Essays in Macroeconomics and Labor Supply

by

Preston W. Mui

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Benjamin Schoefer, Chair
Professor Jón Steinsson
Professor Yuriy Gorodnichenko

Spring 2022

Essays in Macroeconomics and Labor Supply

Abstract

Essays in Macroeconomics and Labor Supply

by

Preston W. Mui

Doctor of Philosophy in Economics

University of California, Berkeley

Professor Benjamin Schoefer, Chair

In this dissertation, I study the role of labor supply in macroeconomic fluctations and the movement of employment in response to these fluctuations. The first chapter is a theoretical and empirical study of the role of firm-specific labor supply in amplifying business cycles. The second chapter focuses on measuring the aggregate labor supply elasticity at the extensive margin, using a novel survey approach. Finally, in the third chapter I measure the effects of government policies in the early stages of the COVID-19 pandemic on employment, using decentralized implementation of these policies.

In the first chapter, I assess the role of labor market monopsony—finitely-elastic firm-specific labor supply—in the context of a New Keynesian model. First, I modify a basic New Keynesian model to include firm-specific labor and calibrate the labor supply elasticities to micro-empirical estimates. Consistent with this mechanism serving as a source of real rigidity, firm-specific labor substantially reduces the slope of the Phillips curve relative to the perfectly competitive labor market benchmark. However, this depends strongly on the elasticity chosen, and requires distinguishing the firm-specific and aggregate labor supply elasticities, which previous work often fails to do. Second, I provide a cross-sectional empirical test for this mechanism. I estimate the firm-specific labor supply elasticity by industry in the Survey of Income and Program Participation using a dynamic monopsony model. I then estimate industry responses to monetary policy shocks. Contrary to the New Keynesian model, I find no evidence that industry differences in firm-specific labor supply elasticities lead to different industry price responses to monetary policy shocks. My results do not support the theory that firm-specific labor is a source of real rigidity.

The second chapter is an innovative investigation into measuring the aggregate labor supply curve using survey methods. I measure extensive-margin labor supply (employment) preferences in two representative surveys of the U.S. and German populations. In the survey, I elicit "reservation raises": the percent wage change that renders a given

individual indifferent between employment and nonemployment. It is equal to their reservation wage divided by their actual, or potential, wage. The reservation raise distribution is the nonparametric aggregate labor supply curve. Locally, the curve exhibits large short-run elasticities above 3, consistent with business cycle evidence. For larger upward shifts, arc elasticities shrink towards 0.5, consistent with quasi-experimental evidence from tax holidays. Existing models fail to match this nonconstant, asymmetric curve.

Finally, in the third chapter, I investigate the labor market ramifications of government-imposed lockdowns in the early stages of the COVID-19 pandemic. I use the high-frequency, decentralized implementation of Stay-at-Home orders in the U.S. to disentangle the labor market effects of SAH orders from the general economic disruption wrought by the COVID-19 pandemic. I find that each week of SAH exposure increased a state's weekly initial unemployment insurance (UI) claims by 1.9% of its employment level relative to other states. A back-of-the-envelope calculation implies that, of the 17 million UI claims between March 14 and April 4, only 4 million were attributable to SAH orders. I present a currency union model to provide conditions for mapping this estimate to aggregate employment losses.

Contents

List of Figures

List of Tables

Acknowledgments

The first chapter of my dissertation is adapted from my paper "Labor Market Monopsony in the New Keynesian Model: Theory and Evidence." I thank Joshua Archer for excellent research assistance. I thank David Berger, Sydnee Caldwell, Ryan Chahrour, Yuriy Gorodnichenko, Pierre-Olivier Gourinchas, Chen Lian, Peter McCrory, Emi Nakamura, David Romer, Benjamin Schoefer, Jón Steinsson, Jacob Weber, Vitalia Yaremko, and Chaoran Yu for helpful and insightful comments. I thank seminar participants at the UC-Berkeley GEMS Conference. I particularly thank the members of my dissertation committee, Benjamin Schoefer, Yuriy Gorodnichenko, and Jón Steinsson, for their guidance.

The second chapter of my dissertation is adapted from "Reservation Raises: The Aggregate Labor Supply Curve at the Extensive Margin," coauthored with my advisor and dissertation commmittee chair, Benjamin Schoefer, who has granted permission for its inclusion here. We are grateful to referees of the previous version for constructive feedback, which has greatly improved the paper. We thank Chaoran Yu for excellent research assistance. We thank Mark Aguiar, Mark Bils, Gabriel Chodorow-Reich, Shigeru Fujita, Adam Guren, Greg Kaplan, Loukas Karabarbounis, Marios Karabarbounis, David Lagakos, Ellen McGrattan, Monika Merz, Luigi Pistaferri, Richard Rogerson, Ayşegül Şahin, Sergio Salgado, and Gianluca Violante for useful comments. We thank seminar participants at Amsterdam/Tinbergen, Bonn, the Central Bank of Chile, Edinburgh, the Federal Reserve Bank of Philadelphia, the Federal Reserve Board of Governors, New York University, Princeton, Stanford/SIEPR, the Virtual Australian Macroeconomics Seminar, UC Berkeley, UC San Diego, and VATT. We thank conference audiences at the Bank of Italy-CEPR Labour workshop, the ifo Conference on Macroeconomics and Survey Data, the NBER Dynamic Equilibrium Models meeting, the NBER Summer Institute Micro Data and Macro Models meeting, the NBER Summer Institute Impulse and Propagation Mechanisms meeting, the Society of Economic Dynamics Annual Meeting, and the West Coast Matching Workshop (Federal Reserve Bank of San Francisco). For financial support of our custom survey, we thank the Clausen Center and the Institute for Research on Labor and Employment at UC Berkeley, and SIEPR at Stanford. We thank the survey design teams of NORC at UChicago, and GSOEP/Kantar. The survey received IRB approval from UC Berkeley (protocol number 2019-02-11815).

The third chapter of my dissertation is a reprint of "Unemployment Effects of Stay-At-Home Orders: Evidence from High Frequency Claims Data," coauthored with my officemates and dear friends ChaeWon Baek, Peter McCrory, and Todd Messer — all of whom have granted permission for its inclusion here. This article has been published at the *Review of Economics and Statistics*, which retains first publication credit, © 2020 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology. All of us are grateful to Christina Brown, Bill Dupor, Yuriy Gorodnichenko, Christina Romer, Maxim Massenkoff, and Benjamin Schoefer, as well as the editor and two anonymous referees, for helpful comments and suggestions.

I am incredibly thankful for the personal and emotional support of all of the friends, family, and colleagues that have supported me throughout my studies. A complete enumeration of the people who have supported me at UC Berkeley would be too long to list here. First and foremost I thank my parents, Kim-ping Mui and Wing-yin Mui, for everything they have done for me not only during my graduate studies but also my entire life. Among my friends and colleagues, in particular I would like to highlight the support of Sam Leone, Peter McCrory, Christina Brown, Todd Messer, and Chaewon Baek, who were other graduate students in my cohort. I am proud to call them colleagues, coauthors, and most importantly, friends.

I would like to thank the staff of of the UC Berkeley Economics Department. Their professionalism is unparalleled, and their underrated work in the Department allows me, as well as the rest of the faculty and graduate students, to focus on our research. I would also like to thank my union, UAW 2865, whose tireless work on behalf of graduate student workers at the University of California is invaluable. I also gratefully acknowledge the financial support of the Clausen Center and the Institute for Research on Labor and Employment at UC Berkeley.

Throughout the years spent conducting research for this dissertation, numerous faculty and graduate students at the UC Berkeley Economics Department provided kind and helpful feedback and suggestions on the work presented in this dissertation. I would to especially thank members of my committee. Yuriy Gorodnichenko provided particularly helpful advice on what empirical tests to conduct, and Jón Steinsson provided crucial feedback on how to frame my work in the larger context of macroeconomic research.

Finally, I am especially thankful to my committee chair, Benjamin Schoefer, who was exceedingly patient and generous during his time as my advisor. His work ethic, brilliance, and most importantly kindness humbles me greatly. He demonstrates the best of what an economic researcher can be, both in his amazing work and his interactions with others. I am incredibly lucky to have had him as my coauthor, advisor and mentor.

Any and all mistakes remain mine.

Chapter 1

Labor Market Monopsony in the New Keynesian Model: Theory and Evidence

1.1 Introduction

Models of business cycles often require strong real rigidities—forces that reduce the responsiveness of firms' desired real prices to change in aggregate demand—to explain short-run fluctuations (Ball and Romer, 1990). Previous work has theorized that firm-specific, rather than homogeneous, labor is a potentially powerful source of real rigidity in New Keynesian models. Firm-specific labor theoretically induces real rigidity because it steepens the marginal cost function of the firm, especially if the labor supply elasticity to the firm is low. At the same time, there is a substantial literature attempting to measure firm-specific labor supply elasticities, often finding that firms exercise considerable monopsony power in the labor market. Despite this, existing work in macroeconomic theory lacks a connection to the micro-empirical estimates of the firm-specific labor supply elasticity (the elasticity of labor supplied to an individual firm to the wage paid by the firm). In addition, there is scant direct empirical evidence, for or against, the existence of this real rigidity mechanism. In this paper, I provide a comprehensive and dedicated theoretical and empirical treatment of the role that firm-specific labor plays in generating real rigidity in New Keynesian models.

First, I clarify the theoretical role of firm-specific labor supply by constructing and calibrating a simple New Keynesian model with firm-specific labor supply. Importantly, I draw a distinction between the aggregate and firm-specific labor supply elasticities, and calibrate both elasticities to values consistent with the micro-empirical literature on labor supply elasticities. This distinction is important because the two elasticities play *opposing* roles in determining the extent of real rigidity in a model. A higher aggregate labor supply elasticity increases real rigidity, while a higher firm-specific labor supply elasticity lowers real rigidity.

Second, I use cross-sectional industry variation in the firm-specific labor supply elasticity to empirically evaluate the firm-specific labor real rigidity mechanism. I use dynamic monopsony methods to estimate firm-specific labor supply elasticities at the industry level, and estimate responses of industry variables to monetary policy shocks using local projections. I do not find any evidence that high-elasticity industries exhibit larger price responses to monetary policy shocks. I also find that high-elasticity industries exhibit larger employment falls in response to monetary policy shocks, contrary to the model's predictions. This evidence casts doubt on the notion that firm-specific labor is a significant source of real rigidity.

Previous work exploring the modeling implications of firm-specific labor has suggested that firm-specific labor is a potentially important source of real rigidity. Woodford (2003) and Woodford (2005) show that in a simple monetary model, factor specificity matters more than both variable optimal markups and intermediate inputs in terms of generating real rigidities. Matheron (2006) estimates Phillips curves on Euro Area data and shows that modeling labor markets as firm-specific yields estimates of price reset probabilities that are much more consistent with empirical estimates of price reset frequencies than estimates based on models without firm-specific labor. Carvalho and Nechio (2016) analyze New Keynesian models with varying degrees of labor specificity (economy-wide, sector-specific, and firm-specific) and find that firm- and sector-specific labor induces greater strategic complementarity and flatter Phillips curves than economy-wide labor markets.[1]

The intuition for how this mechanism operates is as follows. Consider a firm evaluating a given price increase in response to an increase in aggregate demand. The price increase would decrease demand for the firm's output, decreasing the firm's marginal cost due to, e.g., decreasing returns to scale. This decrease in marginal cost attenuates the firm's desire to raise its price—real rigidity. This attenuation is stronger when a firm faces a labor supply curve that is less elastic, since the firm's marginal cost curve is steeper with respect to its own output. The role of firm-specific labor supply is opposite that of aggregate labor supply. The aggregate labor supply elasticity affects how the aggregate wages, and thus all firm's marginal costs, respond to changes in *aggregate* output as opposed to a firm's *own* output. The aggregate labor supply is an example of what Leahy (2011) calls "type 1" real rigidity, which encompasses model features that affect the response of marginal cost to aggregate output. By contrast, firm-specific labor is a type of "type 2" real rigidity, since it affects how a firm chooses its optimal price in response to changes in its marginal cost. A lower aggregate labor supply elasticity leads to weaker real rigidities, whereas lower firm-specific labor supply elasticities lead to stronger real rigidities.

The different roles of the elasticities means that it is important to distinguish between

[1]These papers are part of a larger literature that incorporates firm-specific factor markets into the New Keynesian model to better match inflation dynamics. In particular, the role of firm-specific capital has been used by, e.g. Altig et al. (2011) and Woodford (2005). The intuition behind the two mechanisms is similar—both steepen the relationship between a firm's marginal cost and its own output. Matheron (2006) finds that estimating New Keynesian Phillips curves with firm-specific labor alone yields estimates of price reset probabilities that are consistent with micro empirical estimates of these probabilities, but firm-specific capital alone does not.

the two elasticities in a modeling context, including calibration. Much of the theoretical work on firm-specific labor has centered on where to put "the" labor elasticity; i.e., whether to aggregate labor supply of different types before or after applying the labor disutility transformation.[2] However, as I will explore in this paper, using one elasticity to play both roles risks overstating the impact of firm-specific labor in generating real rigidity, especially at higher elasticities. This is because if (as intuition might suggest) firm-specific elasticities are higher than aggregate elasticities, using the high firm-specific elasticity calibration for both roles induces a significant amount of real rigidity through the "type 1" aggregate labor supply channel. A firm-specific setup with a single elasticity is a feature of many papers that study or use models that feature firm-specific labor, (e.g. Nakamura and Steinsson, 2014; Gorodnichenko and Weber, 2016; Carvalho and Nechio, 2016).

There exists a substantial microempirical literature measuring both elasticities. Micro estimates of the aggregate labor supply elasticity generally find very low elasticities; in a meta-study, Chetty et al. (2012) finds average aggregate Frisch labor supply elasticities of 0.32 (extensive margin) and 0.86 (aggregate hours), although Mui and Schoefer (2021) find much higher local elasticities, around 3, using a survey-based approach. As for the firm-specific labor supply elasticity, the literature is too large to fully enumerate here—in a meta-analysis, Sokolova and Sorensen (2020) collect 1,320 estimates from 53 studies. Generally speaking, estimates of the firm-specific labor supply elasticity are higher than those of the aggregate labor supply elasticity.[3] In particular, elasticity estimates using the dynamic monopsony methods from Manning (2013) find particularly low elasticities. Using LEHD data, Webber (2015) finds an average firm-specific labor supply elasticity of 1.08. Using corporate income tax changes as a source of identification, Berger, Herkenhoff and Mongey (2021) find short-run labor supply elasticities between 1 and 2, depending on the firm's labor market share, while Bassier, Dube and Naidu (2020) find an elasticity of labor supply to firm AKM fixed effects of 3 using matched employer-employee data from Oregon. These estimates are well in the range of elasticities that would generate a substantial amount of real rigidity, relative to a perfectly competitive labor market, but are still above that of most estimates of the aggregate labor supply elasticity.

If firm-specific labor supply elasticities are low, and firms exercise monopsony power in the labor market, the real rigidity mechanism is theoretically very powerful. In the first part of the paper, I show that modifying the simple New Keynesian model from Galí (2008) to include finite firm-specific labor supply elasticities (in addition to, not instead of, the aggregate labor supply elasticity) that are consistent with some empirical estimates can

[2]For example, in a separable utility function with labor types N_i and labor supply elasticity θ, the distinction would be whether to model labor disutility as $\left(\sum_i N_i\right)^{1+\theta}$ or $\left(\sum_i N_i^{1+\theta}\right)$. Note that in both examples, θ is the aggregate labor supply elasticity, whereas the firm-specific labor supply elasticity is infinite in the first case and θ in the second.

[3]There are a few studies in particular narrow contexts which yield estimates of the firm-specific labor supply elasticity which are extremely low. For example, using a natural experiment arising from Veterans Affairs compensation changes, Staiger et al. (2010) finds a short-run elasticity for nurses around 0.1. Dube et al. (2018) finds that labor supply for MTurk tasks are also around 0.1.

lower the slope of the Phillips curve by as much as three-quarters, substantially muting the inflation response and amplifying the output response to demand shocks. The strength of the mechanism depends heavily on the firm-specific labor supply elasticity the model is calibrated to and is much weaker at higher levels (holding the aggregate labor supply elasticity constant).

In light of the theoretical work arguing that firm-specific labor is a strong source of real rigidity and evidence of substantial monopsony power in the labor market, it is surprising that there is little empirical evidence, for or against, as to whether the mechanism actually exists. Previous evidence on this mechanism is indirect; for example, Matheron (2006) calibrates models with and without firm-specific labor and finds that the calibrations with firm-specific labor imply more realistic price reset frequencies. However, no existing work provides direct evidence as to whether or not this real rigidity mechanism is at play.

In the second part of this paper, I empirically investigate assess the existence and strength of the firm-specific real rigidity mechanism by using cross-sectional industry variation in firm-specific labor supply elasticities. A multi-sector version of the model with firm-specific labor predicts that sectors with different elasticities exhibit behavior analogous to that of economies with different firm-specific labor supply elasticities; that is, sectors with higher elasticities should experience larger price decreases and smaller output and employment falls in response to contractionary monetary policy shocks.

I estimate firm-specific labor supply elasticities at the industry level using the dynamic monopsony model from Manning (2013) and data from the Survey of Income and Program Participation (SIPP). Across industries, I find a median firm-specific labor supply elasticity of 1.45 which, compared to a perfectly competitive labor market, would induce a Phillips curve slope that is approximately one-third as steep. There is significant heterogeneity in these elasticities across sectors. The lowest elasticity, in NAICS code 316 ("Leather and Allied Product Manufacturing"), is 0.47; the highest elasticity, in NAICS code 491 ("Postal Services and Contractors"), is 2.93.

To test the mechanism, I estimate "differential impulse-response functions (IRFs)" of industry variables (prices, output, employment, and wages) to monetary policy shocks, which measure how much industry IRFs vary due to differences in firm-specific labor supply elasticities. I find no support for the theory that firm-specific labor supply generates real rigidities. Contrary to the predictions of New Keynesian theory, I do not find that the prices of industries with less elastic firm-specific labor supplies (i.e., less competitive sectors) fall less in response to a contractionary monetary policy shock; the firm-specific labor supply elasticity appears to have no detectable effect on industry price responses to contractionary monetary policy shocks. I also find no difference in output and wage responses across industries with different firm-specific labor supply elasticities. I do find that industries with larger firm-specific labor supply elasticities experience larger drops in employment, contrary to the predictions of my model. Overall, my results cast doubt on the firm-specific labor real rigidity mechanism.

This empirical strategy used in this paper is analogous to other work studying New Keynesian mechanics using cross-sectional variation. For example, there is a literature that uses cross-sectional variation in price change frequencies to assess the importance of nominal rigidities, such as Bils, Klenow and Kryvstov (2003), who compares responses to monetary policy shocks of goods with flexible and sticky prices, and Gorodnichenko and Weber (2016), who compares stock returns of firms with high and low frequency of price adjustment. The methodology of this section is closest to that of Dedola and Lippi (2005) and Henkel (2020), who construct industry-level impulse response functions and project them onto industry characteristics (but not the firm-specific labor supply elasticity).

The rest of the paper proceeds as follows. In Section 1.2, I construct and analyze the New Keynesian model with firm-specific labor. In Section 1.3, I estimate firm-specific labor supply elasticities and calibrate the multi-sector version of the model. I present the cross-sectional analysis of industry responses to monetary policy shocks in Section 1.4. Section 1.5 concludes.

1.2 New Keynesian Firm-Specific Labor Supply

In this section, I embed firm-specific labor supply into a standard New Keynesian model (Galí, 2008). The only modification I make to the model is to model labor services as firm-specific, rather than homogeneous. The representative household has CES preferences over the labor services provided to the firms, in addition to convex disutility with respect to aggregate labor supply. This setup allows me to draw a distinction between the aggregate labor supply elasticity and the firm-specific labor supply elasticity by having both parameters present in the model.

When calibrated to micro-empirical estimates of the firm-specific labor supply elasticity, the slope of the Phillips curve is half the slope of the Phillips curve in a model with homogeneous labor. The model with firm-specific labor exhibits smaller price responses and greater output responses to monetary policy shocks, relative to the model with homogeneous labor. However, this difference depends greatly upon the elasticity chosen.

New Keynesian Model with Firm-specific Labor Supply

Households. An infinitely-lived representative household maximizes

$$\max_{\{C_{it}\},\{L_{it}\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma}}{1-\sigma} - \frac{L_t^{1+1/\eta}}{1+1/\eta} \right), \tag{1.1}$$

where C_t and L_t are CES aggregates of consumption from and labor employed by a continuum of firms, indexed by $i \in [0,1]$, in period t, given by

$$C_t = \left(\int_0^1 C_{it}^{1-1/\epsilon} di \right)^{\frac{1}{1-1/\epsilon}} \tag{1.2}$$

$$L_t = \left(\int_0^1 L_{it}^{1+1/\theta} di \right)^{\frac{1}{1+1/\theta}}, \tag{1.3}$$

where C_{it} and L_{it} denote the quantity of goods (labor) consumed by (provided by) the household from (to) firm i in period t. In this setup, the firm-specific labor supply elasticity is θ and the aggregate labor supply elasticity is η. The period budget constraint is:

$$\int_0^1 P_{it}C_{it}di + Q_tB_t = \int_0^1 W_{it}L_{it}di + B_{t-1} + D_t, \tag{1.4}$$

where P_{it} is the price of good i, W_{it} is the wage paid to labor at firm i, Q_t is the price of zero coupon bond B_t, and D_t are dividends from ownership of firms. Aggregate price indices for goods and services are given by

$$P_t = \left(\int_0^1 P_{it}^{1-\epsilon}di\right)^{\frac{1}{1-\epsilon}} \tag{1.5}$$

$$W_t = \left(\int_0^1 W_{it}^{1+\theta}di\right)^{\frac{1}{1+\theta}}. \tag{1.6}$$

As in the standard model with homogeneous labor, the log-linear versions of the optimal consumption and labor supply decisions are

$$w_t - p_t = \sigma c_t + \eta l_t \tag{1.7}$$

$$c_t = \mathbb{E}_t c_{t+1} - \frac{1}{\sigma}\left(i_t - \mathbb{E}_t\pi_{t+1} + \log\beta\right), \tag{1.8}$$

where lower-case letters denote log deviations from steady state, the nominal interest rate is $i_t \equiv -\log Q_t$ and the inflation rate is $\pi_t \equiv p_t - p_{t-1}$. The firm-specific labor supply decision is

$$l_{it} = \theta(w_{it} - w_t) + l_t \tag{1.9}$$

Firms. There is a continuum of goods-producing firms in the economy, indexed by $i \in [0,1]$. Each firm i produces a specific variety and hires firm-specific labor. Output Y of a particular firm is given by

$$Y_{it} = Z_t\left(L_{it}\right)^{1-\alpha}. \tag{1.10}$$

Firms face Calvo pricing frictions and reset their prices with probability $1-\gamma$ every period. Derivation of the firm's pricing decision, and the Phillips curve, is similar to the standard model and relegated to Appendix Section A.1. The key difference is that the firm internalizes the fact that, as labor hired changes, so too does the required wage to hire that labor.

Monetary Policy. Finally, monetary policy follows a simple interest rate rule:

$$i_t = \rho + \phi_\pi\pi_t + \phi_y(y_t - y_t^n) + v_t, \tag{1.11}$$

Table 1.1: Summary of Parameters

Parameter	Description	Value
β	Discount rate (quarterly)	0.99
σ	Risk aversion parameter	1.0
γ	Frequency of price adjustment	0.75
η	Aggregate labor supply elasticity	0.2
θ	Firm-specific labor supply elasticity	*Varies*
$1-\alpha$	Output elasticity to labor	0.75
ϵ	Product demand elasticity	9.0
ϕ_π	Taylor rule coefficient on inflation	1.5
ϕ_y	Taylor rule coefficient on output	0.125
ρ_v	Persistence of monetary policy shock	0.5

where y_t^n is the natural level of output and v_t is an exogenous monetary policy shock that follows an AR(1) process with persistence ρ_v and a shock term ν_t^v.
Calibration. I calibrate the model at a quarterly frequency. The parameter values follow Galí (2008) for all parameters aside from the firm-specific labor supply elasticity θ, which I will vary to explore the role of this parameter. I report the calibrated parameter values in Table 1.1.

The Phillips Curve With Firm-Specific Labor

Phillips Curve. The Phillips curve with firm-specific labor is:

$$\pi_t = \beta E_t \pi_{t+1} + \frac{(1-\gamma)(1-\beta\gamma)}{\gamma} \frac{\sigma + \frac{\alpha+1/\eta}{1-\alpha}}{1+\epsilon\frac{\alpha+1/\theta}{1-\alpha}} (y_t - y_t^n). \tag{1.12}$$

Compared to the standard model, the Phillips curve is modified by the addition of the firm-specific labor supply elasticity parameter, θ, in the denominator of the $\frac{\sigma+\frac{1/\eta+\alpha}{1-\alpha}}{1+\epsilon\frac{\alpha+1/\theta}{1-\alpha}}$ term in Equation (1.12). One can think of this term as a "real rigidity" term in the Phillips curve, as it captures real rigidity mechanisms in the model. The numerator, $\sigma + \frac{1/\eta+\alpha}{1-\alpha}$, captures the Leahy (2011) "type 1" real rigidity forces, i.e., those that increase a firm's marginal cost when aggregate output increases. When aggregate output increases, marginal costs rise because of increasing aggregate wages (due to falling marginal utility of consumption and convex aggregate labor supply disutility) and decreasing returns to scale in the production function.

The denominator, $1+\epsilon\frac{\alpha+1/\theta}{1-\alpha}$, captures the "type 2" real rigidity forces, which reduce a firm's desired real price change in response to a change in its own marginal cost. The "type 2" real rigidity in this model arises from the interaction of the steepness of the marginal cost curve and demand elasticity. An increase in a firm's price decreases demand for its product; this decreases demand for its product reduces quantity and thus also the marginal cost. The reduction in marginal cost attenuates the desired price increase, and this attenuation is stronger the steeper the firm's marginal cost curve.

This difference highlights the different roles the two labor supply elasticities play in determining the extent of real rigidity in the economy. The aggregate labor supply elasticity appears in the "type 1" rigidity term; the *higher* the aggregate labor supply elasticity, the less aggregate wages, and therefore marginal costs, increase in response to increases in aggregate output. This induces greater real rigidity, and a lower Phillips curve slope. Meanwhile, the firm-specific labor supply elasticity appears in the expression capturing "type 2" rigidity; the *lower* the firm-specific labor supply elasticity, the steeper a firm's marginal cost curve and the less a firm will want to change its own price in response to the change in its marginal cost. This also induces greater real rigidity and a lower Phillips curve slope.

The exact difference between the slopes of the Phillips curve in the finite and infinite firm-specific labor supply elasticity depends greatly on the calibration of the firm-specific labor supply elasticity parameter. The dashed line in Figure 1.1 shows the slope of the Phillips curve, relative to the perfectly competitive case, as a function of the firm-specific labor supply elasticity (holding the other parameters constant). An elasticity of 1.0—a low estimate, but similar to the estimates from Webber (2015)—yields a Phillips curve slope that is 25% that of the competitive case. On the higher end of estimates, an elasticity of 20 (consistent with the results of stock-based estimations of the labor supply elasticity (see Sokolova and Sorensen, 2020) yields a Phillips curve with a slope that is 87% times that of the competitive case.

Comparison with a single-elasticity setup. Note that the relative Phillips curve slopes in Figure 1.1 are holding the *aggregate* labor supply elasticity constant at $\eta = 0.2$. This is not possible if the household preference structure is modeled with a single elasticity, as in Woodford (2003), Woodford (2005), Matheron (2006), and Carvalho and Nechio (2016). Such a single-elasticity setup can dramatically overstate the real rigidity importance of firm-specific labor as moving to the firm-specific case induces real rigidity once through the firm-specific channel as well as a second time if the aggregate labor supply elasticity increases as a result of calibrating the model to firm-specific labor supply elasticities. To see this, consider a model that is identical except that the representative household maximizes

$$\max_{\{C_{it}\},\{L_{it}\}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\sigma}}{1-\sigma} - L_t \right), \tag{1.13}$$

Figure 1.1: Relative Phillips Curve Slope by Firm-Specific Labor Supply Elasticities

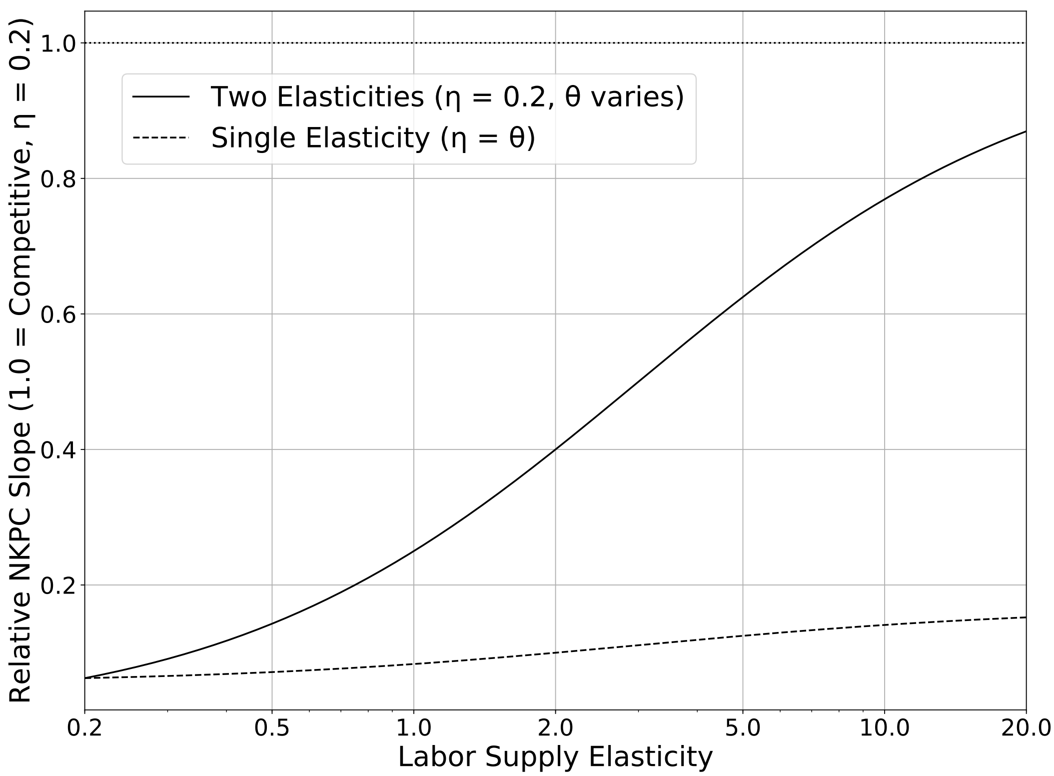

Note: Lines denote the slopes of Phillips curves relative to the case where the aggregate labor supply elasticity is $\eta = 0.2$ and the labor market is perfectly competitive (the firm-specific labor supply elasticity is $\theta = \infty$). The solid line denotes the case where $\eta = 0.2$ and θ varies. The dashed line denotes the case where θ varies and is both the aggregate and firm-specific labor supply elasticity.

where C_t and L_t are consumption and labor indices given by

$$C_t = \left(\int_0^1 C_{it}^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}} \tag{1.14}$$

$$L_t = \int_0^1 L_{it}^{\frac{1+\theta}{\theta}} di, \tag{1.15}$$

in which case θ serves as both the firm-specific and the aggregate labor supply elasticity. Under this preference structure, the Phillips curve is:

$$\pi_t = \beta E_t \pi_{t+1} + \frac{(1-\gamma)(1-\beta\gamma)}{\gamma} \frac{\sigma + \frac{\alpha+1/\theta}{1-\alpha}}{1+\epsilon\frac{\alpha+1/\theta}{1-\alpha}} (y_t - y_t^n) \tag{1.16}$$

In this Phillips curve, θ plays a role in both the "type 1" real rigidity (since it determines the aggregate labor supply elasticity) as well as the "type 2" real rigidity (through the firm-specific labor supply channel) in equation (1.12).[4]

The dashed line in Figure 1.1 plots the Phillips curve slope of this single-elasticity setup, relative to the $\eta = 0.2, \theta = \infty$ two-elasticity case. Like the two-elasticity case, the Phillips curve is substantially flatter at low levels of θ. However, for higher elasticities, the Phillips curve continues to be much flatter for the single-elasticity setup. This is because at higher elasticities the reduction in the firm-specific labor channel of real rigidity is in large part offset by the increase in the aggregate labor supply elasticity channel.

Model Responses to Monetary Policy Shocks. In the two-elasticity model, the real rigidity differences between the competitive case and low levels of the firm-specific labor supply elasticity produces large differences in the models' responses to demand shocks. In Figure 1.2, I plot the IRFs for various macroeconomic variables to monetary policy shocks (specifically, a 100bp annualized interest rate increase) for the standard New Keynesian model and the model augmented with a firm-specific labor supply elasticity of 1.08 (following Webber, 2018).

The model with finite firm-specific labor supply elasticity exhibits larger responses of output, employment and wages than that of the model with a competitive labor market. The differences are especially pronounced in the prices, with inflation and price level responses for the competitive model approximately nearly three times as large as the finite elasticity model. In response to a contractionary monetary policy shock, the on-impact inflation rate in the competitive model is -0.35%, and in the firm-specific labor model it is -0.13%. On-impact, real output falls by 0.26% in the competitive model while it falls by 0.35% in the firm-specific labor model. Similarly, employment and wages fall by more in the firm-specific labor model, because output dictates labor utilization (through the production function). In Figure 1.3, I plot the responses of the firm-specific model relative to that of the perfectly competitive model as the firm-specific labor supply elasticity θ varies. As implied by the differences in the Phillips curve slopes, a lower firm-specific labor supply elasticity leads to larger responses of output and smaller responses of prices and inflation.

1.3 Industry Heterogeneity in Firm-Specific Labor Supply Elasticities

In this section, I estimate firm-specific labor supply elasticities in the U.S. economy using a dynamic monopsony approach. With an eye towards the empirical cross-sectional

[4]The single-elasticity household preference setup is identical to the two-elasticity household preference where the firm-specific and aggregate labor supply elasticities are identical, since $\int_0^1 L_{it}^{\frac{1+\theta}{\theta}} di = \left[\left(\int_0^1 L_{it}^{\frac{1+\theta}{\theta}} di\right)^{\frac{\theta}{1+\theta}}\right]^{\frac{1+\theta}{\theta}}$.

Figure 1.2: Model Impulse-Responses to Monetary Policy Innovations

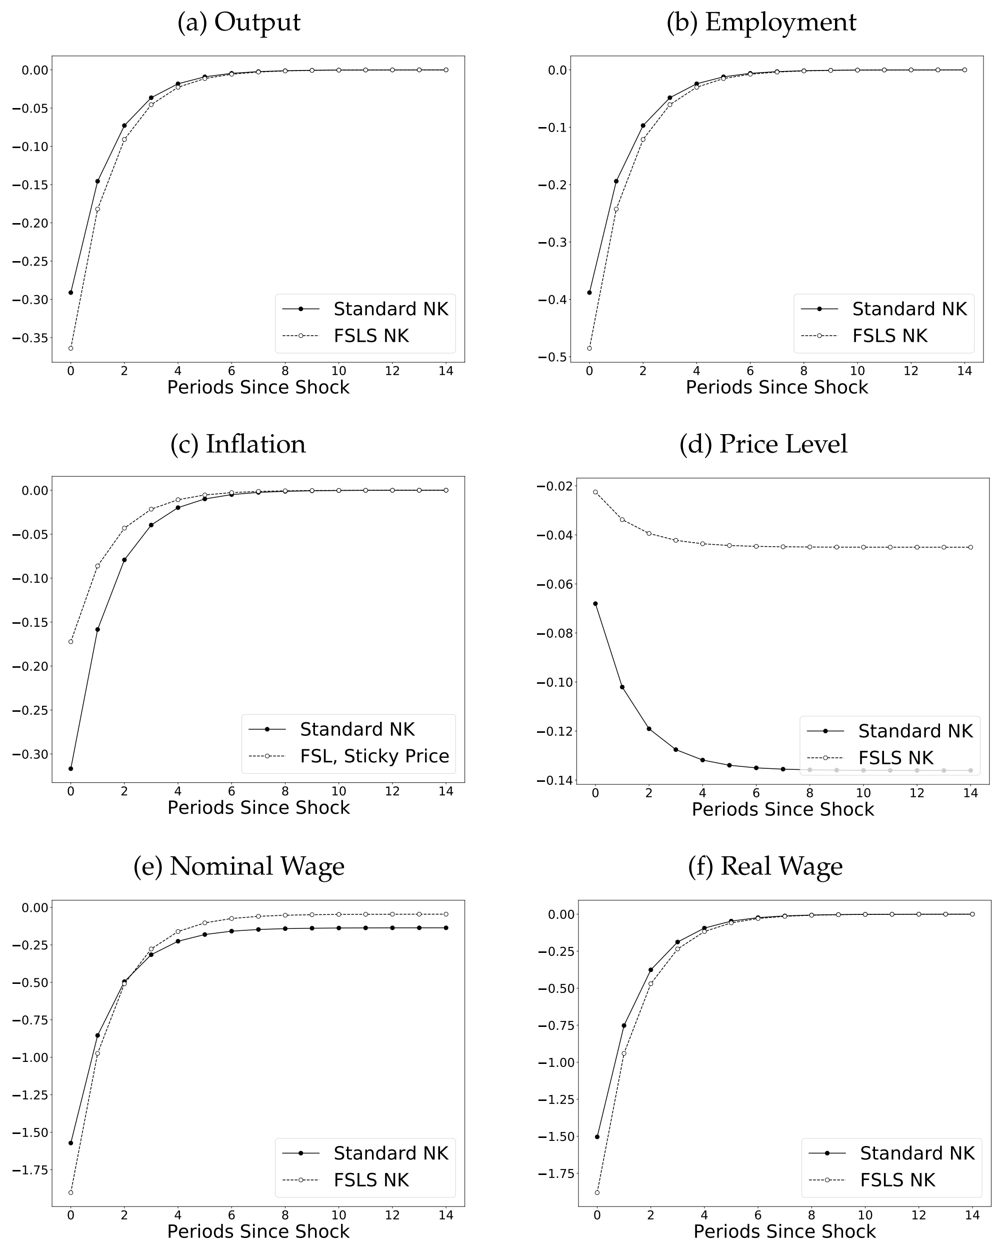

Note: IRFs are in percentage point responses to a 100 basis point (annualized) positive shock to the policy rate. Calibration of the model follows Table 1.1. The "Standard NK" model refers to the model with an infinite firm-specific labor supply elasticity; the "FSLS NK" model refers to the model with the firm-specific labor supply elasticity calibrated to 1.08.

Figure 1.3: Relative Impulse-Responses to Monetary Policy Shocks

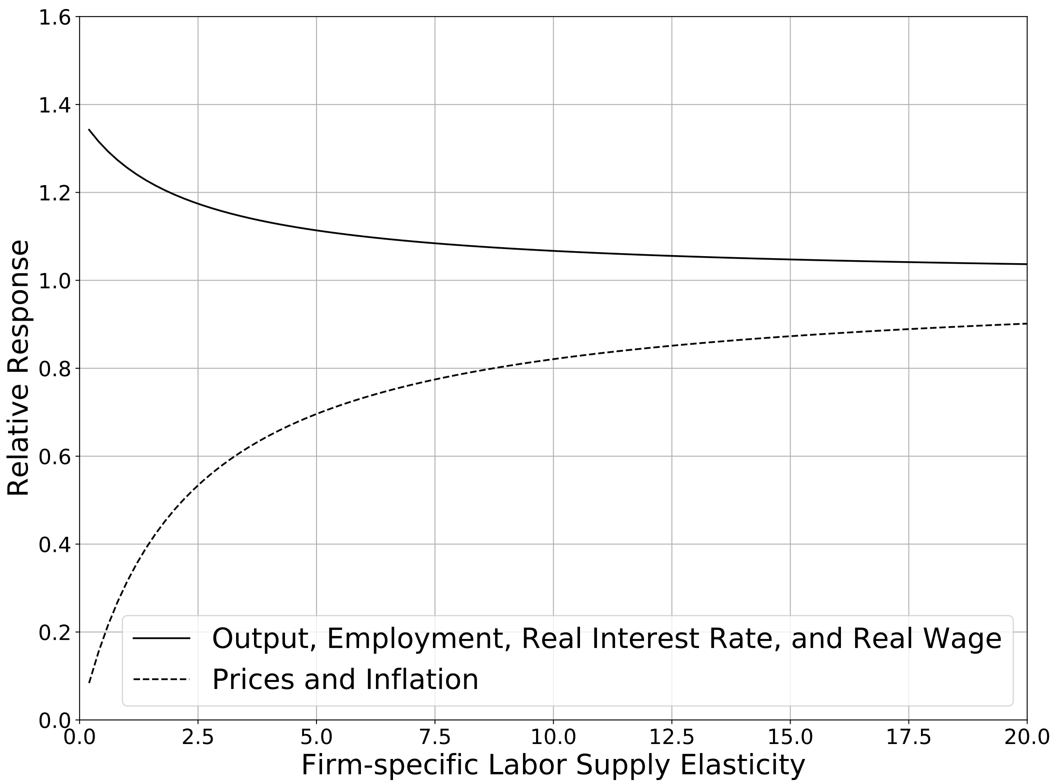

Note: A relative IRF of x means that the response of the variable is x times that of the response in the model with a perfectly competitive labor market. The relative IRFs are stable at all horizons. Calibration of the model follows Table 1.1.

approach, I estimate these elasticities by industry. I begin by describing the theory behind the dynamic monopsony approach to estimating firm-specific labor supply elasticities. The estimation method is from Manning (2013), and is based on the model in Burdett and Mortensen (1998). I then describe the SIPP data, which I use to estimate the firm-specific labor supply elasticities.

Then, I modify the model presented in Section 1.2 to include multiple sectors that are heterogeneous in the firm-specific labor supply elasticity parameter. I calibrate this model using the industry-level firm-specific labor supply elasticities to confirm that cross-sectional differences in responses to monetary policy shocks due to differences in firm-specific labor supply elasticities may be informative about the mechanism in general; industries with different firm-specific labor supply elasticities behave qualitatively like one-sector economies with different firm-specific labor supply elasticities. Lower-elasticity sectors exhibit smaller price decreases and greater output, labor, and wage decreases in response to monetary policy shocks than their higher-elasticity counterparts. Therefore, cross-sectional variation in the firm-specific labor supply elasticity and responses to monetary policy shocks may be informative about the existence of the mechanism in the aggregate.

Dynamic Monopsony Estimation of Firm-specific Labor Supply Elasticities

The Manning (2013) approach to estimating the firm-specific labor supply elasticity is based on the Burdett and Mortensen (1998) wage-posting model with on-the-job search. In this wage-posting model, a finite labor supply elasticity arises from search frictions and a finite arrival rate of job offers. When firms post wages, higher wages are associated with a higher arrival rate, since the job is more attractive to on-the-job searchers (relative to that of a lower wage job), as well as a lower separation rate, since a higher-paid worker is less likely to encounter a more attractive job while searching. Here, I outline the estimation procedure of the labor supply elasticity in this model (a complete proof of the methodology is in Manning, 2013).

For a firm paying (and posting) wage w, let $L(w)$ denote the labor employed by the firm, $R(w)$ denote the flow of recruits to the firm, and $S(w)$ denote the separation rate from the firm. In steady state, the labor supply to the firm can be written as

$$L(w) = R(w)/S(w). \tag{1.17}$$

In elasticity terms, this is

$$\theta_{L,w} = \theta_{R,w} - \theta_{S,w}, \tag{1.18}$$

where $\theta_{V,w}$ refers to the elasticity of $V \in \{L, R, S\}$ to w. The labor supply elasticity can be broken down into

$$\theta_{L,w} = \sigma_R \theta_{R,w}^{E\to E} + \left(1 - \sigma^R\right) \theta_{R,w}^{N\to E} - \sigma_S \theta_{S,w}^{E\to E} - (1 - \sigma_S)\, \theta_{S,w}^{E\to N}, \tag{1.19}$$

where σ_R and σ_S denote the fraction of recruits that are from other employers and the fraction of separations that are to other employers, respectively; $\theta_{R,w}^{E\to E}$ and $\theta_{R,w}^{N\to E}$ denote the elasticities of recruitment from other employers and non-employed workers, respectively; and $\theta_{S,w}^{E\to E}$ and $\theta_{S,w}^{E\to N}$ denote the elasticities of separation to other employers and non-employment, respectively. Manning (2013) shows that the elasticity of recruitment from employment can be written as

$$\theta_{R,w}^{E\to E} = -\frac{\sigma_S}{\sigma_R} \theta_{S,w}^{E\to E}, \tag{1.20}$$

and that the elasticity of recruitment from non-employment can be written as

$$\theta_{R,w}^{N\to E} = \theta_{R,w}^{E\to E} - w \frac{\sigma'_R(w)(1 - \sigma_R(w))}{\sigma_R(w)}, \tag{1.21}$$

where the latter term in Equation (1.21) can be thought of as the bargaining premium that an employed worker receives while searching while employed. Manning (2013) shows that if one estimates a logistic regression of the probability that a worker is a recruit from employment and the log wage is included as one of the regressors, the coefficient on the log wage is equivalent to this bargaining premium term. Thus, in order to estimate the elasticity of labor supply to the firm, one needs to estimate the separation elasticities to employment and non-employment, this bargaining premium term, and the shares of separations and recruits from and to other employers.

Estimation of Elasticities using the Survey of Income and Program Participation

I estimate the firm-specific labor supply elasticity using data from the Survey of Income and Program Participation (SIPP). The SIPP is a household-based survey in the United States comprised of a series of panels. Panels collect information from households in 4-month waves and last between 8 and 16 waves. In this paper, I use the eight survey panels between 1990 and 2008, since the questions used to collect information on job spells were similar throughout this time period, but changed significantly in 2012 in a way that coarsened the information available.

Households are surveyed every four months at the end of each wave. During the survey, respondents are asked about their employment history over the past four months. For each of the four reference months in the wave, respondents report the hours and wage rate or salary of any jobs they held during the month. Jobs are matched between reference months and waves using a unique employer ID number that is constant over the survey panel, as well as a reported job start and end month.[5] For the purposes of estimating separations and recruitment elasticities, I designate someone as employed at a job in a given month if they reported positive earnings at the job in the reference month. Non-employed are those who did not report any earnings at a job in the reference month.

Following Manning (2013), I estimate the separation elasticities $\theta_{S,w}^{E\to E}$ and $\theta_{S,w}^{E\to N}$ by modeling the instantaneous separation rate independently as $S^{ee}(x) = \exp(\beta^{ee}x)$ and $S^{en}(x) = \exp(\beta^{en}x)$, where x is a vector of controls (see below) and the log hourly wage. The elasticity of separations is the coefficient on the log hourly wage. I estimate these using maximum likelihood. The individual log-likelihood contribution is

$$\begin{aligned}\log L = {} & y^{E\to E}\ln\Big[1-\exp(-S^{E\to N}(x))\Big] + (1-y^{E\to E})\ln\Big[\exp(-S^{E\to N}(x))\Big] \\ & + (1-y^{E\to N})\Big[y^{E\to E}\ln[1-\exp(-S^{E\to E}(x))] + (1-y^{E\to E})\ln[\exp(-S^{E\to E}(x))]\Big],\end{aligned} \tag{1.22}$$

where $y^{E\to E}$ and $y^{E\to N}$ are dummies indicating separations to employment or non-employment, respectively. I define separations as those that are not employed at the same

[5]For the earlier 1990-1993 panels, I apply the fix to the erroneous job id coding described in Stinson (2003).

job in the next month. To allow for short breaks between employment spells, I define separations to non-employment as not being employed at any job in the next four months, while separations to employment are defined as separations in which the respondent is employed at some other job at any point in the next four months. Similarly, for the logistic regression used to estimate the bargaining premium term in Equation (1.21), I define a recruit from employment as an observation where the worker is employed at a job, not employed at that same job in the previous month, but employed in some job in the previous four months.[6]

I only include observations (defined as a person in a month) during which the surveyed participant only holds a single job. I drop any spell in which the wage is top-coded or if the hourly wage is under the federal minimum wage at the time. I also drop any reported employment spells if the hours are top-coded (above 98) or if hours are below 10 a week. Finally, I drop employment spells that are indicated as self-employed business, self-employed, or the armed forces. If I drop an employment spell for a given month, I do not treat the individual as non-employed during the month, to avoid erroneous coding of transitions. Rather, I keep the observation for purposes of defining transitions in other months, but drop the observation when estimating the separation elasticities. I use 7,248,517 observations in the separations likelihood estimation. Of these, 143,320 are separations to other employment, and 59,540 are separations to non-employment. The recruitment from non-employment likelihood estimation is estimated on the sample of newly employed workers; that is, observations that report employment, meet the sample criteria above, and were observed as not employed in the same job during the preceding month. Recruits are defined as those who are employed in a job that were not employed in the same job in the previous month; there are 224,757 recruits observed, 156,524 of which are from other employment, and 68,233 of which are from non-employment.

I use the same log hourly wage measure and the same set of controls in the estimation of the separations elasticities and the recruitment from non-employment elasticity. I construct the hourly wage variable using a combination of the reported hourly wage and salary. When there is an hourly wage reported, I use the hourly wage. For jobs in which there is only a reported salary, I impute an hourly wage using the reported monthly earnings divided by the number of weeks in the month multiplied by the reported hours worked per week. For controls, I include gender, marriage status, race, a set of education dummies (high school degree, some college, and college), year dummies, and a dummy indicating which reference month (1 - 4) the observation takes place in as controls. It is important to include the reference month as a control because there is a well-known "seam effect" in the SIPP, where respondents are more likely to report job changes between waves rather than within waves. To obtain sector-specific estimates of the elasticities of separation and recruitment, I interact the log wage variable with dummies indicating

[6] I use this "four-month rule" to avoid the risk of counting a recruit (separation) as from (to) non-employment if there is a temporary break between jobs. The estimated elasticities are similar if I count those transitions as involving non-employment.

NAICS 3-digit industries. Jobs in SIPP are encoded using Census industry codes, which are typically the equivalent of NAICS 4-, 5-, and 6-digit codes. I concord these to 3-digit NAICS codes.[7]

Results: Elasticity Estimates. In Figure 1.4, I plot a histogram of the industry-level estimates of the firm-specific labor elasticity estimates. Notably, the estimates imply a significant degree of monopsony power in the labor market. The median industry has a firm-specific labor supply elasticity of 1.59, and the range of estimates is from 0.47 to 2.93. In Appendix Table A.1, I report the estimated firm-specific labor supply elasticities for each industry, as well as the estimated components of the elasticity (the elasticities of separation to other employment and non-employment, the search premium term, and the shares of separations and recruits to and from other employment).

These estimates are low, but not relative to other work that has estimated firm-specific labor elasticities using dynamic monopsony methods. Webber (2015), using LEHD data, finds an average labor supply elasticity of 1.08 among U.S. firms. Sánchez et al. (2020) find average elasticities of 0.61 and 0.36 for men and women, respectively, using matched firm-worker data from Chile. Barth and Dale-Olsen (2009) find average elasticities between 0.84 and 1.71 depending on the gender and specification using Norwegian establishment data.

Multi-Sector New Keynesian Model

I now extend the model from Section 1.2 to include multiple sectors in order to create testable predictions about how industries with heterogeneous firm-specific labor supply elasticities respond to monetary policy shocks. I briefly summarize the model here and discuss the relevant model differences. Sectors with different labor supply elasticities exhibit differences in business cycle dynamics that are analogous to the differences between one-sector model economies calibrated to different elasticities.

The only modification to the model is that there is now a finite number of sectors, indexed by $j = 1, \ldots, J$, within which are a unit mass of firms, indexed by $i \in [0, 1]$. Households have two-tiered CES preferences over consumption goods and labor supply to firms. The within-sector elasticities of substitution are ϵ_j and θ_j for consumption and labor, respectively; the intersectoral elasticities are ζ and λ. Sectors are also potentially heterogeneous in the returns to scale parameter α_j and the price reset probability parameter γ_j. Firm pricing decisions give rise to a sectoral Phillips curve, the derivation of which is

[7] Because the concordance between the Census industry codes and NAICS 3-digit industries is not unique, I group some NAICS 3-digit codes and estimate their elasticities as if they were one industry. The grouped industries are listed in Appendix Table A.1.

Figure 1.4: Distribution of Firm-specific Labor Supply Elasticities across Industries

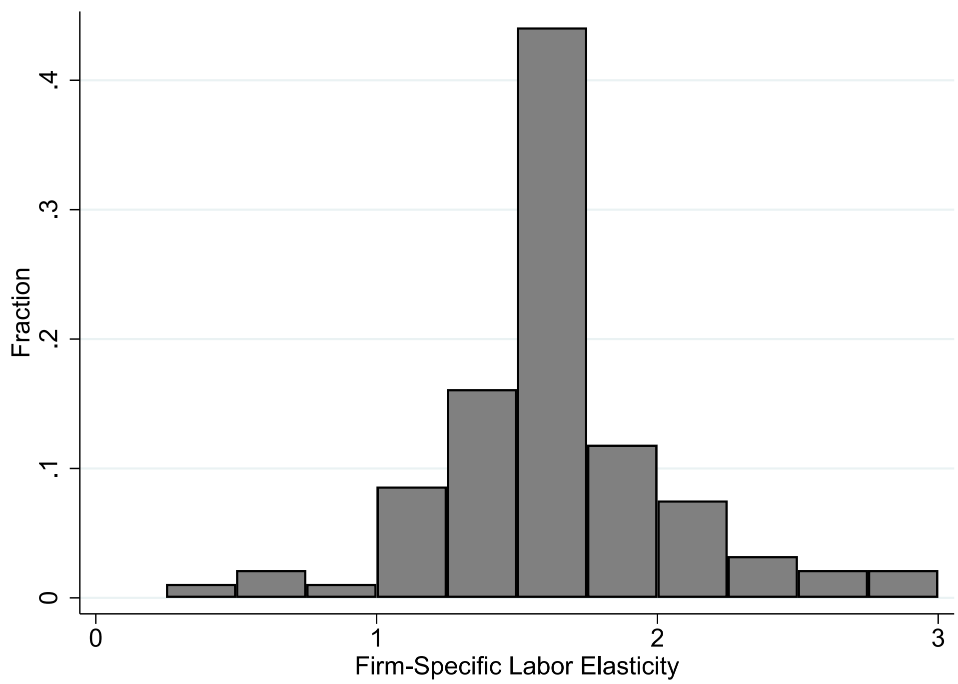

Note: This figure shows the distribution of firm-specific labor supply elasticities, where each observation is a separate NAICS 3-digit industry. The firm-specific labor supply elasticities and the underlying estimates of their components is in Appendix Table A.1.

available in Appendix A.1:

$$\pi_{jt} = \beta \mathbb{E}_t \pi_{j,t+1} + \frac{(1-\gamma_j)(1-\beta\gamma_j)}{\gamma_j} \frac{1}{1+\epsilon_j \frac{\alpha_j + 1/\theta_j}{1-\alpha_j}} \times \left[\left(\frac{\alpha_j}{1-\alpha_j} \right) \check{y}_{jt} - \left(\frac{1+1/\lambda}{1-\alpha_j} \right) \check{z}_{jt} + \sigma \check{y}_t + \frac{1}{\zeta} \left(\check{y}_{jt} - \check{y}_t \right) + \frac{1}{\eta} \check{l}_t + \frac{1}{\lambda} \left(\check{l}_{jt} - \check{l}_t \right) \right]. \tag{1.23}$$

The sectoral Phillips curve in the multi-sector model is analogous to the one-sector Phillips curve. Since firms compete in the product market against other firms in the same sector, what matters for price setting is expectations of sectoral, not aggregate inflation. The term $\frac{1}{1+\epsilon_j \frac{\alpha_j+1/\theta_j}{1-\alpha_j}}$ is a direct analog of the denominator in Equation (1.12); that is, it captures how the firm responds to changes in the marginal cost of its competitors. As in the single-sector model, the firm-specific labor supply elasticity appears in this term, lowering the slope of the sectoral Phillips curves as the elasticity decreases. The bracketed term in Equation (1.23) captures how the marginal cost at the sectoral level evolves. As in the one-sector model, this depends on aggregate output through diminishing marginal utility of consumption and increasing aggregate labor disutility. It also depends on sectoral output and labor *relative* to aggregate output and labor, since consumption and labor are imperfect substitutes across sectors.

Calibration. I calibrate the multi-sector model using the same parameter values as in Section 1.2 for the one-sector model, with the exception of the inter- and intra-sectoral demand and and labor supply elasticities. For the intersectoral labor supply elasticity, I follow Berger, Herkenhoff and Mongey (2021), who estimate the intersectoral labor supply substitution elasticity using changes in corporate tax rates, and set ζ to 0.31. They define the upper-level labor sector as NAICS 3-digits by commuting zones, so the elasticities at the NAICS 3-digit level are likely to be lower. However, lower upper-level labor elasticities do not qualitatively change the results of the multi-sector model.

I set the intersectoral product demand elasticity λ to 3.0, which Hobijn and Nechio (2017) estimate as the sectoral-level elasticity using long-run changes in relative prices in response to changes in value-added tax rates. This elasticity may be too small for NAICS 3-digit sectors, since their sectoral definitions are somewhat larger than NAICS 3-digit sectors. On the other hand, these are elasticities estimated off of long-run changes, not short-run elasticities, which may be lower and more relevant to sectoral responses within a few quarters. I calibrate the model to 92 sectors, each corresponding to a NAICS 3-digit industry for which I have an estimated firm-specific labor supply elasticity from Section 1.3. To isolate the effect of the heterogeneous labor supply elasticity, I keep homogeneous $\alpha_j = 0.25$, $\epsilon_j = 9.0$ and $\gamma_j = 0.75$ for all j. I simulate the economy's response to the same 100 basis point (annualized) monetary policy shock as above.

In Figure 1.5, I plot the sectoral responses of output, prices, labor, and wages for the sectors with the lowest, 10th percentile, median, 90th percentile, and highest firm-specific labor supply elasticities (0.47, 1.18, 1.59, 2.07, and 2.93, respectively). As with the differences between one-sector economies calibrated to different firm-specific labor supply elasticities, monetary policy shocks induce smaller responses of prices and larger responses of output, employment, and wages in sectors with lower elasticities. The response of sectors with heterogeneous firm-specific labor supply elasticities thus resembles the difference between one-sector economies with different firm-specific labor supply elasticities.

There are, however, important quantitative differences in how the firm-specific labor supply elasticity affects the behavior of different one-sector economies and different sectors in a multi-sector economy. In Figure 1.6, I plot the impulse-responses of the lowest, median, and highest-elasticity sectors along with the impulse-responses of one-sector economies calibrated to the same firm-specific labor supply elasticities as those sectors. Relative to their single-sector counterparts, the differences in sectoral responses are less pronounced in prices, and more pronounced in output, employment, and wages. Over time, as the

Figure 1.5: Comparison of Model Sectoral Responses to Monetary Policy Shocks

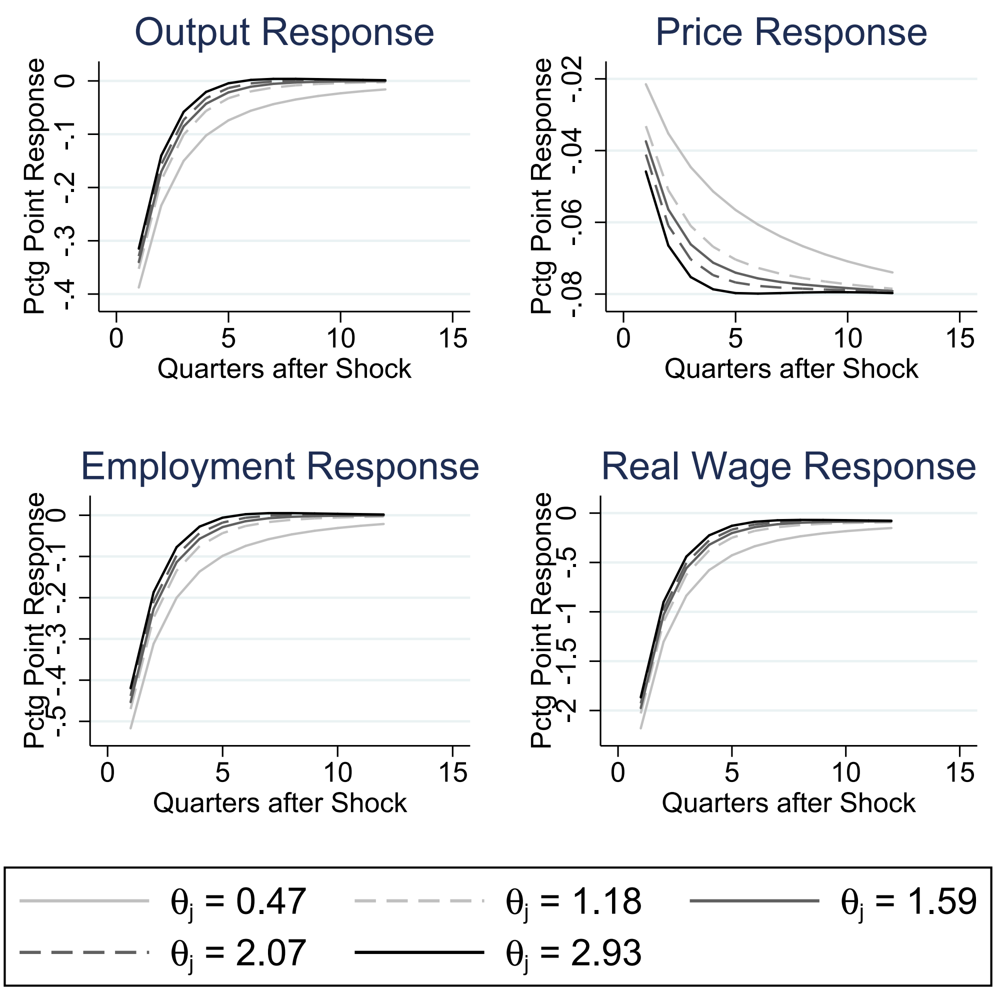

Note: IRFs are in percentage point responses to a 100 basis point (annualized) positive shock to the policy rate; each time unit represents one quarter.

monetary policy shock wears off, the difference in price responses in the multi-sector model disappear as relative sectoral price return to parity. The differences arise from the presence of intersectoral substitution in the product and labor markets, which is not present when comparing one-sector models with different firm-specific labor supply elasticities.

1.4 Results: Empirical Industry Responses to Monetary Policy Shocks

In this section, I test the predictions of the multi-sector New Keynesian model with firm-specific labor supply in the industry cross-section. I estimate IRFs of industry variables, industry-by-industry, and project those IRFs onto the firm-specific labor supply elasticity estimates as well as other industry characteristics. I find no cross-sectional evidence that differences in firm-specific labor supply elasticities are associated with differences in real rigidity between industries. I do not find any differential effect of monetary policy shocks on industry outcomes due to differences in firm-specific labor supply elasticity that support the hypothesis that lower firm-specific labor supply elasticity generates more real rigidity. Industries with differing firm-specific labor supply elasticities do not experience differential price responses to monetary policy shocks. I also do not find any consistent evidence that low-elasticity industries experience greater responses of output, employment, or wages. In fact, industries with higher elasticities actually face more negative responses of output and employment to contractionary policy shocks, contrary to the real rigidity story. These results are robust to the inclusion of various industry characteristics as controls as well as an alternate specification.

Empirical Strategy: Estimating Responses to Monetary Policy Shocks

To estimate the differential effect of monetary policy shocks on industries with different firm-specific labor supply elasticities, I estimate IRFs of industry variables (prices, output, employment, and wages) using a series of Jorda local projections. I then project the IRFs at different horizons on industry characteristics.[8] For each industry i and variable y, I estimate a series of local projections for horizons $h = \{1, \ldots, H\}$:

$$\log y_{i,t+h} - \log y_{i,t-1} = \alpha^{y,h} + \sum_{j=1}^{J} \beta_j^{i,y,h} \Delta \log y_{i,t-j} + \sum_{k=0}^{K} \gamma_k^{i,y,h} shock_{t-k} + \nu_t^{i,y,h}. \tag{1.24}$$

The response of industry variable y in a industry i to the monetary policy shock h periods out is equal to the coefficient $\gamma_0^{i,y,h}$ in Equation (1.24). I estimate the local projections on monthly data up to $H = 24$ horizons and set $I = J = 12$. Then, for each variable y and each horizon h, I then regress the estimated IRFs on the industries' estimated firm-specific

[8]This methodology was previously used by Dedola and Lippi (2005) and Henkel (2020), although not to measure the effect of firm-specific labor supply elasticities on industry outcomes.

Figure 1.6: Comparison of Multi-sector and Single-sector Responses to Monetary Policy Shocks

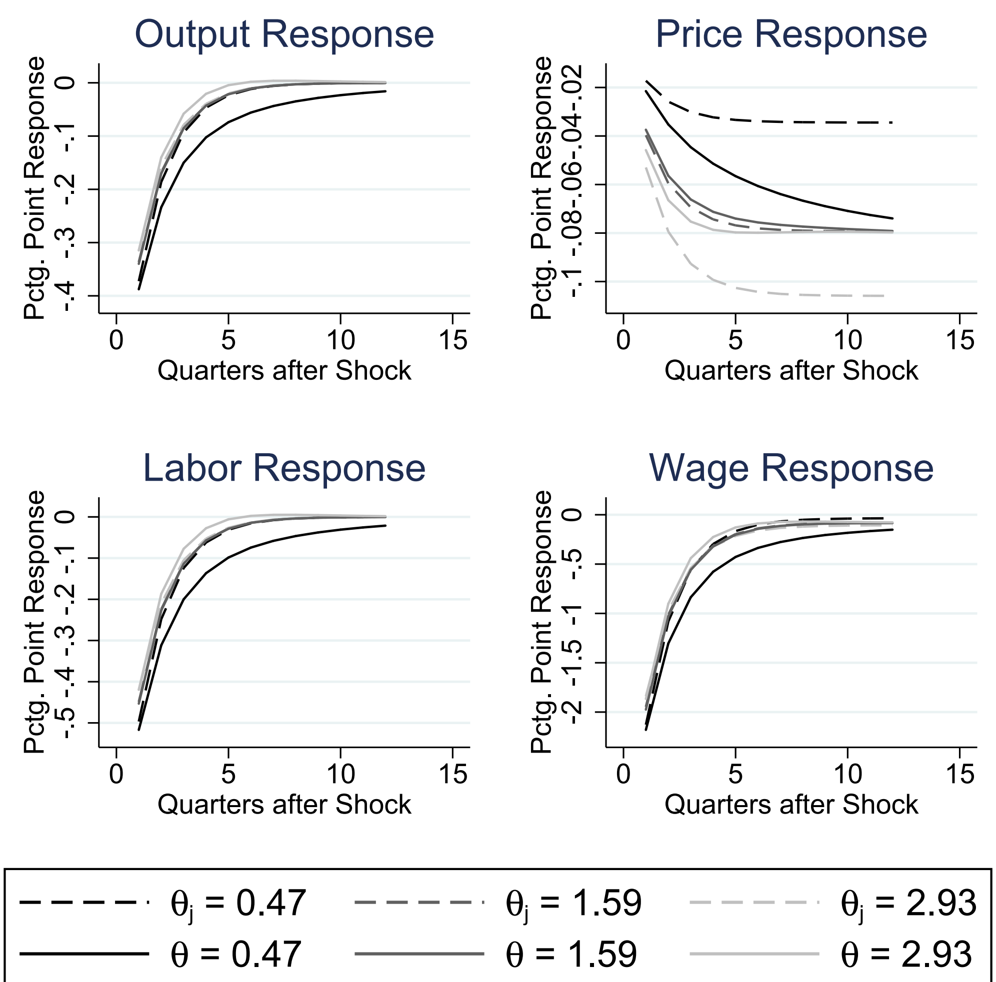

Note: IRFs are in percentage point responses to a 100 basis point (annualized) positive shock to the policy rate. Solid lines indicate the impulse response of the variable for a sector in the multi-sector model; the dashed lines refer to responses of the corresponding one-sector models where the firm-specific labor supply elasticity has been calibrated to the corresponding sector. Time periods correspond to quarters. Light grey lines correspond to the sector (in the multi-sector model) or the one-sector model calibrated to the lowest elasticity sector ($\theta_j = 0.47$); grey lines correspond to the sector (in the multi-sector model) or the one-sector model calibrated to the median elasticity sector ($\theta_j = 1.59$); black lines correspond to the sector (in the multi-sector model) or the one-sector model calibrated to the highest elasticity sector ($\theta_j = 2.93$).

labor supply elasticities and a vector of controls:

$$\hat{\gamma}_0^{i,y,h} = a^{y,h} + b^{y,h} \log(\hat{\theta}_i) + C_i X_i + e_i^{y,h}. \tag{1.25}$$

The coefficient of interest, $b^{y,h}$, which I call the "differential IRF," captures how the industry IRFs relate to industry heterogeneity in the firm-specific labor supply elasticity. The monetary policy shock, which is described below, is scaled so that a positive value corresponds to a contractionary monetary policy shock. To recap the predictions from Section 1.3, the New Keynesian model would predict less negative responses of prices and more negative responses of output, employment, and wages; that is, the theory predicts negative $\hat{b}^{price,h}$ to be negative and positive $\hat{b}^{output,h}$, $\hat{b}^{employment,h}$, and $\hat{b}^{wages,h}$.

Industry Outcome Variables. For prices, I use the monthly Producer Price Index (PPI) data from the Bureau of Labor Statistics (BLS). I use the data available at the NAICS 3-digit level, and do not attempt to replace or construct NAICS 3-digit level data for 3-digit series where it is not available. Most industries are available from 2004 onwards, although some go back further. Unfortunately, the concordance between the earlier SIC-based series (pre-2004) and the later NAICS-based data is insufficiently clean for use.

For output, I construct a monthly output series using data from the Board of Governors of the Federal Reserve System's monthly real industrial production series (G.17) and the Census Bureau's monthly retail and wholesale trade reports. The latter two report nominal sales for retail and wholesale trade industries, which I deflate using the corresponding PPI.[9]

For labor market variables, I use the Current Employment Statistics (CES) from the BLS. For both employment and wages, there are several choices; I use production employment and average weekly real production earnings as employment and wage measures, respectively, although results using other employment and wage series are similar. These data are available on a monthly basis from 1990 onwards.

Monetary Policy Shocks. I use the monetary policy shock series from Bu et al. (2021). The shocks are derived using a Fama-Macbeth two-step procedure in which the authors first estimate the sensitivity of interest rates across the maturity spectrum, and then recover the monetary policy shock from a cross-sectional regression of interest rate changes on the sensitivity estimates. Importantly, this shock series does not show any evidence of an information effect à la Nakamura and Steinsson (2018) and produces conventionally signed impulse-response of aggregate production and prices. The shock series is available at a monthly basis from 1994 onwards.

---

[9] As an alternative, I use the real gross output and real value-added series from the Bureau of Economic Analysis (BEA) Industry Economic Accounts Data. This data is generally available at a NAICS 3-digit level on a quarterly basis from 2000q1 forward. Results are similar between the monthly and quarterly regressions.

Results: Industry Responses and Firm-Specific Labor Supply Elasticities

First, I estimate Equation (1.25) without controls. I plot the estimated differential IRFs, the estimates of $b^{y,h}$, in Figure 1.7. To reiterate, the estimate of $b^{y,h}$ is the coefficient on the firm-specific labor supply elasticity in Equation (1.24), and measures how different the IRFs are due to differences in firm-specific labor supply elasticities. According to the model, industries with higher firm-specific labor supply elasticities should see negative differential IRFs for prices, but positive differential IRFs for output, wages, and employment. I also plot the IRFs of the aggregate counterparts of the industry variables in Figure 1.7.

Contrary to the prediction of the New Keynesian model with firm-specific labor, I find no significant effect of firm-specific labor supply elasticity on responses of industry prices. Industries with higher firm-specific labor supply elasticities do not appear to experience significantly different responses of prices to the monetary policy shock, with the point estimates of the differential IRF for prices near zero. Neither do I find any evidence for the real rigidity predictions in the responses of real output or wages, although the standard errors on the former are quite large and the latter measure is not composition-adjusted.

The only industry outcome which appears to be affected by the firm-specific labor supply elasticity is production employment. Here, the sign of the estimated differential IRF is the opposite of that predicted by the New Keynesian model, which predicted larger employment falls in low-elasticity industries, not high-elasticity industries. The empirical industry responses find the opposite, with employment falling more in higher-elasticity industries, with the difference significantly different (at the 90% level) from zero starting at 20 months from the monetary policy shock. Relative to the aggregate response of production employment, the difference in employment is substantial. At 24 months, the point estimate of the differential IRF is -0.224. The difference in log elasticities for the 25th- and 75th-percentile industries (raw elasticities of 1.18 and 2.07, respectively) is 0.56; these estimates imply a difference in employment responses to the monetary policy shock of -0.136 percentage points. Compared to the aggregate response of employment, this difference is substantial. The aggregate response of production employment to the monetary policy shock at 24 months is -.065 percentage points, and the peak response, at 9 months, is -0.122 percentage points.

Controlling for Other Industry Characteristics

A potential concern is that the firm-specific labor supply elasticity is not randomly assigned between industries. This poses a threat to identification if there are industry characteristics that affect industry responses to monetary policy shocks that are also correlated with the firm-specific labor supply elasticity. For example, it may be the case that industries vary in firm size, which could affect both the firm-specific labor supply elasticity as well as the borrowing capacities of those firms. This could lead to omitted variable bias if monetary policy is stronger in industries with firms that are more financially constrained

Figure 1.7: Differential Effects of Monetary Policy Shocks on Industry Variables (No Controls)

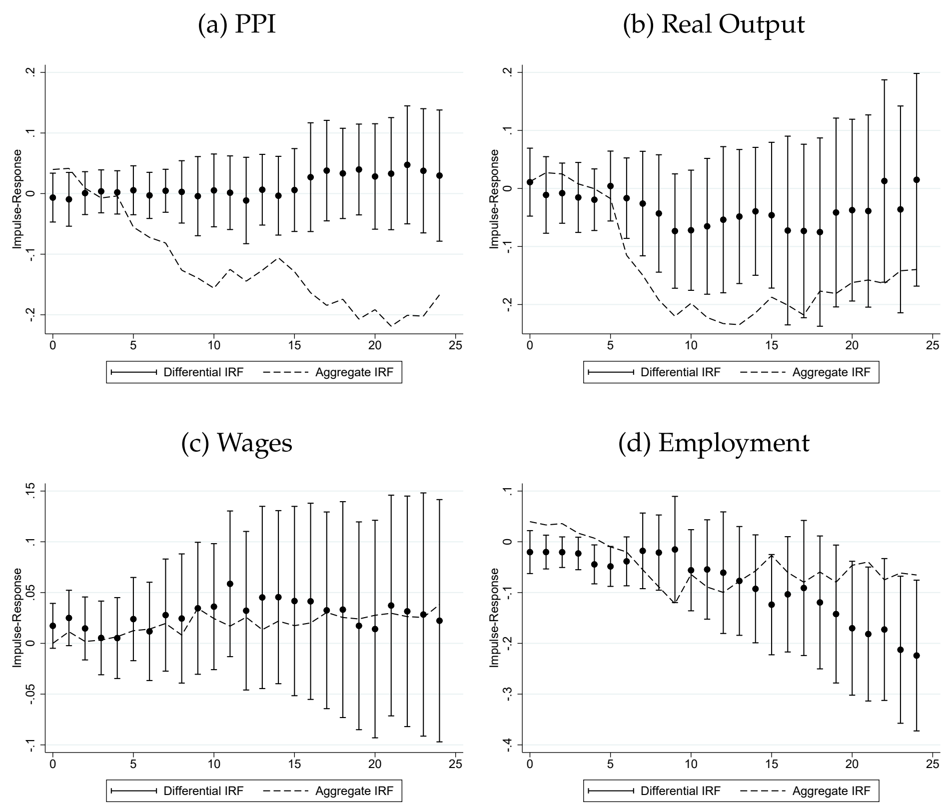

Note: Each point represents the estimate of $\hat{b}^{y,h}$ from estimating equation (1.24), with no controls. Error bars represent 90% confidence intervals. The wage measure is average weekly real production employee earnings and the employment measure is production employees. The impulses-responses of the aggregate variables are estimated using the same local projections as the industry variables in Equation (1.24), but with aggregate variables instead of industry variables. The aggregate variables used are for the FRED series PPIACO (aggregate PPI, for prices), INDPRO (real industrial production, for output), CES0500000006 (production employment, for employment) and CES0500000030 (average weekly real production earnings, for wages).

(as Dedola and Lippi, 2005, find); in this particular case, the differential IRF would be biased downwards. To address the issue of omitted variable bias, I estimate Equation (1.24) with a set of industry controls.

Description of Controls. First, I control for the frequency of price adjustment. Previous work (Bils, Klenow and Kryvstov, 2003; Henkel, 2020) has found differential responses between goods and industries with respect to the frequency of price adjustment. I use the frequency of price adjustment data reported by Nakamura and Steinsson (2008), derived from the BLS data underlying the PPI and CPI. For their PPI data (Table 23 of their online supplement), I manually concord each item to a NAICS 3-digit manufacturing code to cover NAICS codes beginning with 31, 32, and 33. In addition, I supplement this data with their CPI-based frequency of price adjustment data for retail trade and services (Table 20). For both data sources, I use the price change frequency with substitutions as a measure of the frequency of price adjustment.

For the manufacturing NAICS industries (NAICS codes beginning in 31, 32, and 33), I use the PPI-based frequency of price adjustment data, matching PPI product codes to NAICS manufacturing industries. For the retail trade industries, I draw on frequency of price change data from the CPI, matching CPI Entry Level Items (ELIs) to the appropriate NAICS retail sector. For example, "Girl's Dresses" is matched to the NAICS 3-digit code 448 ("Clothing and Clothing Accessory Stores)." For a number of non-manufacturing and non-trade industries, I am also able to obtain a measure of the frequency of price change using the reported CPI data. For example, I use the frequency of price change of the "Airline Fare" item in the CPI data as the measure of the frequency of price change for NAICS 481 ("Air Transportation"). Overall, I am able to map the frequency of price adjustment data to 49 of the NAICS industries. In the case where I have multiple items mapped to a NAICS sub-sector, I take the median frequency of price adjustment of all items matched to that NAICS 3-digit industry.

A second set of controls includes other industry characteristics that, while not present in the canonical New Keynesian model, may affect industry responses to monetary policy shocks and also be related to industry variation in firm-specific labor supply elasticities. First, I construct a measure of interest rate exposure using data from Compustat. For each industry, I compute measures of the interest rate burden (interest expenses over sales), the leverage ratio (total debt over assets), and the short-term debt ratio (short term debt over assets). These measures are computed annually and then averaged over years 2004 through 2019. Second, I construct a measure of the fraction of establishments with under fifty employees in each industry using the Quarterly Census of Employment and Wages from the QCEW. Finally, I include a dummy variable for industries producing durable goods.[10]

A final set of controls is related to the other determinants of real rigidity in the canonical New Keynesian model explored earlier in this paper, returns to scale in the production function and the elasticity of product demand. Real rigidities are theoretically increasing in the product demand elasticity and decreasing in the returns to scale parameter. In particular, it is plausible that industries that are monopsonistic in the labor market may also be monopolistic in the product market (it may be the case that these are related to firm concentration, which could appear in both the product and labor market; or, search frictions in the labor market may be correlated with similar search frictions in the product market). Normally, these might be controlled for by using the profit share and the labor share, respectively; however, if monopsony power allows firms to set wages below the

[10] NAICS codes starting with 33, and codes 321, 327, and 423.

marginal revenue product of labor level, the profit share and labor share are controls that are outcomes variables of the firm-specific labor supply elasticity, thus making them "bad controls."

In the absence of better controls for the returns to scale and the elasticity of product demand, I estimate Equation (1.24) with and without the labor share and profit shares in the set of controls X_i. For the labor share, I use the industry's average compensation over value added and for the profit share I use the average net operating surplus over value added, as reported in the BEA's GDP by Industry statistics. The results are robust to the inclusion or exclusion of the labor and profit shares as controls. In the main text, I report the results without these controls. In Appendix Table A.3 - A.6, I compare the results with and without these labor and profit shares as controls; the estimates are very similar.

Differential IRFs with Controls. I present the differential IRFs, estimated with controls, in Figure 1.8, along with the differential IRFs estimated without controls from Section 1.4 for comparison. I also report the estimated differential IRFs, with and without controls, in Appendix Table A.2 at 3-month intervals. As before, the results with controls provide no evidence for the real rigidity mechanism. Industries with different firm-specific labor supply elasticities do not exhibit significantly different responses of prices to monetary policy shocks. The differential IRFs for output become more negative, contrary to that predicted by the New Keynesian model (which would have predicted positive differential IRFs), although again the standard errors for those estimates are large. Real wages also continue to show no evidence of differential effects of monetary policy. The results for employment become even more negative, further contradicting the notion that firm-specific labor supply generates real rigidity.

Alternative One-Step Specification

For robustness, I estimate the differential IRFs using a one-step estimation procedure. To estimate the differential effect of monetary policy shocks on industries with different firm-specific labor supply elasticities, I again estimate a series of local projections, but instead of estimating the local projections industry-by-industry, for each variable I estimate a single projection using all industries, and interact the monetary policy shocks with industry characteristics. For each industry variable y and for each horizon in $h = 0, \ldots, H$, I estimate:

$$\log y_{i,t+h} - \log y_{i,t-1} = \alpha^{y,h} + \sum_{j=1}^{J} \beta_j^{y,h} \Delta \log y_{i,t-j} + \sum_{k=0}^{K} \gamma_k^{y,h} shock_{t-k} \tag{1.26}$$

$$+ \sum_{k=0}^{K} \delta_k^{y,h} \left(shock_{t-k} \times \log \hat{\theta}_i \right) + \sum_{k=0}^{K} \left(shock_{t-k} \times X_i \right) Z^{y,h} + \nu_{i,t}^{y,h}, \tag{1.27}$$

where X_i is the vector of industry controls. The coefficient of interest is $\hat{\delta}_0^{y,h}$, which gives the differential impulse-response of a an industry variable y to a monetary policy shock h periods out as the firm-specific labor supply elasticity of the industry changes. I estimate

Figure 1.8: Differential Effects of Monetary Policy Shocks on Industry Variables (with Controls)

Note: Each point represents the estimated differential impulse response function at the indicated horizon, ($\hat{b}_0^{y,h}$ from estimating equation (1.24), using the frequency of price adjustment, durables, and financial constraint variables as controls). Error bars represent 90% confidence intervals using robust standard errors. The wage measure is average weekly real production employee earnings and the employment measure is production employees.

the local projections on monthly data up to $H = 24$ horizons and set $I = J = 12$. To recap the predictions from Section 1.4, the New Keynesian model would predict more negative responses of output, employment, and wages, and less negative responses of prices; that is, one would expect $\hat{\delta}_0^{y,h}$ to be negative for prices and positive for output, employment, and wages. In Figure 1.9, I plot the estimated coefficients $\hat{b}^{y,h}$ from this one-step procedure alongside the differential IRF estimates from the two-step procedure in the previous section for comparison.

As with the two-step estimation, I do not find any evidence that low firm-specific labor supply elasticities are associated with stronger real rigidities. Prices do not fall by appreciably less, and output, wages and employment do not fall by significantly more, in

Figure 1.9: Comparing One-Step vs. Two-Step Specifications

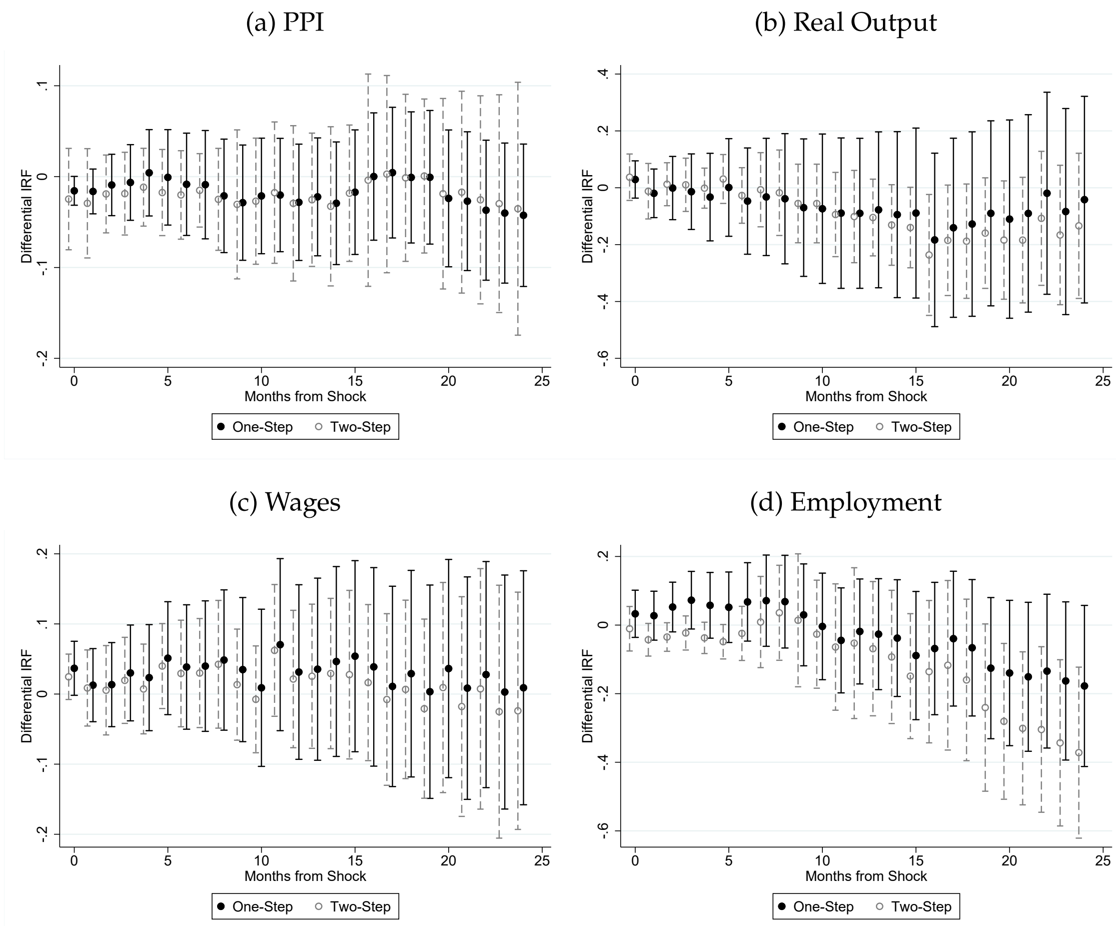

Note: This Figure plots the estimated differential IRFs from the one-step and the two-step specifications. The two-step plot uses the estimates of $b^{y,h}$ from estimating Equation (1.24); the one-step plot uses the estimates of $\delta_0^{y,h}$ from estimating Equation (1.27). Both estimates are with the same set of controls (durables dummy, frequency of price adjustment, interest expense over sales, short-term debt ratio, and total debt over assets) as described in Section 1.4. Error bars indicate 90% confidence intervals, using robust standard errors in the one-step procedure.

low-elasticity industries. Similar to the previous results, I find, if anything, that employment falls by more in the higher-elasticity industries, contrary to the predictions of the model.

1.5 Discussion and Conclusion

My results cast doubt on the theory that firm-specific labor generates real rigidities. Despite the theoretical argument that firm-specific labor is a strong source of real rigidity, I consistently fail to find any evidence that industries with higher firm-specific labor supply elasticities experience more negative price responses in response to contractionary monetary policy shocks; rather, the firm-specific labor supply elasticity appears to have no effect on industry price responses. When it comes to other industry variables, I find evidence in the opposite direction that the real rigidity story would suggest. Industries with higher firm-specific labor supply elasticities actually experience more negative employment responses to monetary policy shocks, as opposed to less negative responses as the model from Section 1.3 predicted. This difference in employment responses is large, relative to the aggregate response of employment. These results are consistent across the two empirical specifications, as well as the inclusion or exclusion of control variables.

There are some statistical caveats to these results. The first is the estimated firm-specific labor supply elasticities are themselves subject to measurement error. This measurement error biases the estimated differential IRFs towards zero, and may make it difficult to detect any differential effects of monetary policy shocks on the price level. Monetary policy shocks, especially during this period, may simply lack the power to detect the differential effects of monetary policy. Finally, if the intersectoral elasticities of demand and labor supply are large, one might not expect large difference in industry price responses, even if firm-specific labor causes real rigidity in the aggregate. If firm-specific labor does indeed induce real rigidity, these are reasons why a cross-sectional approach may fail to detect any differences.

One consistent result is that industries with higher firm-specific labor supply elasticities tend to experience sharper declines in employment that those with lower elasticities. This is contrary to the prediction of the New Keynesian model. However, it is important to note that, in the model, firms compete for workers with other firms in the same industry, since the firm-specific labor supply elasticity parameter θ_j arises from the household's elasticity of substitution of labor supply within firms in the same sectors. Therefore, in the model, differences in sectoral responses of employment arise because differential sectoral responses of prices lead to differential responses of output, which then pass through to differential responses of labor demand. Thus, the firm-specific labor supply elasticity affects employment responses indirectly through its effect on firm's pricing decisions. However, in reality, firms compete for labor with firms outside their industry as well as non-employment. It could be the case that during a negative demand shock, firms with high labor supply elasticities lose workers to firms with low labor supply elasticities, or lose more workers to non-employment than firms with low labor supply elasticities, explaining why I find that high-elasticity industries experience larger employment falls in response to contractionary monetary policy shocks.

If it is the case that firm-specific labor is not a source of real rigidity, why not? In theory, the reason why firm-specific labor matters is that it steepens the marginal cost curve of the firm, which takes this into account when setting prices. It may be the case that the elasticity of labor supply to the firm has little relevance to the marginal cost curve of the firm. This could be the case if firms adjust production using other margins of input into the production process, such as materials or hours. Or, it may be the case that short-term wage stickiness means that the labor supply curve to the firm is uninformative about its marginal cost.

Future research may be able to refine the empirical strategy used in this paper. Larger datasets may permit measuring firm-specific labor supply elasticities by finer industries, geographic-industry variation, or even at the firm-level, as in Webber (2015). This cross-sectional variation may be able to detect real rigidity effects of firm-specific labor that the industry cross-sectional design used in this paper is unable to. More broadly, labor market monopsony may affect business cycle dynamics through channels other than real rigidity. The lack of attention to labor market monopsony's role in business cycle models is surprising, given the centrality of labor supply to macroeconomic models and the recent work showing the extent and rise of labor market monopsony. As such, firm-specific labor's role in business cycle dynamics remains an understudied topic.

Chapter 2

Reservation Raises: The Aggregate Labor Supply Curve at the Extensive Margin

[†] Joint with Benjamin Schoefer

2.1 Introduction

Business cycle fluctuations in total hours largely reflect employment shifts, i.e., they occur along the extensive margin (see, e.g., Heckman, 1984). Hence, the shape of the short-run aggregate labor supply curve at the extensive margin—the total number of individuals desiring to work as a function of prevailing wages—is a crucial factor in business cycle models. In market-clearing equilibrium models, this curve forms the iron link between wages and employment, with business cycles implying large elasticities, i.e., a large mass of marginal individuals (Hansen, 1985; Rogerson, 1988). In models of wage bargaining and search frictions, the curve enters workers' outside options and reservation wages (Jäger, Schoefer, Young and Zweimüller, 2020; Koenig, Manning and Petrongolo, 2020), so that large employment fluctuations again imply a large mass of marginal individuals (Hagedorn and Manovskii, 2008; Ljungqvist and Sargent, 2017). In New Keynesian models, the aggregate labor supply curve shapes the Phillips curves for wages and prices, which imply large elasticities or the presence of frictions (Galí, 2011). The local elasticity of the curve, and hence the mass of marginal individuals, also determines the cyclical amplitude of potential labor market disequilibria and their welfare costs (Shimer, 2009). It also speaks to the employment effects of earnings subsidies (Card and Hyslop, 2005; Kleven, 2019) and tax reforms (Martinez, Saez and Siegenthaler, 2021). Finally, the short-run, Frisch elasticity is also an upper bound for the Hicksian elasticity (Chetty, 2012), which in turn guides the long-run labor supply effects of taxation (Prescott, 2004; Saez et al., 2012).

[†]This chapter is an adaptation (with permission) of "Reservation Raises: The Aggregate Labor Supply Curve at the Extensive Margin," a working paper coauthored with Benjamin Schoefer.

The existing strategies to measure the extensive-margin aggregate labor supply curve are threefold. First, a long literature has structurally estimated specific models with participation choices—making parametric assumptions about functional forms, including the joint distribution of tastes and wages, on the basis of observational data.[†] Second, recent quasi-experimental studies of income tax holidays have disciplined specific arc elasticities—albeit with respect to net-of-tax wage changes an order of magnitude larger than those over business cycles.[†] They may also capture equilibrium, compensated, and frictional effects, rather than purely preferences. Third, a small strand of research has elicited labor supply preferences in surveys—albeit in specific and selected samples such as of unemployed job seekers or older workers.[†]

Yet, no survey evidence exists on the extensive-margin labor supply preferences of a representative sample of individuals from *all* labor force groups: employed, unemployed, and out of the labor force. Such a comprehensive sample is necessary to measure the aggregate labor supply curve for an entire economy and to discipline macro models.

We fill this gap by eliciting extensive-margin labor supply preferences from a representative sample of individuals from all labor forces groups, and on that basis provide a nonparametric estimate of the global aggregate labor supply curve. As a convenient scalar measure, we capture these preferences in the form of *reservation (pay) raises* (or cuts): the hypothetical percent shift in an individual's actual/potential labor earnings required to render her indifferent between employment and nonemployment. It equals the ratio of an individual's actual/potential wage to her reservation wage. It is therefore a close cousin of the standard reservation wage (which enters aggregate labor supply, as in Chang and Kim, 2006, 2007). A convenient property is that, by being normalized by an individual's idiosyncratic actual/potential wage, the reservation raise collapses these two dimensions of heterogeneity into a scalar.[†] Reservation raises then give the aggregate labor supply curve as a univariate function: the cumulative distribution function (CDF) of the reservation raises. Its argument is the *prevailing aggregate raise*, a homogeneous

[†]For examples of structurally estimated labor supply models with participation margins, see, e.g., Heckman and MaCurdy (1980); Chang and Kim (2007); Gourio and Noual (2009); Blundell, Pistaferri and Saporta-Eksten (2016); Chang and Kim (2006); Park (2020); Attanasio, Levell, Low and Sánchez-Marcos (2018); Beffy, Blundell, Bozio, Laroque and To (2019).

[†]For estimates of employment effects of income tax holidays and the implied arc elasticities, see Bianchi, Gudmundsson and Zoega (2001); Chetty, Guren, Manoli and Weber (2012); Martinez, Saez and Siegenthaler (2021); Sigurdsson (2018).

[†]For studies of reservation wages of the unemployed, see Feldstein and Poterba (1984); Krueger and Mueller (2016); Le Barbanchon, Rathelot and Roulet (2019); Kneip, Merz and Storjohann (2020). Mas and Pallais (2019) study the employment preferences of job searchers applying to jobs at a call center. Ameriks, Briggs, Caplin, Lee, Shapiro and Tonetti (2020) do so for the retirement margin of older workers with a focus on job flexibility. Kimball and Shapiro (2008) measure income effects on labor supply to hypothetical wealth shocks in a survey.

[†]Some existing research on reservation wages of the unemployed (Feldstein and Poterba, 1984; Krueger and Mueller, 2016) has constructed the "reservation wage ratio" to describe empirical observations, but not interpreted it through the lens of an economic model or in the context of neoclassical labor supply preferences.

proportionate wage shifter, which stands in for specific experiments such as aggregate productivity shocks or linear tax reforms.

We elicit reservation raises in two representative surveys covering all three labor force groups in the U.S. and Germany. Our first survey covers 2,071 U.S. respondents as part of the AmeriSpeak Omnibus Survey run by NORC at UChicago, in the spring of 2019. Our second survey is a custom questionnaire we integrated into German Socio-Economic Panel (GSOEP) conducted in the fall of 2019, covering 3,510 individuals. We specify our survey to invoke a transitory wage change lasting one month, and hence identify the short-run aggregate labor supply curve that, e.g., speaks to business cycle fluctuations or transitory tax changes.

The two surveys yield strikingly congruent aggregate labor supply curves. In each case, the empirical distribution of reservation raises exhibits a large mass around one—where an individual's reservation wage equals her actual/potential wage. This large mass of marginal individuals generates a large *local* Frisch elasticity above 3, as implied by business cycle evidence (Hansen, 1985; Rogerson, 1988; Hagedorn and Manovskii, 2008; Ljungqvist and Sargent, 2017).

Globally, however, the empirical curves feature nonconstant arc elasticities, and considerable asymmetries. For wage *decreases*, the arc elasticities remain high. Here, considerable shares of employed workers require only moderate wage cuts to prefer temporary nonemployment. By contrast, for wage *increases*—where the curve eats into individuals out of the labor force—arc elasticities drop quickly, to around 0.5. This low value in this portion of the curve is consistent with the quasi-experimental evidence for small employment responses to large net wage increases following income tax holidays (Bianchi, Gudmundsson and Zoega, 2001; Chetty, Guren, Manoli and Weber, 2012; Martinez, Saez and Siegenthaler, 2021; Sigurdsson, 2018).

Overall, therefore, while isoelasticity is a standard assumption in empirical and modeling practice, our survey strategy reveals that no single constant elasticity would capture the global shape of the aggregate labor supply curve of either country. Both curves feature high local elasticities, which would guide business cycles, and, at the same time, low arc elasticities to large wage increases, which are relevant to, e.g., tax holidays. Moreover, we show that no existing calibrated model generates a curve that comes close to the empirical shape (although any given model could be reverse-engineered to match it).

Our survey-based research design aims to isolate *preferences* about *desired* employment status as a function of wage shifts. In the presence of labor market frictions, the labor supply curve need not perfectly guide realized employment changes (analogously to the intensive-margin argument by Keane and Rogerson, 2015; Chetty, 2012). For instance, with frictions, job loss in recessions need not follow the elastic pecking order prescribed by the reservation raise ranking, but may hit high-surplus individuals.

In Section, 2.2, we define aggregate labor supply on the basis of reservation raises. In Section 2.3, we describe the surveys, and discuss the empirical labor supply curves. In Section 2.4, we compare the empirical curves to those of existing macro models with an extensive margin. Section 2.5 concludes with questions our study leaves open.

2.2 Conceptual Framework

We define the individual-level reservation raise and show that its cumulative distribution function (CDF) gives the aggregate labor supply curve.

Individual-level Employment At the individual level, the extensive-margin (employment) status is binary, $e_{it} \in \{0,1\}$. For each individual indexed by $i \sim \mathcal{U}(0,1)$ at time t, *desired* extensive-margin labor supply can be formulated as a standard reservation wage rule. To abstract from hours choices, we cast the rule in terms of reservation earnings y^r_{it} (as an hourly wage w will be featured in Section 2.4) compared to her potential earnings y_{it}:

$$e^*_{it} = \mathbb{1}(y_{it} \geq y^r_{it})$$

This standard reservation wage rule characterizes desired employment in rich settings, including those with dynamic considerations, adjustment costs, search frictions, borrowing constraints, or human capital considerations. Hence, besides spot labor markets, it also applies in frictional models, where desired and actual employment status need not coincide (for labor supply under search frictions, see Krusell et al., 2017). This section intentionally does not spell out detailed models; we present a simple spot labor market setup and specific cases thereof in Section 2.4.

Aggregate desired employment rate E^*_t equals the fraction of individuals with $y_{it} \geq y^r_{it}$:

$$\begin{aligned} E^*_t(\cdot) &= \int_i e^*_{it} di \\ &= \int_i \mathbb{1}(y_{it} \geq y^r_{it}) di \\ &= \int_{y^r} \int_y \mathbb{1}(y \geq y^r) f^{y|y^r}(y|y^r) f^{y^r} g(y^r) dy dy^r, \end{aligned}$$

where an interior employment rate requires heterogeneity in either y_{it} or y^r_{it}, or in both.

We reformulate this standard reservation wage setup by introducing two concepts.

The Aggregate Labor Earnings Shifter First, we define an *aggregate prevailing raise* $1+\Xi_t$. It is a *homogeneous* labor income shifter of potentially *heterogeneous* baseline labor earnings y_{it}—which are always defined *gross* of this aggregate raise, so that the allocative, net-of-raise potential earnings are $(1+\Xi_t)y_{it}$. The shifter $1+\Xi_t$ operationalizes the question: how much would aggregate labor supply change if all labor earnings shifted by a percent amount given by raise $1+\Xi_t$? It stands in for specific experiments such as aggregate wage fluctuations, changes in productivity (e.g., Chang and Kim, 2006), or changes in labor taxes. (For convenience, we will refer to multiplier $1+\Xi_t$ as the raise, rather than Ξ_t.)

The Reservation (Pay) Raise Second, we define an individual's *reservation raise* $1+\xi^*_{it}$ as the *hypothetical* aggregate prevailing raise $1+\Xi_t$ that would render her *marginal*:[†]

$$(1+\xi^*_{it})y_{it} = y^r_{it} \tag{2.1}$$

$$\Leftrightarrow 1+\xi^*_{it} = \frac{y^r_{it}}{y_{it}} \tag{2.2}$$

The reservation raise is a measure of rent, or surplus, from employment *as a fraction of the idiosyncratic earnings*, and hence the individual's distance from entering or leaving employment (relative to her idiosyncratic potential earnings).

Individual-level labor supply is then a cutoff rule of the reservation vs. the prevailing aggregate raise:

$$e^*_{it} = \begin{cases} 0 & \text{if} \quad 1+\xi^*_{it} > 1+\Xi_t \\ 1 & \text{if} \quad 1+\xi^*_{it} \leq 1+\Xi_t. \end{cases} \tag{2.3}$$

Aggregate Labor Supply: the CDF of Reservation Raises Finally, aggregate labor supply can then be reformulated as a univariate function of the aggregate prevailing raise $1+\Xi_t$, with the function given by the reservation raise CDF, evaluated at a given aggregate prevailing raise, corresponding to the fraction of individuals for whom $1+\xi^*_{it} \leq 1+\Xi_t$:

$$E^*_t(1+\Xi_t; F_t) = \int \mathbb{1}(1+\xi^* \leq 1+\Xi_t) dF_t(1+\xi^*) \tag{2.4}$$

$$= \underbrace{F_t(1+\Xi_t)}_{\substack{\text{CDF of reservation raises,} \\ \text{evaluated at aggregate} \\ \text{prevailing raise } 1+\Xi_t}} \tag{2.5}$$

Comparison to the Reservation Wage The reservation raise of course simply equals the idiosyncratic reservation wage normalized by the idiosyncratic potential wage—thereby collapsing both dimensions of heterogeneity into a scalar statistic for an individual's desired employment status (to be paired with an aggregate prevailing raise).

The incremental added value of the reservation raise over the standard reservation wage concept is that it provides a standard labor supply curve: a univariate function drawing on the one-dimensional ranking of labor suppliers, that, evaluated at any aggregate prevailing raise, gives the desired aggregate employment rate. By contrast, the standard reservation wage distribution alone would not sufficiently rank individuals without simultaneous reference to their idiosyncratic potential earnings—encoded in the

[†]The lower case differentiates the micro reservation raise from the aggregate prevailing raise. The *-symbol denotes indifference, rather than a potential idiosyncratic prevailing micro raise.

joint distribution of potential and reservation wages.[†] Of course, that joint distribution does contain more information than the reservation raise distribution: the former can give the desired employment rate for any shift in the distribution of potential earnings, whereas the latter does so specifically for a homogeneous percent shift; of course, with homogeneous wages, reservation wages are sufficient to characterize extensive-margin labor supply.

Employment Adjustment Employment adjustment to a shift in the aggregate prevailing raise from $(1+\Xi_t)$ to $(1+\Xi_t')$ is driven by the mass of nearly marginal individuals (for whom $1+\Xi_t < 1+\xi_{it}^* \leq 1+\Xi_t'$), and amounts to $F_t(1+\Xi_t') - F_t(1+\Xi_t)$.

Aggregate Arc Elasticities For discrete raise changes, the arc elasticities of extensive-margin labor supply are:

$$\epsilon_{E_t,(1+\Xi_t)\to(1+\Xi_t')} = \frac{F_t(1+\Xi_t') - F_t\left(1+\Xi_t\right)}{F_t\left(1+\Xi_t\right)} \bigg/ \frac{(1+\Xi_t') - (1+\Xi_t)}{1+\Xi_t}. \tag{2.6}$$

For infinitesimal changes in $(1+\Xi_t)$, the elasticity is:

$$\epsilon_{E_t,1+\Xi_t} = \frac{(1+\Xi_t)}{E_t}\frac{\partial E_t}{\partial(1+\Xi_t)} = \frac{(1+\Xi_t)f_t(1+\Xi_t)}{F_t(1+\Xi_t)}. \tag{2.7}$$

That is, the elasticity reflects the reverse hazard rate, or inverse Mills ratio. It equals the reverse hazard rate exactly for a baseline aggregate shifter of 1.[†]

A Special Case: Constant Elasticity We can now also clarify the distributional conditions delivering *constant* elasticities, a property convenient for calibration and often assumed in modeling practice. Additionally, empirical work often thinks of a single elasticity to be measured, hence taking isoelasticity as the implicit point of departure (e.g., Chetty, Guren, Manoli and Weber, 2012). We find that isoelasticity requires a power law distribution

[†]As one example, Chang and Kim (2006) Figures 3-5 plot model-implied reservation-*wage* CDFs, generated by a model with idiosyncratic heterogeneity in productivity/wages. While Figure 5, which is the inverse CDF of reservation wages, labels the x-axis "participation," strictly speaking, this curve does not represent the aggregate labor supply curve of the underlying model, exactly because it features heterogeneity in idiosyncratic productivity/wages, so that it would only do so in a counterfactual scenario with homogeneous wages. (Chang and Kim (2007) refer to the same concept of reservation wages.) Of course, it would be easy to construct reservation raises and plot those against the specific aggregate shifter in their model, the productivity shock, and, moreover, the actual equilibrium equations and the associated simulations determine employment on the basis of the full joint distribution of heterogeneous idiosyncratic wages and reservation wages.

[†]Away from this case, a similar logic described in the general case above holds; moreover, in any setting where the baseline aggregate shifter is not normalized to one, one can alternatively redefine a net-of-Ξ_t individual-level wage level, and define the percent change in the aggregate earnings shifter. Lastly, as formalized in Equation (2.8) below when moving to the empirical implementation, the framework accommodates pre-existing non-zero levels of the aggregate wage shifters by simply normalizing the idiosyncratic reservation raise and the aggregate wage shifter by the pre-existing level of the shifter.

$G_{1+\xi^*}(1+\xi^*) = \left(\frac{1+\xi^*}{(1+\xi^*)_{\max}}\right)^{\alpha_{1+\xi^*}}$ with shape parameter $\alpha_{1+\xi^*}$ and maximum $(1+\xi^*)_{\max}$.[†] All interior arc elasticities of this reservation raise distribution are constant and equal to $\epsilon_{E_t,1+\Xi_t} = \alpha_{1+\xi^*}$ (using Equation (2.6)). The arc elasticities mechanically shrink once a pertubation is large enough to cross full nonemployment or full employment.

Heterogeneous Shifters While business cycle or tax reforms studies often consider homogeneous income shifters—which directly map into our setup—, it may also be interesting to study heterogeneous shifters, such as heterogeneous exposure to the business cycle could be accommodated by partitioning individuals into groups by shock size. Aggregate labor supply is then equal to the weighted average of the group-specific CDFs each evaluated by their respective group-specific prevailing raise, so the argument is a vector of group-specific prevailing raises. We focus on the aggregate curve in this paper, but will make available the underlying micro data of reservation raises and covariates. Of course, more broadly, the reservation raise distribution itself—and hence the aggregate labor supply curve it implies and the elasticities it features—is of course an outcome of the economic environment, the duration of a potential wage change, and the distribution of wage changes a given environment features.

2.3 Measurement

We now measure the empirical reservation raise distribution by integrating custom questionnaires into two representative surveys in the United States and in Germany. We follow three steps, mirroring the exposition in Section 2.2: (i) elicit individual-level reservation raises $1+\xi^*_{it}$; (ii) construct and plot their CDF $F_t(1+\xi^*)$, the aggregate labor supply curve; (iii) compute arc elasticities from the CDF.

Survey Implementation

We conduct two custom surveys of U.S. and German households comprising all labor force segments (aged 18 and older), of which we ask a tailored question eliciting directly their idiosyncratic reservation raises. To our knowledge, ours is the first attempt to elicit any reservation wage concepts (let alone reservation raises) from non-job-searchers (job searchers make up a selected section of the population, thereby not providing a lever on the aggregate labor supply curve, as in studies cited in Footnote †).

U.S. Survey: NORC at the University of Chicago AmeriSpeak Survey We integrate reservation-raise questions into a nationally representative survey covering 2,071 respondents

[†]Specifically, the distributional assumptions specify a standard power law distribution $F(X) = P(x < X) = a \cdot \left(x/X_{\text{min}}\right)^{-\gamma+1}$ with shape parameter $\gamma > 0$. A comparison with our reservation-raise-based power law distribution $G_{1+\xi^*}(1+\xi^*) = \left(\frac{1+\xi^*}{(1+\xi^*)_{\max}}\right)^{\alpha_{1+\xi^*}}$ clarifies that we require the *inverse* of the reservation raises to follow a power law distribution: $G_{1+\xi^*}(1+\xi^*) = P(X < 1+\xi^*) = \left(\frac{1+\xi^*}{(1+\xi^*)_{\max}}\right)^{\alpha_{1+\xi^*}} \Leftrightarrow P\left(\frac{1}{1+\xi^*} < \frac{1}{X}\right) = \left(\frac{1}{1+\xi^*} / \frac{1}{(1+\xi^*)_{\max}}\right)^{-\alpha_{1+\xi^*}}$, which is a power law distribution of $\frac{1}{1+\xi^*}$ with minimum $\frac{1}{(1+\xi^*)_{\max}}$, and shape parameter $\gamma = \alpha_{1+\xi^*} + 1$.

in the United States aged 18 and older. Our survey was fielded by NORC at the University of Chicago (henceforth "NORC," formerly the National Opinion Research Center), which, e.g., also runs the General Social Survey. NORC integrated our questionnaire in the AmeriSpeak survey program, a large probability-based panel designed to be representative of the U.S. household population, comprising around 35,000 households in 2019, who are recruited by mail, phone and face-to-face interviews. Dennis (2019) describes the AmeriSpeak sampling, recruitment and survey administration. We integrated our survey into two waves of the AmeriSpeak Omnibus program, conducted on the days following March 19th and April 19th, 2019. Each Omnibus wave draws a nationally representative sample of around 1,000 adults age 18 and older from the AmeriSpeak Panel. Interviews are conducted online and by phone. The Omnibus program is designed for shorter questionnaires such as ours and is, e.g., also used by, e.g., the AP-NORC Center for Public Affairs Research. In this survey, we elicit reservation percent changes in the wage directly, rather than the reservation earnings and actual/potential earnings separately. The survey also contains a limited set of characteristics of the respondent.

German Survey: German Socio-Economic Panel (GSOEP) Our second survey covers 3,527 individuals and is a custom questionnaire we integrated into the 2019 wave of the German Socio-Economic Panel (GSOEP). We did so as part of the SOEP Innovation Sample program, which draws on the GSOEP main sample and permits external researchers to integrate tailored questionnaires, based on an application process and collaborative design and piloting (Richter and Schupp, 2015); it is also used to pilot new permanent questions. The sample design and core fieldwork follow that of the GSOEP overall; Zweck and Glemser (2018) discusses the minor differences of the sampling method. The GSOEP is a maximally representative survey, drawing respondents at an address basis, and implementing multi-month recontact strategies to maximize response rates. The Innovation Sample respondents receive a core questionnaire besides the custom questionnaires proposed by external researchers. Zweck and Glemser (2020) describes details of the 2019 Innovation Sample round, part of which our survey was, with comprehensive information on recruitment and response rates; the full questionnaires and data will be made available by GSOEP. Importantly, the survey was fielded in the fall and winter of 2019, and completed before the onset of COVID (with the results shared with the researchers in the summer of 2020).[†] Surveys are conducted by trained interviewers, including in-person interviews, during which answers are recorded by a computer-assisted personal interview equipment. Kantar, a survey company, conducts the field work on behalf of GSOEP, as well as the programming of the survey. Our sample is again workers 18 and older. In this survey, we elicit the reservation earnings and actual/potential earnings separately, and on that basis construct the reservation raise as their ratio. This survey provides a rich set of covariates of the individual and the household, as generally

[†]The main results were obtained between September 17th, 2019 and December, 2019; 82.5% of the households completed the survey by December, 2019; 97.6% had done so by February, 2020 (see Table 2 in Zweck and Glemser, 2020).

contained in the GSOEP survey.

Ideal Measure of the Reservation Raise To fix ideas, we start with the ideal survey question that tightly mirror the theoretical reservation raise:

> You are currently [non]employed. Suppose the following thought experiment: you (and only you) receive an additional temporary linear incremental tax [or subsidy] on your take-home earnings (at whichever positive hours or job you may choose to work). At what incremental tax [or subsidy] rate would you be indifferent between working for this period and not (at whichever positive hours or job would be your best choice at that tax [subsidy] rate)?

This approach invokes an additional tax [subsidy] on top of any potentially pre-existing taxes and frictions, thereby normalizing the exactly marginal individual's reservation raise to one. We thus do not have to take a stance on the *level* of the baseline already-prevailing aggregate labor tax or tax-like factors $\widehat{\Xi}_t$, broadly defined, in the data. Formally, we would elicit a normalized reservation raise $1+\widetilde{\xi}^*_{it}$ corresponding to:

$$1+\widetilde{\xi}^*_{it} = \frac{1+\xi^*_{it}}{1+\widehat{\Xi}_t}, \tag{2.8}$$

and hence our thought experiments consider percent pertubations of the aggregate shifter $\widetilde{\Xi}_t$ around the baseline level $\widehat{\Xi}_t$, giving a normalized shifter $1+\widetilde{\Xi}_t = \frac{1+\Xi_t}{1+\widehat{\Xi}_t}$ that is centered around 1 and corresponds to the normalized reservation raise $1+\widetilde{\xi}^*_{it}$ above.

Actual Survey Implementation of Reservation Raise Measure The actual questions we implement are the result of piloting in online samples (Amazon Mechanical Turk) and iterations with survey administrators from both NORC and GSOEP. These iterations lead us to formulate relatively concrete hypotheticals compared to the aforementioned ideal question. While the ideal formulation permits job switching and reoptimization (as discussed in Section 2.4), we in practice invoke a "job-constant" perspective for a reference job.[†] We specify the frequency of the Frischian wage change to one month—balancing sufficient shortness to induce short-run, plausibly Frischian variation, and sufficient length to still capture a meaningful extensive-margin choice. We detail the questions below, review the results, and then critically discuss limitations in Section 2.3.

The (Print) Appendix presents our NORC and GSOEP reservation raise questions, separately for NORC and for GSOEP. For GSOEP, we report the English translations; (Online) Appendix B.5 reports the German original text. We specify separate questions for each of the three labor force groups (employed, unemployed, out of the labor force). In each survey, we therefore classify workers based on standard definitions about their

[†]Formally, in the setting described in Section 2.4, this yields a job-j-specific reservation raise $1+\widetilde{\xi}^*_{it,j} = \frac{v_{it,j}}{(1+\Xi_t)y_{it,j}\lambda_{it}}$ for some reference job j.

employment status as well as their search behavior and availability to work, and on that basis route them into the survey arms.[†]

We iterated the GSOEP questionnaire in collaboration with the GSOEP/Kantor survey team, and therefore differs slightly from the NORC questions, also permitting us to assess robustness to varying the specific framing, described below as well as in Section 2.3.

Question for the Employed To keep the scenario sufficiently realistic, we allude to unpaid time off in NORC as well as in our baseline scenario in GSOEP. To avoid confusion associated with job mobility (an insight from piloting), the question permits the worker to be able to return to the original job in this specification. We specify that the worker must not take a second job during this time period, to accurately capture nonemployment vs. employment trade-offs. (We do not differentiate questions for multiple-job holders.)

A potential concern is that we paint an overly specific picture about time off from work; away from spot labor markets, the implied return option may not be realistic. In GSOEP, we therefore randomly assign some employed (and unemployed) workers into a survey arm that does *not* specify the return option but brings up explicitly that take-up may require quitting (and find similar results, discussed in Section 2.3).

We elicit reservation raises slightly differently across the two surveys for the employed (and unemployed). In NORC, we directly elicit percent numbers for the reservation raises for the (un-)employed; we do not separately elicit the respondent's corresponding idiosyncratic reservation wages or (potential) earnings. In NORC, our design does not permit the (un-)employed to report positive reservation raises (which would imply a reservation wage above the actual wage). By contrast, the subsequent iterations with the GSOEP survey design team resulted in a questionnaire that separately elicits reservation earnings and actual/potential earnings (we then construct the reservation raise as ratio of the former over the latter, as in Equation (2.2)). As a result, in the GSOEP survey, employed respondents can report requiring a pay raise not to temporarily separate. As we discuss below, we will find few such observations (perhaps reflecting limited opportunities for time off, or measurement error).

Question for the Unemployed While previous work has measured reservation wages of the unemployed, (e.g., Feldstein and Poterba, 1984; Krueger and Mueller, 2016), our comprehensive coverage of all labor force groups requires us to keep the question for the unemployed comparable to the other two groups'. In NORC, we induce a scenario in

[†]In NORC, we define the three labor force statuses as follows: we use the variable on employment status ("EMPLOY") to partition respondents into the employed (working as an employee, self-employed, or on temporary layoff), unemployed (not working but looking for work) and out of the labor force (not working for retirement, disability, or other reasons). In GSOEP, we define the three labor force statuses as follows: we use the variable on employment status to partition respondents into the employed (("PERW" 1–7, including apprentices and part-time); we then split up the nonemployed ("PERW" 9) into the unemployed ("PNERW02" 1–2, i.e., likely or certain to take up work), and the out of the labor force ("PNERW02" 3–4, i.e., sure to not or unlikely to take up work). Our NORC questionnaire features an additional variant of the question for the temporarily laid off that mirrors that of the employed (supposing the respondent is back at the previous job). We do not ask the self-employed, given the missing wage concept.

which a prospective job permits a one-month earlier start date, albeit at a wage reduction for that month. The particular reason is left unspecified, although we clarify that this interim month is to be spent in nonemployment. In GSOEP, we evoke a situation after job acceptance, and ask the respondent to reflect on the question identical to the employed described above (after a short preamble).

Since the unemployed will want to work, we expect the reservation raise—which reflects the *desired* employment status—to be at most one, as for the employed. In NORC, where as for the employed, the respondents report reservation pay cuts. In GSOEP, we again separately elicit reservation and potential earnings (and take their ratio), and here therefore permit the unemployed to report reservation raises above one (which we again find few unemployed will give).

Question for the Out of the Labor Force By self-classification and revealed preference, the out of the labor force likely have reservation wages exceeding their expected potential wages. So for this group, we ask about the required wage *increase* to induce a respondent into employment, for a concrete job that they envision they could realistically be offered if they searched and did attempt to take up employment. Crucially, for our Frischian perspective, this wage change is supposed to only occur for a single month. For concreteness and realism, we implement this scenario in the form of a sign-up bonus on top of the first-month salary. We also specify that the employment relationship is to last for at least one month.

Naturally, the out of the labor force individuals include those least likely to consider taking up employment (including the disabled, the retired, or students), who may hence rarely think about labor markets. However, the out of the labor force do appear to contain some marginal individuals (as evidenced by the worker flows in and out of the labor force, as documented in, e.g., Davis, Faberman and Haltiwanger, 2006). Moreover, to achieve large employment increases to large responses to, e.g., tax holidays, it is the out of the labor force that would need to be crowded in. The reservation raises identifies those marginal individuals.

Response Rates We have high response rate of 80% for NORC and 70% for GSOEP (here defined as respondents giving nonmissing answers out of the participants). Appendix Table B.4 details those numbers, separately by labor force status, for NORC and GSOEP. While the numbers are not directly comparable, the response rates dramatically exceed those in reservation wage surveys of the unemployed, which are around 10% (see, e.g., Feldstein and Poterba, 1984; Krueger and Mueller, 2016). We discuss residual potential effect of missing information below in Section 2.3.

Our NORC survey covered 2,071 individuals (minus 13 for whom we were unable to assign a labor force status, so they were not asked any subsequent question). For 82% (1,679; 809 in March, and 870 in April) of the NORC participants, we have non-missing reservation raise information.

Our GSOEP questionnaire covered 3,527 individuals (minus 17 respondents without labor force status information). Among those, 70% (2,431) participants have non-missing reservation raises. (In the vast majority of missing observations, both reservation and actual/potential wages are missing). We further drop 164 individuals for whom survey weights are missing.[†]

Covariates and Weighting We present summary statistics for the observations with non-missing reservation raises in Table B.1. We present the numbers for the total sample, as well as the analysis sample with nonmissing reservation raises.

In NORC, we weight observations within each labor force status using the accompanying sample probability weights (to match the American adult population, although the survey is designed to be representative). We also rescale the weights in each wave to represent the proportion of the total sample obtained from each wave, although those were similar (see above). The raw sample was close to the February 2019 BLS population shares for employment, labor force participation, and unemployment (see Table B.1 Panel A); to precisely match that important target of our data, we finally reweight the observations with non-missing reservation raises so that the weighted labor force status proportions precisely match the BLS target.[†] In GSOEP, we again use the sampling weights (which made very little difference), and additionally reweight the observations with nonmissing reservation raises to match the shares of the labor force groups in the data in 2019.[†]

Results

Histograms We present histograms of the empirical reservation raises from the reported reservation raises in Figure B.1 Panel (a) for NORC and in Panel (b) for GSOEP. Differential shading separates observations by labor force status.

For both surveys, the empirical histogram of the reservation raise distribution exhibits a large mass around one—where the reservation wage is close to the individual's actual or potential wage, i.e., the location of marginal individuals. To the left, for wage reductions, the employed and unemployed would be crowded into nonemployment; to the right, for wage increases, labor supply would recruit the out of the labor force individuals into employment (strictly so for NORC, and approximately so for GSOEP, discussed above as well as below in Section 2.3).

Globally, however, the distribution is widely dispersed, as most individual derive considerable and tremendously heterogeneous surplus (or, in the case of the out of the

[†]Based on correspondence with the GSOEP survey team, the weights for those survey entrants will become available in 2021 as part of the 2020 data delivery (wave); we will then update our results.

[†]The March 2019 BLS targets give 60.7% (employed), 2.4% (unemployed) and 36.9% (out of the labor force), given by 60.7% employment to population ratio (source: https://fred.stlouisfed.org/series/EMRATIO) and 63.1% labor force participation rate (source: https://fred.stlouisfed.org/series/CIVPART).

[†]For instance, the 2019 labor force participation rate in Germany was 61.9% according to OECD statistics (age 15 and up; 18 and up not available); in our GSOEP sample, it is 61.14% (in the whole dataset) and 61.06% (among those that were asked our questions). In GSOEP, we have also experimented with dropping low earners, and have found similar aggregate labor supply curves and elasticities.

labor force, would suffer considerable net disutility) from employment. For visual clarity, we bunch raises above 2.0 into the 2.0 group (on a secondary y-axis in the histogram).

Lastly, the NORC—but not the GSOEP—histogram exhibits some likely spurious mass points at 0.5 and 1.5, perhaps due to respondents' rounding; we conjecture that smoothing out those bunching points would spread out more evenly would distribute mass towards a locally more elastic and far-away less elastic curve, thereby further accentuating the asymmetries already present. We discuss this and other limitations of the survey in Section 2.3.

Aggregate Labor Supply Curves To trace out the aggregate labor supply curve, we aggregate the micro reservation raises into a cumulative distribution function (CDF) $F(1+\xi^*)$, plotted in Figure B.2 Panel (a). The curve gives the desired employment rate as a function of any given prevailing raise $1+\Xi$. (The empirical reservation raises are measured as the normalized-around-one baseline raise $1+\widetilde{\xi}^*$ defined in Equation (2.8). Since population size is fixed, employment and employment to population ratio elasticities are equal.)

To facilitate visual inspection with regards to elasticities, we additionally take logs of both axes and normalize the employment rates at the baseline level, thereby plotting changes in desired log employment against changes in $\log(1+\Xi)$. We do so in Panel (b) of Figure B.2. This plot zooms into the local range around 0.05 upwards and downwards.

Arc Elasticities We construct a set of arc elasticities over varying aggregate prevailing raise deviations: the share of the population in a given upward or downward distance Ξ' from the prevailing unit raise $1+\Xi=1$, following the definition in Equation (2.6), $\epsilon_{E,(1+\Xi)\to(1+\Xi')} = \frac{F\left(1+\Xi'\right)-F(1+\Xi)}{F(1+\Xi)} \Big/ \frac{(1+\Xi')-(1+\Xi)}{1+\Xi}$. Appendix B.6 details the calculation and the treatment of marginal individuals. Table B.2 reports these arc elasticities (along with the shares of observations). Figure B.3 visualizes the resulting arc elasticities (for the range of 0.20 upwards and downwards).

Large Local Elasticities *Locally*, i.e., for small shifts, we find large elasticities of around 3, and even higher values for tiny shifts. That is, on both sides, lots of individuals prefer to move in or out of employment in response to small percent wage changes. Considering the NORC results in Table B.2, we find that a local 1% increase in the aggregate prevailing raise crowds in nearly 2.26 percent of additional employment (implying an elasticity of $\frac{d(\text{Emp/Pop})}{\text{Emp/Pop}}/0.01 = \frac{0.0226}{0.631}/0.01 = 3.72$). A 1% decrease implies an even larger elasticity of 5.66. Similarly high local elasticities emerge for GSOEP (2.86 and 9.38).

The small shifts upward and downward are those that would drive business cycle fluctuations in employment in equilibrium models, where shifts in labor productivity and hence wages are small (e.g., a quarterly standard deviation of around 2% as in Hansen, 1985). For this reason, many macro models require large Frisch elasticities (Chetty, Guren, Manoli and Weber, 2012). Locally, the concentration of marginal individuals paints such a highly elastic picture in the survey, mirroring intuitions from models of indivisible labor and worker homogeneity (Hansen, 1985; Rogerson, 1988).

Nonconstancy: Smaller Elasticites for Large Shifts Nonlocal perturbations, to large wage changes, imply dramatically lower arc elasticities. Compared to the high local elasticities of 3.72 to a 1% increase, for instance, the arc elasticity falls to 0.96 when considering a larger raise of 10%. Downward, the arc elasticity falls from 5.66 for the 1% wage decrease to 1.68 for a 10% decrease. For GSOEP, a strikingly similar picture emerges, with elasticities to large changes being even somewhat lower throughout, at 0.41 and 1.22 for a 10% increase and decrease, respectively. The nonconstant elasticities are salient in the arc elasticities plot in Figure B.3. Arc elasticities are largest locally around the baseline prevailing raise, and shrink for larger perturbations.

Asymmetry: Smaller Arc Elasticities Far Upward than Far Downward Not only do the curves exhibit nonconstant arc elasticities, but also an asymmetry: arc elasticities stay relatively high downward, as the labor market continues to find employed workers ready to switch into nonemployment. But upward, the out of the labor force appear hard to recruit into employment. This pattern emerges in NORC, but is if anything more pronounced in GSOEP.

An Implication: External Validity of Specific Arc Elasticity Estimates in the Presence of Nonconstant Elasticities Figure B.3 suggests that a constant elasticity would not provide a realistic description of the *global* aggregate extensive-margin labor supply curve. As one concrete implication, the empirical curve suggests that the small arc elasticities identified by large positive increases in net wages may mask large local elasticities. For example, Chetty, Guren, Manoli and Weber (2012) infer a 0.42 Frischian extensive-margin labor supply elasticity by interpreting employment responses to the tax holiday in Iceland studied by Bianchi, Gudmundsson and Zoega (2001), which reduced average tax rates from 14.5% to 0% for one year.[†] In our framework, this experiment corresponds to an increase in $1 + \Xi_t$ from 1.00 to 1.17.

Our survey-implied labor supply curves accommodate this estimate, as it features an arc elasticity of 0.60 for that large an upward raise shift in NORC, and even lower in GSOEP. At the same time, however, in the *global* curves the surveys imply, this small arc elasticity to a large upward shift masks dramatically larger local elasticities.[†] Hence,

[†] Another quasi-experiment reviewed in Chetty, Guren, Manoli and Weber (2012) is the Self Sufficiency Program in Canada, studied by Card and Hyslop (2005), which raised average net of tax rates from 0.25 to 0.83, for 36 months, with an implied employment elasticity of 0.38.

[†] To some degree, the nonconstant elasticity is of course expected, as the employment rate cannot exceed 100%. A priori, the large macro elasticity benchmarks of around 2.5 cited by Chetty, Guren, Manoli and Weber (2012) for cyclical macro contexts would, out of a baseline employment rate of 79.2% in their Icelandic example of a tax holiday, imply employment rates exceeding 100%, similarly for some of the other case studies with large net-of-tax increases the authors discuss. Of course, in the case studies the empirical employment rates do not reach 100% in response to the subsidies, and therefore do not actually hit the full-employment constraint. By contrast, Martinez, Saez and Siegenthaler (2021) also study a large tax holiday, in Switzerland, and find no treatment effects on employment rates, which therefore implies small

such low estimated specific arc elasticities with respect to large upward net-of-tax wage increases need not provide tight bounds on the arc elasticities in the local portions of the curve, which are those relevant to business cyclical fluctuations.

More generally, nonconstant arc elasticities also imply a trade-off between statistical power and overcoming adjustment costs (e.g., Chetty et al., 2011; Chetty, 2012), and measuring the local elasticities relevant for smaller shocks—unless one is willing to maintain the pervasive assumption of isoelasticity, which however our survey-implied labor supply curves imply appears counterfactual.

Robustness Checks and Limitations

Response Quality As with contingent valuation surveys more generally, and specifically standard reservation wage measures among the unemployed, our survey measures may not accurately capture preferences. On the one hand, idiosyncratic noise in the responses raises would generate spurious dispersion, and hence bias downward the measured elasticities. On the other hand, local elasticities would be overestimated with spurious bunching around 1. Indeed, the mass points in the NORC survey at 0.5 and 1.5 reflect bunching at semi-round numbers. However, the NORC mass around 1.0 reflects a healthily spread-out mass, making it unlikely that sharp and strict bunching drives the result. Most importantly, the GSOEP does not feature such mass points, while otherwise featuring a similar curve overall. The absence of bunching in the GSOEP may reflect higher quality responses. Or, it may reflect the design difference in that in the GSOEP, we elicit the potential and reservation earnings separately.

Comparison to Existing Evidence from Unemployed Job Seekers The local mass of marginal individuals is qualitatively consistent with existing evidence from surveys of the unemployed. Some empirical studies of the reservation wages of the unemployed have constructed the "reservation wage ratio" as an informal normalization (Feldstein and Poterba, 1984; Krueger and Mueller, 2016), revealing that the unemployed state on average high reservation wages relative to their wages—which has been interpreted as implausible (see, e.g.. Shimer and Werning, 2007, p. 1160). However, in our setting, such properties need not indicate bugs but may be features consistent with the unemployed comprising mostly marginal individuals (see Figure B.1 Panels (a) and (b)). Moreover, recent studies with high-quality survey data on reservation wages and larger samples have clarified that even the unemployed report considerable gaps between their reservation wage and the past wage (see, e.g., the histogram in Figure 2 Panel A in Le Barbanchon, Rathelot and Roulet, 2021). The discrepancy may be due to the fact that the evidence in Feldstein and Poterba (1984); Krueger and Mueller (2016) stems from recessionary periods, and relatively low response rates of around or below 10%. Moreover, in our survey implementation, we do not use the past wage as a proxy for the reemployment wage, but evoke a scenario that holds fixed a specific, current or prospective, job. An alternative route would be

elasticities across all intermediate arcs.

to validate the labor supply preferences by studying covariates or realized previous and future employment behavior in the surveys.

Adjustment Frictions Our baseline survey formulation in particular for the employed evokes a spot-market scenario without adjustment frictions. For the employed, a post-nonemployment return to work appears at least implicitly permitted. This scenario may lead employed workers to overstate their reservation raises compared to a scenario in which such return is either not possible or would entail, e.g., losses in wages, skill, or job stability.

We have assessed the relevance of this feature in the GSOEP survey. We have randomly allocated, in a 50/50 proportion, the employed and unemployed into two survey arms: one that deliberately did *not* specify the return option—and instead leaves to the worker to consider whether the month nonemployed may require quitting and a subsequent job switch. The other half was presented with the baseline formulation. The (Print) Appendix lists the supplementary survey questions; Online Appendix B.5 lists the associated German original text.

In Appendix Figure B.6, we replicate the reservation raise distribution using only one of the two survey arms (and accordingly reweight the employed and unemployed doubly). The curves and associated arc elasticities are strikingly similar. The robust pattern implies that at least in this specification, the evocation of the seemingly frictionless setting does not drive the large mass of marginal individuals. The congruence of the two curves depicted in Appendix Figure B.6 also permits us to pool both survey arms for the employed and unemployed, which, in fact, the GSOEP distributions throughout the paper have done.

Ultimately, beyond the survey, any such discomfort extends to the standard, predominant neoclassical labor supply and spot labor markets more generally, perhaps in favor of approaches that dissect labor-supply-like behavior in search-frictional settings (see, e.g. Hall, 2009; Krusell, Mukoyama, Rogerson and Sahin, 2017).

Rationed Labor Supply of the Employed Relatedly, it is conceivable that even some employed respondents are overemployed: they may prefer to be (temporarily) nonemployed in a given month, but adjustment frictions prop up their realized employment status. That is, their reservation raise is above one. In the NORC survey, reservation raises above one are not permitted for the employed (or unemployed), as we phrase their questions explicitly as a wage reduction. Still, the histogram suggests that this concern is of limited relevance: there is no sharp bunching at the maximal values among the employed in NORC, but values below 0.99 remain high. In the GSOEP, we separately elicit reservation earnings for the job, and divide by actual earnings to construct the reservation raise. Hence, employed GSOEP respondents can give raises above 1.0. Inspecting the GSOEP histogram reveals only a small fraction (around 15%) of the employed (or unemployed) workers giving such answers, with limited spread. If anything, if we were to move those workers into the employed group (and declared them marginal by winsorizing their raises down to 1.0), we would obtain a higher elasticity downward and a faster decline upward, but still a high upward local elasticity given by the marginal individuals out of the labor force. Appendix B.6 presents the detailed discussion of these issues and presents that calibration. Overall, we therefore conclude that our treatment of potentially overemployed respondents does not drive our main results.

Duration We set the duration of the wage perturbation to one month, balancing sufficient shortness to plausibly induce short-run (e.g., Frischian) variation and sufficient length to capture a meaningful extensive-margin choice. An interesting extension would be to study longer-lasting deviations. On the one hand, potential wealth effects grow with duration. (In Section 2.4, we find that for the calibrated models, uncompensated curves are essentially identical to Frischian ones even for quarter-long durations.) On the other hand, longer durations help overcome adjustment costs (which we however deemphasize).

Missing Observations We can gauge and bound the potential effects of observations with missing reservation raises on measured elasticities by considering three benchmark cases. First, if observations were missing-at-random, all results would stay the same—which is, implicitly, the assumption we have made by studying the non-missing observations. Second, if all missing individuals were marginal (i.e., have reservation raises within the local range for which we construct arc elasticities), we would of course currently underestimate the local elasticities. Third, since we measure relatively high elasticities, the most interesting alternative case to quantify is the extreme case if all missing observations were perfectly inframarginal. Then, we would currently overestimate local elasticities. We can quantify the bound this overestimate as follows. Formally, the latent, population-level distribution $G(1+\xi^*) = (1-m) \cdot F(1+\xi^*) + m \cdot H(1+\xi^*)$ consists of those of non-missing and missing observations, $F(1+\xi^*)$ and $H(1+\xi^*)$, where m denotes the share of missing observations. In the extreme case in which all missings are inframaginal, the density of $H(1+\xi^*)$ is zero in the local intervals we consider. Then, we arrive at the value of the population elasticities by adjusting the measured ones by $1-m$, i.e., one minus the share of missing observations. This adjustment factor $1-m$ would be 80% for NORC and 70% for GSOEP (as the shares of missing observations, m, are 20% and 30% respectively, as discussed above), such that it only moderately compresses the original elasticities.[†] Hence, the treatment of missing reservation raise observations cannot drive our main results.

Snapshot Our surveys elicit a snapshot of the labor supply curve for one cross-section representative of the U.S. and German populations each. The shape of the curve may vary over time, so it would be interesting to elicit the reservation raises in many repeated cross sections or even in a panel of workers. Unfortunately, the labor market upheaval following the pandemic prevented meaningful follow-up studies in 2020 and 2021.

Transfers and Nonemployment Subsidies Our survey questions do not explicitly specify the possibility of transfers or nonemployment subsidies such as unemployment insurance (UI) benefits. It is difficult to extend the concrete institutional features of the UI system into

[†] Appendix Table B.4 also separates the missings by labor force status, additionally permitting the reader to gauge an asymmetric adjustment.

a neoclassical model of labor supply, which lacks a notion of voluntary and involuntary separations that underlie the eligibility for UI in practice. (In the US, workers that quit are not eligible for UI *de jure*, while in Germany, a waiting period for unilateral quits exists that would exceed the one-month spell we evoke.) We can qualitatively consider the potential scenarios of mismatch between our survey design and the empirical context for which we construct the labor supply curve. Our focus is on the employed, and hence the downward direction, on separations in response to negative shocks. First, we suppose that respondents ignore UI in their responses, but that in practice their separations of interest would be eligible for UI. In that case, we expect an even larger mass of workers on the margin, and hence a higher downward elasticity. Second, suppose that respondents have UI in mind when contemplating the one-month separation, but in practice would not be eligible. (We deem this scenario less likely, as we phrase the question closer to a quit.) Then, for real-world decisions that would leave workers ineligible, some of the marginal workers would require a larger wage cut to prefer to quit, reducing the mass of marginal workers. Ultimately, while we suspect that workers are not likely to have UI on their mind when quitting for institutional reasons, our paper leaves open this possibility. As one piece of evidence, Appendix Figure B.6 indicates that in Germany, the scenario in which we do not permit a reemployment possibility (hence perhaps even less likely to give UI eligibility) yields similar curves than the scenario in which we phrase the setting closer to a vacation (even less likely to evoke UI eligibility); an alternative design may exploit heterogeneity in knowledge about UI. Ultimately, we our existing survey questions leave this question for future research.

2.4 Comparison with Model-Implied Curves

We now show that the aggregate labor supply curves of various macro models do not match the global empirical one. For each model, we (i) construct the individual-level reservation raise $1+\xi^*_{it}$; (ii) compute and plot its (steady state) reservation raise distribution $F_t(1+\xi^*)$ (the aggregate labor supply curve), and (iii) compute its arc elasticities. Specifically, we study a representative household with constant Frisch elasticities, a finitely lived atomistic household including an intensive margin, and a heterogeneous agents with wage shocks and incomplete markets.

Leading Case: Frischian Labor Supply in Spot Labor Market

We now specialize the general framework presented in Section 2.2 to a spot labor market. We consider a Frischian context, because it does not require specifying the temporal dimension of the wage shift, because the Frisch elasticity is a key focus of the literature, and to streamline the exposition.

General Setting The labor supply blocks we study are set in spot labor markets. Consider an individual i with time-separable utility $u_i(c_{it}, h_{it})$ from consumption c_{it} and hours worked h_{it}, with budget Lagrange multiplier λ_{it}, and assets a_{it} earning interest rate r_{t-1}:

$$\max_{a_{it},h_{it},c_{it}} \mathbb{E}_t \sum_{s=t}^{t_i^{\max}} \beta^{s-t} u_i(h_{is}, c_{is}) \tag{2.9}$$

$$\text{s.t. } a_{is} + c_{is} \leq a_{i,s-1}(1 + r_{s-1}) + (1 + \Xi_s)\theta_{is}(h_{is}) \;\; \forall t_i^{\max} \geq s \geq t. \tag{2.10}$$

Gross-of-$(1 + \Xi_s)$ earnings at a given hours choice are $\theta_{it}(h_{it})$, for example, a standard linear wage schedule $\theta_{it}(h_{it}) = w_{it}h_{it}$.

Frischian Labor Supply, Indivisible Labor, and Separable Utility We now study the leading case, which will map most closely into the specific models we study below. First, we specialize to separable utility between consumption and labor/leisure, such that $u_i(h_{is}, c_{is}) = u_i^c(c_{is}) - u_i^h(h_{is})$; we discuss nonseparabilities below. Second, labor is indivisible, such that $h_{it} \in \{0, \widetilde{h}_{it}\}$; we permit intensive-margin hours choices below. Third, we study pertubations in the aggregate prevailing wedge that are Frischian, i.e., that leave λ_{it} constant; we permit wealth effects in Section 2.4.

The discrete employment choice compares costs and benefits of working. Working comes at labor supply disutility $v_{it} = u_i^h(0) - u_i^h(\widetilde{h}_{it})$. v_{it} may also include fixed participation costs (Cogan, 1981). On the benefit side, the worker obtains potential earnings $y_{it} = \theta_{it}(\widetilde{h}_{it})$ (and zero otherwise, although the monetary opportunity cost may involve, e.g., unemployment insurance, discussed below).

Optimal labor supply assigns each individual i her desired hours $h_{it}^* \in \{0, \widetilde{h}_{it}\}$, a binary discrete choice due to indivisible labor, according to a cutoff rule—equivalently, it determines the desired employment status $e_{it}^* \in \{0, 1\}$:

$$h_{it}^* = \begin{cases} 0 & \text{if } \; (1+\Xi_t)\theta_{it}(\widetilde{h}_{it})\lambda_{it} < v_{it} \\ \widetilde{h}_{it} & \text{if } \; (1+\Xi_t)\theta_{it}(\widetilde{h}_{it})\lambda_{it} \geq v_{it}. \end{cases} \Leftrightarrow e_{it}^* = \begin{cases} 0 & \text{if } \; (1+\Xi_t)y_{it}\lambda_{it} < v_{it} \\ 1 & \text{if } \; (1+\Xi_t)y_{it}\lambda_{it} \geq v_{it}. \end{cases} \tag{2.11}$$

That is, an individual prefers employment if the benefits, $(1 + \Xi_t)y_{it}\lambda_{it}$, outweigh the cost, v_{it} (such the post-raise earnings exceed the extensive-margin MRS). For marginal—i.e., indifferent—individuals, the condition holds with equality.

The Frischian Reservation Raise with Indivisible Labor in a Spot Labor Market Here, the Frischian (λ-constant) reservation raise $1 + \xi_{it}^*$ for individual i captures the *hypothetical* aggregate prevailing raise $1 + \Xi_t$ that would render her indifferent:

$$1 + \xi_{it}^* \equiv \frac{v_{it}}{y_{it}\lambda_{it}}. \tag{2.12}$$

Here, the reservation raise encodes three elements: potential labor earnings y_{it}, budget multiplier λ_{it}, and labor disutility v_{it}. These elements, in turn, may capture rich model-specific sources of heterogeneity, such as in wealth, borrowing constraints, skills, hours requirements, job amenities, time endowments, or tastes for leisure.

Intensive Margin The approach accommodates intensive-margin choices. Rather than a binary choice set $h_{it} \in \{0, \widetilde{h}_{it}\}$, suppose now a choice of job j with attributes $(y_{it,j}, v_{it,j})$ (nesting hours differences only) from a job menu $J_{it} = \{(y_{it,j}, v_{it,j})\}_j$. Here, the reservation raise is implicitly defined, as the prevailing raise achieving indifference between working and not—conditional on having reoptimized job choice with respect to that raise.[†]

Nonseparable Preferences and Other Components of the Opportunity Cost of Employment In principle, the reservation raise accommodates richer preference and market structures, such as unemployment insurance, nonseparable preferences, or even search frictions and long-term jobs. Such additional terms are featured in the opportunity cost of employment in the context of search and matching models in representative households (as in Hall and Milgrom, 2008; Chodorow-Reich and Karabarbounis, 2016). A variant, simplified to a spot labor market setting, applies here even with atomistic households, and would be $\widehat{v}_{it} = u_i(0, c(0, \lambda_{it})) - u_i(\widetilde{h}_{it}, c(\widetilde{h}_{it}, \lambda_{it})) + \lambda_{it} \cdot (b_{it} - (c(0, \lambda_{it}) - c(\widetilde{h}_{it}, \lambda_{it}))$, where $c(h, \lambda)$ is the consumption level associated with hours choice h and multiplier λ, and b is a nonemployment subsidy such as unemployment insurance benefits. The models reviewed below will not feature any such additional properties.

Frischian Curves from Specific Macro Models

We plot the reservation raise CDFs and arc elasticities of specific models (in logs and normalized to 0 in steady state on both axes), along with the empirical ones, in Figure B.4. We report arc elasticities for various intervals in Table B.2, as with the survey statistics.

Method Details for each model and the calibrations are in Appendix Section B.7. We parameterize each model so that its steady state employment rate is 60.7%, as in the U.S. 16+ civilian employment to population ratio in February 2019 from the BLS (FRED series EMRATIO), similar to the NORC survey.[†] In each model, we normalize the steady steady prevailing rate raise (including potential taxes) to one. We extract the reservation raise distributions from the steady state equilibrium.

Representative Household Models with Full Insurance

A common specification of aggregate labor supply appeals to a large representative household comprised of a unit mass of individual members, with consumption levels and employment statuses assigned by the utilitarian head (Galí, 2011) or by incentive-compatible lotteries (Hansen, 1985; Rogerson, 1988). Full (cross-sectional) insurance and the pooled budget constraint imply homogeneous $\lambda_{it} = \bar{\lambda}_t$. We consider two canonical cases.

[†]Formally, the "inner loop" gives the optimal intensive-margin job choice conditional on any prevailing raise $1 + \Xi_t$ while ignoring the participation constraint: $j^*(1 + \Xi_t) = \text{argmax}_{j \in J_{it}}\{u(.); \text{s.t. BC}|1 + \Xi_t\}$. Second, the "outer loop" implicitly defines the extensive-margin indifference point $1 + \xi^*_{it} = v_{it,j^*(1+\xi^*_{it})}/(y_{it,j^*(1+\xi^*_{it})}\lambda_{it})$.

[†]Rather than restricting the sample to the prime working age population, we target a fuller population definition because our surveys target individuals 18 and older without an upper age limit.

Homogeneity (Hansen, 1985) The perfect homogeneity model of Hansen (1985) yields a degenerate reservation raise distribution and hence corner cases of employment out of steady state. Qualitatively, the high local elasticity in the data mirrors these intuitions, but in an attenuated way; away from the local mass, the empirical reservation raises exhibit tremendous heterogeneity and hence lower arc elasticities.

Isoelasticities (MaCurdy, 1981) A convenient specification with heterogeneity is in the employment disutility, specifically in the parametric way that delivers a constant elasticity, as derived by Galí (2011). We include two 0.32 and 2.5 isoelasticity cases, following Chetty, Guren, Manoli and Weber (2012), who propose 0.32 as the average of quasi-experimental estimates, and 2.5 as that implied by business cycle evidence. For small changes, the empirical arc elasticities are closer to the large isoelasticity. For larger, in particular positive pertubations, the data exhibit smaller arc elasticities towards 0.50, closer to the 0.32 isoelasticity. Hence, neither isoelastic case—in fact, none—accurately describes the global empirical curve.

Heterogeneous Agent Model

In heterogeneous agent models, atomistic individuals with separate budget constraints make individual-level choices. Heterogeneity arises from stochastic wages, which pass through into budget constraints under incomplete markets, and thence into assets, consumption, and λ_{it}. To study this setting, we introduce indivisible labor into the Huggett (1993) model as in Chang and Kim (2006, 2007), and calibrate the 33-state potential-earnings process to mimic that in Kaplan, Moll and Violante (2018) (whose model features only intensive-margin labor supply), which in turn approximates the empirical earnings dynamics documented in Guvenen, Karahan, Ozkan and Song (2015).

Baseline The model generates small local labor supply elasticities (0.12–0.31) upward, but exhibits larger (up to 0.72) elasticities downward, albeit quickly settling in below 0.5 for large pertubations towards 0.10. Yet, quantitatively, the elasticities are too small throughout, although for positive shifts, the arc elasticity gradient asymptotes towards the 0.32 benchmark proposed by Chetty, Guren, Manoli and Weber (2012).

The Role of Incomplete Insurance Since the equilibrium reservation raise distribution inherits the joint distribution of λ and y, the curve is inelastic if low earnings realizations are offset by high λ values. Incomplete markets generate exactly this negative covariance. To see this, we also plot the curve under complete markets—which generate a homogeneous λ.[†] This curve is dramatically more elastic, especially for large downward perturbations.[†]

[†]The underlying sparse discrete Markov process (chosen for computation reasons) would render the full-insurance curve choppy, otherwise smoothed by the asset distribution. For visual clarity, we can here (since λ is homogeneous) instead plot the reservation raise distribution arising from *continuous* earnings process (which Kaplan, Moll and Violante (2018) discretize).

[†]This distributional intuition at the extensive margin differs from incomplete markets attenuating labor supply elasticities at the intensive margin (as in Domeij and Floden, 2006) and from λ shifting with wealth shocks in non-Frischian settings (which we find has a small effect below in Section 2.4).

This exercise illustrates how the reservation raises can serve as a diagnostic tool for the complex labor-supply implications of richer asset market structures.

Lifecycle and Intensive Margin

A model with both intensive and extensive margins is that by Rogerson and Wallenius (2009), which also features lifecycle patterns (studied by Chetty, Guren, Manoli and Weber, 2012, as a leading macro model with an extensive margin, whose parameterization we largely follow).

Baseline The calibrated economy exhibits a high local elasticity. In the upwards direction, it generates a nearly constant elasticity, mirroring the 2.5 isoelasticity line. Arc elasticities range from 2.60 to 3.20, with local elasticities (from 0.01 raise pertubations) between 2.84 and 2.90.[†] Qualitatively, the model generates some asymmetry, but quantitatively, the model misses the steep decline towards 0.5 in the elasticities upwards.

The Role of the Intensive Margin To assess the importance of intensive-margin reoptimization on extensive-margin labor supply, as discussed in Section 2.4, we also plot a second curve, which instead holds hours fixed at the baseline optimal choice. Intuitively, intensive-margin reoptimization weakly raises the benefit of working, and so the flexible-hours curve weakly exceeds the fixed-hours one, but not by much.

The Role of the Wage-Age Profile In the model, wages are a triangular function of age, a convenient but consequential choice. To show this, we recalibrate the wage-age gradient around the marginal ages (labor force entry and exit) while targeting a lower Frisch elasticity, by allowing a higher level of peak lifetime productivity and a steeper slope of the wage-age productivity gradient. While the elasticities fall by around half locally, the global fit remains off. Ultimately, as with the other models, matching (reverse-engineering) the empirical curve globally would require more complex functional forms.

Non-Frischian, Uncompensated Variation

We finally quantitatively evaluate the divergence between Frischian and uncompensated model curves. For each baseline model, we simulate an unexpected aggregate-raise perturbation lasting for one quarter, a useful horizon for business-cycle frequencies, and permit λ adjustments through wealth effects. Computational details are in Appendix Section B.7. Appendix Figure B.8 shows that the uncompensated curves are close to their Frischian counterparts. Larger divergence may arise with richer asset structures such as illiquid assets and adjustment costs therein (as modeled in Kaplan, Violante and Weidner, 2014; Kaplan, Moll and Violante, 2018, which feature intensive margins only).

[†]Consistent with our global clarification, Chetty et al. (2012), who simulate reforms of specific large tax reductions in the model, find it to exhibit large Frisch elasticities.

2.5 Open Questions

We close by highlighting two questions beyond the scope of our paper, which has focused on descriptive measurement of employment preferences.

First, our descriptive exercise leaves open which deep *sources* of heterogeneity or equilibrium mechanisms drive the asymmetric and locally elastic shape of the empirical curve. While specific models can be reverse-engineered to match the empirical curve, and are hence isomorphic from the perspective of aggregate labor supply, observable attributes associated with the micro reservation raises may adjudicate between specific models.

Second, the labor supply curve represents *preferences* over desired labor supply. In the presence of frictions, even a highly elastic pecking order implied by preferences need not guide *realized* employment fluctuations.[†] For some applications, such as predicting the effect of tax reforms, reduced-form elasticities on the basis of realized employment adjustment may be sufficient. Assessing welfare or developing models of the aggregate labor market require the separation of frictions and preferences.

[†]Krusell et al. (2017) present a model of labor supply with search frictions. Empirically diagnosing the efficiency properties of employment adjustment is challenging (see, e.g., Bils, Chang and Kim, 2012; Jäger, Schoefer and Zweimüller, 2021).

Chapter 3

Unemployment Effects of Stay-At-Home Orders: Evidence from High Frequency Claims Data†

Joint with ChaeWon Baek, Peter McCrory, and Todd Messer

3.1 Introduction

To limit the spread and severity of the COVID-19 pandemic, officials around the globe turned to non-pharmaceutical interventions (NPIs), such as shutting down schools, restricting economic activities to those deemed essential, and requiring people to remain at home whenever possible. In mid-March 2020, Ferguson et al. (2020) issued a report projecting that, in the absence of the effective implementation of NPI mitigation strategies, more than 2 million Americans were potentially at risk of death from the COVID-19 respiratory disease, with many more facing uncertain medical complications in the near- and long-run.

Soon after, state and local officials in the United States began announcing Stay-at-Home (SAH) orders, which restricted residents from leaving their homes except for essential activities. The earliest SAH order was implemented in the Bay Area, California on March 16th, 2020. Three days later, the governor of California issued a state-wide SAH order. By March 24th, more than 50% of the U.S. population was under a SAH order (see Figure 3.1). By April 4th, 95% of the U.S. population was under a state or local SAH order, likely substantially reducing the supply of and demand for locally produced goods and services.

†This chapter is a reprint (with permission) of "Unemployment Effects of Stay-at-Home Orders: Evidence from High Frequency Claims Data," coauthored with ChaeWon Baek, Peter McCrory, and Todd Messer. It has been published in the *Review of Economics and Statistics*, which retains first publication credit, © 2020 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

At the same time, there was mounting evidence of substantial disruption to labor markets in the United States. For the week ending March 21st, 2020, the Department of Labor (DOL) reported that more than 3.3 million individuals filed for unemployment benefits.[†] In the subsequent weeks ending March 28th and April 4th, initial claims for unemployment once again hit unprecedented highs of more than 6.9 million claims and 6.7 million claims, respectively. Taken together, total unemployment insurance (UI) claims over this three week period was almost 17 million.

How much of the initially observed increase in UI claims was attributable to the newly implemented SAH orders? This is not a straightforward question to answer since the increase in unemployment claims could plausibly be attributed to a multitude of factors other than SAH orders that occurred at the same time. For example, consumer and business sentiment both declined and economic uncertainty rose as the pandemic worsened. One stark example of this economic uncertainty was the swift drop in the value of the S&P 500 stock market index, which lost roughly 30% of its value between February 20 and March 16, the first day a SAH order was announced in the United States.

In this paper, we disentangle the local effects of SAH orders from the broader economic disruption brought on by the COVID-19 pandemic and other factors affecting all states equally. We do so by providing evidence of a direct causal link between the implementation of SAH orders and the observed increase in UI claims. To the best of our knowledge, this paper is the first systematic study of the causal link between SAH orders and UI claims in the United States. This is our main contribution.

We show that the decentralized implementation of SAH orders across the U.S. induced high-frequency regional variation as to when and to what degree local economies were subject to such orders. We leverage the cross-sectional variation in the length of time that states were exposed to such orders to estimate its effect on UI claims.[†,†]

We find that an additional week of exposure to SAH orders increased UI claims by approximately 1.9% of a state's employment level, relative to unexposed states. The effect is precisely estimated and robust to the inclusion of a battery of controls one might suspect are correlated with both local labor market disruption and SAH implementation, lending it a causal interpretation. The set of controls we consider include the severity of the local exposure to the coronavirus pandemic, state-level political economy factors, and each state's industry composition.

We use our cross-sectional estimate to calculate the implied aggregate effect of SAH

[†]For comparison, in this week one year prior, there were just over 200 thousand initial claims for unemployment insurance. This was also the first time since the DOL began issuing these reports that the flow into unemployment insurance exceeded the number of individuals with continuing claims.

[†]Our variable of interest pertains to the *government* implementation of SAH orders. Our design does not aim to capture the effects of, for example, social distancing behaviors that may have taken place in the absence of a government order.

[†]In this paper, we principally focus on UI claims for three reasons: (1) UI claims are among the highest frequency indicators of real economic activity—especially as it relates to the labor market; (2) These data are consistently reported at a subnational level; (3) The data are publicly and readily available.

orders on the number of new unemployment claims. This exercise yields an estimate of approximately 4 million UI claims attributable to SAH orders through April 4, comprising roughly 24% of total claims over the time period. We refer to this calculation as the relative-implied aggregate estimate of employment losses from SAH orders.

It is well known that cross-sectional research designs, such as the one employed in our paper, hold constant general equilibrium effects as well as other aggregate factors. Simply scaling up our cross-sectional estimate may therefore give a biased impression of the aggregate effect of SAH orders on UI claims in the United States.

To understand the nature of these general equilibrium forces, we present a simplified currency union model to provide conditions under which the relative-implied estimate represents an upper or lower bound on aggregate employment losses. When the SAH shock is viewed primarily as a technology shock—and in the empirically relevant case with sticky prices—our estimate represents an *upper bound* on the aggregate effect. However, when SAH orders are treated as a local demand shock, the interpretation is a bit more subtle and depends upon the persistence of the shock and degree of price flexibility. Across all combinations of price rigidity, persistence and nature of the SAH shock, we find that our back-of-the-envelope estimate, at most, understates aggregate employment losses by a factor of approximately two. With sticky prices and a zero-persistence shock, the relative-implied estimate associated with the SAH-induced local demand shock understates aggregate employment losses by 12%.

Taken together, the model results then imply a (non-binding) *upper bound* on UI claims from SAH orders through April 4, 2020 of approximately 8 million. Thus, relative to the total rise of around 16.5 million, at most around 50% of the total rise in UI claims over this period can be attributed to SAH orders.

Finally, we document the robustness of our empirical results by considering an alternative research design relying upon county-level data. Specifically, we estimate county-level specifications which allow us to control for unobserved state-level factors, such as each state's ability to respond to and process unprecedented numbers of unemployment claims. We find similar results in this case. Appendix C.1 documents the robustness of our headline result to alternative research designs and empirical specifications.

Related Literature

Our paper relates most obviously to the rapidly growing economic literature studying the COVID-19 pandemic, its economic implications, and the policies used to address the simultaneous public health and economic crises. The epidemiology literature has focused on the health effects of NPIs. In a notable study, Hsiang et al. (2020) estimate that, in six major countries, NPI interventions prevented or delayed over 62 million COVID-19 cases.[†] Our focus is, instead, on the macroeconomic effects of the coronavirus pandemic. Broadly speaking, the macroeconomic literature on COVID-19 has split into two distinct yet highly related strands. Here we provide a representative, albeit not exhaustive, review.

[†]The six countries are China, South Korea, Italy, Iran, France, and the United States.

The first strand of research focuses on the relationship between macroeconomic activity, policy, and the unfolding pandemic. Gourinchas (2020) and Atkeson (2020) are early summaries of how the public health crisis and associated policy interventions interact with the economy. Both emphasize the trade-off between flattening the pandemic curve while steepening the recession curve. Similarly, Faria–e–Castro (2020) studies the effect of a pandemic-like event in a quantitative DSGE model in order to assess the economic damage associated with the pandemic along with the fiscal interventions employed in the U.S. to attempt to flatten the recession curve. Eichenbaum et al. (2020) derive an extension of the standard Susceptible-Infected-Recovered (SIR) epidemiological model to incorporate macroeconomic effects, formalizing the relationship between the flattening the pandemic curve and amplifying the recession curve. We view our paper as providing causally identified, empirical support for the claim that flattening the pandemic curve requires steepening the recession curve.

The second strand of research uses high-frequency data to understand the economic fallout wrought by the COVID-19 pandemic. Our paper aligns more closely with this strand of the literature. Baker et al. (2020) show that economic uncertainty measured by stock market volatility, newspaper-based economic uncertainty, and subjective uncertainty in business expectation surveys rose sharply as the pandemic worsened. Lewis et al. (2020) derive a weekly national economic activity index and show that the COVID-19 outbreak had already had a substantial negative effect on the United States economy in the early weeks of the crisis. Hassan et al. (2020) use firm earnings calls to quantify the risks to firms as a result of the COVID-19 crisis. Coibion et al. (2020b) examine how the pandemic affected the labor market in general. Using a repeated large-scale household survey, they show that by April 6th, 2020, 20 millions jobs were lost and the labor market participation rate had fallen sharply.

Our paper also relates to empirical work studying the effect of lockdown policies more specifically. For example, Hartl et al. (2020) study the effect of lockdowns in Germany on the spread of the COVID-19. In contrast to these papers, we use geographic variation to understand the effect of COVID-19 on economic activity. In that respect, our paper can be thought of a high frequency version of Correia et al. (2020), who find that over the long term, NPI policies implemented in response to the 1918 Influenza Pandemic ultimately resulted in faster growth during the recovery following the pandemic.

Other papers employing geographic variation in NPI implementation to understand their contribution to the economic fallout associated with COVID-19 pandemic include the following: Kong and Prinz (2020) use high-frequency Google search data as a proxy for UI claim activity to study the labor market effects of various NPIs; Coibion et al. (2020a) study the effect of lockdowns on employment and macroeconomic expectations; Kahn et al. (2020) document broad declines job market openings in mid-March prior to implementation of SAH orders; Kudlyak and Wolcott (2020) provide evidence that the bulk of UI claims over this period were classified as temporary, suggesting that the long-run costs of lockdowns may be mitigated, so long as worker-firm matches persist until the recovery; and, Sauvagnat et al. (2020) document regional lockdowns depressed the market value of affected firms.

A closely related paper is Friedson et al. (2020), which uses the state-wide SAH order implementation in California along with high frequency data on confirmed COVID-19 cases and deaths to estimate the effect of this policy on flattening the pandemic curve. Unlike our approach, however, the authors in this paper use a synthetic control research design to identify the causal effects on this policy. The authors argue that the SAH order in California reduced the number of cases by 150K over three weeks; the authors perform a back-of-the-envelope calculation to calculate roughly 2-4 jobs lost over a three week period in California per case saved. In contrast to Friedson et al. (2020), we are able to directly estimate the causal effect of SAH orders on UI claims. Taking their benchmark number of cases saved over three weeks, we find that a SAH order implemented over three weeks in California would increase UI claims by 6.4 per case saved.

3.2 Data

State-Level Stay-at-Home Exposure

We construct a county-level dataset of SAH order implementation based on reporting by the *New York Times*. On March 24th, 2020, the *New York Times* began tracking all cities, counties, and states in the United States that had issued SAH orders and the dates that those orders became effective.[†]

We calculate the number of weeks that each county c in the U.S. had been under a SAH order between day $t-k$ and day t (and counting the day that the policy became effective).[†] We denote this variable with $SAH_{c,s,t,t-k}$, where s indicates the state in which the county is located. Except when explicitly stated, we drop the $t-k$ subscript and set k to be large enough so that this variable records the total number of weeks of SAH implementation in county c through time t.

As an example, consider Alameda County, California. Alameda County was among the first counties to be under a SAH order when one was issued on March 16th, 2020. Here, $SAH_{Alameda,CA,Mar.28} = 13/7$, as Alameda County had been under Stay-at-Home policies for thirteen days. Los Angeles County, California, on the other hand, did not issue a SAH order before the State of California did so. We therefore set $SAH_{LosAngeles,CA,Mar.28} = 10/7$ since the state-wide order was issued in California on March 19th, 2020.

[†]The most recent version of this page is available at `https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html`. In a few instances, states implemented the closure of non-essential businesses prior to broader SAH orders that affected businesses and households alike. We show that our results are qualitatively and quantitatively robust to accounting for this occasional discrepancy in timing in Appendix C.1. We choose to rely upon the *New York Times* reporting since it provides sub-state variation. Over time, the *New York Times* stopped separately reporting sub-state orders when a state-wide SAH order was issued. We used the *Internet Archive* to verify the timing and location of SAH orders as reported in the *New York Times*.

[†]When a city implements a SAH order, we assign that date to all counties in which that city is located—unless of course the county had already issued a SAH order.

The previous two examples illustrate how, in some instances, county officials took action before the state in which they were located did. Unfortunately, however, our main outcome of interest, new unemployment claims, is available to us only at the state-level.[†]

To aggregate county-level SAH orders to the state level, we construct a state-level measure of the duration of exposure to SAH orders by taking an employment-weighted average across counties in a given state. Formally, we calculate:

$$SAH_{s,t} \equiv \sum_{c \in s} \frac{Emp_{c,s}}{Emp_s} \times SAH_{c,s,t} \tag{3.1}$$

Employment for each county is the average level of employment in 2018 as reported by the BLS in the Quarterly Census of Employment and Wages (QCEW).[†] One can think of $SAH_{s,t}$ as the average number of weeks a worker in state s was subject to SAH orders by time t.

Figure 3.2 reports $SAH_{s,Apr.4}$ for each state in the U.S. and the District of Columbia. California had the highest exposure to SAH orders at 2.5, indicating that Californian workers were on average subject to SAH orders for two and a half weeks. Conversely, five states (Arkansas, Iowa, Nebraska, Northa Dakota, and South Dakota) had no counties under SAH orders by April 4. The average value across all states of $SAH_{s,Apr.4}$ is 1.2.

Main Outcome Variable: State Initial Claims for Unemployment Insurance

Our main outcome of interest is initial unemployment insurance claims. Initial UI claims is among the highest-frequency real economic activity indicators available. As discussed in the introduction, initial claims for unemployment insurance for the week ending March 21st, 2020 were unprecedented, with more than 3 million workers claiming benefits. By the end of that week, very few states or counties had issued SAH orders. Figure 3.1 shows that by March 21st, only around 20% of the U.S. population was under such directives. This suggests that a substantial portion of the initial economic disruption associated with the COVID-19 crisis may have occurred in the absence of SAH orders.

Let $UI_{s,t}$ indicate new unemployment insurance claims for state s at time t and UI_{s,t_0,t_1} denote cumulative unemployment claims for state s from time t_0 to t_1. In our baseline specification, we consider the effect of SAH orders on cumulative weekly unemployment insurance claims by state from March 14th, 2020 to April 4th, 2020:

$$UI_{s,Mar.21,Apr.4} = UI_{s,Mar.21} + UI_{s,Mar.28} + UI_{s,Apr.4} \tag{3.2}$$

[†]While we lack sufficient data to estimate county-level effects on UI claims, in Section 3.6 we consider county-level regressions in which we estimate the March to April change in log employment and the unemployment rate using data published by the Bureau of Labor Statistics. We find quantitatively similar results even after conditioning on state-level fixed effects. In Appendix C.1 we use this county-level variation to study the impact of SAH orders on retail and workplace mobility, as measured by the Google mobility index.

[†]The annual averages by county in 2019 were, at the time of writing, not yet publicly available.

We then normalize this variable by employment for each state, as reported in the 2018 QCEW, to construct our outcome variable of interest:

$$\frac{UI_{s,Mar.21,Apr.4}}{Emp_s} \tag{3.3}$$

Our choice of April 4th, 2020 as the end date for this regressions is driven by the observation that, by April 4th, 2020, approximately 95% of the U.S. population was under a SAH order. In Section 3.6, we consider 2-week and 4-week horizon specifications and find quantitatively similar results.

3.3 Empirical Specification

We now turn to our research design. Our main design is a state-level, cross-sectional regression:

$$\frac{UI_{s,Mar.21,Apr.4}}{Emp_s} = \alpha + \beta_C \times SAH_{s,Apr.4} + X_s\Gamma + \epsilon_s \tag{3.4}$$

where α is a constant, β_C is the coefficient on state-level exposure to SAH orders, X_s is a vector of controls with associated vector of coefficients Γ, and ϵ_s represents the error term in this equation.

To illustrate the motivation for our empirical design, in Figure 3.3 we compare the evolution of UI claims to state employment of "early adopters," defined as those states being in the top quartile of SAH exposure through April 4, 2020, to that of "late adopters," defined as those states being in the bottom quartile.[†] This figure provides *prima facie* graphical evidence of the main result of our paper: in the first few weeks, early adopters initially had a higher rise in unemployment claims relative to late adopters. By the week ending April 4th, 2020, the relative effect of adopting SAH orders early largely disappears, reflecting the fact that by this point approximately 95% of the U.S. population was under a SAH order, with most having been under the order for the full week ending April 4th.

This figure also suggests that SAH orders alone likely do not account for all of the rise in unemployment claims.[†] In the early weeks, late adopters also experienced historically unprecedented levels of UI claims even though early adopters had higher claims on average. For example, consider the week ending March 28. Here the difference between the median value of the two groups was approximately 1% of state employment; in that week, the median value of initial claims to employment for late adopters was roughly 3%, despite close to zero SAH exposure by this point. By April 4th, this difference almost

[†]The upper and lower edges of the boxes denote the interquartile range of each group, with the horizontal line denoting the median. As is standard, the "whiskers" denote the value representing 1.5 times the interquartile range boundaries.

[†]We thank an anonymous referee for pointing out that this could have the alternative interpretation that local SAH order implementation had substantial negative spillover effects on the rest of the country. See Section 3.5 for a model-driven discussion of such potential spillover effects between states.

completely disappears. Late adopters, who were under SAH orders for a much shorter period of time (or not at all, in some cases), converged to similar levels of unemployment claims relative to employment.

Confounding Factors

In order for our estimate $\hat{\beta}_C$ to have a causal interpretation, it must be the case that the timing of SAH orders implemented at the state and sub-state-level be orthogonal with unobserved factors affecting reported state-level UI claims.[†]

We provide further support for our causal interpretation by testing the magnitude and significance of the estimate $\hat{\beta}_C$ against the inclusion of three sets of important controls. The first set of controls considers the impact that the COVID-19 outbreak itself had on local labor markets. States that chose to implement SAH orders earlier may have done so simply because of the intensity, perceived or otherwise, of the local outbreak. In most macro-SIR models, a larger real outbreak would directly result in a larger drop in consumption due to a higher risk of contracting the virus associated with consumption activity (e.g. Eichenbaum et al. (2020)). To account for this concern, we control for the number of excess deaths, as reported by the Centers for Disease Control and Prevention (CDC), relative to population. We also include the share of the population over 60, as this demographic was more at risk of serious health complications arising from contracting COVID-19.

Additionally, one may be concerned that consumers' perceptions of the outbreak differed from its actual severity. During this time period, the reported number of new confirmed cases was an important statistic reported by the media. This statistic, which suffers from differential testing capability and definitions across states, differs from the measure of excess deaths as it focuses on how local labor markets may have interpreted the severity of the outbreak.[†] We therefore also include the total confirmed cases relative to population.[†] Note that the severity of the outbreak would lead to an upward bias in our estimate $\hat{\beta}_C$ if states were more likely to enact SAH orders when the local outbreak

[†]An additional reason for preferring April 4th is that over longer horizons, there is greater risk of omitted variable bias (i.e. $\mathbb{C}ov[\epsilon_s SAH_{s,Apr.4}] \neq 0$). A salient example is the rollout of the Paycheck Protection Program (PPP) on April 3rd. (The PPP was a central component of the CARES Act, a two trillion fiscal relief package signed into law on March 27, 2020. The PPP authorized $350 billion dollars in potentially forgivable SBA guaranteed loans.) This program provided forgivable loans to small businesses affected by the economic fallout of the pandemic, so long as those loans were used to retain workers. On the margin, PPP incentivizes firms to not lay off their workers, which would tend to lower UI claims for the week after April 4th. Depending upon how this interacts with the differential timing of SAH implementation, the bias could go in either direction.

[†]Evidence from Fetzer et al. (2020) suggests that the arrival of confirmed COVID-19 cases leads to a sharp rise in measures of economic anxiety, which would have an effect on real economic activity through the change in household and firm beliefs about the future state of the economy.

[†]We rely upon confirmed COVID-19 cases as compiled at the county-by-day frequency by USAFacts. USAFacts is a non-profit organization that compiles these data from publicly available sources, typically from daily reports issued by state and local officials. See `https://usafacts.org/visualizations/coronavirus-covid-19-spread-map/` for more details.

was worse or perceived to have been worse, which may itself have led to labor market disruptions.[†]

The second set of controls we consider relates to the political economy of the state government. Some states may have had more generous social safety nets that led workers to separate from firms earlier than in states with less generous policies. Moreover, states with generous policies may also have been more likely to respond earlier to the pandemic, thereby generating bias. To account for this concern, we consider two political economy controls. First, we include the average UI replacement rate in 2019, as reported by the Department of Labor's Employment and Training Administration.[†] Second, we include the Republican vote share in the 2016 presidential election.[†] The first measure is designed to capture the generosity of the social safety net, while the latter is meant to capture political constraints on state and local officials to implement various public health NPIs.

Finally, our last set of controls is intended to address the concern that the timing of SAH implementation may be related to the sectoral composition within each state, and therefore the magnitude of job losses experienced by that state irrespective of SAH orders. To address this concern, we use a measure of predicted state-level UI claims as determined by industry composition within each state and the monthly change in jobs as reported in the national jobs report in March by the BLS. These numbers are based on a survey reference period that concluded on March 14th, 2020—fortuitously for us, two days before any SAH order was announced. Specifically we construct a Bartik-style control:

$$B_s = \sum_i \Delta \ln Emp_{i,March} \times \omega_{i,s} \tag{3.5}$$

where $\Delta \ln Emp_{i,March}$ is the monthly percentage change in employment in industry i (3-digit NAICs) for the month of March. $\omega_{i,s}$ is the share of employment in industry i in the state, as reported in the QCEW for 2018.

We also control for the extent of work-at-home capacity at the state-level. Dingel and Neiman (2020) construct an index denoting the share of jobs that can be done at home by cities, industries, and countries. We construct a state-level index by taking an state employment-weighted average of the Dingel and Neiman (2020) industry-level (2-digit NAICS) work-at-home index. It may be the case that states with a higher capacity to work from home may have been willing to implement SAH orders earlier if the labor market disruption of such policies was perceived to be lower when more workers are able to work from home. If this index is correlated with the number of initial UI claims received by

[†]Our controls for excess deaths and confirmed cases are taken as cumulative sums as of the end of the sample period, which is April 4th in the benchmark analysis. We experimented with using lagged values of these measures as pre-period controls, and they had no effect on the magnitude or significance of our coefficient of interest. These results are available upon request.

[†]See `https://oui.doleta.gov/unemploy/ui_replacement_rates.asp` for more details.

[†]As reported by the *New York Times* at `https://www.nytimes.com/elections/2016/results/president`.

the state in the absence of implementing SAH orders, then failing to include this control would introduce bias.[†]

Causal interpretations aside, the cross-sectional framework is nevertheless constrained in only answering the following question: By how much did UI claims increase in a state that implemented SAH orders *relative* to a state that did not? The constant term absorbs, for example, the general equilibrium effects of stay-at-home orders which would affect all states within the U.S.—not just those implementing SAH orders. To the extent that other states' labor markets were affected in any way by the local imposition of SAH orders, then $\hat{\beta}_C$ will fail to capture the *entire* effect of such policies. We postpone discussion of the mapping between the relative effect of SAH orders and their aggregate effect until after presenting our cross-sectional results.

3.4 Results

Effects of SAH Orders on State-Level UI Claims

In Table 3.1, we present results from estimating Equation (3.4). Column (1) shows the univariate specification, with no controls. The point estimate of approximately 1.9% (SE: 0.67%) implies that a one-week increase in exposure to SAH orders raises the number of claims as a share of state employment by 1.9% relative to states that did not implement SAH orders. Figure 3.4 displays this result graphically. The bubbles are shaded according to the intensity of the confirmed COVID-19 cases per thousand people and the size of the bubbles are proportional to state population.

In Column (2), we control for the number of confirmed COVID-19 cases per one thousand people, excess deaths by state, and the share of state population over the age of 60. As discussed, these are intended to control for factors related to the pandemic that might simultaneously affect both the timing of SAH implementation and the severity of state labor market disruptions. The change in the coefficient is immaterial—economically and statistically. In Column (3) we control for political economy factors: the state's UI replacement rate in 2019 and the 2016 Trump vote share. Our estimate $\hat{\beta}_C$ falls only slightly to 1.8%. In Column (4) we include controls for each state's sectoral composition (and in turn its sensitivity to both the pandemic-induced crisis and timing of SAH implementation). Our point estimate is again largely unchanged.

Finally, in column (5), we select a parsimonious specification that captures dimensions of each set of controls. We control for confirmed cases, excess deaths, the UI replacement rate, and the WAH index (the only significant variable). In this specification, which is our

[†]In unreported regressions, we study whether the effect of SAH orders differentially depends upon the value of the work-at-home index; we find no evidence that this is the case.

preferred specification, the estimate of β_C is still 1.9%.[†,†]

Our results support the idea that policies that work to flatten the pandemic curve also imply a steepening of the recession curve (Gourinchas, 2020). To quantify this steepening of the recession curve, we use our point estimate of the relative effect on state-level UI claims of SAH orders to calculate a back-of-the-envelope estimate of the total implied number of UI claims between March 14 and April 4 attributable to SAH orders. We calculate the relative-implied estimate as follows:[†]

$$\text{Relative-Implied-Aggregate-Claims} = \sum_s \hat{\beta}_C \times SAH_{s,Apr.4} \times Emp_s \tag{3.6}$$

where s indexes a particular state. This is a back-of-the-envelope calculation as it simply scales up the cross-sectional coefficient $\hat{\beta}_C$ according to each state's SAH exposure through April 4, 2020 and each state's level of employment.

This back-of-the-envelope calculation yields an estimate of 4 million UI claims attributable to SAH orders through April 4. Ignoring cross-regional spillovers, this relative-implied estimate suggests that approximately 24% of total claims through April 4, 2020 were attributable to such orders.

This calculation does not incorporate general equilibrium effects or spillovers that may have arisen as a result of local SAH implementation. As we discuss in Section 3.5, when the SAH order is interpreted as a local productivity shock, this represents an upper bound on aggregate employment losses; when, however, the SAH implementation is treated as a local demand shock, the analysis is a bit subtler. Yet, even in this case, we find that at most the relative-implied aggregate multiplier understates true employment aggregate employment losses by a factor of 2. Through the lens of the model, this provides an upper bound on total employment losses attributable SAH orders: 8 million UI claims through April 4, or approximately half of the overall spike in claims during the initial weeks of the economic crisis induced by the COVID-19 pandemic.

An alternative back-of-the-envelope calculation to assess the magnitude of our estimate is to instead focus the relative contribution of SAH orders in terms of typical cross-sectional variation in UI claims in our sample. Our estimates imply that a state which implemented SAH orders one week earlier saw an increase in UI claims by 1.9% of its 2018 employment

[†]In the appendix, we consider three additional robustness exercises at the state-level. We alternate the horizon over which the model is estimated (2 and 4 weeks), estimate the model by weighted least squares, and re-estimate the model dropping one state at a time. The results are quantitatively and qualitatively similar.

[†]In unreported regressions, we find that, when including all regressors, $\hat{\beta}_C$ is somewhat attenuated—albeit statistically indistinguishable from our baseline estimate; however, this attenuation is largely driven by the parametric assumption of linearity on the share of votes for Trump in 2016, which places substantial leverage on Wyoming and West Virginia. Dropping these states from the full specification with all control variables yields a point estimate of 1.8% (SE: 0.75%). These regressions are available upon request.

[†]We use the terminology "relative-implied" because in the cross-section we are only able to identify effects of SAH orders relative to states not implementing SAH orders. We discuss this issue at greater length in Section 3.5.

level relative to a state one week later, which is slightly less than 50% of the cross-sectional standard deviation of employment-normalized claims between weeks ending March 21 and April 4.[†]

3.5 Aggregate Versus Relative Effects

Our empirical strategy relies on cross-sectional variation in the timing and location of SAH orders to identify the relative effect such policies had on labor markets during the initial weeks of the COVID-19 outbreak in the United States. In this section, we discuss in greater detail the sorts of spillovers that are likely to be relevant and the conditions under which the relative-implied aggregate estimate (see equation (3.6)) represents a lower or upper bound on the aggregate effects of SAH orders on UI claims. This is important for how one should interpret our back-of-the-envelope calculation that in the early period of the crisis, approximately only 24% of UI claims through April 4, 2020 were related to SAH orders.

To the extent that there are cross-regional (either positive or negative) spillovers of SAH orders, our estimate will not capture the *aggregate* effect of SAH orders. This limitation is related to the stable unit value (SUTVA) assumption in the causal inference literature, which requires that potential outcomes be independent of the treatment status of other observational units. Because of considerable trade between U.S. states, SUTVA is likely to be violated in our setting.[†]

To guide our discussion, we use a benchmark currency-union model to study the effects of SAH orders on the local economy, the rest of the currency union, and the entire economy as a whole. We present results for an economy characterized either by sticky prices or flexible prices, with SAH orders modeled as either a pure local demand shock or a pure local productivity/supply shock; the evidence from Appendix C.1 suggests that both channels were operative.[†] We then briefly summarize other important cross-regional spillovers not well-captured by the currency model we study. The most salient of these spillovers relate to the *informational* effect of early SAH implementation in some parts of the country.

Currency Union Model: Supply and Demand Shock Implications of SAH Orders

In this section, we consider the implications of local demand or supply shocks in a benchmark currency union model under either sticky or flexible prices. The model we consider is a simpler version of the baseline, separable utility, complete markets

[†]We thank an anonymous referee for this particular recommendation.

[†]SUTVA violations are likely to be more salient in the cross-section when the model is estimated over longer horizons. This is, in part, why we choose as our baseline the 3-week horizon specification.

[†]Additionally, as is discussed in Brinca et al. (2020), it is appropriate to view the COVID-19 pandemic (and associated policy responses) as some combination of demand and supply shocks. We consider pure demand and supply shocks to illustrate the economic implications of each in isolation.

model presented in Nakamura and Steinsson (2014), modified to incorporate productivity shocks and discount rate shocks (to model negative local supply and demand shocks, respectively).[†] We follow Nakamura and Steinsson (2014) in calibrating the model to the U.S. setting. The full model specification is relegated to the Appendix; here we present only those aspects of the model modified to study the effects of SAH orders.

Modeling SAH Orders

Our first model experiment is to treat the implementation of SAH orders as a pure local demand shock. To incorporate this into the model, we introduce a consumption preference shock, δ_t. This preference shock causes home region households to prefer, all else equal, delaying consumption into the future. This may be a reasonable way to model the SAH shock for a variety of reasons. First, to the extent that the drop in retail mobility, as shown in Appendix C.1, represents a decline in goods consumption, households may simply be delaying such purchases until temporarily closed stores reopen. Second, the inability to purchase locally furnished goods and services may lead households to temporarily save more than they might otherwise choose to do, which would be observationally equivalent to a discount rate shock only to consumption.

Households in the home region maximize the present discounted value of expected utility over current and future consumption C_t and labor supply N_t.

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\delta_t \frac{(C_t)^{1-\sigma}}{1-\sigma} - \chi \frac{(N_t)^{1+\psi}}{1+\psi} \right],$$

where β is the rate of time discounting, σ is the inverse intertemporal elasticity of substitution, ψ is the inverse Frisch elasticity of labor supply, and χ is the weight on labor supply. The discount rate shock process follows

$$\log \delta_t = \rho^{\delta} \log \delta_{t-1} + \epsilon_t^{\delta}. \tag{3.7}$$

We close the household side of the model by assuming preferences for varieties are constant elasticity of substitution (CES), which gives rise to the standard CES demand curve via cost minimization.

Alternatively, the SAH orders may be modeled as a local productivity shock. Even if demand for locally produced goods is unchanged, firms may be constrained in supplying the goods and services demanded by local households or by the rest of the currency union. We model this interpretation as a region-level productivity shock for intermediate-goods-producing firms. A firm i in the home region faces the following production function

$$y_{h,t}(i) = A_t N_{h,t}(i)^{\alpha},$$

[†]Implications from a model with different preference structures (e.g. Greenwood et al. (1988) preference) and with incomplete market are qualitatively the same. Unlike the original focus of Nakamura and Steinsson (2014), the model we consider does not incorporate government spending shocks, as that is not our focus in this paper.

where $y_{h,t}(i)$ is the output of a firm i, $N_{h,t}(i)$ is the amount of labor input hired by the firm, and A_t is region-wide technology in the home region. α is the returns to scale parameter on labor. The aggregate supply shock A_t evolves according to the following process:

$$\log A_t = \rho^A \log A_{t-1} + \epsilon_t^A. \tag{3.8}$$

Firms maximize profits subject to demand by households. Nominal rigidities are specified à la Calvo (1983) with associated price-reset parameter θ.

Finally, we close the model by assuming bond markets are complete, labor markets are perfectly competitive, and, when prices are sticky, the monetary authority follows a union-wide Taylor rule. A full derivation is available in the Appendix.

Model Results: Modeling SAH Order Shocks under Flexible and Sticky Prices

We model the implementation of SAH orders as a one-time negative shock with either $\epsilon_t^\delta = -1$ (for local demand shocks) or $\epsilon_t^A = -1$ (for local supply shocks). We choose zero decay parameters on the shock series to illustrate the dynamics of the model in settings in which the shock induced by the SAH order is temporary. Specifically, we set $\rho^A = \rho^\delta = 0$. For the purposes of mapping the relative-implied employment losses to aggregate employment losses, this is without loss for the results for the technology shock but not without loss with respect to the demand shock with sticky prices. Below, we discuss what happens when the demand shock exhibits some persistence.

We calibrate the remaining parameter values according to Nakamura and Steinsson (2014) (see their Section III.D.). When working with the sticky price model, we set the Calvo parameter $\theta = 0.75$. In the flexible price model, we set $\theta = 0$.

We consider each of the two types of shocks in isolation under either sticky prices or fully flexible prices. In each of the four scenarios, we calculate the on-impact responses of home region employment, foreign region employment, and aggregate employment to the local shock. Because the model is calibrated to a quarterly frequency and because our empirical design estimates the relative effect over a short horizon (3-weeks), the relevant horizon for mapping the model to the cross-section is the *on-impact* relative effect between employment in the shocked home region and the non-shocked foreign region.

The results from these exercises are reported in Figure 3.5 and Table 3.2. Figure 3.5 shows the *on-impact* responses of employment in a home region (blue circles) and a foreign region (red crosses), and aggregate employment (black squares) under the four different scenarios. Table 3.2 then compares the relative-implied aggregate employment calculated from the differences between the responses of home and foreign employment and the responses of aggregate employment under different scenarios.[†]

In the model, only three of the four stylized scenarios we consider produce relative effects of SAH orders that are consistent with the positive coefficient we estimate in the

[†]Formally, the relative-implied estimate in the model is calculated as $n(\ell_t - \ell_t^*)$, where ℓ_t and ℓ_t^* represent log deviations from steady state of home and foreign region per-capita employment respectively. n is the size of the home-region. This is exactly the model-analog of the relative-implied estimate reported in equation (3.6).

Table 3.1: Effect of Stay-at-Home Orders on Cumulative Initial Weekly Claims Relative to State Employment for Weeks Ending March 21 thru April 4, 2020

	(1) Bivariate	(2) Covid	(3) Pol. Econ.	(4) Sectoral	(5) All
SAH Exposure thru Apr. 4	0.0194^{***}	0.0192^{**}	0.0178^{**}	0.0209^{***}	0.0187^{**}
	(0.00664)	(0.00742)	(0.00818)	(0.00637)	(0.00714)
COVID-19 Cases per 1K		-0.00213			0.00194
		(0.00621)			(0.00676)
Excess Deaths per 1K		0.0446			0.0480
		(0.109)			(0.113)
Share Age 60+		0.237			
		(0.281)			
Avg. UI Replacement Rate			0.0719		0.0726
			(0.0794)		(0.0787)
2016 Trump Vote Share			-0.0225		
			(0.0508)		
Work at Home Index				-0.331^{+}	-0.388^{+}
				(0.192)	(0.229)
Bartik-Predicted Job Loss				-2.401	
				(7.528)	
Constant	0.0815^{***}	0.0357	0.0621	0.181^{**}	0.182^{**}
	(0.00848)	(0.0543)	(0.0481)	(0.0742)	(0.0821)
Adj. R-Square	0.0829	0.0434	0.0618	0.0966	0.0763
No. Obs.	51	51	51	51	51

Table 3.2: On-Impact Response of Union-Wide Employment and Relative-Implied Aggregate Employment to a Local SAH-induced: (i) Preference Shock with Flexible Prices, (ii) Preference Shock with Sticky Prices, (iii) Technology Shock with Flexible Prices, and (iv) Technology Shock with Sticky Prices

	Flexible				Sticky		
	Total	Implied	Factor		Total	Implied	Factor
Preference Shock	-0.047	-0.021	2.21	$\rho^{\delta} = 0.9$	-0.032	-0.075	0.43
				$\rho^{\delta} = 0.0$	-0.093	-0.083	1.12
Technology Shock	0.003	-0.021	-0.16		0.1642	0.1398	1.18

data. When the SAH orders are modeled as local productivity shocks, only the flexible price equilibrium produces an immediate, relative decline in employment in the home region subject to the shock. When the SAH orders are instead modeled as local demand shocks, both the sticky price and flexible price economies produce a steeper decline in the shocked home region's employment relative to the rest of the economy, as suggested by the cross-sectional evidence presented above.

When SAH orders are modeled as negative productivity shocks with fully flexible prices, the immediate, relative effect of SAH orders is an *upper bound* on the aggregate employment effect over the same horizon. This is because the decline in local employment arising from the SAH order is offset by an increase in employment in the rest of the economy. The mechanism is that in the flexible price case, the negative productivity shock in the home region translates into an improvement in the foreign region's terms of trade. This, in turn, increases labor demand in the foreign region, which increases employment in the foreign region.

In contrast, when prices are fully flexible in response to an SAH-induced home-region demand shock, the relative-implied estimate represents a *lower* bound on aggregate employment losses. This is because employment in both the home and foreign regions fall in response to the shock. With prices being fully flexible, the negative preference shock in the home region leads to a decline in prices for home goods relative to foreign goods, making foreign consumption more expensive. This, in turn, decreases demand for foreign goods, resulting in a decline in foreign employment, which is necessary for market clearing. When prices are fully flexible and the effect of SAH orders is a pure local demand shock, aggregating the relative employment losses understates the aggregate employment losses by a factor of about two (see Table 3.2, Row 1, Column 3).

The case with sticky prices and SAH orders modeled as a pure local demand shock lies in between the previous two scenarios. When the local demand shock is sufficiently persistent, the immediate, relative effect of SAH orders could potentially *overstate* the aggregate employment effect. This is because employment in the foreign region increases on impact. Meanwhile, when the demand shock has essentially no persistence, so that it only affects demand in the home region for a single quarter, employment in the foreign region also falls on impact, implying that the (aggregated) relative employment effect again understates aggregate employment losses, in the quarter of the shock (See Figure 3.5). Regardless, the degree to which this on-impact effect understates aggregate employment losses is bounded above by the response under flexible prices to a local demand shock.

The evidence presented in Appendix C.1 suggests that SAH orders represented a shock to both the supply of and demand for locally produced goods. This on its own implies that the flexible price, preference shock scenario provides a non-binding upper bound on aggregate employment losses. Specifically, in this scenario the relative-implied aggregate estimate would understate employment losses by roughly a factor of two. The distance from this upper bound increases, moreover, with price rigidity and the persistence of the SAH shock. In the baseline calibration, when prices are sticky and the demand shock has no persistence, the relative-implied job losses understates aggregate employment losses by 12%.

Other Cross-Regional Spillovers

The benchmark currency-union model presented in the previous section illustrates how locally implemented SAH orders would affect the local economy, other regions in the currency union and the entire economy as a whole. The spillover forces in the model work through the trade in goods between regions and associated price and expenditure switching effects. However, there may be other important cross-regional spillovers that are not well-captured by the model, but may nevertheless be important for interpreting our empirical results in light of the aggregate effects of SAH orders.

An important example is an *informational effect* of early SAH implementation in some parts of the economy. For example, the early imposition of SAH orders in some regions may signal to the rest of the country that a SAH order is likely to be imposed some time in the near future. This informational channel can be incorporated into the model by assuming that the foreign region learns, on-impact, that a SAH order will be imposed in the foreign region in the subsequent period. We experimented with this specific informational channel of local SAH order implementation and found that the upper and lower bounds provided in the previous subsection continued to hold.[†]

A more subtle informational effect of SAH implementation relates to the credible signal it sends about the severity of the COVID-19 pandemic and the potential economic disruptions it is likely to induce, even in the absence of any additional SAH orders. In this interpretation, the SAH orders have spillover effects on the rest of the economy through the changes they induce to beliefs held by households and firms about the future path of the economy. As opposed to other signals conveyed by public officials about the severity of the pandemic, SAH implementation is a credible signal because it imposes non-trivial costs on the economy. This could, in turn, lead to a reduction in demand as a result of increased economic anxiety and fear of exposure to the COVID-19.

If this second informational effect of local SAH implementation ultimately led to job losses throughout the rest of the country, then our relative-implied estimate would understate the aggregate job losses attributable to SAH orders. Neither the model nor the empirical design takes this particular spillover mechanism into account. We view understanding the role of SAH orders as credibly communicating the severity of the pandemic as an important and interesting avenue for future research.[†]

Another important example is spillovers through firm networks—internal and external.[†] For example, complex supply chains may cause economic activity to decline in parts of the country where SAH orders are not yet enacted if the sourcing of intermediate inputs is affected. Alternatively, national chains may close establishments located in regions

[†]These results are available upon request.

[†]Coibion et al. (2020a) provide evidence that local SAH orders led households in the affected regions to hold more pessimistic views of the future path of the economy. This is a separate, though related, channel than the *aggregate* change in beliefs that may have occurred following the early imposition of SAH orders.

[†]We thank an anonymous referee for pointing this out.

without SAH orders due to losses in other major markets with SAH orders. Arguably, these sorts of spillovers would lead our relative-implied estimate of job losses to understate true aggregate employment losses. However, we believe these channels are minor, as the adjustments would need to occur over a very short period time. The horizon of our empirical specifications is three weeks, during which time existing inventories were likely to be sufficient for production.[†]

3.6 Alternative Specification: County-Level Employment and Unemployment Effects

A major concern with the estimates of Equation (3.4) is that states may have experienced substantial difficulty in scaling up their systems to process the historically unprecedented numbers of unemployment claims. For example, it is well known that some states' unemployment insurance systems rely on archaic computer programming languages.[†] Thus, it is reasonable to be worried that states with more cumbersome systems may systematically report lower UI claims numbers relative to those states with more efficient systems.

A priori, the induced omitted variable bias could go in either direction. On the one hand, states with stronger UI systems may have also been more inclined to respond aggressively to the COVID-19 pandemic with SAH orders, generating an upward bias in our estimates. On the other hand, the severity of labor market disruptions from the COVID-19 pandemic may have both made it more difficult for states to process new claims *and* made them more likely to impose SAH orders earlier—thus, generating a downward bias. While we have already controlled for measures of COVID-19 in our estimates of Equation (3.4), in this subsection we present an alternative design at the county-level using employment and unemployment as outcomes, albeit at a lower frequency. Using total employment, rather than unemployment insurance claims, allows us to sidestep the issue of whether states could meet demand for UI claims. This design also allows for the inclusion of state fixed effects to identify the relative effect of SAH orders using within-state variation in the timing of SAH implementation.

We analyze the effects of SAH orders at the county-level relying upon local area unemployment and employment statistics constructed by the Bureau of Labor Statistics (BLS). The downside is that this data is constructed at the monthly frequency, rather than the weekly frequency in our main specification.[†] The BLS primarily relies upon the Current

[†]It is a well known observation that inventories generally adjust more slowly to changes in sales, consistent with the claim that this particular source of bias is most relevant at lower frequencies and longer horizons. (See Ramey and West, 1999; Bils and Kahn, 2000).

[†]See, for example, "'COBOL Cowboys' Aim To Rescue Sluggish State Unemployment Systems" by NPR (`https://www.npr.org/2020/04/22/841682627/cobol-cowboys-aim-to-rescue-sluggish-state-unemployment-systems`).

[†]In Appendix C.1 we estimate event study specifications using high frequency employment statistics at the county-level for a subset of counties in the U.S. for which these data exist. We find no evidence of differential changes in county-level employment prior to SAH implementation while at the same time finding that SAH orders lowered employment on average by 1.9% after one week.

Population Survey (CPS) as the primary input into constructing estimates of county-level employment and unemployment.† Fortunately, the survey reference periods for the CPS aligns quite nicely with measuring household employment and unemployment just prior to the broad implementation of SAH orders and one month hence. The reference week for the CPS for March 2020 was March 8th through March 14th and the reference week for April was April 12th through April 18th.

We estimate analogs of our state-level regression at the county-level, using as our outcome variable either the log change in employment or the change in the unemployment rate between March 2020 and April 2020. County-level treatment is the weekly SAH exposure through April 15, 2020. Formally, we estimate the following regression by ordinary least squares:

$$\Delta y_{c,s,April} = \alpha_s + \beta^y_{C,county} \times SAH_{c,s,Apr.15} + X_{c,s}\Gamma + \epsilon_{c,s} \tag{3.9}$$

where $y_{c,s,April}$ indicates the monthly change between March and April in either log employment or the unemployment rate. α_s are state-level fixed effects which control for all state-level policies implemented between mid-March and mid-April that may have been systematically related to observed UI claims during that period. We also report results when constraining $\alpha_s = \alpha$ to provide a natural benchmark against our state-level regression. We also control for the number of confirmed COVID-19 cases per thousand people and the WAH index, which are our only controls available at the county-level.†

Because the first outcome variable we consider at the county-level is the log change in county employment, we expect that the estimated relative effect of SAH orders on local employment, $\hat{\beta}^{emp}_{C,county}$, will be comparable to our estimate of the same parameter at the state-level.† If the timing of the decentralized implementation of SAH orders was orthogonal to state-level economic conditions and if there were negligible spillovers from treated counties to untreated counties within the same state, then we would expect to see a relatively stable coefficient regardless of whether we include state fixed effects, α_s, or not.

Table 3.3 provides the results for the effects of SAH orders on employment. The first column shows the results restricting $\alpha_s = \alpha$ (e.g., no state fixed effects). The point estimate suggests that the relative effect of SAH exposure on employment at the county-level is to reduce employment by of -1.8% (SE: .57%). That we use a different outcome variable and different level of disaggregation yet obtain a coefficient of similar magnitude is encouraging.

†For additional details on the methodology employed by the Bureau of Labor Statistics, see `https://www.bls.gov/lau/laumthd.htm`.

†We control for the number of confirmed COVID-19 cases through April 15th to align with the timing of the surveys used by the BLS to construct county-level employment and unemployment statistics.

†Note that because we use the 2018 QCEW to normalize UI claims at the state-level, we should expect the county-level estimates to be slightly lower in magnitude since the state-level regressions calculates the percent change off of a smaller base value.

Figure 3.1: Cumulative Share of Population under Stay-at-Home Orders in the U.S.

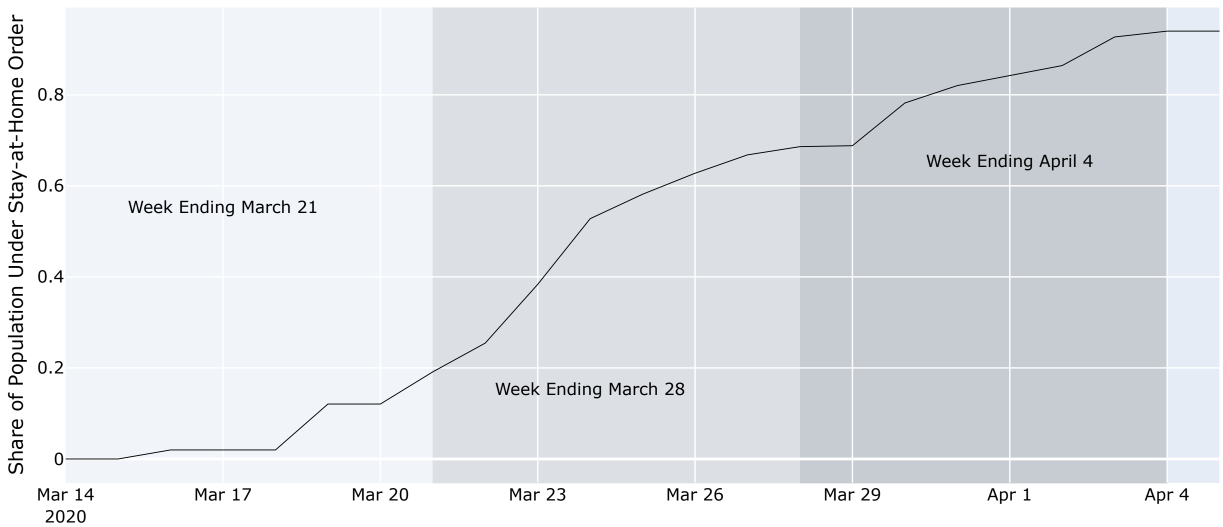

	(1) $\Delta \ln Emp$	(2) $\Delta \ln Emp$	(3) $\Delta \ln Emp$	(4) $\Delta \ln Emp$
SAH Exposure thru Apr. 15	-0.0176***	-0.0124**	-0.0129**	-0.00905**
	(0.00568)	(0.00464)	(0.00453)	(0.00397)
Covid-19 Cases per 1K Emp			-0.0000280	-0.000116
			(0.0000348)	(0.000121)
Work at Home Index			0.0549	0.0547
			(0.0457)	(0.0537)
Constant	-0.0824***	-0.113***	-0.129***	-0.135***
	(0.0147)	(0.00900)	(0.0157)	(0.0139)
Dep Mean	-0.12	-0.14	-0.14	-0.14
States	51.00	12.00	12.00	12.00
State FE	No	Yes	Yes	Yes
CZ FE	No	No	No	Yes
Adj. R-Square	0.10	0.62	0.63	0.74
No. Obs.	3141.00	1116.00	1116.00	453.00

Table 3.3: County-Level Specification: Effect of Stay-at-Home Orders on Local Employment Growth

Columns (2) and (3) focus on the 12 states for which there is variation across counties in the timing of SAH orders. The magnitude of the estimate falls by about one third, regardless of whether we include controls—although this difference is not statistically significant. If, as we argue above, the timing of SAH implementation was orthogonal to policies and economic conditions at the state-level[†], then the decline in the point estimate is suggestive evidence of negative spillovers between treated and untreated counties. While this may be the appropriate interpretation, it appears that the bulk of employment losses were nevertheless concentrated within the labor markets in which SAH orders were implemented.

Finally, in the last column, we include commuting zone fixed effects and find that the coefficient is roughly a third of the effect estimated in column (3).[†] Following a similar logic as in the previous paragraph, this would suggest that not only were the bulk of employment losses concentrated within the labor market, they were moreover concentrated within the specific counties in which the SAH orders were implemented.

Table 3.4 provides the results for the effects of SAH orders on the change in the county-level unemployment rate. As with the employment specification, the first column does not include state fixed effects. In columns (2) and (3) we include state fixed effects; in the final column, we condition further on commuting zone fixed effects. Consider the result reported in column (3), the state fixed effects specification with controls for local COVID-19 pandemic and capacity for the local labor force to work from home: the point estimate is 1.57 (SE: 0.331), implying that each week of SAH exposure at the county-level increased the local unemployment rate by 1.57.

In sum, we view the the county-level results as corroborating evidence of the main result in this paper: that the cross-sectional effect of SAH orders had real costs to the labor markets in the early weeks of the crisis, but that such costs were likely dwarfed by other factors in the early weeks of the crisis. While not inconsistent with our state-level analysis, broadly the county-level design yields somewhat lower point estimates than in our benchmark specification. In this respect, relative to a null that all observed UI claims were attributable to SAH orders, the state-level specification yields the most conservative estimate of the relative effect of such orders on local labor markets. Through the lens of our theoretical model, these cross-sectional estimates imply, at most, a non-binding upper bound of half of total UI claims through April 4, 2020 being attributable to SAH orders.

3.7 Conclusion

While non-pharmaceutical interventions (NPIs) are necessary to slow the spread of viruses such as COVID-19, they likely steepen the recession curve. But to what extent? We provide

[†]And the average treatment effect among counties in the twelve states appearing in columns (2)-(4) is the same as for counties.

[†]We use the United States Department of Agriculture (USDA) 2000 county to commuting zone crosswalk. This is available at `https://www.ers.usda.gov/data-products/commuting-zones-and-labor-market-areas/`.

estimates of how much one prominent NPI disrupted local labor markets in the short run in the U.S. in the early weeks of the coronavirus pandemic.

In particular, we investigate the effect of Stay-at-Home (SAH) orders on new unemployment claims in order to quantify the causal effect of this severe NPI (i.e., flattening the pandemic curve) on economic activity (i.e., steepening the recession curve). The decentralized implementation of SAH orders in the U.S. induced both geographic and temporal variation in when regions were subject to restrictions on economic and social mobility. Between March 14th and April 4th, the share of workers under such orders rose from 0% to almost 95%. This rise was gradual but steady, with new areas implementing SAH orders on a daily basis. We couple this variation in SAH implementation with high-frequency unemployment claims data to quantify the resulting economic disruption.

We find that a one-week increase in stay-at-home orders raised unemployment claims by 1.9% of state-level employment. This estimate is robust to a battery of controls, including the severity of the local COVID-19 pandemic, the local political economy response, and the industry mix of the local economy. A back-of-the-envelope calculation using our estimate implies that SAH orders resulted in a rise of 4 million unemployment insurance claims, about a quarter of the total unemployment insurance claims during this period. A stylized currency union model suggests that in some empirically relevant cases, this estimate can be seen as an upper bound. When it instead represents a lower bound, it at most understates job losses by a factor of two.

While it is beyond the scope of this paper to uncover all determinants of the unprecedented initial rise in unemployment during the COVID-19 pandemic, there is evidence that the economic downturn was already under way by the time that SAH orders were implemented. Even before the national emergency was announced by President Trump on March 13, 2020, households were reallocating their spending away from in-person goods and services.[†] Consistent with this evidence, our estimates imply that a sizeable share of the increase in unemployment in the early weeks of the COVID-19 crisis was due to other channels, such as decreased consumer sentiment, stock market disruptions, and social distancing that would have occurred in the absence of government orders.

Nevertheless, despite representing a minority share of the overall increase in unemployment in the initial three weeks of the crisis, our estimates suggest that over longer horizons SAH orders played a much larger role. Performing an out-of-sample forecast through April 25 of the relative-implied aggregate effect of SAH orders is illustrative: An additional 7.5 million UI claims between April 4 and April 25 are due to SAH orders, little more than half of the additional overall increase in UI claims nationally during that time.[†]

[†]By March 13, grocery spending was up 44%, restaurant spending was down 10%, and entertainment and recreation spending was down 23%, all relative to their respective levels in January 2020. At about the same time—and preceding any reported SAH orders—both national consumer spending and small business revenue began their precipitous declines. Statistics calculated from data available at `https://tracktherecovery.org/`.

[†]This helps to reconcile our estimates with Coibion et al. (2020a) who find a larger contribution of SAH orders to job losses throughout April than we do. In this exercise, we adjust for whether a state

In sum, we see our paper as providing evidence that undoing SAH orders may relieve only a fraction of the economic disruption arising from the COVID-19 pandemic while at the same time exacerbating the public health crisis. This implies that the economic downturn may persist at least until the pandemic itself is resolved. At the same time, we document a large elasticity of unemployment with respect to such lockdown measures, suggesting that the costs of SAH orders are non-trivial in the long-run.

reopened before April 25; not adjusting increases the out-of-sample forecast to 7.6 million claims. See https://www.nytimes.com/interactive/2020/us/states-reopen-map-coronavirus.html for state reopening dates.

Figure 3.2: Employment-Weighted State Exposure to Stay-at-Home Policies Through Week Ending April 4

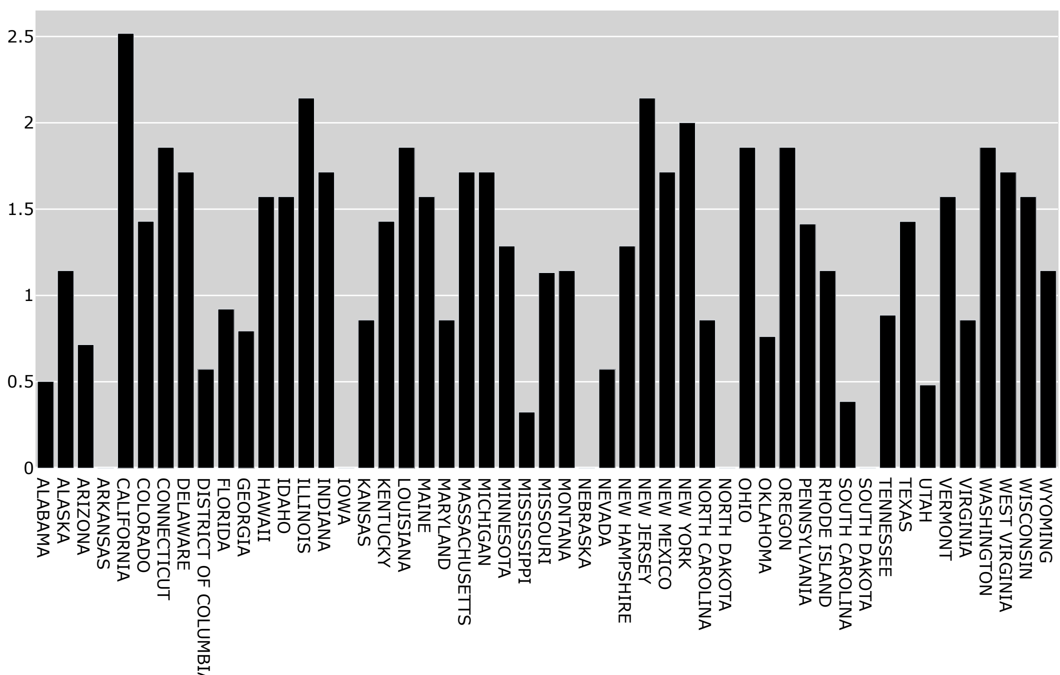

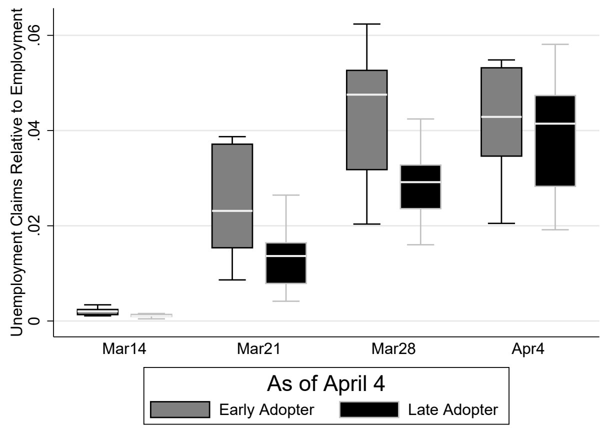

Figure 3.3: Box Plots by Week of Initial UI Claims Relative to Employment for Early and Late Adopters of SAH orders

Figure 3.4: Scatterplot of SAH Exposure to Cumulative Initial Weekly Claims for Weeks Ending March 21 thru April 4

Figure 3.5: On-Impact Response of Home Employment, Foreign Employment, and Union-Wide Employment to a Local SAH-induced: (i) Technology Shock with Flexible Prices, (ii) Technology Shock with Sticky Prices, (iii) Preference Shock with Flexible Prices, and (iv) Preference Shock with Sticky Prices

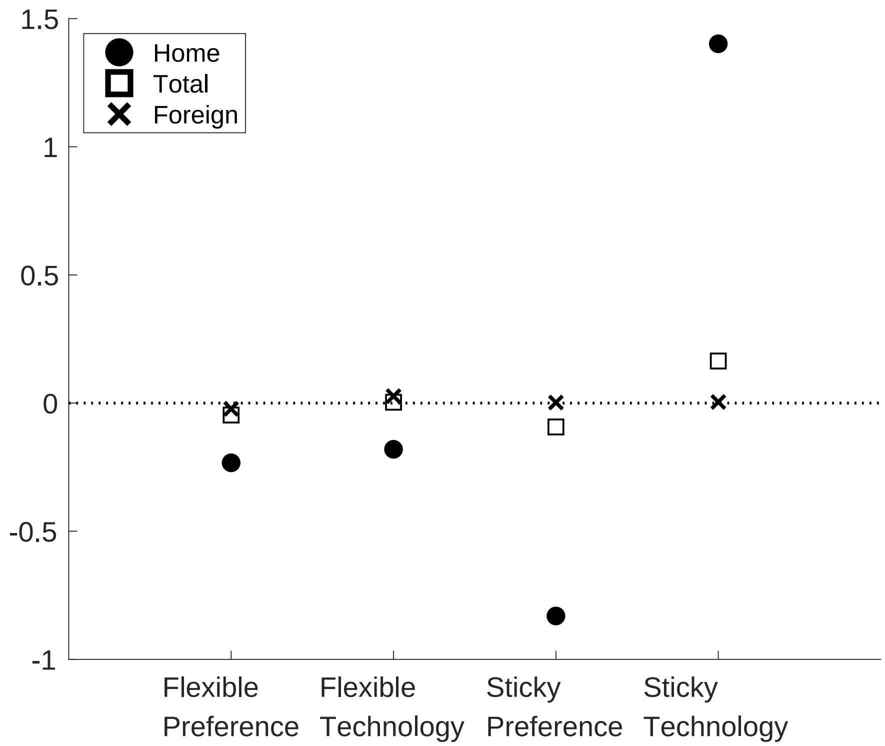

	(1)	(2)	(3)	(4)
	ΔUR	ΔUR	ΔUR	ΔUR
SAH Exposure thru Apr. 15	1.574***	1.382***	1.570***	0.944***
	(0.400)	(0.331)	(0.331)	(0.216)
Covid-19 Cases per 1K Emp			-0.000239	0.0110
			(0.00468)	(0.00806)
Work at Home Index			-12.29**	-5.437
			(5.336)	(5.089)
Constant	4.114***	4.425***	7.922***	6.689***
	(0.888)	(0.642)	(2.005)	(1.863)
Dep Mean	7.69	7.11	7.11	7.32
States	51.00	12.00	12.00	12.00
State FE	No	Yes	Yes	Yes
CZ FE	No	No	No	Yes
Adj. R-Square	0.13	0.39	0.40	0.59
No. Obs.	3141.00	1116.00	1116.00	453.00

Table 3.4: County-Level Specification: Effect of Stay-at-Home Orders on Local Unemployment Rate

Bibliography

Altig, David, Lawrence J. Christiano, Martin Eichenbaum, and Jesper Linde, "Firm-specific Capital, Nominal Rigidities and the Business Cycle," *Review of Economic Dynamics*, 2011, pp. 225–247.

Ameriks, John, Joseph Briggs, Andrew Caplin, Minjoon Lee, Matthew Shapiro, and Christopher Tonetti, "Older Americans Would Work Longer if Jobs Were Flexible," *American Economic Journal: Macroeconomics*, 2020, *12* (1), 174–209.

Atkeson, Andrew, "What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios," No. w26867, National Bureau of Economic Research 2020.

Attanasio, Orazio, Peter Levell, Hamish Low, and Virginia Sánchez-Marcos, "Aggregating Elasticities: Intensive and Extensive Margins of Female Labour Supply," *Econometrica*, 2018, *86* (6), 2049–2082.

Baker, Scott R, Nicholas Bloom, Steven J Davis, and Stephen J Terry, "COVID-Induced Economic Uncertainty," No. w26983, National Bureau of Economic Research 2020.

Ball, Laurence and David Romer, "Real Rigidities and the Non-neutrality of Money," *Review of Economic Studies*, 1990.

Barbanchon, Thomas Le, Roland Rathelot, and Alexandra Roulet, "Unemployment Insurance and Reservation Wages: Evidence from Administrative Data," *Journal of Public Economics*, 2019, *171*, 1–17.

_ , _ , and _ , "Gender Differences in Job Search: Trading Off Commute Against Wage," *The Quarterly Journal of Economics*, 2021, *136* (1), 381–426.

Barth, Erling and Harald Dale-Olsen, "Monopsonistic Discrimination, Worker Turnover and the Gender Wage Gap," *Labour Economics*, 2009, pp. 589 – 597.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu, "Monopsony In Movers: the Elasticity of Labor Supply to Firm Wage Policies," *Working Paper*, 2020.

Beffy, Magali, Richard Blundell, Antoine Bozio, Guy Laroque, and Maxime To, "Labour Supply and Taxation with Restricted Choices," *Journal of Econometrics*, 2019, *211* (1), 16–46.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor Market Power," *NBER Working Paper*, 2021.

Bianchi, Marco, Bjorn Gudmundsson, and Gylfi Zoega, "Iceland's Natural Experiment in Supply-Side Economics," *American Economic Review*, 2001, *91* (5), 1564–1579.

Bils, Mark and James A Kahn, "What inventory behavior tells us about business cycles," *American Economic Review*, 2000, *90* (3), 458–481.

_ , Peter Klenow, and Oleksiy Kryvstov, "Sticky Prices and Monetary Policy Shocks," *Federal Reserve of Minneapolis Quarterly Review*, 2003.

_ , Yongsung Chang, and Sun-Bin Kim, "Comparative Advantage and Unemployment," *Journal of Monetary Economics*, 2012, *59* (2), 150–165.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten, "Consumption Inequality and Family Labor Supply," *American Economic Review*, 2016, *106* (2), 387–435.

Brinca, Pedro, João B. Duarte, and Miguel Faria e Castro, "Measuring Sectoral Supply and Demand Shocks During COVID-19," FRB St. Louis Working Paper No. 2020-011 2020.

Bu, Chunya, Wenbin Wu, and John Rogers, "A Unified Measure of Fed Monetary Policy Shocks," *Journal of Monetary Economics*, 2021, *118*, 331–349.

Burdett, Kenneth and Dale Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, *39.*

Calvo, Guillermo, "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 1983, *12* (3), 383–398.

Card, David and Dean Hyslop, "Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers," *Econometrica*, 2005, *73* (6), 1723–1770.

Carvalho, Carlos and Fernanda Nechio, "Factor Specificity and Real Rigidities," *Review of Economics Dynamics*, October 2016.

Chang, Yongsung and Sun-Bin Kim, "From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy," *International Economic Review*, 2006, *47* (1), 1–27.

_ and _ , "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations," *American Economic Review*, 2007, *97* (5), 1939–1956.

Chetty, Raj, "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 2012, *80* (3), 969–1018.

_ , Adam Guren, Day Manoli, and andrea Weber, "Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? a Meta-analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual 2012*, 2012.

_ , John Friedman, Nathaniel Hendren, Michael Stepner, and the Opportunity Insights Team, "The Economic Impacts of COVID-19: Evidence from a New Public Database Built from Private Sector Data," Working Paper May 2020.

_ , _ , Tore Olsen, and Luigi Pistaferri, "Adjustment Costs, Firm Responses, and Micro Vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records," *The Quarterly Journal of Economics*, 2011, *126* (2), 749–804.

Chodorow-Reich, Gabriel, "Regional Data in Macroeconomics: Some Advice for Practitioners," *Journal of Economic Dynamics and Control*, Forthcoming.

_ and Loukas Karabarbounis, "The Cyclicality of the Opportunity Cost of Employment," *Journal of Political Economy*, 2016, *124* (6), 1563–1618.

Cogan, John, "Fixed Costs and Labor Supply," *Econometrica*, 1981, pp. 945–963.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber, "The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending," Working Paper 27141, National Bureau of Economic Research May 2020.

_ , _ , and _ , "Labor Markets during the COVID-19 Crisis: A Preliminary View," No. w227017, National Bureau of Economic Research 2020.

Correia, Sergio, Stephan Luck, Emil Verner et al., "Fight the Pandemic, Save the Economy: Lessons from the 1918 Flu," No. 20200327, Federal Reserve Bank of New York 2020.

Davis, Steven, Jason Faberman, and John Haltiwanger, "The Flow Approach to Labor Markets: New Data Sources and Micro-macro Links," *Journal of Economic Perspectives*, 2006, *20* (3), 3–26.

Dedola, Luca and Francesco Lippi, "The Monetary Transmission Mechanism: Evidence From the Industries of Five OECD Countries," *European Economic Review*, 2005, *49*, 1543–1569.

Dennis, Michael, "Technical Overview of the AmeriSpeak Panel NORC's Probability-Based Household Panel," *NORC at the University of Chicago*, 2019.

Dingel, Jonathan I and Brent Neiman, "How many jobs can be done at home?," No. w26948, National Bureau of Economic Research 2020.

Domeij, David and Martin Floden, "The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates are Biased," *Review of Economic Dynamics*, 2006, *9* (2), 242–262.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri, "Monopsony In Online Labor Markets," *Working Paper*, 2018.

Eichenbaum, Martin S, Sergio Rebelo, and Mathias Trabandt, "The macroeconomics of epidemics," No. w26882, National Bureau of Economic Research 2020.

Faria–e–Castro, Miguel, "Fiscal Policy during a Pandemic," No. 2020-006D, Federal Reserve bank of St. Louis 2020.

Feldstein, Martin and James Poterba, "Unemployment Insurance and Reservation Wages," *Journal of Public Economics*, 1984, *23* (1-2), 141–167.

Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, Zulma Cucunubá, Gina CuomoDannenburg, Amy Dighe, Ilaria Dorigatti, Han Fu, Katy Gaythorpe, Will Green, Arran Hamlet, Wes Hinsley, Lucy C Okell, Sabine van Elsland, Hayley Thompson, Robert Verity, Erik Volz, Haowei Wang, Yuanrong Wang, Patrick GT Walker, Caroline Walters, Peter Winskill, Charles Whittaker, Christl A Donnelly, Steven Riley, and Azra C Ghani, "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand," manuscript, Imperical College 2020.

Fetzer, Thiemo, Lukas Hensel, Johannes Hermle, and Christopher Roth, "Coronavirus Perceptions and Economic Anxiety," *Review of Economics and Statistics, forthcoming*, 2020.

Friedson, Andrew I, Drew McNichols, Joseph J Sabia, and Dhaval Dave, "Did California's Shelter-in-Place Order Work? Early Coronavirus-Related Public Health Benefits," Working Paper 26992, National Bureau of Economic Research April 2020.

Galí, Jordi, *Monetary Policy, Inflation, and the Business Cycle*, Princeton University Press, 2008.

Galí, Jordi, "The Return of the Wage Phillips Curve," *Journal of the European Economic Association*, 2011, *9* (3), 436–461.

Gorodnichenko, Yuriy and Michael Weber, "Are Sticky Prices Costly? Evidence From the Stock Market," *American Economic Review*, 2016, *106*, 165–199.

Gourinchas, Pierre-Olivier, "Flattening the pandemic and recession curves," Working Paper 2020.

Gourio, François and Pierre-Alexandre Noual, "The Marginal Worker and the Aggregate Elasticity of Labor Supply," *Working Paper*, 2009.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory W Huffman, "Investment, capacity utilization, and the real business cycle," *The American Economic Review*, 1988, pp. 402–417.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What do data on millions of US workers reveal about life-cycle earnings risk?," Technical Report 2015.

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

Hall, Robert, "Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor," *Journal of Political Economy*, 2009, *117* (2), 281–323.

_ **and Paul Milgrom**, "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, August 2008, *98* (4), 1653–1674.

Hansen, Gary, "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 1985, *16* (3), 309 – 327.

Hartl, Tobias, Klaus Wälde, and Enzo Weber, "Measuring the impact of the German public shutdown on the spread of COVID19," Covid economics, Vetted and real-time papers, CEPR press, 1, 25-32 2020.

Hassan, Tarek Alexander, Stephan Hollander, Laurence van Lent, and Ahmed Tahoun, "Firm-level Exposure to Epidemic Diseases: Covid-19, SARS, and H1N1," No. w26971, National Bureau of Economic Research 2020.

Heckman, James, "Comments on the Ashenfelter and Kydland Papers," in "Carnegie-Rochester Conference Series on Public Policy," Vol. 21 Elsevier 1984, pp. 209–224.

_ **and Thomas MaCurdy**, "A Life Cycle Model of Female Labour Supply," *The Review of Economic Studies*, 1980, *47* (1), 47–74.

Henkel, Lukas, "Sectoral Output Effects of Monetary Policy: Do Sticky Prices Matter?," *European Central Bank Working Paper Series*, 2020.

Hobijn, Bart and Fernanda Nechio, "Sticker Shocks: Using Vat Changes to Estimate Upper-level Elasticities of Substitution," *Federal Reserve Bank of San Francisco Working Paper Series*, 2017.

Hsiang, Solomon, Daniel Allen, Sébastien Annan-Phan, Kendon Bell, Ian Bolliger, Trinetta Chong, Hannah Druckenmiller, Luna Yue Huang, Andrew Hultgren, Emma Krasovich, Peiley Lau, Jaecheol Lee, Esther Rolf, Jeanette Tseng, and Tiffany Wu, "The effect of large-scale anti-contagion policies on the COVID-19 pandemic," *Nature*, 2020.

Huggett, Mark, "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 1993, *17* (5-6), 953–969.

Jäger, Simon, Benjamin Schoefer, and Josef Zweimüller, "Marginal Jobs and Job Surplus: A Test of the Efficiency of Separations," *NBER Working Paper*, 2021.

Jäger, Simon, Benjamin Schoefer, Samuel Young, and Josef Zweimüller, "Wages and the Value of Nonemployment," *The Quarterly Journal of Economics*, 2020, *135* (4), 1905–1963.

Kahn, Lisa B, Fabian Lange, and David G Wiczer, "Labor Demand in the Time of COVID-19: Evidence from Vacancy Postings and UI Claims," Working Paper 27061, National Bureau of Economic Research April 2020.

Kaplan, Greg, Benjamin Moll, and Giovanni Violante, "Monetary Policy According to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

_ , Giovanni Violante, and Justin Weidner, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, *2014* (1), 77–138.

Keane, Michael and Richard Rogerson, "Reconciling Micro and Macro Labor Supply Elasticities: A Structural Perspective," *Annual Review of Economics*, 2015, *7* (1), 89–117.

Kimball, Miles and Matthew Shapiro, "Labor Supply: Are the Income and Substitution Effects both Large or Both Small?," *NBER Working Paper*, 2008.

Kleven, Henrik, "The EITC and the Extensive Margin: A Reappraisal," *NBER Working Paper*, 2019.

Kneip, Alois, Monika Merz, and Lidia Storjohann, "Aggregation and Labor Supply Elasticities," *Journal of the European Economic Association*, 2020, *18* (5), 2315–2358.

Koenig, Felix, Alan Manning, and Barbara Petrongolo, "Reservation Wages and the Wage Flexibility Puzzle," 2020.

Kong, Edward and Daniel Prinz, "The Impact of Non-Pharmaceutical Interventions on Unemployment During a Pandemic," No. 3581254, SSRN 2020.

Kropko, Jonathan and Robert Kubinec, "Interpretation and identification of within-unit and cross-sectional variation in panel data models," *PLOS ONE*, 04 2020, *15* (4), 1–22.

Krueger, Alan and Andreas Mueller, "A Contribution to the Empirics of Reservation Wages," *American Economic Journal: Economic Policy*, 2016, *8* (1), 142–79.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Aysegul Sahin, "Gross Worker Flows over the Business Cycle," *American Economic Review*, November 2017, *107* (11), 3447–3476.

Kudlyak, Marianna and Erin Wolcott, "Pandemic Layoffs," Working Paper May 2020.

Leahy, John, "A Survey of New Keynesian Theories of Aggregate Supply and Their Relation to Industrial Organization," *Journal of Money, Credit and Banking*, 2011.

Lewis, Daniel, Karel Mertens, and James H Stock, "US Economic Activity during the Early Weeks of the SARS-Cov-2 Outbreak," No. w26954, National Bureau of Economic Research 2020.

Ljungqvist, Lars and Thomas Sargent, "Do Taxes Explain European Employment? Indivisible Labor, Human Capital, Lotteries, and Savings," *NBER Macroeconomics Annual*, 2006, *21*, 181–246.

_ **and** _ , "The Fundamental Surplus," *American Economic Review*, 2017, *107* (9), 2630–65.

MaCurdy, Thomas, "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, 1981, *89* (6), 1059–1085.

Manning, Alan, *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton University Press, 2013.

Martinez, Isabel, Emmanuel Saez, and Michael Siegenthaler, "Intertemporal Labor Supply Substitution? Evidence from the Swiss Income Tax Holidays," *American Economic Review*, 2021, *111* (2), 506–46.

Mas, Alexandre and Amanda Pallais, "Labor Supply and the Value of Non-Work Time: Experimental Estimates from the Field," *American Economic Review: Insights*, 2019, *1* (1), 111–26.

Matheron, Julien, "Firm-specific Labor and Firm-specific Capital: Implications For the Euro-data New Phillips Curve," *International Journal of Central Banking*, 12 2006.

Mui, Preston and Benjamin Schoefer, "Reservation Raises: The Aggregate Labor Supply Curve at the Extensive Margin," *NBER Working Paper*, 2021.

Nakamura, Emi and Jón Steinsson, "Five Facts About Prices: a Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 2008, *123*, 1415–1464.

_ **and** _ , "Fiscal Stimulus In a Monetary Union: Evidence From Us Regions," *American Economic Review*, 2014, *104*, 753–792.

_ **and** _ , "High-frequency Identification of Monetary Non-neutrality: the Information Effect," *Quarterly Journal of Economics*, 2018, *133*, 1283–1330.

Park, Choonsung, "Consumption, Reservation Wages, and Aggregate Labor Supply," *Review of Economic Dynamics*, 2020, *37*, 54–80.

Prescott, Edward, "Why Do Americans Work So Much More Than Europeans?," *Federal Reserve Bank of Minneapolis Quarterly Review*, 2004, *28* (1), 2–13.

Ramey, Valerie A, "Ten Years After the Financial Crisis: What Have We Learned from the Renaissance in Fiscal Research?," *Journal of Economic Perspectives*, 2019, *33* (2), 89–114.

_ **and Kenneth D West**, "Inventories," in "Handbook of macroeconomics," Vol. 1, Elsevier, 1999, chapter 13, pp. 863–926.

Richter, David and Jürgen Schupp, "The SOEP Innovation Sample (SOEP IS)," *Schmollers Jahrbuch: Journal of Applied Social Science Studies/Zeitschrift für Wirtschafts-und Sozialwissenschaften*, 2015, *135* (3), 389–400.

Rogerson, Richard, "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, January 1988, *21* (1), 3–16.

_ **and Johanna Wallenius**, "Micro and Macro Elasticities in a Life Cycle Model With Taxes," *Journal of Economic Theory*, 2009, *144* (6), 2277–2292.

Saez, Emmanuel, Joel Slemrod, and Seth Giertz, "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review," *Journal of Economic Literature*, 2012, *50* (1), 3–50.

Sánchez, Rafael, Javier Finot, and Maurcio G. Villena, "Gender Wage Gap and Firm Market Power: Evidence From Chile," Technical Report November 2020.

Sauvagnat, Julien, Jean-Noël Barrot, and Basile Grassi, "Estimating the costs and benefits of mandated business closures in a pandemic," Working Paper May 2020.

Shimer, Robert, "Convergence in Macroeconomics: The Labor Wedge," *American Economic Journal: Macroeconomics*, January 2009, *1* (1), 280–97.

_ **and Ivan Werning**, "Reservation Wages and Unemployment insurance," *The Quarterly Journal of Economics*, 2007, *122* (3), 1145–1185.

Sigurdsson, Jósef, "Labor Supply Responses and Adjustment Frictions: A Tax-Free Year in Iceland," *Working Paper*, 2018.

Sokolova, Anna and Todd Sorensen, "Monopsony In Labor Markets: a Meta-analysis," *ILR Review*, 2020, *74*, 27–55.

Solon, Gary, Steven J Haider, and Jeffrey M Wooldridge, "What are we weighting for?," *Journal of Human resources*, 2015, *50* (2), 301–316.

Staiger, Douglas O., Joanne Spetz, and Ciaran S. Phibbs, "Is There Monopsony In the Labor Market? Evidence From a Natural Experiment," *Journal of Labor Economics*, 2010.

Stinson, Martha H., "Technical Description of Sipp Job Identification Number Editing In the 1990-1993 SIPP Panels," Technical Report, U.S. Census Bureau 07 2003.

Webber, Douglas, "Firm Market Power and the Earnings Distribution," *Labour Economics*, 2015, pp. 123–134.

_ , "Employment Adjustment Over the Business Cycle: the Impact of Competition In the Labor Market," *Working Paper*, 2018.

Woodford, Michael, *Interest and Prices*, Princeton University Press, 2003.

_ , "Firm-specific Capital and the New Keynesian Phillips Curve," *International Journal of Central Banking*, 9 2005.

Zweck, Bettina and Axel Glemser, "SOEP-IS 2016-Methodenbericht zum Befragungsjahr 2016 des SOEP-Innovationssamples," *SOEP Survey Papers*, 2018.

_ **and** _ , "SOEP-IS 2019: Survey Report on the 2017 SOEP Innovation Sample," *SOEP Survey Papers*, 2020.

Appendix A

Appendix to Chapter 1

A.1 Theoretical Appendix

The Phillips curve with Firm-Specific Labor

Here I derive the Phillips curve of the single-sector New Keynesian model with firm-specific labor supply in Section 1.2.

Firm Pricing Decision. Let P_t^* be the optimal reset price of a firm that resets their price in time t. Then, the pricing decision follows

$$\max_{P_t^*} \mathbb{E}_t \sum_{s=t}^{\infty} \gamma^s \left[Q_{t,s} \left(P_t^* Y_{s|t} - \Sigma_s(Y_{s|t}) \right) \right], \tag{A.1.1}$$

where $Q_{t,s}$ is the S.D.F. between t and s, and Σ_s is the nominal total cost function in time s, as a function of output. For a firm that sets their price in t, demand for the firm's output $Y_{s|t}$ is given by $Y_{s|t} = Y_s \left(\frac{P_t^*}{P_s} \right)^{-\theta_p}$, implying $\frac{\partial Y_{s|t}}{\partial P_t^*} = -\epsilon \frac{Y_{s|t}}{P_t^*}$. The firm's optimal pricing condition is given by

$$\mathbb{E}_t \sum_{s=t}^{\infty} \gamma^k \left[Q_{t,s} Y_{s|t} \left(P_t^* - \frac{\epsilon}{\epsilon - 1} \Sigma'(Y_{s|t}) \right) \right] = 0, \tag{A.1.2}$$

so on average (weighted) deviations from the steady state markup are zero. A first-order Taylor expansion of (A.1.2) yields

$$p_t^* - p_{t-1} = (1 - \beta\gamma) \sum_{s=t}^{\infty} (\beta\gamma)^k \mathbb{E}_t \left[mc_{s|t} + p_s - p_{t-1} \right], \tag{A.1.3}$$

where $mc_{s|t}$ is the log deviation of the firm's real marginal cost (deflated by the sectoral price index) from steady state.

The Firm's Marginal Cost Function. The firm-specific nominal marginal cost function is comprised of labor costs:

$$\Sigma_{s|t} = W_{s|t} L_{s|t}, \tag{A.1.4}$$

and the marginal cost function takes into account the increase in the firm's wages necessary to increase labor supply at the firm level:

$$\begin{aligned}
\frac{\partial \Sigma_{s|t}}{\partial Y_{s|t}} &= \frac{\partial W_{s|t}}{\partial L_{s|t}} \frac{\partial L_{s|t}}{\partial Y_{s|t}} L_{s|t} + W_{s|t} \frac{\partial L_{s|t}}{\partial Y_{s|t}} && \text{(A.1.5)} \\
&= \left(\frac{\partial W_{s|t}}{\partial L_{s|t}} L_{s|t} + W_{s|t} \right) \frac{\partial L_{s|t}}{\partial Y_{s|t}} && \text{(A.1.6)} \\
&= \frac{1+\theta}{\theta} W_{s|t} \frac{\partial L_{s|t}}{\partial Y_{s|t}}, && \text{(A.1.7)}
\end{aligned}$$

where the substitution $\frac{\partial W_{s|t}}{\partial L_{s|t}} L_{s|t} = \frac{1+\theta}{\theta} W_{s|t}$ in equation (A.1.7) arises from the household's labor supply to the firm, $W_{s|t} = \left(\frac{L_{s|t}}{L_s} \right)^{1/\theta} W_s$. In log deviations, equation A.1.7 can be written as

$$\begin{aligned}
mc_{s|t} &= w_{s|t} - mpl_{s|t} - p_t && \text{(A.1.8)} \\
&= w_{s|t} - mpl_{s|t} - p_t, && \text{(A.1.9)}
\end{aligned}$$

where $mpl_{s|t}$ is the (log) marginal cost of the resting firm in period s. In log deviations, labor supply to the firm is

$$w_{s|t} = w_s + \frac{1}{\theta} \left(l_{s|t} - l_s \right), \tag{A.1.10}$$

and labor demand at the firm level is determined by output, which is determined by the firm's relative price and aggregate output:

$$\begin{aligned}
l_{s|t} &= \frac{1}{1-\alpha} \left(y_{s|t} - z_s \right) && \text{(A.1.11)} \\
&= \frac{1}{1-\alpha} \left[-\epsilon (p_t^* - p_s) + y_s - z_s \right], && \text{(A.1.12)}
\end{aligned}$$

and aggregate labor demand l_s is, up to a first-order approximation, given by

$$l_s = \frac{1}{1-\alpha} \left(y_s - z_s \right). \tag{A.1.13}$$

Plugging this into equation (A.1.10), the firm's wages are given by

$$w_{s|t} = w_s - \epsilon \frac{1/\theta}{1-\alpha} \left(p_t^* - p_s \right). \tag{A.1.14}$$

The marginal product of labor of the firm is given by

$$mpl_{s|t} = z_s - \alpha l_{s|t} \tag{A.1.15}$$
$$= z_s - \frac{\alpha}{1-\alpha}\left[-\epsilon(p_t^* - p_s) + y_s - z_s\right], \tag{A.1.16}$$

where the substitution for $l_{s|t}$ is from (A.1.12). So, substituting in the expressions for $w_{s|t}$ and $mpl_{s|t}$ into (A.1.9), the firm's real marginal cost is

$$mc_{s|t} = w_s - \epsilon\frac{\alpha + 1/\theta}{1-\alpha}\left(p_t^* - p_s\right) - z_s - \frac{\alpha}{1-\alpha}\left[y_s - z_s\right] \tag{A.1.17}$$
$$= mc_s - \epsilon\frac{\alpha + 1/\theta}{1-\alpha}\left(p_t^* - p_s\right), \tag{A.1.18}$$

where mc_s is the aggregate average marginal cost analog $mc_s = w_s - z_s - \alpha l_s$. Substituting (A.1.18) into (A.1.3) and rearranging terms yields

$$p_t^* = (1-\beta\gamma)\sum_{s=t}^{\infty}(\beta\gamma)^k \mathbb{E}_t\left[p_s - \Omega mc_s\right], \tag{A.1.19}$$

where $\Omega = \frac{1}{1+\epsilon\frac{\alpha+1/\theta}{1-\alpha}}$. Following Galí (2008), this can be written recursively and in terms of the output gap to yield the Phillips curve,

$$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + \frac{(1-\gamma)(1-\beta\gamma)}{\gamma}\frac{\sigma + \frac{1/\eta+\alpha}{1-\alpha}}{1+\epsilon\frac{\alpha+1/\theta}{1-\alpha}}.(y_t - y_t^n) \tag{A.1.20}$$

Sectoral Phillips Curves with Firm-specific Labor

In this section, I derive the sector-specific Phillips curves in Section 1.3.

Let $P_{j,t}^*$ be the optimal reset price of a firm in sector j that resets their price in time t. Then, the pricing decision follows

$$\max_{P_{j,t}^*}\mathbb{E}_t\sum_{s=t}^{\infty}\gamma_j^s\left[Q_{t,s}\left(P_{j,t}^* Y_{j,s|t} - \Sigma_{j,s}(Y_{j,s|t})\right)\right] \tag{A.1.21}$$

from here on, the derivation follows the single-sector Phillips curve derivation in Appendix Section A.1 from equations (A.1.2) through (A.1.19), except that sector-specific aggregates take the place of economy-wide aggregates; e.g. $Y_{j,s}$ instead of Y_s. Doing so yields the firm's pricing choice

$$p_{j,t}^* = (1-\beta\gamma_j)\sum_{s=t}^{\infty}(\beta\gamma_j)^k \mathbb{E}_t\left[p_{j,s} - \Omega_j mc_{j,s}\right], \tag{A.1.22}$$

where $p_{j,s}$ is the log deviation of the sector's price index, $mc_{j,s}$ is the sector real marginal cost analog $mc_{j,s} = w_{j,s} - z_{j,s} - \alpha l_{j,s}$, and $\Omega = \frac{1}{1+\epsilon_j \frac{\alpha_j + 1/\theta_j}{1-\alpha_j}}$.

The sectoral real marginal cost analog, $rmc_{j,t}$, can be written as

$$rmc_{j,t} = w_{j,t} - p_{j,t} + y_{jt} - l_{jt} \tag{A.1.23}$$

$$= w_{j,t} - p_{j,t} - \frac{1}{1-\alpha_j}\left(z_{j,t} - \alpha_j y_{j,t}\right) \tag{A.1.24}$$

$$= \frac{1}{\lambda}\left(l_{j,t} - l_t\right) + w_t + \frac{1}{\eta}\left(y_{j,t} - y_t\right) - p_t - \frac{1}{1-\alpha_j}\left(z_{j,t} - \alpha_j y_{j,t}\right) \tag{A.1.25}$$

$$= \frac{1}{\lambda}\left(\frac{y_{jt} - z_{jt}}{1-\alpha_j} - l_t\right) + w_t + \frac{1}{\eta}\left(y_{j,t} - y_t\right) - p_t - \frac{1}{1-\alpha_j}\left(z_{j,t} - \alpha_j y_{j,t}\right) \tag{A.1.26}$$

$$= \left(\frac{1/\lambda + \alpha_j}{(1-\alpha_j)} + \frac{1}{\eta}\right) y_{jt} - \left(\frac{\lambda+1}{\lambda(1-\alpha_j)}\right) z_{jt} + \left(\sigma - \frac{1}{\eta}\right) y_t + \left(\frac{1}{\eta} - \frac{1}{\lambda}\right) l_t, \tag{A.1.27}$$

where $l_{j,t}$ has been substituted out for $\frac{1}{1-\alpha_j}(y_{jt} - z_{jt})$ using a first-order approximation. So, the sectoral Phillips curve is:

$$\pi_{jt} = \beta \mathbb{E}_t \pi_{j,t+1} + \frac{(1-\gamma_j)(1-\beta\gamma_j)}{\gamma_j} \frac{1}{1+\epsilon_j \frac{\alpha_j + 1/\theta_j}{1-\alpha_j}} \times \tag{A.1.28}$$

$$\left[\left(\frac{1/\lambda + \alpha_j}{(1-\alpha_j)} + \frac{1}{\eta}\right) \check{y}_{jt} - \left(\frac{\lambda+1}{\lambda(1-\alpha_j)}\right) \check{z}_{jt} + \left(\sigma - \frac{1}{\eta}\right) \check{y}_t + \left(\frac{1}{\eta_t} - \frac{1}{\lambda}\right) \check{l}_t\right]. \tag{A.1.29}$$

A.2 Empirical Appendix

Table A.1: Estimates of Firm-specific Labor Supply Elasticities

NAICS Code(s)	EE Sep. Elasticity	EN Sep. Elasticity	Search Premium	Recruits Share EE	Separations Share EE	FSLS Elasticity
111	-0.937	-1.348	1.012	0.592	0.629	1.672
	(0.114)	(0.168)	(0.195)			
112	-1.243	-1.697	1.336	0.656	0.663	2.192
	(0.204)	(0.305)	(0.327)			
113	-0.875	-1.073	0.386	0.650	0.697	1.738
	(0.202)	(0.273)	(0.294)			
114	-1.318	-3.115	0.954	0.725	0.729	2.867
	(0.496)	(1.135)	(1.126)			
212	-0.479	-0.434	0.933	0.717	0.673	0.650
	(0.241)	(0.360)	(0.467)			
211	-0.310	-0.387	0.469	0.770	0.739	0.520
	(0.165)	(0.267)	(0.317)			
23	-0.486	-0.633	0.500	0.716	0.710	0.868
	(0.027)	(0.044)	(0.048)			
311	-1.095	-1.273	0.871	0.680	0.678	1.966
	(0.077)	(0.113)	(0.135)			
312	-0.811	-0.727	1.296	0.774	0.724	1.253
	(0.204)	(0.325)	(0.471)			
314	-1.048	-1.208	-0.534	0.743	0.750	2.284
	(0.386)	(0.761)	(0.754)			
313, 315	-1.047	-1.057	0.550	0.706	0.676	1.891
	(0.104)	(0.155)	(0.189)			
322	-1.137	-1.193	0.821	0.727	0.698	2.023
	(0.137)	(0.201)	(0.261)			
325	-1.066	-0.971	0.799	0.763	0.719	1.853
	(0.100)	(0.156)	(0.184)			
324	-0.920	-1.549	0.845	0.726	0.709	1.769
	(0.254)	(0.385)	(0.423)			
326	-1.165	-1.340	0.775	0.727	0.741	2.187
	(0.119)	(0.213)	(0.211)			
316	-0.459	-1.805	2.818	0.696	0.692	0.472
	(0.437)	(0.645)	(0.837)			
321	-1.156	-1.240	1.019	0.714	0.730	2.069
	(0.145)	(0.242)	(0.276)			
337	-0.793	-0.508	0.275	0.722	0.704	1.405
	(0.145)	(0.223)	(0.260)			
327	-0.943	-0.734	0.646	0.733	0.695	1.601
	(0.142)	(0.226)	(0.256)			

Continued on next page

Estimates of Firm-specific Labor Supply Elasticities (ctd)

NAICS Code(s)	EE Sep. Elasticity	EN Sep. Elasticity	Search Premium	Recruits Share EE	Separations Share EE	FSLS Elasticity
331	-1.182	-0.797	0.234	0.778	0.705	2.087
	(0.128)	(0.201)	(0.244)			
332	-0.957	-0.988	0.403	0.747	0.736	1.806
	(0.086)	(0.149)	(0.165)			
333	-0.893	-0.987	0.636	0.748	0.718	1.618
	(0.082)	(0.132)	(0.156)			
334, 335	-0.731	-0.660	0.485	0.758	0.724	1.291
	(0.078)	(0.120)	(0.137)			
336	-0.812	-0.876	0.372	0.771	0.721	1.504
	(0.063)	(0.103)	(0.121)			
339	-0.681	-1.097	0.833	0.720	0.700	1.235
	(0.117)	(0.187)	(0.217)			
482	-0.987	-1.211	1.280	0.756	0.724	1.681
	(0.278)	(0.432)	(0.506)			
485, 487	-0.947	-1.100	1.151	0.756	0.733	1.625
	(0.120)	(0.203)	(0.252)			
484, 492	-0.854	-1.148	0.827	0.736	0.728	1.560
	(0.076)	(0.126)	(0.145)			
493	-1.160	-1.026	1.056	0.695	0.720	2.003
	(0.219)	(0.375)	(0.338)			
491	-1.670	-1.668	1.055	0.717	0.669	2.930
	(0.150)	(0.227)	(0.263)			
483	-0.767	-1.335	1.925	0.713	0.689	1.132
	(0.246)	(0.385)	(0.561)			
481	-0.946	-0.610	0.325	0.747	0.716	1.675
	(0.110)	(0.180)	(0.197)			
488	-0.760	-0.841	0.668	0.705	0.698	1.339
	(0.141)	(0.221)	(0.268)			
515	-0.860	-1.329	-0.040	0.729	0.764	1.883
	(0.176)	(0.298)	(0.314)			
517	-1.092	-0.587	0.375	0.747	0.717	1.901
	(0.107)	(0.169)	(0.199)			
221, 562	-0.894	-0.820	0.686	0.760	0.715	1.549
	(0.102)	(0.161)	(0.186)			
423	-0.927	-1.251	0.862	0.737	0.714	1.691
	(0.082)	(0.131)	(0.148)			
424	-1.025	-1.314	0.976	0.716	0.701	1.837
	(0.077)	(0.117)	(0.139)			
444	-1.155	-1.221	0.897	0.678	0.688	2.058
	(0.099)	(0.145)	(0.169)			
452	-0.968	-1.201	0.944	0.639	0.685	1.738
	(0.060)	(0.090)	(0.108)			

Continued on next page

Estimates of Firm-specific Labor Supply Elasticities (ctd)

NAICS Code(s)	EE Sep. Elasticity	EN Sep. Elasticity	Search Premium	Recruits Share EE	Separations Share EE	FSLS Elasticity
445	-1.063	-1.533	1.208	0.620	0.680	1.921
	(0.048)	(0.077)	(0.091)			
441	-0.879	-0.847	0.796	0.734	0.735	1.539
	(0.085)	(0.136)	(0.153)			
447	-0.845	-1.175	0.579	0.669	0.706	1.643
	(0.142)	(0.224)	(0.252)			
442, 722	-0.642	-1.306	0.613	0.666	0.684	1.306
	(0.164)	(0.257)	(0.277)			
443, 446, 451, 453	-0.751	-1.365	0.886	0.663	0.684	1.422
	(0.052)	(0.083)	(0.094)			
448, 454	-0.876	-1.131	0.847	0.637	0.671	1.576
	(0.086)	(0.128)	(0.149)			
521, 522	-0.884	-0.978	0.828	0.769	0.735	1.563
	(0.091)	(0.143)	(0.171)			
523, 525	-1.392	-0.725	-0.241	0.767	0.756	2.659
	(0.208)	(0.330)	(0.320)			
524	-0.872	-0.799	0.886	0.741	0.723	1.473
	(0.105)	(0.165)	(0.184)			
531	-0.986	-0.924	0.805	0.701	0.711	1.728
	(0.091)	(0.136)	(0.161)			
814	-0.721	-0.607	0.572	0.569	0.568	1.145
	(0.157)	(0.204)	(0.255)			
721	-0.839	-1.192	1.128	0.642	0.671	1.428
	(0.080)	(0.124)	(0.137)			
812	-0.814	-0.906	0.979	0.647	0.659	1.330
	(0.111)	(0.166)	(0.186)			
115, 323, 511, 512, 518, 519,	-0.829	-1.053	0.496	0.697	0.714	1.592
532, 533, 541, 561, 811	(0.024)	(0.039)	(0.042)			
711, 713	-0.885	-1.463	1.296	0.565	0.619	1.512
	(0.077)	(0.102)	(0.129)			
621	-0.612	-0.941	0.688	0.720	0.716	1.122
	(0.046)	(0.077)	(0.084)			
622	-0.559	-0.992	0.751	0.750	0.746	1.037
	(0.039)	(0.069)	(0.076)			
611	-0.838	-1.361	0.978	0.778	0.771	1.571
	(0.037)	(0.066)	(0.067)			

Note: This table shows the firm-specific labor supply and the estimates of the underlying components (the elasticities of separation to employment and non-employment and the search premium term). Standard errors in parentheses. For rows where multiple NAICS 3-digit industries are listed, these industries were estimated together due to imperfect concordance from Census industry codes and NAICS 3-digit codes, i.e. there was one dummy variable representing those industries.

Estimates of Firm-specific Labor Supply Elasticities (ctd)

NAICS Code(s)	EE Sep. Elasticity	EN Sep. Elasticity	Search Premium	Recruits Share EE	Separations Share EE	FSLS Elasticity
623, 624	-0.629	-1.180	0.821	0.675	0.696	1.179
	(0.042)	(0.072)	(0.076)			
712	-1.001	-1.162	0.983	0.583	0.597	1.681
	(0.197)	(0.243)	(0.304)			
813	-0.934	-0.810	0.593	0.666	0.640	1.589
	(0.131)	(0.173)	(0.205)			
921, 922	-0.869	-1.308	1.034	0.742	0.710	1.561
	(0.068)	(0.101)	(0.124)			
923	-0.970	-1.081	0.881	0.710	0.720	1.728
	(0.142)	(0.226)	(0.243)			
924, 925	-1.280	-1.563	0.667	0.715	0.722	2.460
	(0.217)	(0.290)	(0.334)			
926, 927	-0.789	-1.543	1.610	0.705	0.708	1.327
	(0.158)	(0.222)	(0.268)			
928	-1.206	-1.038	1.518	0.747	0.726	1.948
	(0.148)	(0.255)	(0.327)			

Note: This table shows the firm-specific labor supply and the estimates of the underlying components (the elasticities of separation to employment and non-employment and the search premium term). Standard errors in parentheses. For rows where multiple NAICS 3-digit industries are listed, these industries were estimated together due to imperfect concordance from Census industry codes and NAICS 3-digit codes, i.e. there was one dummy variable representing those industries.

Table A.2: Estimates of Differential IRFs

IRF Horizon	Price Level	Price Level	Real Output	Real Output	Employment	Employment	Wages	Wages
0 months	-0.007	-0.025	0.011	0.037	-0.020	-0.011	0.017	0.025
	(0.025)	(0.034)	(0.036)	(0.050)	(0.026)	(0.039)	(0.013)	(0.020)
3 months	0.004	-0.019	-0.015	0.010	-0.023	-0.023	0.005	0.020
	(0.022)	(0.028)	(0.037)	(0.057)	(0.019)	(0.030)	(0.022)	(0.037)
6 months	-0.003	-0.020	-0.017	-0.027	-0.039	-0.024	0.012	0.029
	(0.023)	(0.030)	(0.042)	(0.059)	(0.029)	(0.048)	(0.029)	(0.046)
9 months	-0.004	-0.031	-0.073	-0.055	-0.015	0.014	0.035	0.013
	(0.040)	(0.050)	(0.060)	(0.084)	(0.064)	(0.118)	(0.040)	(0.048)
12 months	-0.011	-0.029	-0.054	-0.101	-0.061	-0.053	0.032	0.021
	(0.043)	(0.052)	(0.077)	(0.099)	(0.073)	(0.134)	(0.047)	(0.060)
15 months	0.006	-0.018	-0.046	-0.140	-0.124**	-0.149	0.042	0.028
	(0.042)	(0.046)	(0.076)	(0.086)	(0.060)	(0.111)	(0.057)	(0.073)
18 months	0.033	-0.001	-0.075	-0.188	-0.119	-0.160	0.033	0.007
	(0.045)	(0.056)	(0.099)	(0.122)	(0.080)	(0.143)	(0.065)	(0.077)
21 months	0.033	-0.017	-0.039	-0.184	-0.182**	-0.301**	0.037	-0.018
	(0.056)	(0.068)	(0.101)	(0.134)	(0.080)	(0.136)	(0.066)	(0.095)
24 months	0.030	-0.035	0.015	-0.134	-0.224**	-0.372**	0.022	-0.024
	(0.066)	(0.085)	(0.111)	(0.155)	(0.090)	(0.152)	(0.073)	(0.103)
Controls	Yes	No	Yes	No	Yes	No	Yes	No

Note: Standard errors in parentheses. The differential IRFs are the estimates of $b^{y,h}$ in Equation (1.24). Controls included were the log frequency of price adjustment, durables dummy, industry interest expenses over sales, leverage ratio, and short-term debt ratio as described in Section 1.4. Coefficients are significantly different from zero at the 1% (***), 5% (**), and 10% (*).

Table A.3: Robustness to Controls: Price Level

Industry Characteristic	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log Labor Supply Elasticity	0.030	-0.000	0.036	0.030	0.013	-0.035	-0.026
	(0.066)	(0.078)	(0.072)	(0.067)	(0.064)	(0.085)	(0.090)
Log Freq. Price Adj.		-0.038				-0.037	-0.050
		(0.032)				(0.037)	(0.041)
Interest Rate Burden			-0.506			0.676	-0.146
			(0.935)			(2.143)	(2.419)
Leverage Ratio			0.072			0.226	0.305
			(0.231)			(0.356)	(0.377)
Short-term Debt Ratio			-0.024			-0.161	-0.265
			(0.642)			(0.768)	(0.795)
Durable Dummy				-0.009		-0.015	-0.002
				(0.048)		(0.053)	(0.057)
Frac. Small Estabs.					-0.340*	-0.435	-0.459
					(0.164)	(0.237)	(0.255)
Log Labor Share							-0.132
							(0.145)
Log Profit Share							-0.023
							(0.046)
N	49	40	49	49	49	40	40
R^2	0.00	0.04	0.02	0.01	0.09	0.15	0.18

Note: The estimates reported are $\hat{b}^{y,h}$ from Equation (1.24), the coefficient on the log firm-specific labor supply elasticity, for $h = 24$. The construction of the control variables are described in Section 1.4. Standard errors in parentheses; coefficients are significantly different from zero at the 1% (***), 5% (**), and 10% (*) levels.

Table A.4: Robustness to Controls: Real Output

Industry Characteristic	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log Labor Supply Elasticity	0.015	0.057	-0.114	0.021	0.009	-0.134	-0.102
	(0.111)	(0.132)	(0.144)	(0.107)	(0.114)	(0.155)	(0.129)
Log Freq. Price Adj.		0.016				0.077	-0.070
		(0.080)				(0.087)	(0.084)
Interest Rate Burden			-0.774			5.212	-1.815
			(4.869)			(8.504)	(7.294)
Leverage Ratio			1.167			0.889	0.698
			(0.812)			(0.896)	(0.773)
Short-term Debt Ratio			-0.811			-0.909	-1.116
			(1.311)			(1.257)	(1.053)
Durable Dummy				-0.135		-0.130	-0.118
				(0.072)		(0.084)	(0.073)
Frac. Small Estabs.					0.141	-0.444	0.075
					(0.376)	(0.491)	(0.447)
Log Labor Share							-0.136
							(0.231)
Log Profit Share							0.182*
							(0.085)
N	33	28	33	33	33	28	28
R^2	0.00	0.01	0.10	0.11	0.01	0.27	0.55

Note: The estimates reported are $\hat{b}^{y,h}$ from Equation (1.24), the coefficient on the log firm-specific labor supply elasticity, for $h = 24$. The construction of the control variables are described in Section 1.4. Standard errors in parentheses; coefficients are significantly different from zero at the 1% (***), 5% (**), and 10% (*) levels.

Table A.5: Robustness to Controls: Production Employment

Industry Characteristic	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log Labor Supply Elasticity	-0.224*	-0.325*	-0.215*	-0.220*	-0.223*	-0.372*	-0.369*
	(0.090)	(0.154)	(0.091)	(0.087)	(0.091)	(0.152)	(0.158)
Log Freq. Price Adj.		0.039				0.020	-0.004
		(0.047)				(0.047)	(0.048)
Interest Rate Burden			0.756			3.393	1.419
			(0.502)			(2.789)	(3.011)
Leverage Ratio			0.002			-0.643	-0.371
			(0.187)			(0.418)	(0.444)
Short-term Debt Ratio			-0.629			-0.199	-0.472
			(0.849)			(0.991)	(0.991)
Durable Dummy				-0.113*		-0.144*	-0.128
				(0.046)		(0.060)	(0.063)
Frac. Small Estabs.					-0.051	0.218	0.266
					(0.165)	(0.261)	(0.278)
Log Labor Share							-0.226
							(0.169)
Log Profit Share							0.009
							(0.053)
N	60	34	60	60	60	34	34
R^2	0.10	0.13	0.14	0.18	0.10	0.31	0.39

Note: The estimates reported are $\hat{b}^{y,h}$ from Equation (1.24), the coefficient on the log firm-specific labor supply elasticity, for $h = 24$. The construction of the control variables are described in Section 1.4. Standard errors in parentheses; coefficients are significantly different from zero at the 1% (***), 5% (**), and 10% (*) levels.

Table A.6: Robustness to Controls: Average Real Weekly Production Earnings

Industry Characteristic	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log Labor Supply Elasticity	0.022	0.029	0.037	0.025	0.023	-0.024	0.044
	(0.073)	(0.122)	(0.067)	(0.071)	(0.073)	(0.103)	(0.094)
Log Freq. Price Adj.		-0.018				0.002	-0.023
		(0.037)				(0.032)	(0.029)
Interest Rate Burden			-0.107			0.751	-0.577
			(0.371)			(1.893)	(1.793)
Leverage Ratio			0.468**			0.663*	0.877**
			(0.138)			(0.284)	(0.264)
Short-term Debt Ratio			-0.087			-0.170	-0.371
			(0.628)			(0.673)	(0.590)
Durable Dummy				-0.076*		-0.057	-0.016
				(0.037)		(0.041)	(0.038)
Frac. Small Estabs.					-0.024	-0.439*	-0.530**
					(0.132)	(0.177)	(0.166)
Log Labor Share							-0.328**
							(0.100)
Log Profit Share							-0.060
							(0.031)
N	60	34	60	60	60	34	34
R^2	0.00	0.01	0.19	0.07	0.00	0.42	0.60

Note: The estimates reported are $\hat{b}^{y,h}$ from Equation (1.24), the coefficient on the log firm-specific labor supply elasticity, for $h = 24$. The construction of the control variables are described in Section 1.4. Standard errors in parentheses; coefficients are significantly different from zero at the 1% (***), 5% (**), and 10% (*) levels.

Appendix B

Appendix to Chapter 2

B.1 NORC at UChicago Survey Questions

Question for the Employed The following is a hypothetical situation we ask you to think about regarding your current job, so please read [listen] carefully and try to think about what you would do if presented with this choice.

Suppose, for reasons unrelated to you, your employer offers you the following choice: Either you take unpaid time off from work for one month, or you stay in your job for that month and only receive a fraction of your regular salary. No matter what choice you take, after the month is over, your salary will return to normal.

In this hypothetical scenario, you cannot take an additional job to make up for the lost income during that month.

Assume this choice is real and you have to make it. At what point would the cut in your salary be just large enough that you would choose the unpaid month of time off over working for the month at that lower salary?

For example, an answer of 5% means that a 5% wage cut would be the point where you would choose to take unpaid time off for the month instead of working for 5% lower pay during that month. But if the wage cut was less than 5%, you would instead choose to work for that than take unpaid time off. Choose any percentage between 1% to 100%, where the cut wage cut is just large enough that you would prefer to not work at all for no pay than work at reduced pay for that month.

Question for the Unemployed The following is a hypothetical situation we ask you to think about a potential job you may be looking for, so please read [listen] carefully and try to think about what you would do if presented with this choice.

Suppose you have found the kind of job you are looking for and the employer would like to hire you. The regular start date for the job is one month away. As an alternative, your employer offers you the option to start working immediately, rather than waiting a month.

However, if you chose to start work immediately, for that first month, you will only receive a fraction of the regular salary. The job is otherwise exactly the same. No matter what choice you take, after the month is over, the salary will then resume at the regular salary.

In this hypothetical scenario, you cannot take an additional job to make up for the lost income during that month.

Assume this choice is real and you have to make it. At what point would the cut in your salary be just large enough that you would choose the waiting a month without working and without the salary over starting the job immediately for the first month at that lower salary?

For example, an answer of 5% means that a 5% wage cut would be the point where you would choose to wait a month without working instead of working for % lower pay during that month. But if the wage cut was less than 5%, you would instead choose to work at that wage than wait a month without working. Choose any percentage between 1% to 100%, where the cut wage cut is just large enough that you would prefer to not work at all for no pay than work at reduced pay for that month.

Question for the Out of the Labor Force The following is a hypothetical situation that may not have anything to do with your actual situation, but please read [listen] carefully and try to think about what you would do if presented with this choice.

Think of the range of jobs that you would realistically be offered if you searched for jobs (even if you currently are not looking for a job and may not accept any of these potential jobs).

Suppose you had such job offers in hand. Currently you would likely not take such jobs, at least not at the usual salary. However, suppose the employer were nevertheless trying hard to recruit you, specifically by offering an additional sign-up bonus. The requirement to receive the bonus is that you will work for at least one month. The bonus comes as a raise of the first month's salary. This sign-up bonus will only be paid in the first month (on top of the regular salary that month), afterwards the salary returns to the regular salary.

Assume this choice is real and you have to make it. We would like to learn whether there is a point at which the bonus in the first month is just high enough that you would take the job.

5% means you would take the job if your employer paid a bonus of just 5% of the regular salary in the first month. 100% means you would require a bonus as large as the regular salary. 500% would mean you require a bonus equal to five times as large as the regular salary.

Choose any percentage bonus that would be just high enough that you would take the job. You can enter a high number (e.g., 100,000%) if you think you would not take any job, even if it paid a lot.

B.2 German Socio-Economic Panel Survey Questions (English Translations)

The questions below are the (authors') English translations; the German original text is in Appendix B.5.

Questions for the Employed

Potential (Here: Actual) Earnings (Q434) What was your labor income [salary] in the last month?
If you had special payments, e.g., vacation pay or retroactive payments, please do not include such payments in your calculations. By contrast, do include overtime pay. In case you are self-employed: Please estimate your monthly profit before and after taxes.
Please report if possible both:
- the gross salary, that is, the wages or the salary before deducting taxes and social insurance
- the net salary, that is, the wages or the salary after deducting taxes and contributions to pension, unemployment and health insurances.

[We use the "pnett" variable, i.e., the net salary.]

Baseline: Reservation Earnings Please imagine the following hypothetical scenario:
Your employer cuts, for instance because of a situation of reduced demand, your salary for one month.
After that month, your salary will return to its normal level.
How high would the net salary have to be for that month, for you to still go to work at that reduced salary, rather than preferring to take unpaid vacation?

Variant: Reservation Earnings [Identical to baseline question except for the last sentence:] How high would the net salary have to be for that month, for you to still go to work at that reduced salary, rather than preferring to interrupt your job, e.g., by taking vacation days or by giving up the job, e.g., by quitting?

Calculation of Reservation Raise We calculate the reservation raise as the ratio of the reservation earnings over the actual earnings.

Questions for the Unemployed

Potential Earnings You have responded that you currently do not have a job, but are open to accepting a job.
Please now imagine a job that would be realistic for you and that appropriate for your qualifications.
How high would your monthly net salary be, if you had such a position to accept?

Baseline: Reservation Earnings Now please imagine the following hypothetical scenario:
You have found this job and accepted it.
In the course of the job, your employer cuts, for instance because of a situation of reduced

demand, your salary for one month.
After that month, your salary will return to its normal level (that is, [the number the respondent gave above as the salary for this job]).
How high would the net salary have to be for that month for you to still go to work at that reduced salary, rather than preferring to take unpaid vacation?

Variant: Question Giving Reservation Earnings [Identical to baseline question except for the last sentence:]
How high would the net salary have to be for that month for you to still go to work at that reduced salary, rather than preferring to interrupt your job, e.g., by taking vacation days or by giving up the job, e.g., by quitting?

Calculation of Reservation Raise We calculate the reservation raise as the ratio of the reservation earnings over the anticipated potential earnings.

Questions for the Out of the Labor Force

Potential Earnings You have responded [in a previous labor force status question] that you are currently not employed and are also not looking for a job.
Please now nevertheless imagine a job that could be realistic for you and would be appropriate for your qualifications.
Additionally, imagine which salary would be realistic for such a job.
What would you estimate as your monthly net salary for such a job?

Reservation Earnings Now please imagine the following hypothetical scenario:
Right now, we know that you would likely not accept this job.
However, please imagine now that the employer would, for this job, guarantee a one-time special payment as a sign-up bonus at the end of the first month.
Following the first month, the salary falls back to the normal level (that is, [the number the respondent gave above as the salary in this job]).
How high would this one-time special payment need to be in order for you to accept this job and work for at least the full first month?

Calculation of Reservation Raise We calculate the reservation raise as the ratio of the reservation earnings (which are the sum of the estimated potential earnings plus the reservation level of the sign-up bonus) over the anticipated potential earnings.

B.3 Tables

Table B.1: Summary Statistics of Survey and Sub-samples of Survey

	GSOEP (German)		NORC (U.S.)	
	Survey	Sample	Survey	Sample
Employed[1]	58.0%	58.5%	62.1%	60.7%
Unemployed	8.8%	8.4%	5.2%	2.4%
Out of Labor Force	32.6%	33.1%	31.9%	36.9%
Age (Mean)	52.1	51.1	47.4	48.1
Age (Median)	53	52	47	48
Age (Std. Dev)	18.3	17.5	17.8	17.6
Pctg. Female	51.9%	50.7%	51.6%	50.8%
Partnered	62.9%	65.2%	57.8%	59.6%
H.S. Diploma	19.7%	18.1%	28.6%	28.9%
Some College	N/A	N/A	28.2%	29.3%
Vocational	58.1%	59.2%	N/A	N/A
College or Higher	22.2%	22.8%	32.3%	33.4%
Annual Household Income[2]	37,554.04	37,930.80	62,181.58	62,951.76
Number of Respondents	3,346	2,431	2,071	1,679

Note: The table reports summary statistics (means and standard deviations) for the NORC (U.S.) survey and the GSOEP (German) survey. “Survey” refers to the summary statistics of all respondents in the survey that were asked our questions. “Sample” refers to the subset of respondents for which we have nonmissing reservation raise statistics. All statistics use survey weights (with the exception of the “Respondents” row), with the “Sample” column reweighted to replicate overall proportions of labor force groups in the whole survey (for GSOEP, in turn mirroring the OECD numbers for 2019) or 2019 BLS labor force status statistics (in NORC).

[1] For GSOEP, the weights on the three labor force groups in the “Survey” column do not add up to 100% because a small number of respondents do not cleanly fall into any of the labor force statuses.

[2] For GSOEP, household income figure is *net* household income, reported monthly in Euros and multiplied by 12 to achieve annual net household income. For NORC, the household income figure is gross, and reported in bins. We calculate the mean household income using the bottom of these bins; (for example, a respondent in the $50,000 to $60,000 bin is treated as having $50,000 in gross annual income). The average household income in the table is therefore likely an underestimate.

Table B.2: Mass of Marginal Agents and Local Arc Elasticities: Reservation Raise Distribution Around 1.00 for Surveys and Calibrated Models

	Increase in Raise ($1+\Xi' > 1$)		Decrease in Raise ($1+\Xi' < 1$)	
	$\frac{dEmp}{Pop} \times 100$	Elasticity	$\frac{dEmp}{Pop} \times 100$	Elasticity
Survey or Model	Panel A: **Raise Interval: 0.01**			
Data: U.S. (NORC)	2.26	3.72	4.61	7.59
Data: Germany (GSOEP)	1.70	2.86	5.56	9.38
Hansen	100.0	∞	100.0	∞
Constant: 0.32	0.20	0.32	0.20	0.32
Constant: 2.5	1.53	2.52	1.51	2.48
Heterogeneous Agent	0.11	0.18	0.43	0.72
Rogerson-Wallenius	1.73	2.84	1.76	2.90
Panel B: **Raise Interval: 0.03**				
Data: U.S. (NORC)	2.31	1.27	5.55	3.05
Data: Germany (GSOEP)	1.90	1.07	5.80	3.26
Hansen	100.0	∞	100.0	∞
Constant: 0.32	0.58	0.32	0.59	0.32
Constant: 2.5	4.66	2.56	4.45	2.44
Heterogeneous Agent	0.42	0.23	1.04	0.58
Rogerson-Wallenius	5.01	2.79	5.40	2.96
Panel C: **Raise Interval: 0.05**				
Data: U.S. (NORC)	4.11	1.35	14.36	4.73
Data: Germany (GSOEP)	2.13	0.72	6.00	2.02
Hansen	100.0	∞	100.0	∞
Constant: 0.32	0.96	0.32	0.99	0.33
Constant: 2.5	7.87	2.59	7.31	2.41
Heterogeneous Agent	0.93	0.31	1.52	0.51
Rogerson-Wallenius	8.30	2.74	9.18	3.02
Panel D: **Raise Interval: 0.10**				
Data: U.S. (NORC)	5.81	0.96	22.35	3.68
Data: Germany (GSOEP)	2.43	0.41	7.22	1.22
Hansen	100.0	∞	100.0	∞
Constant: 0.32	1.89	0.31	2.02	0.33
Constant: 2.5	16.33	2.69	14.06	2.32
Heterogeneous Agent	1.39	0.23	2.70	0.45
Rogerson-Wallenius	15.85	2.61	19.37	3.19

Note: The table presents shares and arc elasticities of the reservation raise distributions for the data (U.S. (NORC) as well as German (GSOEP) discussed in Section 2.3), as well as for the models presented in the model meta-analysis in Section 2.4. The associated aggregate labor supply curves and arc elasticities are plotted in Figure B.4. The left columns present the share of marginal agents (those with reservation raise levels around one) for various intervals around one, above one, ("+", e.g., 1.00 and 1.01), and below one ("$-$", e.g., 0.99 and 1.00). The right columns present the implied local arc elasticities for each interval.

B.4 Figures

Figure B.1: Empirical Distribution of Reservation Raises

(a) U.S. (NORC)

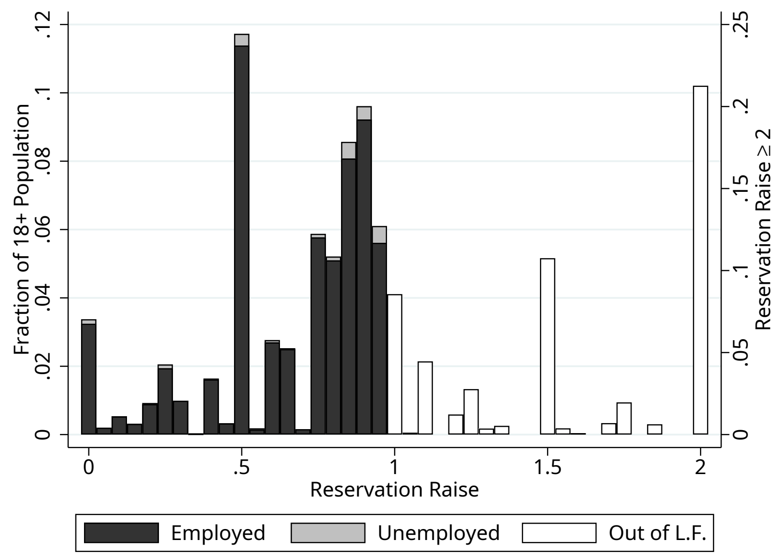

(b) Germany (GSOEP)

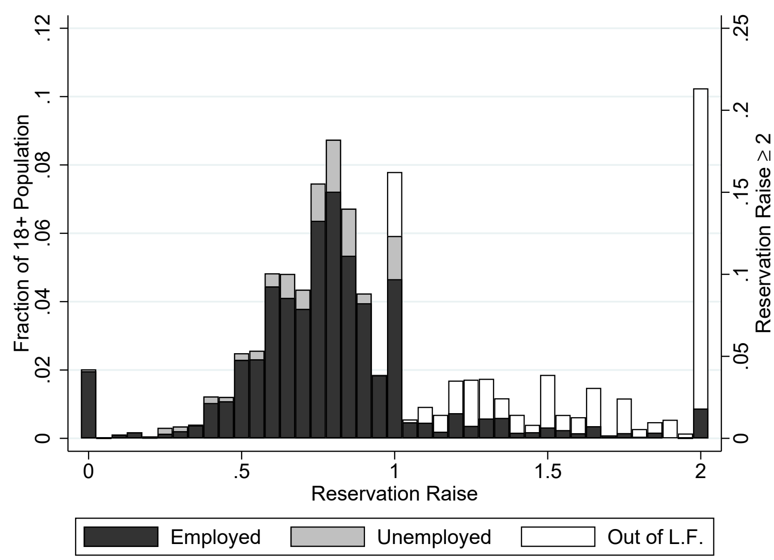

Note: The figure plots histograms of the empirical distribution of reservation raises in a representative sample of the U.S. population (NORC) in Panel (a), as well as of the German population (GSOEP) in Panel (b). Both histograms separate out the observations by their labor force status. For visual clarity, the histograms bunch raises above 2.0 into the 2.0 group, and report this share on the secondary y-axis.

Figure B.2: Empirical Local Distribution of Reservation Raises

(a) Full Cumulative Distribution Functions = Global Aggregate Labor Supply Curves

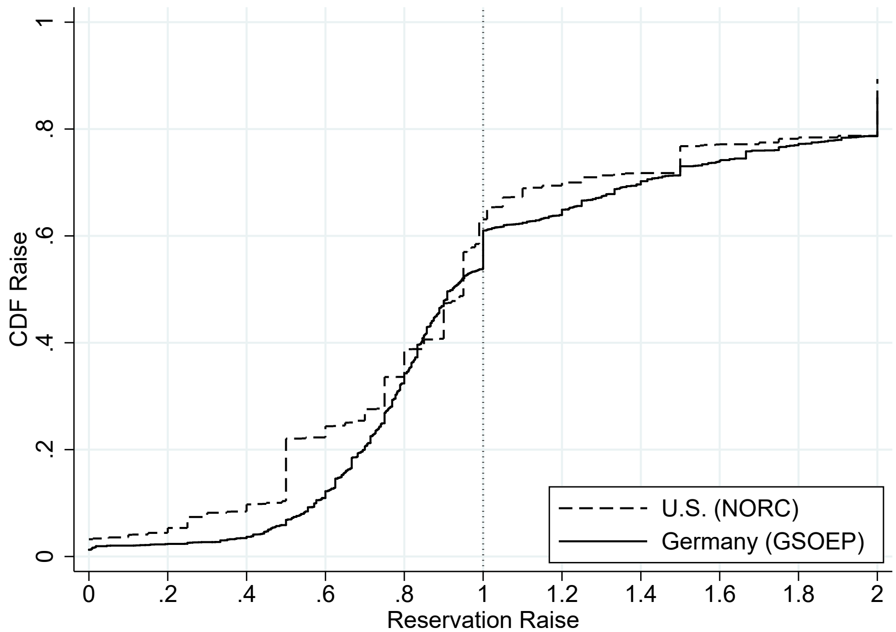

(b) Zoomed-in Version: Local Behavior

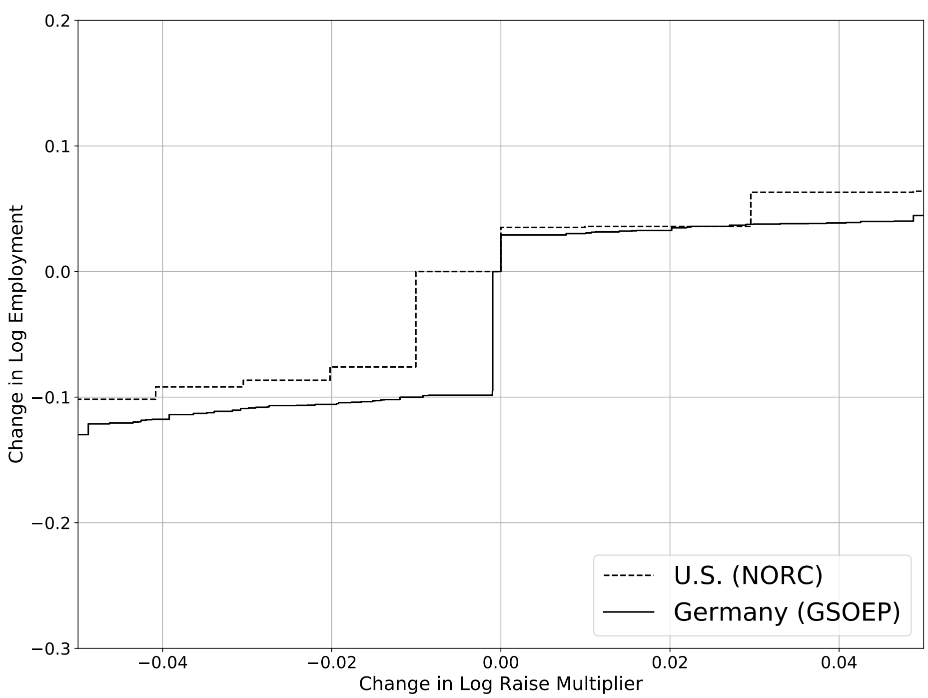

Figure Note on next page.

Note for Figure B.2:

The figure plots cumulative distribution functions of the empirical distribution of reservation raises in a representative sample of the U.S. population (NORC) by the dashed line, as well as of the German population (GSOEP) by the solid line. Panel (a) does so for the full CDF. This CDF is (when evaluated at the cutoff set to the prevailing aggregate raise) the aggregate labor supply curve at the extensive margin. For visual clarity, the CDFs bunch raises above 2.0 into the 2.0 group. Panel (b) takes logs on both sides and zooms into a 0.05 range of the aggregate prevailing raise.

Figure B.3: Empirical Arc Elasticities

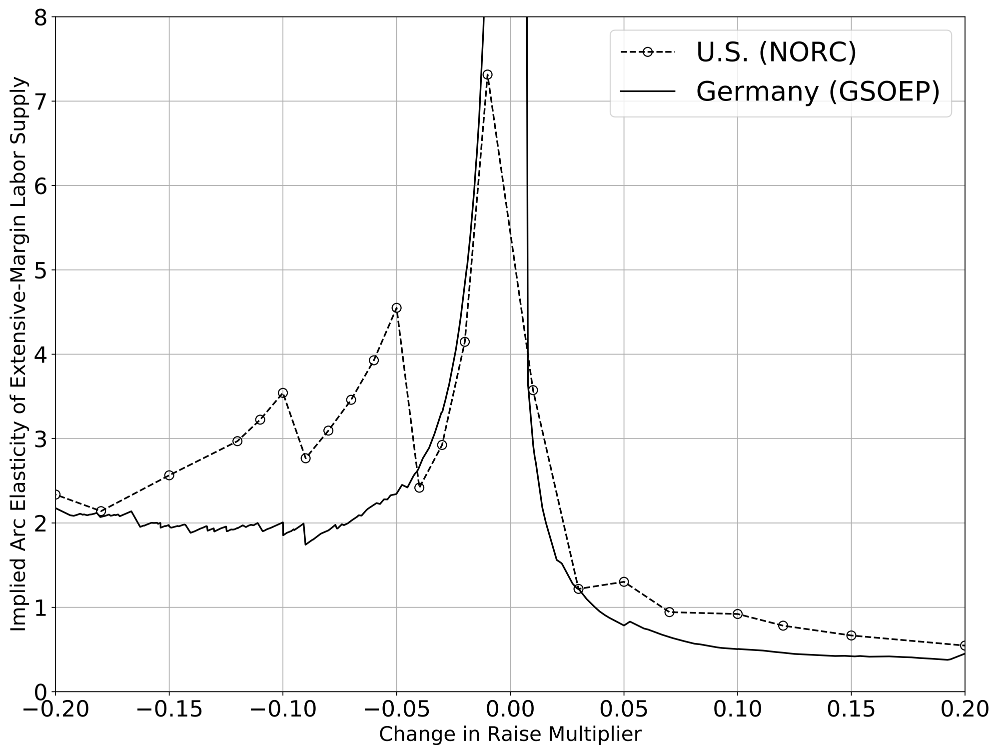

Note: The figure compares the arc elasticities of aggregate desired labor supply in a representative sample of the U.S. population (NORC) (in hollow circles denoting increments with observations) by a dashed line, as well as of the German population (GSOEP) by a solid line. Since the GSOEP permits, and features, mass points of exactly marginal respondents with a reservation raise of exactly one, the elasticity is infinite at this point (so we cap the y-axis), and drops to a finite level at the first non-unit observation. We linearly interpolate the elasticities between points with empirical observations.

Figure B.4: Comparing Models and the Data

(a) Aggregate Labor Supply Curves (CDFs of Reservation Raises)

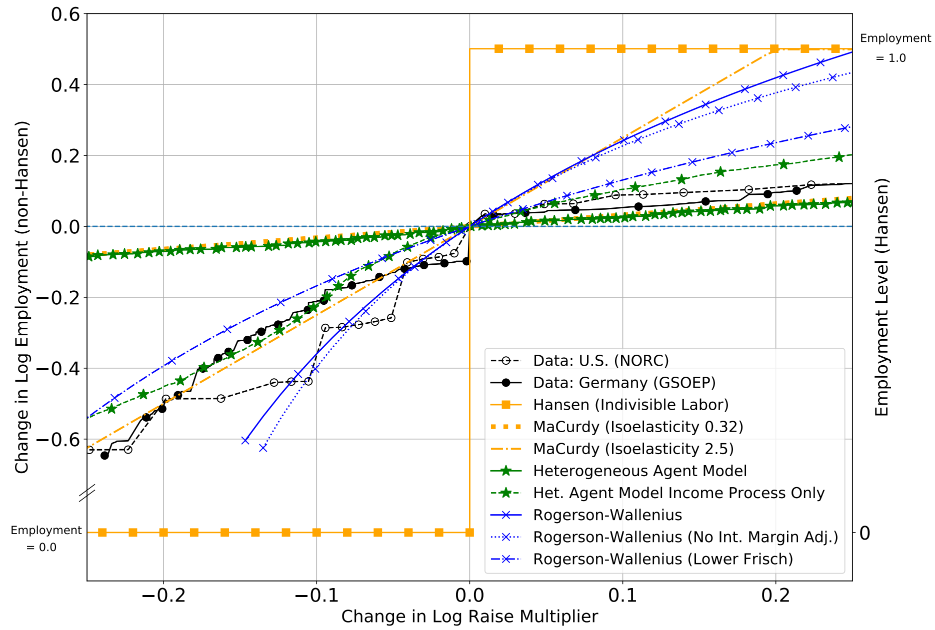

(b) Arc Elasticities

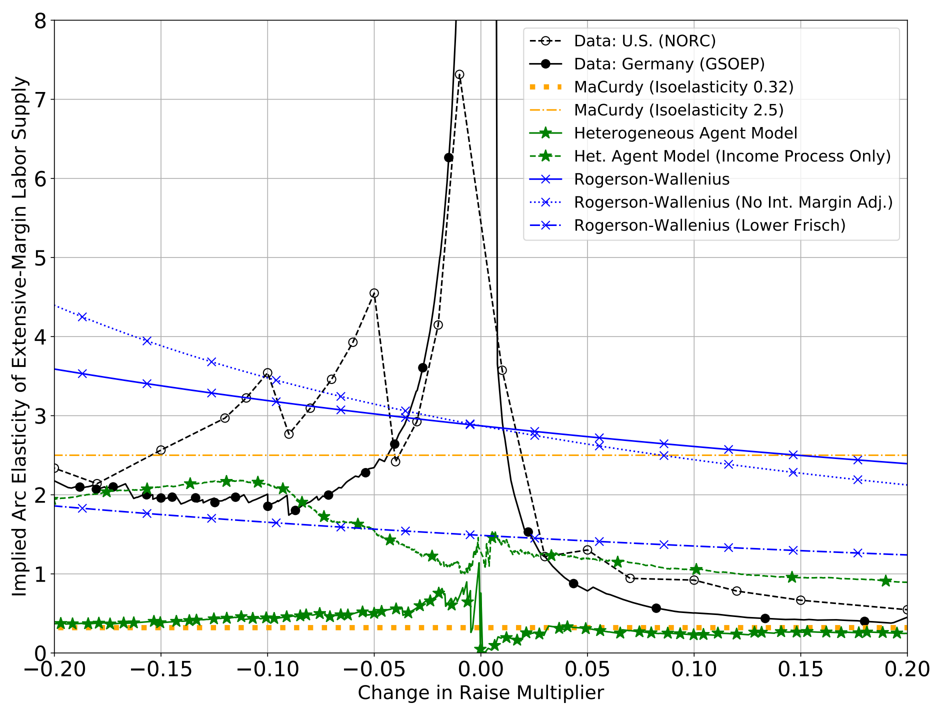

Figure note on next page.

Note for Figure B.4:

The figure compares the empirical and model-implied aggregate labor supply curves at the extensive margin building on our reservation raise approach. Panel (a) plots the labor supply curve as follows: it plots the deviation in the log (desired) employment rate (y-axis) against deviations in the aggregate prevailing raise (x-axis). Hence, it corresponds to the CDFs of the reservation raise distribution, logging both axes, and plotting deviations from baseline levels (a unit reservation raise and the employment levels harmonized across models by calibration). The Hansen indivisible labor model is plotted on a secondary y-axis denoting the employment level (rather than in log deviations). Panel (b) plots arc elasticities of the employment rate with respect to deviations of the aggregate prevailing raise $1+\Xi$, for range of deviations of the raise around the baseline level (the x-axis). The arc elasticities are calculated as $\frac{d\text{Emp}}{\text{Emp}}/\frac{d(1+\Xi)}{1+\Xi}$, from the baseline employment level (harmonized across models by calibration) and from a corresponding baseline net of raise rate $1+\Xi$ normalized to 1.0.

B.5 German Socio-Economic Panel Survey Questions: German Original Text

The questions below are the German original text; the English translations are in the in Appendix B.2.

Questions for the Employed

Potential (Actual) Earnings (Q434w) Wie hoch war Ihr Arbeitsverdienst im letzten Monat? // Wenn Sie im letzten Monat Sonderzahlungen hatten, z.B. Urlaubsgeld oder Nachzahlungen, rechnen Sie diese bitte nicht mit. Entgelt für überstunden rechnen Sie dagegen mit. // Falls Sie selbständig sind: Bitte schätzen Sie Ihren monatlichen Gewinn vor und nach Steuern. // Bitte geben Sie nach Möglichkeit beides an: // -den Bruttoverdienst, das heißt Lohn oder Gehalt vor Abzug der Steuern und Sozialversicherung // -und den Nettoverdienst, das heißt den Betrag nach Abzug von Steuern und Beiträgen zur Renten-, Arbeitslosen- und Krankenversicherung. //

[We use “pnett”, i.e., the net salary.]

Baseline: Question Giving Reservation Earnings Stellen Sie sich bitte nun folgendes hypothetisches Szenario vor:
Ihr Arbeitgeber kürzt beispielsweise aufgrund einer verringerten Auftragslage für einen Monat lang Ihr Gehalt.
Im Anschluss an diesen Monat geht Ihr Gehalt wieder auf sein normales Niveau zurück.
Wie hoch müsste der Nettoverdienst in diesem Monat mindestens sein, damit Sie auch zu diesem gekürzten Gehalt weiterhin arbeiten gehen statt lieber unbezahlten Urlaub zu nehmen?

Variant: Question Giving Reservation Earnings [Identical to baseline question except for the last sentence:]
Wie hoch müsste der Nettoverdienst in diesem Monat mindestens sein, damit Sie auch zum gekürzten Gehalt weiterhin arbeiten gehen statt lieber Ihre Stelle zu unterbrechen, z.B. durch Urlaub nehmen, oder aufzugeben, z.B. durch Kündigen?

Question for the Unemployed

Question Giving Reservation Earnings Sie haben angegeben, dass Sie derzeit keine Arbeitsstelle haben, sich aber vorstellen können, eine Stelle anzutreten.
Stellen Sie sich bitte nun eine Arbeitsstelle vor, die für Sie realistisch wäre und Ihren Qualifikationen entspräche.
Was meinen Sie: Wie hoch wäre Ihr monatliches Nettogehalt, wenn Sie eine solche Stelle annehmen würden?

Baseline: Question Giving Reservation Earnings Stellen Sie sich bitte nun folgendes hypothetisches Szenario vor:
Sie haben die Stelle gefunden und angenommen.

Im Verlauf des Arbeitsverhältnisses kürzt Ihr Arbeitgeber beispielsweise aufgrund einer verringerten Auftragslage für einen Monat lang Ihr Gehalt.
Im Anschluss an diesen Monat geht Ihr Gehalt wieder auf sein normales Niveau zurück.
Wie hoch müsste der Nettoverdienst in diesem Monat mindestens sein, damit Sie auch zum gekürzten Gehalt weiterhin arbeiten gehen statt lieber unbezahlten Urlaub zu nehmen?

Variant: Question Giving Reservation Earnings Stellen Sie sich bitte nun folgendes hypothetisches Szenario vor:
Sie haben die Stelle gefunden und angenommen.
Im Verlauf des Arbeitsverhältnisses kürzt Ihr Arbeitgeber beispielsweise aufgrund einer verringerten Auftragslage für einen Monat lang Ihr Gehalt.
Im Anschluss an diesen Monat geht Ihr Gehalt wieder auf sein normales Niveau zurück.
Wie hoch müsste der Nettoverdienst in diesem Monat mindestens sein, damit Sie auch zum gekürzten Gehalt weiterhin arbeiten gehen statt lieber Ihre Stelle zu unterbrechen, z.B. durch Urlaub nehmen, oder aufzugeben, z.B. durch Kündigen?

Question for the Out of the Labor Force

Question Giving Potential Earnings Sie haben angegeben, dass Sie derzeit nicht berufstätig sind und auch keine Arbeitsstelle suchen.
Stellen Sie sich nun bitte trotzdem eine Stelle vor, die für Sie realistisch sein könnte und Ihren Qualifikationen entspräche.
Stellen Sie sich auch vor, welches Gehalt für eine solche Stelle realistisch wäre.
Was meinen Sie: Wie wäre Ihr monatliches Nettogehalt für eine solche Stelle?

Question Giving Potential Earnings Stellen Sie sich bitte nun folgendes hypothetisches Szenario vor:
Derzeit würden Sie diese Stelle ja wahrscheinlich nicht annehmen. Stellen Sie sich nun aber vor, dass der Arbeitgeber für diese Stelle am Ende des ersten Monats eine einmalige Sonderzahlung als Einstiegsbonus garantiert. Im Anschluss an den ersten Monat fällt das Gehalt wieder auf das normale Niveau zurück.
Wie hoch müsste diese einmalige Sonderzahlung sein, damit Sie diese Arbeitsstelle annehmen und zumindest für den ganzen ersten Monat lang arbeiten würden?

B.6 Empirical Appendix

Calculation of Arc Elasticities

In the main text, we construct the arc elasticity for non-infinitesimal changes in the aggregate raise as:

$$\epsilon_{E_t,(1+\Xi_t)\to(1+\Xi'_t)} = \frac{F_t(1+\Xi'_t) - F_t\left(1+\Xi_t\right)}{F_t\left(1+\Xi_t\right)} \Big/ \frac{(1+\Xi'_t)-(1+\Xi_t)}{1+\Xi_t}. \tag{B.6.1}$$

Our empirical analog of the arc elasticity (normalized so that the prevailing aggregate raise is $1-\Xi_t = 1$) is:

$$\hat{\epsilon}_{E_t,(1)\to(1+\Xi'_t)} = \frac{\tilde{F}_t\left(1+\Xi'_t\right) - \tilde{F}_t(1)}{\tilde{F}_t(1)} \Big/ \Xi'_t. \tag{B.6.2}$$

We now discuss how we construct the the CDF at which we evaluate the counterfactual employment level $\tilde{F}_t\left(1+\Xi'_t\right)$, and then discuss the baseline employment level $\tilde{F}_t(1)$.

Breaking Ties We have two tie-breaking choices to make at points of indifference. Formally, our tie-breaking rule depends on the direction of the shift (whether $1+\Xi'_t > 1$, < 1, or $= 1$) as follows:

$$\tilde{F}_t(1+\Xi'_t) = \begin{cases} \sum_{i=1}^{n} \omega_i \mathbb{1}_{(1+\xi_{it}<1+\Xi'_t)} & \text{if } \ 1+\Xi'_t < 1 \\ \sum_{i=1}^{n} \omega_i \mathbb{1}_{(1+\xi_{it}\le 1+\Xi'_t)} & \text{if } \ 1+\Xi'_t > 1 \\ \text{given by labor force status (see below)} & \text{if } \ 1+\Xi'_t = 1, \end{cases} \tag{B.6.3}$$

where ω_i is the survey weight on respondent i (in NORC, the sampling weights and the labor force weights as described in the main text; in GSOEP, the labor force weights as described in the main text).

First, we determine how to allocate workers exactly indifferent at the original prevailing aggregate raise, i.e., for whom $1+\xi^*_{it} = 1$. The empirical analog of desired employment at the prevailing raise ($\tilde{F}_t(1)$)—and therefore the mass of individuals crowded into employment or nonemployment depending on the direction of the pertubation—depends on the survey used. For the U.S. (NORC) survey, $\tilde{F}_t(1)$ is simply the (weighted) fraction of respondents that choose a reservation raise less than than 1 (by survey construction, all employed and unemployed respondents are restricted to reporting a reservation raise lower than 1, so there is no empirical difference between less than or less than equal with this survey). For the German (GSOEP) survey, survey respondents are not prevented from reporting $1+\xi^*_{it} = 1$; that is, they can report being exactly indifferent to employment or non-employment. Economically speaking, a respondent with $1+\xi^*_{it} = 1$ is exactly indifferent, and so the exact level of desired aggregate labor supply is undefined if there is a mass of respondents with $1+\xi^*_{it} = 1$. To break this tie, we use labor force status: we calculate $\tilde{F}_t(1)$ as including the exactly marginal ($1+\xi^*_{it} = 1$) employed or unemployed workers, but *not* the exactly marginal out-of-the-labor force respondents. This implies that for upward perturbations $1 + \Xi'_t > 1$, there is a mass of out of the labor force (those with $1 + \xi^*_{it} = 1$ being crowded into employment (plus the nearly marginal other respondents); conversely, the exactly marginal employed and unemployed respondents will be crowded into nonemployment for any downward perturbation $1 + \Xi'_t < 1$ (plus the nearly marginal other respondents).

Second, we need to break ties for how to treat respondents that are exactly marginal at the *new* aggregate prevailing wedge $1 + \xi^*_{it} = 1 + \Xi'_t$, to compute the counterfactual employment level in the numerator $F_t(1 + \Xi'_t)$ (an issue largely present in NORC due to the percentage point increments). In words, the above formula clarifies that when calculating downwards elasticities, we count individuals who report indifference at the perturbation, i.e., for whom $1 + \xi^*_{it} = 1 + \Xi'_t$, as being induced into nonemployment out of employment; when calculating upward elasticities, we count such individuals as being induced into employment out of nonemployment.

Winsorizing the (Un-)Employed In Table B.3, we report the proportions of exactly and nearly marginal respondents by labor force status for each survey. Across all three each labor force groups, there is a substantial proportion of respondents who report being exactly or nearly marginal, especially for the employed. In NORC, by design these respondents are bunched at the lowest increment below 1.0 (0.99) for those in the labor force, and above 1.00 (1.01) for those out of the labor force. In GSOEP, the mass of exactly marginal individuals is permitted to occur at exactly 1.00 (when workers' reservation earnings equal the actual earnings). For the employed, the nearly (but not exactly) marginal respondents are predominantly below 1.0, consistent with our decision to allocate these workers to desiring employment.

Figure B.5 Panel (a) plots the alternative CDF and Panel (b) shows the corresponding arc elasticities that would arise if we winsorized the employed and unemployed with reservation raises above 1 to exactly 1 (and for the arc elasticities treat them following the tie-breaking rule described above). Of course, this procedure dramatically raises local elasticities downward, and somewhat attenuates upward elasticities. Of course, some of those answers may reflect measurement error (e.g., the employed may report their actual earnings noisily or have another number in mind when then giving the reservation earnings), so that some of these observations need not be exactly marginal.

Table B.3: Shares of Exactly and Nearly Marginal Respondents

L.F. Status	Reservation Raises				
	< 0.99	$[0.99, 1.00)$	1.00	$(1.00, 1.01]$	> 1.01
Panel A: NORC					
Empl.	56.26%	4.44%	0.00%	0.00%	0.00%
Unempl.	2.23%	0.17%	0.00%	0.00%	0.00%
OOLF	0.00%	0.00%	0.00%	2.26%	34.64%
Panel A: GSOEP					
Empl.	46.52%	0.28%	4.10%	0.05%	7.56%
Unempl.	7.15%	0.00%	1.26%	0.00%	0.00%
OOLF	0.00%	0.00%	1.70%	0.00%	31.37%

Note: The table presents respondents that are exactly and nearly marginal by employment status. The reported percentages are the (weighted) percentages of the whole sample.

Figure B.5: Rationed Labor Supply of the Employed and Unemployed: Winsorizing Reservation Raise Above 1 (GSOEP)

(a) CDFs of Reservation Raises

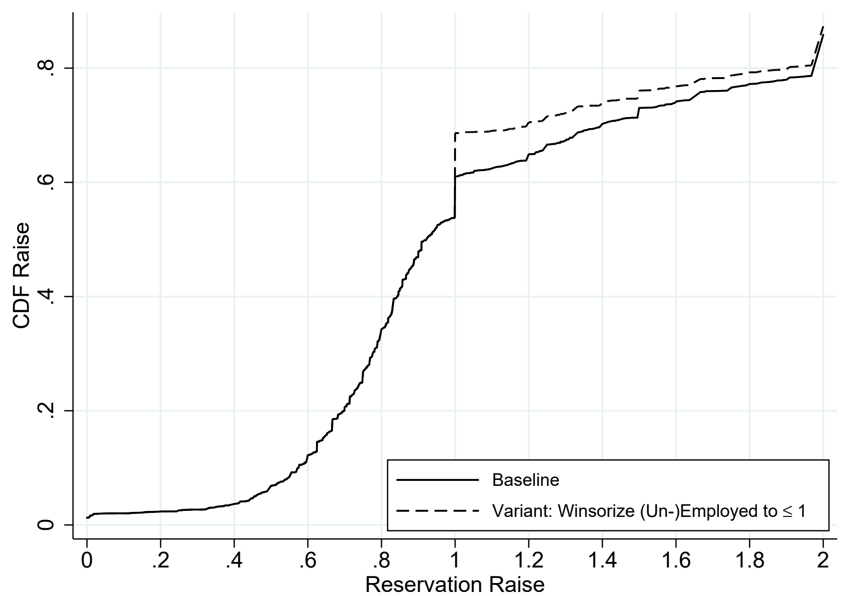

(b) Arc Elasticities

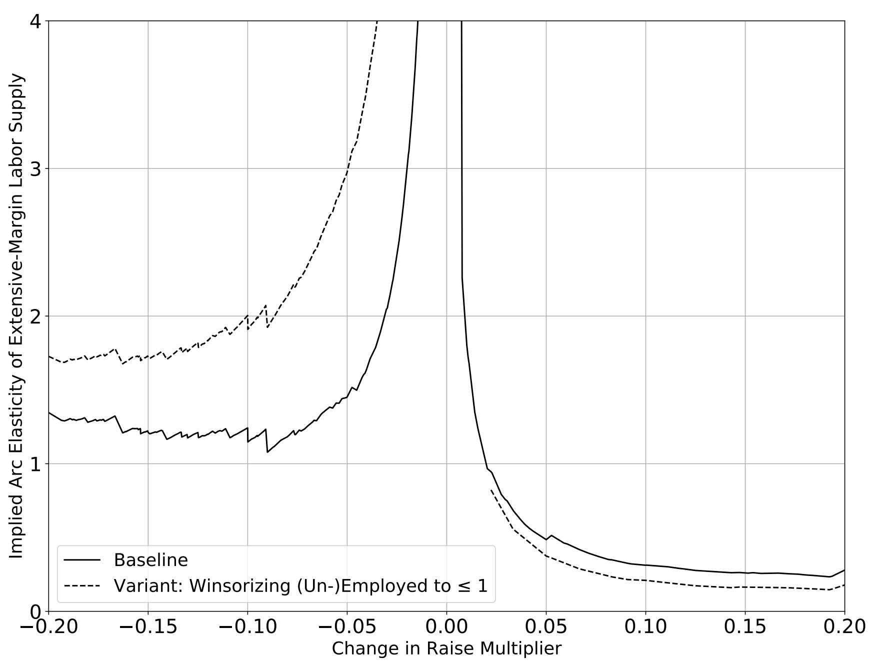

Note: The figure compares the GSOEP reservation raise distributions (CDF in Panel (a), arc elasticities in Panel (b)) under our baseline treatment of the employed versus treating those who are employed (or unemployed) and report a reservation raise greater than 1 as exactly marginal, i.e., we winsorize those observations and reassign them a value of exactly 1.

Missing Observations

Table B.4: Missing Observations

Panel A: U.S. (NORC)

L.F. Status	Missing Reservation Raise	Total (Incl. Nonmissing)
Employed	247	1284
Unemployed	10	83
OOLF	122	691
Any L.F.	392	2,058
Missing L.F. Status		13

Panel B: Germany (GSOEP)

L.F. Status	Missing Res. Raise	Missing Salary	Missing Res. Res. Salary	Total, incl. non-missing (Fraction of Total)
Employed (Baseline)	194 (22.1%)	58 (6.3%)	148 (18.7%)	933 (29.4%)
Employed (Variant)	174 (19.2%)	60 (6.5%)	148 (16.2%)	926 (28.7%)
Unemployed (Baseline)	39 (22.2%)	31 (17.4%)	36 (18.4%)	144 (4.6%)
Unemployed (Variant)	39 (25.0%)	32 (20.2%)	39 (24.9%)	129 (4.2%)
OOLF	618 (47.0%)	482 (36.1%)	605 (46.1%)	1,360 (32.7%)
Any L.F.	1,062 (29.6%)	662 (27.2%)	991 (17.3%)	3,493 (99.5%)
Missing L.F. status				17 (0.46%)
Missing Weight				164

Note: The table presents the fraction of respondents with missing information on reservation raises (and for GSOEP, separately for potential and reservation earnings), by employment status.

Robustness and Alternative Specifications (GSOEP)

Figure B.6: The Role of Adjustment Frictions: The Reservation Raise Distribution, Baseline vs. Variant (GSOEP)

(a) CDFs of Reservation Raises

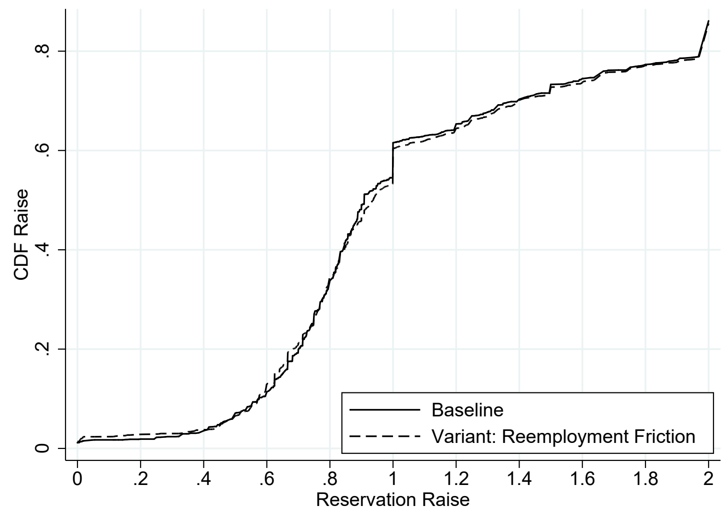

(b) Arc Elasticities

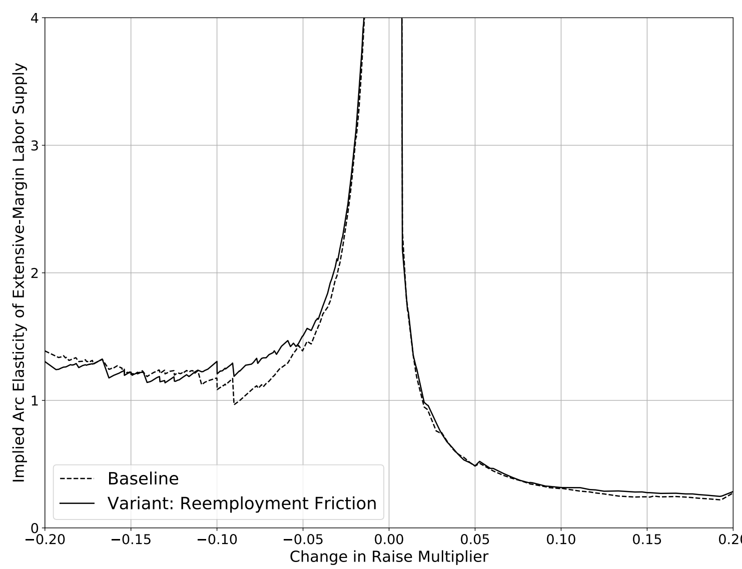

Note: The figure compares the GSOEP reservation raise distributions (CDF in Panel (a), arc elasticities in Panel (b)) under the baseline question (which evokes frictionless labor supply adjustment during the nonemployment month, akin to a vacation) and the variant question (which explicitly eludes to the possibility that a quit from the job may be necessary to achieve that month in nonemployment). The scenarios apply to the employed and unemployed. The resulting graphs reweight the employed and unemployed to match their shares in the full sample.

B.7 Theoretical Appendix for Model Meta-Analysis

We detail each model's derivations, and then cover computational details.

Detailed Derivation and Discussion

We now present a detailed model-by-model meta-analysis applying the reservation-raise approach as a unifying bridge between structurally different labor supply blocks. The parameters for our calibrated models in this meta-analysis are in Appendix Table B.5. Appendix Figure B.7 plots additional model-specific reservation raise histograms and supplementary items.

Representative Household: Full Insurance and "Command" Labor Supply

A common specification appeals to a large representative household, comprised of a unit mass of individual members, which we explicitly index by $i \in [0,1]$. Micro utility $u_i(c_{it}) - e_{it}v_{it}$ is separable, where $e_{it} \in \{0,1\}$ is an employment indicator. Potential earnings are y_{it}. There is potentially some uncertainty over the path of wages and interest rates. The large household has a *pooled* budget constraint and assigns consumption levels and employment statuses to its individual members:[†]

$$\max_{\{c_{it},e_{it}\}_i, A_t} \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \int_0^1 \left[u_i(c_{is}) - e_{is}v_{is}\right] di \tag{B.7.1}$$

$$\text{s.t. } A_s + \int_0^1 c_{is} di \leq A_{s-1}(1 + r_{s-1}) + \int_0^1 (1 + \Xi_s) y_{is} e_{is} di + T_s \ \ \forall s \geq t. \tag{B.7.2}$$

Full (cross-sectional) insurance implies that the marginal utility of consumption is optimally set homogeneous across households, equal to the multiplier on the pooled budget constraint,

$$\overline{\lambda}_t = \frac{\partial u_i(c_{it})}{\partial c_{it}} \ \ \forall\, i, \tag{B.7.3}$$

[†]We take a perspective, as, e.g., Galí (2011), that the household head directly assigns allocations. Hansen (1985) and Rogerson (1988) present incentive-compatible lotteries. The Hansen (1985) set-up is equivalent to a representative household with utility function $U(c_t, E_t) = \log(c_t) - \overline{v}E_t$, with intratemporal first-order condition $\overline{\lambda}_t \overline{w}_t = \overline{v}$.

Figure B.7: Further Details on Model Reservation Raises Distributions

(a) Hansen, Histogram

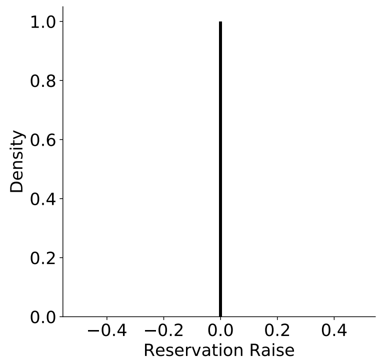

(b) MaCurdy: 0.32 Isoelasticity

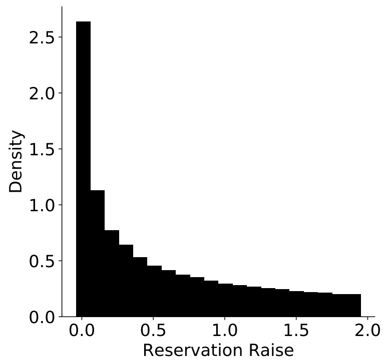

(c) Het-Agent: 2 Wage States

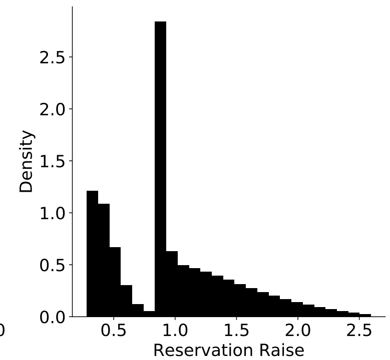

(d) Het-Agent: 2 Wage Statess

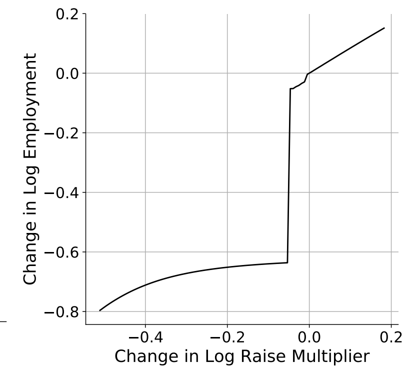

(e) Het-Agent 33 Wage States

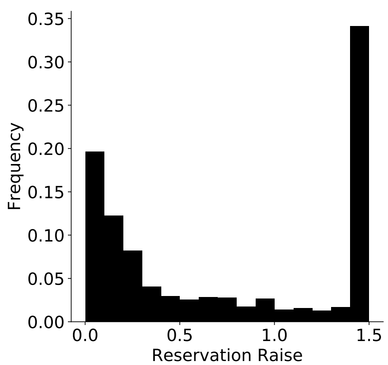

(f) Het-Agent Model: Isolating 3/33 Wage States

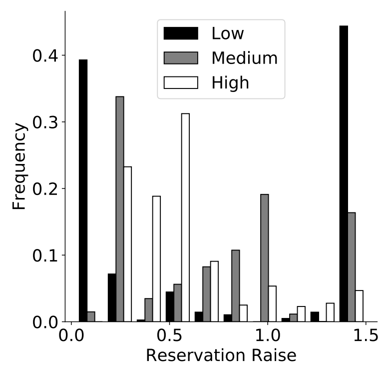

(g) Het-Agent Model (33 Wage States): Het. vs. Homog. λ

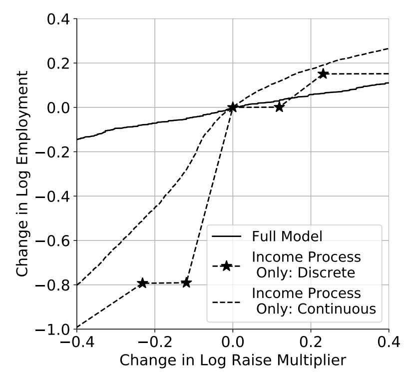

(h) Rogerson-Wallenius

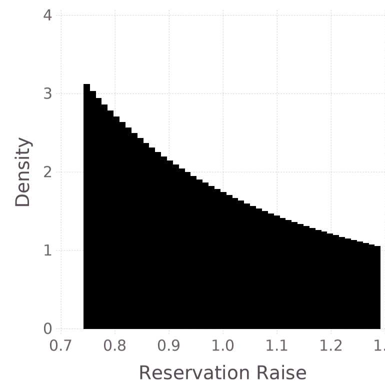

Figure Note on next page.

Note for Figure B.7:

The figure plots additional simulated data from the models reviewed in the meta-analysis in Section 2.4. Panel (a) plots the histogram of the Hansen (1985) model's reservation raises. Panel (b) plots the histogram of the reservation raises that would emerge in an isoelastic representative household setting with an elasticity of 0.32. Panel (c) plots the reservation raise histogram from the two potential-earnings states heterogeneous agent model; Panel (d) plots the associated aggregate labor supply curve. Panel (e) plots the histogram of reservation raises in the 33 potential-earnings states heterogeneous agent model following the realistic earnings process. Panel (f) provides reservation raises for three earnings states; the low (1876.61), medium (24,489.68), and high (117,080.23) potential-earnings levels. The densities are normalized so that the total density by earnings level sums to one; however, there is 0.395, 0.164, and 0.033 of density for the low, medium, and high earnings levels that have reservation raises above 1.5 (which we censor in this histogram). Panel (g) plots the original aggregate labor supply curve for the 33-state heterogeneous agent economy and heterogeneous borrowing constraint multiplier λ, but adds curves for two full-insurance models by setting the borrowing constraint multiplier λ homogeneous, for the original coarse discrete earnings process as well as a richer continuous-state version. It thereby highlights the role of the covariance of potential earnings and the shadow value of income in shaping the inelastic labor supply curve. Panel (h) plots the reservation raise histogram for the calibrated Rogerson and Wallenius (2009) model.

Table B.5: Parameters of Macro Models with an Extensive-Margin of Labor Supply

Parameter	Symbol	Value (by Variant)	
Panel A: Hansen (Indivisible Labor)			
Employment disutility	$\overline{v}$	1.0	
Potential earnings	$\overline{y}$	1.0	
MUC	$\overline{\lambda}$	1.0	
Panel B: MaCurdy (Isolesticity)			
		Low Frisch (0.32)	High Frisch (2.50)
CRRA cons. param.	σ	1.00	"
Potential earnings	$\overline{y}$	1.00	"
Shape parameter of labor disutility distutility	α_v	0.32	2.50
Max. labor disutility	$v_{\max}$	4.759	1.221
Panel C: Heterogeneous Agent Model			
		Toy Model	HANK Earnings Process
Potential-earnings states		$[y_1, y_2] = [0.0797, 0.15]$	33-State Markov process from
Transition probabilities		$[\lambda_{12}, \lambda_{21}] = [0.1, 0.2]$	Kaplan, Moll and Violante (2018)*
CRRA cons. param.	γ	2.0	2.0
Interest rate	r	0.03	0.03
Discount rate	β	0.95	0.97
Labor disutility	$\overline{v}$	3.0	2.083×10^{-5}
UI benefit/nonemp. payoff	b	0.06	0.00
Asset grid: min. assets (& borrowing limit)	$a_{\min}$	-0.02	-1.775
Asset grid: max. assets	$a_{\max}$	0.75	5,000,000
Panel D: Rogerson-Wallenius			
		Baseline	Low-Frisch Variant
Interest rate	r	0.0	"
CRRA cons. param.	σ	1.0	"
Labor disutility shifter	Γ	42.492	40.000
Minimum hours	$\overline{h}$	0.258	0.272
Maximum prod.	$\hat{w}_0$	1.000	1.112
Prod.-age slope	$\hat{w}_1$	0.851	1.320
Intensive-margin Frisch elasticity	γ	0.5	"
Tax rate	τ	26.0%	"

Note: The table presents the parameters for the models with an extensive margin of labor supply presented in the model meta-analysis Section 2.4, generating the calibrated aggregate labor supply curves plotted in Figure B.4. *: We describe the 33-state earnings process in Appendix Section B.7.

which eliminates λ as a source of cross-sectional variation in reservation raises even with heterogeneity in consumption utility function $u_i(\cdot)$. Due to spot jobs, expectations and intertemporal aspects are subsumed in $\overline{\lambda_t}$. Going forward, $\overline{x}_t$ denotes idiosyncratic variables x_{it} that are homogeneous in the cross-section.

First, we define the allocative micro reservation raise in this large-household structure, here rendering the household *head* indifferent between sending that marginal member i to employment rather than nonemployment, where we can index an individual i by her disutility-earnings type vy:

$$1 + \xi_{it}^* = \frac{v_{it}}{\overline{\lambda}_t y_{it}} = 1 + \xi_{vyt}^*. \tag{B.7.4}$$

Second, we trace out the *aggregate* labor supply curve from the distribution of the reservation raises, which in turn subsumes the detailed heterogeneity in wages and labor supply disutilities:

$$E_t(1+\Xi_t) = F_t(1+\Xi_t) = P\left(1+\xi_{it} \leq 1+\Xi\right) = P\left(\frac{v_{it}}{y_{it}\overline{\lambda}_t} \leq 1+\Xi_t\right) = P\left(\frac{v_{it}}{y_{it}} \leq (1+\Xi_t)\overline{\lambda}_t\right) \tag{B.7.5}$$

$$= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \mathbb{1}\left[\frac{v}{y} \leq (1+\Xi_t)\overline{\lambda}_t\right] dG_t(v,y), \tag{B.7.6}$$

where $G_t(v, y)$ is the CDF of the joint distribution of v and y.

Third, the arc elasticities follow the definition in Equation (2.6) and depend on the joint distributions of v and y.

Below we review specific cases of this representative-household class of labor supply model block, to study more concrete curves.

Hansen (1985) The setup nests the model of indivisible labor and homogeneous households by Hansen (1985), where specifically $\overline{y}_t = \tilde{h}w_{it}$ and $v_{it} = \overline{v}\ \forall i$ (which in the original paper is $A\ln(1-h_{it})$), with one exogenous hours option $h_{it} \in \{0, \widetilde{h} > 0\}$.

First, all individuals have the same reservation raise —i.e., all are exactly marginal:

$$1 + \xi_{it}^* = 1 + \overline{\xi}_t^* = \frac{\overline{v}}{\overline{\lambda}_t \overline{y}_t}. \tag{B.7.7}$$

Second, the reservation raise distribution (Appendix Figure B.7 Panel (a)), is degenerate.

Third, the Frisch elasticity is locally infinite at $1 + \Xi_t$. Interior solutions are obtained through λ_t (decreasing marginal utility from consumption).

Heterogeneity Only in Disutility of Labor We now maintain wage homogeneity, but disutility of labor v is distributed between individuals according to CDF $G_t^v(v)$. First, each individual i is now characterized by their type $v(i)$, and the household head maximizes:

$$\max_{\{c_{vt}, e_{vt}\}, A_t} \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \int \left[u(c_{vs}) - e_{vs}v_s\right] g(v)dv \tag{B.7.8}$$

$$\text{s.t. } A_s + \int c_{vs}g(v)dv \leq A_{s-1}(1+r_{s-1}) + (1+\Xi_s)y_s \int e_{vs}g(v)dv + T_s \ \ \forall s \geq t. \tag{B.7.9}$$

First, we define each individual's reservation raise ,characterized by their type $v(i)$:

$$1+\xi_{it}^{*}=\frac{v_{it}}{\overline{y}_t\overline{\lambda}_t}=1+\xi_{vt}^{*}. \tag{B.7.10}$$

Second, aggregate labor supply, i.e., distribution of $1+\xi_{it}^{*}$, will follow directly from $G_t^v(v)$ since consumption and wages are homogeneous. The household head sends off members with $1+\xi_{it}^{*}<1+\Xi_t$ to employment, and all others to nonemployment:

$$E_t(1+\Xi_t)=F_t(1+\Xi_t)=P\left(1+\xi_{it}^{*}\leq 1+\Xi_t\right)=P\left(v_{it}\leq\frac{1+\Xi_t}{\overline{y}_t\overline{\lambda}_t}\right)=G_t^v\left(\frac{1+\Xi_t}{\overline{y}_t\overline{\lambda}_t}\right). \tag{B.7.11}$$

Alternatively, pointwise optimization would lead to a disutility cutoff rule $v_t^*=(1+\Xi_t)\overline{y}_t\overline{\lambda}_t$: $v_{it}\geq v_t^*$ types work, $v_{it}<v_t^*$ types stay at home.

Third, the elasticity is given by $\left[(1+\Xi_t)g_t^v\left(\frac{1+\Xi_t}{\overline{y}_t\overline{\lambda}_t}\right)\right]/\left[1-G_t^v\left(\frac{1+\Xi_t}{\overline{y}_t\overline{\lambda}_t}\right)\right]$.

MaCurdy (1981) Isoelastic Preferences A common representative household setup (with a pooled budget constraint and homogeneous wages) applies the familiar isoelastic *intensive*-margin MaCurdy (1981) preferences to the extensive margin:

$$\frac{C_t^{1-\sigma}}{1-\sigma}-\Psi\frac{E_t^{1+1/\eta}}{1+1/\eta}. \tag{B.7.12}$$

We now *reverse-engineer* a distribution of disutility $G_t^v(v)$ that delivers this labor supply specification. The micro reservation raise is again given by (B.7.10). Suppose v follows a power law distribution $G_t^v(v)=\left(\frac{v}{v_{\max}}\right)^{\alpha_v}$ with shape parameter α_v over support $[0,v_{\max}]$. Then, aggregate employment is (building on Section 2.2, assuming positive nonemployment by all types):

$$E_t(1+\Xi_t)=F_t(1+\Xi_t)=P\left(\frac{v_{it}}{\overline{y}_t\overline{\lambda}_t}\leq 1+\Xi_t\right)=G_t^v\left((1+\Xi_t)\overline{y}_t\overline{\lambda}_t\right)=\left(\frac{(1+\Xi_t)\overline{y}_t\overline{\lambda}_t}{v_{\max}}\right)^{\alpha_v}. \tag{B.7.13}$$

The reservation raise distribution then too is a power law distribution inheriting shape parameter α_v—giving the constant extensive margin Frisch elasticity:

$$\epsilon_{E_t,1+\Xi_t}=\frac{(1+\Xi_t)F_t(1+\Xi_t)}{F_t(1+\Xi_t)}=\frac{(1+\Xi_t)\alpha_v(1+\Xi_t)^{-1}\left(\frac{(1+\Xi_t)\overline{y}_t\overline{\lambda}_t}{v_{\max}}\right)^{\alpha_v}}{\left(\frac{(1+\Xi_t)\overline{y}_t\overline{\lambda}_t}{v_{\max}}\right)^{\alpha_v}}=\alpha_v. \tag{B.7.14}$$

Hence, the representative household can be written with MaCurdy preferences, by simply rearranging the aggregate labor supply curve (B.7.13):

$$v_{\max} E_t^{\frac{1}{\alpha_v}} = (1+\Xi_t)\overline{y}_t\overline{\lambda}_t, \tag{B.7.15}$$

which is the FOC of objective function (B.7.12) for $\eta = \alpha_v$ and $\Psi = v_{\max}$.[†]

In Appendix Figure B.7 Panel (b), we plot the density of reservation raises for a MaCurdy model with potential earnings $\overline{y}$ and marginal utility of consumption $\overline{\lambda}$ are normalized to one, and the Frisch elasticity is 0.32. The maximum micro labor supply disutility is set to $0.607^{-1/0.32}$ for an equilibrium employment rate at 60.7%.

Heterogeneous (Sticky) Wages and Isoelasticity (Galí, 2011) The New Keynesian model presented in Galí (2011) (which also microfound the isoelasticity) additionally features wage heterogeneity. Individuals are a unit square indexed by $(l, n) \in [0,1] \times [0,1]$. l denotes the type of labor, paid wage y_{lt}, which may diverge across types due to wage stickiness. n indexes labor disutility, $n^{1/\eta}$. The household head maximizes:

$$\max_{c_t, \{E_{lt}\}_l} \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t}\Big(\frac{c_s^{1-\sigma}-1}{1-\sigma} - \Psi \int_0^1 \overbrace{\int_0^{E_{ln}} n^{1/\eta}}^{E_{lt}^{1+1/\eta}/(1+1/\eta)} \; dndl\Big) \tag{B.7.17}$$

$$\text{s.t. } A_t + \int_0^1 c_{lt} dl \le A_{t-1}(1+r_{t-1}) + (1+\Xi_t) y_{lt} E_{lt} + T_t \;\; \forall s \ge t, \tag{B.7.18}$$

where the l-specific employment rate is $E_{lt} = \int_0^1 e_{lt} dl$.

We now cast this setting into the reservation raise framework. First, we define the micro reservation raise, characterizing individual i by type nl:

$$1 + \xi^*_{nlt} = \frac{\Psi n^\eta}{y_{lt}\overline{\lambda}_t}. \tag{B.7.19}$$

[†]Alternatively, we can directly derive total disutility of labor $V(E_t)$ from employment rate $E_t \in [0,1]$, where the head optimally sorts the members by their disutility of labor up until $v = \mu(E_t)$, a threshold defined as the disutility of working of the marginal individual for total employment $E_t = G^v(\mu(E_t)) = \left(\frac{\mu(E_t)}{v_{\max}}\right)^{\alpha_v}$, which gives quantile function $\mu(E_t) = v_{\max} E_t^{1/\alpha_v}$, and hence:

$$V(E_t) = \int_0^{\mu(E_t)} v dG_t^v(v) = \frac{\alpha_v}{v_{\max}{}^{\alpha_v}} \int_0^{\mu(E_t)} (v)^{\alpha_v} dv = \frac{\alpha_v}{v_{\max}{}^{\alpha_v}} \frac{v^{1+\alpha_v}}{1+\alpha_v}\Bigg|_0^{\mu(E_t)} = v_{\max}\frac{E_t^{1+1/\alpha_v}}{1+1/\alpha_v}, \tag{B.7.16}$$

which again mirrors MaCurdy utility function (B.7.12) for $\eta = \alpha_v$ and $\Psi = \overline{v}$.

Second, $1+\xi^*_{nlt}$ follows (with *some* nonemployment within each *wage*-type l), a power law distribution with maximum $\Psi\left(\left(\int_0^1 {y_{lt}}^\eta dl\right)^{1/\eta}\overline{\lambda}_t\right)$ and shape parameter η.† This implies the following aggregate labor supply curve:

$$E_t(1+\Xi_t) = F_t(1+\Xi_t) \qquad = P\left(\frac{\Psi s^{1/\eta}}{y_{lt}\overline{\lambda}_t} \leq 1+\Xi_t\right) \tag{B.7.21}$$

$$= \int_0^1 \left(\frac{(1+\Xi_t)y_{lt}\overline{\lambda}_t}{1/\eta}\right)^\eta dl \tag{B.7.22}$$

$$= \left(\frac{(1+\Xi_t)}{\Psi\Big/\left(\left(\int_0^1 y_{lt}^\eta dl\right)^{1/\eta}\overline{\lambda}_t\right)}\right)^\eta. \tag{B.7.23}$$

Third, the elasticity is again precisely η.

Heterogeneous Agent Models

We now move to heterogeneous agent models, where atomistic households make labor supply and consumption decisions with separate budget constraints potentially facing incomplete markets. These class of models can feature heterogeneity in λ_{it}, which is determined in equilibrium.

A useful classification of heterogeneity is whether it is permanent or transitory.

Permanent Heterogeneity With atomistic agents with separate budget constraints but *permanent* heterogeneity, a mass point of marginal individuals endogenously emerges. Specifically, in this setting individuals choose a lifetime fraction of working l_i, or equivalently a probability of working in a given period ϕ_{it} s.t. $\int_{t=0}^{\infty}\phi_{it} = l_i$, as in the time-averaging approach of Ljungqvist and Sargent (2006). Permanent heterogeneity in tastes, endowments or wages affects the average employment probability, yet at each given point in time, these "interior" households are marginal. This local mass of marginal actors makes up

†Intuitively, the distribution of the reservation raise is power law distributed with the same parameter within each labor type. As a result, changes in $1+\Xi_t$ elicit the same proportional employment changes from each labor type, and the aggregate employment elasticity inherits that homogeneous elasticity. Our expression holds for $1+\Xi_t$ small enough that $1+\xi^*_{nlt} > 1+\Xi_t$ holds for some n within *all* labor types l, i.e., the aggregate net of raise rate must be high enough that *some* workers in each labor type are nonemployed. Otherwise, there is full employment from some labor types, and the labor response from those labor types is zero, so the aggregate Frisch elasticity is lower than η, and the CDF (labor supply curve) is:

$$E_t(1+\Xi_t) = F_t(1+\Xi_t) = P\left(n \leq \left(\frac{(1+\Xi_t)y_{lt}\overline{\lambda}_t}{\Psi}\right)^\eta\right) = \int_0^1 \min\left\{\left(\frac{(1+\Xi_t)y_{lt}\overline{\lambda}_t}{\Psi}\right)^\eta, 1\right\} dl. \tag{B.7.20}$$

one minus the fraction of households that either never or always work—implying an empirically uninteresting locally infinite elasticity.[†]

We therefore next move to more realistic models with time-varying heterogeneity, starting with stochastic wages below, then moving to deterministically time-varying wage-age profile in Appendix Section B.7

Time-Varying Heterogeneity: Stochastic Wages (Huggett, 1993) We now consider the popular case where the heterogeneity between households arises from stochastic productivity. Incomplete financial markets mean that income shocks pass through into budget constraints, and thence into consumption/savings policies, assets, consumption, and λ_{it}. To study this setting through the lens of the reservation raise framework, we introduce indivisible labor into the Huggett (1993) model as in Chang and Kim (2006, 2007).

There is a continuum of infinitely lived individuals, in discrete time. Assets a_{it} earn interest r_t. An individual chooses consumption c_{it} and indivisible labor supply $e_{it} \in \{0, 1\}$. Potential earnings y_{it} follow an exogenous Markov process. She faces borrowing limit $a_{\text{min}} < 0$ (set so that positive consumption is always feasible if working even at the lowest earnings level and when at the borrowing constraint), with discount factor $\beta \leq 1$:

$$\max_{c_{it}, e_{it}, a_{it}} \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} \left[\frac{c_{is}^{1-\sigma}}{1-\sigma} - \overline{v} e_{is} \right] \tag{B.7.24}$$

$$\text{s.t.} \quad a_{i,s} = (1 + \Xi_s) y_{is} e_{is} + (1 + r_s) a_{i,s-1} - c_{is} \ \ \forall s \geq t \tag{B.7.25}$$

$$a_{is} \geq a_{\text{min}} \ \ \forall s \geq t. \tag{B.7.26}$$

First, we calculate the reservation raise by individual, indexed by states a and y:

$$1 + \xi^*_{ay} = \frac{\overline{v}}{\lambda_{ay} y}. \tag{B.7.27}$$

[†]To see how permanent heterogeneity can generate trivial reservation raise dispersion (in continuous time), consider a household (indexed by $i \in [0, 1]$) characterized by disutility v_i, initial endowments a_{0i}, and wages w_i (and consumption tastes $u_i(c_{it})$), with stable interest rates $r = \rho$ and no borrowing constraint. So the household's problem is $\max_{c_{it}, e_{it}, a_{it}} \mathbb{E}_t \int_{s=t}^{\infty} e^{-\rho(s-t)} \left[u_i(c_{is}) - v_i e_{is} \right] ds$ subject to a lifecycle budget constraint $\dot{a}_{is} = (1 + \Xi_s) y_i e_{is} + r a_{is} - c_{is} + \mathbb{1}(s = t) \cdot a_{it} \forall s \geq t \Leftrightarrow \int_{s=t}^{\infty} e^{-r(s-t)} c_{is} ds = \int_{s=t}^{\infty} e^{-r(s-t)} (1 + \Xi_s) y_i e_{is} ds + a_{it}$. First, labor supply is an employment policy e^*_{it} characterized by a constant-over-the-lifecycle reservation raise $1 + \xi^*_{it} = \frac{v_i}{\lambda_i y_i} = 1 + \xi^*_i$. Second, the distribution of the reservation raise (labor supply curve) is $E_t(1 + \Xi_t) = F(1 + \Xi_t) = \int_i \mathbb{1}[1 + \xi^*_i \leq 1 + \Xi_t] di$. The constant raise structure implies that for a given prevailing raise $1 + \Xi_t$, there are three reservation raise regions. Two inframarginal regions denote workers that do not work even for (small) net of raise rate increases, as well as those that always work even for small net of raise rate declines. The third set is the set of marginal individuals, who endogenously are *exactly* indifferent, and hence will *all* drop out of work for small net of raise rate declines, and *all* move into employment for small net of raise rate increases. Hence, if there is a mass point of these marginal individuals at the prevailing raise, the labor supply curve will exhibit an infinite Frisch elasticity at the extensive margin.

Second, we calculate the reservation raise distribution (CDF) from the joint distribution of assets and productivities, yielding the labor supply curve:

$$E_t(1+\Xi_t) = F_t(1+\Xi_t) = \sum_{y\in Y}\int_{a_{\min}}^{\infty} \mathbb{1}[1+\xi^*_{ay} \leq 1+\Xi] g_t(a,y)da, \tag{B.7.28}$$

where $g(a, y)$ is the density of agents with assets a and potential earnings y.

Third, the arc elasticities follow Equation (2.6), and depend on the joint distributions of λ and y.

Below we assess these properties with two concrete earnings processes. We solve for consumption and labor supply rules, as well as the joint distribution of assets and productivity states, for an exogenous and constant interest rate $r_s = r \;\; \forall s \geq t$.

Two-State Potential-Earnings Process We start by describing a simple economy with a two-state Markov process for potential earnings, jumping from y_1 to $y_2 > y_1$ (y_2 to y_1) with probability λ_{12} (λ_{21}). Our goal here is to convey intuitions, and to illustrate the complexity of aggregate labor supply already with only two wage states—and how reservation raises can unveil and organize the obscure labor supply curve. The parameters are not picked to match any empirical moments, except for an equilibrium employment rate of 60.7% when $1+\Xi_t = 1$. We plot the distribution of the reservation raises in Appendix Figure B.7 Panel (c).

In the model, for both wage levels, $1+\xi^*_{ay}$ is increasing in assets, since λ_{ay}, the individual's budget multiplier, is decreasing in assets. As expected, $1+\xi^*_{ay_2} < 1+\xi^*_{ay_1}$ for any given asset level a, since higher wages raise consumption and the opportunity cost of not working. For $1+\Xi_t = 1$, all high earners work for any holdings in the asset grid (i.e., $1+\xi^*_{ay_2} < 1 \; \forall a \in [a_{\min}, a_{\max}]$). Low earners work if assets (and consumption) are below threshold $a^*_{y_1}$ s.t. $1+\xi^*_{a^*_{y_1}y_1} = 1$.

The implied labor supply curve is plotted in Appendix Figure B.7 Panel (d), and exhibits complex behavior even with only two wage types, due to the asset distribution. When the labor raise is at $1+\Xi_t = 1$, the marginal individual is a low-wage worker with a relatively high asset level. As $1+\Xi_t$ falls, low-earners drop out of employment in descending order of their assets holdings, with lower and lower density. At some point, the marginal individual is a low-wage earner with assets at the borrowing limit. Since there is a *mass* of such individuals, the labor supply curve is locally infinitely elastic at that point (echoing locally the logic in the models of homogeneity of Hansen, 1985; Rogerson, 1988). As $1+\Xi_t$ falls further, all low-wage individuals become nonemployed, and the marginal individual is now a high earner (and again the pecking order is given by asset holdings).

Realistic Earnings Process We now apply a realistic 33-state potential-earnings process, mimicking that in Kaplan, Moll and Violante (2018) (whose model features only intensive-margin labor supply), which in turn approximates the empirical patterns documented in Guvenen, Karahan, Ozkan and Song (2015). We detail the construction of that variant in Appendix Section B.7. The computational details for the full model are again described in Appendix Section B.7, and the full set of parameters are in Appendix Table B.5.

We plot the distribution of the reservation raises in Appendix Figure B.7 Panel (e). To further illustrate the compositional origins of the reservation raise distribution, Panel (f) plots the reservation raise distribution for three particular out of the 33 total values of potential-earnings states. High-potential-earnings individuals tend to have lower reservation raises, as expected, but the states themselves are not completely informative without reference to the Markov process that guides expected earnings dynamics and equilibrium assets distributions, further highlighting the benefit of the reservation raises as the scalar statistic capturing the heterogeneity relevant to the labor supply choice.

Overall, in the heterogeneous agent model calibrated to a realistic earnings process, the reservation raise distribution is widely dispersed. Specifically and as a result, the model generates a *small* local Frisch elasticity. For a 0.01 pertubation, the downward arc elasticity is 0.72 on the high side, but much smaller upwards (0.18). For large pertubations towards 0.10, the elasticities quickly settle in below 0.5. The equilibrium reservation raise distribution and hence labor supply curve inherit the joint distribution of λ and y, so that the curve is particularly inelastic if low earnings realizations are offset by associated high λ values.

Intensive and Extensive Margins, and Lifecycle Dynamics: the Rogerson and Wallenius (2009) Model

As in the general intensive-margin case in Section 2.4, permitting hours choices preserves the reservation raise logic. A leading model with both margins is that by Rogerson and Wallenius (2009), which also features lifecycle patterns (and the Frischian behavior of which Chetty, Guren, Manoli and Weber, 2012, studied as a leading macro model with an extensive margin). We discuss our parameterization in Appendix Section B.7, largely following the paramaterization choices of Chetty, Guren, Manoli and Weber (2012), but we change the tax rate and apply a 60.7% employment rate target for consistency with all our models and the survey, all hence matching our U.S. 2019 broad population benchmark.

The overlapping generations economy is set in continuous time and has a unit mass of individuals born at every instant, denoted by i, and each lives for a length of time equal to one. Age at time t is denoted by $d_{it} \in (0, 1)$. (In our calibration, we will set the discount rate to zero, and individuals can save and borrow at zero interest rate.) The individual freely chooses hours worked h_{it} and consumption c_{it} at some utility $u(c_{is})$, which is separable from disutility of hours, here following the MaCurdy isoelastic sturcture with $v(h_{it}) = \Gamma \frac{h_{it}^{1+1/\gamma}}{1+1/\gamma}$. Earnings $\theta_{is}(h_{is})$ depend on hours subject to a nonconvexity and age, as we discuss below. The optimization problem at time t for individual i of age d (with remaining lifetime $1 - d_{it}$) is:

$$\max_{c_{it},h_{it}} \mathbb{E}_t \int_{s=t}^{t+(1-d_{it})} e^{-\rho(s-t)} \left[u(c_{is}) - v(h_{is})\right] ds \tag{B.7.29}$$

$$\text{s.t. } c_{is} + \dot{a}_{is} = r_s a_{is} + (1 + \Xi_s) y_{is} \;\; \forall t + (1 - d_{it}) \geq s \geq t. \tag{B.7.30}$$

Earnings $\theta_{is}(h_{is})$ are structured as follows. Hourly wages $w_{it} = w_{d_{it}}$ are a triangular, single-peaked function of age d, generating lifecycle aspects. Moreover, rather than $y = hw$, to generate an extensive margin, $\theta_{is}(h_{is})$ features a nonconvexity of earnings in hours, in form of fixed hours cost: labor hours are productive, and hence are paid wages w_d, only above hours threshold $\underline{h}$:

$$\theta_{it}(h_{it}) = w_{d_{it}} \cdot \max\{h_{it} - \underline{h}, 0\}. \tag{B.7.31}$$

Absent this fixed cost, the marginal disutility at $h = 0$ hours is zero, and so everyone works positive hours (provided positive wages)—eliminating the extensive margin, as in our intensive-margin example in Section 2.4.

First, in a given period t, heterogeneity in reservation raises solely reflect heterogeneity in age d, so we can write reservation raises and choices indexed by age types d. Hours choices $h^*_{dt}(1+\Xi_t)$ are given by $(1+\Xi_t)w_d\lambda_{dt} = \Gamma[h^*_{dt}(1+\Xi_t)]^{1/\gamma}$. Since our context features an intensive margin, this reservation raise is implicitly defined as a fixed point, as in our general job-choice case in Section 2.4:

$$1 + \xi^*_{dt} = \frac{v\left(h^*_{dt}(1+\xi^*_{dt})\right)}{\lambda_{dt}\theta_{dt}(h^*_{dt}(1+\xi^*_{dt}))} = \frac{\frac{\Gamma}{1+\frac{1}{\gamma}}\left(\frac{\lambda_{dt}(1+\xi^*_{dt})w_d}{\Gamma}\right)^{\gamma+1}}{\lambda_{dt}w_d\left(\left[\frac{\lambda_{dt}(1+\xi^*_{dt})w_d}{\Gamma}\right]^{\gamma} - \underline{h}\right)}. \tag{B.7.32}$$

That is, individuals work when the (hourly) wage is above some threshold w^*.[†] Also, setting $\underline{h} = 0$ nests the MaCurdy intensive-margin-only setting, with $1 + \xi^*_{dt} = 0$ for all workers and ages, as in our general intensive-margin job choice in Section 2.4.

Second, Appendix Figure B.7 Panel (h) plots the histogram of the reservation raise distribution, which also gives the aggregate labor supply curve:

$$E_t(1+\Xi_t) = P\left(\frac{\Gamma\left(\underline{h}(1/\gamma+1)\right)^{1/\gamma}}{\lambda_{dt}w_d} \leq 1 + \Xi_t\right) \tag{B.7.33}$$

$$= P\left(\frac{1}{w_d} \leq \frac{1+\Xi_t}{\Gamma\left(\underline{h}(1/\gamma+1)\right)^{1/\gamma}/\lambda_{dt}}\right). \tag{B.7.34}$$

In the nonstochastic steady state, the budget constraint (B.7.30) reduces into a lifecycle budget constraint, generating a homogenous λ, so the distribution of reservation raises solely inherits that of $1/w_d$, a feature we discuss in detail below.

[†] In fact, without uncertainty and perfect capital markets and hence a lifetime budget constraint $\overline{\lambda}$, we could then solve for the age-specific reservation raise explicitly as $1+\xi^*_d = \frac{\Gamma\left(\underline{h}(1/\gamma+1)\right)^{1/\gamma}}{\overline{\lambda}w_d}$, and therefore also solve for the reservation wage and hence marginal ages.

Third, we compute the extensive-margin arc elasticities. The Frisch arc elasticities range from 2.60 to 3.20 in this particular calibration, with local elasticities (from 0.01 net of raise rate pertubations) between 2.84 and 2.90.[†]

Further Computational Details

We describe additional details of the models discussed in Section 2.4 and above.

The Representative Household: A Short-Lived, Uncompensated Shock

Here we describe how we model and quantify the uncompensated labor supply response of a representative household with MaCurdy-style convex labor supply disutility and shared consumption, depicted in Appendix Figure B.8.

We consider a household that maximizes

$$\max_{C_s,E_s} \sum_{s=t}^{\infty} \beta^{s-t} \frac{C_s^{1-\sigma}}{1-\sigma} - \Psi \frac{E_s^{1-\eta}}{1-\eta} \tag{B.7.35}$$

$$\text{s.t.} \sum_{s=t}^{\infty} \frac{1}{1+r} C_s \leq \sum_{s\geq t}^{\infty} \frac{1}{1+r} (1+\Xi_s) y_s E_s, \tag{B.7.36}$$

so that wages are constant at $y_t = y \forall t$. We also consider the case were $\beta(1+r) = 1$ so that $C_t = C \forall t$. We also have assumed that initial assets A_0 are zero, which implies the largest wealth effect among the range of nonnegative initial asset holdings, thereby providing the largest difference between the Frischian and uncompensated setting (away from the representative household being borrowing-constrained, a setting covered by our heterogeneous agent model).

We study partial equilibrium, i.e., hold aggregate equilibrium variables (interest rates, net of aggregate prevailing raise potential earnings/wages) fixed. We first construct the employment baseline for the unperturbed setting. Denote $\bar{E}$ and $\bar{C}$ as the employment and consumption levels in a stable setting in which $1 + \Xi_t = 1 \forall t$. The intratemporal substitution condition and the budget constraint imply, respectively $y\bar{C}^{-\sigma} = \Psi \bar{E}^{1/\eta}$ and $\bar{C} = y\bar{E}$. Solving these conditions for $\bar{E}$ delivers $\bar{E} = \left[\frac{y^{1-\sigma}}{\Psi}\right]^{\frac{\eta}{1+\eta\sigma}}$.

Second, we turn to labor supply under a pertubation of the raise of size $1 + \Xi$ lasting T periods. In our uncompensated experiment, we set the baseline aggregate prevailing raise $1 + \Xi_t = 1 + \Xi$ for $t = 1, \ldots, T$, potentially diverging at a constant level from the baseline raise subsequently reset to unity at $1 + \Xi_t = 1$ for $t > T$. The labor response we plot is labor supply in period 1 under the initial raise level $1 + \Xi_t = 1 + \Xi$.

[†]In principle, we could obtain the elasticity analytically from the reservation raise distribution. Our method to measure the arc elasticities on the basis of the reservation raise distribution complements the construction of the Frisch elasticity by Chetty, Guren, Manoli and Weber (2012), who simulate a small, short-lived one percentage point tax change, which requires repeatedly solving the model for each generation, may include non-Frischian features, and only isolates one arc elasticity.

Figure B.8: Frischian vs. Uncompensated Quarter-Long Deviation in the Aggregate Prevailing Raise: Extensive-Margin Aggregate Labor Supply Responses in Three Calibrated Models

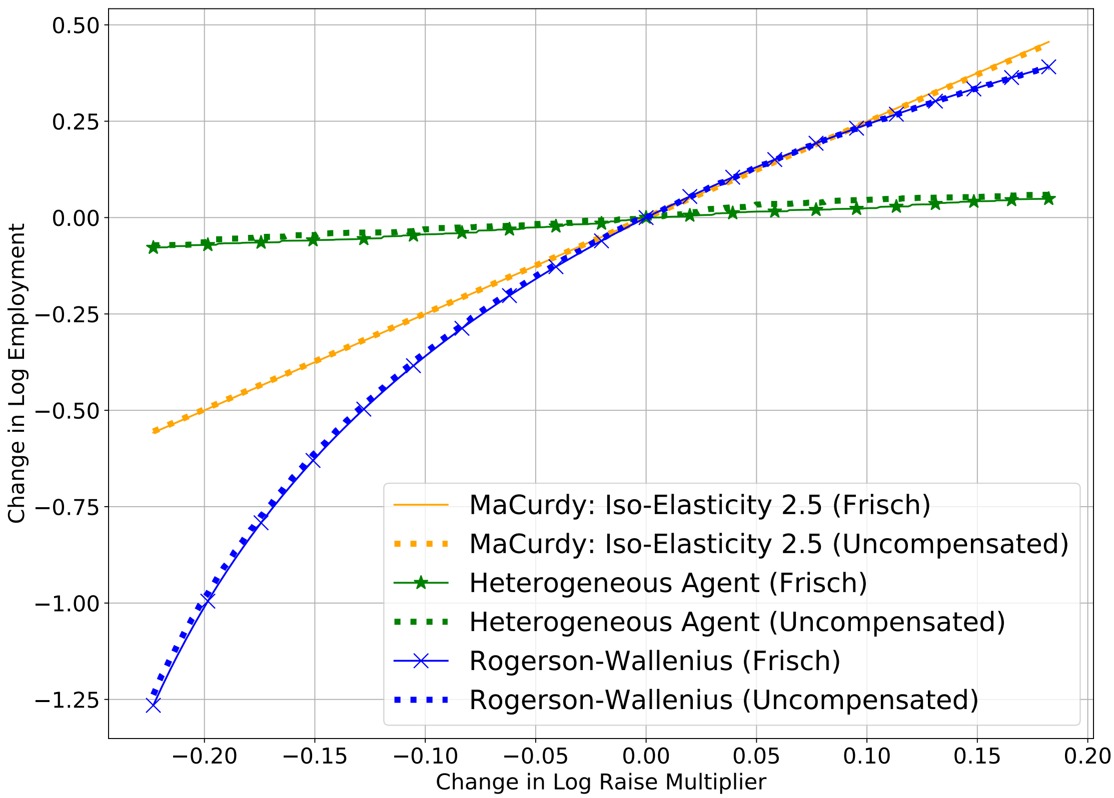

Note: The figure compares aggregate labor supply curves that are purely Frischian (from our reservation raise distributions) and from non-Frischian, uncompensated pertubations in the aggregate prevailing raise that are short-lived and last one quarter in each model. The three curves are output from simulating three of the models we discuss in detail in Section 2.4: a representative household model with an isoelasticity of 2.5, a heterogeneous agent mode with a realistic 33-state earnings process, and the Rogerson-Wallenius model with lifecycle aspects and an intensive-margin hours choice. The specific quantitative experiments are detailed in Appendix Section B.7 for each model. Model parameters are in Appendix Table B.5.

Let E' and E'' denote labor supply when $1+\Xi_t = 1+\Xi$ and $1+\Xi_t = 1$ respectively. Then, optimal intratemporal labor supply implies

$$yC^{-\sigma} = \Psi E''^{1/\eta} \tag{B.7.37}$$

$$(1+\Xi)yC^{-\sigma} = \Psi E'^{1/\eta}. \tag{B.7.38}$$

Hence, initial equals eventual labor supply times the Frisch-elasticity-scaled raise:

$$\implies E' = (1+\Xi)^{\eta} E''. \tag{B.7.39}$$

The budget constraint then implies for consumption C in this raise series (or for λ):

$$\sum_{t=T+1}^{\infty}\left(\frac{1}{1+r}\right)^t yE'' + \sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^t (1+\Xi)wE' = \sum_{t=0}^{\infty}\left(\frac{1}{1+r}\right)^t C \tag{B.7.40}$$

$$\sum_{t=T+1}^{\infty}\left(\frac{1}{1+r}\right)^t yE'' + \sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^t (1+\Xi)^{1+\eta}wE'' = \frac{1+r}{r}C \tag{B.7.41}$$

$$\frac{1+r}{r}yE'' - \sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^t \left(1-(1+\Xi)^{1+\eta}\right) yE'' = \frac{1+r}{r}C \tag{B.7.42}$$

$$yE'' - \frac{r}{1+r}\sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^t \left(1-(1+\Xi)^{1+\eta}\right) yE'' = C \tag{B.7.43}$$

$$\left[1 - \frac{r}{1+r}\sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^t \left(1-(1+\Xi)^{1+\eta}\right)\right] yE'' = C. \tag{B.7.44}$$

Let $m(T, 1+\Xi) \equiv \left[1 - \frac{r}{1+r}\sum_{t=1}^{T}\left(\frac{1}{1+r}\right)^t \left(1-(1+\Xi)^{1+\eta}\right)\right]$. Combining the above with the intratemporal substitution condition (B.7.38), one can solve for L' in particular as a function of baseline employment level $\bar{E}$ in the unperturbed setting, duration of the pertubation T, and raise deviation $1+\Xi$:

$$E' = \left[(1+\Xi)^{\eta} m(T, 1+\Xi)^{-\sigma\eta/(1+\sigma\eta)}\right]\left(\frac{y^{1-\sigma}}{\Psi}\right)^{\frac{\eta}{1+\sigma\eta}} = \left[(1+\Xi)^{\eta} m(T, 1+\Xi)^{-\sigma\eta/(1+\sigma\eta)}\right]\bar{E}. \tag{B.7.45}$$

The model is calibrated so that the period length corresponds to one month, so this experiment simulates a one-quarter shift in the prevailing aggregate labor raise by implementing a three-period duration of the shift. The quarterly interest rate is set to 0.764% (implying an annual discount factor of 0.97).

The Heterogeneous Agent Model with Extensive-Margin Labor Supply

We describe the model the solution algorithm, and how we simulate the short-lived uncompensated shock. We also describe the 33-state potential-earnings process.

The Model In this section we describe our modification to Huggett (1993), with endogenous labor supply, which occurs along the extensive margin only.

Individuals solve

$$\max_{c_{it}, e_{it}\in\{0,1\}, a_{it}} \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t}\left[\frac{c_{is}^{1-\sigma}}{1-\sigma} - \overline{v}e_{is}\right] \tag{B.7.46}$$

$$\text{s.t.} \quad a_{i,s} = (1+\Xi_s)y_{is}e_{is} + b(1-e_{is}) + (1+r_s)a_{i,s-1} - c_{is} \quad \forall s \geq t \tag{B.7.47}$$

$$a_{is} \geq a_{\min} \quad \forall s \geq t, \tag{B.7.48}$$

where $y_{i,t}$ follows the Markov process described in Appendix Section B.7 below. Households endogenously choose their labor supply e_{it}, which is restricted to 0 or 1. As described in the main text, since individuals within the same asset and productivity levels face the same problem, consumption and labor supply decisions (and hence reservation raises) can be written as a function of assets and productivity.

The first-order condition on consumption is, as in the standard case,

$$u_c(c^*(a,y), e^*(a,y)) = V_a(a,y), \tag{B.7.49}$$

where V is the value function for someone at asset level a and earnings state y. The optimality condition on labor supply is

$$e^*(a,w) = \begin{cases} 1 & \text{if } V_a(a,y)y > \overline{v} \\ 0 & \text{if } V_a(a,y)y < \overline{v}. \end{cases} \tag{B.7.50}$$

A similar optimality condition solves the problem at the binding constraint a_{min}:

$$e^*(a_{\text{min}},y) = \begin{cases} 1 & \text{if } \frac{(y+ra)^{1-\sigma}}{1-\sigma} - v > \frac{(ra)^{1-\sigma}}{1-\sigma} \\ 0 & \text{otherwise }. \end{cases} \tag{B.7.51}$$

If $a_{\text{min}} < 0$, individuals at the borrowing constraint hence are always employed.

Solution Algorithm We solve the model with parameters $\sigma = 2$, $r = 0.03$, $\beta = 0.97$, and unemployment insurance $b = 0$. We set the borrowing constraint at $a_{\text{min}} = -z_1 r + 0.001$, so that positive consumption is possible at the lowest productivity and asset levels if the individual works. We choose the labor supply disutility shifter $\bar{v}$ to match the equilibrium employment rate 60.7%.

We use a grid of assets comprising a discrete set of asset levels A with minimum a_{min} and maximum $a_{\text{max}} = 50000000$. We place fifty asset levels equally spaced between a_{min} and 0, 450 levels between 0 and 1000000, and 500 levels between 1000000 and 5000000. We solve the consumption and labor supply rules using value function iteration:

$$V^{n+1}(a,y) = \max_{a' \in A, e \in \{0,1\}} \left\{ u(ye + (1+r)a - a') + \beta \sum_{y'} T_{y,y'} V(a',y') \right\}, \tag{B.7.52}$$

where $T_{y,y'}$ is the transition probability between productivity levels y and y'. Consumption is given by $c(a,y) = ye^*(a,y) + (1+r)a - a'^*(a,y)$, where e^* and a'^* are the solutions to the maximization problem in (B.7.52) for an individual characterized by asset and productivity states (a,y).

Consumption and labor supply rules solved for, we calculate the equilibrium joint distribution of a and y, $g(a,y)$, by solving the system of equations:

$$g(a,y) = \sum_{\tilde{y}} \sum_{\tilde{a} \text{ s.t. } a'^*(\tilde{a},\tilde{y}) = a} g(\tilde{a},\tilde{y}) T_{\tilde{y},y}. \tag{B.7.53}$$

With the joint distribution of assets and productivity assets, value functions, and consumption choices, we can solve for the distribution of reservation raises, and therefore the labor supply curve.

A Short-Lived, Uncompensated Shock We describe how we obtain the uncompensated labor supply curve in response to a quarter-long raise pertubation depicted in Appendix Figure B.8. The purpose of this exercise is to simulate the aggregate extensive-margin labor supply response of a heterogeneous agent economy under an *uncompensated* (non-Frischian) one-period change in the benefit of working i.e., the prevailing aggregate labor raise. We study partial equilibrium, i.e., hold aggregate equilibrium variables (interest rates, potential earnings) fixed.

Consider an individual with assets a and productivity y. That individual faces a temporary prevailing aggregate raise $1 + \Xi_s = 1 + \Xi$ during some period s, which then returns to a raise of level $1 + \Xi_t = 1$ for $t > s$. Then, that individual solves

$$\max_{c, e \in \{0,1\}} \left\{ u(c, e) + \beta \sum_{y'} T_{y,y'} V(a', y') \right\} \tag{B.7.54}$$

$$\text{s.t. } a' = (1 + r)a - c + ey, \tag{B.7.55}$$

where $u(c, e) = \frac{c^{1-\sigma}}{1-\sigma} - \bar{v}e$, and where V is the value function from the solution to the equilibrium with the baseline unit raise in all periods.

For a given prevailing labor raise, the solution is easily found by maximizing the utility over a grid of consumption points under employment and nonemployment, since the problem is not recursive. We then measure the labor supply response as the difference in the measure of individuals who choose employment under the temporary first-period-only labor raise $1 + \Xi$ versus the measure of individuals who choose employment in the baseline economy with the unit raise.

The model is calibrated so that the period length corresponds to one quarter, so this experiment simulates a one-quarter shift in the prevailing aggregate raise.

The Potential Earnings Process

We now apply a realistic 33-state potential-earnings process, mimicking that in Kaplan, Moll and Violante (2018) (whose model features only intensive-margin labor supply), which in turn approximates the empirical patterns documented in Guvenen, Karahan, Ozkan and Song (2015). Our Markov process represents an underlying process modeled as the sum of two independent components $\log y_{it} = \log y_{1,it} + \log y_{2,it}$, with the log of each component $y_{j,it}$ evolving according to a jump-drift process. Jumps arrive at a Poisson rate γ_j, and trigger new draws of the earnings component from a mean-zero normal distribution. Between jumps, the process drifts toward zero at rate β_j.[†] Kaplan, Moll

[†]Of course, in our model not all individuals will work; we do not estimate a latent potential earnings process such that the modeled realized earnings, taking into account labor supply decisions, would generate realized empirical earnings dynamics.

and Violante (2018) implement this process as two finite-state continuous time Markov processes for each independent component. In our application, we do so as a single discrete-time Markov process in which the income states are hence combinations of the states of the two income processes.[†] We discretize the continuous time transition rates between states by using the matrix exponential; i.e., the discrete time transition matrix for income component j is calculated as $T^{j,d} = \exp T^{j,c} = \sum_{k=0}^{\infty} \frac{1}{k!} \left(T^{j,d}\right)^k$, where $T^{j,c}$ is the continuous time transition matrix for component j. The continuous time transition rates are measured with quarters as the unit of time, so the discrete time transition matrix is also in quarters. Then, we collapse the discrete time transition matrices for the two components into a single transition matrix between one-dimensional income states. $T^d_{y,y'}$, the transition probability between the single-dimension income state y to y' for which $\log y = \log y_1 + \log y_2$ and $\log y' = \log y'_1 + \log y'_2$, is then equal to $T^{1,d}_{y_1,y'_1} T^{2,d}_{y_2,y'_2}$. (For this process and the income levels chosen, conveniently each y state is uniquely defined by one (y_1, y_2) combination.)

The Rogerson and Wallenius (2009) Model

We describe the solutions, and the simulations of the shock.

Parameterizing the Baseline Model The original Rogerson and Wallenius (2009) distribution of the hourly wage w_d (labor efficiency) arises from a uniform age distribution and a triangular wage-age gradient (single-peaked at $d = 1/2$ with $w_{d=1/2} = \hat{w}_0$ as the maximum wage level, and generally $w(d) = \hat{w}_0 - \hat{w}_1|d - 0.5|$). We approximate the continuum of generations with 1,000,000 equally spaced discrete generations, and solve the model following the Technical Appendix of Chetty, Guren, Manoli and Weber (2012). We choose the utility function parameters (Γ, the labor disutility shifter, γ, the labor supply intensive-margin elasticity), effective labor supply parameters ($\overline{h}$, the minimum number of hours worked, and $\hat{w}_1$, the slope of the wage-age gradient) and the tax rate at which the model equilibrium is calculated. We assume CRRA log consumption utility ($\sigma = 1$). We set the initial tax (raise) rate at 26%, which was the average net tax rate faced by an average single worker in 2017. We set the labor supply intensive-margin elasticity to 2.0. From this point, we conduct two paramaterizations. In the first, we choose the remaining three parameters, Γ, $\overline{h}$, and $\hat{w}_1$, to match three equilibrium targets, as in Chetty, Guren, Manoli and Weber (2012): the employment rate (60.7%, as in the other model exercises), the maximum intensive-margin hours choice (0.45), and the ratio of the lowest wage to the highest wage received over the lifecycle (0.5). This paramaterization sets $\Gamma = 42.492$, $\overline{h} = 0.258$, and $\hat{w}_1 = 0.851$.

For each generation/age, indexed by d, we calculate hours at each age, h^*_d, and then calculate the reservation raises using $1 + \xi^*_d = \frac{(1-\tau)w_d(h^*_d - \overline{h})u'(c_d)}{v(h^*_d)}$. This formulation of the

[†]Inconsequential for quantities, we normalize the earnings state levels so that the average steady-state potential earnings are equal to the 2015 U.S. average personal income.

reservation raise is "normalized" so that the relevant wage is the after-tax wage, and so the indifferent worker is that of the age d such that $1+\xi_d^*=1$.

This, combined with the (uniform) distribution of age, gives the distribution of reservation raises, from which we can compute the arc elasticities.

Low-Frisch Elasticity Parameterization Second,, we also set the peak of the wage-age profile to target a lower extensive-margin Frisch elasticity. This paramaterization sets $\Gamma = 40.000$, $\overline{h} = 0.248$, $\hat{w}_1 = 1.319$, and lifetime peak productivity at 1.110.

Shutting off the Intensive Margin We also add a variant that shuts off intensive-margin reoptimization. We do so by simply solving for the optimal policies, extracting the reservation raises, and then computing alternative reservation raises that hold hours fixed at the corresponding unit raise point, such that $1+\xi_d = \frac{v(h^*_{d,1+\Xi=1})}{\theta_d(h^*_{d,1+\Xi=1})\lambda}$.

A Short-Lived, Uncompensated Shock We simulate the labor supply response under a temporary, short, but noninstantaneous (and therefore non-Frischian) shift in the prevailing aggregate raise. As in the other models, we again study partial equilibrium, i.e., hold aggregate equilibrium variables (e.g., interest rates) fixed. We suppress calendar time subscripts in what follows. We continue to solve the model in continuous time, i.e., in the context of considering a time interval corresponding to a month-long duration, one could work for part of the period rather than having a period-long policy. Households are subject to our aggregate prevailing labor raise $1+\Xi$ for a time interval of duration m. After this interval, the raise returns to unity. The raise shock is unanticipated, and once occurring, the households perfectly foresee that the raise deviation will last exactly m time units before returning to unity. Upon realization of the shock, households will re-optimize their planned consumption and labor supply paths.

To solve for assets, we first solve for assets at age d before the raise shock. Currently held assets are determined by past earnings, government transfers (which are equal to $\tau\bar{c}$, where $\bar{c}$, taken as parametric by the household, is the equilibrium consumption level in turn equal to average income and hence $\tau\bar{c}$ is the average labor income raise payment and also government rebate), and consumption c:

$$\int_0^d \big((1-\tau)e_{\tilde{d}}y_{\tilde{d}} + \tau\bar{c} - c\big)\, d\tilde{d}, \tag{B.7.56}$$

where $e_{\tilde{d}}$ is desired employment and $y(\tilde{d})$ is potential gross earnings at age $\tilde{d}$. For $\tilde{d} \in [d_{\text{min}}, d_{\text{max}}]$, where d_{min} and d_{max} are the (endogenous) entry and exit ages,

$$e_{\tilde{d}}y_{\tilde{d}} = w_{\tilde{d}}(h_{\tilde{d}} - \bar{h}) = w_{\tilde{d}}(h_0\hat{w}_0^{-1/\gamma}w_{\tilde{d}}^{1/\gamma} - \bar{h}) = [h_0\hat{w}_0^{-1/\gamma}w_{\tilde{d}}^{1+1/\gamma} - \bar{h}w_{\tilde{d}}] \tag{B.7.57}$$

$$= \begin{cases} [h_0\hat{w}_0^{-1/\gamma}\left(\hat{w}_0 - 0.5\hat{w}_1 + \hat{w}_1\tilde{d}\right)^{1+1/\gamma} - \bar{h}\left(\hat{w}_0 - 0.5\hat{w}_1 + \hat{w}_1\tilde{d}\right)] & \text{if } \tilde{d} < 0.5 \\ [h_0 e_0^{-1/\gamma}\left(\hat{w}_0 + 0.5\hat{w}_1 - \hat{w}_1\tilde{d}\right)^{1+1/\gamma} - \bar{h}\left(\hat{w}_0 + 0.5\hat{w}_1 - \hat{w}_1\tilde{d}\right)] & \text{if } \tilde{d} \geq 0.5, \end{cases} \tag{B.7.58}$$

and 0 if $\tilde{d} \notin [d_{\min}, d_{\max}]$. The lifetime gross-of tax/raise labor income up to age d is:

$$\int_0^d e_{\tilde{d}} y_{\tilde{d}} d\tilde{d} = \begin{cases} 0 & \text{if } d < d_{\min} \\ \left(\frac{h_{\tilde{d}}}{(2+\frac{1}{\gamma})\hat{w}_1} - \frac{\bar{h}}{2\hat{w}_1}\right) w_{\tilde{d}}^2 \Big|_{d_{\min}}^{d} & \text{if } d_{\min} \leq d < 0.5 \\ \left(\frac{\tilde{d}}{(2+\frac{1}{\gamma})\hat{w}_1} - \frac{\bar{h}}{2\hat{w}_1}\right) w_{\tilde{d}}^2 \Big|_{d_{\min}}^{0.5} & \\ \qquad +(1-\tau)\left(-\frac{h_{\tilde{d})}}{(2+(1-\tau)\frac{1}{\gamma})\hat{w}_1} + \frac{\bar{h}}{2\hat{w}_1}\right) w_{\tilde{d}}^2 \Big|_{0.5}^{d} & \text{if } 0.5 \leq d < d_{\max} \\ \left(\frac{h_{\tilde{d}}}{(2+\frac{1}{\gamma})\hat{w}_1} - \frac{\bar{h}}{2\hat{w}_1}\right) w_{\tilde{d}}^2 \Big|_{d_{\min}}^{0.5} + \left(-\frac{h_{\tilde{d}}}{(2+\frac{1}{\gamma})\hat{w}_1} + \frac{\bar{h}}{2\hat{w}_1}\right) w_{\tilde{d}}^2 \Big|_{0.5}^{d_{\max}} & \text{if } d \geq d_{\max}, \end{cases} \tag{B.7.59}$$

from which follows lifetime net income if multiplied by $1 - \tau$.

Consider an individual of age d. Let m denote the length of the temporary raise change. One solves for optimal consumption and labor supply by finding the consumption level $c_{\Xi,d}$ that balances the income's lifetime budget constraint, subject to (a) their labor income being subjected to a multiplier and (b) the individual adjusting the remainder of their lifetime's labor supply to meet extensive and intensive-margin labor supply optimality conditions. In our experiment, for each given age level d, the time series of the aggregate prevailing raise will be given by

$$1 + \Xi_{\tilde{d}} = \begin{cases} 1 + \Xi & \text{if } \tilde{d} \in [d, d+m] \\ 1 & \text{if } \tilde{d} > d + m. \end{cases} \tag{B.7.60}$$

For a proposed consumption level $c_{\Xi,d}$ (where subscript d denotes the time at which the raise pertubation started, rather than the period during which the consumption occurs, as consumption is constant across all post-raise ages $\tilde{d} > d$), during the ages $\tilde{d} > d$, let $h_{\tilde{d},d}$ be the age $\tilde{d} > d$ labor supply choice of an individual that was age d when the temporary labor raise shift began.

For working ages $\tilde{d}$, intensive-margin labor supply implies that

$$\Gamma h_{\tilde{d},d}^{\gamma} = (1-\tau)(1+\Xi_{\tilde{d}})u'(c_{\Xi,d})w_{\tilde{d}}. \tag{B.7.61}$$

As in the standard setup, there will be cutoff rules that dictate extensive-margin labor supply. Under a temporary $1 + \Xi$ shift, one cannot dictate age cut-offs since the benefit of working does not follow the same single-peaked shape as the original model. However, one can determine raise-productivity cutoffs in $(1 + \Xi_{\tilde{d}})w_{\tilde{d}}$.

At ages $\hat{d}$ at which the individual is indifferent to extensive-margin labor supply (conditional on optimizing on the intensive margin if working), the intensive and extensive-margin conditions imply respectively:

$$\Gamma h_{\hat{d},d}^{\gamma} = (1-\tau)(1+\Xi_{\hat{d}})u'(c_{\Xi,d})w_{\tilde{d}} \tag{B.7.62}$$

$$\Gamma \frac{h_{\hat{d},d}^{1+\gamma}}{1+\gamma} = (1-\tau)(1+\Xi_{\hat{d}})u'(c_{\Xi},d)w_{\tilde{d}}(h_{\hat{d},d}-\bar{h}). \tag{B.7.63}$$

Combining these two implies an hours choice at the marginal age of $h_{\hat{d},d} = \frac{(1+\gamma)}{\gamma}\bar{h}$, on the basis of which we solve for the marginal age (productivity) as follows:

$$\Gamma\left(\frac{(1+\gamma)}{\gamma}\bar{h}\right)^{\gamma} = (1-\tau)(1+\Xi_{\hat{d}})u'(c_{\Xi,\hat{d}})w_{\tilde{d}} \implies (1+\Xi_{\hat{d}})w_{\tilde{d}} \ = \frac{\Gamma\left(\frac{(1+\gamma)}{\gamma}\bar{h}\right)^{\gamma}}{(1-\tau)u'(c_{\Xi,\hat{d}})}. \tag{B.7.64}$$

The individual will prefer working over nonworking at age $\tilde{d}$ if $(1+\Xi_{\tilde{d}})w(\tilde{d}) \geq \Gamma\left(\frac{(1+\gamma)}{\gamma}\bar{h}\right)^{\gamma} / \left((1-\tau)u'(c_{\Xi,d})\right)$. From this cutoff, one can compute optimal planned extensive-margin supply for every age $\tilde{d} > d$. For a proposed candidate for the consumption level, one can then compute the balance of the individual's lifetime budget constraint given both the change in consumption and the lifetime extensive- and intensive-margin labor supply responses.[†] The solution to the individual's problem is the consumption level $c_{\Xi,d}$ that balances the individual's lifetime budget constraint. Repeating this for every individual in the economy (i.e., repeating this for every age $d \in [0,1]$) delivers the aggregate labor supply response. We measure the labor supply response to this temporary (but noninstantaneous) raise shift using the change in labor supply upon impact of the raise.

We set the length of the uncompensated raise shift to 1/240, to represent the length of one quarter out of a 60-year adult lifespan.

[†]We isolate the labor supply responses, and therefore hold fixed in our partial-equilibrium experiment all aggregate variables except for the prevailing raise (i.e., government transfers and taxes, so the government budget is unbalanced in this exercise).

Appendix C

Appendix to Chapter 3

C.1 Additional Empirical Results

Panel Specification

One concern with the cross-sectional specifications is that there may be some unobserved aggregate factor that induced large increases in UI claims at the same time that states and local municipalities implemented SAH orders. Alternatively, there may be time-invariant state-specific factors that drove both increases in unemployment claims and SAH orders. To address these concerns, we employ a panel specification, which allows us to control for week and state fixed effects.

We modify the specification so that the outcome variable is the flow value of initial claims on date t and the SAH order treatment is the share of the *current week* that a state was subject to SAH orders, where we take a weighted average of county-level exposure as before.[†]

$$\frac{UI_{s,t}}{Emp_s} = \alpha_s + \phi_t + \beta_P \times SAH_{s,t,t-7} + \mathbf{X}_{s,t}\Gamma + \epsilon_{s,t} \tag{C.1.1}$$

We consider a variety of state-time controls. We include two lags of $SAH_{s,t,t-7}$ to account for dynamics in the effect of SAH orders on unemployment claims. Additionally, we include the share of the population that works from home, the number of confirmed cases per one thousand people, and the Bartik-style employment control from before. Each of these three controls is interacted with a dummy equal to one for weeks ending March 21st, 2020 and onward.[†] We estimate the following fixed effects panel regression on weekly observations for the week ending January 4 through the week ending April 11.[†]

[†]Because in our sample no state or local municipality reopened, once $SAH_{s,t,t-7} = 1$ it remains equal to one for all remaining weeks.

[†]Note that because our measures of work-from-home and employment loss are constant across time, we are controlling for the relative effect of each from before the week ending March 21st.

[†]We drop the first two weeks in all specifications to ensure the sample size is constant throughout.

Table C.1 provides our estimate of $\hat{\beta}_P$ for the contemporaneous effect and two lags. Column (1) presents the results with no lags. The point estimate of 0.90% (SE: 0.35%) suggests that a full week of SAH order exposure increased unemployment claims by .90% of total state-level employment. In column (2), we include two lags of SAH orders. The point estimate on the contemporaneous effect is little changed, though it rises slightly. Importantly, neither of the coefficients on the first nor the second lag is significant. This result suggests that, in our sample, that SAH orders have constant, contemporaneous effects on UI claims. At longer horizons, we would suspect non-linearities to eventually kick in, with the effect of SAH orders declining. Finally, our point estimates are little changed when including additional controls in Column (3).

Our estimates $\hat{\beta}_P$ in the first three columns tend to be somewhat lower than what we find in our benchmark, cross-sectional design. In particular, the panel design implies that each week of SAH exposure increased UI claims by 1% of state employment; in contrast, our estimates of $\hat{\beta}_C$ imply that each week of SAH exposure increased UI claims by approximately 1.9% of state employment. While, at first glance, β_C and β_P aim to estimate the same moment, the inclusion of state and time fixed effects imply that they are not directly comparable.[†] In column (4), we consider the panel specification in which we drop state fixed effects, to make the panel and cross-sectional regressions comparable: the point estimate rises to 1.2% and is statistically indistinguishable from what we find in the cross-section.

High Frequency Effects on Proxies for Local Economic Activity

In this subsection, we provide additional evidence that the SAH orders had immediate and highly localized effects on daily indicators of economic activity. This exercise is important because of concerns that the state-level effects we estimate above simply reflect differential labor market disruptions that would have occurred in the absence of SAH orders in precisely those places most likely to implement SAH orders earliest.

We estimate the local effect of SAH using high frequency proxies for economic activity from Google's Community Mobility Report, which measures changes in visits to establishments in various categories, such as retail and work.[†] Early on in the COVID-19 pandemic, Google began publishing data documenting how often its users were visiting different types of establishments. The data are reported as values relative to the median visitation rates by week-day between January 3, 2020 and February 6, 2020.[†,†]

[†]See Kropko and Kubinec (2020) for a discussion of the proper interpretation of two-way fixed effect estimators in relation to one-way fixed effect estimators.

[†]`https://www.google.com/covid19/mobility/`

[†]One possible limitation of this data is that the sample of accounts included in the surveys is derived from only those with Google Accounts who opt into location services. We believe sample selection bias is unlikely to be a major concern given Google's broad reach (there are over 1.5 billion Gmail accounts, for example).

[†]Note that for privacy reasons, data is missing for some days for some counties. When possible, we carry forward the last non-missing value. Excluding counties with missing values yields the same result;

We use the retail and workplace mobility indices because these two indices are consistently recorded for the time sample we study. Failing to find an effect on these proxies for local economic activity would call into question the results we find in the aggregate, at the state-level. We interpret retail mobility as broadly representing "demand" responses to SAH orders and workplace mobility as broadly representing "supply," at least on-impact.[†] Over longer-horizons, workers laid off because of demand-side disruptions will, naturally, cease commuting to and from work.

Formally, we estimate event studies of the following form:

$$Mobility_{c,t} = \alpha_c + \phi_{CZ(c),t} + \sum_{k=\underline{K}}^{\overline{K}} \beta_k SAH_{c,t+k} + X_{c,t} + \underline{D}_{c,t} + \overline{D}_{c,t} + \varepsilon_{c,t} \tag{C.1.2}$$

where $Mobility_{c,t}$ represents either the retail or workplace mobility index published by Google for county c on day t, and $SAH_{c,t}$ is a dummy variable equal to 1 on the day a county imposes SAH orders. We set $\underline{K} = -17$ and $\overline{K} = 21$ so that the analysis examines three weeks prior and two and a half weeks following the imposition of SAH orders.[†] The event study is estimated over the period February 15th through April 24th, 2020. We non-parametrically control for county size by discretizing county employment into fifteen equally sized bins and interacting each bin with time fixed effects. α_c refers to the inclusion of county fixed effects. To isolate the local effect of SAH orders on economic activity, we also include commuting zone-by-time fixed effects.[†] This implies that our event-study estimates are identified only off of differential timing of SAH implementation among counties contained within the same commuting zone.

Results for retail mobility are presented in Figure C.1. The day SAH orders went into effect, there was an immediate decline of approximately 2% in retail mobility. This falls further to 7% the day after SAH order implementation, before slowly recovering to approximately 2% lower retail mobility two and a half weeks following the SAH order imposition.[†] The large transitory dip may reflect sentiment among consumers to shut-in before revisiting grocery stores and pharmacies. Alternatively, given our inclusion of commuting zone-by-time fixed effects, the transitory nature of the shock may reflect

this figure is available from the authors upon request.

[†]Of course, both indicators are equilibrium outcomes of both supply and demand shocks. The on-impact effect on work-place mobility at the very least reflects disruptions to each firm's ability to produce. Similarly, the on-impact effect on retail mobility is indicative of a decline in retail demand by consumers since, presumably, the supply of retail goods is at least fixed in the very short-run.

[†]Because our sample is necessarily unbalanced in event-time, we also include "long-run" dummy variables, $\underline{D}_{c,t}$ and $\overline{D}_{c,t}$. $\underline{D}_{c,t}$ is equal to 1 if a county imposed SAH orders at least $\overline{K}$ days prior. $\overline{D}_{c,t}$ is equal to 1 if a county will impose SAH at least $\underline{\text{K}}$ periods in the future.

[†]We use the United States Department of Agriculture (USDA) 2000 county to commuting zone crosswalk. This is available at `https://www.ers.usda.gov/data-products/commuting-zones-and-labor-market-areas/`.

[†]Restricting the sample to exclude never-takers yields the same result. This design identifies the mobility effects off of counties that ultimately implemented SAH orders but at different times.

negative, within-labor market spillovers of SAH orders. Regardless, the lack of a pre-trend is noticeable and provides additional support for a causal interpretation.

SAH orders may have affected firms' ability to produce by preventing workers from accessing their places of employment. To investigate whether SAH orders may have affected firms' productive capacity through this channel, we re-estimate our event study using workplace mobility as the outcome variable.[†]

Figure C.2 shows the result. As with the retail mobility event study, the workplace mobility index exhibits no differential pre-trend prior to the county-level imposition of SAH orders. In the first two days following the imposition of SAH orders, workplace mobility declined sharply relative to non-treated counties within its commuting zone. This relative decline in workplace mobility persists for nearly two and a half weeks following.

We draw three conclusions from these high-frequency event studies. First, the lack of pre-trends in the event studies suggest that the timing of SAH orders can be seen as plausibly randomly assigned with respect to local labor market conditions. This provides corroborating evidence for our cross-sectional identification strategy. In particular, it suggests that there were real effects of the SAH orders on local economies. Second, with the important caveat that both mobility indices are equilibrium objects, SAH orders appear to have had *both* local supply and local demand effects. Both retail mobility and workplace mobility fell substantially on impact and remained persistently low for at least two weeks following implementation of SAH orders. Third, given that overall workplace and retail mobility in the U.S. fell by 48 and 40 percent through April 24th relative to their baseline levels, our results bolster the claim that alternative mechanisms were responsible for the majority of job losses in the early weeks of the crisis; upon SAH implementation, relative workplace and retail mobility fell by, at most, 2 and 7 percent, respectively.

Alternative Cross-Sectional Specifications

The first type of robustness check we do is varying the horizon over which the cross-sectional regression is estimated, considering two natural alternative specifications: a two week horizon and a four week horizon. For the two week horizon specification, we consider cumulative initial claims between March 14 and March 28 regressed on SAH exposure over the same window; for the four week specification, the end date is April 11. We include the same set of controls as in our benchmark specification (Table 3.1, Column (5)).

Columns (1) and (2) of Table C.2 report the results from varying the horizon over which the model is estimated. Relative to our baseline result of 1.9%, estimating the model over

[†]An obvious concern with simply replacing the outcome variable is that changes in workplace mobility, unlike retail mobility, is highly dependent on the ability of individuals to work from home. The timing of SAH orders may be partially driven by the ability of workers in some regions to transition to working at home. In unreported regressions, we also non-parametrically control for this possibility by partitioning the WAH variable into 15 equally sized bins and interacting each bin with time fixed effects. The event study is essentially unchanged.

Figure C.1: County Retail Mobility Event Study

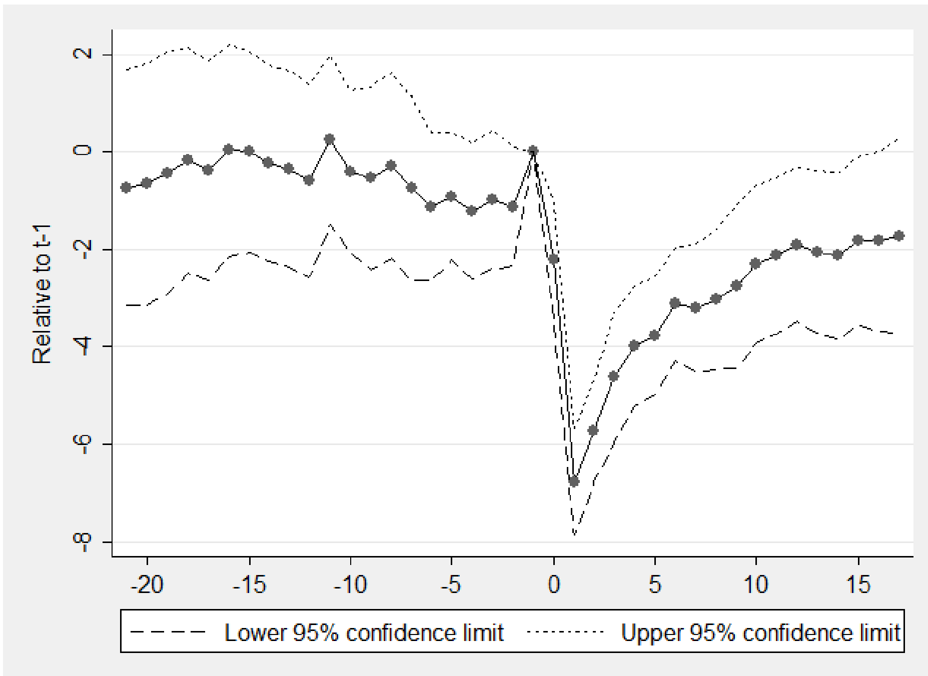

just two weeks lowers the point estimate slightly to 1.83% (SE: 0.91%). Conversely, when the model is estimated over a four week horizon, the point estimate is 1.7% (SE: 0.59%).

In Column (3) of Table C.2 we estimate the effect of SAH exposure on UI claims, over the same three week horizon as in the benchmark case, weighting observations by state-level employment from the QCEW in 2018 (an approached advocated for by some papers in the local multiplier literature).[†] Again, we consider the same set of controls as in our benchmark specification. The point estimate from the WLS regression is elevated slightly: 2.10% (SE: 0.54%). Regardless, weighting delivers quantitatively similar estimates.

Influence of Specific States

One may also be concerned that individual states' responses, either in terms of rising unemployment claims or SAH orders, is driving our results. To understand whether this is the case, we replicate our benchmark specification (column (5) in Table 3.1) from above, dropping one state at a time. The resulting coefficient estimates for β_C are available in Figure C.3, along with 90 percent confidence intervals constructed from robust standard errors.

[†]For arguments in either direction, see Ramey (2019) and Chodorow-Reich (Forthcoming), respectively. See also Solon et al. (2015).

Table C.1: Panel Specification: Effect of Stay-at-Home Orders on Initial Weekly Claims Relative to State Employment

	(1)	(2)	(3)	(4)
SAH Exposure Current Week	0.00919**	0.0101***	0.00997***	0.0125***
	(0.00350)	(0.00321)	(0.00329)	(0.00353)
SAH Exposure First Lag		-0.00293	-0.00367	-0.00299
		(0.00359)	(0.00358)	(0.00372)
SAH Exposure Second Lag		0.00245	-0.00115	0.000809
		(0.00230)	(0.00302)	(0.00332)
State FE	Y	Y	Y	N
Week FE	Y	Y	Y	Y
Post-March 21 X Work at Home Index	N	N	Y	Y
Post-March 21 X Excess Deaths per 1K	N	N	Y	Y
Post-March 21 X COVID-19 Cases per 1K	N	N	Y	Y
Post-March 21 X Avg. UI Replacement Rate	N	N	Y	Y
Adj. R-Square	0.826	0.822	0.831	0.801
No. Obs.	765	663	663	663

Table C.2: Effect of Stay-at-Home Orders on Cumulative Initial Weekly Claims Relative to State Employment: (i) 2-Week Horizon, (ii) 4-Week Horizon, (iii) Weighted Least Squares

	(1) Thru Mar. 28	(2) Thru Apr. 11	(3) WLS
SAH Exposure (varied horizons)	0.0183**	0.0166***	0.0209***
	(0.00908)	(0.00592)	(0.00541)
COVID-19 Cases per 1K	0.00197	0.000854	-0.00472
	(0.0109)	(0.00463)	(0.00306)
Excess Deaths per 1K	-0.0819	0.0691	0.214**
	(0.0959)	(0.0787)	(0.106)
Work at Home Index	-0.152	-0.587**	-0.486+
	(0.184)	(0.261)	(0.258)
Constant	0.111+	0.303***	0.242**
	(0.0649)	(0.0920)	(0.0921)
Adj. R-Square	0.0125	0.129	0.172
No. Obs.	51	51	51

Figure C.2: County Workplace Mobility Event Study

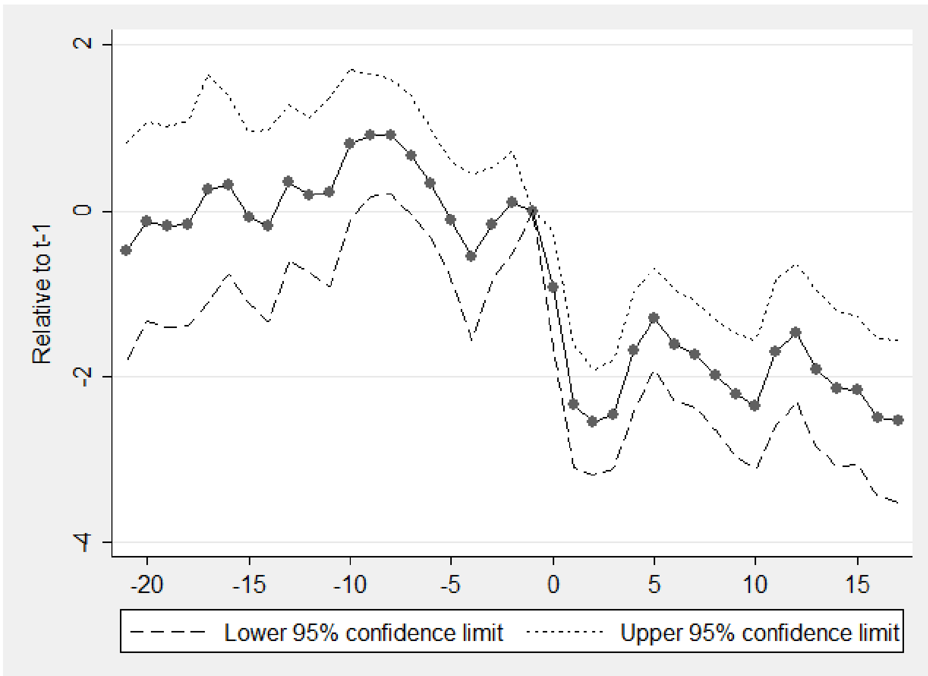

Figure C.3: Benchmark Specification Estimated Dropping One State at a Time

Pre-SAH Determinants of UI Claims

In this subsection, we broaden our analysis to adjust for determinants of state-level UI claims that may have been correlated with the timing of SAH implementation at the local level, as reported by the *New York Times*.

The first change that we make, relative to the results presented in Table 3.1, is to control for the March 7 to March 14 change in consumer spending. Because consumption is a leading indicator, changes to consumer spending tend to precede changes to employment. Thus, this allows us to control for leading determinants—as manifested in changes to state-level consumer spending—of employment losses that may have also been correlated with the timing of the implementation of SAH orders.

To do so, we rely upon the newly available, daily consumer spending index constructed by Chetty et al. (2020). These high frequency indicators of state-level economic activity is constructed from proprietary private sector microdata and made publicly available at `https://tracktherecovery.org`.

The second adjustment made in this subsection relates to the timing of state-level SAH implementation. In a few notable instances, the closure of non-essential businesses by state and local officials did not coincide with the broader SAH orders requiring all individuals to remain at home except for essential activities.[†] For example, on March 19 the governor of Pennsylvania issued a statewide executive order that required non-essential, in-person business activity to cease. This preceded by nearly a week the full statewide SAH order that was put into effect on March 23. A similar discrepancy between SAH dates and non-essential business closure occurred in Nevada.

This is potentially important since both Pennsylvania and Nevada experienced larger cumulative increases in UI claims to employment than the rest of the country through April 4. If the discrepancy between non-essential business closure and SAH implementation (as reported by the *New York Times*) was systematically correlated with the severity of job losses, then our estimate of β_C may be biased. In particular, if the pattern for Pennsylvania and Nevada holds more generally—large UI claims increase and relatively early non-essential business closure—then our estimates of β_C in Table 3.1 will be biased downwards, leading us to understate both the relative employment effect of SAH orders and their implied aggregate effect.

We adjust for the discrepancy between SAH implementation as reported in the *New York Times* and non-essential business closures by constructing a combined SAH/business closure treatment variable:

$$SAHBIZ_{s,t} = \max\left\{SAH_{s,t}, BIZ_{s,t}\right\}, \tag{C.1.3}$$

where $BIZ_{s,t}$ is the number of weeks state s was subject to a non-essential business closure through date t.[†].

Table C.3 records the results after incorporating the March 7 to March 14 change in the consumer spending index and adjusting the treatment variable to handle discrepancies between reported SAH implementation dates and dates of non-essential business closures. This table is structured identically to Table 3.1 except for the aforementioned changes.

[†]The closure of non-essential businesses is a prominent feature of most SAH orders.

[†]We use the state-level non-essential business closure dates compiled in Kong and Prinz (2020).

Both qualitatively and quantitatively the effect on unemployment of SAH orders is essentially unchanged relative to the benchmark specification. Consider Column (5): The point estimate of 1.9% (SE: 0.88%) implies that each additional week that a state was subject to a SAH order and/or non-essential business closures increased unemployment claims by 1.9% of the state's employment level.

While this point estimate is the same as our benchmark estimate, the relative-implied aggregate estimate of employment losses due to SAH orders through April 4, 2020 needs to be slightly adjusted. Incorporating non-essential business closure dates weakly increases each state's degree of SAH exposure. Recalculating equation (3.6) with the model estimated in Column (5) of Table C.3 yields an estimate of 4.6 million claims through April 4 attributable to SAH orders or approximately 27% of the overall increase in UI claims over the same period.[†]

County-Level Event Study Employment Specification

In Subsection 3.6 we use BLS-reported, month-to-month changes in county employment and unemployment to estimate the effect of SAH orders after controlling for state fixed effects. In what follows, we use county-level, high frequency employment indices to provide additional evidence that SAH orders had highly localized effects on county-level employment.[†]

Not only is the effect we estimate in this subsection consistent with our central finding, but by using high frequency, county-level data we are able to directly assess our assumption that the timing of local SAH implementation was uncorrelated with the relative severity of the local economic downturn. Consistent with the evidence presented in Subsection C.1, we find no evidence of differential pre-trends in employment around the implementation of SAH orders.

For the subset of counties for which the high-frequency employment indices are available, we estimate the following event study specification:

$$EmpIDX_{c,t} = \alpha_c + \phi_{state(c),t} + \sum_{k=\underline{K}}^{\overline{K}} \beta_k SAH_{c,t+k} + X_{c,t} + \underline{D}_{c,t} + \overline{D}_{c,t} + \varepsilon_{c,t} \tag{C.1.4}$$

[†]The two controls we consider in this section each slightly alter the estimated coefficient for the specification analogous to our benchmark specification. Controlling only for the change in the consumer spending index attenuates the point estimate to 1.4% (SE: 0.80%). Only adjusting for the discrepancies between non-essential business closure dates and reported SAH dates amplifies the point estimate somewhat to 2.4% (SE: 0.68); however, this latter effect appears to be driven almost entirely by Pennsylvania and Nevada. Dropping these states from the estimation yields a point estimate of 1.9% (SE: 0.68). These results are available upon request.

[†]The county-level employment indices we use were constructed by Chetty et al. (2020) and are available at `https://tracktherecovery.org` The county-level employment statistics we use are built out from anonymized microdata from private companies. See Chetty et al. (2020) for a fuller description of the data construction and for evidence that these series tend to track lower-frequency, publicly available series constructed from representative surveys.

where $EmpIDX_{c,t}$ represents the county-level, employment index available at `https://trackthereecovery.org`, $SAH_{c,t}$ is a dummy variable equal to 1 on the day a county imposes SAH orders, and $\phi_{state(c),t}$ is a state-by-time fixed effect. As in Subsection C.1, we set $\underline{K} = -17$ and $\overline{K} = 21$; the analysis thus examines three weeks prior and two and a half weeks following the imposition of SAH orders.[†] The event study is estimated over the period February 15th through April 24th, 2020. For this event study specification, we include no additional controls beyond county fixed effects and state-by-time fixed effects.

The results of this exercise are reported in Figure C.4. In the three weeks prior to the implementation of SAH orders, there is no statistically discernible pre-trend in employment.[†] However, there is a clear decline in employment after SAH orders were put into place. By one week following the SAH implementation, the employment index was down by 1.9% (SE: 0.5%). Two weeks following SAH implementation, the county-level index was down by by nearly twice as much.

For this analysis, we rely upon a subset of counties for which we have a high frequency measure of employment changes and for which there exist within-state variation. Nevertheless, despite relying upon a different subset of the variation for identification, the weekly effect on employment we estimate here is remarkably consistent with our state-level analysis, in terms of both magnitude and linearity of the effect. We view this as strongly corroborating our baseline finding and allaying concerns that the timing of SAH implementation was differentially correlated with the severity of each labor markets economic downturn.

[†]Our sample is necessarily unbalanced in event time, so we include "long-run" dummy variables $\underline{D}_{c,t}$ and $\overline{D}_{c,t}$ which are equal to 1 if a county imposed a SAH order at least $\overline{K}$ days prior or will impose a SAH order at least $\underline{\text{K}}$ days in the future, respectively.

[†]While not statistically meaningful, there appears to be a slight inflection point approximately one week prior to SAH implementation. However, even this is likely a statistical artifact, since the county-level employment statistics we rely upon are primarily reliant upon weekly payroll data from the company Paychex. Chetty et al. (2020) write: We convert the weekly Paychex data to daily measures of employment by assuming that employment is constant within each week.

Figure C.4: County Employment Event Study

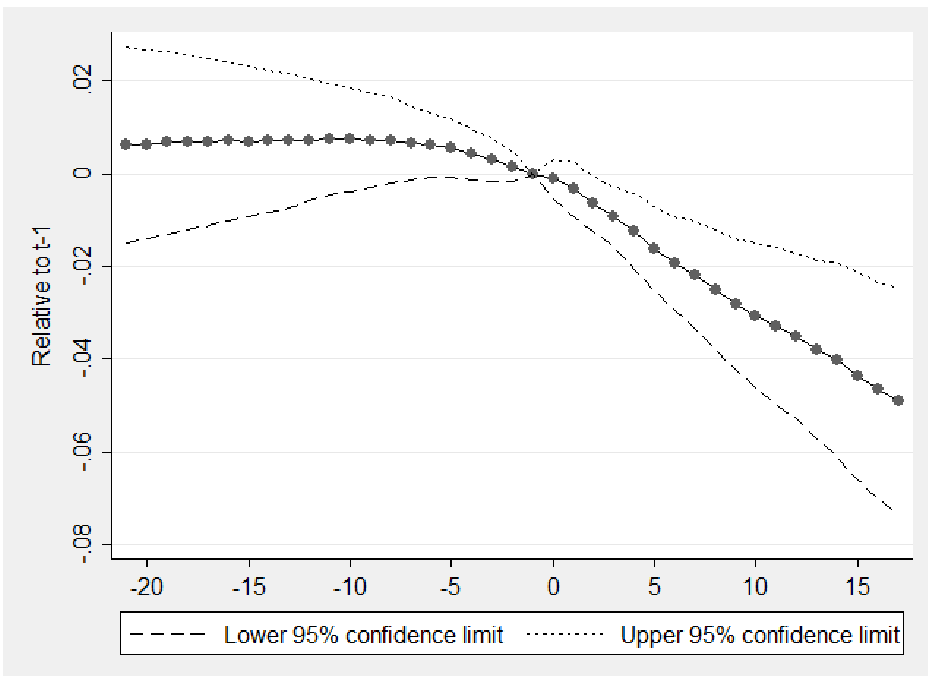

Table C.3: Effect of Stay-at-Home Orders on Cumulative Initial Weekly Claims Relative to State Employment for Weeks Ending March 21 thru April 4, 2020 After Accounting for Additional Pre-SAH Determinants of UI Claims.

	(1) Bivariate	(2) Covid	(3) Pol. Econ.	(4) Sectoral	(5) All
SAH/Business Closure Exposure	0.0214**	0.0218**	0.0215**	0.0224**	0.0191**
	(0.00855)	(0.00916)	(0.00972)	(0.00882)	(0.00884)
Mar. 7 to Mar. 14 Spending Change	-0.158	-0.183	-0.183	-0.310	-0.351
	(0.293)	(0.289)	(0.289)	(0.272)	(0.279)
COVID-19 Cases per 1K		-0.00295			0.00249
		(0.00579)			(0.00592)
Excess Deaths per 1K		0.0537			0.0637
		(0.120)			(0.109)
60+ Ratio to Total Population		0.308			
		(0.266)			
Avg. UI Replacement Rate			0.0740		0.0751
			(0.0764)		(0.0754)
2016 Trump Vote Share			0.00881		
			(0.0589)		
Work at Home Index				-0.500***	-0.563***
				(0.184)	(0.187)
Bartik-Predicted Job Loss				1.219	
				(7.388)	
Constant	0.0743***	0.0144	0.0372	0.259***	0.239***
	(0.0152)	(0.0517)	(0.0536)	(0.0793)	(0.0764)
Adj. R-Square	0.131	0.107	0.106	0.186	0.179
No. Obs.	51	51	51	51	51

C.2 Local SAH Orders in a Currency Union Model

We develop a framework to help us interpret the "relative effect"—which we estimate in the data—as compared to the "aggregate effect" of stay-at-home orders. To that end, we use a simple version of Nakamura and Steinsson (2014) of a two-country monetary union model, albeit abstracting from government spending as that is not the focus of our paper.

Households

Consider a currency union comprised of two regions: a home region of size n, and a foreign region of size $1-n$. In each region, there are infinitely many households with *identical* preferences and initial wealth.

A household j in home region has the following preferences:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\delta_t \frac{\left(C_t^j\right)^{1-\sigma}}{1-\sigma} - \chi \frac{\left(N_t^j\right)^{1+\psi}}{1+\psi} \right]$$

where

$$C_t^j = \left[\phi_H^{\frac{1}{\eta}} \left(C_{H,t}^j\right)^{\frac{\eta-1}{\eta}} + \phi_F^{\frac{1}{\eta}} \left(C_{F,t}^j\right)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \text{ with } \phi_H + \phi_F = 1,$$

$$C_{H,t}^j = \left(\int_0^n \left(\frac{1}{n}\right)^{\frac{1}{\epsilon}} c_{h,t}^j(i)^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}i \right)^{\frac{\epsilon}{\epsilon-1}}, \quad C_{F,t} = \left(\int_n^1 \left(\frac{1}{1-n}\right)^{\frac{1}{\epsilon}} c_{f,t}^j(i)^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}i \right)^{\frac{\epsilon}{\epsilon-1}}.$$

Total consumption of a household j in a home region is a CES aggregator of a *bundle* of home goods, $C_{H,t}^j$ and a *bundle* of foreign goods, $C_{F,t}^j$. Here, ϕ_F denotes the steady state share of the foreign goods imported from by a household in the home region. When $\phi_H = 1-\phi_F > n$, there is home bias.† η is the elasticity of substitution between home goods and imported goods from a foreign region, and ϵ denotes the elasticity of substitution across differentiated goods. β is discount factor and δ_t denotes consumption-preference shock in a home region, which evolves according to the following law of motion:

$$\log \delta_t = \rho^\delta \log \delta_{t-1} + \epsilon_t^\delta.$$

Then optimal allocations of expenditures (per household) are given by

$$C_{H,t}^j = \phi_H \left(\frac{P_{H,t}}{P_t}\right)^{-\eta} C_t^j, \quad C_{F,t} = \phi_F \left(\frac{P_{F,t}}{P_t}\right)^{-\eta} C_t^j,$$

$$c_{h,t}^j(i) = \left(\frac{p_{h,t}(i)}{P_{H,t}}\right)^{-\epsilon} C_{H,t}^j, \quad c_{f,t}^j(i) = \left(\frac{p_{f,t}(i)}{P_{F,t}}\right)^{-\epsilon} C_{F,t}^j,$$

†In the baseline calibration following Nakamura and Steinsson (2014), we calibrate $\phi_H = 0.69$ and $n = 0.1$, so that there is significant home bias.

with price indices defined as follows:

$$P_t = \left[\phi_H P_{H,t}^{1-\eta} + \phi_F P_{F,t}^{1-\eta}\right]^{\frac{1}{1-\eta}},$$
$$P_{H,t} = \left[\frac{1}{n}\int_0^n p_{h,t}(i)^{1-\epsilon}\mathrm{d}i\right]^{\frac{1}{1-\epsilon}},$$
$$P_{F,t} = \left[\frac{1}{1-n}\int_n^1 p_{f,t}(i)^{1-\epsilon}\mathrm{d}i\right]^{\frac{1}{1-\epsilon}}.$$

Here, P_t denotes consumper price index of a home region, and $P_{H,t}$ ($P_{F,t}$) is producer price index of home (foreign) goods.

In our baseline specification, we assume identical households in a given region with the same initial wealth and *complete* financial markets, which makes aggregation straightforward. Thus, we have

$$c_{h,t}(i) \equiv \int_0^n c_{h,t}^j(i)\mathrm{d}j = \left(\frac{p_{h,t}(i)}{P_{H,t}}\right)^{-\epsilon} C_{H,t}, \quad c_{f,t}(i) \equiv \int_0^n c_{f,t}^j(i)\mathrm{d}j = \left(\frac{p_{f,t}(i)}{P_{F,t}}\right)^{-\epsilon} C_{F,t}$$
$$C_{H,t} = \int_0^n C_{H,t}^j \mathrm{d}j = \phi_H \left(\frac{P_{H,t}}{P_t}\right)^{-\eta} C_t, \quad C_{F,t} = \int_n^1 C_{F,t}^j \mathrm{d}j = \phi_F \left(\frac{P_{F,t}}{P_t}\right)^{-\eta} C_t,$$
$$C_t = \int_0^n C_t^j \mathrm{d}j = nC_t^j,$$

where variables without j superscript are aggregate variables in a home region.

With the optimal allocations, we can write household j's budget constraint (in real terms with the home region's CPI as a numeraire) as follows:

$$C_t^j + \mathbb{E}_t\left[M_{t,t+1}B_{t+1}^j\right] \leq B_t^j + \frac{W_t}{P_t}N_t^j + \int_0^1 \frac{\Xi_{h,t}^j(i)}{P_t}\mathrm{d}i - \frac{T_t^j}{P_t}.$$

Note that W_t is home region's nominal wage, and N_t^j is a household j's labor supply. Here, we assume perfect immobility across the regions, meaning wages will be determined at the regional level. B_{t+1}^j is a household j's state-contingent asset holdings and note again that we assume complete financial markets. Here P_t denotes price index that gives the minimum price of one unit of consumption good, C_t. *i.e.* P_t is the Consumer Price Index (CPI) in the home region.

Optimality conditions for $j \in (0, n]$ are

$$\chi\left(N_t^j\right)^{\psi} = \delta_t \left(C_t^j\right)^{-\sigma}\frac{W_t}{P_t},$$
$$\delta_t\left(C_t^j\right)^{-\sigma} = \beta\mathbb{E}_t\left[\delta_{t+1}\left(C_{t+1}^j\right)^{-\sigma}\frac{1+i_t}{1+\pi_{t+1}}\right],$$

where i_t is one-period nominal spot interest rate which satisfies $\mathbb{E}_t[M_{t,t+1}] = 1/(1+i_t)$.

Households in the foreign region are symmetric relative to those in the home region, and we use $*$ to denote foreign variables. So we have

$$C_t^{*j} = \left[(\phi_H^*)^{\frac{1}{\eta}} \left(C_{H,t}^{*j}\right)^{\frac{\eta-1}{\eta}} + (\phi_F^*)^{\frac{1}{\eta}} \left(C_{F,t}^{*j}\right)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \text{ with } \phi_H^* + \phi_F^* = 1.$$

For *aggregate* optimal allocations in the foreign region, we have

$$c_{h,t}^*(i) \equiv \int_n^1 c_{h,t}^{*j}(i)\mathrm{d}j = \left(\frac{p_{h,t}^*(i)}{P_{H,t}^*}\right)^{-\epsilon} C_{H,t}^*, \quad c_{f,t}^*(i) \equiv \int_n^1 c_{f,t}^{*j}(i)\mathrm{d}j = \left(\frac{p_{f,t}^*(i)}{P_{F,t}^*}\right)^{-\epsilon} C_{F,t}^*$$

$$C_{H,t}^* = \int_n^1 C_{H,t}^{*j}\mathrm{d}j = \phi_H^* \left(\frac{P_{H,t}^*}{P_t^*}\right)^{-\eta} C_t^*, \quad C_{F,t}^* = \int_n^1 C_{F,t}^{*j}\mathrm{d}j = \phi_F^* \left(\frac{P_{F,t}^*}{P_t^*}\right)^{-\eta} C_t^*,$$

$$C_t^* = \int_n^1 C_t^{*j}\mathrm{d}j = (1-n)C_t^{*j}.$$

Optimality conditions for foreign households for $j \in [n, 1)$ are

$$\chi \left(N_t^{s,j*}\right)^{\psi} = \delta_t^* \left(C_t^{j*}\right)^{-\sigma} \frac{W_t^*}{P_t^*},$$

$$\delta_t^* \left(C_t^{j*}\right)^{-\sigma} = \beta \mathbb{E}_t \left[\delta_{t+1}^* \left(C_{t+1}^{j*}\right)^{-\sigma} \frac{1+i_t}{1+\pi_{t+1}^*}\right].$$

Terms of Trade, and Real Exchange Rate

Before moving on to firms in each region, let us define terms showing the relationships between various price measures. First, we define terms of trade, S_t as

$$S_t \equiv \frac{P_{F,t}}{P_{H,t}}.$$

From this, we can write the relationship between CPI and Producer Price Index (PPI) in a home region as:

$$g(S_t) \equiv \frac{P_t}{P_{H,t}} = \left[\phi_H + \phi_F S_t^{1-\eta}\right]^{\frac{1}{1-\eta}}, \quad \frac{P_t}{P_{F,t}} = \frac{P_t}{P_{H,t}} \frac{P_{H,t}}{P_{F,t}} = \frac{g(S_t)}{S_t}.$$

For the case of the foreign region, we have

$$g^*(S_t) \equiv \frac{P_t^*}{P_{H,t}^*} = \left[\phi_H^* + \phi_F^* S_t^{1-\eta}\right]^{\frac{1}{1-\eta}}, \quad \frac{P_t^*}{P_{F,t}^*} = \frac{P_t^*}{P_{H,t}^*} \frac{P_{H,t}^*}{P_{F,t}^*} = \frac{g^*(S_t)}{S_t}.$$

Finally, we write the real exchange rate in terms of $g(S_t)$ and $g^*(S_t)$ as follows:

$$Q_t = \frac{P_t^*}{P_t} = \frac{g^*(S_t)}{g(S_t)}.$$

Firms

We assume that there is a continuum of intermediate-goods-producing firms in each region, producing differentiated intermediate goods by using labor as input. We assume a competitive labor market.

Production technologies of each intermediate-goods-producing firms are given by

$$y_{h,t}(i) = A_t N_{h,t}(i)^{\alpha}, \quad \alpha < 1,$$
$$y_{f,t}(i) = A_t^* N_{f,t}^*(i)^{\alpha}, \quad \alpha < 1,$$

where $y_{h,t}(i)$ $(y_{f,t}(i))$ is the production output of a firm i in the home (foreign) region, $N_{h,t}(i)$ $(N_{f,t}^*(i))$ is the amount of labor input hired by a firm i in the home (foreign) region, and A_t (A_t^*) is region-wide technology in the home (foreign) region. Both technology processes evolve according to the following laws of motion:

$$\log A_t = \rho^A \log A_{t-1} + \epsilon_t^A,$$
$$\log A_t^* = \rho^{A*} \log A_{t-1}^* + \epsilon_t^{A*}$$

This implies that region-wide labor demand can be written as

$$\begin{aligned}
N_t &= \int_0^n N_{h,t}(i)\mathrm{d}i = \int_0^n \left(\frac{y_{h,t}(i)}{A_t}\right)^{\frac{1}{\alpha}} \mathrm{d}i = \left(\frac{1}{A_t}\right)^{\frac{1}{\alpha}} \int_0^n y_{h,t}(i)^{\frac{1}{\alpha}}\mathrm{d}i \\
&= \left(\frac{Y_{H,t}}{A_t}\right)^{\frac{1}{\alpha}} \int_0^n \frac{1}{n} \left(\frac{p_{h,t}(i)}{P_{H,t}}\right)^{-\frac{\epsilon}{\alpha}} \mathrm{d}i = \left(\frac{Y_{H,t}}{A_t}\right)^{\frac{1}{\alpha}} \Delta_t^{\frac{1}{\alpha}}, \\
N_t^* &= \int_0^n N_{f,t}^*(i)\mathrm{d}i = \int_n^1 \left(\frac{y_{f,t}(i)}{A_t^*}\right)^{\frac{1}{\alpha}} \mathrm{d}i = \left(\frac{1}{A_t^*}\right)^{\frac{1}{\alpha}} \int_n^1 y_{f,t}(i)^{\frac{1}{\alpha}}\mathrm{d}i \\
&= \left(\frac{Y_{F,t}}{A_t^*}\right)^{\frac{1}{\alpha}} \int_n^1 \frac{1}{1-n} \left(\frac{p_{f,t}(i)}{P_{i,t}}\right)^{-\frac{\epsilon}{\alpha}} \mathrm{d}i = \left(\frac{Y_{F,t}}{A_t^*}\right)^{\frac{1}{\alpha}} (\Delta_t^*)^{\frac{1}{\alpha}},
\end{aligned}$$

by defining $\Delta_t \equiv \frac{1}{n}\int_0^n \left(\frac{p_{h,t}(i)}{P_t}\right)^{-\epsilon} \mathrm{d}i$, and $\Delta_t^* \equiv \frac{1}{1-n}\int_n^1 \left(\frac{p_{f,t}(i)}{P_t^*}\right)^{-\epsilon} \mathrm{d}i$ as price dispersion terms in each region.

Firms are subject to Calvo-type pricing frictions, so they solve the following problem:

$$\max_{p_{h,t}^{\#}(i)} \mathbb{E}_t \left[\sum_{k=0}^{\infty} Q_{t,t+k}\theta^k \left(p_{h,t}^{\#}(i) - MC_{h,t+k|t}(i)\right) y_{h,t+k|t}(i)\right]$$

subject to $y_{h,t+k|t}(i) = \left(\frac{p_{h,t}^{\#}(i)}{P_{H,t}}\right)^{-\epsilon} \left(C_{H,t} + C_{H,t}^*\right)$, and with $Q_{t,t+k} = \beta^k \frac{\delta_{t+k}u'(C_{t+k})}{\delta_t u'(C_t)}$. Note that here, $C_{H,t}^*$ denotes a composite index of foreign consumption of home goods, and $MC_{h,t+k|t}(i)$ is nominal marginal cost.

Then optimality conditions for pricing are given by

$$p_{h,t}^{\#}(i) = \frac{\epsilon}{\epsilon - 1} \frac{\mathbb{E}_t \sum_{k=0}^{\infty} (\beta\theta)^k \delta_{t+k} u'(C_{t+k}) mc_{h,t+k|t}(i) P_{H,t+k}^{\epsilon} \left(C_{H,t} + C_{H,t}^*\right)}{\mathbb{E}_t \sum_{k=0}^{\infty} (\beta\theta)^k \delta_{t+k} u'(C_{t+k}) P_{H,t+k}^{\epsilon-1} \left(C_{H,t} + C_{H,t}^*\right)},$$

with $mc_{h,t+k|t}(i)$ is real marginal cost of a firm i in terms of PPI, $P_{H,t}$.

Aggregate real marginal cost with $\alpha < 1$ can be written as follows:

$$\begin{aligned} mc_{h,t}(i) &= \frac{W_t/P_{H,t}}{\alpha A_t N_{h,t}(i)^{\alpha-1}} = \frac{w_t}{\alpha A_t} N_{h,t}(i)^{1-\alpha} \\ &= \frac{w_t}{\alpha A_t} \left(\frac{y_{h,t}(i)}{A_t}\right)^{\frac{1-\alpha}{\alpha}} = \frac{w_t}{\alpha A_t} \left(\frac{Y_{H,t}}{A_t}\right)^{\frac{1-\alpha}{\alpha}} \left(\frac{y_{h,t}(i)}{Y_{H,t}}\right)^{\frac{1-\alpha}{\alpha}} \\ &= mc_{H,t} \left(\frac{p_{h,t}(i)}{P_{H,t}}\right)^{-\frac{\epsilon(1-\alpha)}{\alpha}}, \\ mc_{H,t} &\equiv \frac{w_t}{\alpha A_t} \left(\frac{Y_{H,t}}{A_t}\right)^{\frac{1-\alpha}{\alpha}}. \end{aligned}$$

with $w_t \equiv W_t/P_{H,t}$.
Combining this with the previous optimal pricing equation then generates

$$p_{h,t}^{\#}(i)^{1+\frac{\epsilon(1-\alpha)}{\alpha}} = \frac{\epsilon}{\epsilon - 1} \frac{\mathbb{E}_t \sum_{k=0}^{\infty} (\beta\theta)^k u'(C_{t+k}) mc_{H,t+k} P_{H,t+k}^{\epsilon/\alpha} Y_{H,t+k}}{\mathbb{E}_t \sum_{k=0}^{\infty} (\beta\theta)^k u'(C_{t+k}) P_{H,t+k}^{\epsilon-1} Y_{H,t+k}}.$$

We have similar conditions for intermediate-goods-producing firms in the foreign region.

International Risk Sharing Condition and Market Clearing Conditions

Combining each region's Euler equation gives

$$\delta_t \left(\frac{1}{n} C_t\right)^{-\sigma} = \kappa \delta_t^* \left(\frac{1}{1-n} C_t^*\right)^{-\sigma} \frac{1}{\mathcal{Q}_t},$$

with complete markets and symmety of initial conditions, $\kappa = 1$, generating

$$\delta_t^{-\frac{1}{\sigma}} C_t = \frac{n}{1-n} \delta_t^{*-\frac{1}{\sigma}} C_t^* \mathcal{Q}_t^{\frac{1}{\sigma}},$$

with $\mathcal{Q}_t \equiv P_t^*/P_t$ for the real exchange rate.

Goods market clearing conditions in each region are:

$$Y_{H,t} = C_{H,t} + C_{H,t}^* = \phi_H \left(\frac{P_{H,t}}{P_t}\right)^{-\eta} C_t + \phi_H^* \left(\frac{P_{H,t}^*}{P_t^*}\right)^{-\eta} C_t^*,$$

$$Y_{F,t} = C_{F,t} + C_{F,t}^* = \phi_F \left(\frac{P_{F,t}}{P_t}\right)^{-\eta} C_t + \phi_F^* \left(\frac{P_{F,t}^*}{P_t^*}\right)^{-\eta} C_t^*.$$

Finally, we close the model by imposing the following monetary policy rule:

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)(\phi_\pi \pi_t^{agg} + \phi_y \hat{y}_t^{agg}),$$

where π_t^{agg} is a union-wide inflation rate and $\hat{y}_t^{agg}$ is union-wide output gap.

Modelling Stay-at-Home Orders

We model the imposition of SAH orders in two ways: (i) as a local supply shock, and (ii) as a local demand shock. When we model the SAH as a local productivity shock, we introduce the negative productivity shock for intermediate-goods-producing firms by setting negative values for ϵ_t^A. Alternatively, we also model the imposition of SAH orders via a negative preference shock, since SAH orders may directly reduce consumption by limiting retail mobility, as discussed in Subsection C.1. In this case, we introduce negative shocks to ϵ_t^δ.

C.3 Data Appendix

Table C.4 reports all sources used in this paper.

Table C.4: Data Sources

Variable	Source
Initial Unemployment Claims (Accessed 6/17/2020)	FRED (Mnemonic *ICLAIMS, where * indicates state abbreviation)
County Employment Data	BLS `https://www.bls.gov/lau` (Accessed 6/4/2020)
Stay-at-Home Orders (Accessed with *Internet Archive*)	*New York Times* `https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html`
Covid Confirmed Cases (Accessed 6/5/2020)	UsaFacts `https://usafacts.org/visualizations/coronavirus-covid-19-spread-map/`
State Excess Deaths (Accessed 6/4/2020)	CDC `https://www.cdc.gov/nchs/nvss/vsrr/covid19/excess_deaths.htm`
Share Age 60+ (Accessed 6/16/2020)	Census Bureau `https://www.census.gov/data/tables/time-series/demo/popest/2010s-state-detail.html`
Average UI Replacement Rate (Accessed 6/16/2020)	Department of Labor's Employment and Training Administration `https://oui.doleta.gov/unemploy/ui_replacement_rates.asp`
2016 Trump Vote Share (Accessed 6/17/2020)	*New York Times* `https://www.nytimes.com/elections/2016/results/president`
Work at Home Index	Dingel and Neiman (2020)
March Employment Losses for Bartik (Accessed 4/10/2020)	BLS `https://download.bls.gov/pub/time.series/ce/ce.industry`
Google Mobility Reports (Accessed 5/21/2020)	`https://www.google.com/covid19/mobility/`
Daily Consumer Spending and Employment	Track the Recovery `https://tracktherecovery.org`
State Non-Essential Business Closure Dates	Kong and Prinz (2020)