $j \in \{n, K, H\}$ with respect to the tax rate τ_j by

$$e_{\tau_j}^j(\theta) = \frac{\partial \log j(\theta)}{\partial \tau_j}$$

where $j(\theta)$ is the optimal choice of j for a given θ, holding other choices fixed. Hence, $e_\theta^j(\theta)$ measures the percent change in demand for j for a one unit change in τ_j, holding l and all $j' \in \{n, K, H\} \setminus \{j\}$ fixed.

3.3.2 Wedges in the Optimal Allocation

In order to understand features of the optimal allocation, it is intuitive to analyze the distortions it creates relative to the choices each household would make in the absence of intervention. Such comparisons can be made concrete via the concept of wedges, which can be thought of as locally linear marginal tax rates. More precisely, consider the following household problem when the household is subject to linear tax rates on income, children and child investments, τ_y, τ_n, τ_K and τ_H:

$$\max_{y,K,H,n} u\left(y\left(1-\tau_y\right)-K\left(1+\tau_K\right)-\tau_H H-\tau_n n\right)-\phi\left(\frac{y}{\theta},H\right)+v\left(n,K,H,\theta\right)$$

From the first order conditions, it is simple to show that the household's optimal choices are such that the following equations hold:

$$\begin{aligned}
\tau_y &= 1-\tfrac{\phi_l(y/\theta,H)}{\theta u'(c)} \\
\tau_n &= \tfrac{v_n(n,K,H,\theta)}{u'(c)} \\
\tau_K &= \tfrac{v_K(n,K,H,\theta)}{u'(c)}-1 \\
\tau_H &= \tfrac{v_H(n,K,H,\theta)-\phi_H(y/\theta,H)}{u'(c)}
\end{aligned}$$

These wedges provide a convenient way to quantify the distortions created by the planner's optimal allocation, as I now demonstrate. To avoid cluttered notation, all functions and their derivatives are implicitly evaluated at the optimal allocation, e.g. $v_n\left(\theta\right)=v_n\left(n\left(\theta\right),K\left(\theta\right),H\left(\theta\right),\theta\right)$. All proofs of results are contained in the appendix.

3.4 Optimal Fertility Policies Under Quasi-linearity

It is instructive to begin by assuming that the household has quasi-linear preferences, so that $u\left(c\right)=c$. This is reminiscent of Diamond (1998) who analyzed the income taxation problem under risk neutrality, and was able to provide clear intuition for results as a consequence. I start by discussing the income wedge itself, and then move onto the main wedges of interest on child quantity, and goods and time investment.

3.4.1 The Optimal Income Wedge, $\tau_y^*(\theta)$

Proposition 12. *The optimal income wedge is given by*

$$\frac{\tau_y^*(\theta)}{1-\tau_y^*(\theta)} = \eta(\theta)\frac{1-F(\theta)}{\theta f(\theta)}\left(1+\frac{1}{\epsilon^c(\theta)}\right)$$

This expression is standard in the optimal taxation literature (see Saez (2001) for a detailed discussion), and highlights the usual three forces that determine the optimal income wedges. $\eta(\theta) = \frac{1}{1-F(\theta)}\int_\theta^\infty (1-\alpha(v))f(v)\,dv \geq 0$ captures the planner's preference for redistribution via the Pareto weights. $\frac{1-F(\theta)}{\theta f(\theta)}$ is a standard hazard rate term, and trades off the gain in revenue from the wedge (numerator) against the total distortion it creates (denominator). Finally, the elasticity term captures the usual behavioral effects of income taxation, where higher taxation leads to lower labor supply and hence less income to tax. The key difference is that the elasticity $\epsilon^c(\theta)$ implicitly accounts for the fact that labor maybe more or less elastically supplied depending on the fertility choices of the household.

3.4.2 The Optimal Child Quantity Wedge, $\tau_n^*(\theta)$

Proposition 13. *The optimal child quantity wedge is given by*

$$\tau_n^*(\theta) = \eta(\theta)\frac{1-F(\theta)}{\theta f(\theta)}\left(-\frac{1}{e^n_{\tau_n}(\theta)}\right)\epsilon^n_\theta(\theta)$$

The first two terms have the same interpretation as before. Focusing on the novel term $\left(-\frac{1}{e^n_{\tau_n}(\theta)}\right)\epsilon^n_\theta(\theta)$ indicates that the shape and sign of the optimal child quantity wedge depends on how households adjust their choice of child quantity with respect to changes in both the wedge and ability. Noting that $e^n_{\tau_n}(\theta) < 0$ since it captures a pure substitution effect, it is intuitive that the wedge is declining in the absolute size

of this semi-elasticity since if households' child quantity choices are more sensitive to the wedge, then it is a blunter instrument for redistribution, and has larger efficiency costs. This intuition closely resembles the role of ϵ^c in the optimal income wedge expression, as discussed by Saez (2001).

The appearance of $\epsilon^n_\theta(\theta)$ is novel to the literature. I begin by exploiting the fact that all other terms are positive, to establish that the elasticity $\epsilon^n_\theta(\theta)$ is a sufficient statistic for the sign of $\tau^*_n(\theta)$.

Corollary 1. *$\tau^*_n(\theta)$ is negative if and only if*

$$\epsilon^n_\theta(\theta) < 0$$

To grasp the intuition for this result suppose that households with higher earnings ability prefer to have fewer children (the elasticity is negative) so that child quantity will be subsidized on the margin, and having an additional child will lower a household's total tax bill. The optimality of this subsidy stems from how it affects the household incentive constraint. In a standard income tax model incentive constraints bound the amount of redistribution that is possible since if income taxes are too high it becomes better for high earnings ability households to masquerade as lower ability households by choosing to earn less income and hence pay less income tax. However, since child quantity is observable, when households also make fertility choices, for a high ability household to successfully imitate a lower ability household, she would have to choose not only a lower income, but also a higher number of children (given that the elasticity is negative). Optimal taxes can then exploit the fact that there are now two margins of adjustment for each household. In particular, by subsidizing child quantity, lower income households will have even more children, thus making imitation more costly for higher ability households. The subsidy therefore relaxes the incentive constraints of each household since imitation is less desirable. Relaxation

of the incentive constraint is valuable because it can then be tightened by increasing the amount of income tax that higher ability households have to pay, thus allowing more redistribution to take place.

While the sign of ϵ_θ^n determines the sign of the optimal child quantity wedge, its magnitude also plays a distinct role. Specifically, the larger the absolute value of ϵ_θ^n, the larger is the absolute value of the optimal wedge. This relationship stems from the fact that the planner would like to redistribute resources as efficiently as possible, and that the elasticity ϵ_θ^n directly measures how efficient a child quantity distortion is relative to an income distortion. To see this, note that by dividing the expressions for the optimal income and child quantity wedges, I obtain

$$\tau_n^*(\theta) / \frac{\tau_y^*(\theta)}{1-\tau_y^*(\theta)} = \epsilon_\theta^n(\theta) \left(-\frac{1}{e_{\tau_n}^n(\theta)}\right) / \left(1+\frac{1}{\epsilon^c(\theta)}\right)$$

so that for fixed responses to tax rates, $e_{\tau_n}^n(\theta)$ and $\epsilon^c(\theta)$, the relative size of the child quantity wedge is increasing in the magnitude of $\epsilon_\theta^n(\theta)$. Intuitively, when $\epsilon_\theta^n(\theta)$ is larger in absolute value, distorting child quantity is more efficient than distorting income because the incentive constraint is less binding in that direction. Put a different way, as $|\epsilon_\theta^n(\theta)|$ increases, the difference in child quantity for households with different abilities rises for any given tax system. Therefore, distorting a household's child quantity choice is relatively more efficient as imitation is more costly for households along this dimension.

3.4.3 The Optimal Child Goods Investment Wedge, $\tau_K^*(\theta)$

Proposition 14. *The optimal child goods investment wedge is given by*

$$\tau_K^*(\theta) = \eta(\theta) \frac{1-F(\theta)}{\theta f(\theta)} \left(-\frac{1}{e_{\tau_K}^K(\theta)}\right) \epsilon_\theta^K(\theta)$$

Focusing on the novel term $\left(-\frac{1}{e^K_{\tau_K}(\theta)}\right)\epsilon^K_\theta(\theta)$, I note that $e^K_{\tau_K}(\theta) < 0$ by the substitution effect, and that the optimal wedge declines as this tax response increases in magnitude. Given this, it is clear that ϵ^K_θ is a sufficient statistic for the sign of $\tau^*_K(\theta)$.

Corollary 2. *$\tau^*_K(\theta)$ is positive if and only if*

$$\epsilon^K_\theta(\theta) > 0$$

The intuition for this result is similar to the case of child quantity. Suppose that households with higher earnings ability prefer to invest more goods in their children (the elasticity is positive) so that goods investment will be taxed on the margin, and additional investing will raise a household's total tax bill. The optimality of this tax again stems from how it affects the household incentive constraint. Since goods investment is observable, for a high ability household to successfully imitate a lower ability household, she would have to choose not only a lower income, but also a lower investment level. Hence, by taxing goods investment, lower income households will invest less, thus making imitation more costly for higher ability households. The tax therefore relaxes the incentive constraints of each household since imitation is less desirable. Relaxation of the incentive constraint is valuable because it can then be tightened by increasing the amount of income tax that higher ability households have to pay, thus allowing more redistribution to take place.

Like the child quantity wedge, the optimal child goods investment wedge also depends on the size of the elasticity $\epsilon^K_\theta(\theta)$. Again, this link captures the pure efficiency of taxing goods investment relative to taxing any other margin. For example, suppose that $\epsilon^K_\theta(\theta) > 0$ so that $\tau^*_K(\theta) > 0$. Then, $\tau^*_K(\theta)$ increases with $\epsilon^K_\theta(\theta)$ because the steeper the gradient of K with respect to θ, the less the incentive constraint binds in that direction, making taxation of this choice more efficient.

3.4.4 The Optimal Child Time Investment Wedge, $\tau_H^*(\theta)$

Proposition 15. *The optimal child time investment wedge is given by*

$$\tau_H^*(\theta) = \eta(\theta)\frac{1-F(\theta)}{\theta f(\theta)}\left(-\frac{1}{e_{\tau_H}^H(\theta)}\right)\epsilon_\theta^H(\theta)$$

Noting that $e_{\tau_H}^H < 0$ by the substitution effect, I again conclude that the optimal wedge is decreasing in the magnitude of the tax response, and that ϵ_θ^H is a sufficient statistic for the sign of $\tau_H^*(\theta)$.

Corollary 3. *$\tau_H^*(\theta)$ is positive if and only if*

$$\epsilon_\theta^H > 0$$

The intuition for this result is similar to the case of goods investment. Suppose that households with higher earnings ability prefer to invest more time in their children (the elasticity is positive) so that time investment will be taxed on the margin, and additional investing will raise a household's total tax bill. The optimality of this tax again stems from how it affects the household incentive constraint. Since time investment is observable, for a high ability household to successfully imitate a lower ability household, she would have to choose not only a lower income, but also a lower investment level. Hence, by taxing time investment, lower income households will invest less, thus making imitation more costly for higher ability households. The tax therefore relaxes the incentive constraints of each household since imitation is less desirable. Relaxation of the incentive constraint is valuable because it can then be tightened by increasing the amount of income tax that higher ability households have to pay, thus allowing more redistribution to take place.

Like the previous wedges, the optimal child time investment wedge also depends on the size of the elasticity $\epsilon_\theta^H(\theta)$. Again, this link captures the pure efficiency of taxing

goods investment relative to taxing any other margin. For example, suppose that $\epsilon_\theta^H(\theta) > 0$ so that $\tau_H^*(\theta) > 0$. Then, $\tau_H^*(\theta)$ increases with $\epsilon_\theta^H(\theta)$ because the steeper the gradient of H with respect to θ, the less the incentive constraint binds in that direction, making taxation of this choice more efficient.

3.5 Optimal Fertility Policies Without Quasi-linearity

3.5.1 Elasticity Concepts

I now relax the assumption that households have quasi-linear preferences over private consumption, so that $u'(c) > 0$, and $u''(c) < 0$. While this is certainly more realistic, the cost is that the expressions for the optimal wedges become substantially more complicated. In order to maintain a sufficient statistics approach, I must define the following elasticities of consumption.

Define the elasticity of consumption with respect to ability θ holding all choices except $j \in \{n, K, H\}$ fixed by

$$\epsilon_\theta^{c,j} = \frac{\partial \log c}{\partial \log \theta}$$

Define the semi-elasticity and elasticity of consumption with respect to the tax rate τ_j holding all choices except $j \in \{n, K, H\}$ fixed by

$$e_{\tau_j}^{c,j} = \frac{\partial \log c}{\partial \tau_j},\ \epsilon_{\tau_j}^{c,j} = \frac{\partial \log c}{\partial \log \tau_j}$$

As shown in the appendix, these elasticities can be expressed using the elasticities already defined in the risk neutral case. However, for brevity and intuition, I use the consumption notation in the analysis that follows.

Finally, it is useful to define the elasticity of marginal utility of consumption,

$$\upsilon(\theta) = -\frac{u''(c(\theta))\,c(\theta)}{u'(c(\theta))}$$

3.5.2 Optimal Wedges

The Optimal Income Wedge, $\tau_y^*(\theta)$

Proposition 16. *The optimal income wedge is given by*

$$\frac{\tau_y^*(\theta)}{1-\tau_y^*(\theta)} = \eta(\theta)\frac{1-F(\theta)}{\theta f(\theta)}u'(\theta)\left(\frac{1+\epsilon^u(\theta)}{\epsilon^c(\theta)}\right)$$

With a slight abuse of notation, I redefine $\eta(\theta) = \frac{1}{1-F(\theta)}\int_\theta^\infty \left(\frac{1}{u'(v)} - \frac{\alpha(v)}{\lambda}\right) f(v)\,dv \geq 0$ to capture redistribution preferences with household risk aversion. $\lambda = 1/\mathbb{E}[1/u'(\theta)]$ is the social cost of public funds (equal to unity in the risk neutral case). Finally, the elasticity term now captures both the usual behavioral cost of income taxation (ϵ^c) and the benefit via the income effect ($1+\epsilon^u$) since households facing higher taxes will actually supply more labor. Again, these elasticities must be estimated taking fertility choices into account, in contrast to the existing literature that tends to ignore these choices.

The Optimal Child Quantity Wedge, $\tau_n^*(\theta)$

Proposition 17. *The optimal child quantity wedge is given by*

$$\tau_n^*(\theta) = \eta(\theta)\frac{1-F(\theta)}{\theta f(\theta)}u'(\theta)\left(-\frac{\epsilon_\theta^n}{e_{\tau_n}^n}\left(1-\upsilon(\theta)\,\epsilon_{\tau_n}^{c,n}\right) - \frac{\epsilon_\theta^{c,n}}{e_{\tau_n}^{c,n}}\upsilon(\theta)\,\epsilon_{\tau_n}^{c,n}\right)$$

I begin by noting that this expression nests the risk neutral case in which $u' = 1$ and $\upsilon = 0$. Therefore, risk aversion introduces an additional term that crucially plays a role in determining the sign of $\tau_n^*(\theta)$.

Corollary 4. *The optimal child quantity wedge is negative if and only if*

$$-\frac{\epsilon_\theta^n}{e_{\tau_n}^n}\left(1-\upsilon(\theta)\,\epsilon_{\tau_n}^{c,n}\right) < \left(\upsilon(\theta)\,\epsilon_{\tau_n}^{c,n}\right)\frac{\epsilon_\theta^{c,n}}{e_{\tau_n}^{c,n}}$$

Although more complicated than the risk neutral case, intuition for this condition can be obtained by considering each the role of each term separately. First, note that $-\epsilon_\theta^n/e_{\tau_n}^n$ captures exactly the same incentive and efficiency considerations as in the risk neutral case: ϵ_θ^n measures how much distorting n relaxes the incentive constraint, while $e_{\tau_n}^n < 0$ measures the behavioral response of a household's child quantity choice. Second, consider the term $\epsilon_\theta^{c,n}/e_{\tau_n}^{c,n}$. As defined, this term simply measures the mechanical effects that a change in child quantity choice has on consumption via the budget constraint, $c = \theta l(1-\tau_y) - K(1+\tau_K) - \tau_H H - \tau_n n$, holding all other choices fixed. I stress that these elasticities can hence be computed as functions of elasticities already studied. They are not independent parameters. In isolation, this term is irrelevant from the planner's point of view. However, what matters is how these changes in consumption affect the household's optimal choice of labor supply and hence income.

In the risk neutral case, a household's optimal choice of income depends only on ability θ, and the income wedge τ_y. As such, changes in consumption due to changes in child quantity have no effect on household income. Mathematically, this is represented by $\upsilon = 0$, so that the above condition becomes the one discussed in the risk neutral section. However, when the household is risk averse, optimal income choices also depend on all the other choices and wedges, in particular the child quantity choice and wedge. In this case, changes in consumption due to changes in child quantity

directly affect the income choice of a household. For example, if child quantity is subsidized, then having another child offers an alternative source of income to the household. Therefore, a household might optimally choose to work less and have another child instead. But if a household works less, then there will be less tax revenue from the income tax, and therefore the subsidies to child quantity will have to decrease to respect the aggregate resource constraint faced by the planner.

The strength of this effect is measured by the term $\upsilon(\theta)\,\epsilon^{c,n}_{\tau_n}$. $\epsilon^{c,n}_{\tau_n}$ measures by what percentage consumption changes as a result of a 1% change in the child quantity wedge τ_n (acting through the choice of child quantity in the budget constraint). $\upsilon(\theta)$ is an inverse measure of how much a household is willing to substitute children for consumption. In the risk neutral case, $\upsilon(\theta) = 0$ so that a household is infinitely willing to substitute consumption for children and so does not adjust her income at all in the face of a child quantity wedge. As $\upsilon(\theta)$ increases, the household becomes less willing to substitute, and so adjusts her income accordingly to maintain the same level of consumption in the face of a child quantity wedge.

Therefore, when choosing whether to tax of subsidize child quantity, the planner must trade-off the benefits of relaxing incentives against this new cost of discouraging household labor effort. The above condition shows that when the gradient of child quantity with respect to ability is steep (and negative) enough, the benefits of subsidization outweigh the costs, and $\tau_n^* < 0$ is optimal.

The Optimal Child Goods Investment Wedge, $\tau_K^*(\theta)$

Proposition 18. *The optimal child goods investment wedge is given by*

$$\tau_K^*(\theta) = \eta(\theta)\frac{1-F(\theta)}{\theta f(\theta)}u'(\theta)\left(-\frac{\epsilon^K_\theta}{e^K_{\tau_K}}\left(1-\upsilon(\theta)\,\epsilon^{c,K}_{1+\tau_K}\right)-\frac{\epsilon^{c,K}_\theta}{e^{c,K}_{\tau_K}}\upsilon(\theta)\,\epsilon^{c,K}_{1+\tau_K}\right)$$

Rearranging the final term makes it clear that a similar result holds concerning the sign of $\tau_K^*(\theta)$.

Corollary 5. *The optimal child goods investment wedge is positive if and only if*

$$-\frac{\epsilon_\theta^K}{e_{\tau_K}^K}\left(1-\upsilon(\theta)\,\epsilon_{1+\tau_K}^{c,K}\right) > \left(\upsilon(\theta)\,\epsilon_{1+\tau_K}^{c,K}\right)\frac{\epsilon_\theta^{c,K}}{e_{\tau_K}^{c,K}}$$

The intuition for this condition is very similar to the child quantity case. A positive wedge on child goods investment will relax incentives when $\epsilon_\theta^K > 0$, but also mechanically reduces consumption via the budget constraint. To the extent that households are unwilling to substitute consumption for goods investment, this will lead households to actually earn more income to make up for the increased cost of K. In this sense, there is actually a larger force for taxing goods investment on the margin since it has both an incentives and tax revenue benefit.

The Optimal Child Time Investment Wedge, $\tau_H^*(\theta)$

Proposition 19. *The optimal child time investment wedge is given by*

$$\tau_H^*(\theta) = \eta(\theta)\frac{1-F(\theta)}{\theta f(\theta)}u'(\theta)\left(-\frac{\epsilon_\theta^H}{e_{\tau_H}^H}\left(1-\upsilon(\theta)\,\epsilon_{\tau_H}^{c,H}\right)-\frac{\epsilon_\theta^{c,H}}{e_{\tau_H}^{c,H}}\upsilon(\theta)\,\epsilon_{\tau_H}^{c,H}\right)$$

Again, the sign of $\tau_H^*(\theta)$ is determined by the final term in brackets.

Corollary 6. *The optimal child time investment wedge is positive if and only if*

$$-\frac{\epsilon_\theta^H}{e_{\tau_H}^H}\left(1-\upsilon(\theta)\,\epsilon_{\tau_H}^{c,H}\right) > \left(\upsilon(\theta)\,\epsilon_{\tau_H}^{c,H}\right)\frac{\epsilon_\theta^{c,H}}{e_{\tau_H}^{c,H}}$$

Once again, the intuition is similar to that of the goods investment case.

3.6 Connection to Optimal Commodity Taxation

One of Becker's original insights was that fertility choices could be modeled using standard consumer theory. It is therefore unsurprising that the optimal tax treatment of these choices bares some similarities to the theoretical results on optimal commodity taxation. However, I now emphasize that although similar in flavor, my results stem from different underlying mechanisms. Further, I also describe ways in which my approach is somewhat richer than the usual treatment of commodity taxation.

Since Corlett and Hague (1953), it has been well understood that it is optimal to tax those commodities that are substitutable for labor in the utility function. This stems from the simple logic that by encouraging labor supply, the planner can achieve more redistribution via the income tax. More recently, Boadway and Jacobs (2014) have derived explicit formulas for these optimal commodity taxes that spell out this intuition, and feature relevant elasticities, much like I do here. However, there are two key differences between the usual approach to commodity taxation and the method I adopt in this paper.

First, commodity taxes are almost always assumed to be from a restricted class. The fact that the planner cannot in general observe how much of a particular commodity each household purchases means that taxation must be anonymous. Therefore, commodity taxes are restricted to be linear. By contrast, fertility choices can much more plausibly be assumed to be observed by the planner. In reality, many governments collect information on household births for social security reasons, and surveys such as the Consumer Expenditure Survey and American Time Use Survey in the USA contain special sections devoted to investments in children. Given this observability, non-linear taxation of fertility choices becomes feasible and motivates the approach I have taken in this paper. As a result of this non-linearity, the forces determining optimal fertility taxes and their relationships to the optimal income tax schedule differ

markedly from the case of commodity taxation. As I have shown, the optimal taxes on fertility choices depend on social welfare weights and the distribution of abilities in the economy; both of these forces are missing in Boadway and Jacobs (2014) when linear commodity taxes are considered. Perhaps surprisingly, the non-linearities I consider actually lead to simpler (indeed closed-form) expressions for the optimal taxes since I can express everything in terms of ability, while linear taxes lead to complicated expressions involving the so-called "index of discouragement" (Mirrlees, 1976) and averages over all other taxes being chosen.

When analyzing the optimal taxes on fertility choices, a key insight I develop is the role that these taxes play in relaxing incentive constraints. On a superficial level, this result seems very similar to the original point made by Corlett and Hague (1953). However, the actual mechanism underlying my result is very different to the usual mechanism from commodity taxation, which relies on non-separabilities in the household utility function between commodity choices and labor supply (Atkinson and Stiglitz, 1976). In particular, my utility function is separable so that the usual mechanism for commodity taxation is shut down. The key difference, and the underlying force for taxation of fertility choices in my model, is that the ability parameter θ enters the utility function twice. The first is in the disutility of labor supply function ϕ, reflecting the standard interaction between income and ability. The second is in the sub-utility function v, which then generically directly links ability to the fertility choices. As described earlier, this interaction reflects a wide range of potential mechanisms governing the patterns of fertility choices in the population, from quantity-quality trade-off to heterogeneous tastes for children. It is this link that drives the optimality of distorting household fertility choices. An advantage of my analysis is that the exact mechanisms behind the link is irrelevant for the properties of the optimal taxes. All that matters are the elasticities that the mechanisms generate.

3.7 Implementation

While the wedges just analyzed are useful to build intuition for how households' choices are distorted by the planner in the optimal allocation, they do not give insights into how taxes might be used to implement the allocation itself. The design of such taxes is the topic of this section.

Formally, a tax system is a function $T : \mathbb{R}^4_+ \to \mathbb{R}$ that maps household choices (y, K, H, n) into a tax payable to the planner $T(y, K, H, n)$, where $T(y, K, H, n) < 0$ means that the household receives funds from the planner. A tax system implements the optimal allocation if and only if

$$\begin{aligned}(c(\theta), y(\theta), K(\theta), H(\theta), n(\theta)) &= \arg\max_{c,y,K,H,n} u(c) - \phi\left(\tfrac{y}{\theta}, H\right) + v(n, K, H, \theta) \\ &\text{s.t. } c = y - K - T(y, K, H, n)\end{aligned}$$

for all household types $\theta \in \Theta$, i.e. if the tax system induces each household to choose the bundle prescribed to them in the optimal allocation chosen by the planner.

In order to define the set of tax systems that will implement the optimal allocation, I adapt a procedure described by Werning (2011). The idea is to find the "smallest" tax system that will implement the allocation, where "smallest" refers to the size of the taxes specified by the function T.

Proceeding constructively, first fix a type $\theta \in \Theta$, and consider a type specific tax system that ensures that a household of type θ always wants to choose her prescribed bundle $(c(\theta), y(\theta), K(\theta), H(\theta), n(\theta))$. In other words, given any bundle (y, K, H, n), set the tax on that bundle such that the household is indifferent between that bundle and her prescribed choice:

$$u\left(y - K - T^{\theta}(y, K, H, n, \theta)\right) - \phi\left(\frac{y}{\theta}, H\right) + v(n, K, H, \theta) = U(\theta)$$

where $U(\theta)$ is the household's utility from choosing her prescribed bundle in the optimal allocation. Rearranging this equation yields an expression for the "smallest" type-specific tax that will induce the type θ household to choose the correct bundle, where "smallest" refers to the indifference condition:

$$T^{\theta}(y, K, H, n, \theta) = y - K - u^{-1}\left(U(\theta) + \phi\left(\frac{y}{\theta}, H\right) - v(n, K, H, \theta)\right)$$

Note that for $(y, K, H, n) = (y(\theta), K(\theta), H(\theta), n(\theta))$, the tax also satisfies the condition

$$c(\theta) = y(\theta) - K(\theta) - T^{\theta}(y(\theta), K(\theta), H(\theta), n(\theta), \theta)$$

Having defined a tax system for each type of household, the dependence on θ can be removed by taking the upper envelope of taxes at every bundle (y, K, H, n):

Proposition 20. *Given a set of "smallest" type specific tax systems* $\left\{T^{\theta}(y, K, H, n, \theta)\right\}_{\Theta}$, *the tax system defined by*

$$T^{*}(y, K, H, n) = \sup_{\theta'} T^{\theta}(y, K, H, n, \theta')$$

will implement the optimal allocation.

The proof builds on the fact that at any bundle outside of the optimal allocation, the tax is constructed to be large enough to deter any household from wanting to choose it, while at any $(y(\theta'), K(\theta'), H(\theta'), n(\theta'))$ in the optimal allocation,

$$\begin{aligned} T^{*}(y(\theta'), K(\theta'), H(\theta'), n(\theta')) &= T^{\theta}(y(\theta'), K(\theta'), H(\theta'), n(\theta'), \theta') \\ &= y(\theta') - K(\theta') - c(\theta') \end{aligned}$$

which together with incentive compatibility of the optimal allocation ensures that each household type chooses her prescribed bundle.

It is clear that T^* itself is the "smallest" tax system that will implement the optimal allocation since it was constructed from a set of "smallest" type specific tax systems. Therefore, any tax system that is at least as large at T^* everywhere and is equal to T^* on the optimal allocation will also implement it.

Proposition 21. *If T is a function of (y, K, H, n) such that $T(y, K, H, n) \geq T^*(y, K, H, n)$ everywhere, and $T(y, K, H, n) = T^*(y, K, H, n)$ for all (y, K, H, n) in the optimal allocation, then T will implement the optimal allocation.*

This result provides a clean way to assess the optimality of real policies that redistribute resources based on income, child quantity, and child investments, to which I now turn.

3.8 Quantitative Exercise

In this section, I implement my model numerically, in order to quantitatively analyze the optimal tax policies. I first construct a baseline economy and calibrate it to U.S. data, and then proceed to analyze the optimal taxes under this calibration.

3.8.1 Baseline Calibration

In the baseline economy, there is no social planner, and households simply face linear taxes and subsidies on income, child investments, and child quantity. For simplicity, I abstract from the child time-investment decision. Using the following functional form for utility, a household of ability θ solves

$$\max_{c,l,K,n} \frac{c^{1-\upsilon}}{1-\upsilon} - \eta \frac{l^{1+\frac{1}{\epsilon}}}{1+\frac{1}{\epsilon}} + \psi \frac{(\theta K)^{1-\sigma_K}}{1-\sigma_K} + \frac{(\theta n)^{1-\sigma_n}}{1-\sigma_n} - \phi n$$

subject to

$$c + (1 + \tau_K) K + \tau_n n = (1 - \tau_y) \theta l$$

The main advantage of this specification of utility is that it provides clean expressions for the key elasticities mentioned in the theoretical analysis. In particular, it is simple to show the following:

Proposition 22. *Using this utility specification, the relevant elasticities can be computed as*

$$\epsilon^c(\theta) = \epsilon$$

$$\epsilon^K_\theta(\theta) = -\frac{\sigma_K - 1}{\sigma_K}$$

$$\epsilon^n_\theta(\theta) = -\frac{\sigma_n - 1}{\sigma_n}$$

$$\upsilon(\theta) = \upsilon$$

This tight link between preference parameters and elasticities yields a transparent calibration method that I now describe.

Exogenously Calibrated Parameters

I set the coefficient of relative risk-aversion, $\upsilon = 1$, and the Frisch elasticity of labor supply, $\epsilon = 1$. I also set the linear tax on labor income $\tau_y = 0.2$ and the linear tax on child investment goods $\tau_K = 0$ so that it has the same relative price as the consumption good since a majority of goods are consumed by both parents and children within the household. All exogenously set parameters are summarized in Table 3.1.

Parameter	Value	Target
υ	1	CRRA = 1
ϵ	1	Frisch elasticity = 1
τ_y	0.2	
τ_K	0	

Table 3.1: Exogenously Set Parameters

Endogenously Matched Parameters

I assume that the distribution of ability is from the Pareto-lognormal family, and choose its parameters (mean μ, standard deviation Σ, and Pareto tail parameter α) so that the income distribution matches the income distribution of households with children from the US Census Bureau (2014). Relatedly, I set η so that on average, households spend a third of their total time endowment (normalized to 1) working.

I set the remaining parameters so that the baseline economy replicates salient features of the data related to fertility choices. Since ϕ scales the utility cost of having children, I set it so that the mean number of children born to a household is 1.9 (Sommer, 2016). Since σ_n uniquely determines the relationship between fertility and ability (and hence income), I set σ_n to match an income elasticity of fertility $\frac{\partial \log n}{\partial \log y} = -0.1$, which is in line with the evidence discussed in Jones et al. (2008) and the US Census Bureau (2017). In a similar manner to ϕ and σ_n, the parameters ψ and σ_K determine the scale and gradient of child expenditures. Therefore I set these parameters to match two moments documented by Lino (2014): households in the the bottom third of the income distribution spend 25% of their income on each child while households in the middle third spend 16% of their income per child. Finally, the linear tax on children is set at $\tau_n = -0.02$ (i.e. it is a subsidy) so that the fraction of GDP spent by the government on households with children is 0.5% in the baseline economy (OECD, 2017). All endogenously set parameters are summarized in Table 3.2.

Parameter	Value	Target	Source
μ	0.0011	Income Distribution	Census (2014)
Σ	0.82	Income Distribution	Census (2014)
α	40	Income Distribution	Census (2014)
η	8	$\mathbb{E}[l] = \frac{1}{3}$	
ϕ	0.23	$\mathbb{E}[n] = 1.9$	Sommer (2016)
σ_n	1.25	$\frac{\partial \log n}{\partial \log y} = -0.1$	Census (2017)
ψ	0.6	$\mathbb{E}\left[\frac{K}{ny}\right] = 0.25$ for bottom third	Lino (2014)
σ_K	1.09	$\mathbb{E}\left[\frac{K}{ny}\right] = 0.16$ for middle third	Lino (2014)
τ_n	-0.02	$\mathbb{E}[\tau_n n] = 0.5\%$ GDP	OECD (2017)

Table 3.2: Endogenously Matched Parameters

Pareto Weights

In order to compute optimal taxes, I need to choose a set of Pareto Weights. In my main results, I choose utilitarian weights, so that $\alpha(\theta) = 1$ for all $\theta \in \Theta$. This means that all redistribution is driven by the planner's desire to insure households against the risk of having low ability in earning and child rearing.

3.8.2 Results

The Optimal Wedges on Child Quantity, $\tau_n^*(\theta)$

Figure 3.1 shows the optimal wedges for child quantity, where, to ease interpretation, I have plotted the wedge as a fraction of median income. Three features stand out. First, the wedges are always negative which means that the optimal allocation features marginal subsidies to child quantity. The optimality of subsidies is justified by appealing to the intuition behind corollary 4. In the calibration, $\sigma_n > 1$ so that $\epsilon_\theta^n < 0$, i.e. higher ability households have fewer children. Furthermore, $\upsilon = 1$ so that

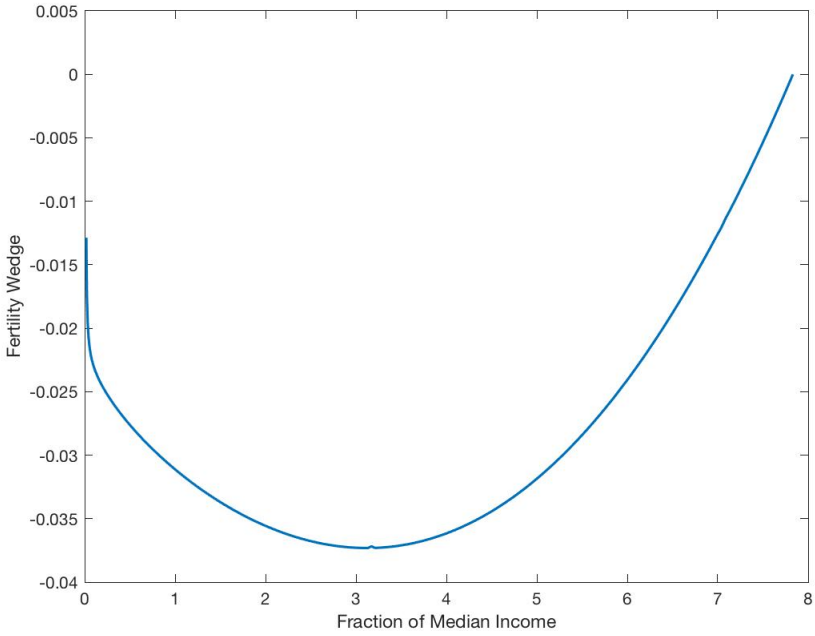


Figure 3.1: Optimal Wedges on Child Quantity

income effects on labor supply are relatively weak. Therefore, the benefit of loosening incentive constraints by subsidizing child quantity outweighs the cost of lower labor supply. Second, the wedges have a pronounced U-shape, which is determined by the pattern of the term $\eta(\theta)\frac{1-F(\theta)}{\theta f(\theta)}$ in the expression for the optimal wedge: the subsidy is largest where the benefit from redistribution most outweighs the cost of distorting household choices. Finally, the wedges are quantitatively meaningful and take an average value of 3% of median annual income, which translates to \$2400/year using my empirical income distribution. This magnitude is comparable to real child credit policies in the US and other countries.

The Optimal Wedges on Child Investment, $\tau_K^*(\theta)$

Figure 3.2 shows the optimal wedges for child investment, where again, I have plotted the wedge as a fraction of median income. Similar to the child quantity wedge, there are three key features. First, the wedges are always negative which means that the

optimal allocation features marginal subsidies to child investment. The optimality of subsidizing child investment on the margin follows a similar logic to the child quantity wedge. In the calibration, $\sigma_K > 1$ so that $\epsilon_\theta^K < 0$, i.e. higher ability households spend less in total on their children. Combining this with the weak income effect on labor supply implies that subsidizing child investment is optimal, since the benefit of loosening incentive constraints outweighs the cost of reducing labor supply. Second, and unlike the quantity wedges, the child investment wedges are predominantly increasing in household income, so that richer households receive a smaller marginal subsidy than poorer households. In this sense, the child investment wedges are progressive. This progressivity is driven by the fact that income effects matter more to the planner for households with higher ability and hence income: discouraging the labor supply of a high ability household has a larger effect on the aggregate resource constraint (and hence total redistribution) than discouraging the labor supply of a lower ability household. Therefore it is optimal to set smaller subsidies of child investment for higher ability households. Finally, the wedges are quantitatively significant, and average 9% of median annual income, which translates to \$7,300/year using my empirical income distribution.

3.9 Implementation and Policy Comparisons

3.9.1 Implementation: computing a version of $T^*(y, K, n)$

As discussed earlier, the wedges alone are insufficient to understand what taxes can implement the optimal allocation. Furthermore, while the function T^* provides a kind of lower bound on how large taxes must be at every possible bundle (y, K, n), it is a fairly abstract object with little real-world motivation. Therefore, I now consider a

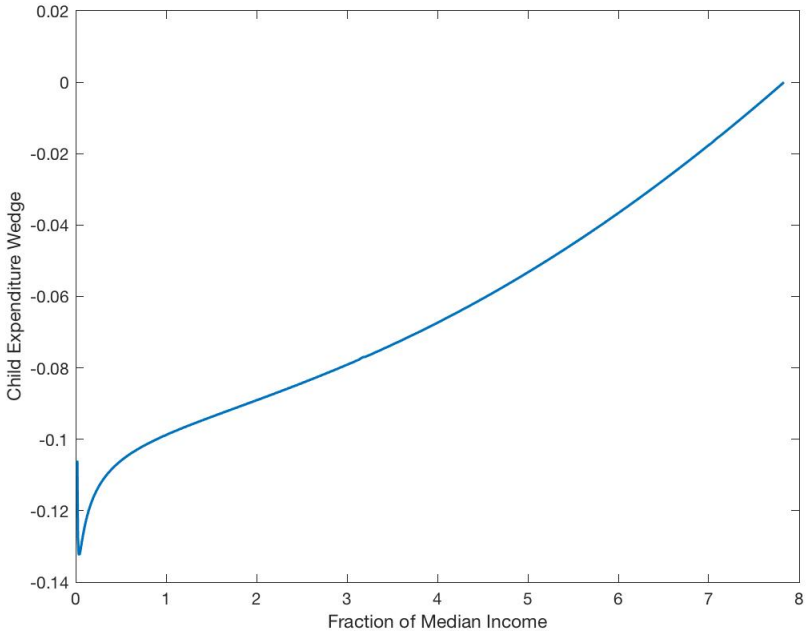


Figure 3.2: Optimal Wedges on Child Investment

tax function T that is everywhere equal to T^*, but can much more easily be connected to policies we observe in reality.

In the USA, a household's total tax bill is computed by subtracting any available credits away from the income tax bill. Therefore I include as the first component of my tax function T, an income tax function $T_0(y)$, that matches the real income tax function used for Head of Household filers in the US in 2016 (Tax Foundation). The table below shows how these income taxes are computed.

The second component, T_1, is simply computed as the residual tax bill,

$$T_1(y, K, n) = T^*(y, K, n) - T_0(y)$$

and hence represents total credits received by the household in the case $T_1 < 0$. Note that T_1 must depend on all three choices (y, K, n) to ensure that equality can be obtained for all possible bundle choices.

Lower Bracket ($)	Lower Bracket ($)	Marginal Tax Rate
0	13,250	0.1
13,250	50,400	0.15
50,400	130,150	0.25
130,150	210,800	0.28
210,800	413,350	0.33
413,350	441,000	0.35
441,000	$+\infty$	0.396

Table 3.3: US Income Tax Parameters (2016)

In sum, I consider a tax function

$$T(y, K, n) = T_0(y) + T_1(y, K, n)$$

that is everywhere equal to $T^*(y, K, n)$ and hence implements the optimal allocation.

3.9.2 A Review of Tax Policy Towards Families in the USA

Tax policy towards families in the USA can broadly be split into two key categories: income tax credits, and health care subsidies. In each of these categories, the most prominent programs are the EITC and CTC, and Medicaid.

EITC and CTC As described in the introduction, the EITC has significant dependence on both income and child quantity within a household. Analysis by Hoynes and Rothstein (2016) shows that the program is targeted at low income households, with credits increasing in the number of children. The CTC is similar to the EITC in design, but has much larger reach in terms of income eligibility - households with income in excess of $100,000/year are still eligible for CTC payments. In the same

paper, Hoynes and Rothstein comment that this scope seems puzzling and perhaps goes against the redistributive aims of the government.

Medicaid Medicaid provides subsidized healthcare coverage for families whose incomes are below 138% of the Federal Poverty Line for the relevant household size. If a household is eligible, I model Medicaid as a payment of $2577 per child to the household, which was the estimated expenditure per child in the USA in 2014 (Henry J Kaiser Family Foundation).

3.9.3 Results

Figure 3.3 compares the tax credits implied by the function T_1 to the tax credits implied by the the real policies just described. The figure plots the credits as a function of household income, holding the number of children fixed at $n = 2$. Since the function T_1 also depends on child expenditures, I fix K at its optimal value for each income. For completeness, other cases are considered in the appendix. The top panel shows results for all households, while the bottom panel shows only the bottom 90% of households in the income distribution.

The key result is that for all income levels, credits in the model are substantially larger than the credits implied by real policies. While the large credits at the top of the distribution are an artifact of the exogenous income tax function I imposed as part of my implementation, the large credits at the bottom are robust to other income tax functions since income taxes are necessarily close to zero when income is close to zero.[4]

[4]The calibrated distribution of ability implies that optimal distortions should be zero at the top of the distribution. This means that the optimal marginal tax rate on income goes to zero for the highest income levels, which is in stark contrast to the income tax function I use to implement the allocation. As a result, the credit function features large credits to these households to offset the excessive income tax that they pay.

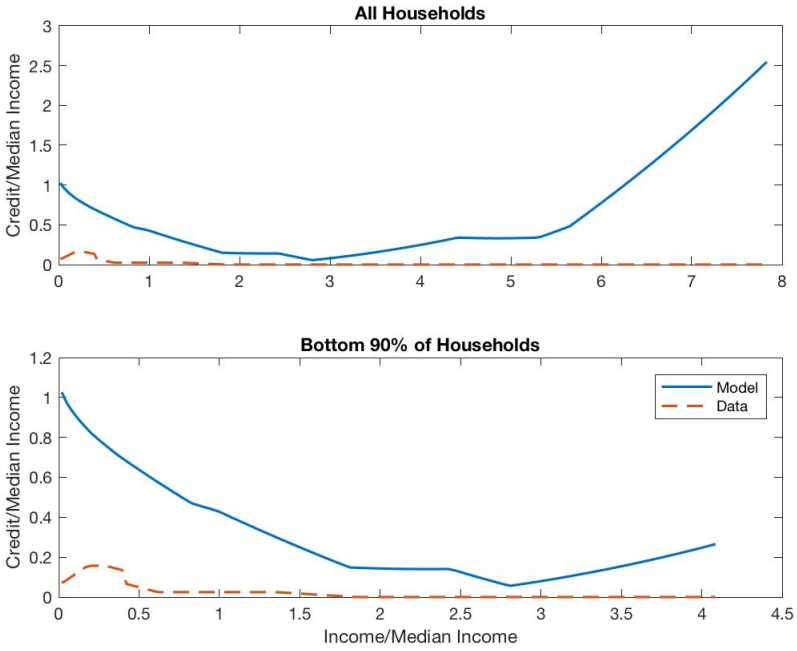


Figure 3.3: Tax Credits for households with two children

The large magnitude of credits in the optimal allocation follows from the logic outlined in the analysis of the optimal wedges. Subsidizing child quantity and expenditures on the margin not only mechanically increases the total credits received by poorer households, but also loosens the incentive constraints of higher ability households, which allows the planner to increase marginal income tax rates, and afford even more distribution towards lower ability households. Quantitatively, $\sigma_K > \sigma_n$ implies that subsidizing child expenditure is a relatively more efficient way to achieve this redistribution, though both choices are distorted in the optimal allocation.

This logic suggests that real credits are too small relative to the optimum because they do not fully exploit the gains from subsidizing both child quantity and child investment. In particular, none of the real policies mentioned have any dependence on child expenditure, and thus completely forgo the associated gains from subsidization, while the marginal subsidies of child quantity seem too small.

3.10 Welfare Gains From Simple Policies

The numerical exercises suggest that there are large welfare gains from subsidizing child-related investments and child quantity choices. While child quantity is arguably verifiable by a tax authority, it is much harder to make the same case for child expenditures, which in reality may span a wide range of goods and services, some of which will naturally overlap with personal expenditures (e.g. food and housing). When child expenditures are treated as anonymous by the social planner, non-linear taxation becomes infeasible, and linear taxes or subsidies must be considered instead. In this section, I investigate how much of the welfare gain from non-linear taxation of child expenditures can be achieved using a linear tax rate.

3.10.1 Augmented Planning Problem

The planner chooses an allocation $\{c(\theta), y(\theta), K(\theta), n(\theta)\}_{\theta\in\Theta}$ subject to both the same information constraints as before, and also behavioral constraints that describe how a household makes choices when confronted with a linear tax on child expenditures. Following Boadway and Jacobs (2014), it is useful to consider a household choosing consumption and child expenditures for given choices of income and child quantity,

$$\max_{c,K} u(c) - \phi\left(\frac{y}{\theta}\right) + v(n, K, \theta)$$

subject to

$$c + (1 + \tau_K) K = y - T(y, n)$$

where T is a tax function that depends only on income and child quantity, and arises from the implementation of the solution to the planning problem.

Solving this problem yields conditional demand functions, $\tilde{c}(\theta, y, n, \tau_K)$ and $\tilde{K}(\theta, y, n, \tau_K)$, that specify how a household will choose her consumption and child

expenditures for given levels of labor productivity, income, child quantity, and the tax rate on child expenditures. The planner must take this behavior into account when solving the planning problem. Specifically, the planner chooses an allocation $A = \{c(\theta), y(\theta), n(\theta), K(\theta)\}_{\Theta}$ and a tax rate τ_K that solve

$$\max_{A,\tau_K} \int_{\Theta} \alpha(\theta) U(\theta) f(\theta) d\theta$$

subject to

$$\begin{aligned}
U(\theta) &= u(c(\theta)) - \phi\left(\tfrac{y(\theta)}{\theta}\right) + v(n(\theta), K(\theta), \theta) \\
\theta &= \arg\max_{\theta'} U(\theta) \\
c(\theta) &= \tilde{c}(\theta, y, n, \tau_K) \\
K(\theta) &= \tilde{K}(\theta, y, n, \tau_K) \\
\int_{\Theta} (c(\theta) + K(\theta)) f(\theta) d\theta &\leq \int_{\Theta} y(\theta) f(\theta) d\theta
\end{aligned}$$

3.10.2 Quantitative Results

The optimal linear tax is $\tau_K^* = -0.3$, so that child expenditure is subsidized, as in the non-linear case. This number is also very close to the mean subsidy rate in the unrestricted model, which turns out to be -0.29. Figure 3.4 plots the optimal allocations in the cases when taxes are left unrestricted and when the tax on child expenditures is constrained to be linear. It is clear that the allocations are very similar even when non-linear taxation of child expenditures is ruled out, indicating that a large fraction of the welfare gains from taxing these choices can be achieved through the simple policy of a linear subsidy on child expenditures, accompanied by

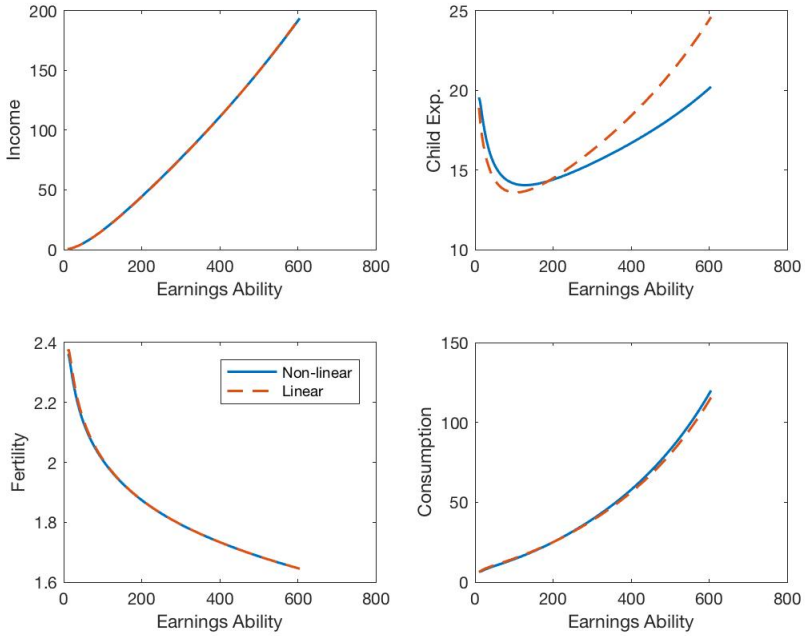


Figure 3.4: Optimal Allocations

non-linear taxation of income and child quantity. Indeed, comparing social welfare under each regime indicates that the restricted taxes obtain 99% of the welfare gain from non-linear taxation.

3.11 Conclusion

I have presented a general framework for analyzing the optimal tax treatment of choices related to fertility. My model is the first to adopt a truly Mirrleesian approach to fertility taxation, and as such offers multiple advantages over the existing literature. First, in the face of a plethora of potential underlying mechanisms governing the link between fertility choices and other household choices such as income, my model delivers expressions for optimal wedges that are robust to a wide range of mechanisms by expressing wedges in terms of reduced form elasticities that are estimable given the relevant data. Second, my expressions give clear intuition for the shape and sign of

the optimal distortions on child quantity and investments in goods and time, and how these properties relate to the elasticities. Finally, the quantitative exercise indicates that large welfare gains are achievable by taxing child-related choices in an optimal manner, and that most of these gains are obtainable using feasible policies.

A natural limitation of this study stems from the Mirrleesian approach I take towards analyzing optimal taxes. In the model I present, all choice variables are deterministic and monotonic functions of the household type parameter θ. This means that child quantity is a deterministic and decreasing function of household income. While this relationship captures the general negative correlation between income and fertility well, it imposes too much structure on the data that the model generates, since it becomes impossible for some high income households to have as many children as some lower income households, an outcome that is certainly observed in the data. Further, reinterpreting the variable n as the average number of children in households of a fixed type is unsatisfactory since it would make the corresponding taxes a function of average child quantity, which then becomes difficult to reconcile with real policies. A possible solution to this issue would be to let child quantity be a random variable that the household can only partially control, for example by choosing the mean value. This might reflect the uncertainty associated with child conception and would certainly allow the model to generate more realistic relationships between child quantity and other outcomes. The taxes could also then be adjusted so that taxes are a function of number of children actually born rather than the household's choice of the mean value. This is an interesting extension that I hope to pursue in later work.

References

Alvarez, F.: 1999, Social Mobility: The Barro-Becker Children Meet the Laitner-Loury Dynasties, *Rev. Econ. Dyn.* **2**(1), 65–103.

Angrist, J. D., Lavy, V. and Schlosser, A.: 2010, Multiple experiments for the causal link between the quantity and quality of children, *J. Labor Econ.* **28**(4), 773–824.

Attanasio, B. O., Low, H. and Virginia, S.: 2008, Explaining Changes in Female Labor Supply in a Life-Cycle, *Am. Econ. Rev.* **98**(4), 1517–1552.

Balestrino, A., Cigno, A. and Pettini, A.: 2002, Endogenous Fertility and the Design of Family Taxation, *Int. Tax Public Financ.* **9**(2), 175–193.

Baughman, R. and Dickert-Conlin, S.: 2009, The earned income tax credit and fertility, *J. Popul. Econ.* **22**(3), 537–563.

Becker, G. S. and Lewis, H. G.: 1974, Interaction between Quantity and Quality of Children, *NBER Chapters* **I**, 81–90.

Becker, G. S., Murphy, K. M. and Tamura, R.: 2016, Human Capital, Fertility, and Economic Growth, **98**(5).

Black, S. E., Devereux, P. J. and Salvanes, K. G.: 2005, The More the Merrier? The Effect of Family Size and Birth Order on Children's Education*, *Q. J. Econ.* **120**(2), 669–700.

Blundell, R., Pistaferri, L. and Saporta-Eksten, I.: 2015, Children, Time Allocation and Consumption Insurance, **284024**(March), 1–38.

Boca, D. D. and Wiswall, M.: 2010, Household Choices and Child Development, *Work. Pap.* (5155), 1–56.

Boca, D. D. and Wiswall, M.: 2014, Transfers to Households with Children and Child Development, (8393).

Cigno, A.: 2001, Comparative advantage, observability, and the optimal tax treatment of families with children, *Int. Tax Public Financ.* **8**(4), 455–470.

Cigno, A., Luporini, A. and Pettini, A.: 2003, Transfers to families with children as a principal-agent problem, *J. Public Econ.* **87**(5-6), 1165–1177.

Cigno, A. and Pettini, A.: 2002, Taxing family size and subsidizing child-specific commodities?, *J. Public Econ.* **84**(1), 75–90.

Córdoba, J. C. and Ripoll, M.: 2014, The Elasticity of Intergenerational Substitution, Parental Altruism, and Fertility Choice, (15).

Cremer, H., Dellis, A. and Pestieau, P.: 2003, Family size and optimal income taxation, *J. Popul. Econ.* **16**(1), 37–54.

Cremer, H., Gahvari, F. and Pestieau, P.: 2006, Pensions with endogenous and stochastic fertility, *J. Public Econ.* **90**(12), 2303–2321.

Cremer, H., Gahvari, F. and Pestieau, P.: 2011, Fertility, human capital accumulation, and the pension system, *J. Public Econ.* **95**(11-12), 1272–1279.

Cunha, F., Heckman, J. and Schennach, S.: 2010, Estimating the Technology of Cognitive and Noncognitive Skill Formation, *Econometrica* **78**(3), 883–931.

D'Addio, A. C. and Mira d'Ercole, M.: 2005, Policies, Institutions and Fertility Rates: a Panel Data Analysis for OECD Countries, (41), 7–43.

Dahl, G. B. and Lochner, L.: 2012, The Impact of Family Income on Child Achievement : Evidence from the Earned Income Tax Credit, **102**(5), 1927–1956.

De La Croix, D., Doepke, M., Azariadis, C., Cole, H., Farmer, R., Hansen, G., Ohanian, L. and Schneider, M.: 2002, Inequality and Growth: Why Differential Fertility Matters.

Department of Health & Human Services: 2016, 2016 Federal Poverty Level Chart, *Technical report*.

Diamond, P. A.: 1998, Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates, *Am. Econ. Rev.* **88**(1), 83–95.

Guryan, J., Hurst, E. and Kearney, M.: 2008, Parental Education and Parental Time with Children, *Source J. Econ. Perspect.* **22**(3), 23–46.

Hanushek, E.: 1992, The Trade-Off between Child Quality and Quantity.

Ho, C. and Pavoni, N.: 2016, Efficient child care subsidies, *Work. Pap.* .

Hotz, V. J., Klerman, J. A. and Willis, R. J.: 1997, The Economics of Fertility in Developed Countries: a Survey, *Handb. Popul. Fam. Econ.* (7), 275–347.

Hoynes, H. and Rothstein, J.: 2016, Tax Policy Toward Low-Income Families.

Jacobs, B. and Boadway, R.: 2014, Optimal linear commodity taxation under optimal non-linear income taxation, *J. Public Econ.* **117**, 201–210.

Jones, L. E. and Schoonbroodt, A.: 2008, Fertility Theories: Can They Explain the Negative Fertility-Income Relationship?

Lino, M.: 2014, Expenditures on Children by Families, 2013, (1528).

Marr, C., Huang, C.-c., Sherman, A. and Debot, B.: 2015, EITC and Child Tax Credit Promote Work, Reduce Poverty and Support Children's Development, **000**, 23–27.

Mirrlees, J.: 1971, Exploration in optimum income the theory, *Rev. Econ. Stud.* **38**(2), 175–208.

OECD: 2017, Public spending on family benefits Public, *Technical report*.

Rosenzweig, M. R. and Zhang, J.: 2009, Do population control policies induce more human capital investment? Twins, birth weight and China's "one-child" policy, *Rev. Econ. Stud.* **76**(3), 1149–1174.

Saez, E.: 2001, Using Elasticities to Derive Optimal Income Tax Rates, *Rev. Econ. Stud.* pp. 205–229.

Sommer, K.: 2016, Fertility choice in a life cycle model with idiosyncratic uninsurable earnings risk, *J. Monet. Econ.* **83**, 27–38.

Stantcheva, S.: 2015, Optimal Taxation and Human Capital Policies over the Life Cycle, *Work. Pap.* .

Thevénon, O.: 2011, Family Policies in OECD countries: A comparative Analysis., *Popul. Dev. Rev.* **37**(1), 57–87.

Werning, I.: 2011, Nonlinear Capital Taxation, **2011**, 1–20.

Appendix

3.A The Planning Problem

Recall the baseline planning problem using the First Order Approach:

$$\max_{A} \int_{\Theta} \alpha(\theta) U(\theta) f(\theta) d\theta$$

subject to

$$U(\theta) = u(c(\theta)) - \phi\left(\frac{y(\theta)}{\theta}, H(\theta)\right) + v(n(\theta), K(\theta), H(\theta), \theta)$$

$$U_{\theta}(\theta) = \frac{y(\theta)}{\theta^2}\phi_l\left(\frac{y(\theta)}{\theta}, H(\theta)\right) + v_{\theta}(n(\theta), K(\theta), H(\theta), \theta)$$

$$\int_{\Theta} (c(\theta) + K(\theta)) f(\theta) d\theta \leq \int_{\Theta} y(\theta) f(\theta) d\theta$$

This is a standard optimal control problem and so can be solved using Hamiltonian techniques: substitute the first constraint into the aggregate resource constraint replacing the $c(\theta)$ term, and attach multipliers $\mu(\theta)$ to the differential equation

constraint, and λ to the resource constraint. The Hamiltonian then reads

$$\begin{aligned}\mathcal{H} &= \alpha(\theta)U(\theta)f(\theta) \\ &+ \lambda(y(\theta)-K(\theta))f(\theta) \\ &- \lambda u^{-1}\left(U(\theta)+\phi\left(\tfrac{y(\theta)}{\theta},H(\theta)\right)-v\left(n(\theta),K(\theta),H(\theta),\theta\right)\right)f(\theta) \\ &+ \mu(\theta)\left(\tfrac{y(\theta)}{\theta^2}\phi_l\left(\tfrac{y(\theta)}{\theta},H(\theta)\right)+v_\theta\left(n(\theta),K(\theta),H(\theta),\theta\right)\right)\end{aligned}$$

Adopting abbreviated notation for derivatives, the First Order Conditions (FOCs) are as follows:

$$\begin{aligned}U(\theta): &\quad \alpha(\theta)f(\theta)-\lambda\tfrac{1}{u'(\theta)}f(\theta) &&= -\mu_\theta(\theta) \\ y(\theta): &\quad \lambda\left(1-\tfrac{\phi_l(\theta)}{\theta u'(\theta)}\right)f(\theta)+\mu(\theta)\left(\tfrac{1}{\theta^2}\phi_l(\theta)+\tfrac{y(\theta)}{\theta^3}\phi_{ll}(\theta)\right) &&= 0 \\ n(\theta): &\quad \lambda\left(\tfrac{\beta v_n(\theta)}{u'(\theta)}\right)f(\theta)+\mu(\theta)v_{n\theta}(\theta) &&= 0 \\ K(\theta): &\quad \lambda\left(\tfrac{\beta v_K(\theta)}{u'(\theta)}-1\right)f(\theta)+\mu(\theta)v_{\theta K}(\theta) &&= 0 \\ H(\theta): &\quad \lambda\left(\tfrac{\beta v_H(\theta)-\phi_H(\theta)}{u'(\theta)}\right)f(\theta)+\mu(\theta)v_{\theta H}(\theta) &&= 0\end{aligned}$$

Noting that $\lim_{\theta\to+\infty}\mu(\theta)=\lim_{\theta\to 0}\mu(\theta)=0$, and recalling that $\mathbb{E}[\alpha(\theta)]=1$, integrating the first condition over all of Θ yields $\frac{1}{\lambda}=\mathbb{E}\left[\frac{1}{u'(c(\theta))}\right]$. Similarly, integrating from any θ to $+\infty$ yields $\mu(\theta)=-\int_\theta^\infty\left(\frac{\lambda}{u'(c(v))}-\alpha(v)\right)f(v)\,dv$. Given this, define $\eta(\theta)=-\mu(\theta)/(1-F(\theta))\lambda$. To derive the elasticity terms, consider a household of

ability θ facing wedges τ_y, τ_n, τ_K, and τ_H. From the main text, I have that

$$\begin{aligned}
\tau_y &= 1 - \tfrac{\phi_l(y/\theta,H)}{\theta u'(c)} \\
\tau_n &= \tfrac{v_n(n,K,H,\theta)}{u'(c)} \\
\tau_K &= \tfrac{v_K(n,K,H,\theta)}{u'(c)} - 1 \\
\tau_H &= \tfrac{v_H(n,K,H,\theta) - \phi_H(y/\theta,H)}{u'(c)}
\end{aligned}$$

In each case, I now consider the partial derivative of variable $j \in (y, n, K, H)$ with respect to θ and τ_j, holding all other terms fixed.

For income, it is useful to recall that $y = \theta l$, and treat l as the variable of interest. Following Stantcheva (2014), I then define

$$\epsilon^u(\theta) = \frac{\phi_l(l,H)/l + \theta^2 u''(c)}{\phi_{ll}(l,H) - \theta^2 u''(c)}$$

$$\epsilon^c(\theta) = \frac{\phi_l(l,H)/l}{\phi_{ll}(l,H) - \theta^2 u''(c)}$$

which immediately lead to the expressions in the text, noting that $\epsilon^u = \epsilon^c$ when $u'' = 0$.

In the case of n, I differentiate the household FOC with respect to θ,

$$v_{nn}(\theta)\frac{\partial n}{\partial \theta} + \beta v_{n\theta}(\theta) = \tau_n u''(c)\frac{\partial c^n}{\partial \theta}$$

where

$$\frac{\partial c^n}{\partial \theta} = l(1-\tau_y) - \tau_n \frac{\partial n}{\partial \theta}$$

Similarly for τ_n,

$$v_{nn}(\theta)\frac{\partial n}{\partial \tau_n} = u'(c) + \tau_n u''(c)\frac{\partial c^n}{\partial \tau_n}$$

where

$$\frac{\partial c^n}{\partial \tau_n} = -n - \tau_n \frac{\partial n}{\partial \tau_n}$$

Combining these expressions yields

$$\theta v_{n\theta}(\theta) = \theta\tau_n u''(c)\frac{\partial c^n}{\partial \tau_n}\left(\left(\frac{\partial c^n}{\partial \theta}/\frac{\partial c^n}{\partial \tau_n}\right) - \left(\frac{\partial n}{\partial \theta}/\frac{\partial n}{\partial \tau_n}\right)\right) - \theta u'(c)\frac{\partial n}{\partial \theta}/\frac{\partial n}{\partial \tau_n}$$

which can be expressed in terms of the relevant elasticities,

$$\theta v_{n\theta}(\theta) = u'(c)\left(-\epsilon^n_\theta/e^n_{\tau_n} - \left(-\frac{u''(c)\,c}{u'(c)}\right)\epsilon^{c,n}_{\tau_n}\left(\left(\epsilon^{c,n}_\theta/e^{c,n}_{\tau_n}\right) - \left(\epsilon^n_\theta/e^n_{\tau_n}\right)\right)\right)$$

Substituting into the FOC of the program yields the expression for the optimal wedge.

In the case of K, first differentiate with respect to θ,

$$v_{KK}(\theta)\frac{\partial K}{\partial \theta} + \beta v_{\theta K}(\theta) = (1+\tau_K)u''(c)\frac{\partial c^K}{\partial \theta}$$

and then with respect to $1+\tau_K$,

$$_{KK}(\theta)\frac{\partial K}{\partial(1+\tau_K)} = u'(c) + (1+\tau_K)u''(c)\frac{\partial c}{\partial(1+\tau_K)}$$

Combining yields

$$\beta\theta v_{\theta K}(\theta) = u'(c)\left(-\frac{\epsilon^K_\theta}{e^K_{\tau_K}} - \left(\frac{-u''(c)\,c}{u'(c)}\right)\epsilon^{c,K}_{1+\tau_K}\left(\frac{\epsilon^{c,K}_\theta}{e^{c,K}_{\tau_K}} - \frac{\epsilon^K_\theta}{e^K_{\tau_K}}\right)\right)$$

which can be substituted into the relevant FOC of the program to yield the expression for the optimal wedge.

Finally, consider H. Differentiating with respect to θ,

$$v_{HH}(\theta)\frac{\partial H}{\partial \theta}+\beta v_{\theta H}(\theta)-\phi_{HH}(\theta)\frac{\partial H}{\partial \theta}=\tau_H u''(c)\frac{\partial c^H}{\partial \theta}$$

and then with respect to τ_H,

$$v_{HH}(\theta)\frac{\partial H}{\partial \tau_H}-\phi_{HH}(\theta)\frac{\partial H}{\partial \tau_H}=u'(c)+\tau_H u''(c)\frac{\partial c^H}{\partial \tau_H}$$

Combining yields

$$\theta v_{\theta H}(\theta)=u'(c)\left(-\frac{\epsilon_\theta^H}{e_{\tau_H}^H}-\left(-\frac{u''(c)c}{u'(c)}\right)\epsilon_{\tau_H}^{c,H}\left(\frac{\epsilon_\theta^{c,H}}{e_{\tau_H}^{c,H}}-\frac{\epsilon_\theta^H}{e_{\tau_H}^H}\right)\right)$$

which can be substituted into the relevant FOC of the program to yield the expression for the optimal wedge.

3.B Proof of Proposition 9

Proof. First consider a bundle (y,K,H,n) not in the optimal allocation. By construction, the tax payable by a household of type θ is at least the tax payable under the type specific tax system so that the household is weakly better off by choosing her bundle in the optimal allocation. Now, consider a bundle in the optimal allocation, $(y(\theta'),K(\theta'),H(\theta'),n(\theta'))$ for some $\theta'\in\Theta$. I claim that

$$T^*(y(\theta'),K(\theta'),H(\theta'),n(\theta'))=T(y(\theta'),K(\theta'),H(\theta'),n(\theta'),\theta')$$

i.e. that the θ' type household faces the largest tax bill at the θ'-bundle in the optimal allocation. By incentive compatibility of the optimal allocation, I know that

for another type $\theta'' \neq \theta'$

$$u\left(y\left(\theta'\right) - K\left(\theta'\right) - T\left(y\left(\theta'\right), K\left(\theta'\right), H\left(\theta'\right), n\left(\theta'\right), \theta'\right)\right)$$

$$-\phi\left(\frac{y(\theta')}{\theta''}, H\left(\theta'\right)\right) + v\left(n\left(\theta'\right), K\left(\theta'\right), H\left(\theta'\right), \theta''\right) \qquad \leq \quad U\left(\theta''\right)$$

where

$$T\left(y\left(\theta'\right), K\left(\theta'\right), H\left(\theta'\right), n\left(\theta'\right), \theta'\right) = y\left(\theta'\right) - K\left(\theta'\right) - c\left(\theta'\right)$$

by construction of the type specific tax system for θ'. Therefore, to make any θ'' type household indifferent between this bundle and her optimal allocation bundle, it must be that

$$T\left(y\left(\theta'\right), K\left(\theta'\right), H\left(\theta'\right), n\left(\theta'\right), \theta''\right) \leq T\left(y\left(\theta'\right), K\left(\theta'\right), H\left(\theta'\right), n\left(\theta'\right), \theta'\right)$$

thus establishing the claim. In particular, this means that any household of a type $\theta'' \neq \theta'$ is weakly better off choosing her own bundle in the optimal allocation, which completes the implementation proof. □

3.C Policy Comparisons

Figures 3.5 and 3.6 compare the optimal credits from the model to the policies in the US for the cases of three children and one child respectively. The patterns are similar to the case of two children, as described above.

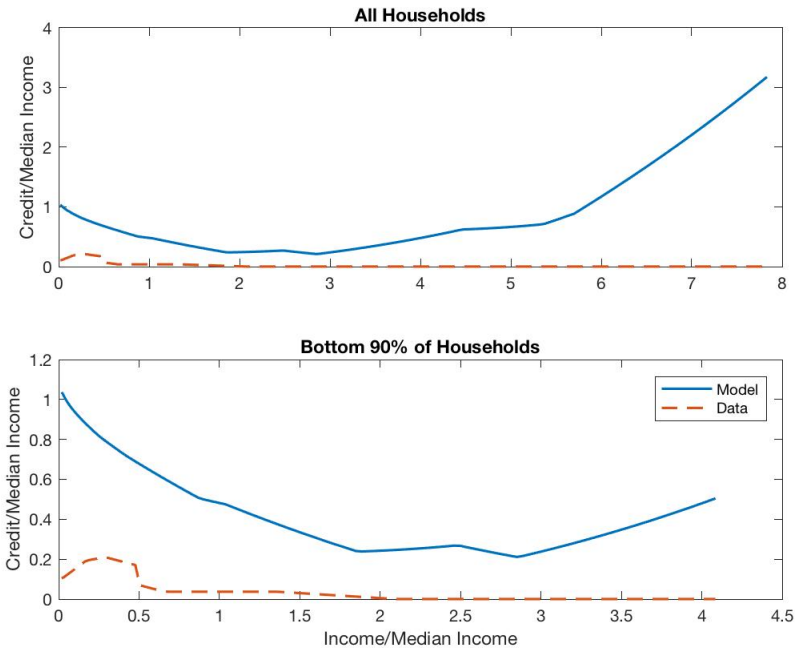


Figure 3.5: Tax Credits for households with three children

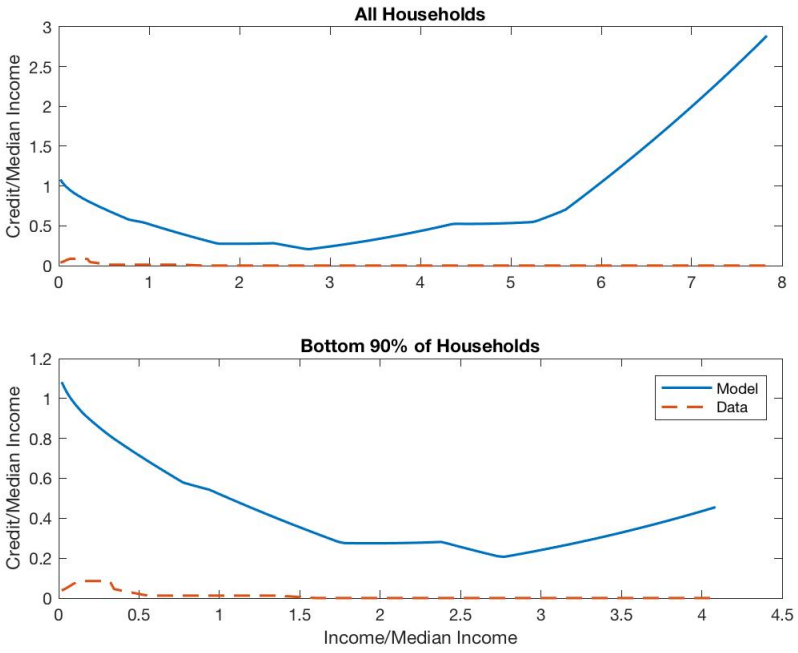


Figure 3.6: Tax Credits for households with one child